Discrete Comput Geom 28:449–465 (2002)
DOI: 10.1007/s00454-002-2881-6

Discrete & Computational
Geometry

Visibility Graphs and Oriented Matroids*

James Abello[1] and Krishna Kumar

[1]Information Visualization Research, Shannon Laboratories, AT&T Labs-Research,
180 Park Avenue, Florham Park, NJ 07932, USA
abelloj@optonline.net

Abstract. We describe a set of necessary conditions for a given graph to be the visibility graph of a simple polygon. For every graph satisfying these conditions we show that a uniform rank 3 oriented matroid can be constructed in polynomial time, which if affinely coordinatizable yields a simple polygon whose visibility graph is isomorphic to the given graph.

1. Introduction

Visibility graphs are fundamental structures in computational geometry. They find applications in areas such as graphics [15], [23] and robotics [18], yet very little is known about their combinatorial structure. This paper addresses the question of characterizing internal visibility graphs of simple plane polygons, henceforth simply called visibility graphs. Two vertices of a simple polygon P are called visible if the open line segment between them is either a boundary edge of P, or is completely contained in the interior of the polygon. Note that in this setting, two vertices are considered to be invisible if the open line segment between them passes through a third vertex of the polygon. The visibility graph of a polygon is the graph whose vertices correspond to the vertices of the polygon and edges correspond to visible pairs of vertices in the polygon. From the

* Historical Note: The main result of this paper was presented at the Graph Drawing Conference held in Princeton in 1994 and an extended abstract version appeared in the corresponding proceedings [5]. Based on that extended abstract, Streinu and O'Rourke [21] developed pseudo-visibility graphs which turned out to be precisely the same class of graphs characterized here. Subsequently, Ghosh proposed in [13] a new necessary condition for visibility graphs which we conjecture now produces again the same class of graphs introduced in this work, namely, *quasi-persistent* graphs whose shortest paths satisfy three new conditions that we call local-separability, path symmetry and path consistency. On a more personal level, at the time of this writing the geographical coordinates of the second author are unknown. Any "visible" information will be appreciated.

computational standpoint, the complexity of the recognition problem for visibility graphs is only known to be in PSPACE [10]. It is not known to be in NP nor is it known to be NP-complete.

Visibility graphs do not lie in any of the well known classes of graphs such as planar graphs, chordal, circle or perfect graphs [10], [13]. The first set of necessary conditions for a graph to be a visibility graph was obtained by Ghosh [12]. However, it was shown by Everett [10] that these conditions were not sufficient. Further necessary conditions were developed by Coullard and Lubiw [9], but they also showed that they are not sufficient. Abello et al. [4] have strengthened these results by showing that the proposed conditions are not sufficient, even for triconnected graphs, and in the case of the conditions of [10], even for planar graphs.

In this paper we develop stronger necessary conditions for a graph to be a visibility graph.

In order to show that a given set of conditions on a graph are sufficient for the graph to be a visibility graph, one must demonstrate that every graph satisfying the conditions can be realized as the visibility graph of a simple polygon in the plane. However, this reconstruction problem appears to be quite difficult in the general case. In this paper we solve a combinatorial version of the reconstruction problem for general visibility graphs. We prove new necessary conditions for visibility graphs and show that these conditions are sufficient to construct a uniform oriented matroid of rank 3 corresponding to each graph in this class. These oriented matroids are combinatorial representations of simple polygons realizing the graphs, in the sense that any affine realization of the oriented matroids yields a simple polygon whose visibility graph is isomorphic to the given graph. It would be sufficient to show that each of these oriented matroids is affinely realizable, in order to obtain a characterization of visibility graphs of simple polygons. The main results of the paper are the following:

- In Section 3 a class of graphs called *quasi-persistent* graphs is defined and it is shown that visibility graphs are properly contained in this class. Quasi-persistent graphs are recognizable in polynomial time.
- Several new necessary conditions are proven for a given quasi-persistent graph to be a visibility graph (Section 4). These conditions strengthen Ghosh's original necessary conditions for visibility graphs.
- For each quasi-persistent graph satisfying these necessary conditions, a uniform oriented matroid of rank 3 is constructed (in polynomial time) such that any affine realization of the oriented matroid yields a simple polygon whose visibility graph is isomorphic to the given graph (Section 5).

We have made a conscious effort to make the paper as self-contained as possible.

2. Definitions

It is clear that every visibility graph is Hamiltonian and we therefore restrict our attention to Hamiltonian graphs. We further assume that the graphs considered are undirected, loopless and do not have multiple edges.

Let $G = (V, E)$ be a Hamiltonian graph with a prescribed Hamiltonian cycle H. The vertices of G are labeled along H from 0 to $n-1$. The vertex labeled i is denoted v_i. v_{i+1} and v_{i-1} denote the predecessor and successor of v_i on H. All subscript arithmetic is modulo n. It will be convenient to think of G as being embedded in the plane so that H forms a simple closed curve. In this setting, a traversal of H from v_i to v_j in the order $v_i, v_{i+1}, \ldots, v_j$ may be thought of as a counterclockwise traversal of H, and the traversal that goes from v_i to v_j in the order $v_i, v_{i-1}, \ldots, v_{j+1}, v_j$ corresponds to clockwise traversals. In this paper, unless specified otherwise, traversals of H are implicitly assumed to be in counterclockwise order.

For any two vertices v_i and v_j, the ordered set $\{v_i, v_{i+1}, \ldots, v_{j-1}, v_j\}$ of vertices encountered in traversing H from v_i to v_j, is called the chain from v_i to v_j and is denoted chain$[v_i, v_j]$. This set of vertices constitutes a simple path in G. The chain from v_{i+1} to v_{j-1} is denoted chain(v_i, v_j). We also use chain$[v_i, v_j)$ and chain$(v_i, v_j]$ to indicate left and right closed "intervals." We emphasize that chain(v_i, v_k) and chain(v_k, v_i) are always disjoint. We write $v_i < v_j < v_k$ if v_j lies on chain(v_i, v_k).

- *Blocking vertices*
  Two vertices v_i and v_j of G are said to be invisible if $v_i v_j \notin E$. For an invisible pair (v_i, v_k), a vertex v_j is called an *inner blocking* vertex [12], relative to H, if v_j lies on chain(v_i, v_k) and $v_x v_y \notin E$ for all v_x on chain$[v_i, v_j)$ and v_y on chain$(v_j, v_k]$. Similarly, a vertex v_j is called an *outer blocking* vertex relative to H, for the invisible pair $v_i v_k$, if v_j lies on chain(v_k, v_i) and $v_x v_y \notin E$ for all v_x on chain$(v_j, v_i]$ and v_y on chain$[v_k, v_j)$. In general, v_j is called a blocking vertex for the invisible pair $v_i v_k$, if it is either an inner or an outer blocking vertex for this pair. As an example, in Fig. 1, vertices 0, 3, 6 and 9 are blocking vertices for the invisible pairs (1, 9), (2, 6), (3, 7) and (0, 8), respectively.
- *Ordered paths and separability*
  A simple path $P = u_0 u_1 \cdots u_r$ is called an *ordered path* relative to H, if the vertices in P are encountered in the order $u_0, u_1, \ldots, u_r$ when H is traversed from u_0. Similarly, a simple cycle $C = u_0 u_1 \cdots u_r u_0$ is called an *ordered cycle* relative to H if the vertices in C are encountered in the order $u_0, u_1, \ldots, u_r, u_0$ (or its reverse) when H is traversed from u_0. Two pairs $v_i v_j$ and $v_k v_l$ are said to be *separable* [12] with respect to a vertex v_p, if both v_i and v_j are encountered before

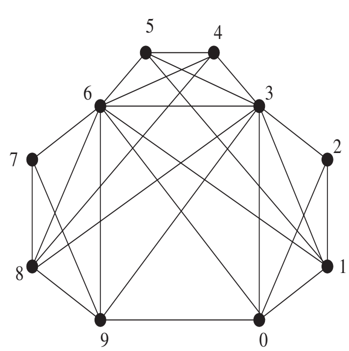

Fig. 1. A q-persistent graph.

v_k and v_l (or vice versa), when H is traversed from v_p. In this case we say that $v_i v_j$ and $v_k v_l$ are v_p-separable; otherwise we say that they are v_p-inseparable. Note that two pairs $v_i v_j$ and $v_k v_l$ are separable with respect to v_p, when the two pairs do not interlace on the boundary, (i.e. $v_i < v_j < v_k < v_l$) and v_p lies on chain(v_j, v_k) or on chain(v_l, v_i)).

In Fig. 1 the pairs (2, 6) and (1, 9) are not separable with respect to 3 and 0 but they are with respect to 7. On the other hand, (7, 3) and (7, 0) are separable with respect to 2.

3. Quasi-Persistent Graphs

We now introduce a class of graphs called quasi-persistent graphs, and show that visibility graphs of simple polygons are contained in this class. This class is a natural generalization of persistent graphs, a class originally introduced by Abello and Egecioglu [1].

- *Definition of quasi-persistence*
  A graph G with a Hamiltonian cycle H is said to be *quasi-persistent* (or *q-persistent*) relative to H, if for every triple of vertices $v_i < v_p < v_q$ such that $v_i v_p$ and $v_i v_q \in E$ and $v_i v_j \notin E$, for all v_j in chain(v_p, v_q), the following conditions hold:
  — v_p is adjacent to v_q;
  — for every v_j in chain(v_p, v_q), at least one of the vertices v_p or v_q is a blocking vertex for $v_i v_j$.
  The graph in Fig. 1 is a q-persistent graph.
- *First neighbor before and after a pair of non-consecutive vertices*
  For a pair $v_i v_j$ of non-consecutive vertices, let v_p be the first vertex adjacent to v_i that is encountered on a clockwise traversal of H starting from v_{j-1}. v_p is called the first neighbor of v_i before v_j and is denoted by pn$(v_i v_j)$. Similarly, the first vertex v_q adjacent to v_i, encountered on a counterclockwise traversal of H from v_{j+1}, is called the first neighbor of v_i after v_j and is denoted sn$(v_i v_j)$. Since G is Hamiltonian, pn$(v_i v_j)$ and sn$(v_i v_j)$ exist for every invisible pair $v_i v_j$, and they are distinct. Also, note that the definition is not symmetric, i.e. it is not necessary that pn$(v_i v_j)$ and sn$(v_i v_j)$ be the same as pn$(v_j v_i)$ and sn$(v_j v_i)$, respectively. The q-persistence conditions imply that for any invisible pair $v_i v_j$, the vertices pn$(v_i v_j)$ and sn$(v_i v_j)$ are adjacent in G, and at least one of them is a blocking vertex for $v_i v_j$.
  In Fig. 1 the pair (2, 7) has 0 and 3 as first neighbors.

3.1. *Ears in Q-Persistent Graphs*

Assume an ordering $\{v_1, v_2, \ldots, v_n\}$ of a Hamiltonian cycle H in a graph G (subscript arithmetic is module n). A vertex v_r is called an *ear* of G with respect to H if $v_{r-1}v_{r+1}$ is in $E(G)$. Every vertex in Fig. 1, except $0, 3, 6$ and 9, are ears with respect to the Hamiltonian cycle $\{0, 1, \ldots, 9, 0\}$.

We show next a graph generalization of Meisters' two ear theorem for simple polygons and then we prove that q-persistent graphs (with at least four vertices) are closed under ear deletion.

Ears Theorem. *Any Hamiltonian graph G that satisfies the first q-persistence condition has at least two ears.*

Proof. Any graph G that satisfies the first q-persistence condition must have two vertices v_i and v_k, non-consecutive in the Hamiltonian cycle, such that $v_i v_k$ is in $E(G)$. Moreover, if no vertex in chain(v_i, v_k) is an ear it implies that the subgraph induced by chain$[v_i, v_k]$ is a path. This implies the existence of an ear on chain(v_i, v_k) and one on chain(v_k, v_i). □

Ear Deletion. *Q-Persistent graphs (with at least four vertices) are closed under ear deletion.*

Proof. Let G_r denote the Hamiltonian graph obtained from G by deleting an ear v_r of G. Suppose that G_r is not q-persistent. This means that there exists a triple of vertices $v_i < v_p < v_q$ in G_r such that v_i is adjacent to v_p and v_q but not adjacent to any vertex in chain(v_p, v_q), for which the q-persistence of G_r is violated. This implies that v_r must lie on chain(v_p, v_q) otherwise the q-persistence of G will be violated too. By the same reason v_i and v_r must be adjacent in G. By applying the q-persistence conditions to the triples $v_i < v_p < v_r$ and $v_i < v_r < v_q$ in G one concludes that $v_p v_r$ and $v_r v_q$ must be edges of G.

Suppose now that the first q-persistence condition is violated, i.e. $v_p v_q$ is not an edge of G_r and this means that it is also not an edge of G. Therefore considering the invisible pair $v_q v_p$ we conclude that the vertex pn$(v_q v_p)$ and the vertex sn$(v_p v_q)$ lie on chain$(v_p, v_i]$ and chain$(v_p, v_r]$, respectively (note that v_q is adjacent to v_r and v_i).

We show now that $v_y = v_r$. Suppose, to the contrary, that v_y lies on chain$(v_p, v_{r-1}]$. If $v_x \neq v_i$, note that neither v_x nor v_y can be a blocking vertex for $v_p v_q$, since the pairs $v_p v_i$ and $v_p v_r$ are in E. This contradicts the second q-persistence condition for the invisible pair $v_p v_q$. Thus $v_x = v_i$. However, then, by the first q-persistence condition, $v_i v_y$ is an edge in E. This is turn implies that $v_i v_y$ is also an edge of G_r, which means that v_i is adjacent to a vertex distinct from v_r in chain$(v_p v_q)$ of G_r and this is contrary to our hypothesis.

So we may assume that $v_y = v_r$. Following arguments similar to those in the last paragraph, we see that $v_x = v_i$. Now, since $v_y = v_r$ is an ear of G, it cannot be a blocking vertex for $v_p v_q$. By the second q-persistence condition it then follows that $v_x = v_i$ must be a blocking vertex for this pair. It then follows that no vertex in chain(v_r, v_i) is adjacent to a vertex in chain(v_i, v_r). This means that $v_{r-1} v_{r+1}$ is not an edge in G, contradicting the assumption that v_r was an ear of G. This concludes the case when the first q-persistence condition is violated.

Now suppose that the second q-persistence condition is violated, i.e. neither v_p nor v_q is a blocking vertex for an invisible pair $v_i v_j$ in G_r, with v_j lying on chain$(v_p v_q)$ of G_r. This implies by definition that they are not blocking vertices for this pair in G. By assumption, v_r is not a blocking vertex in G. Consider the pair $v_i v_j$ in G. If v_j lies on chain(v_p, v_r), then $v_p = \text{pn}(v_i v_j)$ and $v_i = \text{sn}(v_i v_j)$. On the other hand, if v_j lies on chain(v_i, v_q), then $v_i = \text{pn}(v_i, v_j)$ and $v_q = \text{sn}(v_i v_j)$. For either of these pair of choices, the second q-persistence condition is violated in G for the invisible pair $v_i v_j$. □

3.2. *Q-Persistent Graphs and Ghosh's Conditions*

Ghosh [12] gave the first set of necessary conditions for a given graph to be a visibility graph. These conditions which we henceforth call Ghosh's original conditions are summarized below.

Proposition 1 (Ghosh). *If a graph G is the visibility graph of a simple polygon, then*:

1. *G has a Hamiltonian cycle H.*
2. *Every ordered cycle (relative to H) of length ≥ 4 has a chord.*
3. *Every invisible pair in G has a blocking vertex relative to H.*
4. *If two invisible pairs are separable with respect to a vertex v_p, then v_p cannot be the only blocking vertex for both the invisible pairs.*

Our q-persistent graphs satisfy the first and third conditions of Proposition 1 by definition. In fact, the second q-persistence condition (*ordered chordality*) appears, at first glance, to be much stronger than Proposition 1's third condition. However, it can be shown that the class of q-persistent graphs is equivalent to the class of graphs that satisfies the first three conditions. This is stated in Theorem 2 below ([16] contains a detailed proof).

Theorem 2. *A graph G with a Hamiltonian cycle H is q-persistent relative to H if and only if every ordered cycle of length ≥ 4 has a chord and every invisible pair has a blocking vertex (relative to H).*

Thus, q-persistent graphs are not a fundamentally new class of graphs. The main advantage of the above formulation is that the simpler structure of the definition makes it easier to analyze and prove properties of the resulting class. It is interesting to note the relationship between the two q-persistence conditions and conditions 2 and 3 of Proposition 1. The first q-persistence condition is a "weaker" version of ordered chordality, in the sense that graphs that are Hamiltonian and ordered chordal are *properly* contained in the class of (Hamiltonian) graphs satisfying the first q-persistence condition. On the other hand, the second q-persistence condition is a stronger version of condition 3 since Hamiltonian graphs that satisfy the second q-persistence condition are properly contained in the class of Hamiltonian graphs satisfying condition 3. However, when both pairs of conditions are considered together, the classes become equivalent!

Since visibility graphs satisfy all four of Ghosh's conditions, it follows from Theorem 2 that visibility graphs are properly contained in the class of *q-persistent graphs*. The following sections develop additional conditions for a q-persistent graph to be a visibility graph and show that they encode enough information that allow us to recover an *oriented matroid* (*chirotope*). If the obtained oriented matroid is realizable it will mean that our set of conditions is a characterization of visibility graphs of polygons.

Remark. We mention in closing that Theorem 2 together with the fact that conditions 2 and 3 of Proposition 1 can be checked in quadratic time [13] tell us that q-persistent graphs can be recognized in polynomial time.

4. Novel Necessary Conditions for Visibility Graphs

We assume throughout this section that P is a simple polygon in the Euclidean plane and that a q-persistent graph G is its visibility graph. Arbitrary points of the plane are denoted as p_x, p_y, etc. For two points p_x and p_y, the ray from p_x in the direction of p_y is denoted r_{xy}. For a vertex v_i of G the corresponding vertex of P is denoted v_i^*. We also use r_{ij} to denote the ray from v_i^* in the direction of v_j^*. A polygon P whose visibility graph is G is called a realization of the graph G. A given q-persistent graph that is a visibility graph can have many different realizations.

4.1. *Geometric Interpretation of Q-Persistence*

Suppose $W = w_0, \ldots, w_k$ is the sequence of neighbors of a vertex v_i in G, obtained in traversing H, with $w_0 = v_{i+1}$ and $w_k = v_{i-1}$. In a polygon P realizing G, $\angle w_k^* v_i^* w_{j-1}^* < \angle w_k^* v_i^* w_j^*$ for $1 \leq j \leq k-1$ (this is a well known fact). The second q-persistence condition can now be interpreted geometrically: let $v_i < v_p < v_q$ be a triple of vertices in G such that $v_i v_p, v_i v_q \in E$ and $v_i v_j \notin E$ for all v_j on chain(v_p, v_q). For the corresponding triple of points v_i^*, v_p^*, and v_q^* in a realization P of G, there exists a unique segment $v_k^* v_{k+1}^*$ on the boundary of P, such that v_k and v_{k+1} lie on chain$[v_p, v_q]$, and for *any* ray r_{ix} such that $\angle v_{i-1}^* v_i^* v_p^* < \angle v_{i-1}^* v_i^* p_x^* < \angle v_{i-1}^* v_i^* v_q^*$, the first segment on the boundary of P that is intersected by ray r_{ix} is $v_k^* v_{k+1}^*$. The fact to be emphasized here is that the segment so obtained does not depend on the specific ray fixed, but only on the triple of points involved (see Fig. 2).

For any vertex v_j on chain$(v_p, v_k]$, the vertex v_p is a blocking vertex (in G) for $v_i v_j$ and for any vertex v_j on chain$[v_{k+1}, v_q)$, the vertex v_q is a blocking vertex for $v_i v_j$. The segment $v_k^* v_{k+1}^*$ is called the *split segment* for the triple of points $v_i^* v_p^* v_q^*$. The corresponding edge in the graph is called the *split edge*. Intuitively, the split edge determines which one of the points v_p^* and v_q^* is involved in "physically" blocking a given pair $v_i v_j$ on chain(v_p, v_q) in a given polygon whose visibility graph is G. In general, the split edge is not determined by the visibility graph alone. Different polygons with the same underlying visibility graph may have different split edges for the same triple of vertices in G.

The q-persistence conditions stipulate that for any invisible pair $v_i v_j \notin E$ of a q-persistent graph G, at least one of the vertices pn$(v_i v_j)$ or sn$(v_i v_j)$ must be a blocking vertex for the invisible pair. However, according to the discussion above, in any fixed realization of G, at most one of these vertices "physically" blocks the corresponding invisible pair of points in the realization. This motivates the following definitions.

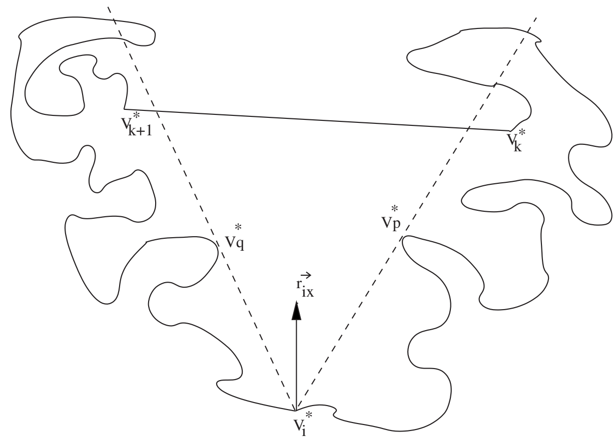

Fig. 2. Geometric interpretation of q-persistence.

4.2. *Blocking Vertex Assignments*

A vertex v_p is called a *primary blocking vertex* for an invisible pair $v_i v_j$ if v_p is a blocking vertex for $v_i v_j$ and $v_i v_p \in E$. By the definition of blocking vertices, the only possible choices for the primary blocking vertices for $v_i v_j$ are $\text{pn}(v_i v_j)$ and $\text{sn}(v_i v_j)$. Therefore, if either of the vertices $\text{pn}(v_i v_j)$ or $\text{sn}(v_i v_j)$ is a blocking vertex for the invisible pair $v_i v_j$ in a q-persistent graph G, then it is called a primary blocking vertex for $v_i v_j$. The q-persistence conditions imply that every invisible pair has at least one primary blocking vertex. Also, the primary blocking vertices of the pair $v_i v_j$ are not necessarily the same as those for $v_j v_i$.

A **blocking vertex assignment**[1] for a q-persistent graph G is a function β: $\bar{E} \to V$ such that, for all $v_i v_j \notin E$, $\beta(v_i v_j)$ is a primary blocking vertex for $v_i v_j$. Any q-persistent graph has at least one blocking vertex assignment. If G is a visibility graph, then every fixed realization, P of a given q-persistent graph, determines a particular blocking vertex assignment for G as follows. For a triple $v_i < v_p < v_q$ of vertices of G, such that $v_i v_p$, $v_i v_q \in E$ and $v_i v_j \notin E$ for all v_j on $\text{chain}(v_p, v_q)$, let $v_k^* v_{k+1}^*$ be the split segment in P for the triple $v_i^* v_p^* v_q^*$. We set $\beta(v_i v_j) = v_p$, for all v_j on $\text{chain}(v_p, v_k]$. For all v_j on $\text{chain}[v_{k+1}, v_q)$ we set $\beta(v_i v_j) = v_q$. From the discussion in the last section, it follows that β is a blocking vertex assignment for G. This blocking vertex assignment is called a canonical blocking vertex assignment for G determined by the realization P.

We now consider the following problem: Given a q-persistent graph together with a blocking vertex assignment β, determine the conditions under which there exists a polygon P whose visibility graph is G, and such that the canonical assignment on G determined by P is β. Such conditions will clearly yield a characterization of visibility graphs. It turns out that blocking vertex assignments on q-persistent graphs must satisfy four additional necessary conditions in order to be canonical assignments.

---

[1] Everett, in [10], also defines a similar notion, but the requirement that vertices in the image of the function be primary blocking vertices makes the definition given here strictly stronger than the one in [10].

A blocking vertex assignment is said to be **locally inseparable** if any two invisible pairs $v_i v_j$ and $v_k v_l$ such that $\beta(v_i v_j) = \beta(v_k v_l) = v_p$ are v_p-inseparable (see definition in the last paragraph of Section 2). The following is a necessary condition for a blocking vertex assignment to be a canonical blocking vertex assignment:[2]

Necessary Condition 1 (Local Separability). *If β is a canonical blocking vertex assignment for a q-persistent graph G, determined by a realization P, then β is locally inseparable.*

Graph Occluding Paths versus Geometric Shortest Paths. In order to state the remaining necessary conditions we need to introduce the following definition. Given an invisible pair $v_i v_k$ in G, an **occluding path** generated by β, between v_i and v_k, denoted $\text{path}_\beta(v_i, v_k)$ is a path $v_i u_0 \cdots u_r v_k$ in G, such that $u_0 = \beta(v_i v_k)$, $u_j v_k \notin E$, and $u_{j+1} = \beta(u_j, v_k)$ for $0 \leq j \leq r - 1$. It is readily seen that a given blocking vertex assignment determines a unique occluding path between every invisible pair of vertices. It can also be shown that this path is simple and that every internal vertex on this path is a blocking vertex for the invisible pair. For notational convenience, we identify $\text{path}_\beta(v_i, v_k)$ with its underlying set of vertices.

When a graph G is the visibility graph of a polygon P, the graph theoretical notion of occluding path corresponds to the geometric notion of shortest path under the geodesic metric. This fact is stated in the following proposition.

Proposition 3. *Let β be the canonical blocking vertex assignment for a visibility graph G, determined by a fixed realization P. A vertex v_x lies on* $\text{path}_\beta(v_i, v_j)$ *if and only if v_i^* lies on the Euclidean shortest path in P between v_i^* and v_j^*.*

Path symmetry and path consistency. The remaining necessary conditions arise as a result of this correspondence between occluding and Euclidean shortest paths. A blocking vertex assignment is called **path-symmetric** if for every invisible pair $v_i v_k$ such that $\text{path}_\beta(v_i, v_k) = v_i u_0 \cdots u_r v_k$, we have $\text{path}_\beta(v_k, v_i) = v_k u_r \cdots u_0 v_i$. We denote this as $\text{path}_\beta(v_k, v_i) = \text{path}^R_\beta(v_i, v_k)$. In other words, even though a blocking vertex assignment is not symmetric in the blocking vertices it assigns to invisible pairs $v_i v_k$ and $v_k v_i$, it must ensure the symmetry of the occluding paths generated under the assignment between every invisible pair of vertices. Since Euclidean shortest paths between two points inside a simple polygon are unique, it readily follows that canonical blocking vertex assignments must be path-symmetric; this is Necessary Condition 2 below.

Necessary Condition 2 (Path Symmetry). *If β is a canonical blocking vertex assignment for a q-persistent graph determined by a realization P, then β is path-symmetric.*

Path consistency conditions. The two remaining necessary conditions reflect the constraints imposed on occluding paths, generated by canonical blocking vertex assign-

[2] Everett [10] conjectures a similar result. However, since our definition of a blocking vertex assignment is stricter, our *local separability* Necessary Condition 1 is stronger.

ments, because of their correspondence with Euclidean shortest paths. A blocking vertex assignment satisfying these two conditions is called a **path-consistent** assignment.

Necessary Condition 3. *If β is a canonical blocking vertex assignment for a q-persistent graph determined by a realization P, and if* $u_x \in \mathrm{path}_\beta(v_i, u_y)$ *and* $u_y \in \mathrm{path}_\beta(u_x, v_k)$, *then* $u_x, u_y \in \mathrm{path}_\beta(v_i, v_k)$.

Necessary Condition 4. *If β is a canonical blocking vertex assignment for a q-persistent graph determined by a realization P, and*

1. *If* $v_p \in \mathrm{path}_\beta(v_i, v_k)$ *is an inner blocking vertex for* $v_i v_k$, *then for all* v_x *on* $\mathrm{chain}[v_i, v_p)$ *and* v_y *on* $\mathrm{chain}(v_p, v_k]$, $v_p \in \mathrm{path}_\beta(v_x, v_y)$.
2. *If* $v_p \in \mathrm{path}_\beta(v_i, v_k)$ *is an outer blocking vertex for* $v_i v_k$, *then for all* v_x *on* $\mathrm{chain}[v_p, v_i)$ *and* v_y *on* $\mathrm{chain}(v_k, v_p)$, $v_p \in \mathrm{path}_\beta(v_x, v_y)$.

The proofs of Necessary Conditions 3 and 4 are based on the fact that Euclidean shortest paths satisfy the above combinatorial conditions. It is natural to ask whether all the above four conditions are independent of each other. It can be shown that in fact they are. Namely, for any subset of these conditions, there exist q-persistent graphs for which blocking vertex assignments can be constructed that satisfy only that subset and no others. On the other hand, to contrast these conditions with those in Proposition 1, Everett has exhibited a graph that satisfies Ghosh's original conditions and yet it is not a visibility graph. It can be shown that Everett's example is a q-persistent graph that does not have a blocking vertex assignment that satisfies condition 1. The graph in Fig. 1 is a q-persistent graph that was shown not to be a visibility graph in [2]. It can be shown that this graph has one blocking vertex assignment that is locally separable, and another that is path symmetric and path consistent, but no one that satisfies all conditions simultaneously.

4.3. *Q-Persistence Ensures Symmetry and Consistency*

Before moving along, one may wonder if there are q-persistent graphs which do not have path-consistent and path-symmetric assignments. It is a little bit surprising that the answer is no as stated in the following theorem. The proof relies on the fact that q-persistent graphs are closed under ear deletion.

Theorem. *If G is a q-persistent graph with respect to a Hamiltonian cycle H, then G has a blocking vertex assignment, with respect to H, that is path symmetric and path-consistent.*

Proof. Assume that $w_0, w_1, \ldots, w_s$ is the ordered sequence of neighbors of a vertex v_r that is obtained by traversing the Hamiltonian cycle H, with $w_0 = v_{r+1}$ and $w_s = v_{r-1}$. The first thing to notice is that q-persistence implies that any invisible pair $w_i w_k$ $(i < k)$ has blocking vertex assignments only in the set $\{v_r, w_{i+1}, \ldots, w_{k-1}\}$.

Let v_r be an *ear* of G and let $v_r v_t$ be an invisible pair with $v_t \in \mathrm{chain}(w_j, w_{j+1})$ for some $j, 0 \le j < s$. If β is a path-symmetric and path-consistent blocking assignment for

G_r, the subgraph of G obtained by deleting the ear v_r, then the only possible choices for $\beta(v_r v_t)$ lie on the set $\{w_j, w_{j+1}\}$. It can be shown that at least one of these two vertices can always be chosen to extend β to a path-consistent and path-symmetric blocking assignment for G. By the *ears theorem*, G_r is also q-persistent and the result follows by induction. □

Remark. The proof of the last theorem indicates that any set of conditions that characterizes a visibility graph of simple polygons shall be preserved under ear deletion and shall be recoverable from the neighborhood of the ear and its interaction with the ear deleted subgraph. In this regard, the *local separability* condition is the only condition that remains to be subject to further scrutiny. We point out later, some developments in this direction.

In summary, we have shown that the conditions of Proposition 1 can be strengthened to state that if a graph is a visibility graph, then it is persistent and has a blocking vertex assignment that is *simultaneously* path-symmetric, path-consistent and locally separable. The key question is whether these conditions are sufficient. In the next section we provide a partial answer to this question. Namely, we show how to obtain a simplicial chirotope associated with any q-persistent graph that satisfies our necessary conditions. It is worth mentioning that these conditions have been proven to be sufficient for k-spiral polygons when $k \leq 2$, and for funnel and staircase polygons. In the case that $n \leq 7$, the conditions are also sufficient since every simplicial chirotope on that many elements is affinely coordinatizable over the reals [14].

5. Q-Persistent Graphs and Oriented Matroids

We consider the problem of determining, given as input a q-persistent graph with a blocking vertex assignment satisfying the conditions of the previous section, a combinatorial representation of a potential polygon whose visibility graph is isomorphic to the given graph. The main result of this section is that such a combinatorial reconstruction appears to be significantly easier than the actual reconstruction of the polygon.

Oriented matroids are a well studied combinatorial representation [7], [11], [17] for point configurations. In the following we adopt the conventions of [7] and identify oriented matroids with their representations by *chirotopes*. An equivalence proof for this representation and the classical definition in terms of signed circuits of matroids may be found in [17]. We are concerned here only with the definition of oriented matroids of rank 3. Let τ_n, $n \geq 3$, denote the set of increasing triples from the set $\{0, \ldots, n-1\}$ (i.e. 3-tuples (i, j, k) where $i < j < k$). A mapping $\chi\colon \tau_n \to \{-1, +1, 0\}$ (that can be extended by alternation to the set of all ordered triples from $\{0, \ldots, n-1\}$) is called a **chirotope** if for all i, $0 \leq i \leq n-1$, and all 4-tuples, $0 \leq j \leq k \leq l \leq m \leq n-1$, from $\{0, \ldots, n-1\}$ the set

$$\left\{\begin{array}{c} \chi(i,j,k)\chi(i,l,m) \\ -\chi(i,j,l)\chi(i,k,m) \\ \chi(i,j,m)\chi(i,k,l) \end{array}\right\}$$

either contains $\{-1, +1\}$ or equals $\{0\}$. The chirotope is called simplicial if its image is contained in the set $\{-1, +1\}$.

A chirotope is called **coordinatizable** if there exists an $n \times 3$ matrix M such that for any triple $(i, j, k) \in \tau_n$, $\chi(i, j, k)$ agrees with the sign of the corresponding 3×3 subdeterminant of M. The chirotope associated with a point configuration assigns to each triple of points its orientation (given by the signed area). The fact that these subdeterminants obey the chirotope conditions above follows from the well known Grassman–Plucker identities (see [7]). Deciding if a given rank 3 oriented matroid is coordinatizable is known to be NP-hard [2]. It is also polynomially equivalent to the decision problem for the existential theory of the reals [19] and thus in PSPACE [8].

We now establish the existence of a simplicial chirotope, corresponding to every q-persistent graph G with a blocking vertex assignment β that is path-symmetric, path consistent and locally separable. Call such β a *feasible* blocking vertex assignment. The chirotope has the property that any of its coordinatizations defines a simple polygon whose visibility graph is isomorphic to the input graph and induces a canonical blocking vertex assignment on G which is exactly β. This chirotope, called the *Normal Chirotope*, for the pair (G, β) can be constructed in polynomial time given G and β.

Normal chirotope construction. Given a q-persistent graph G with a feasible blocking vertex assignment β, we define a function $\chi_{G,\beta}\colon \tau_n \to \{-1, +1\}$ (that can be extended by alternation to the set of all ordered triples in $\{0, \dots, n-1\}$), where

$$\chi(i, j, k) = \begin{cases} -1 & \text{if there exists an occluding path generated by } \beta \text{ that contains} \\ & \text{the vertices } v_i, v_j \text{ and } v_k, \\ +1 & \text{otherwise.} \end{cases}$$

It is clear that χ can be constructed from G and β in $O(n^4)$ time. Moreover, it defines a simplicial chirotope. This constitutes the next result.

Theorem 4. *If G is a q-persistent graph and β is a blocking vertex assignment that is path-symmetric*, *path-consistent and locally separable*, *then $\chi_{G,\beta}$ is a simplicial chirotope*.

Proof. The proof consists in checking from the definitions that for all i, $1 \le i \le n$, and all 4-tuples, $1 \le j \le k \le l \le m \le n$, from $\{1, 2, \dots, n\}$ the set

$$\left\{ \begin{matrix} \chi(i, j, k)\chi(i, l, m) \\ -\chi(i, j, l)\chi(i, k, m) \\ \chi(i, j, m)\chi(i, k, l) \end{matrix} \right\},$$

denoted $\langle i \mid jklm \rangle$, either contains $\{-1, +1\}$ or equals $\{0\}$. We say that a five subset $\langle i \mid jklm \rangle$ satisfies the chirotope condition if the above property holds. It is also useful to note, by the alternating property, that a 5-subset satisfies the chirotope condition iff it satisfies it for any permutation of the five elements in the subset. To simplify notation we indicate $\chi([i, j, k])$ simply as $[i, j, k]$.

We first claim that if any five vertices v_i, v_j, v_k, v_l, v_m lie on an occluding path generated by β, then $\langle i \mid jklm \rangle$ satisfies the chirotope condition.

To see this, let $w_0, w_1, \dots, w_r$ be an occluding path between w_0 and w_r. By the alternating property we can assume that they occur in this specified order and prove that

$\langle i \mid jklm \rangle$ satisfies the chirotope condition in that case. Since β is *feasible*, it follows that any two consecutive vertices in the sequence v_i, v_j, v_k, v_l, v_m, are invisible and that the intermediate vertices lie on the occluding path between them (this follows from properties of blocking vertex assignments in q-persistent graphs). Thus $[x, y, z] = -1$ for any three vertices v_x, v_y, v_z that occur in that order in the path. So the sign of each of the triples involved in the chirotope condition is -1 and this implies the set contains $\{-1, +1\}$ as required. We invoke this claim several times in what follows.

Suppose that there exists a 5-subset for which $\langle i \mid jklm \rangle$ does not satisfy the chirotope conditions. By the alternating property we can assume that these vertices occur in the same order of the traversal of the Hamiltonian cycle H.

We consider a number of cases and derive a contradiction on each case. Suppose the set does not contain -1. In this case, each product in the set evaluates to $+1$. In particular, $[i, j, l][i, k, m] = -1$. Without loss of generality, assume $[i, j, l] = -1$ and $[i, k, m] = +1$.

We have three possible scenarios. $v_l \in \text{path}_\beta(v_i, v_j)$ or $v_j \in \text{path}_\beta(v_i, v_l)$ or $v_i \in \text{path}_\beta(v_j, v_l)$. We provide the argument for just one of these cases since the others are completely symmetric.

If $v_l \in \text{path}_\beta(v_i, v_j)$, the path consistency condition implies that v_l is also in $\text{path}_\beta(v_i, v_k)$ and thus $[i, k, l] = -1$. Since, $[i, j, m][i, k, l] = 1$ this implies $[i, j, m] = -1$. Again, three choices arise:

1. $v_m \in \text{path}_\beta(v_i, v_j)$. This implies by path consistency that $v_m \in \text{path}_\beta(v_i, v_k)$ which means that $[i, k, m] = +1$, contradicting our earlier supposition.
2. $v_j \in \text{path}_\beta(v_i, v_m)$. This implies $v_j \in \text{path}_\beta(v_i, v_l)$, contradicting our assumption that $v_l \in \text{path}_\beta(v_i, v_j)$.
3. $v_i \in \text{path}_\beta(v_j, v_m)$. Since $v_l \in \text{path}_\beta(v_i, v_j)$ this implies that $v_l \in \text{path}_\beta(v_j, v_i)$. Thus, $v_i \in \text{path}_\beta(v_l, v_m)$ and $[i, l, m] = -1$. This means that $[i, j, k] = -1$ and again we have three choices:
   - In case $v_k \in \text{path}_\beta(v_i, v_j)$, this forces $v_k \in \text{path}_\beta(v_l, v_j)$. We have now that $v_i \in \text{path}_\beta(v_m, v_l)$, $v_l \in \text{path}_\beta(v_i, v_k)$ and $v_k \in \text{path}_\beta(v_l, v_j)$. Invoking the first path consistency condition we get that all five vertices now lie on the occluding $\text{path}_\beta(v_m, v_j)$, contradicting our earlier claim about any five vertices that lie on an occluding path.
   - In the case that $v_j \in \text{path}_\beta(v_i, v_k)$ we get that $v_i \in \text{path}_\beta(v_m, v_l)$, $v_l \in \text{path}_\beta(v_i, v_j)$ and $v_j \in \text{path}_\beta(v_l, v_k)$ which together imply that all five vertices lie on $\text{path}_\beta(v_m, v_k)$, which is a contradiction.
   - Finally, if $v_i \in \text{path}_\beta(v_j, v_k)$ observe that $v_i \in \text{path}_\beta(v_l, v_m)$. This means that $v_j v_k$ and $v_l v_m$ are v_i-separable, contradicting that β was feasible to start with.

The cases when $v_j \in \text{path}_\beta(v_i, v_l)$ and $v_i \in \text{path}_\beta(v_j, v_l)$ are completely symmetric.

Of course we do not need to consider when the set does not contain $+1$. Here each product evaluates to -1. In this case, too, the arguments are quite similar to the last case. By exhaustively enumerating all the possibilities is not difficult to derive a contradiction in each case. □

Although the above proof is tedious the methods used are elementary. Also, the construction is simple and intuitive. If we could establish the coordinatizability of the above chirotopes in each case, this would characterize visibility graphs of simple polygons completely.

Now, we show that if the *Normal Chirotope* is realizable, then the corresponding realization yields a polygon whose visibility graph is isomorphic to G. Moreover, the canonical assignment induced on the graph by the polygon is precisely β.

Theorem 5. *If the chirotope $\chi_{G,\beta}$ of the previous theorem is realizable, then the necessary conditions discussed in this paper fully characterize visibility graphs of simple polygons.*

Proof. Suppose that $\chi_{G,\beta}$ is affinely coordinatizable. First note that the points $v_0^*, \ldots, v_{n-1}^*$ in a plane realization together with the segments $v_i^* v_{i+1}^*$ (mod n) constitute a simple polygon P. To see this notice that if β is *feasible*, it is impossible that $\chi(i, i+1, j) \cdot \chi(i, i+1, j+1) = -1$ and $\chi(j, j+1, i) \cdot \chi(j, j+1, i+1) = -1$ when $|i-j| > 1$. In the realization this implies that no two segments of the polygon intersect, ensuring simplicity. Also, note that since the chirotope is simplicial, the resulting point configuration is non-degenerate.

Now consider a triple $v_i v_p v_q$ in G such that v_i is adjacent to no vertex in chain(v_p, v_q) but is adjacent to both v_p and v_q. Let $v_k v_{k+1}$ be the *split edge* for the triple v_i, v_p, v_q that is determined by β. Interpreting the signs that β assigns to the ordered triples as orientations of the corresponding triples of points, is not difficult to check the following:

- The interior of the triangle $v_i^* v_p^* v_q^*$ contains no points of P.
- The points corresponding to chain(v_p, v_k) and the points corresponding to chain$[v_{k+1}, v_q]$ lie on opposite half-spaces of the line containing $v_i^* v_p^*$. Similarly, the points corresponding to chain$[v_p, v_k]$ and those corresponding to chain$[v_{k+1}, v_q)$ lie on opposite half-spaces of the line containing $v_i^* v_q^*$.

Note also that by *local-inseparability* v_i cannot lie both on an occluding path from v_p to a vertex on chain(v_i, v_p) and also on a path from v_q to one on chain(v_q, v_i). This together with the first item above, allows us to claim that v_i^*, v_p^* and v_q^*, are visible from each other. On the other hand, the second item above is the key to conclude that v_i^* is invisible from all the points corresponding to those on chain(v_p, v_q).

Therefore, v_p^* and v_q^* are successive neighbors of v_i^* and $v_k^* v_{k+1}^*$ is the *split* segment for this triple. A similar argument shows the converse case, that is when v_p^* and v_q^* are the succesive neighbors of a vertex v_i^*, then the corresponding vertexes are all adjacent to each other. Also, if the *split segment* determined by the realization is $v_k^* v_{k+1}^*$, then the corresponding *split edge* determined by β for this triple is $v_k v_{k+1}$. Repeating the argument for each such "minimal" triple shows that the coordinatization gives a simple polygon P whose visibility graph is G and determines the canonical vertex assignment β on G.

From the previous discussion and the fact that visibility graphs of simple polygons are q-persistent and have *blocking assignments* that are simultaneously *locally separable*, *path symmetric and path consistent*, the result follows. □

6. Relation with Pseudo-Visibility and Stretchability

Based on a previous version of some of the results presented here [5], Streinu and O'Rourke [21] define a notion of visibility among vertices and edges of pseudo-polygons. They characterize the corresponding class of bipartite graphs and call it Vertex-Edge Visibility graphs. This class encodes precisely the notions of *split segments* and *split edges* as introduced in Section 4.1 where the geometric interpretation of q-persistence is discussed. Their characterization corresponds precisely to choosing a *blocking vertex assignment* that is *locally separable*, *path-symmetric and path-consistent*. Therefore, it is not surprising that they are able to recover also (by different means) an oriented matroid. In their case they work with pseudo-line arrangements but these are equivalent to oriented matroids. Questions of straight line realizability are avoided in that work since they start by defining visibility on pseudo-line arrangements. In summary, Theorem 4 is equivalent to their characterization of pseudo-visibility graphs. All this suggests that the intrinsic difficulty of characterizing visibility graphs of simple polygons lies in our lack of understanding of the "real" obstructions to stretchability of the Normal Chirotope introduced in Section 5. We know at this point that this chirotope is realizable for certain specialized classes of polygons which include spiral and 2-spiral polygons, funnel polygons, monotone polygons and convex fans. We will describe these and related realizability results in a forthcoming paper [1].

7. An Alternative Necessary Condition for Visibility

The q-persistent graph depicted in Fig. 1 has exactly four blocking vertexes (i.e. 0, 3, 6 and 9). Therefore all the invisible pairs must use them to block their visibility. Moreover, because visibility graphs are hereditary with respect to ordered cycles any necessary condition must behave accordingly. We know that q-persistent graphs are hereditary by our Ears Theorem (Section 3). Moreover, they always have symmetric and path-consistent block assignments, therefore the reason why any q-persistent graph will fail to be a visibility graph is because the local-separability property, even if it is satisfied for the entire graph, may fail when restricted to the subgraph induced by an ordered subcycle. In [13] Ghosh has suggested a new condition for visibility graphs that exploits this fact. The condition basically says that for any ordered subcycle C, the restriction of the blocking assignment to the invisible pairs of C cannot assign more than $|C| - 3$ vertices. This makes perfect sense since the sum of internal angles of any simple polygon on $|C|$ vertices is $(|C| - 2)180^\circ$. Ghosh uses this very succinct condition to offer an alternate explanation of why the graph in Fig. 1 is not a visibility graph. The reason is that any global blocking vertex assignment forces vertices 0, 3, 6 and 9 to be reflex vertices and then the subpolygon 0, 2, 3, 6, 7, 9, 0 has six vertices out of which four are reflex vertices and this is impossible. Notice however that even though this graph has a local-separable blocking assignment (in our sense) it does not have one that satisfies *simultaneously* the path symmetry and path consistency conditions. This together with the fact that q-persistent graphs are ordered cycle hereditary suggests the following conjecture:

Conjecture. *Graphs satisfying the conditions of Proposition* 1 *plus Ghosh's new condition discussed above* (*proposed in* [13]) *are not sufficient to recognize visibility graphs of simple polygons*.

Acknowledgments

Thanks to Sandra Sudarsky for her continued support. To Laszlo Szekely and Farhad Shahrokhi, our appreciation for organizing a very stimulating special session at the AMS regional conference, in Columbia, South Carolina, March 2001. This meeting provided a relaxed forum to think again about beautiful questions that had been patiently waiting.

References

1. J. Abello. Visibility and Stretchability. In preparation
2. J. Abello and O. Egecioglu. Visibility Graphs of Staircase Polygons with Uniform Step Length. *International Journal of Discrete and Computational Geometry*, **3** (1993), 27–37.
3. J. Abello, O. Egecioglu and K. Kumar. Visibility Graphs of Staircase Polygons and the Weak Bruhat Order: From Visibility Graphs to Maximal Chains. *Discrete and Computational Geometry*, **14**(3) (1995), 331–358.
4. J. Abello, L. Hua and C. Pisupati. On Visibility Graphs of Simple Polygons. *Congressus Numerantium*, **90** (1992), 119–128.
5. J. Abello and K. Kumar. Visibility Graphs and Oriented Matroids. In *Proceedings of Graph Drawing*, *Princeton*, pp. 147–158. LNCS, 894. Springer-Verlag, Berlin, 1994.
6. B. Aronov. On the Geodesic Voronoi Diagram of Point Sites in a Simple Polygon. *Algorithmica*, **4** (1989), 109–140.
7. J. Bokowski and B. Sturmfels. *Computational Synthetic Geometry*. Lecture Notes in Mathematics, 1335. Springer-Verlag, New York, 1989.
8. J. Canny. Some Algebraic and Geometric Computations in PSPACE. In *Proceedings of the* 20*th ACM Symposium on Theory of Computing*, pp. 460–467, 1988.
9. C. Coullard and A. Lubiw. Distance Visibility Graphs. In *Proceedings of the ACM Symposium on Computational Geometry*, pp. 290–302, 1991.
10. H. Everett. Visibility Graph Recognition. Ph.D. Dissertation, Department of Computer Science, University of Toronto, 1990.
11. J. Folkman and J. Lawrence. Oriented Matroids. *Journal of Combinatorial Theory*, *Series B*, **25** (1978), 199–236.
12. S.K. Ghosh. On Recognizing and Characterizing Visibility Graphs of Simple Polygons, pp. 132–139. LNCS, 318. Springer-Verlag, Berlin, 1988.
13. S.K. Ghosh. On Recognizing and Characterizing Visibility Graphs of Simple Polygons. *Discrete and Computational Geometry*, **17** (1997), 143–162.
14. J.E. Goodman and R. Pollack. Semispaces of Configurations, Cell Complexes of Arrangements. *Journal of Combinatorial Theory*, *Series A*, **37** (1984), 84–102.
15. N. Greene, M. Kass and G. Miller. Hierarchical Z-buffer Visibility. *SIGGRAPH Proceedings*, pp. 231–238, 1993.
16. K. Kumar. Combinatorial Aspects of Point Visibility. Ph.D. Dissertation, Department of Computer Science, Texas A&M University, August 1993.
17. J. Lawrence. Oriented Matroids and Multiple Ordered Sets. *Linear Algebra and its Applications*, **48** (1982), 1–12.
18. T. Lozano-Perez and M.A. Wesley. An Algorithm for Planning Collision-Free Paths Among Polyhedral Obstacles. *Communications of the ACM*, **22** (1979), 560–570.

19. N.E. Mnev. The Universality Theorems on the Classification Problem of Configuration Varieties and Convex Polytope Varieties, pp. 527–544. Lecture Notes in Mathematics, 1346. Springer-Verlag, New York, 1988.
20. J. O'Rourke. *Art Gallery Theorems and Algorithms*. Oxford University Press, Oxford, 1987.
21. J. O'Rourke and I. Streinu. Vertex-Edge Pseudo Visibility Graphs. *Proceedings of the ACM Symposium on Computational Geometry*, pp. 119–128, Nice, 1997.
22. P.O. Shor. Stretchability of Pseudolines Is NP-Hard, pp. 531–554. DIMACS Series in Discrete Mathematics and Theoretical Computer Science, 4. AMS, Providence, RI, 1991.
23. S. Teller and P. Hanrahan. Global Visibility Algorithms for Illumination Computations. *SIGGRAPH Proceedings*, pp. 239–246, 1993.

Received April 22, 2001, *and in revised form November* 7, 2001. *Online publication October* 29, 2002.