Perfect Graphs and Orthogonally Convex Covers

Rajeev Motwani †
Arvind Raghunathan ‡
Huzur Saran ‡

573 Evans Hall
Computer Science Division
U.C. Berkeley, CA 94720.

ABSTRACT

We consider the problem of covering simple orthogonal polygons with convex orthogonal polygons. In the case of horizontally or vertically convex polygons we show that the polygon covering problem can be reduced to the problem of covering a permutation graph with minimum number of cliques.

In general, orthogonal polygons can have concavities (dents) with four possible orientations. In the case where the polygon has three dent orientations, we show that the polygon covering problem can be reduced to the problem of covering a weakly triangulated graph with a minimum number of cliques. Since weakly triangulated graphs are perfect, we obtain the following duality relationship: the minimum number of orthogonally convex polygons needed to cover an orthogonal polygon P with at most three dent orientations is equal to the maximum number of points of P, no two of which can be contained together in an orthogonally convex covering polygon.

Finally, we show that in the case of orthogonal polygons with all four dent orientations, the above duality relationship fails to hold.

December 15, 1987

† Supported by the National Science Foundation under grant DCR-8411954.
‡ Supported by the Semiconductor Research Corporation under grant SRC-52055.

Perfect Graphs and Orthogonally Convex Covers

Rajeev Motwani †
Arvind Raghunathan ‡
Huzur Saran ‡

573 Evans Hall
Computer Science Division
U.C. Berkeley, CA 94720.

1. Introduction

One of the most well-studied class of problems in computational geometry concerns the notion of *visibility*. Two points in the plane are said to be visible from each other in the presence of obstacles (which are generally polygonal) if there exists a straight-line path between the two points which does not meet any of the obstacles. Other notions of visibility involve paths which are not straight lines, e.g. rectilinear or staircase paths. There is an intimate connection between visibility problems and polygon covering problems. In his recent book on the Art Gallery Problem, O'Rourke [20] states that the fundamental problems involving visibility in computational geometry will not be solved until the combinatorial structure of visibility is more fully understood. In this paper (and a companion paper [19]) we attempt to study this combinatorial structure. A *visibility graph* has vertices which correspond to geometric components, such as points or lines, and edges which correspond to the visibility of these components from each other. Here, we will be concerned with the visibility graphs for regions inside a simple orthogonal polygon. We show that certain special classes of these visibility graphs are perfect. We use this property of the visibility graphs to devise polynomial algorithms for a class of polygon covering problems that are NP-hard in general.

An *orthogonal* (or *rectilinear*) polygon (OP), P, is a polygon with all its edges parallel to one of the co-ordinate axes. A polygon is said to be *simple* if it has no holes, i.e. the polygon boundary is a closed, connected curve. Let n denote the number of edges on the boundary of P. Here, we are only concerned with simple orthogonal polygons. An orthogonal polygon is said to be *horizontally convex* (or *vertically convex*) if its intersection with every horizontal (resp. vertical) line segment is either empty or a single line segment. An *orthogonally convex polygon* (OCP) is both horizontally and vertically convex. Every OP induces a grid which is constructed by drawing horizontal and vertical lines through all its vertices. A collection of polygons, $C = \{P_1, P_2, \cdots, P_r\}$ where $P_i \subseteq P$, is said to cover a polygon P if the union of all the polygons in C is P. Whenever we speak of a set of covering polygons for an arbitrary polygon P, it will be assumed that each covering polygon is totally contained in P.

The following classification of orthogonal polygons is due to Culberson and Reckhow [7]. Consider the traversal of the boundary of P in the clockwise direction, i.e. ensuring that the interior is always to the right. At each corner (vertex) of P, we either turn 90^o right (outside corner) or 90° left (inside corner). A *dent* is an edge of the perimeter of P, both of whose endpoints are inside corners. The direction of traversing a dent gives its *orientation*: for instance, a dent traversed from west to east has a **N** orientation. We will use the natural definition of the compass direction, i.e. the positive direction along the y-axis will be referred to as the north direction and so on. Figure 1 illustrates the **N**, **S**, **E** and **W** dents. For a dent D, $o(D)$ indicates its orientation. Two dents D_1 and D_2 are said to be similarly oriented if $o(D_1) = o(D_2)$. D_1 and D_2 are said to be oppositely oriented if $o(D_1) = \mathbf{N}$, $o(D_2) = \mathbf{S}$ or, if $o(D_1) = \mathbf{E}$, $o(D_2) = \mathbf{W}$. Otherwise, D_1 and D_2 are said to have orthogonal orientations. An OP is classified according the orientations of its dents. A class k OP has dents of k different orientations. A class 0 OP does not have dents and is an OCP. A vertically or horizontally convex polygon is a class 2 OP which has only an opposing pair of dents, i.e. either **N** and **S** or **E** and **W** (see Figure 2). A class 3 OP without **N** dents is shown in Figure 3.

The problem of covering general (non-orthogonal) polygons by simpler components has received considerable attention in the literature [5, 6, 15, 16, 25]. It turns out, however, that most of these problems are NP-hard, especially in the case where the polygons are allowed to have holes [1, 25]. The various kinds of orthogonal coverings studied earlier include coverings by rectangles, orthogonally convex polygons and star-shaped orthogonal polygons. Several algorithmic results have been obtained for covering orthogonal polygons. For instance, Franzblau and Kleitman have an $O(n^2)$ algorithm for covering a vertically convex orthogonal polygon without holes with a minimum number of rectangles [8]. Keil has provided an $O(n^2)$ algorithm for covering similar polygons with a minimum number of orthogonally convex polygons [17]. Reckhow and Culberson [21] later provided an $O(n^2)$ algorithm for covering a class 2 orthogonal polygon with a minimum number of orthogonally convex polygons.

In this paper, we are more concerned with combinatorial results concerning covering problems for orthogonal polygons. Let an *independent set* of points in a polygon P, with respect to a class of covering polygons C, denote a set of points in P, no two of which can be covered by any covering polygon from the class C. A *duality* theorem for covering problems is of the following form: the size of the minimum cover by polygons from class C is equal to the size of the maximum independent set of points with respect to the class C. Many interesting duality theorems have been obtained for polygonal covering problems. Chv`atal [8] conjectured that a duality theorem holds for the problem of covering orthogonal polygons by rectangles. This conjecture was shown to be false by Szemeredi and Chung (cited in [4]). However, Chaiken et al [4] showed that the duality theorem holds for polygons that are orthogonally convex. Győri [12] then showed that the duality relationship holds even if the polygon is only vertically (or horizontally) convex. For the case of horizontally convex polygons, an $O(n^2)$ algorithm was devised by Franzblau and Kleitman [8], based on an observation due to A. Frank. Later, Saks [23] showed that a graph determined by the boundary squares of the grid induced by the vertices of an OCP is perfect. Other related work includes that of Shearer [24], Boucher [3] and Albertson & O'Keefe [2].

The purpose of this paper is three-fold: we first show that the visibility graph of a vertically convex orthogonal polygon is a permutation graph [10, 18]. We then show that the visibility graph of a class 3 polygon is a weakly triangulated graph [13]. Thus, we reduce the polygon covering problem to the problem of covering a weakly triangulated graph with a minimum number of cliques. Since weakly triangulated graphs are perfect [13], we get the following duality relationship for a class 3 polygon P: the minimum number of orthogonally convex polygons needed to cover an orthogonal polygon P is equal to the maximum number of points of P, no two of which can be contained together in an orthogonally convex covering polygon. Further, the Ellipsoid Method of Grötschel, Lov'asz and Schrijver [11] gives us a polynomial algorithm for the covering problem for polygons with three dent orientations. Recently, Hayward and Hoang [14] have obtained an $O(n^5)$ algorithm for the minimum clique cover problem for weakly triangulated graphs, thus providing us with a purely combinatorial algorithm for the polygon covering problem under consideration. Since this algorithm is combinatorial, it does not suffer from the same drawbacks as the Ellipsoid method. Finally, we show that the visibility graph is not perfect for general (class 4) polygons. Further, we show that the above duality relationship fails to hold for general (or class 4) orthogonal polygons.

Reckhow and Culberson [7] independently reached the conclusion that the problem of covering an orthogonally convex polygon with two dent orientations can be reduced to that of finding a minimum clique cover in a comparability graph [10, 18]. However, permutation graphs are a subset of comparability graphs, and, as such, we believe that our result is stronger. Again, in the case of covering a polygon with three dent orientations, Robert Reckhow [22] has independently shown that the visibility graph of such a polygon satisfies the Strong Perfect Graph Conjecture [10, 18]. He also provides an $O(n^2)$ geometric algorithm to cover a class 3 polygon with a minimum number of orthogonally convex polygons.

This paper is organized as follows. In Section 2, we discuss the theoretical framework for this problem, as discussed in [7, 21]. Section 3 discusses the connection between the covering problem for vertically convex polygons and permutation graphs. In Section 4, we state our main results for covering class 3 polygons. Section 5 gives a proof of a technical lemma called the Crossing Lemma, which we use to show that the covering problem for class 3 polygons reduces to the clique covering problem for weakly triangulated graphs in Sections 6 and 7, giving us the duality relationship mentioned above. Section 8 shows why we feel that the above techniques will probably not extend to the more general problem of covering a class 4 polygon. We also show that the duality relationship fails to hold for these polygons. In Section 9 we will consider possible extensions of our results.

2. Preliminaries

In this section we develop some of the tools required to analyze the problem of finding minimum orthogonally convex covers for orthogonal polygons. Some of the definitions and observations stated here are due to Culberson and Reckhow [21, 7]. Throughout this paper, P refers to the simple orthogonal polygon to be covered.

2.1. Staircase Paths

A *maximal OCP* in P is an OCP contained in P, but not contained in any other OCP contained in P. A *staircase path* in P corresponds to a sequence of points $u = x_0, x_1, \cdots, x_r = v$ contained in P such that, (a) each adjacent pair of points, x_i and x_{i-1}, determine a vertical or horizontal line segment which is contained in P, and (b) in traversing the staircase path from u to v the edges corresponding to the adjacent pairs of point are traversed in at most two of the four possible compass directions. More informally, a staircase path is a connected sequence of horizontal and vertical edges such that the path alternates between left and right turns. We say $u \equiv v$ (read as *u sees v*) if there exists a staircase path joining u and v. We will denote by $s(u,v)$ a fixed staircase path with $u,v \in P$ as its two end-points. The following observations demonstrate the inherent relationship between staircase paths and covers by OCP's.

Observation 1. For any two points $u, v \in P$, $u \equiv v$ if and only if some OCP includes them both.

Observation 2. Any covering of P by OCP's can be made into a covering of P by the same number of maximal OCP's.

We say that a staircase path from u to v goes southwest if, in traversing it from u to v, we go west on all the horizontal segments and south on all vertical segments. Thus, staircase paths between u and v can be of four possible orientations: northeast, northwest, southeast, southwest. However, depending upon the direction of traversal, the same staircase path might be viewed as a northwest/southeast path in one case, or a northeast/southwest path in the other. Thus, we classify staircase paths into two types: In a type I staircase path, one may travel northwest or southeast on it and in a type II staircase path, one may travel northeast or southwest on it. A vertical or horizontal staircase path (line) is both of type I and II.

2.2. Dent Lines and Zones

For each dent edge D, we construct a *dent line* $\vec{D}$ by extending D in both directions until it meets the perimeter of P. The orientation of $\vec{D}$ is the same as the orientation of D. $\vec{D}$ divides P into three *zones*. Two of these zones, called $B_l(\vec{D})$ and $B_r(\vec{D})$ (see Figure 4), are said to be *below* the dent, and are the two connected components of P which lie to the left in a clockwise traversal of $\vec{D}$. We will refer to the region $B_l(\vec{D}) \cup B_r(\vec{D})$ as the B zone or $B(\vec{D})$. The third zone, $A(\vec{D})$, is said to lie *above* $\vec{D}$, and is the connected component of P which lies to the right in a clockwise traversal of $\vec{D}$. For any two points $u \in B_l(D)$, $v \in B_r(D)$, there will not exist any staircase path between u and v and, thus, $u \not\equiv v$. We now observe the following facts about dents and zones. Those stated without proof can be found in [21].

Observation 3. Let u and v be two points in P. If $u \not\equiv v$ then there exists a dent D such that $u \in B_l(D), v \in B_r(D)$ or $v \in B_l(D), u \in B_r(D)$.

In this case, we say that D *separates* u and v, and D itself is called a separating dent for u and v. In general, there may be more than one dent separating two points in P.

Observation 4. Let u, v and w be three points in P such that a dent D separates u from v. If $u \equiv w$ and $v \equiv w$, then $w \in A(D)$.

The following observations are stated for particular dent orientations. However, it is not very hard to see that they hold in all of their reflection and rotation symmetric versions.

Observation 5. Let D be a dent and u, v be points such that $u \in B(D)$ and $v \in A(D)$. Without loss of generality, let $o(D) = \mathbf{N}$, then u lies to the north of $\vec{D}$ and v lies to the south of $\vec{D}$.

Proof: Every staircase path from v to u must at first lie totally in $A(D)$, then cross $\vec{D}$ and then lie totally in $B(D)$. When the path leaves $A(D)$ it must be traveling north. If v lies to the north of $\vec{D}$, then the path must first travel south to get below $\vec{D}$. But no staircase path can travel both north and south. A similar argument establishes that u lies to the north of $\vec{D}$.

Q.E.D.

Observation 6. Let u, v and w be points in P such that v lies to the northwest of u and w lies to the northeast of u. If $u \equiv v$, $u \equiv w$ and $v \not\equiv w$ then there is a **N** dent separating v from w.

2.3. The Region DAG, the Visibility Graph and the Source Graph

The set of all dent lines of P subdivides P into *regions*. Reckhow and Culberson [21] construct a region DAG (directed acyclic graph) as follows: The nodes of the region DAG are the regions, and there is an arc from region u to region v if u and v share a common border $\vec{D}$ and u is *below* $\vec{D}$ (see Figure 5). In the following, we relate the notion of visibility by staircase paths with the covering problem.

Definition 1. Let u and v be regions in P. We say that $u \equiv v$ (read as region u *sees* region v) if and only if some OCP (contained in P) includes both u and v.

Observation 7. [21] Let u and v be regions in P, and let q_u and q_v be arbitrary points in u and v respectively. Then, there is a staircase path between q_u and q_v if and only if $u \equiv v$.

The *visibility graph*, $G(V, E)$, for the polygon P is an undirected version of the closure of the region DAG. More formally, the vertex set V of G contains a vertex corresponding to every region in P. Two vertices u and v are adjacent in the graph G if the corresponding regions in P can be covered by a single OCP. We will use the same notation for a region of P and the corresponding vertex in V. Thus, we have that $<u,v> \in E$ if and only if $u \equiv v$. Clearly, if $<u,v> \in E$ then there is a staircase path from each point in the region u to each point in the region v.

A *source* is a region of zero in-degree in the region DAG (see Figure 5). It is easy to see that each source has at most one dent line of a given orientation as its border: two dent lines of the same orientation are parallel to each other, and no region of the polygon with both of them as border lines can be a source. Thus a source can be called an order k source, $1 \le k \le 4$, where k denotes the number of dent lines that border the source. It is clear that a class 3 polygon does not have order 4 sources. From this, Reckhow and Culberson [21] show that a class 3 polygon has only $O(n)$ sources. The following result from [21] shows the importance of sources.

Lemma 1. If β is a set of maximal orthogonally convex polygons that includes every source of P, then β includes every region of P.

We can now construct the source graph $G_s(V_s, E_s)$ as follows (see Figure 6). The vertex set V_s has a vertex corresponding to each source region of P. As before, we have the edge $<u,v>$ in E_s if $u \equiv v$. Clearly, $V_s \subseteq V$ and the source graph G_s is an vertex induced subgraph of the visibility graph G. The following lemma [21] provides the relationship between the covering problem for P and the graph G_s.

Lemma 2. Let $H(V', E')$ be a complete subgraph (clique) of G_s. Then, the sources of V' can be covered by a single maximal orthogonally convex polygon.

Since the sources u and v, such that $u \not\equiv v$, cannot both be covered by an orthogonally convex polygon, a minimum clique cover of G_s (that is, a minimum cardinality set of cliques of G_s with every vertex of G_s belonging to some clique) corresponds exactly to a minimum cover of P by orthogonally convex polygons. Finding a minimum clique cover is NP-hard for general graphs [9], and this formulation of the problem does not give us an algorithm immediately. However, there is an important subclass of graphs (called perfect graphs) for which the minimum clique cover problem can be solved in polynomial time [18]. We will show in a later section that the visibility graph for a class 3 polygon is perfect, implying that the source graph is perfect too.

3. Vertically Convex Polygons and Permutation Graphs

In this section, we show that the visibility graph G of a vertically convex polygon P is a permutation graph [10, 18]. A vertically convex polygon belongs to the class of class 2 polygons. Another kind of class 2 polygons has dents of two orthogonal orientations. It is our belief that these kinds of class 2 polygons also have the property that their visibility graphs are permutation graphs. The proof will be omitted for the sake of brevity. By definition, a vertically convex polygon can have dents of **N** or **S** orientations only.

A *comparability graph* is one which can be obtained from a partially ordered set Q by taking the elements of Q as its vertices and joining two elements if and only if they are comparable. In other words, it is the undirected version of the transitive closure of Q. We now define permutation graphs. Although permutation graphs were originally defined differently, we provide an equivalent definition [18] that is suitable for this paper.

Definition 2. A graph G is a *permutation graph* if both G and its complement are comparability graphs.

Comparability graphs are known to be in the class of graphs called perfect graphs [18]. It follows that permutation graphs are also perfect graphs.

For every region u of P, we pick an arbitrary representative point q_u and argue about this set of points. Recall that, by Observation 7, regions u and v see each other if and only if $q_u \equiv q_v$. For notational simplicity, we denote region u and its representative point q_u by the same name u.

The following lemma shows that visibility in a vertically convex polygon is in some sense a transitive property.

Lemma 3. Let points u, v and w be in P. Let $u \equiv v$, such that u lies to the south of v, and let $v \equiv w$, such that v lies to the south of w. Then, $u \equiv w$.

Proof: Without loss of generality, let the staircase from u to v be northeast. If the staircase from v to w is also northeast, then $u \equiv w$. Assume to the contrary that $u \not\equiv w$. Thus, the staircase from v to w is northwest. By Observation 6, a **W** dent separates u and w, a contradiction.

Q.E.D.

Given the visibility graph G, we construct a directed graph H_G from G as follows. The undirected version of H_G is G. Edge $<u, v>$ of G is oriented from u to v in H_G (denoted by $u \to v$) if point u is to the south of point v. Since edges are directed from south to north, H_G is acyclic. In order to prove that H_G is the transitive closure of some partial order, all we need to show is that if edges $u \to v$ and $v \to w$ exist in H_G, then edge $u \to w$ exists in H_G. By the orientation of edges, u is to the south of v and v is to the south of w. In addition, $u \equiv v$ and $v \equiv w$, implying, by Lemma 3, that $u \equiv w$. Moreover, u is to the south of w, and hence, edge $u \to w$ is in H_G. We have thus shown the following lemma.

Lemma 4. The visibility graph G of a vertically convex polygon P is a comparability graph.

We now show that the complement graph G^c is also a comparability graph. As before, we construct a directed graph H_{G^c} from G^c, such that the undirected version of H_{G^c} is G^c. The orientation of an edge $<u, v>$ is from u to v if u is to the west of v. As before, H_{G^c} is acyclic. Now, we show transitivity, i.e., if $u \to v$ is in H_{G^c} and $v \to w$ is in H_{G^c}, then $u \to w$ is in H_{G^c}. We first need the following observation.

Observation 8. Let D be a dent in a vertically convex polygon P such that $o(D) = \mathrm{S}$ (N). Then every point in $B_l(D)$ is to the west (east, resp.) of every point in $B_r(D)$.

Proof: Trivial.

Q.E.D.

Lemma 5. If $u \to v$ and $v \to w$ are edges in H_{G^c}, then $u \to w$ is an edge in H_{G^c}.

Proof: Since $<u, v> \in E(G^c)$, $u \not\equiv v$. Thus, dent D_u separates u and v. Without loss of generality, let $o(D_u) = \mathrm{S}$. By Observation 8, $u \in B_l(D_u)$ and $v \in B_r(D_u)$. Similarly, we can argue that there exists a dent D_v separates v and w.

Case 1: $o(D_v) = \mathbf{S}$.

By Observation 8, $v \in B_l(D_v)$ and $w \in B_r(D_v)$. If $\vec{D}_v$ is to the north of $\vec{D}_u$, then $u \in B_l(D_v)$, and hence, $u \not\equiv w$. If $\vec{D}_u$ is to the north of $\vec{D}_v$, then $w \in B_r(D_u)$, and hence, $u \not\equiv w$. Thus, $u \rightarrow w$ is an edge in H_{G^c}.

Case 2: $o(D_v) = \mathbf{N}$.

We have, by Observation 8, that $v \in B_r(D_v)$ and $w \in B_l(D_v)$. Suppose there exists another dent D'_u separating u and v, such that $o(D'_u) = \mathbf{N}$, then by the proof of Case 1, we are done. Therefore, assume that there is no **N** dent separating u and v. Similarly, assume that there is no **S** dent D'_v separating v and w. This implies that v is to the south of $\vec{D}_u$, and that $\vec{D}_v$ is to the south of $\vec{D}_u$. Thus, $u \in B_r(D_v)$, implying that $u \not\equiv w$. We have now shown that $u \rightarrow w$ is an edge in H_{G^c}.

Q.E.D.

The preceding arguments have established the following lemma.

Lemma 6. G^c, the complement of the visibility graph G, is a comparability graph.

Now, Lemmata 4 and 6 together imply Theorem 1.

Theorem 1. The visibility graph G of a vertically convex polygon is a permutation graph.

4. Polygons with Three Dent Orientations

In this section, we state our main results concerning orthogonal polygons with three dent orientations (class 3 polygons). We first assert that the visibility graph of a class 3 polygon is perfect. In a perfect graph G, the size of a minimum clique cover of every subgraph G' is equal to the size of a maximum independent set of G'. We then state the duality relationship for class 3 polygons. We first need the following definition.

Definition 3. A graph G is *weakly triangulated* if neither G nor G^c, the complement of G contain induced cycles of length greater than four.

The following theorem, proved by Hayward [13], will be useful.

Theorem 2. (Hayward) Weakly triangulated graphs are perfect.

We now state the following theorem, the proof of which is contained in the next three sections.

Theorem 3. The visibility graph of a class 3 orthogonal polygon P is weakly triangulated.

Theorem 2, together with Hayward's theorem, provides us with the following duality relationship:

Corollary 1. (The Duality Relationship) The minimum number of orthogonally convex polygons needed to cover an orthogonal polygon P with at most three dent orientations is equal to the maximum number of points of P, no two of which can be contained together in an orthogonally convex covering polygon.

Since the source graph is an induced subgraph of the visibility graph, by Theorem 2, it must be perfect. Moreover, the induced subgraphs of a weakly triangulated graph are also weakly triangulated. By Lemmata 1 and 2, we have that a minimum clique cover of the source graph corresponds to a minimum cover of P by maximal OCP's. Recently, Hayward and Hoang [14] have devised an $O(n^5)$ algorithm to compute the minimum clique cover of a weakly triangulated graph. This is a significant improvement over the Ellipsoid method generally used for perfect graphs. The main advantage of working with source graphs is that, in a class 3 polygon the number of sources must be $O(n)$. It is not very hard to see that a minimal cover by OCP's can be efficiently constructed given the minimum clique cover of G_s. Thus, we have a polynomial algorithm for the convex cover problem for class 3 polygons.

5. The Crossing Lemma.

In this section we prove a technical lemma, called the Crossing Lemma, which will be required in the proof of Theorem 3. Two staircase paths are said to *cross* if they meet at some point. We will only be considering pairs of staircase paths which cross and have distinct endpoints. Observations 9 and 10 are concerned with pairs of crossing staircase paths which are of the same type and of different types, respectively. These observations are valid in all their reflection and rotation symmetric versions.

Observation 9. Let points $u,v \in P$ be such that $u \equiv v$ and u lies southwest of v. Let points $u',v' \in P$ be such that $u' \equiv v'$ and u' lies southwest of v'. If staircase paths $s(u,v)$ and $s(u',v')$ cross then $u' \equiv v$ and $u \equiv v'$. Moreover, u and u' lie southwest of v' and v, respectively.

Proof: $s(u,v)$ travels southwest from v to u. $s(u',v')$ travels southwest from v' to u', and meets $s(u,v)$ at some point, say p. Thus, we can take $s(u,v)$ from v to p and then take $s(u',v')$ from p to u', establishing a southwest staircase from v to u'. Similarly, there is a southwest staircase from v' to u.

Q.E.D.

Observation 10. Let points $u,v \in P$ be such that $u \equiv v$ and u lies southwest of v. Let points $u',v' \in P$ be such that $u' \equiv v'$ and u' lies southeast of v'. If $s(u,v)$ and $s(u',v')$ cross and $u \not\equiv u'$ then a S dent separates u from u'.

Proof: As before, let $s(u,v)$ meet $s(u',v')$ at p. Now, p sees u to its southwest and u' to its southeast. By Observation 6, there is a S dent separating u and u'.

Q.E.D.

Let $G'(V',E')$ be a subgraph of the visibility graph $G(V,E)$. For every vertex $v \in V'$, fix a point (also called v) which lies in the region of P corresponding to v. Note that this notation should cause no confusion, by Observation 7. For every edge $<u,v> \in E'$, fix a staircase path

$s(u,v)$ from point u to point v. We call this collection of points and staircase paths an *instantiation* of the subgraph G'.

Let C be k-cycle in the graph G such that $V(C)=\{v_0,v_1, \cdots ,v_{k-1}\}$ and $<v_i,v_{i+1}> \in E(C)$, for each i (all indices here and in the rest of the paper are modulo k). Consider C', an instantiation of C.

Lemma 7. (*Crossing Lemma*) If C is an induced cycle of G (i.e. C has no chords) then some pair of non-adjacent staircase paths must cross in C'.

Proof: Assume to the contrary, that none of the non-adjacent pairs of staircase paths cross. Let v_0 be the northernmost point of $V(C)$. Let l denote the line segment connecting the neighbors of v_0 in C, viz. v_1 and v_{k-1}. Let R denote the region in P which is enclosed by $s(v_{k-1},v_0)$, $s(v_0,v_1)$ and l (see Figure 7). Since $v_{k-1} \not\equiv v_1$, an edge of the boundary of P must intersect l. Thus, there are points of the boundary of P in R. Let p denote a northernmost point in R which lies on the boundary of P. Since P is a simple polygon and the set of staircase paths of C' forms a closed, non-crossing path, the region enclosed by the closed path cannot include p. Thus, some point, say $v_i \in \{v_2,v_3, \cdots v_{k-2}\}$ must lie in R to the north of p. Then, $v_i \equiv v_0$, implying that v_i has a chord to v_0, and we have a contradiction.

Q.E.D.

Lemma 8. Suppose C is an induced k-cycle of G, where $k>4$. Any pair of crossing staircase paths in C', $s(v_i,v_{i+1})$ and $s(v_j,v_{j+1})$ say, must be of different types if they have distinct endpoints.

Proof: Let $i,j \in \{0, \cdots ,k-1\}$ be such that $|i-j|>1 \ mod\ k$. Assume to the contrary that $s(v_i,v_{i+1})$ and $s(v_j,v_{j+1})$ cross and are of the same type, say type I. Assume, without loss of generality, that v_i and v_j lie to the southeast of v_{i+1} and v_{j+1}, respectively. By Observation 10, we have that $v_i \equiv v_{j+1}$ and $v_j \equiv v_{i+1}$. Thus, $<v_i , v_{j+1}> \in E(C)$ and $<v_j , v_{i+1}> \in E(C)$. Since $k>4$, one of these edges will cause a chord in C and give us a contradiction.

Q.E.D.

6. Induced Cycles in the Visibility Graph.

In this section we prove the first part of Theorem 3. We show that for a class 3 polygon P the visibility graph G has no induced k-cycles, for $k>4$.

Lemma 9. Suppose P is a class 3 polygon (with **W**, **S** and **E** dent orientations). Then, the visibility graph, G, cannot have induced k-cycles where $k>4$.

Proof: Assume to the contrary that C is an induced k-cycle in G. Let $V(C)=\{v_0,v_1, \cdots v_{k-1}\}$ denote the set of vertices of C in the cyclic order. Let C' be an instantiation of C. By the Crossing Lemma, some pair of non-adjacent staircase paths, $s(v_i,v_{i+1})$ and $s(v_j,v_{j+1})$, must cross. By Lemma 8, the two staircase paths must be of different types. Assume, without loss of generality, that $s(v_i,v_{i+1})$ is of type I and $s(v_j,v_{j+1})$ is of type II. We may further assume that v_i lies to the southeast of v_{i+1} and that v_{j+1} lies to the southwest of v_j. Since P has only three dent orientations: **W**, **S** and **E**, we assert that $v_{i+1} \equiv v_j$. If this were not the case, then, by Observation 10, there would be a **N** dent separating the two points, giving us a contradiction. To prevent chords, we have that $v_i \not\equiv v_{j+1}$, $v_i \not\equiv v_j$ and $v_{i+1} \not\equiv v_{j+1}$. By Observation 10, there must be a **W** dent, D_W, separating v_i from v_j. Similarly, there must be a **S** dent, D_S, separating v_i from v_{j+1} and an **E** dent, D_E, separating v_{i+1} from v_{j+1} (see Figure 8).

First, note that $\vec{D}_E$ must lie in $A(D_W)$ and $\vec{D}_W$ must lie in $A(D_E)$. This implies that $\vec{D}_E$ must lie to the east of $\vec{D}_W$. We now know that $v_{j+1} \in B_r(D_W)$ and $v_i \in B_l(D_E)$. Moreover, it is clear that $v_i \notin B_r(D_W)$ and $v_{j+1} \notin B_l(D_E)$.

Consider the set of staircase paths which remains after the removal of $s(v_{i+1}, v_j)$ from the instantiation of C. In the sequence of paths from v_{j+1} to v_i, let $s(v_k, v_{k+1})$ be the first staircase path to intersect with $\vec{D}_W$. We will assume that $k+1 \neq i$; otherwise, $k \neq j+1$ and we can argue symmetrically. Let the intersection point of $s(v_k, v_{k+1})$ and $\vec{D}_W$ be called p (see Figure 9). Note that $v_k \in B_r(D_W)$ and $v_{k+1} \in A(D_W)$.

Since v_{i+1} sees v_i to its southeast, and p is to the south of the point of intersection of $\vec{D}_W$ and $s(v_i, v_{i+1})$, there is a northwest staircase path from p to v_{i+1}. If the staircase path $s(v_k, v_{k+1})$ is of type I (that is, southeast from v_k to v_{k+1}), then v_{k+1} has a northwest staircase path from v_{k+1} to p, and thence to v_{i+1}. Since $k+1 \neq i$ and $k+1 \neq j$, we have a chord from v_{k+1} to v_{i+1}, a contradiction. If, on the other hand, $s(v_k, v_{k+1})$ is of type II, then again $v_{k+1} \equiv v_{i+1}$; otherwise, there must be a **N** dent separating v_{k+1} from v_{i+1}, by Observation 6. Since **N** dents do not occur in P we again have a chord from v_{k+1} to v_{i+1}, a contradiction.

Q.E.D.

7. Induced Cycles in the Complement of the Visibility Graph

In this section, we establish the other part of the proof of Theorem 3, namely that the complement graph G^c of the visibility graph cannot have any induced cycles of length 5 or more. The following definitions and claims will prove useful in establishing this result. Let C be a k-cycle in the graph G^c, where $k \geq 5$. Let $V(C) = \{v_0, v_1, \cdots, v_{k-1}\}$ denote the vertices of C in a cyclic order. In other words, for each $v_a \in V(C)$ we have $\{v_{a-1}, v_{a+1}\} \subseteq N(v_a, G^c)$, the neighbor set of v_a in G^c. Recall that $v_b \in N(v_a, G^c)$ if and only if $v_a \not\equiv v_b$. Hence, C is an induced k-cycle if and only if for each v_a it is the case that $N(v_a, G^c) \cap V(C) = \{v_{a-1}, v_{a+1}\}$.

From Observation 4, we have that if $v_a \not\equiv v_b$ then there must be a dent which separates the two vertices. Let D_a denote the dent which separates the vertices v_a and v_{a+1} which are consecutive in the cycle C (recall that all indices are modulo k). Thus, the k-cycle C determines a sequence of k dents corresponding to the k cycle edges in C (see Figure 10). We first claim that if any two of these k dents are distinct and of the same orientation then there exists a chord for the cycle C. Assume, without loss of generality, that the three dent orientations in P are **N**, **W** and **S**.

Claim 1. Suppose the dents D_a and D_b are distinct but of the same orientation, then C cannot be an induced cycle.

Proof: Suppose, for the moment, that $\{a, a+1\} \cap \{b, b+1\} = \emptyset$. We will relax this condition later. Assume, without loss of generality, that $o(D_a) = o(D_b) = S$ and the dent line $\vec{D}_a$ lies to the north of the the dent line $\vec{D}_b$. We can always renumber the vertices of C to ensure that $a=0$, which would imply that $2 \le b$ and $b+1 \le k-1$. Thus, we are assured that the edges $<v_{a+1}, v_{b+1}>$ and $<v_a, v_b>$ cannot be present in G^c since they would be chords for the cycle C.

Since $k \ge 5$, there must be a vertex in $V(C)$ which is adjacent to neither v_a nor v_{a+1} in G^c, otherwise there would be chords in C. This would imply that there is a vertex which is adjacent to both v_a and v_{a+1} in the visibility graph G. By Observation 5, we have that both v_a and v_{a+1} must lie to the south of the dent line $\vec{D}_a$. A similar argument also shows that both v_b and v_{b+1} must lie to the south of the dent line $\vec{D}_b$, and, thus, also to the south of the dent line $\vec{D}_a$.

We now assert that $v_b \notin A(D_a)$. Suppose v_b did lie in the zone $A(D_a)$. We know that the edge $<v_a, v_b>$ is present in G, therefore $v_a \equiv v_b$. We also know that v_a is in one of the B zones of D_a. By Observation 5, v_b must be north of $\vec{D}_a$ and, hence, of $\vec{D}_b$ also, which is a contradiction. A similar argument shows that $v_{b+1} \notin A(D_a)$ since the edge $<v_{a+1}, v_{b+1}>$ must be present in the visibility graph G and v_{b+1} lies to the south of the dent line $\vec{D}_a$.

Assume, without loss of generality, that $v_a \in B_l(D_a)$ and that $v_{a+1} \in B_r(D_a)$. It then follows from the above argument that $v_b \in B_l(D_a)$ and $v_{b+1} \in B_r(D_a)$, since there are staircase paths from v_a to v_b and from v_{a+1} to v_{b+1} (see Figure 11). Recall that if there is a staircase path between two points inside the polygon then they cannot lie in different B zones of some dent.

Finally, we observe that the dent line $\vec{D}_b$ must lie entirely in the zone $A(D_a)$. Suppose this were not the case, then $\vec{D}_b$ must lie either in $B_l(D_a)$ or in $B_r(D_a)$. Clearly, if $\vec{D}_b$ were inside the zone $B_l(D_a)$ then both $B_l(D_b)$ and $B_r(D_b)$ would also lie inside the zone $B_l(D_a)$. This cannot be the case since $v_{b+1} \in B_r(D_a)$ and $v_{b+1} \notin A(D_b)$. A similar argument establishes that $\vec{D}_b$ cannot lie in $B_r(D_a)$.

Consider a path which starts from $v_b \in B_l(D_a)$ and first goes to the dent line $\vec{D}_a$. It then follows the dent line until it reaches the boundary of the zone $B_r(D_a)$. Finally, it traverses this zone until reaches the point $v_{b+1} \in B_r(D_a)$. Clearly, this path never crosses the dent line $\vec{D}_b$ since it never entered the zone $A(D_a)$. In other words, we have exhibited a path between v_b and v_{b+1} which lies entirely inside the polygon P and never crosses the dent line of the dent separating the two points (see Figure 11). This contradicts the definition of the dent line and completes the proof for the case where $\{a, a+1\} \cap \{b, b+1\} = \emptyset$.

It is easy to extend the above argument to the case where $\{a, a+1\} \cap \{b, b+1\} \neq \emptyset$. First, observe that the only case in which the intersection is non-empty is where $a+1 = b$. In this case, we can argue that the edge $<v_a, v_{b+1}>$ must be absent in G^c since it would be a chord for C. Observe that the edge $<v_a, v_{b+1}>$ cannot be a cycle edge for C since the length of the cycle is greater than 3. Thus, we have that $v_a \equiv v_{b+1}$. As before, we can argue that v_a and v_b must lie to the south of the dent line $\vec{D}_a$ while v_b and v_{b+1} lie to the south of the dent line $\vec{D}_b$. Again, we only consider the case where the dent line $\vec{D}_b$ lies to the south of the dent line $\vec{D}_a$. It is now easy to see that v_{b+1} cannot lie in the zone $A(D_a)$.

Without loss of generality, we only consider the situation where both v_a and v_{b+1} lie in $B_l(D_a)$ while v_b lies in the zone $B_r(D_a)$. Note that v_a and v_{b+1} cannot lie in different B zones of the dent D_a since $v_a \equiv v_{b+1}$. Again, it is easy to verify that the dent line $\vec{D}_b$ must lie entirely in the zone $A(D_a)$. Thus, we are now able to construct a path from v_{b+1} to v_b which never crosses the dent line $\vec{D}_b$ (see Figure 12). This contradicts the definition of the dent line and completes the proof.

Q.E.D.

We now show that if a dent D corresponds to two non-adjacent edges in the cycle C then there must be a chord in the cycle C.

Claim 2. Suppose $\{a,a+1\} \cap \{b,b+1\} = \emptyset$ then it cannot be the case that $D_a = D_b$ unless C has a chord.

Proof: Suppose to the contrary that $D_a = D_b = D$ and that C has no chords in G^c. Thus, the dent D separates v_a from v_{a+1} and v_b from v_{b+1}. Assume, without loss of generality, that $v_a \in B_l(D)$ and $v_{a+1} \in B_r(D)$. Similarly, assume that $v_b \in B_l(D)$ and $v_{b+1} \in B_r(D)$. Now, the dent D must also separate v_a from v_{b+1} and v_b from v_{a+1} (see Figure 13). Thus, the two edges $<v_a, v_{b+1}>$ and $<v_b, v_{a+1}>$ are not present in G and they are both present in G^c. One of these two edges will always be a chord for C when $k \geq 5$. Since C is an induced cycle (by assumption) this cannot happen.

Q.E.D.

We are now ready to prove the following lemma.

Lemma 10. If G is the visibility graph of a class 3 polygon P then G^c cannot have an induced cycle of length 5 or more.

Proof: Assume to the contrary that C is an induced k-cycle in the graph G^c. It is clear that $k \geq 6$ since an induced 5-cycle in G^c would imply the existence of an induced 5-cycle in the visibility graph G (see Figure 14). From Claim 1, it is clear that in the multi-set $D(C) = \{D_0, D_1, \cdots, D_{k-1}\}$, there cannot be more than one dent of each of the three allowed orientations. Moreover, Claim 2 states that the same dent cannot correspond to two non-adjacent edges of C. Thus, we cannot have an induced k-cycle where $k \geq 7$ since only three dent orientations are permitted. We now complete the proof of the lemma by showing that, given Claims 1 and 2, even induced 6-cycles are not possible.

Consider the case where $k=6$. Given Claims 1 and 2, it is clear that we can renumber the vertices of C to ensure that $D_0 = D_1 = D_N$, $D_2 = D_3 = D_W$ and $D_4 = D_5 = D_S$, where $o(D_N) = \mathbf{N}$, $o(D_W) = \mathbf{W}$ and $o(D_S) = \mathbf{S}$ (see Figure 15). We will show that the edge $<v_5, v_1>$ is not present in G and, thus, is a chord for C.

There must exist vertices in C which are adjacent (in G) to both v_5 and v_0, e.g. v_3. Observation 5 then implies that the points v_5 and v_0 must lie to the south of the the dent line $\vec{D}_S$. Similarly, we can show that the points v_0 and v_1 must both lie to the north of the dent line $\vec{D}_N$. This can only happen if the dent line $\vec{D}_N$ lies to the south of the dent line $\vec{D}_S$ (see Figure 16). Now the point v_3 must have staircase paths to both v_0 and v_1, otherwise the chords $<v_3, v_0>$ or $<v_3, v_1>$ will be present in G^c. Thus, v_3 must lie to the south of the dent line $\vec{D}_N$. Similarly, points in v_3 must have staircase paths to points in both v_5 and v_0 and, thus, v_3 must lie to the north of the dent line $\vec{D}_S$. Since the dent line $\vec{D}_N$ lies to the south of the dent line $\vec{D}_S$ this gives a contradiction.

Q.E.D.

8. Orthogonal Polygons with Four Dent Orientations

In this section, we demonstrate arbitrarily large induced odd cycles in the source graph of an orthogonal polygon with four dent orientations. This would show that the source graph, and hence, the visibility graphof a class 4 polygon, is not perfect [10, 18], and also imply that the duality relationship of Corollary 1 does not hold for class 4 polygons.

Consider the polygon P_5, shown in Figure 17. There are exactly five sources, but the source graph is a 5-cycle without chords, which is not perfect, and hence, the duality relationship fails to hold.

Note that $B_l(D_2) \subset B_l(D_1)$. If we now modify $B_l(D_2)$ to obtain the polygon P_7, shown in Figure 18, we find that the source graph is a 7-cycle without chords. A similar construction to P_7 would give a polygon whose source graph is a 9-cycle with no chords, and so on to obtain arbitrarily large induced odd cycles.

9. Further Work

The main contribution of this paper has been the demonstration of the intimate connection between orthogonal polygon covering and classes of perfect graphs, and deriving the duality relationship of section 4. The main tool of our analysis has been the visibility graph for regions inside an orthogonal polygon. We have demonstrated certain interesting combinatorial properties of these kinds of graphs. It is our belief that a careful examination of the combinatorial structure of different kinds of visibility graphs may lead to the solution of other open problems in computational geometry. In particular, we feel that our approach might prove fruitful in solving the open problems listed below.

(1) To find a duality relationship and/or a polynomial algorithm for class 4 polygons. Is the source graph a member of some other class of graphs for which the minimum clique cover problem can be solved in polynomial time?

(2) Will similar techniques work for the problem of covering orthogonal polygons with a minimum number of rectangles? We know that Chvatal's conjecture is false in general for orthogonal polygons without holes. Culberson and Reckhow [7] describe a scheme to obtain

source graphs for this case. Is the source graph a member of some other class of graphs for which the minimum clique cover problem can be solved in polynomial time?

Acknowledgements

The authors wish to thank Joe Culberson and Bob Reckhow for sending them accounts of their work.

References

1. A. Aggarwal, "The Art Gallery Theorem: Its Variations, Applications and Algorithmic Aspects," *Ph.D. Thesis*, Dept. of Elecl. Engg. and Computer Science, Johns Hopkins University, 1984.
2. Michael O. Albertson and Claire J. O'Keefe, "Covering Regions with Squares," *SIAM J. Alg. Disc. Meth.*, vol. 2, pp. 240-243, 1981.
3. A. Boucher, "It's hard to color antirectangles," *SIAM J. Alg. Disc. Meth.*, vol. 5, 1984.
4. S. Chaiken, D. J. Kleitman, M. Saks, and J. Shearer, "Covering Regions by Rectangles," *SIAM J. Alg. Disc. Meth.*, vol. 2, pp. 394-410, Dec. 1981.
5. B. Chazelle and D. Dobkin, "Decomposing a Polygon into its Convex Parts," *Proc. of the 11th. Annual ACM Symposium on Theory of Computing*, pp. 38-48, 1979.
6. B. Chazelle, "Computational Geometry and Convexity," *Ph.D. Thesis*, Dept. of Computer Science, Carnagie-Mellon University, Pittsburgh, Pa, July, 1980.
7. J. Culberson and R. Reckhow, "Dent Diagrams: A Unified Approach to Polygon Covering Problems," *Tech. Rep. TR 87-14, Dept. of Computing Science, Univ. of Alberta*, July, 1987.
8. D. S. Franzblau and D. J. Kleitman, "An Algorithm for Constructing Regions with Rectangles: Independence and Minimum Generating Sets for Collection of Intervals," *Proc. Sixteenth Annual ACM Symposium on Theory of Computing*, pp. 167-174, 1984.
9. M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, pp. 53-56, W. H. Freeman and Company, San Francisco, 1979.
10. M. C. Golumbic, *Algorithmic Graph Theory and Perfect Graphs*, Academic Press, New York, 1980.
11. M. Grötschel, L. Lov´asz, and A. Schrijver, "The Ellipsoid Method and its Consequences on Combinatorial Optimization," *Combinatorica 1*, pp. 169-197, 1981.
12. E. Gy´ori, "A Minimax Theorem of New Type," *to appear.*
13. R. B. Hayward, "Weakly Triangulated Graphs," *J. of Comb. Theory, Series B 39*, pp. 200-209, 1985.
14. R. B. Hayward and C. Hoang, *(Personal Communication)*, Dec. 1987.
15. J. M. Keil and J. R. Sack, "Minimum Decopmpositions of Polygonal Objects," *Computational Geometry (G. T. Toussaint, ed.)*, North Holland, 1985.

16. J. M. Keil, "Decomposing a Polygon into Simpler Components," *SIAM Journal on Computing*, vol. 14, 4, pp. 799-817, 1985.

17. J. M. Keil, "Minimally Covering a Horizontally Convex Orthogonal Polygon," *Proc. 2nd. Annual ACM Symp. on Computational Geometry*, pp. 43-51, Yorktown Heights, New York, June 1986.

18. L. Lov´asz, "Perfect Graphs," *Selected Topics in Graph Theory, 2*, pp. 55-85, Academic Press, Inc., London, 1983. Lowell W. Beineke and Robin J. Wilson, eds.

19. R. Motwani, A. Raghunathan, and H. Saran, "Covering Orthogonal Polygons with Star Polygons: The Perfect Graph Approach," *submitted to the Fourth Annual ACM Symp. on Computational Geometry*.

20. J. O'Rourke, *Art Gallery Theorems and Algorithms*, Oxford University Press, New York, 1987.

21. R. Reckhow and J. Culberson, "Covering a Simple Orthogonal Polygon with a Minimum Number of Orthogonally Convex Polygons," *Proc. 3rd. Annual ACM Symposium on Computational Geometry*, pp. 268-277, June 1987.

22. R. Reckhow, "Covering Orthogonally Convex Polygons with Three Orientations of Dents," *Tech. Rep. TR 87-17*, Dept. of Computing Science, Univ. of Alberta, Aug. 1987.

23. Michael Saks, "A Class of Perfect Graphs associated with Planar Rectilinear Regions," *SIAM J. Alg. Disc. Meth.*, vol. 3, pp. 330-342, 1982.

24. James B. Shearer, "A Class of Perfect Graphs," *SIAM J. Alg. Disc. Meth.*, vol. 3, pp. 281-284, 1982.

25. J. O'Rourke and K. Supowit, "Some NP-Hard Polygon Decomposition Problems," *IEEE Trans. Info. Theory*, vol. IT-29(2), pp. 181-190, 1983.

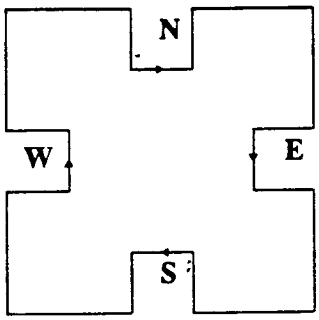

Fig. 1 Orientation of dents

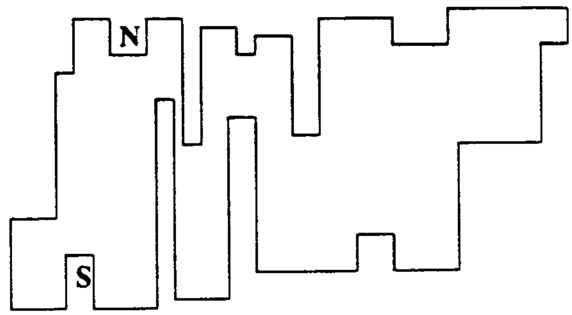

Fig. 2 A Vertically Convex Polygon

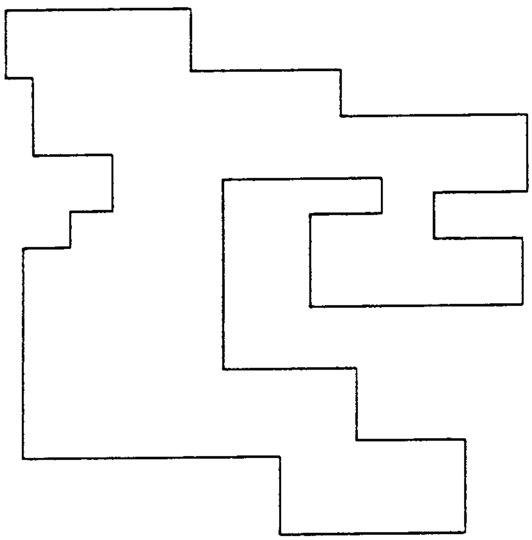
Fig. 3 A Class 3 Polygon (No **N** dents)

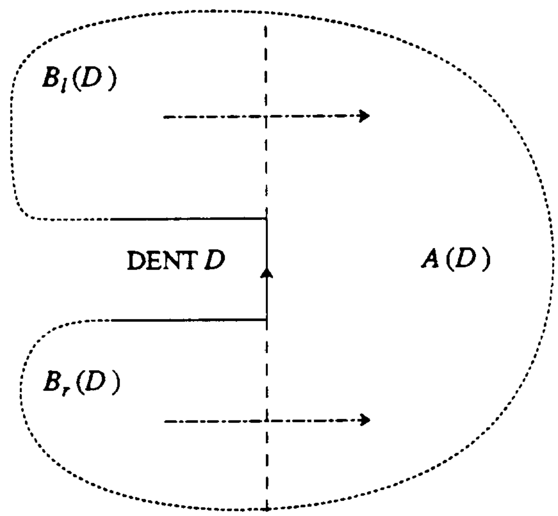

Fig. 4 Dent lines and zones

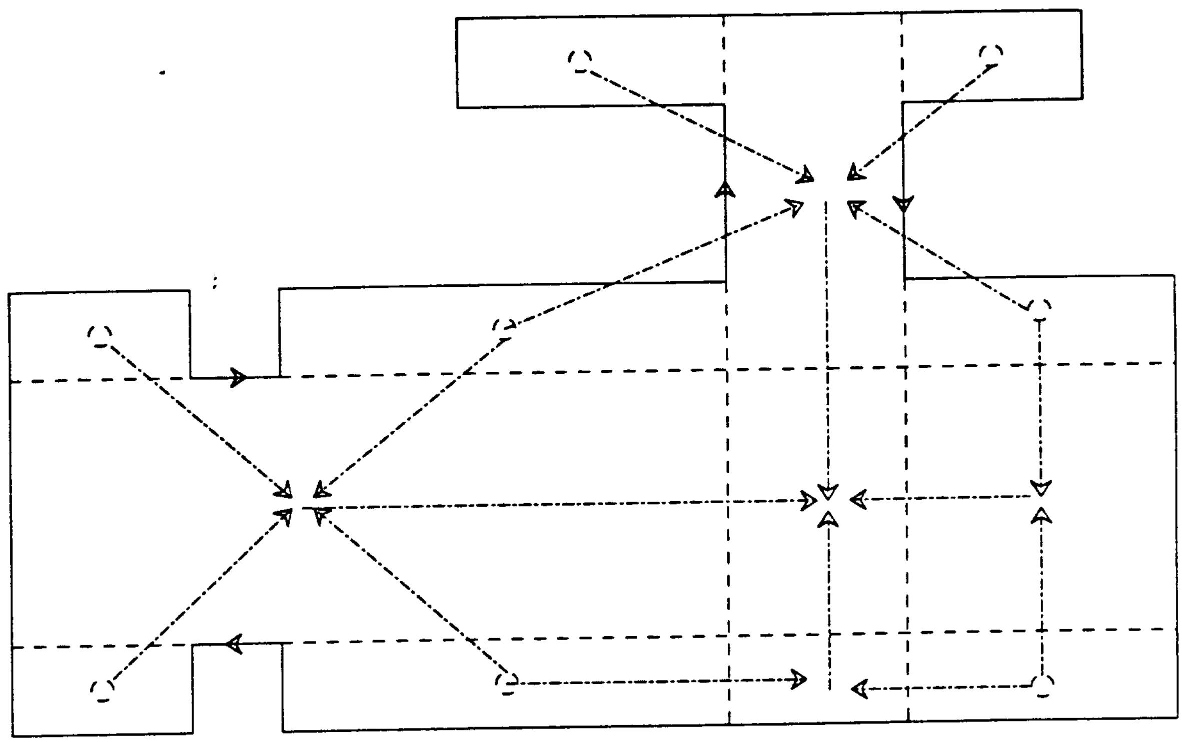

Figure 5: Identifying sources

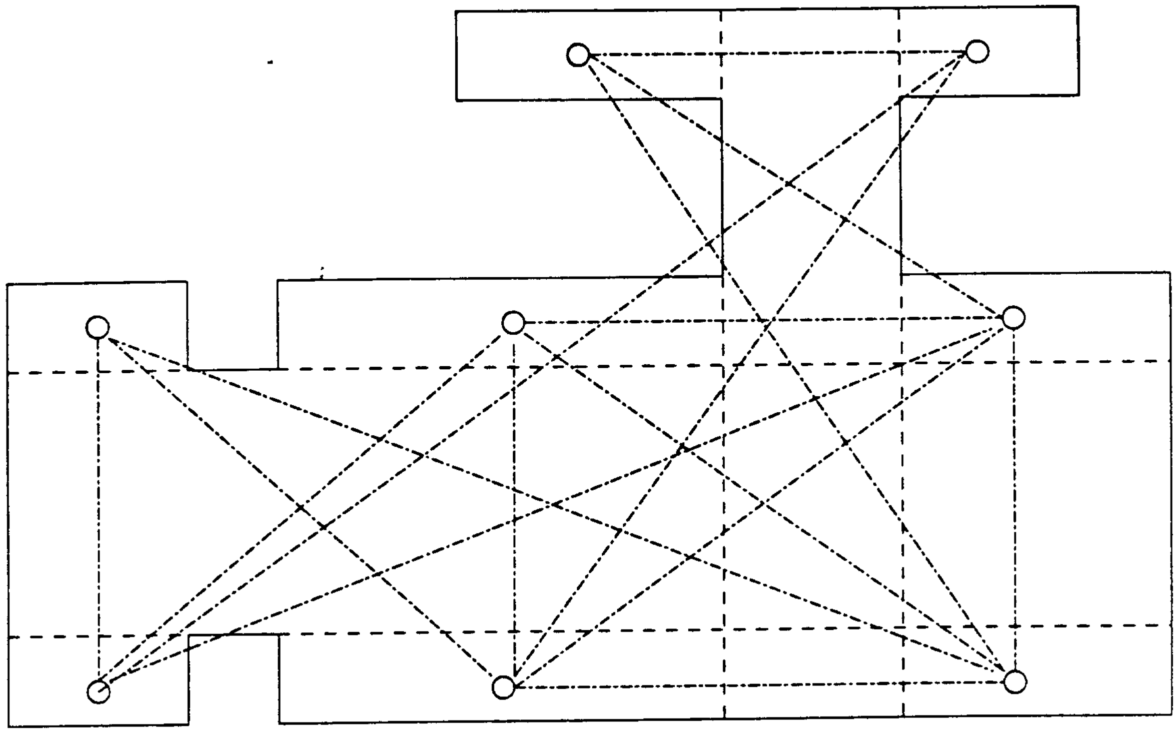

Figure 6: The Source Graph

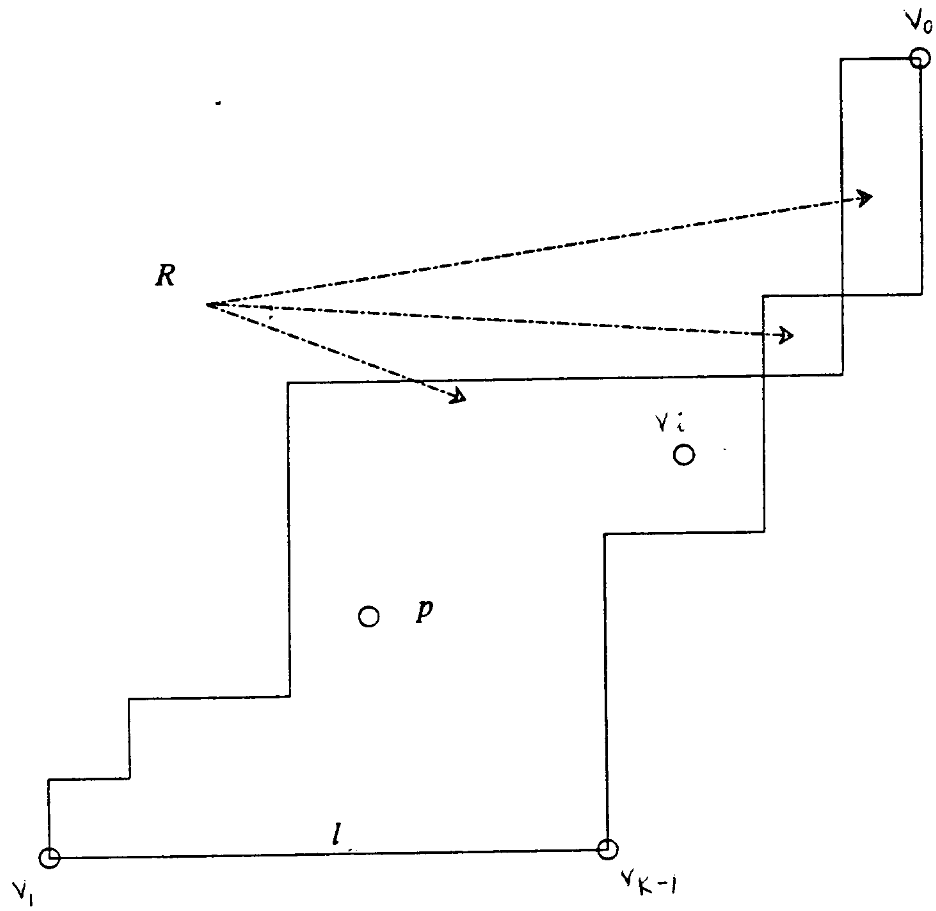

Figure 7: Proof of Lemma 7

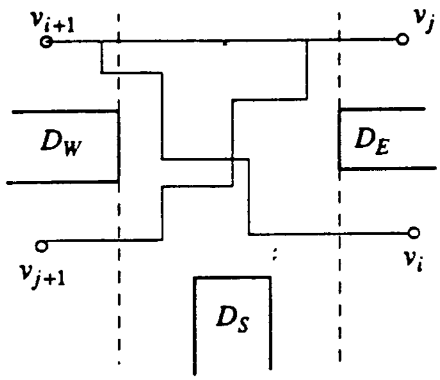

Fig. 8: Proof of Lemma 9

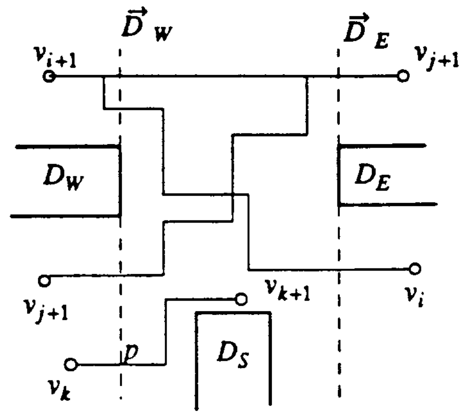

Fig. 9 Proof of Lemma 9

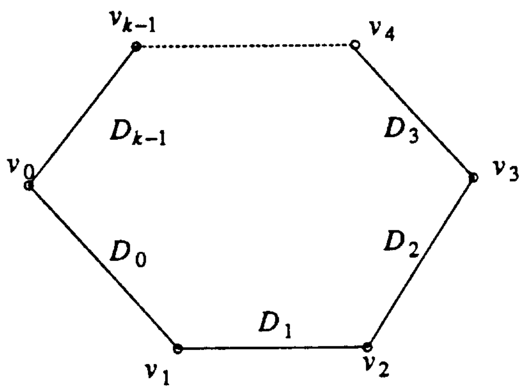

Fig. 10 A k-cycle in G^c

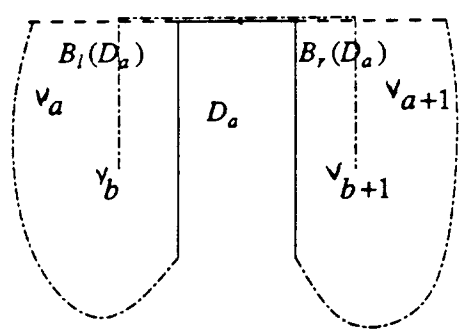

Fig. 11 Proof of Claim 1

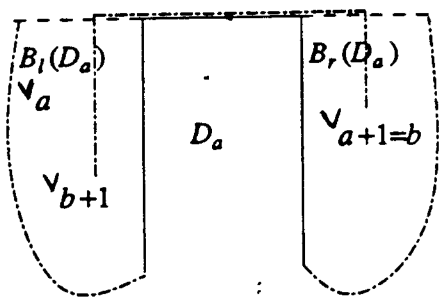

Fig. 12 Proof of Claim 1

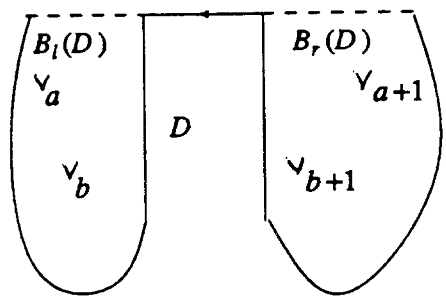

Fig. 13 Proof of Claim 2

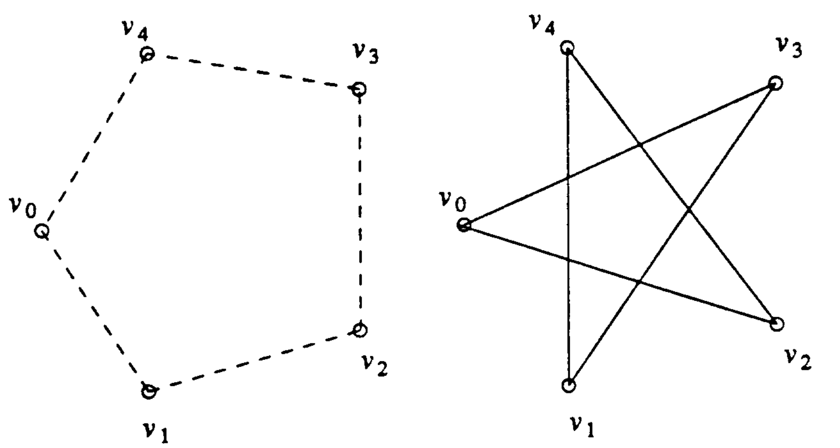

Fig. 14 A 5-cycle in G_c and its complement in G

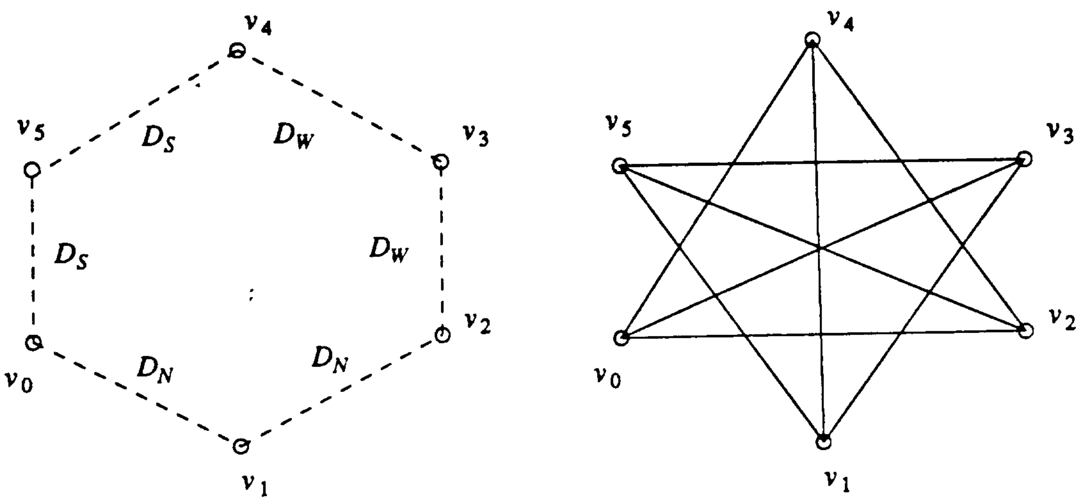

Fig. 15 Proof of Lemma 10

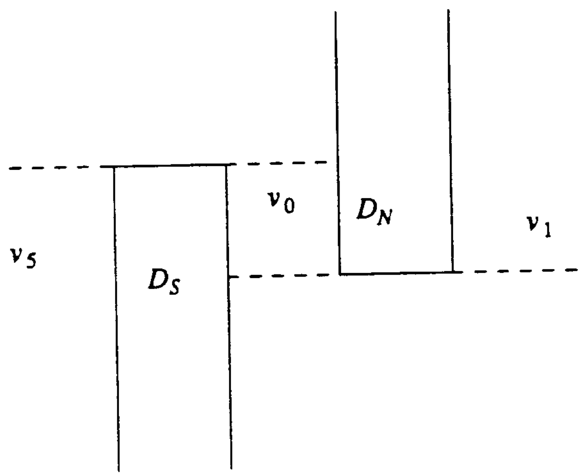

Fig. 16 Proof of Lemma 10

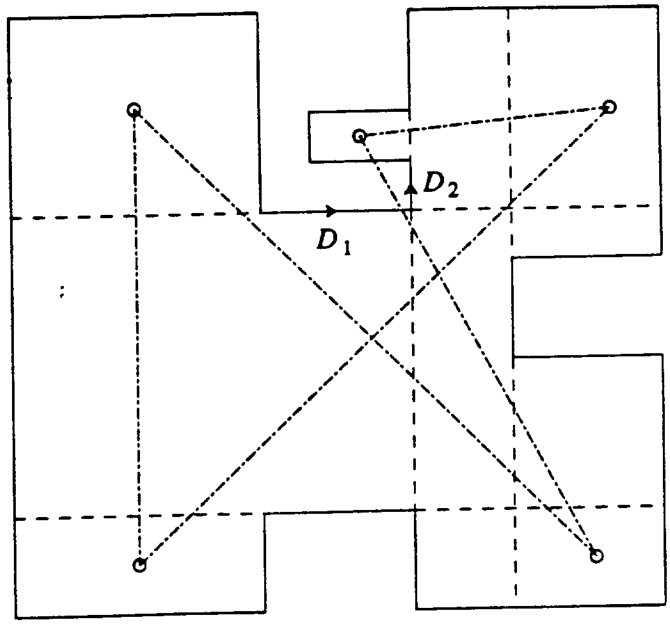

Fig. 17. P_5, the 5-cycle

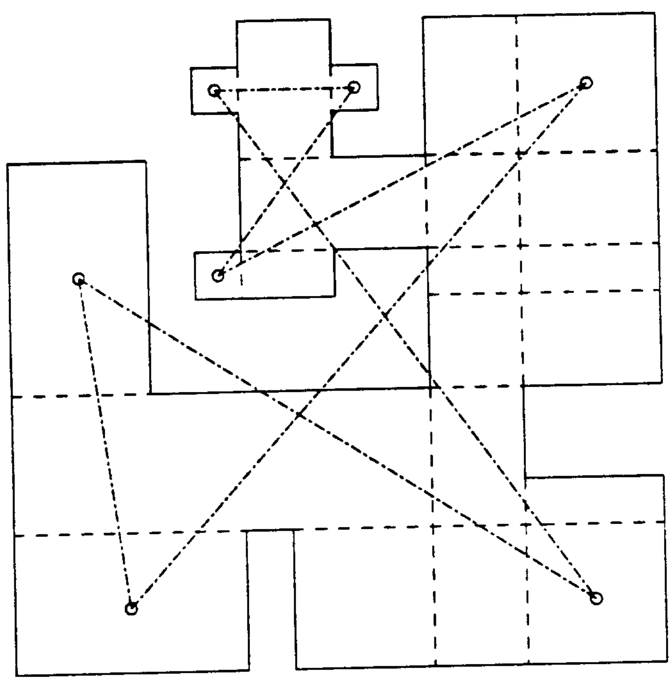

Fig. 18. P_7, the 7-cycle