Smith ScholarWorks

Computer Science: Faculty Publications | Computer Science

1-1-1997

Vertex-Edge Pseudo-Visibility Graphs: Characterization and Recognition

Joseph O'Rourke
Smith College, jorourke@smith.edu

Ileana Streinu
Smith College, istreinu@smith.edu

Follow this and additional works at: https://scholarworks.smith.edu/csc_facpubs

Part of the Computer Sciences Commons

Recommended Citation

O'Rourke, Joseph and Streinu, Ileana, "Vertex-Edge Pseudo-Visibility Graphs: Characterization and Recognition" (1997). Computer Science: Faculty Publications, Smith College, Northampton, MA.
https://scholarworks.smith.edu/csc_facpubs/328

This Conference Proceeding has been accepted for inclusion in Computer Science: Faculty Publications by an authorized administrator of Smith ScholarWorks. For more information, please contact scholarworks@smith.edu

Vertex-Edge Pseudo-Visibility Graphs: Characterization and Recognition

Joseph O'Rourke
Ileana Streinu*

Abstract

We extend the notion of polygon visibility graphs to pseudo-polygons defined on *generalized configurations of points*. We consider both vertex-to-vertex, as well as vertex-to-edge visibility in pseudo-polygons.

We study the characterization and recognition problems for vertex-edge pseudo-visibility graphs. Given a bipartite graph G satisfying three simple properties, which can all be checked in polynomial time, we show that we can define a generalized configuration of points and a pseudo-polygon on it, so that its vertex-edge pseudo-visibility graph is G. This provides a full characterization of vertex-edge pseudo-visibility graphs and a polynomial-time algorithm for the decision problem. It also implies that the decision problem for vertex visibility graphs of pseudo-polygons is in NP (as opposed to the same problem with straight-edge visibility, which is only known to be in PSPACE).

1 Introduction

Characterizing visibility graphs has remained an elusive problem [O'R93]. Ghosh [Gho88, Gho97] proposed a set of necessary conditions as a starting point. Everett [Eve90] proved their insufficiency and proposed new conditions. She also placed the recognition problem in PSPACE by reducing it to the existential theory of the reals. Abello and Kumar [AK95] expanded the set of conditions and first related the problem with oriented matroid theory. Their conditions, plus realizability (stretchability) of a certain oriented matroid associated with the graph, would allow a characterization of visibility graphs. But realizability of oriented matroids is a very strong condition. It has been shown by Mnëv [Mnë91] (see also Shor [Sho91]) that it is as hard as the existential theory of the reals, for which so far only exponential algorithms are known. In this context the problem seems too hard to attack.

In this paper we deal with the characterization and recognition problems for a class of polygon visibility graphs, the ve-graphs introduced in [OS97]. Elsewhere we deal with the reconstruction (drawing) problem [Str96a].

Our approach introduces two innovations. First, instead of the usual vertex-vertex visibility graph (v-graph), we study the *vertex-edge visibility graph* (ve-graph) G_{VE} of a polygon. We introduced this concept in a previous paper [OS97]. There we showed that this combinatorial structure contains more geometric information than the vertex visibility graph for straightline polygons. Second, we generalize the notion of straightline visibility to visibility along pseudolines. Here we reinterpret and clarify some of the results of Abello and Kumar [AK95], as well as taking them a step farther. We mix our two ideas by starting with the ve-graph of a pseudo-polygon, and showing that such graphs can be characterized by three simple properties recognizable in polynomial time. It follows that recognition of vertex visibility graphs for pseudo-polygons is in NP. In a companion paper [Str96b], it is shown that the class of straightline ve-graphs is properly contained in the class of pseudoline ve-graphs, and similarly for v-graphs. In particular it follows that Abello and Kumar's oriented matroids may not be stretchable, and thus their characterization fails to completely capture straightline v-graphs. Similarly, our characterization of pseudo ve and v-graphs would be incomplete if taken to their straightline counterparts.

By stating the problem in the pseudo-visibility context, we isolate the combinatorial (oriented matroid) structure of the problem from the stretchability issue, in a manner similar to [AK95]. Focus on ve-graphs rather than v-graphs, and explicit focus on pseudo-

*Department of Computer Science, Smith College, Northampton, MA 01063, USA.
{orourke,streinu}@cs.smith.edu. Supported by NSF grant CCR-9421670.

Permission to make digital/hard copies of all or part of this material for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication and its date appear, and notice is given that copyright is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires specific permission and/or fee
Computational Geometry 97 Nice France
Copyright 1997 ACM 0-89791-878-9/97/06 ..$3.50

visibility, result in a simpler set of conditions than have been obtained previously. The relative simplicity of these conditions has led, as just mentioned, to the first examples of non-stretchable visibility graphs [Str96b].

The paper has two main sections. In Section 2, we start with a pseudo-polygon and its ve-graph, and derive its properties. These properties are all relatively unsurprising, direct generalizations of those established for straight-line ve-graphs in [OS97]. The definition of vertex-edge visibility in the pseudoline context is, however, not completely straightforward. In order to concentrate on the more novel characterization, we do not include proofs in Section 2; see [OS96, OS97]. In Section 3, we start from abstract properties of the graph and construct a pseudo-polygon that realizes it. This involves the construction of an acyclic uniform rank-3 oriented matroid. We establish that the matriod has the claimed properties using Knuth CC-system axioms [Knu92]. Other systems of axioms (such as co-circuits) have been considered in previous versions of this paper, but they led to a much more involved case analysis.

2 Pseudo-Visibility in Pseudo-Polygons

Generalized Configurations of Points

Our generalization of straightline visibility to pseudo-visibility depends the the notion of a "generalized configurations of points" introduced by Goodman and Pollack [GP84].[1] Recall that an *arrangement of pseudolines* $\mathcal{L}$ is a collection of simple curves, each of which separates the plane, such that each pair of lines of $\mathcal{L}$ meet in exactly one point, where they cross.

Definition 2.1 *Let* $V = \{v_0, v_1, \ldots, v_{n-1}\}$ *be a set of points in the Euclidean plane* $\mathbb{R}^2$*, and let* $\mathcal{L}$ *be an arrangement of* $\binom{n}{2}$ *pseudolines such that every pair of points* v_i *and* v_j *lie on exactly one pseudoline* $l_{ij} \in \mathcal{L}$*, and each pseudoline in* $\mathcal{L}$ *contains exactly two points of* V*. Then the pair* $(V, \mathcal{L})$ *is a* generalized configuation of points *in general position.*

The phrase "in general position" indicates that no three points of V lie on one line of $\mathcal{L}$.

Pseudo-Polygon

Two points a and b on a pseudoline $l \in \mathcal{L}$ determine a unique (closed) *segment ab* consisting of those points on l that lie between the two points. For $0 \leq i \leq n-1$, let $e_i = \overline{v_i v_{i+1}}$ be the segment determined by v_i and v_{i+1} on $l_{i,i+1}$.[2]

Definition 2.2 *The segments* $e_i = v_i v_{i+1}$ *form a pseudo-polygon iff:*

1. *The intersection of each pair of segments adjacent in the cyclic ordering is the single point shared between them:* $e_i \cap e_{i+1} = v_{i+1}$*, for all* $i = 0, 1, \ldots, n-1$.
2. *Nonadjacent segments do not intersect:* $e_i \cap e_j = \emptyset$*, for all* $j \neq i + 1$.

See Fig. 1a for an example. Throughout we let P denote a pseudo-polygon, with V its vertices labeled $V = (v_0, v_1, \ldots, v_{n-1})$ in counterclockwise (ccw) order, and E its set of edges similarly labeled. Note edges are closed segments. We use the term *exterior* to designate points of the plane strictly exterior to P (and so not on its boundary).

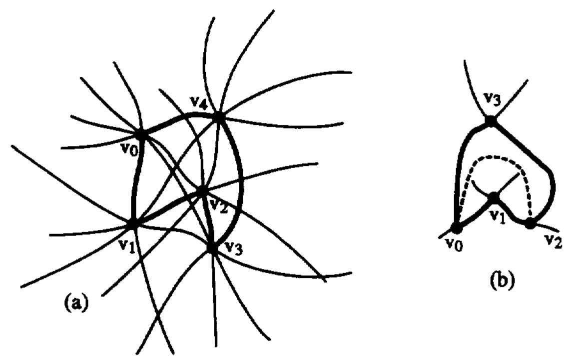

Figure 1: (a) A pseudo-polygon (not every intersection between pseudolines is shown); $|V| = 5$, $|\mathcal{L}| = 10$. (b) v_0 cannot see v_2.

Vertex-vertex Pseudo-Visibility

Pseudo-visibility is determined by the underlying arrangement $\mathcal{L}$: lines-of-sight are along pseudolines in $\mathcal{L}$.

Definition 2.3 *Vertex* v_i sees *vertex* v_j $(v_i \leftrightarrow v_j)$ *iff either* $v_i = v_j$*, or they lie on a line* $l_{ij} \in \mathcal{L}$ *and the segment* $v_i v_j$ *is nowhere exterior to* P.

Note that our definition of pseudo-visibility is dependent upon $\mathcal{L}$: it does not make sense to ask if two points of V see one another without providing the underlying arrangement $\mathcal{L}$. Dependence upon $\mathcal{L}$ means there is not complete freedom to assign which vertex sees which. For example, in Fig. 1b, v_0 could not be arranged to see v_2, because the pseudoline l_{02} would have to intersect l_{01} (and l_{12}) twice, violating the definition of a pseudoline arrangement.

Definition 2.4 *The* vertex pseudo-visibility graph $G_V(P)$ *of a polygon is a labeled graph with node set* V*, and an arc between two vertices iff they can see one another (according to Def. 2.3).*

[1] Their definition is for the projective plane, and includes a special line l_∞. We use the Euclidian plane.

[2] All index arithmetic is mod n throughout the paper.

We will often abbreviate $G_V(P)$ to G_V. Note that G_V is Hamiltonian: the arcs corresponding to the polygon boundary form a Hamiltonian circuit $(v_0, \ldots, v_{n-1})$. And also note that since G_V is labeled by V, which we assumed was labeled in a ccw boundary traversal order, the Hamiltonian circuit is provided by the labeling of the graph.

Vertex-edge Pseudo-Visibility

We need to define when a vertex sees an edge. In [OS97], we defined v to see e in a straightline context if v sees an open interval of e. Here we extend this notion to the pseudo-visibility context without adding any new points and pseudolines to the generalized configuration, to keep the definitions purely combinatorial. We start with the notion of a "witness." Let $r_j^i \subset l_{ij}$ be the ray along l_{ij} starting at and including v_j, directed away from (and therefore excluding) v_j.

Definition 2.5 *Vertex v_j is a* witness *for the vertex-edge pair (v_i, e) iff either*

1. *v_i is an endpoint of e, and v_j is also (here we permit $v_j = v_i$); or*
2. *v_i is not an endpoint of e, and*
   - *(a) v_i sees v_j; and*
   - *(b) the ray r_j^i intersects e at a point p,*
   - *(c) either $v_j = p$, or the segment v_jp is nowhere exterior.*

We will refer to the line l_{ij} in the above definition as the *witness line*.

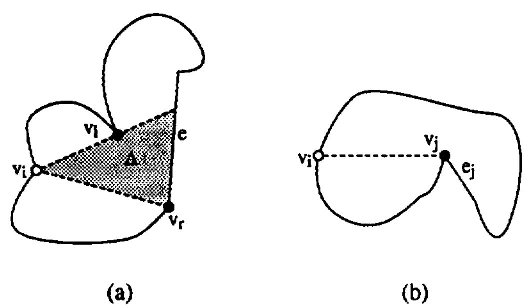

Figure 2: (a) v_i sees e; (b) v_i does not see e_j although v_j is a witness.

Definition 2.6 *Vertex v* sees *edge e ($v \rightarrow e$) iff there are at least two witnesses v_r and v_l for (v, e).*

We call v_r the *right witness* and v_l the *left witness* if their ccw ordering is (v_r, e, v_l). Thus v_r is to the right from the viewpoint of v, and v_l to the left. Note that an endpoint of an edge sees that edge, because then both endpoints are witnesses.

The intent of this definition is illustrated in Fig. 2a: v_i sees an open interval of e. The reason we demand two witnesses is that one witness does not suffice, as is clear from Fig. 2b.

Lemma 2.7 *Under the general position assumption, if v_i sees e, there are exactly two witnesses for (v_i, e).*

Definition 2.8 *The* vertex-edge pseudo-visibility graph *(ve-graph) G_{VE} of a polygon is a labeled bipartite graph with node node set $V \cup E$, and an arc between $v \in V$ and $e \in E$ iff v can see e (according to Def. 2.5).*

Ve-graph Properties

Our aim now is to obtain characterizing properties of G_{VE}. The key property concerns how a gap in v_k's view of the polygon's boundary can occur: it can only occur in one of the two ways illustrated in Fig. 3. We first state this condition using both v-v and v-e visibility, and later (Theorem 2.13) remove the v-v information.

We use the following notation to specify parts of the polygon boundary: $P[i,j]$ is the closed subset of the polygon boundary ccw from v_i to v_j. $P(i,j]$ excludes v_i; $P[i,j)$ excludes v_j; and $P(i,j)$ excludes both.

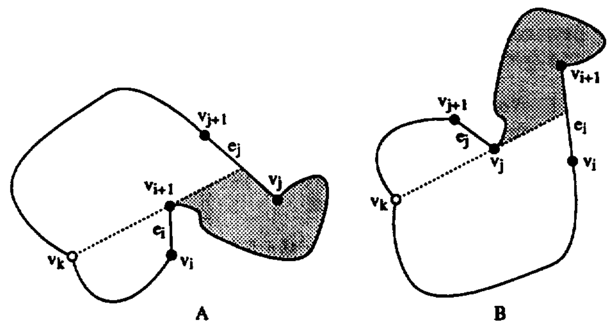

Figure 3: Two cases for v_k seeing e_i followed by e_j.

Lemma 2.9 *If v_k sees non-adjacent edges e_i and e_j and no edge between, $v_k \in P[j+1, i]$, then exactly one of Case A or B holds (see Fig. 3):*

Case A

1. *v_k sees v_{i+1} but not v_j; and*
2. *v_{i+1} is the right-witness for (v_k, e_j); and*
3. *v_{i+1} sees e_j but v_j does not see e_i.*

Case B

1. *v_k sees v_j but not v_{i+1}; and*
2. *v_j is the left-witness for (v_k, e_i); and*
3. *v_j sees e_i but v_{i+1} does not see e_j*

The structure established by the preceding lemma is best captured by the notion of "pockets":

Definition 2.10 *If v_i sees e_j and v_r and v_l are the right and left witnesses respectively (cf. Fig. 4), then $P[i, r)$ and $P(l, i]$ are the right and left* near pockets *of (v_i, e_j), and $P(r, j]$ and $P[j+1, l)$ are the right and left* far pockets *of (v_i, e_j), respectively.*

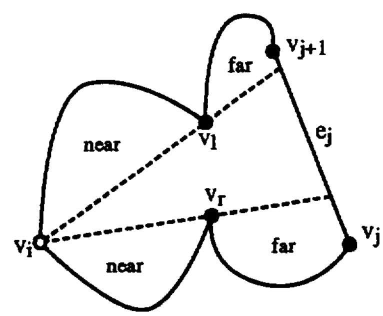

Figure 4: Definition of near and far pockets.

Note that: v_i is in both near pockets; e_j is not part of either far pocket, and the witnesses are not in any of the four pockets. If either witness is an endpoint of e_j, then the corresponding pocket is empty.

Lemma 2.11 *If v_i sees e_j and v_r and v_l are the right and left witnesses respectively, then*

1. *No vertex in the right near pocket sees an edge in the right far pocket.*
2. *No vertex in the right far pocket sees an edge in the right near pocket.*

Symmetric claims hold for the left pockets.

Lemma 2.11 leads immediately to Lemma 2.12, which conveys the same import in more graph-theoretic terms.

Lemma 2.12 *If v_i sees e_j and v_r and v_l are the right and left witnesses respectively, then v_r is an articulation point of the subgraph of G_{VE} induced by $P[i, j]$, and symmetrically, v_l is an articulation point of the subgraph induced by $P[j+1, i]$.*

We may now state our characterization of ve-graphs by discarding from some of the previous lemmas all but vertex-edge visibility information.

Theorem 2.13 *If G_{VE} is the vertex-edge visibility graph of a pseudo-polygon P, then it satisfies these two properties:*

1. *If v_k sees non-adjacent edges e_i and e_j and no edge between, $v_k \in P[j+1, i]$, then exactly one of these holds:*
   - **A.** $(v_{i+1}, e_j) \in G_{VE}$, *or*
   - **B.** $(v_j, e_i) \in G_{VE}$.
2. *In the two cases above, additionally:*
   - **A.** *v_{i+1} is an articulation point of the subgraph of G_{VE} induced by $P[k, j]$.*
   - **B.** *v_j is an articulation point of the subgraph of G_{VE} induced by $P[j+1, k]$.*

We will prove in Section 3 that these properties basically provide a complete characterization of vertex-edge visibility graphs.

Information in the Vertex-Edge Visiblity Graph

We have established the key low-level properties of G_{VE} in Theorem 2.13, but they give little insight into higher-level properties of the graph. To carry out the proof of Theorem 3.2 in the next section, we will need to derive from the lower-level concepts a number of additional combinatorial concepts, analogs of the geometric notions of convexity of polygon vertices, partial local sequences, and shortest paths trees. These have been shown to be derivable from G_{VE} in the straight-line case in [OS97], and the generalization to pseudo-visibility is along similar lines. We will omit most of the details here and give just the definitions that are needed to understand the abstract counterparts introduced later in the proof of the main theorem. See [OS96] for details.

Given a set of points in the plane, rotate a directed line around each point and record the ordered list of the other points as they are encountered by the rotating line. In addition, assign a sign to each point: positive if it is encountered by the forward ray from the center of rotation, negative if by the backward ray. The infinite sequence thus obtained is called the *i-sequence* for the vertex. It is periodic, fully characterized by one half-period. The half-period is a signed permutation of all the vertices different from the point of rotation, $\alpha_1\alpha_2\cdots\alpha_{n-1}$. A circular rotation of this permutation, with a change of the sign of the element sent from the beginning to the end of the permutation, also characterizes the same *i*-sequence.

As an example, consider the points in Fig. 5, ignoring (temporarily) the polygon boundary, and imagining a full complement of fully-extended pseudolines. The *i*-sequences of the points could be as follows, where negative points are indicated with a bar:

v_0 : $214\bar{3}\bar{6}5$
v_1 : $2\overline{506}34$
v_2 : 341506
v_3 : 506412
v_4 : $50\bar{3}612$
v_5 : 602143
v_6 : 214305

The collection of *i*-sequences for all the points in the set is a version of what Goodman and Pollack [GP84] called a *cluster of stars* (see also [Str96a]).

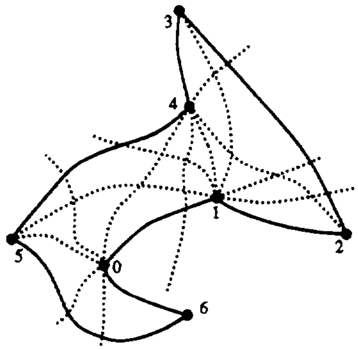

Figure 5: A pseudo-polygon used to illustrate i-sequences.

From an i-sequence, one can read off *chirotope* information, i.e., whether a triple i, j, k makes a right or left turn at j. Using a different terminology (vortex-free tournaments, pre-CC and CC-systems), Knuth [Knu92] has given a system of axioms equivalent to uniform acyclic oriented matroids of rank 3. Knuth's CC-systems can be interpreted as characterizing the i-sequences of generalized configurations of points. We will make use of his results in the proof of our main theorem.

Rotating a directed pseudoline around a vertex in a polygon and recording in a circular (signed) sequence only the visible vertices induces a *partial i-sequence*. The partial i-sequences of a polygon are uniquely determined by the ve-graph, and can be easily computed from it as soon as the convexity properties of the vertices of the polygon are determined. For example, for the vertices of Fig. 5, the partial i-sequences for each vertex are as follows:

v_0 : $14\overline{6}5$
v_1 : $2\overline{50}34$
v_2 : 341
v_3 : 412
v_4 : $50\overline{3}12$
v_5 : 6014
v_6 : 05

The edges of the v-graph meeting at a common vertex may form convex or reflex angles. This information is uniquely determined from the ve-graph via computing an extended relation of visibility between vertices of the polygon and edges of the v-graph.

For any two distinct vertices of a simple polygon, there exists a unique shortest path between them. For a fixed source vertex v_i, the set of shortest paths from v_i to all other vertices induces a *shortest-path tree* (sp-tree) rooted at v_i. The tree is ordered: there is a natural ordering between the subtrees rooted at each vertex, as given by the ccw traversal of the boundary of the polygon. Each internal node of a sp-tree is a reflex vertex and a *turn* is associated with it: left or right, according to how the shortest paths from the root through that vertex turn there.

Next we define two canonical circular orderings of vertices around each vertex of a polygon. These orderings will be used in the proof of the Theorem 3.2. The intuition behind them is that for every pseudo-polygon there exists another pseudo-polygon in a *normal form* having the same ve-graph. We will introduce a concept of i-sequences for a pseudo-polygon: the normal form has the property that the polygon i-sequences are identical to the i-sequences of the underlying generalized configuration of points. The existence of the normal form will be a consequence of the Theorem 3.2.

We define a total ordering on all the internal vertices of the sp-tree rooted at v_i (and hence on all the shortest paths from the root to those vertices). If v_j and v_k are two children of the same node, with v_i, v_j, v_k occuring in this order in a ccw traversal of the boundary of the polygon, then all the vertices in the tree rooted at v_j are listed *before* all the vertices in the tree rooted at v_k. If v_j is a right (left) turn, then vertex v_j itself is listed *after* (*before*) all the children in the subtree rooted at v_j. For example, in Fig. 5, the sp-tree rooted at 1 makes a left turn at vertex 0. The tree rooted at 6 makes a right turn at 0, a left turn at 4, and a right turn at 1. The total order of the vertices of the polygon in the figure derived this way is as follows:

v_0 : 214356
v_1 : 234506
v_2 : 341506
v_3 : 506412
v_4 : 506123
v_5 : 621043
v_6 : 214305

Lastly, combining the total order of the shortest paths (vertices) around each vertex with the partial i-sequence information, one can define a canonical circular order of *signed* vertices around each vertex v_i of the polygon, as follows. If the vertex v_i is convex, the ordering is the same as the one induced by the sp-tree. If it is reflex, take the signed permutation representing its partial i-sequence. Each vertex in this permutation is a child of the root (v_i) in the sp-tree for the current vertex. Create a new signed permutation from the partial i-sequence by replacing each vertex with the ordered list of the vertices in the subtree rooted at v_j, with the same sign as the root v_j. For example, consider $v_i = v_4$ in Fig. 5, a reflex vertex whose partial i-sequence is v_4 : $50\overline{3}12$. We replace $v_j = v_0$ (0 in the sequence) with 06, resulting in v_4 : $506\overline{3}12$.

This canonical circular ordering will be called the *polygon i-sequence* for vertex v_i in P; it is in general different from the i-sequence for vertex v_i in the context of the point configuration of the vertices of the polygon. Note, for example, that the polygon i-sequence for v_4 just obtained is different from the i-sequence of v_4 in the point configuration ($50\overline{3}612$), because 6 was encountered before $\overline{3}$ when spinning about v_4 in the configuration, but all of 0's children are sorted along with 0 without regard to where blocked lines of visibility (e.g., l_{43}) might encounter them. The normal form mentioned previously forces, in this case, v_6 to lie left rather than right of l_{43}.

3 Abstract ve-graphs

We will work with bipartite graphs G_{VE} defined on two circularly ordered lists $V = (v_0, \ldots, v_{n-1})$ of *vertices* and $E = (e_0, \ldots, e_{n-1})$ of *edges*. Most of our terminology and notation carries over from the geometric setting, but we repeat here to emphasize that in the abstract setting, all definitions must be combinatorial. Two edges e_i and e_j are *adjacent* if $j = i+1$ or $i = j+1$; vertex v_i is *adjacent* to the edges e_{i-1} and e_i; and edge e_i is adjacent to its "endpoints" v_i and v_{i-1}. The polygon boundary between vertices v_i and v_j is defined as a list $P[i,j] = (v_i, e_i, v_{i+1}, \ldots, e_{j-1}, v_j)$. As before, $P[i,j)$ and $P(i,j]$ exclude "endpoint" vertices. We use $v_i \to e_j$ as alternate notation for $(v_i, e_j) \in G_{VE}$: *v_i sees e_j*. For a vertex $v_k \in P[j+1, i]$, two edges e_i and e_j are called *consecutive from v_k* if v_k sees e_i and e_j but does not see any $e_l \in P[i,j]$.

Definition 3.1 *An* abstract ve-graph *is a bipartite graph G_{VE} on sets of vertices V and edges E, both circularly labeled as above, with $|V| = |E| \geq 3$, satisfying the two properties of Theorem 2.13, plus the additional condition that each vertex sees its adjacent edges: $\forall i$, $v_i \to e_i$ and $v_i \to e_{i-1}$.*

Our main goal is to prove that this definition captures ve-graphs of pseudo-polygons:

Theorem 3.2 (Main Theorem) *If G_{VE} is an abstract ve-graph, then there exists a generalized configuration of points $\{p_0, \cdots, p_{n-1}\}$ and a pseudo-polygon specified by this ordering of the vertices, whose ve-graph is the same as G_{VE}.*

The plan of the proof is as follows. Our ultimate goal is to prove the configuration of points exists by constructing a corresponding uniform rank-3 acyclic oriented matroid. This requires specifying the orientation of all triples of points. We use shortest paths to define these signed triples. With these in hand, we will need to prove the collection of triples do constitute the appropriate type of matroid. We use Knuth's "counterclockwise" CC-axioms for this purpose. Several of his axioms can be established from the properties of i-sequences; the remaining we prove directly.

Note that we start only with a definition of v-e visibility. From here we will define a notion of v-v visibility, as well as abstract counterparts of convex and reflex vertices, pockets, shortest paths, shortest-path trees, and i-sequences. At this point we will be ready to define the triples and construct the matroid. Most of the definitions through shortest paths mimic in the abstract setting those in [OS97] for straight-line polygons. The additional factor here is that we have to prove that the definitions are consistent and lead to objects satisfying the properties one would expect from them.

We define v-v visibility according to which of Case A and B holds in the defining properties from Theorem 2.13: in Case A, we define $v_k \to v_{i+1}$, and in case B, we define $v_k \to v_j$, just as in Lemma 2.9. If $j = i+1$, we define $v_k \to v_{i+1}$. Moreover, if case A holds, we say that v_{i+1} is a *right* articulation point for $P[k, i+1)$ (the *near*) and $P(i+1, j]$ (the *far*) *right pockets*. If case B holds, v_j is the left articulation point for $P(j,k]$ (the near) and $P[i+1, j)$ (the far) *left pockets*. In either case, the pockets and articulation points are relative to visibility from v_k (or to the triple v_k, e_i, e_j).

Lemma 3.3 *$\forall i$, v_i sees at least one other edge e_j different from e_{i-1} and e_i.*

Lemma 3.4 *$\forall i$, exactly one of the following two cases holds:*

A. *$v_{i-1} \to e_i$ and $v_{i+1} \to e_{i-1}$, or*

B. *$v_{i-1} \not\to e_i$ and $v_{i+1} \not\to e_{i-1}$.*

In the first case we say v_i is *convex*, and in the second *reflex*.[8]

Proof: Assume this is not true. Then either $v_{i-1} \to e_i$ but $v_{i+1} \not\to e_{i-1}$ or $v_{i-1} \not\to e_i$ but $v_{i+1} \to e_{i-1}$. Assume the first case holds (the other can be treated similarly). Then, taking v_{i+1} to play the role of v_k in property 1 from Theorem 2.13, we know that $v_{i+1} \to e_i$ by Def. 3.1. Because $v_{i+1} \not\to e_{i-1}$, there must be an edge e_j, $j \in P[i+2, i-1)$, such that v_{i+1} sees e_j and e_i consecutively (such an edge must exist, by Lemma 3.3). See Fig. 6. Moreover, Case **B** of Theorem 2.13 must hold. Then e_{i-1} lies in the far pocket for the triplet v_{i+1}, e_j, e_i. But then, by property 2, it cannot be the case that v_{i-1}, which is in the far pocket, sees e_i, which is in the near pocket. This contraction establishes the lemma. □

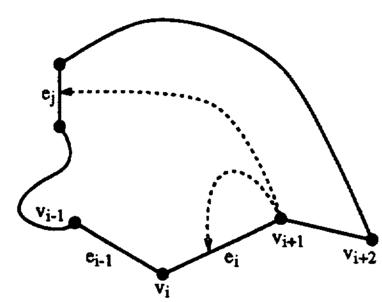

Figure 6: Lemma 3.4: $v_{i+1} \to e_i$ and $v_{i+1} \to e_j$ and v_{i+1} sees no edge between.

The next lemma shows that the articulation point properties of the ve-graph carry through to v-v visibility.

[8] Recall our general position assumption removes worry about the intermediate case.

Lemma 3.5 *No vertex in a near pocket can see any vertex in the corresponding far pocket and vice-versa.*

The next step is to show the v-v visibility is symmetric:

Lemma 3.6 *If $v_i \to v_j$ then $v_j \to v_i$.*

This lemma guarantees that the v-graph associated to an abstract ve-graph is indeed an undirected graph, and is the base case for the inductive proof of the more general property of symmetry of paths between pairs of vertices v_i and v_j (Lemma 3.15).

As an aside, note that some of the properties that we prove can be found as axioms in [AK95] (e.g., path symmetry is their "Necessary Condition 2"): it is indeed surprising that our simple conditions are enough to imply these high level properties. The property in the following lemma is one of their defining axioms.

Lemma 3.7 *If $v_k \to v_i$ and $v_k \to v_j$ consecutively, then $v_i \to v_j$.*

From the ve-graph we can also define an extended relation of visibility between vertices and v-graph edges, as well as an abstract angle (convex or reflex) between two v-edges with a common endpoint. We say that vertex v_k sees visibility edge v_iv_j if $v_k \to e_l$ and $i \in P(k,l]$, $v_j \in P[l+1,k)$. We define the abstract *angle* $\angle v_iv_kv_j$ between two adjacent v-edges v_kv_i and v_kv_j as being *convex* if $v_i \to v_kv_j$ and $v_j \to v_kv_i$, reflex if none of these two conditions holds. Then we can construct the abstract partial i-sequence for a vertex v_k using the following algorithm.

Algorithm 1 (Construction of the partial i-sequence for vertex v_k from convex/reflex angle information)

Start with vertex v_{k+1} in a list.

Record in the list the consecutive vertices v_i for which $\angle v_{k+1}v_kv_i$ is convex. Let v_m be the last such vertex. If we have not yet reached v_{k-1} (i.e., $m \neq k-1$), continue with the rest of the vertices v_j for which the angle $\angle v_{k+1}v_kv_j$) is reflex. For each such vertex v_j, determine a pair (v_{i_1}, v_{i_2}) of previously listed consecutive vertices, $v_{i_1} \in P[k+1,m)$, $v_{i_2} \in P(i_1,m]$, so that v_j does not see v_{i_1} and sees v_{i_2} consecutive to v_k. Then, if two consecutive vertices v_{j_1} and $v_{j_2} \in P(m,k-1]$ lie between the same pair of consecutive vertices v_{i_1}, v_{i_2}, list them in the same order in the partial i-sequence, between v_{i_1} and v_{i_2}. All the vertices v_i with $v_{k+1}v_kv_i$ convex are listed as positive, the others as negative, in the signed partial i-sequence.

Some simple lemmas guarantee that the above definition is consistent and the algorithm is correct.

The next result may seem rather trivial, although the proof is not.

Lemma 3.8 *There exists at least one convex vertex.*

We have defined pockets and articulation points earlier. The following lemma shows that pockets embed nicely: this is the main tool used to define abstract shortest paths and sp-trees.

Lemma 3.9 (Subpocket embedding)

Let $v_i \to v_k$ with v_k a right articulation point for visibility from v_i. Let e_m be the next edge visible from v_i after v_k and let v_j be in the far pocket $P(k,m]$. Assume that $v_k \not\to v_j$ and let v_j be in a far pocket P' from v_k. Then $P' \subset P(k,m]$. A similar embedding property holds when v_k is a left articulation point.

For example, in Fig. 5, consider $i=6$, $k=0$, $m=4$, and $j=3$: $v_i = v_6 \to v_k = v_0$, $v_k = v_0 \not\to v_j = v_3$, and $v_j = v_3 \in P(0,4]$. Then $P' = P[3,4) \subset P(0,4]$.

We need here some terminology which will help in formulating and proving symmetry properties for pockets and shortest paths. If v_j is in the right far pocket $P(k,m]$ from v_i and in the right far pocket $P(p,s]$ from v_k we will say that the embedding pattern of the far pocket containing v_j invisible from v_i is of type RR, with first articulation point v_k of type R and second articulation point v_p of type R. Similarly we define type RL, LR, LL embedding patterns and articulation points, according to whether the first pocket is right and the second left, etc. The following lemma shows that there is a symmetry in the embedding patterns seen from the two endpoints. In what follows we will also use the notation $\overline{R} = L$ and $\overline{L} = R$. The following corrollary is a direct consequence of the proof of Lemma 3.9.

Corollary 3.10 (Pocket symmetry) *If $v_i \not\to v_j$ and the embedding pattern of the far pockets from v_i to v_j is AB $(A, B \in \{L, R\})$ with articulation points v_k of type A and v_p of type B in this order, then the embedding pattern of the pockets from v_j to v_i is reversed, BA, with articulation points v_p of type $\overline{B}$ and v_k of type $\overline{A}$ in this order.*

The previous lemmas allow us to correctly define the *pocket embedding tree* rooted at a vertex v_i. This is a tree with a fixed ordering on the children of each internal node and with a sign (corresponding to a right/left turn) associated to each node other than the root. The tree has exactly n nodes. Each node is labelled with a signed vertex of V, and the set of vertices in a subtree correspond to the set of vertices in a subpocket with articulation point given by the vertex labelling the root of the subtree. Moreover, the sign of the root of the subtree is right/left, corresponding to whether it is a left or right articulation point for visibility from its parent.

The pocket embedding tree is defined recursively as follows. Its root is the vertex v_i and its children correspond to the far pockets from v_i in the order induced by the circular order of V. Each node is labelled by the corresponding articulation point and signed by its right/left type. Lemma 3.9 ensures that once we get into a pocket we can continue subdividing the pockets by looking at the far subpockets of invisibility from the root of the subtree, restricted only to the vertices in the current pocket (subtree). This in turn guarantees that this process generates a tree.

The labels of the paths from the root v_i to the other nodes of this tree correspond intuitively (i.e., in the case of a polygon) to the shortest paths from v_i to all the other vertices of the polygon. This motivates our using the terminology *shortest paths* for them in the combinatorial setting of abstract ve-graphs. Let v_i and v_j be two vertices of G_{VE}. We will define the *abstract shortest path* from v_i to v_j, $sp(v_i, v_j)$ as an ordered list of vertices starting with v_i and ending with v_j. In what follows, *concat* denotes the function that concatenates two lists and *reverse* the function that reverses an ordered list.

The following definition is just a rigorous formalization of this concept.

$sp(v_i, v_j)$ is constructed recursively as follows.

if $v_i \to v_j$ **then** $sp(v_i, v_j) = (v_i, v_j)$

else $sp(v_i, v_j) = (v_i)$ *concat* $sp(v_k, v_j)$, where v_k is the articulation point of the far pocket from v_i containing v_j.

The shortest paths contain more information that just the simple ordered list of vertices from a source to a destination. They also capture the *left* or *right* turns on the path, at articulation points along the way. First let us fix some notation for this. If $sp(v_i, v_j) = (v_i, v_k, \cdots, v_j)$ and v_k is a right (left) articulation point for v_i, then we say that the shortest path $sp(v_i, v_j)$ *makes a right (left) turn* at v_k. From v_k on, the left/right turns are defined on the subsequent subpaths towards v_j. We will denote by $ssp(v_i, v_j)$, the *signed shortest path* from v_i to v_j, to be a signed list obtained from $sp(v_i, v_j)$ and attaching the appropriate signs to its vertices (+ for right turn, − for left turn). The following lemma shows that the signs are already determined by the shortest path.

Lemma 3.11 *The right/left turns on a shortest path are determined by the indices of the articulation points. More precisely, if v_k is a right articulation point from v_i to the far pocket containing v_j, then all the vertices $v_{j'}$ in the far pocket have indices $j' > k$ (circularly, i.e. $v_j' \in P(k, i)$). If it is a left articulation point, all the vertices $v_{j'}$ in the far pocket have indices $j' < k$ (circularly, i.e. $v_{j'} \in P(i, k)$).*

Lemma 3.12 *The signed list of vertices on a path from the root v_i in a pocket tree rooted at v_i to an internal node v_j is equal to the $ssp(v_i, v_j)$.*

The following lemmas show that the shortest paths glue together nicely. We note that these properties were among the axioms in [AK95].

Lemma 3.13 If $v_k \in sp(v_i, v_j)$ then $ssp(v_i, v_k)$ *is the signed sublist of $ssp(v_i, v_j)$ starting at v_i and ending at v_k.*

Lemma 3.14 If $v_k \in sp(v_i, v_j)$ then $sp(v_i, v_k)$ concat $sp(v_k, v_j) = sp(v_i, v_j)$.

Lemma 3.15 $sp(v_i, v_j) =$ *reverse* $sp(v_j, v_i)$.

To summarize, up to this point we have defined for every vertex v_i of V a tree rooted at v_i and labelled with all the vertices of V. There is an ordering of the children of all internal nodes and a sign (*turn*) associated with each internal node. The paths from the root to any vertex define *abstract shortest paths* from the root to that vertex. Two shortest paths trees are compatible in the sense of Lemma 3.13. We will now define a circular ordering of the shortest paths $sp(v_i, v_j)$ around a vertex v_i by combining the partial i-sequences for v_i with the shortest path tree rooted at v_i, exactly as described at the end of Section 2.

The last step is to define a predicate on any triple of vertices of V and show that it satisfies Knuth's CC-system axioms. It has been shown by Knuth that CC-systems are equivalent to uniform rank-3 acyclic oriented matroids [Knu92, p. 40], and it is well known (see e.g., [BLW$^+$93]) that these in turn are equivalent with Goodman and Pollack's generalized configurations of points (in general position).[4]

We will write $i < j < k$ iff the indices i, j and k occur in this order in V (as usual, indices are taken mod n).

The predicate will be denoted as ijk. It is first defined for indices $i < j < k$ then extended as usual for other permutations of three points: odd number of inversions change the sign, even number keep the same sign (see Knuth's axioms on next page). We say that ijk holds iff in the i-sequence for vertex v_i, the vertices v_j and v_k appear positively in this order in a half-period. For example, the polygon i-sequence for v_1 in Fig. 5 is v_1 : 234506. Thus $5 < 6$ occur in this order in the sequence, and so 156 holds.

The following lemma matches the chirotope definition in [AK95].

Lemma 3.16 *If $i < j < k$ then ijk holds iff i, j, k do not occur on a common abstract shortest path. Equivalently, $\neg ijk$ iff i, j, k belong to the same shortest path.* Continuing with the same example, 156 holds because $1 < 5 < 6$ and they do not occur on a shortest path; and $\neg 160$ because $1 < 6 < 0$ and they occur on the shortest path (v_1, v_0, v_6).

Proof: We first prove sufficiency. Let $i < j < k$ and ijk, and assume for contradiction that i, j, k are on the same shortest path. Then either $j \in sp(i, k)$, in which case j has to be a right turn and contradicts the definition of the i-sequence for vertex i (k would occur before j and not vice-versa); or $i \in sp(k, j)$ which again implies that in the i-sequence for i, j is followed by $\overline{k}$, contradicting the hypothesis; or $k \in sp(j, i)$, which again implies that k occurs after j in the i-sequence for i, contradiction.

To prove the necessity, assume i, j, k are not on the same shortest path. Then in the sp-tree rooted at i, j and k are on different branches, and if those branches originate at i, they do not form a reflex angle. The definition of the i-sequences then guarantees that a vertex of smaller index (j) is encountered in the i-sequence before a vertex of higher index (k). □

Note that the conditions given by this lemma constitute the definition in [AK95] for their oriented matroid defined via chirotope axioms. We have found that verifying these or any other equivalent oriented

[4] The modifier "uniform" indicates general position.

matroid axioms involves a tedious case analysis, which we circumvent via i-sequences and an appropriate interpretation of Knuth's CC-axioms [Knu92, p. 4]:

Axiom 1. (Cyclic symmetry): $ijk \Rightarrow jki$.

Axiom 2. (Antisymmetry): $ijk \Rightarrow \neg ikj$.

Axiom 3. (Nondegeneracy): $ijk \vee ikj$.

Axiom 4. (Interiority): $ijk \wedge ikl \wedge ilj \Rightarrow jkl$.

Axiom 5. (Transitivity): $ijk \wedge ijl \wedge ijm \wedge ikl \wedge ilm \Rightarrow ikm$.

To finish the proof we need to verify that Knuth's axioms hold, that we can define a pseudo-polygon on the corresponding generalized configuration of points (i.e., edges do not cross), and then show that its ve-graph is exactly the graph that we started with. First, we notice that the five axioms of Knuth can be grouped into three categories. Axioms 2, 3, and 5 are equivalent to the existence of circular i-sequences for each point of the system (this is a consequence the corrolary on [Knu92, p. 12]), which we have already established. Axiom 1 requires that the signs of triples do not change under circular permutations of the points, which is a direct consequence of Lemma 3.16. This leaves Axiom 4, the acyclicity condition, which ensures that the points can be realized as an affine generalized configuration of points.

Lemma 3.17 *Axiom 4 holds:* $ijk \wedge ikl \wedge ilj \Rightarrow jkl$.
Proof: Let a, b, c and d be four arbitrary points occuring in this order in V: $a < b < c < d$. Without loss of generality we can assume that $i = a$. The symmetry in Axiom 4 tells that we have only two cases to verify: (1) $j = b, k = c, l = d$, or (2) $j = b, l = c, k = d$.

Case 1. $i < j < k < l$. The premiss, interpreted through Lemma 3.16, says that: i, j, k and i, k, l are not on the same shortest path, and, using Axiom 2 ($ilj \Rightarrow \neg ijl$), that i, j, l are on the same shortest path. The consequent says that j, k, l are not on the same path. Assume for the sake of a contradiction that j, k, l are on the same path. Three cases are possible: j on $sp(l, k)$ and is a right turn; k on $sp(j, l)$ and is a right turn; or l on $sp(k, j)$ and is a right turn. In the first case (Fig. 7a), i is in the near right pocket from l, and k in the far right pocket from l, contradicting the first premiss ijk. In the second case, we have to use the premiss $\neg ijl$ to get a contradiction. Either j is on $sp(i, l)$ (Fig. 7b), in which case it follows that k is on $sp(i, l)$, contradicting ikl; or i on $sp(l, j)$, in which case either ijk or ikl have to be on the same sp, contradicting the first premisses; or, finally, l on $sp((j, i)$. In this case we obtain a contradiction with ikl. The third case may be treated similarly.

Case 2. $i < j < l < k$ is treated similarly. □

This concludes the proof that the i-sequences defined from the ve-graph satisfy the CC-system axioms of Knuth, and hence they form an affine generalized configuration of points in general position. We now have to verify that the order of V induces a pseudo-polygon.

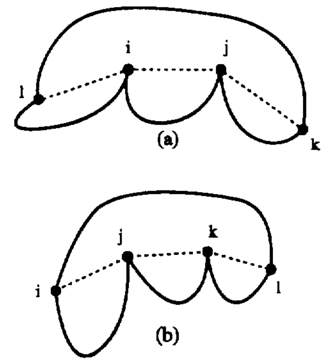

Figure 7: Lemma 3.17, $i < j < j < l$, i, j, l on a shortest path: (a) $j \in sp(l, k)$; (b) $k \in sp(j, l)$ and $j \in sp(i.l)$.

Lemma 3.18 *No two edges* $v_i v_{i+1}$ *and* $v_j v_{j+1}$ *of* $\overline{P}$ *cross.*
Proof: Assume the contrary. Then (1) v_j and v_{j+1} are on separate sides of $v_i v_{i+1}$, and (2) v_i and v_{i+1} are on separate sides of $v_j v_{j+1}$. There are two ways (1) can happen: either v_i sees e_j after e_i, in which case $v_{i+1} \rightarrow e_j$, so $v_{i+1} v_j v_{j+1}$ is a left turn, contradicting (2); or v_{i+1} sees e_j, in which case a contradiction is similarly derived. □

We will denote by $\overline{P}$ the pseudo-polygon obtained this way from the generalized configuration of points associated with G_{VE}.

Lemma 3.19 *The ve-graph of* $\overline{P}$ *coincides with* G_{VE}.

This is true since the construction of the generalized configuration of points preserved all the vertex-to-edge visibilities and invisibilities.

To conclude, let us notice that the conditions characterizing abstract ve-graphs can be easily verified in polynomial time. This proves the main result that the recognition problem for ve-graphs is in P. Also, as a consequence of the relationship between ve-graphs and v-graphs, if one is given a graph one can easily "guess" a ve-graph and verify that it matches the given one and satisfies the ve-graph properties. This places the pseudo v-graph recognition problem in NP.

Finally we mention some open problems raised by our work:

- *Given a v-graph* G_V*, in polynomial time find a ve-graph* G_{VE} *compatible with it (in the sense that there is a pseudo-polygon* P *such that* G_V *and* G_{VE} *are its pseudo-visbility v- and ve-graphs respectively), or show none exists.* Such a polynomial-time algorithm would place pseudo v-graph recognition problem in P.
- *Characterize the realizable rank-3 acyclic uniform oriented matroids produced by the construc-*

tion in Theorem 3.2. The class of oriented matroids obtained from the construction in Theorem 3.2 is a strict subclass of all the acyclic uniform rank-3 oriented matroids. It has been shown elsewhere [Str96b] that not all are stretchable (i.e., realizable with straight lines). However, this beeing such a restricted class, it might be possible to characterize or recognize them with an algorithm of a complexity smaller than the known PSPACE. The alternative is to show that they are as complex as pseudoline stretchablity.

- *Remove the general position assumption.* A more detailed definition of ve-visibility is necessary to deal with degeneracies. See [OS97] for a hint of the complications.

Acknowledgements

We are grateful to James Abello for explaining his work to us.

References

[AK95] J. Abello and K. Kumar. Visibility graphs and oriented matroids. In R. Tamassia and I. G. Tollis, editors, *Graph Drawing (Proc. GD '94)*, volume 894 of *Lecture Notes in Computer Science*, pages 147–158. Springer-Verlag, 1995.

[BLW+93] A. Björner, M. Las Vergnas, N. White, B. Sturmfe ls, and G. Ziegler. *Oriented Matroids*. Cambridge University Press, Cambridge, 1993.

[Eve90] H. Everett. *Visibility Graph Recognition*. PhD thesis, Dept. Comput. Sci., Univ. Toronto, Toronto, ON, 1990. Tech. Rep. 231/90.

[Gho88] S. K. Ghosh. On recognizing and characterizing visibility graphs of simple polygons. In *Proc. 1st Scand. Workshop Algorithm Theory*, volume 318 of *Lecture Notes in Computer Science*, pages 96–104. Springer-Verlag, 1988.

[Gho97] S. K. Ghosh. On recognizing and characterizing visibility graphs of simple polygons. *Discrete Comput. Geom.*, 17(2):143–162, 1997. Revision of [Gho88].

[GP84] J. E. Goodman and R. Pollack. Semispaces of configurations, cell complexes of arrangements. *J. Combin. Theory Ser. A*, 37:257–293, 1984.

[Knu92] D. E. Knuth. *Axioms and Hulls*, volume 606 of *Lecture Notes in Computer Science*. Springer-Verlag, Heidelberg, Germany, 1992.

[Mnë91] N. E. Mnëv. The universality theorem on the oriented matroid stratification of the space of real matrices. In J. E. Goodman, R. Pollack, and W. Steiger, editors, *Discrete and Computational Geometry: Papers from the DIMACS Special Year*, volume 6 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 237–243. AMS/ACM, 1991.

[O'R93] J. O'Rourke. Computational geometry column 18. *Internat. J. Comput. Geom. Appl.*, 3:107–113, 1993. Also in *SIGACT News* 24:1 (1993), 20–25.

[OS96] J. O'Rourke and I. Streinu. Pseudo-visibility graphs in pseudo-polygons: Part I. Technical Report 041, Dept. Comput. Sci., Smith College, Northampton, MA, January 1996. Revised Apr. 1996.

[OS97] J. O'Rourke and I. Streinu. The vertex-edge visibility graph of a polygon. *Comput. Geom. Theory Appl.*, 1997. To appear.

[Sho91] P. W. Shor. Stretchability of pseudolines is NP-hard. In P. Gritzmann and B. Sturmfels, editors, *Applied Geometry and Discrete Mathematics: The Victor Klee Festschrift*, volume 4 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 531–554. AMS Press, 1991.

[Str96a] I. Streinu. Clusters of stars. Technical Report, Dept. Comput. Sci., Smith College, Northampton, MA, 1996. See also these proceedings.

[Str96b] I. Streinu. Non-stretchable pseudo-visibility graphs. Technical Report, Dept. Comput. Sci., Smith College, Northampton, MA, 1996.