On Fully Characterizing Terrain Visibility Graphs

by

Noushin Saeedi Bidokhti

B.Sc., Sharif University of Technology, 2008

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF SCIENCE

in

The Faculty of Graduate Studies

(Computer Science)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

March 2012

Abstract

A *terrain* is an x-monotone polygonal line in the xy-plane. Two vertices of a terrain are mutually visible if and only if they are adjacent or the line segment that connects them lies above the terrain (except at its endpoints). A graph whose vertices represent terrain vertices and whose edges represent mutually visible pairs of terrain vertices is called a *terrain visibility graph*. Understanding the graph theoretic properties of all terrain visibility graphs may help us understand the combinatorial structure of these geometric objects and help to address problems related to geometric visibility. The properties that are true of all terrain visibility graphs are called necessary properties. The set of properties that, if true, imply that a graph is a terrain visibility graph is called a sufficient set. We would like to find properties that are both necessary and sufficient; that is, we would like to characterize, graph theoretically, terrain visibility graphs.

Abello et al. [Discrete and Computational Geometry,14(3):331–358,1995] studied the *core* induced subgraphs of the visibility graphs of *staircase polygons*, which are exactly the class of terrain visibility graphs. They showed that the visibility between vertices in such structures implies some ordering requirements on the slopes of the lines that connect pairs of vertices in any realization. They approached the problem of whether certain graph properties are sufficient by creating a slope order on the lines that connect all pairs of polygon vertices (in a possible realization) such that the slope order is consistent with the desired visibility graph. The main contribution of our work is to give a much simpler proof that these properties guarantee such a slope ordering. Our proof consequently gives a faster algorithm for constructing this slope order. Our approach attempts to clarify the implications of the graph theoretic properties on the ordering of the slopes, and may be interpreted as defining properties on an underlying oriented matroid.

Table of Contents

List of Figures

Acknowledgements

First and foremost I would like to thank my supervisor, Prof. William Evans. His inspiration and support gave me the opportunity and motivation to explore very interesting research problems. I learned a lot from our discussions on the problems. I thank him for his understanding, great advice, patience, generosity with his time, and his invaluable support during my studies. I have enjoyed working with him a lot and I am truly indebted to him for being a great mentor, both professionally and personally.

I would like to extend my gratitude to my second reader, Prof. David Kirkpatrick, for his helpful comments. I thank the faculty members and my colleagues in the Beta Lab who have made my M.Sc. studies at the University of British Columbia more enjoyable. I also thank Holly Kwan and Lara Hall for always being helpful.

I would also like to thank Prof. Janos Pach, Radoslav Fulek, and other people in the Discrete and Computational Geometry Group at EPFL (École Polytechnique Fédérale de Lausanne), where I did an internship last year. It has been a pleasure working with them. It helped me improve my research abilities, and made me more determined in my career. I thank Janos for his valuable advice.

Many people supported me during my stay in Vancouver and helped make life more fun. Thanks to Monir Hajiaghayi, Ali Mahmoudzadeh, Parnian Alimi, Niousha Bolanzadeh, Azin Dastpak, Amirhossein Mehrabian, Zohreh Jabbari, Pantea Jabbari, Babak Bostan, Ida Karimfazli, Behrooz Yousefzadeh, Nima Taherinejad, Mona Haraty, Syavash Nobarany, and Yassaman Sefidgar. I would also very much like to thank my aunt Pari and my cousins Taymaz and Tara for their true kindness and invaluable support during my time in Vancouver.

Finally, I owe my especial thanks to my dearest parents Sima and Mahmoud and my caring sister Shirin for their unconditional love and support. I have learned the most from them. Their inspiration, love, support, advice, and trust have made me who I am today, and I am deeply grateful to them for each day of my life, and for all their sacrifices.

For my parents and sister

Chapter 1

Introduction

Problems related to geometric visibility have arisen from applications in graphics and motion planning in robotics. In graphics, for example, the hidden line problem for computer-drawn polyhedra is to determine which edges, or parts of edges, of a polyhedra are visible from a given vantage point. In robotics or more specifically motion planning, we would like to find the shortest path between two positions for a robot without hitting the objects around. Visibility problems include the well-known art-gallery problems: planning a path for a guard so that it can protect the entire art gallery, or determining the minimum number of the guards which together can observe the whole art gallery.

It is conceivable that understanding the combinatorial structure of visibility is needed before addressing problems involving visibility in computational geometry. A *visibility graph* is a fundamental combinatorial structure which has proven useful for addressing such problems. The vertices of a visibility graph correspond to geometric components such as points or line segments. There is an edge between two vertices of the graph if the components are visible to each other. There are variations in the definition of visibility depending on the underlying application. For the class of art gallery problems, the vertex visibility graph of polygons is more commonly used. Here, the geometric components are the vertices of the polygon, and two vertices are *visible* if and only if the line segment connecting them is in the interior or along the boundary of the polygon. Throughout this thesis, by the term visibility graph, we mean the vertex visibility graph of the geometric object unless otherwise stated.

Visibility graphs have been used to find the shortest path between two vertices of a polygon that goes through the inside of the polygon. Such a path is a subgraph of the visibility graph [46]. Thus, we may construct the visibility graph of the polygon and simply find the shortest path between two vertices in the graph. Since computing the visibility graph takes linear time [41] in the size of the graph, this is a reasonably efficient method.

Ideally, we would like to fully understand the combinatorial properties of visibility in polygons. The properties that are true of all visibility graphs are called necessary properties. The set of properties that, if true, imply that a graph is a visibility graph is called a sufficient set. If a graph satisfies the sufficient graph theoretic properties, it is *realizable* as a visibility graph of a polygon. We would like to find properties that are both necessary and sufficient; that is, we would like to *characterize*, graph theoretically, visibility graphs. There is a close relation between characterization and *recognition* of a class of graphs. A recognition algorithm for a class of graphs is an algorithmic characterization of that class. Studying the recognition and characterization of visibility graphs may help us find more efficient algorithms for problems related to geometric visibility.

Although there are some partial results for restricted polygons, no characterization of visibility graphs of simple polygons is available in general. It may be useful to separate the characterization problem into two parts: First find combinatorial conditions on point sets representing the vertices of a visibility graph that obeys certain graph theoretic properties; then decide if a point set with these combinatorial conditions is realizable. Such combinatorial conditions may be interpreted in the language of oriented matroids. Several important and hard realizability problems of combinatorial geometry can be reduced to the realizability problem of oriented matroids. The realizability problem for oriented matroids in general is NP-hard. However, the graph theoretic properties of polygon visibility graphs make the family of the oriented matroid representations restricted. It may be possible that the realizability problem for such oriented matroids is easier to address.

1.1 Problem statement

A one-dimensional *terrain* is an x-monotone polygonal line in the xy-plane. The endpoints of the terrain line segments mark terrain vertices. Two vertices of a terrain are mutually visible if and only if they are adjacent in the terrain or the line segment that connects them lies above the terrain (except at its endpoints). The visibility graph of a terrain with n vertices is the undirected graph with vertex set $[n] \equiv \{1, 2, \ldots, n\}$ and edges $\{\{a, b\} \mid$ the ath and bth terrain vertices are mutually visible$\}$. The ath vertex of the terrain is the vertex with the ath smallest x-coordinate. (The vertex numbers in the terrain visibility graph give the ordering of vertices along the terrain.)

We want to use graph theoretic properties to fully characterize terrain visibility graphs; that is we would like to find properties that are both necessary and sufficient for a graph to be a terrain visibility graph. It is relatively simple to show that all terrain visibility graphs satisfy three properties described in Chapter 3. It is much harder to show that all graphs with these three properties are terrain visibility graphs. In fact, though Abello, et al. [4] claimed this result, a complete proof has never been published.

1.2 Description of results

We give a streamlined proof of a step towards characterizing terrain visibility graphs.

Abello et al. [Discrete and Computational Geometry,14(3):331–358,1995] studied the *core* induced subgraphs of the visibility graphs of *staircase polygons*, which are exactly the class of terrain visibility graphs. (See the definitions at the start of Section 2.2.) They showed that the visibility between vertices in such structures implies some ordering requirements on the slopes of the lines that connect pairs of vertices in any realization. They approached the problem of whether certain graph properties are sufficient by creating a slope order on the lines that connect all pairs of polygon vertices (in a possible realization) such that the slope order is consistent with the desired visibility graph. The main contribution of our work is to give a much simpler proof that these properties guarantee such a slope ordering. Our proof consequently gives a faster algorithm for constructing this slope order.

Our approach is to establish an orientation on every triple of terrain vertices such that they are consistent with the desired visibility graph. For three points $a < b < c$, the orientation of the triple (a, b, c) determines whether b lies above or below the line through a and c. Such a triple orientation implies an ordering on the slopes of the three lines connecting these points. We define our orientation such that the resulting slope ordering requirements are consistent with the desired visibility graph, and show that the requirements induced by the set of all these triple orientations are globally consistent and form a global slope order on all lines.

We clarify the implications of the graph theoretic properties on the ordering of the slopes. Moreover, we show that our orientation has specific properties that forbid certain substructures in the point set realizing it. This may help prove realizability. In Chapter 5, we interpret our orientation as defining properties on an underlying oriented matroid.

Chapter 2

Related work

There have been extensive studies to better understand visibility graphs. The visibility may be defined among the vertices of a polygon, line segments in the plane, or various other geometric objects in two or higher dimensions. O'Rourke [49], and Ghosh and Goswami [37] give an excellent review of research on visibility graphs. Ghosh's book [36] is also a good detailed source of information on visibility graphs.

Considerable literature exists on construction of visibility graphs. For a simple polygon, there are efficient algorithms to compute the visibility graph. Asano et al. [10] and Welzl [64] have $O(n^2)$ algorithms to find the visibility graph of an n-vertex polygon. Hershberger [41] presents an $O(m + n \log \log n)$ algorithm where m is the number of edges in the visibility graph. The term $O(n \log \log n)$ reflects the time to triangulate a polygon using the algorithm of Tarjan and Van Wyk [62]. Chazelle [20] gives a linear triangulation algorithm, thus Hershberger's result becomes $O(m + n)$ which is in fact $O(m)$.

Other visibility problems that have been studied involve the characterization and recognition problems. There is no polynomial algorithm known to recognize visibility graphs (that is, to decide whether a given graph is the visibility graph of some polygon). Nor is the problem known to be NP-hard, or even in NP. Everett [29] expresses the recognition problem for visibility graphs as a sentence in the existential theory of the reals, and shows that the problem can be solved in PSPACE.

Characterizing visibility graphs and finding algorithms to recognize them seem challenging. The class of visibility graphs does not lie in any of the well-known classes of graphs such as planar graphs, chordal, circle or perfect graphs [29, 35]. Everett [29] shows that there is no finite set of forbidden induced subgraphs which characterizes visibility graphs.

Results on visibility graphs so far have involved restricting the class of graphs, or restricting the class of polygons, or adding extra information to the graph.

2.1 Restricting the class of graphs

ElGindy [28] shows that any *maximal outerplanar graph* is a visibility graph. A graph is *outerplanar* if it has a planar representation with all the vertices on the outer face. A *maximal outerplanar graph* is an outerplanar graph such that the addition of any arc results in a graph that is not outerplanar. A triangulation of a simple polygon is a maximal outerplanar graph. Thus the visibility graph of any simple polygon on n vertices is a supergraph of an n vertex maximal outerplanar graph [29]. ElGindy provides a linear time embedding algorithm to construct a *uni-monotone* polygon such that its visibility graph is identical to any desired maximal outerplanar graph. A *monotone polygonal chain* is a set of vertices in the plane, ordered along some direction, with consecutive vertices joined. A polygon is *monotone* if it can be broken into two monotone polygonal chains (both ordered along the same direction). A *uni-monotone* polygon is a monotone polygon where one chain consists of a single edge. A terrain, whose visibility properties we study in this thesis, is essentially the same as a monotone polygonal chain (the only difference is that we consider terrains to be monotone in the horizontal direction).

Colley [24] extends the range of graphs that can be identified as visibility graphs. He defines a new class of graphs, *tree of cliques graphs*, and presents an algorithm to recognize these graphs. A *tree of cliques graph* is formed by replacing each face of a two-connected outerplanar graph (except for the outerface) with a clique on the same vertices. He proves that all tree of cliques graphs are visibility graphs by giving an embedding algorithm that constructs a uni-monotone polygon with the required visibility. His algorithm is an extension of ElGindy's algorithm. Both ElGindy's and Colley's algorithms can be modified to create a *uniform* uni-monotone polygon (that is, a uni-monotone polygon whose vertices are uniformly spaced in the direction of monotonicity). Not every uniform uni-monotone polygon has a tree of cliques visibility graph.

2.2 Restricting the class of polygons

A *staircase polygon* is a polygon consisting of alternate horizontal and vertical edges which is formed by a polygonal chain leading down and to the right plus the two additional edges required to complete a polygon. These polygons are also known as *orthogonal convex fans*. A *convex fan* is a polygon that has a convex vertex visible to all other vertices (that is, has a convex *kernel* vertex). An *orthogonal convex fan* is a convex fan consisting of only horizontal and vertical edges. The *core* of a staircase polygon is formed by the reflex vertices plus the vertices adjacent to the kernel vertex. A uniform step length staircase polygon is a staircase polygon whose vertices are evenly spaced with respect to the horizontal direction.

Abello and Egecioglu [2] give a polynomial time recognition algorithm, using linear programming, to recognize visibility graphs of uniform step length staircase polygons. Using this method, they show that there exist graphs that are visibility graphs of staircase polygons but realizable only if not all polygon vertices are uniformly spaced.

Colley [24, 25] proves a strong relationship between the visibility graphs of uni-monotone polygons and staircase polygons, which turns up in the recognition problems. Colley defines *skew transformations*, which move the points in the xy-plane vertically an amount proportional to their x-coordinate, and shows that skew transformations do not affect visibility (as they do not affect the collinearity of points). He uses the result to show that recognizing the visibility graph of a staircase polygon is equivalent, under linear-time reduction, to the problem of recognizing the visibility graph of a monotone polygonal chain, where the additional information of the order of the vertices on the chain is given. The vertex ordering of the core of a staircase polygon is determined by the visibility graph of the staircase polygon. A crucial part of his argument is that the core induced subgraph of the visibility graph of a staircase polygon is identical to the visibility graph of a monotone polygonal chain (or equivalently a terrain visibility graph), and vice versa. Using this and the result from Abello and Egecioglu [2], Colley shows that the visibility graphs of uniform uni-monotone polygons are a *strict* subset of the visibility graphs of general uni-monotone polygons, if the outside face of the polygon is fixed.

There have been few results that give a graph theoretic characterization of visibility graphs for restricted classes of polygons. Everett and Corneil [30] characterize the visibility graphs of 1-*spiral* polygons and show they are a subclass of *interval graphs*, and hence are *perfect*. A polygon is k-*spiral* if its boundary contains at most k concave chains. They give a linear time algorithm for recognizing the visibility graphs of 1-spiral polygons. Everett [29] shows that visibility graphs of 2-spiral polygons are perfect[1]. This is not true of k-spiral polygons for $k > 2$. She also shows that the visibility graph

[1]She shows that they are perfect if the strong perfect graph conjecture is true. The conjecture was proposed by Berge [13] in 1963 and some forty years later, it was proved by Chudnovsky et al. [23].

of a convex polygon with a single convex hole is a *circular arc graph*. A *circular arc graph* is the *intersection graph* of a collection of arcs on a circle.

Colley et al. [26] characterize the visibility graphs of *towers*. A *tower* is a polygon formed by two reflex chains sharing one common endpoint, plus one edge joining the other endpoints of the chains. They present a linear time algorithm to recognize visibility graphs of towers, and tie these visibility graphs to *bipartite permutation graphs*. There is an exponential lower bound and a doubly exponential upper bound on the grid size required to realize the visibility graphs of towers, by results from Lin and Skiena [45] and Colley et al. [26]. Visibility graphs of towers are also characterized by Choi et al. [22], where they use the term *funnel* instead of tower. They give a linear time algorithm to recognize the visibility graph of a funnel and to reconstruct the funnel from its visibility graph. They show that the visibility graphs of funnels are *weakly triangulated* and therefore *perfect*.

Abello and Kumar [6] give a characterization of the visibility graphs of the class of 2-spiral polygons when extra information is added to the graph. We discuss this result later in the next section.

2.3 Adding extra information

Many results on visibility graphs involve adding extra information to the graphs. Originally Ghosh [34] conjectured three necessary conditions for recognizing visibility graphs of simple polygons, when the Hamiltonian cycle forming the polygon boundary is specified. Everett [29] gave a counterexample to this conjecture and suggested a stronger version of the third necessary condition, which Srinivasaraghavan and Mukhopadhyay [56] proved is indeed necessary. Abello et al. [9] showed that, even with the stronger version of the third condition, the necessary conditions are insufficient. Ghosh [35] identified another necessary condition to circumvent the new counterexample. (Meanwhile, Abello and Kumar [7] proposed four other necessary conditions, but Ghosh [35] showed that all of them follow from his third and fourth conditions.) Ghosh conjectured that the four necessary conditions are sufficient. The first three conditions can be tested in polynomial time [29, 35, 37]. Although Vijay [63] claimed to have shown a polynomial time algorithm to test the fourth condition, whether this property can be tested in polynomial time is still open [37]. Later, Streinu [59] proved Ghosh's conjecture [35] is false.

Coullard and Lubiw [27] introduced further necessary conditions for a graph to be a visibility graph. They developed a new structural property of visibility graphs: Each 3-connected component of a visibility graph has a vertex ordering in which every vertex is adjacent to a previous 3-clique; that is, each 3-connected component of a visibility graph has a 3*-clique ordering*. The weaker result that each vertex is adjacent to a previous 2-clique is a consequence of polygon triangulation. The 3-clique ordering property is not sufficient. The property can be tested in polynomial time, and is used to give an algorithm for the *distance visibility graph problem*, which is the problem of whether an edge-weighted graph is the visibility graph of a simple polygon with the weights as Euclidean distances [27]. They proved each 3-connected component of a visibility graph is a visibility graph, and that every 3-connected visibility graph has a 3-clique ordering. They used this to reconstruct unique polygon representations for the 3-connected components (each vertex is adjacent to at least three old fixed vertices so the position of the new vertex is fixed), and showed how to reconstruct a polygon representation for the whole graph by pasting them together.

ElGindy [28] conjectured a characterization of visibility graphs of convex fans. He suggested a *decomposition strategy* and gave an algorithm to check whether a graph is the visibility graph of a convex fan, when the Hamiltonian cycle forming the boundary is known. However, the reconstruction appears tricky and the correctness of the algorithm is not clear.

Abello et al. [3, 4] claimed to have characterized visibility graphs of staircase polygons (or equivalently orthogonal convex fans) in a series of two papers, only one of which has appeared in the literature. The preliminary results by Abello et al.[8] are precursors to these two papers. Abello et al.[4] presented a necessary property for such visibility graphs, which was called the *persistent* property. They showed that the visibility between vertices in such structures implies some ordering requirements on the slopes of the lines that connect pairs of vertices in any realization. They approached the problem of whether the persistent property is sufficient, given the Hamiltonian cycle forming the boundary, by creating a slope order on the lines that connect all pairs of polygon vertices (in a possible realization) such that the slope order is consistent with the desired visibility graph. The definition of the persistent property in the initial papers [4, 8] was stated incorrectly and could not guarantee this result. Abello [1] corrected the definition in 2004. Abello et al. [4, 8] claimed that the proposed slope order is realizable by a point set but a complete proof has not been published. We describe their work in more detail in Section 4.2.

Following Ghosh's result [34], Abello and Kumar [7] study visibility graphs of simple polygons, when the Hamiltonian cycle forming the boundary of the polygon is given. They define a new class of graphs called *quasi-*

persistent graphs, which they show is equivalent to the class of graphs satisfying the first two necessary conditions of Ghosh [34]. This implies that visibility graphs of simple polygons are all quasi-persistent. They also show that quasi-persistent graphs are closed under *ear* deletion. An *ear* is a vertex whose successor and predecessor vertices (with regards to the Hamiltonian cycle) are connected. Using a geometric interpretation of a polygon realizing a quasi-persistent graph, they determine vertices that block the line of sight between pairs that are not mutually visible. This gives a *blocking vertex assignment*. (Different polygons with the same visibility graph may have different blocking vertex assignments.) They show that a blocking vertex assignment, if determined by a polygon realization, satisfies four necessary conditions. The last three conditions are based on properties of Euclidean shortest paths (as the shortest path between a pair that is not mutually visible is determined by the blocking vertex assignment of the pairs on the path). Ghosh [35] shows all four necessary conditions of Abello and Kumar follow from his third and fourth conditions. Abello and Kumar show that every quasi-persistent graph has a blocking vertex assignment that satisfies the last three conditions, and suggest that the first condition is what restrains quasi-persistent graphs to be visibility graphs. To address the realizability problem, they introduce an *oriented matroid* approach: first find the combinatorial properties on the point set corresponding to the vertices of the visibility graph (that is represented by an oriented matroid or equivalently a *simplicial chirotope*); then decide whether such a point set is realizable. In particular, they show how to construct a *uniform rank 3 oriented matroid* for every quasi-persistent graph satisfying the four conditions, which if affinely realizable yields a simple polygon with the desired visibility graph.

Abello and Kumar [6] show that these conditions are sufficient to reconstruct a polygon from the graph when restricted to 2-spiral polygons. The characterization of visibility graphs of 2-spiral polygons relies on the existence of an ear whose removal leaves a smaller graph in the same class such that *any* realization of the smaller graph can be extended to a realization of the initial graph. This is not true for k-spiral polygons where $k \geq 3$. Thus, in the general case, a reconstruction algorithm needs to take into consideration that only some of the realizations of an induced subgraph may be extensible to the realization of the initial graph. This motivates understanding the *realization space* of a given visibility graph. It has been shown that there are simple configurations of points with *disconnected realization space* [16, 40, 54].

O'Rourke and Streinu [51] introduce a new polygon visibility graph, the *vertex-edge visibility graph*, that represents visibility between vertices and edges. They suggest that the additional geometric information such graphs provide may simplify the problem of characterizing them. They show that a vertex-edge visibility graph determines the convexity of vertices, the vertex visibility graph, and, for each vertex, the *partial local sequence* (definition below) and the *shortest path tree*. Moreover, they show that it contains the same information as an *edge visibility graph* (which shows edge-to-edge visibility in a polygon), and that an external vertex-edge visibility graph determines the convex hull vertices.

The following definitions are useful to understand the following related works. The *local sequence* for a vertex v is the circular sequence of all other vertices as they are encountered by a rotating line through v. The *partial local sequence* contains only the visible vertices. The collection of all local sequences is called a *cluster of stars*, and forms an *affine* (or *acyclic*) *uniform rank 3 oriented matroid* [38, 57], whose topological representation is a *generalized configuration of points* [38]. Let $\mathcal{P}$ be a set of points in the Euclidean plane, and and let $\mathcal{L}$ be an *arrangement of pseudolines* such that every pair of points in $\mathcal{P}$ lie on exactly one pseudoline, and each pseudoline in $\mathcal{L}$ contains exactly two points of $\mathcal{P}$. Then the pair $(\mathcal{P}, \mathcal{L})$ is a generalized configuration of points in general position ("in general position" indicates that no three points of $\mathcal{P}$ lie on the same pseudoline of $\mathcal{L}$). A *pseudoline* is a simple curve that separates the plane. An *arrangement of pseudolines* is a collection of pseudolines, such that each pair of them meet in exactly one point, where they cross. The set of all local sequences of a generalized configuration of points determines the *chirotope* information, or equivalently the set of all triple orientations, of the point set. The orientation of a triple (i, j, k) shows whether k is to the right or to the left of the pseudoline through i to j. The same is true when pseudolines are straight-lines. The set of all triple orientations determines the *order type*.

O'Rourke and Streinu [50] generalize the notion of straight-line visibility to visibility along pseudolines. The idea of pseudo-visibility comes from the concept of duality between *pseudoline arrangements* and *generalized configurations of points*. They give a complete characterization of vertex-edge (pseudo) visibility graphs of *pseudo-polygons*. They show every vertex-edge pseudo-visibility graph satisfies three properties recognizable in polynomial time. From these properties, they derive a number of additional combinatorial concepts (such as convexity of vertices, partial local sequences and shortest path trees), and use them to define a predicate on any triple of vertices (the predicates match the chirotope definition of Abello and Kumar [7]). They show the predicates satisfy Knuth's *CC system* axioms [44]. CC systems are equivalent to uniform rank 3 acyclic oriented matroids [44], which in turn are equivalent to Goodman and Pollack's generalized configurations of points in general position [14]. They show that the order of the corresponding generalized configuration of points induces a pseudo-polygon (as the edges do not cross) with the desired vertex-edge pseudo-visibility graph, and conclude that their properties are also sufficient. As a consequence of the relationship between vertex-edge visibility graphs and vertex visibility graphs, they show that the recognition problem for vertex visibility graphs of pseudo-polygons is in NP (the same problem with straight-line visibility is only known to be in PSPACE).

Streinu [57] concludes the work of O'Rourke and Streinu [50] by providing efficient algorithms for recognizing and drawing clusters of stars, which consequently solves the pseudo-polygon drawing problem. She gives a characterization of clusters of stars that are realizable as generalized configurations of points. She gives an $O(n^2)$ time drawing algorithm that uses an $O(n^2)$ space data structure to keep the order type of the generalized configuration of points, from which the orientation of each triple can be retrieved in constant time. Her characterization conditions suggest that an orientation (clockwise or counterclockwise) of every triple that is consistent over a set of local sequences (that is, obeys a generalized transitivity law) is realizable as a generalized configuration of points, for some ordering of the point set. Knuth's CC systems [44] can also be interpreted as characterizing the local sequences of generalized configurations of points.

The concept of pseudo-visibility detaches the *stretchability* question from the combinatorial aspects of the problem. A pseudoline arrangement (or equivalently an affine rank 3 oriented matroid) is *stretchable* or *realizable* if it is isomorphic to a line arrangement. The main difficulty in fully characterizing vertex-edge visibility graphs of (straight-line) polygons is to decide whether a certain class of affine rank 3 oriented matroids is stretchable. It is well-known that stretchability of pseudoline arrangements, in general, is NP-hard [47, 55]. Mnëv [47] shows determining the stretchability of a pseudoline arrangement is equivalent to the existential theory of the reals (via which one can decide if a graph is a visibility graph [29]). Shor [55] proves the NP-hardness by reducing the *monotone* 3-*SAT* problem, a variant of the 3-SAT problem, to the stretchability problem. Using *Pappus* and *Desargues* configurations, he constructs a pseudoline arrangement which is stretchable if and only if the given monotone 3-SAT formula is satisfiable. However, there exist various techniques to prove stretchability for particular instances [15–19, 52, 53]. O'Rourke and Streinu [50] suggest that the class of acyclic uniform rank 3 oriented matroids generated by the vertex-edge pseudo-visibility graphs is a strict subclass of all acyclic uniform rank

3 oriented matroids, so it may be possible to characterize or recognize them (with an algorithm of complexity smaller than PSPACE).

Streinu [58] shows that the class of vertex-edge (straight-line) visibility graphs is properly contained in the class of vertex-edge pseudo visibility graphs. She introduces *star-like pseudo-polygons* and shows that a star-like vertex-edge pseudo-visibility graph may not be realizable. She gives a complete combinatorial characterization for the stretchability problem of this subclass. However, since vertex-edge pseudo-visibility graphs contain more geometric information than pseudo-visibility graphs, it is possible that a pseudo-visibility graph is associated with several (possibly exponentially many) vertex-edge pseudo-visibility graphs, some of which are stretchable and some not [59]. For instance, all star-like pseudo-visibility graphs are realizable [58]. Later, Streinu [59] shows an infinite family of pseudo-visibility graphs that are not realizable. She shows that any graph in this class induces a unique vertex-edge pseudo-visibility graph all of whose compatible pseudo-polygons are necessarily constructed on top of a generalized configuration of points containing a nonstretchable substructure. Since the graphs in this family satisfy the necessary conditions of O'Rourke and Streinu [50], Abello and Kumar [7], and Ghosh [35], this implies that these necessary conditions are not sufficient to characterize (straight-line) visibility graphs.

Other combinatorial objects describing the structure of simple polygons (that contain more information than a visibility graph) have been studied by Everett et al. [31]. They introduce *stabbing information* of a simple polygon, which basically stores certain information about the intersections that each line connecting two vertices makes with the other edges of the polygon. They define three variations of the stabbing information, namely *weak*, *strong*, and *labelled*. They show the strong stabbing information is sufficient to recover the convex hull, the internal and external visibility graphs and to determine the reflex vertices; but is not sufficient to establish the order type of the vertex set. The labelled stabbing information is equivalent to the order type.

Horizontal and vertical visibility information in orthogonal polygons (or equivalently, the *orthogonal edge visibility graph*) has been studied by O'Rourke [48], and Jackson and Wismath [42]. O'Rourke [48] shows that horizontal and vertical (internal) visibility information consists of two disjoint trees. A tree is *realizable* if there is an orthogonal polygon that has the tree as one of the two components of its visibility graph. Two trees can *mesh* if they are jointly realizable by the same polygon. O'Rourke shows that every tree is realizable but not every pair of them are jointly realizable. He characterizes realizable labelled trees and the meshable labellings of tree pairs. But the characterization in terms of the structure of (unlabelled) tree pairs remains unsolved. The simpler problem of *bar visibility graphs* has been characterized by Wismath [65], and independently by Tamassia and Tollis [61]. Bar visibility graphs differ from orthogonal polygon visibility trees: the bars do not have to connect into a polygon, and they have two "sides"; This considerably widens the class of *bar representable* graphs. Later, Jackson and Wismath [42] show how to reconstruct an orthogonal polygon from both the internal and external horizontal and vertical visibility information. They use the horizontal visibility information to determine the convexity of vertices. Using the convexity information and both horizontal and vertical visibility information, they compute two partial orders on the x and y dimensions and subsequently assign integer coordinates to the polygon vertices such that the polygon is consistent with the input.

Chapter 3

Necessary properties for terrain visibility graphs

We first define the graph properties that we use, and then present the set of conditions satisfied by all terrain visibility graphs.

3.1 Graph property definitions

We use the following notations for integer intervals:

$(i..j) = \{x \in \mathbf{N} \mid i < x < j\}$,

$(i..j] = \{x \in \mathbf{N} \mid i < x \leq j\}$,

$[i..j) = \{x \in \mathbf{N} \mid i \leq x < j\}$, and

$[i..j] = \{x \in \mathbf{N} \mid i \leq x \leq j\}$.

Definition 1. *A graph $G = ([n], E)$ has the* X-property *if for every four vertices $a < b < c < d$, if $\{a, c\} \in E$ and $\{b, d\} \in E$ then $\{a, d\} \in E$.*

This property is called *inversion complete* by Abello [1].

Definition 2. *A graph $G = ([n], E)$ has the* bar-property *if for every edge $\{a, c\} \in E$ where $a + 1 < c$, there exists a vertex $b \in (a..c)$ such that $\{a, b\} \in E$ and $\{b, c\} \in E$.*

This property is an ordered version of *chordality* [1, 35]: every ordered cycle of length at least four has a chord.

Definition 3. *A graph $G = ([n], E)$ that has both the X-property and the bar-property and contains the Hamiltonian path $1, 2, \ldots, n$ is called* persistent.

Abello et al. [4] initially defined persistent graphs in 1995 in a different and slightly incorrect manner so that the X-property was not guaranteed (incorrect in the sense that some properties that they assume from their definition are not true). Abello modified the definition in 2004 [1] to include this property (which he called inversion completeness). His subsequent definition gives the same class of graphs as our definition does.

3.2 Necessary conditions

Now, we present necessary conditions for terrain visibility graphs and prove that terrain visibility graphs are all persistent. Colley [24] shows a strong relationship between terrain visibility graphs and visibility graphs of staircase polygons. He proves that the core induced subgraph of the visibility graph of a staircase polygon is identical to a terrain visibility graph, and vice versa. (See pages 5–6 of Section 2.2 for more details.) Abello et al. [4] show that the core induced subgraph of the visibility graph of a staircase polygon (and as a result a terrain visibility graph), with respect to its ordering, is persistent (for their persistent definition). Here, for completeness of results, we re-prove the necessary conditions for terrain visibility graphs by a simpler and more geometric approach. Our proof relies on two lemmas. The first is called the *Order Claim* in the literature [12, 21, 43], and the second is called the *Midpoint Claim* by King [43].

Lemma 1. *Terrain visibility graphs have the X-property.*

Proof. Consider four vertices $a < b < c < d$ in a terrain visibility graph $G = ([n], E)$, such that $\{a, c\} \in E$ and $\{b, d\} \in E$. Since $\{a, c\} \in E$, we know the terrain does not touch the half-strip above the segment $\overline{ac}$. Denote the half-strip above the segment $\overline{uv}$ by $H^+(\overline{uv})$. Similarly, we know the terrain does not touch $H^+(\overline{bd})$. Hence the ray above b crosses the line segment $\overline{ac}$ and the ray above c crosses the line segment $\overline{bd}$, which means the two segments $\overline{ac}$ and $\overline{bd}$ intersect in-between. Thus the line segment $\overline{ad}$ lies in the region $H^+(\overline{ac}) \cup H^+(\overline{bd})$. Therefore, the terrain does not penentrate $H^+(\overline{ad})$, which implies $\{a, d\} \in E$. See Figure 3.1. □

Lemma 2. *Terrain visibility graphs have the bar-property.*

Proof. Consider a terrain visibility graph $G = ([n], E)$. For any edge $\{a, c\} \in E$ where $a + 1 < c$, the terrain induced on the vertices in $[a..c]$ together with the edge $\{a, c\}$ create a simple polygon. The fact that every simple polygon admits a triangulation implies the existence of a vertex $b \in (a..c)$ such that $\{a, b\} \in E$ and $\{b, c\} \in E$. See Figure 3.2. □

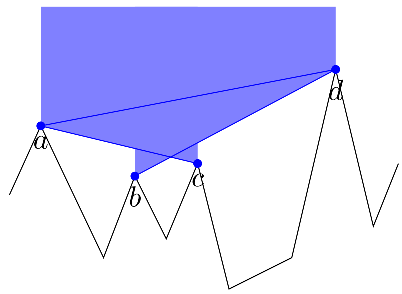

Figure 3.1: X-property in terrain visibility graphs.

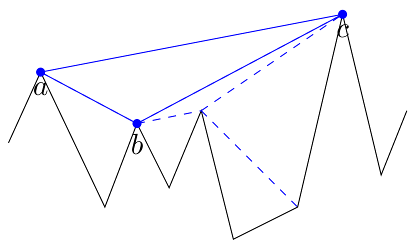

Figure 3.2: Bar-property in terrain visibility graphs.

We know terrain visibility graphs contain the Hamiltonian path $1, 2, \ldots, n$. Moreover, by Lemma 1 and Lemma 2, we have terrain visibility graphs satisfy both the X-property and the bar-property. Thus we conclude the following theorem (which is equivalent to Theorem 3.5 by Abello et al. [4]).

Theorem 1. *Terrain visibility graphs are persistent.*

Chapter 4

On the sufficiency of the persistent property

Determining whether a graph with certain properties is the visibility graph of of a simple polygon is known as visibility graph *realizability* or *recognition*. The actual drawing of the point set whose visibility graph is the desired graph is called visibility graph *reconstruction*. Even though there have been extensive studies on understanding the combinatorial properties of visibility graphs, the realizability and reconstruction problems appear difficult and the results are very limited (even for special classes of simple polygons).

Ideally, we would like to show that the persistent property is sufficient to imply that the graph is a terrain visibility graph, and to be able to recover a terrain from a given persistent graph; but this seems challenging. Abello et al. [4] attempted characterizing the core induced subgraphs of the visibility graphs of staircase polygons, which are exactly the class of terrain visibility graphs. They showed that the visibility information of a given vertex in such structures implies some ordering requirements on the slopes of the lines connecting this vertex to the others in any realization. They showed that there is a global order on all slopes that is consistent with all these local slope orders, for visibility graphs satisfying the persistent property. However, this does not imply that there is a point set that realizes the resulting slope order (that is, the properties are sufficient).

The main contribution of our work is to give a much simpler proof that the persistent property guarantees such a global slope ordering. Our proof consequently gives a faster algorithm for constructing this global slope order. Our approach attempts to clarify the implications of the graph theoretic properties on the ordering of the slopes, and may be interpreted as defining properties on an underlying oriented matroid.

Here, we first introduce the terminology that we use for the representation of the slope ordering. We next describe the overall approach of Abello et al. [4] briefly. Then, we give two different much simpler proofs to show that the persistent property guarantees a global slope order. The first proof is a proof by contradiction and is based on the implications of certain slope ordering requirements, which are imposed by the persistent property. The second proof uses a combinatorial result from Felsner and Weil [32].

4.1 Representation of the slope order

We use terminology similar to that used by Abello et al. [4] for the representation of the slope order of the lines connecting pairs of points.

A *tableau* of size n is a two-dimensional array of $n-1$ columns (indexed from 1 to $n-1$) where the ath column contains $n-a$ entries (indexed from $a+1$ to n), and whose entries are the integers $1, 2, \ldots, \binom{n}{2}$. For a tableau T, we refer to the entry in the ath column and bth row as $T[a, b]$. Note that $a < b$.

Consider a non-degenerate point set $[n]$. We may represent the slope ordering of the lines connecting all pairs of the points by a tableau T of size n, such that $T[i, j]$ is the *rank* of the slope of the line through i and j. (In other words, $T[i, j] = s$ if and only if the rank of the slope of the line through i and j is s.) We know that every three points $a < b < c$ form either a *positive orientation* (that is, b lies below the segment $\overline{ac}$) or a *negative orientation* (that is, b lies above the segment $\overline{ac}$). This implies that, in the tableau T representing the slope ordering, we have either $T[a, b] < T[a, c] < T[b, c]$ or $T[a, b] > T[a, c] > T[b, c]$. (This establishes Lemma 3.1 of Abello et al. [4].)

For three integers $a < b < c$, we say the *triple* $T[a, b]$, $T[a, c]$, and $T[b, c]$ is *oriented positively* if and only if $T[a, b] < T[a, c] < T[b, c]$. Similarly, the triple is *oriented negatively* if and only if $T[a, b] > T[a, c] > T[b, c]$. The triple is *balanced* if either

$$T[a, b] < T[a, c] < T[b, c], \text{ or}$$
$$T[a, b] > T[a, c] > T[b, c].$$

A *balanced tableau* is a tableau whose triples are all balanced.

The *skeleton* S_T of a tableau T is a two-dimensional array of the same dimensions as T where

$$S_T[a, c] = \begin{cases} 1 & \text{if } c = a + 1 \text{ or } T[a, c] > T[a, b] \text{ for all } b \in (a..c), \\ 0 & \text{otherwise.} \end{cases}$$

Figure 4.1 shows a balanced tableau and its skeleton.

An *n-triangle* is the strict lower triangle of an $n \times n$ zero-one matrix. A *persistent n-triangle* is the n-triangle of the adjacency matrix of a persistent graph.

1							
1	2						
2	14	3					
3	5	4	4				
6	8	7	21	5			
9	11	10	13	12	6		
15	16	17	19	18	20	7	
22	23	24	25	26	27	28	8

1							
1	2						
1	1	3					
1	0	1	4				
1	0	1	1	5			
1	0	1	0	1	6		
1	1	1	0	1	1	7	
1	1	1	1	1	1	1	8

Figure 4.1: A balanced tableau and its skeleton.

A skeleton S_T of a balanced tableau T of size n can be interpreted as the n-triangle of an adjacency matrix for an undirected graph with vertices $1, 2, \ldots, n$ and edges $\{\{a, c\} \mid S_T[a, c] = 1\}$. This graph is the *skeleton graph* of the balanced tableau. Note that the skeleton graph of a balanced tableau representing the slope order of a terrain point set is identical to the terrain visibility graph. (We know $S_T[a, c] = 1$ if and only if $c = a + 1$ or the slope rank of line $\overline{ac}$ is greater than the slope ranks of all lines $\overline{ab}$, where $b \in (a..c)$. This means that $S_T[a, c] = 1$ if and only if the terrain vertices a and c are adjacent in the terrain, or no terrain vertex in $(a..c)$ lies above the line $\overline{ac}$; that is, terrain vertices a and c are mutually visible. See Lemma 3.2 of Abello et al. [4].)

A tableau representing the slope order of a simple configuration of points is balanced (See Lemma 3.1 of Abello et al. [4]). This and Theorem 1 imply that the skeleton graph of a balanced tableau representing a terrain is persistent.

To characterize terrain visibility graphs, we would like to show that persistent graphs are visibility graphs of terrains. That is equivalent to showing that every persistent graph is the skeleton of a balanced tableau representing the slope order of a terrain point set. We have not proved that but here we show that every persistent graph is the skeleton graph of a balanced tableau representing a slope order. The slope order may or may not be realizable (as a terrain). In other words, we prove the following theorem. (The converse of that is true by Lemma 3.4 of Abello et al. [4].)

Theorem 2. *If M is a persistent n-triangle, then there exists a balanced tableau whose skeleton is identical to M.*

We will prove this theorem in Section 4.3.

4.2 Abello et al.'s approach

Abello et al. [4] prove that a graph is persistent if and only if it is the skeleton graph of a balanced tableau. They argue that a skeleton graph of a balanced tableau is persistent directly by using the definitions. Here we only describe their approach for the reverse direction (that is, every persistent graph is the skeleton graph of a balanced tableau).

For a persistent graph $G = (V, E)$, they define an edge e to be *reversible* if and only if the graph $G = (V, E \setminus e)$ remains persistent. They use the idea of reversible edges to partially order persistent graphs so that they can generate any of them in a canonical manner. Given a persistent graph of order n, they give an algorithm that starts from a clique of the same order and successively removes a reversible edge until the desired persistent graph is generated. They use the basis of this algorithm to reconstruct a balanced tableau from a given persistent graph. Namely, they start from a balanced tableau whose skeleton represents a clique, and perform operations on the tableau entries so that the underlying skeleton becomes incrementally closer to the desired graph. They define *flush* and *augmentation* operations, which are compositions of Coxeter type I and type II transformations, on a balanced tableau. Each flipping of a one entry in the skeleton (that is, removing a reversible edge) is done through a sequence of flush and augmentation operations on the tableau and is complicated. They establish a loop invariant to show the correctness of their proposed algorithm. However, the loop invariant is not intuitive and requires a complex proof. The overall complexity of their algorithm is $O(n^5)$.

4.3 Our main result

We reprove that every persistent graph is the skeleton graph of a balanced tableau but in a much simpler way. By the definition of a tableau's skeleton, we know that the skeleton entries (whether zero or one) imply certain inequality relations amongst the tableau entires. The following lemma summarizes these relations.

Lemma 3. *For a tableau T and an n-triangle M, $S_T = M$ if and only if*

1. *If $M[a, c] = 1$, then $c = a + 1$ or $T[a, c] > T[a, b]$ for all $b \in (a..c)$, and*
2. *If $M[a, c] = 0$, then there exists $b \in (a..c)$ such that $T[a, b] > T[a, c]$.*

Proof. It follows directly from the definition of the skeleton of a tableau. □

Using the combinatorial properties of a persistent graph G, we would like to derive a balanced tableau satisfying the conditions in Lemma 3, where M is the n-triangle of the adjacency matrix of G. We show that the inequality relations required by Lemma 3 and by the balanced property give a partial order on the tableau entries; and as a result any tableau that realizes this partial order is balanced, with a skeleton graph identical to the given persistent graph.

Our approach is to orient *all* tableau triples so that they are consistent with Lemma 3. Since our orientation is guaranteed to be balanced, to show the existence of such a balanced tableau, we only need to show that our orientation forms a partial order on the tableau entries.

4.3.1 Orienting the triples

We infer combinatorial properties on the structure of a persistent n-triangle, and use these structural properties in determining the orientation of a tableau triple.

Definition 4. *Let M be an n-triangle. The collection of the entries $M[a;[b..c]] \cup M[[a..b];c]$, where $a < b < c$, is called the abc*-hook*, and is denoted by $hook_M(abc)$. The* corner *of $hook_M(abc)$ is $M[a,c]$. The* column-arm *of $hook_M(abc)$ is $M[a;[b..c]]$, and the* row-arm *of $hook_M(abc)$ is $M[[a..b];c]$.*

The half-strict *abc*-rectangle *is the submatrix $M[[a..b);(b..c]]$, and is denoted by $rect_M(abc)$. (Figure 4.2 illustrates the abc-hook and the half-strict abc-rectangle.)*

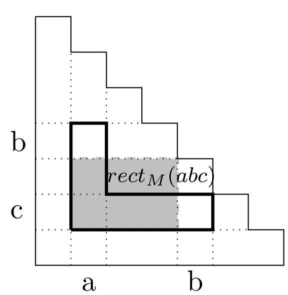

Figure 4.2: The *abc*-hook (bold outline) and the half-strict *abc*-rectangle (shaded).

Note that for every $a < b < c < d$ we have $rect_M(abc)$ is contained in $rect_M(abd)$ and $rect_M(bcd)$ is contained in $rect_M(acd)$. We refer to this property as *rectangle containment*.

The following lemma identifies the structure of the abc-hook in a persistent n-triangle M.

Lemma 4. *If M is a persistent n-triangle then each $hook_M(abc)$ is of one of the following forms (See Figure 4.3):*

- *The corner is one, or*
- *Either the row-arm or the column-arm consists of all zeros, or*
- *$M[a,b] = M[b,c] = 1$ and all other entries in $hook_M(abc)$ are zeros.*

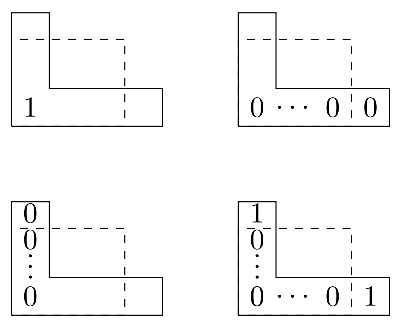

Figure 4.3: Possible structures of an abc-hook in a persistent n-triangle.

Proof. Suppose to the contrary that there exists an abc-hook that does not have any of the above properties. Thus, we have

1. $M[a,c] = 0$;
2. there exists $i \in (a..b]$ such that $M[i,c] = 1$; and
3. there exists $j \in [b..c)$ such that $M[a,j] = 1$;

such that either $i \neq b$ or $j \neq b$. But this contradicts M having the X-property on the four points $a < i < j < c$, which is impossible.

□

It is easy to deduce the possible structures of a half-strict abc-rectangle in a persistent n-triangle, from Lemma 4 and the X-property. Figure 4.4 illustrates these possible structures with regards to the associated abc-hook. As shown in the picture, in the four cases above each half-strict abc-rectangle contains a one entry whereas in the four cases below each half-strict abc-rectangle consists of all zero entries.

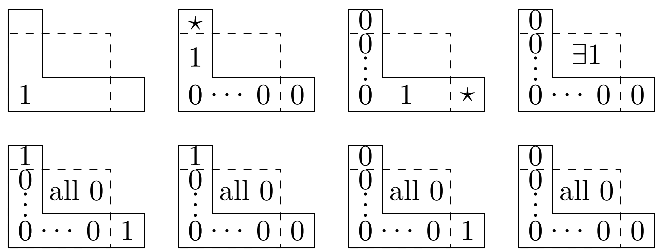

Figure 4.4: Possible structures of a half-strict abc-rectangle and its abc-hook in a persistent n-triangle.

From the bar-property we conclude the following lemma.

Lemma 5. *Let M be a persistent n-triangle. For every $a < b < c < d$ in $[n]$, if both $rect_M(abc)$ and $rect_M(bcd)$ consist of all zero entries then $M[[a..b);(c..d]]$ is a zero matrix.*

Proof. The *k-upper triangle* of a matrix is the collection of the entries above the *kth diagonal*. Suppose to the contrary that $M[[a..b);(c..d]]$ contains a one entry. Let $M[a',c']$ be a one entry lying at the kth diagonal in the submatrix $M[[a..b);(c..d]]$ such that the k-upper triangle of the submatrix contains only zero entries. It is easy to see that such an entry exists. We have $M[a',c'] = 1$ with $a' < b < c < c'$. Since M is persistent, by the bar-property we know there exists $b' \in (a'..c')$ such that $M[a',b'] = 1$ and $M[b',c'] = 1$. Since the k-upper triangle in the submatrix $M[[a..b);(c..d]]$ consists of zero entries, we infer $b \le b' \le c$. From knowing that $rect_M(bcd)$ consists of all zero entries, we conclude $b' = c$. So $M[a',b'] = 1$ implies that there is a one entry in $M[[a..b);c]$. Hence $rect_M(abc)$ contains a one entry, which is a contradiction. See Figure 4.5. □

Let M be a persistent n-triangle. We would like to find a balanced tableau whose skeleton is identical to M (if possible). Given M, we define the *orientation function* α_M from the 3-element subsets of $[n]$ to $\{+,-\}$ as follows:

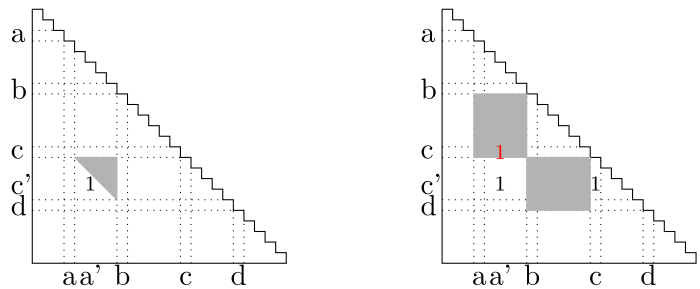

(a) We have $M[a', c'] = 1$. The bar-property implies that there exists $b' \in (a'..c')$ such that $M[a', b'] = 1$ and $M[b', c'] = 1$. Since the grey region consists of zero entries, we conclude $b \leq b' \leq c$.

(b) We have $M[b', c'] = 1$. Since $rect_M(bcd)$ consists of zero entries, we conclude $b' = c$. But this is a contradiction.

Figure 4.5: The bar-property implies $M[[a..b);(c..d]]$ is a zero matrix if both $rect_M(abc)$ and $rect_M(bcd)$ consist of all zero entries, for a persistent M.

For every $a < b < c$ in $[n]$, let

$$\alpha_M(abc) = \begin{cases} + & \text{if } rect_M(abc) \text{ contains a one entry,} \\ - & \text{if } rect_M(abc) \text{ consists of all zero entries.} \end{cases}$$

A tableau T *agrees* with orientation function α_M if for all $a < b < c$,

$$\begin{aligned} T[a,b] < T[a,c] < T[b,c] & \quad \text{if } \alpha_M(abc) = +, \textit{ and} \\ T[a,b] > T[a,c] > T[b,c] & \quad \text{if } \alpha_M(abc) = -. \end{aligned}$$

Figure 4.6 illustrates the orientation function α_M and the inequality relations required in a tableau that agrees with α_M.

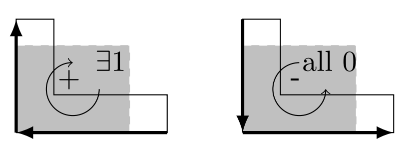

Figure 4.6: Orientation function α_M and the implied inequality relations.

Suppose there exists a tableau T that agrees with α_M. The following lemma shows that our orientation captures both the balanced property and the properties required to satisfy $S_T = M$.

Lemma 6. *Let M be a persistent n-triangle. If there is a tableau T whose entries agree with orientation function α_M, then*

1. *T is balanced, and*
2. *$S_T = M$.*

Proof. We have each tableau triple oriented either positively or negatively, and hence all triples are balanced. Therefore T is a balanced tableau. As illustrated in Figure 4.7, it is easy to observe that having each tableau triple oriented by α_M suffices to have:

1. if $M[a, c] = 1$, then $c = a + 1$ or $T[a, c] > T[a, b]$ for all $b \in (a..c)$, and
2. if $M[a, c] = 0$, then there exists $b \in (a..c)$ such that $T[a, b] > T[a, c]$. (A particular b is $b \in (a..c)$ such that $M[a, b] = 1$ and $M[a, x] = 0$ for all $x \in (b..c)$. Notice that such b always exists because $M[i, i + 1] = 1$ for all $i \in [1..n)$, by the persistent property.)

Thus the triple orientations induce the conditions in Lemma 3, which implies $S_T = M$. □

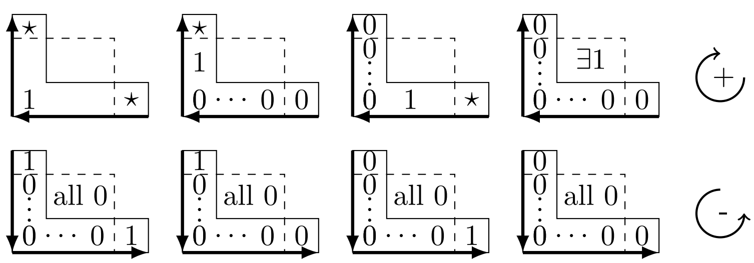

Figure 4.7: The triple orientations agree with the conditions in Lemma 3. Possible triple orientations and their implied inequality relations over a tableau that agrees with the orientation are illustrated with regards to the associated half-strict rectangles.

If a tableau T agrees with α_M then the entries of T have to obey a set of inequalities (as shown in Figure 4.7). From the above lemma, we conclude that if the set of all these inequalities (which are derived from α_M) gives a partial order on the tableau entries, then any tableau realizing this partial order would be a balanced tableau with skeleton M. In the following, we prove that our orientation gives a partial order.

4.3.2 The M-tableau relation digraph D_M

We introduce a directed graph to represent the required inequality relations between the tableau entries so that the tableau agrees with α_M. We show that the resulting digraph is acyclic if M is persistent, which concludes the proof of Theorem 2.

The notation $\binom{[n]}{k}$ represents the k-element subsets of $[n]$.

Definition 5. *Let M be a persistent n-triangle. The* M-tableau relation digraph *is a directed graph $D_M = (V_M, E_M)$ such that*

- $V_M = \{v_{ab} \mid a < b$ *and* $\{a, b\} \in \binom{[n]}{2}\}$ *(that is, the vertices are the 2-element subsets of $[n]$—the vertex v_{ab} represents the tableau entry in the ath column and the bth row), and*
- $E_M = \{(v_{ab}, v_{ac})$ *and* $(v_{ac}, v_{bc}) \mid \alpha_M(abc) = -\} \cup \{(v_{bc}, v_{ac})$ *and* $(v_{ac}, v_{ab}) \mid \alpha_M(abc) = +\}$ *(that is, the edges represent the inequality relations between the tableau entries that are imposed by the orientation function α_M).*

In the remaining, we give two different proofs to show that D_M is acyclic. Both proofs result from the use of rectangle containment and Lemma 5. The first one is a direct proof by contradiction, in which we show how certain sets of edges imply the existence of other edges in the digraph (which subsequently gives a contradiction in the case of a cycle). The second one relates our work to a combinatorial result from Felsner and Weil [32].

4.3.3 The digraph D_M is acyclic — proof 1

We embed the digraph D_M on the n-triangle M such that each vertex v_{xy} is placed at the position of the entry $M[x, y]$ and each edge is a straight line segment. The bth *row* of the graph is the subgraph induced on the vertices $\{v_{xb} \mid x \in [1..b)\}$ (that is, the vertices at row b). The ath *column* of the graph is the subgraph induced on the vertices $\{v_{ay} \mid y \in (a..n]\}$ (that is, the vertices at column a).

The following lemma shows that each row or column of the graph is acyclic.

Lemma 7. *Let D_M be an M-tableau relation digraph where M is a persistent n-triangle. Each row or column of D_M is acyclic.*

Proof. For an edge e that has both its endpoints in the same row or column, we define $e\text{-}rect_M(e)$ as follows: For every $a < b < c$, let $e\text{-}rect_M((v_{ab}, v_{ac})) =$

$e\text{-}rect_M((v_{ac}, v_{ab})) = e\text{-}rect_M((v_{ac}, v_{bc})) = e\text{-}rect_M((v_{bc}, v_{ac})) = rect_M(abc)$. We say $e\text{-}rect_M(e)$ is the half-strict rectangle associated with the edge e.

Suppose there is a cycle that has all its vertices in the same row. Let e_{out} be the out-going cycle edge incident on the rightmost point of the cycle. From the direction of e_{out} we know $e\text{-}rect_M(e_{out})$ contains a one entry. It is easy to see that $e\text{-}rect_M(e_{out})$ is covered by the union of all half-strict rectangles associated with cycle edges in the opposite direction relative to e_{out}. These half-strict rectangles contain only zero entries. This contradicts $e\text{-}rect_M(e_{out})$ having a one entry. See Figure 4.8.

A similar contradiction arises if a cycle occurs in a column (we consider the in-going cycle edge incident on the topmost point in this case). □

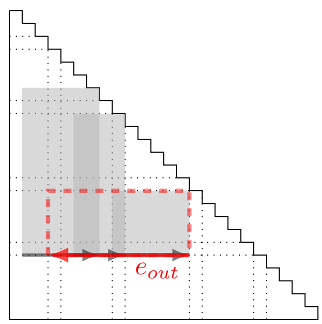

Figure 4.8: No cycle lies completely in the same row. Otherwise $e\text{-}rect_M(e_{out})$ is contained in the union of the half-strict rectangles associated with edges in the opposite direction relative to e_{out}, which is impossible.

Corollary 1 (Transitivity on a row or a column)**.** *Let $D_M = (V_M, E_M)$ be an M-tableau relation digraph where M is a persistent n-triangle. We have:*

- *if (v_{a_1b}, v_{a_2b}) and (v_{a_2b}, v_{a_3b}) are in E_M, then (v_{a_1b}, v_{a_3b}) is in E_M, and*
- *if (v_{ab_1}, v_{ab_2}) and (v_{ab_2}, v_{ab_3}) are in E_M, then (v_{ab_1}, v_{ab_3}) is in E_M.*

Proof. We know that for every $a < b < c$ in $[n]$, we have a directed edge between the pairs $\{v_{ac},\ v_{bc}\}$ and $\{v_{ab},\ v_{ac}\}$ in E_M. Thus the corollary is a direct consequence of Lemma 7. □

The following two lemmas provide the clues to prove that D_M is acyclic.

Lemma 8. *Let $D_M = (V_M, E_M)$ be the M-tableau relation digraph defined on a persistent n-triangle M. For every $a < b < c < d$ in $[n]$ the digraph D_M satisfies the four properties below (see Figure 4.9):*

Property 1: *if $(v_{bd}, v_{bc}) \in E_M$, then $(v_{ad}, v_{ac}) \in E_M$.*

Property 2: *if $(v_{ac}, v_{ad}) \in E_M$, then $(v_{bc}, v_{bd}) \in E_M$.*

Property 3: *if $(v_{ad}, v_{ac}) \in E_M$ and $(v_{ac}, v_{bc}) \in E_M$, then $(v_{bd}, v_{bc}) \in E_M$.*

Property 4: *if $(v_{ac}, v_{bc}) \in E_M$ and $(v_{bc}, v_{bd}) \in E_M$, then $(v_{ac}, v_{ad}) \in E_M$.*

Proof. We prove each property as follows:

Property 1. If $(v_{bd}, v_{bc}) \in E_M$ then $\alpha_M(bcd) = +$ (by Definition 5). Thus $rect_M(bcd)$ contains a one entry. Since $rect_M(bcd)$ is contained in $rect_M(acd)$, we know $rect_M(acd)$ contains a one entry. Therefore $\alpha_M(acd) = +$, which implies $(v_{ad}, v_{ac}) \in E_M$.

Property 2. If $(v_{ac}, v_{ad}) \in E_M$ then $\alpha_M(acd) = -$. Thus $rect_M(acd)$ consists of zero entries. Since $rect_M(bcd)$ is contained in $rect_M(acd)$, we know $rect_M(bcd)$ consists of all zero entries. Hence $\alpha_M(bcd) = -$, which implies $(v_{bc}, v_{bd}) \in E_M$.

Property 3. If $(v_{ad}, v_{ac}) \in E_M$ and $(v_{ac}, v_{bc}) \in E_M$, then $\alpha_M(acd) = +$ and $\alpha_M(abc) = -$. Suppose to the contrary that (v_{bd}, v_{bc}) is not in E_M. This implies that $\alpha_M(bcd) = -$. From $\alpha_M(bcd) = -$ and $\alpha_M(acd) = +$, we conclude $M[[a..b);(c..d]]$ contains a one entry. Thus we have both $rect_M(abc)$ and $rect_M(bcd)$ consist of all zero entries whereas $M[[a..b);(c..d]]$ contains a one entry, which is a contradiction to Lemma 5.

Property 4. If $(v_{ac}, v_{bc}) \in E_M$ and $(v_{bc}, v_{bd}) \in E_M$, then $\alpha_M(abc) = -$ and $\alpha_M(bcd) = -$. Thus both $rect_M(abc)$ and $rect_M(bcd)$ consist of all zero entries. From Lemma 5, we conclude $M[[a..b);(c..d]]$ is a zero matrix. This implies $rect_M(acd)$ consists of zero entries. Therefore $\alpha_M(acd) = -$ and we have (v_{ac}, v_{ad}) in E_M.

□

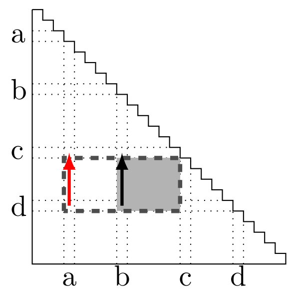

(a) Since $rect_M(bcd)$ is contained in $rect_M(acd)$, we deduce that if (v_{bd}, v_{bc}) is in E_M then (v_{ad}, v_{ac}) (in red) is in E_M.

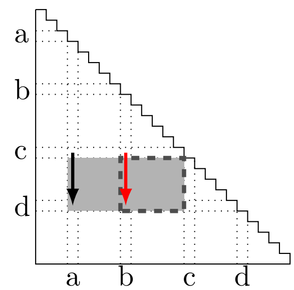

(b) Since $rect_M(bcd)$ is contained in $rect_M(acd)$, we deduce that if (v_{ac}, v_{ad}) is in E_M then (v_{bc}, v_{bd}) (in red) is in E_M.

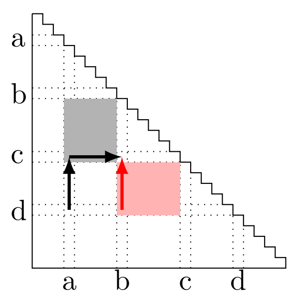

(c) The red region contains a one entry; otherwise it contradicts Lemma 5. This implies (v_{bd}, v_{bc}) (in red) is in E_M.

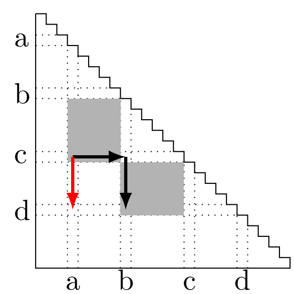

(d) The grey regions consist of all zero entries. Thus Lemma 5 implies (v_{ac}, v_{ad}) (in red) is in E_M.

Figure 4.9: The properties of D_M. The black edges imply the existence of the red edges.

Lemma 9. *Let $D_M = (V_M, E_M)$ be an M-tableau relation digraph where M is a persistent n-triangle. Let P be a path $v_{a_0b_0}, v_{a_1b_1}, \ldots, v_{a_kb_k}$ in D_M such that the vertex-induced subgraph of D_M on the vertex set $\{v_{a_0b_0}, \ldots, v_{a_kb_k}\}$ is acyclic. Assume that b_m is the greatest element in $\{b_0, b_1, \ldots, b_k\}$. If $b_k < b_0$ then $(v_{a_kb_m}, v_{a_kb_k})$ is in E_M. If $b_0 < b_k$ then $(v_{a_0b_0}, v_{a_0b_m})$ is in E_M. (See Figure 4.10.)*

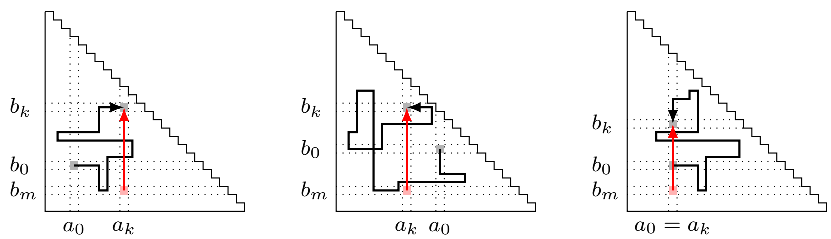

(a) If $b_k < b_0$, then $(v_{a_kb_m}, v_{a_kb_k})$ (in red) is in E_M.

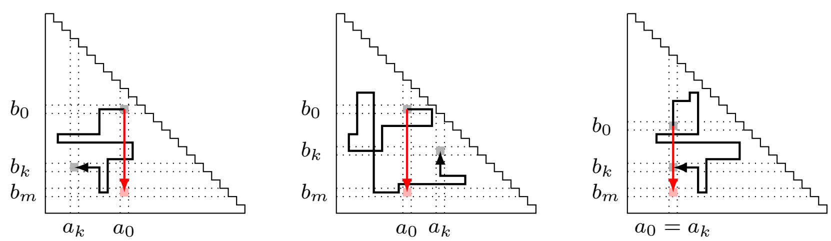

(b) If $b_0 < b_k$, then $(v_{a_0b_0}, v_{a_0b_m})$ (in red) is in E_M.

Figure 4.10: The black path $v_{a_0b_0}, v_{a_1b_1}, \ldots, v_{a_kb_k}$ implies the existence of the red edge.

Proof. First observe that, for every i such that $0 \leq i < k$, either $a_i = a_{i+1}$ or $b_i = b_{i+1}$. We proceed by induction on the number of edges in P. For succinctness, we occasionally use $\mathrm{v_i}$ for $v_{a_ib_i}$. Clearly the statement holds when the length of P equals one. We show that if the statement holds for all paths of length less than k, then the statement holds for all paths of length k.

Suppose that b_m is greater than both b_0 and b_k. If $b_k < b_0$, let P_1 be the path from $\mathrm{v_m}$ to $\mathrm{v_k}$. If $b_0 < b_k$, let P_1 be the path from $\mathrm{v_0}$ to $\mathrm{v_m}$. We know P_1 is a path of length less than k whose vertex-induced subgraph in D_M is acyclic. Thus, by the induction hypothesis, we have P_1 implies the existence of the desired edge in E_M. See Figure 4.11.

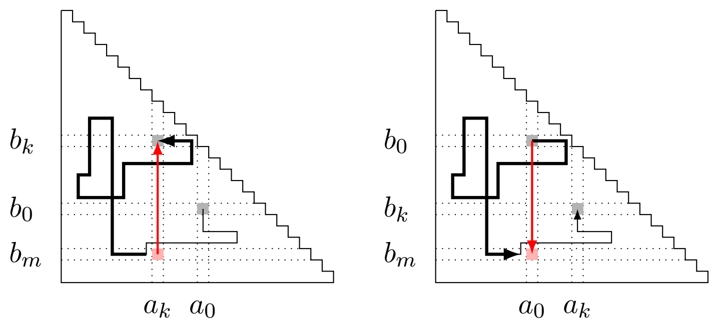

Figure 4.11: The case when b_m is greater than both b_0 and b_k. The black solid path is P. The black thick path is P_1. From the induction hypothesis, we conclude that P_1 implies the existence of the red edge.

Now suppose that $b_m = b_0$. Let P_1 be the path from v_0 to $\mathrm{v}_{\mathrm{k}-1}$. Note that $b_{k-1} \leq b_0$. If $b_{k-1} = b_0$, then the edge $(\mathrm{v}_{\mathrm{k}-1}, \mathrm{v}_{\mathrm{k}})$ is already the desired edge. This is because if $b_{k-1} = b_0$ then we have $b_{k-1} > b_k$ (by the conditions of the lemma, b_k is not b_0), and hence $a_{k-1} = a_k$. Thus $(\mathrm{v}_{\mathrm{k}-1}, \mathrm{v}_{\mathrm{k}}) = (v_{a_k b_m}, v_{a_k b_k})$. So we assume $b_{k-1} < b_0$. By the induction hypothesis on P_1, we infer that $(v_{a_{k-1} b_0}, v_{a_{k-1} b_{k-1}})$ is in E_M. By Corollary 1 and Lemma 8 (properties 1 and 3), it is easy to observe that this and $(\mathrm{v}_{\mathrm{k}-1}, \mathrm{v}_{\mathrm{k}}) \in E_M$ imply that $(v_{a_k b_0}, v_{a_k b_k})$ is in E_M. See Figure 4.12(a).

Finally assume that $b_m = b_k$. The argument is similar to the previous case. Let P_1 be the path from v_1 to v_{k}. Note that $b_1 \leq b_k$. If $b_1 = b_k$ then the edge $(\mathrm{v}_0, \mathrm{v}_1)$ is already the desired edge. This is because $b_1 = b_k$ implies $b_1 > b_0$ (because b_k is not b_0 by the conditions of the lemma), and hence $a_1 = a_0$. Thus $(\mathrm{v}_0, \mathrm{v}_1) = (v_{a_0 b_0}, v_{a_0 b_k})$. So we assume $b_1 < b_k$. By the induction hypothesis on P_1, we infer that $(v_{a_1 b_1}, v_{a_1 b_k})$ is in E_M. By Corollary 1 and Lemma 8 (properties 2 and 4), it is easy to observe that this and $(\mathrm{v}_0, \mathrm{v}_1) \in E_M$ imply that $(v_{a_0 b_0}, v_{a_0 b_k})$ is in E_M. See Figure 4.12(b). This concludes the proof.

□

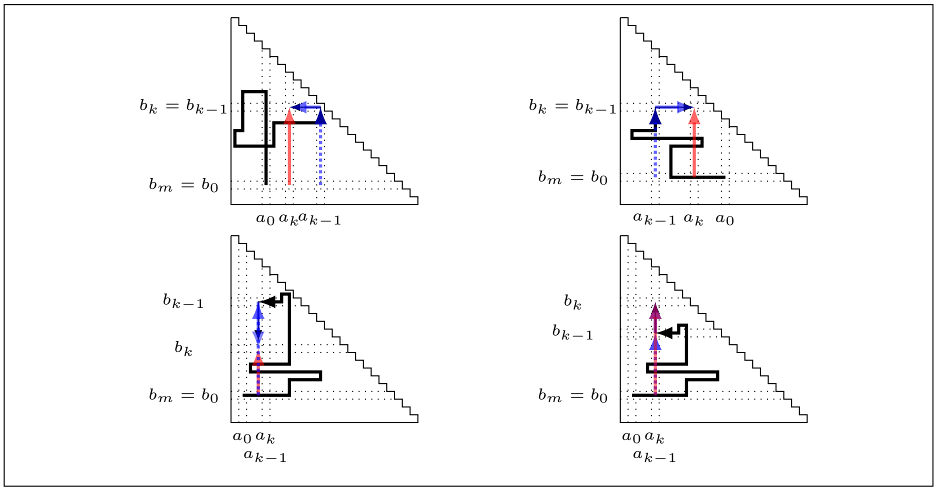

(a) The case when $b_m = b_0$ and $b_{k-1} \neq b_0$. We have $a_k < a_{k-1},\ b_k = b_{k-1}$ or $a_k > a_{k-1},\ b_k = b_{k-1}$ or $a_k = a_{k-1},\ b_k > b_{k-1}$ or $a_k = a_{k-1},\ b_k < b_{k-1} < b_0$. In the first and second subcases, we use the properties 1 and 3 in Lemma 8, respectively. In the third and fourth subcases we use Corollary 1.

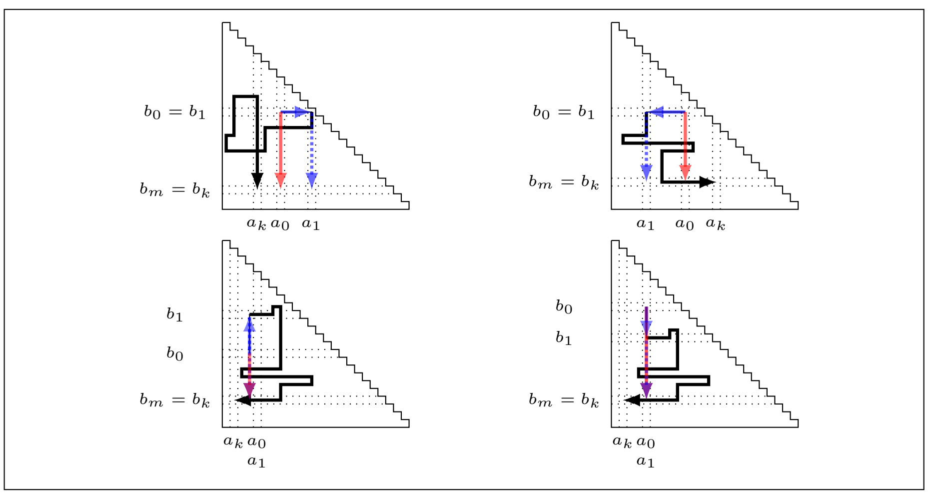

(b) The case when $b_m = b_k$ and $b_1 \neq b_k$. We have $a_0 < a_1,\ b_0 = b_1$ or $a_0 > a_1,\ b_0 = b_1$ or $a_0 = a_1,\ b_0 > b_1$ or $a_0 = a_1,\ b_0 < b_1 < b_k$. In the first and second subcases, we use the properties 4 and 2 in Lemma 8, respectively. In the third and fourth subcases we use Corollary 1. Note that the third subcase is impossible since we know the solid blue edge contradicts the existence of the dotted blue edge.

Figure 4.12: The cases $b_m = b_0$ and $b_m = b_k$. The black thick path is P_1. The path P is composed of P_1 and the blue solid edge. From P_1, by the induction hypothesis, we conclude the existence of the dotted blue edge. The blue solid and dotted edges imply the existence of the red edge.

Lemma 10. *Let $D_M = (V_M, E_M)$ be an M-tableau relation digraph where M is a persistent n-triangle. The digraph D_M is acyclic.*

Proof. Suppose to the contrary that D_M is cyclic. Again, for succinctness, we occasionally use $\mathrm{v_i}$ for $v_{a_i b_i}$. Let $C = v_{a_0 b_0},\ v_{a_1 b_1},\ \ldots,\ v_{a_{k-1} b_{k-1}},\ v_{a_0 b_0}$ be a shortest cycle in D_M. Observe that, for every i such that $0 \leq i \leq k-1$, either $a_i = a_{i+1 \bmod k}$ or $b_i = b_{i+1 \bmod k}$.

From Corollary 1, we know that C is composed of edges that alternate horizontal and vertical alignment. Assume that $(\mathrm{v_0}, \mathrm{v_1})$ is a topmost edge in C (we have $b_0 = b_1$ and no element in $\{b_0, b_1, \ldots, b_{k-1}\}$ is smaller that b_0). Let $(\mathrm{v_m}, \mathrm{v_{m+1}})$ be a bottommost edge in C (so we have $b_m = b_{m+1}$ and no element in $\{b_0, b_1, \ldots, b_{k-1}\}$ is greater that b_m). By Lemma 7, we know that $b_m > b_0$. Now, let $P_1 = \mathrm{v_1},\ \mathrm{v_2},\ \ldots,\ \mathrm{v_m}$ and $P_2 = \mathrm{v_{m+1}},\ \mathrm{v_{m+2}},\ \ldots,\ \mathrm{v_{k-1}},\ \mathrm{v_0}$. Both P_1 and P_2 are of length at least one (because $(\mathrm{v_0}, \mathrm{v_1})$ and $(\mathrm{v_m}, \mathrm{v_{m+1}})$ are in different rows, by Lemma 7). Since the length of the paths P_1 and P_2 are less than the length of C, we know that their vertex-induced subgraphs in D_M are acyclic. Thus, by Lemma 9, we conclude $(v_{a_1 b_1}, v_{a_1 b_m})$ and $(v_{a_0 b_{m+1}}, v_{a_0 b_0})$ are in E_M. Therefore, we have the three edges $(v_{a_0 b_{m+1}}, \mathrm{v_0})$, $(\mathrm{v_0}, \mathrm{v_1})$ and $(\mathrm{v_1}, v_{a_1 b_m})$ in E_M, which contradicts Lemma 8. See Figure 4.13. □

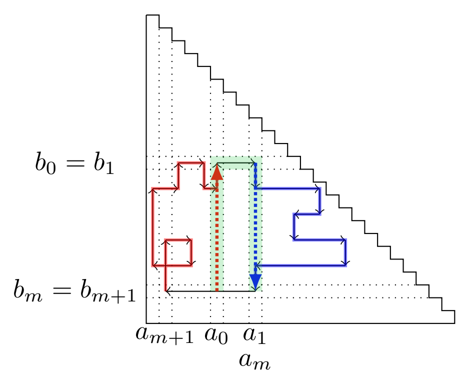

Figure 4.13: The contradiction if D_M were cyclic. The path P_1 (in blue) implies the blue dotted edge. The path P_2 (in red) implies the red dotted edge. But the green path is a contradiction.

Theorem 2 follows directly from this lemma. The lemma shows that our orientation forms a partial order on the tableau entries. By Lemma 6, we guarantee that if a tableau agrees with our orientation, then it is balanced and has the desired skeleton. As a result, *any* tableau that realizes this partial order would be a balanced tableau with the desired skeleton.

4.3.4 The digraph D_M is acyclic — proof 2

Here, we use a result from Felsner and Weil [32] to prove that D_M is acyclic. We first describe their terminology and their related result.

Definition 6. *Let $- \prec +$. For an integer r such that $1 \le r \le n$ an r-* signotope *on $[n]$ is a function α from the r-element subsets of $[n]$ to $\{-,+\}$ such that for every $(r+1)$-element subset P of $[n]$ and all i, j, and k such that $1 \le i < j < k \le r+1$ either $\alpha(P^{\lfloor i \rfloor}) \preceq \alpha(P^{\lfloor j \rfloor}) \preceq \alpha(P^{\lfloor k \rfloor})$ or $\alpha(P^{\lfloor i \rfloor}) \succeq \alpha(P^{\lfloor j \rfloor}) \succeq \alpha(P^{\lfloor k \rfloor})$, where $P^{\lfloor x \rfloor}$ denotes the set P minus the xth largest element of P. We refer to this property as* monotonicity*.*

Felsner and Weil associate an r-signotope α on $[n]$ with a directed graph whose vertices are the $(r-1)$-element subsets of $[n]$, and whose edges are

$$\begin{aligned} \to_\alpha \quad \equiv \quad & \{P^{\lfloor i \rfloor} \to_\alpha P^{\lfloor j \rfloor} \mid P \in \binom{[n]}{r},\ 1 \le i < j \le r,\ \alpha(P) = +\} \\ & \cup \\ & \{P^{\lfloor j \rfloor} \to_\alpha P^{\lfloor i \rfloor} \mid P \in \binom{[n]}{r},\ 1 \le i < j \le r,\ \alpha(P) = -\}. \end{aligned}$$

They prove the following lemma.

Lemma 11 (Felsner and Weil [32])**.** *For an r-signotope α on $[n]$ the graph with vertices $\binom{[n]}{r-1}$ and edges $\to_\alpha$ is acyclic.*

We show that our orientation function α_M is a 3-signotope and hence the associated directed graph, which is identical to the transpose of D_M (consisting of the edges of D_M with their directions reversed), is acyclic. This implies D_M is acyclic.

Felsner and Weil give an abstract combinatorial proof that the graphs associated with signotopes are acyclic in general. Lemma 10 may be viewed as an alternative and perhaps more intuitive proof for the acyclicity of such graphs when restricted to 3-signotopes. While the substantial ideas of both proofs are the same, our proof explains in more detail the manner in which sets of edges imply the existence of other edges. In the case of a cycle, this results in a contradiction.

We restate Definition 6 for the case $r = 3$. (See pages 14–15 of Felsner and Weil [32].)

Definition 7. *A 3-signotope on $[n]$ is a function $\alpha : \binom{[n]}{3} \to \{-,+\}$ such that for every 4-element subset $\{a,b,c,d\}$ of $[n]$ with $a < b < c < d$, the* orientation sequence *$(\alpha(abc), \alpha(abd), \alpha(acd), \alpha(bcd))$ is* monotone *(that is, it has at most one change of sign). In other words, it is one of the eight columns of the table in Figure 4.14.*

$\alpha(abc)$	+	−	−	−	+	+	+	−
$\alpha(abd)$	+	+	−	−	+	+	−	−
$\alpha(acd)$	+	+	+	−	+	−	−	−
$\alpha(bcd)$	+	+	+	+	−	−	−	−

Figure 4.14: Monotone orientation sequences.

We use this definition in proving the following lemma.

Lemma 12. *Let M be a persistent n-triangle. The orientation function $\alpha_M : \binom{[n]}{3} \to \{-,+\}$ where*

$$\alpha_M(abc) = \begin{cases} + & \text{if } rect_M(abc) \text{ contains a one entry} \\ - & \text{otherwise,} \end{cases}$$

is a 3-signotope. Moreover, for every 4-tuple, the resulting orientation sequence excludes $(-,-,-,+)$ and $(+,-,-,-)$.

Proof. We need to show that for any 4-element subset $P = \{a,b,c,d\}$ of $[n]$ the orientation sequence $(\alpha_M(abc), \alpha_M(abd), \alpha_M(acd), \alpha_M(bcd))$, where $a < b < c < d$, is monotone. Let $P^{\lfloor x \rfloor}$ denote the set P minus the xth largest element of P. We say $\alpha_M(P^{\lfloor i \rfloor})$ is at *position* i of the orientation sequence.

We know either $\alpha_M(bcd) = +$ or $\alpha_M(bcd) = -$. Consider the following cases.

Case 1. $\alpha_M(bcd) = +$:
Since $\alpha_M(bcd) = +$, we know there is a one entry in $rect_M(bcd)$. As a result $rect_M(acd)$ contains a one entry, which implies $\alpha_M(acd) = +$. Thus, we may get a non-monotone orientation sequence only if $\alpha_M(abd) = -$ and $\alpha_M(abc) = +$. Assume $\alpha_M(abd) = -$. Hence $rect_M(abd)$ consists of all zero entries and so does $rect_M(abc)$. Therefore we have $\alpha_M(abc) = -$, which concludes this case. See Figure 4.15.

Case 2. $\alpha_M(bcd) = -$:
Suppose we have a non-monotone orientation sequence. Thus the sequence contains at least one + sign. We show that the rightmost + sign in the

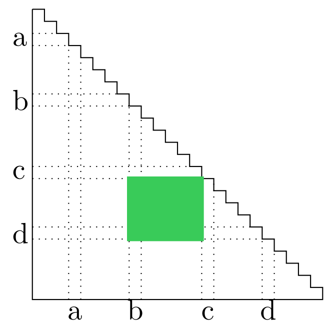

(a) From $\alpha_M(bcd) = +$, we know $rect_M(bcd)$ (in green) contains a one entry.

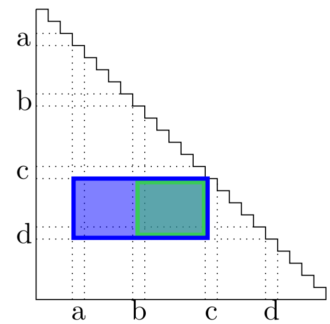

(b) Since $rect_M(bcd)$ (in green) is contained in $rect_M(acd)$ (in blue), we know $\alpha_M(bcd) = +$ implies $\alpha_M(acd) = +$.

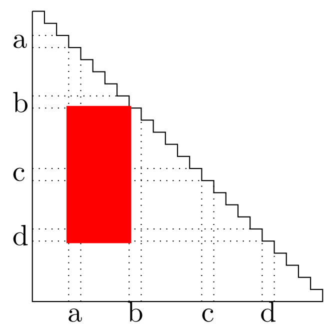

(c) From $\alpha_M(abd) = -$, we know $rect_M(abd)$ (in red) consists of all zero entries.

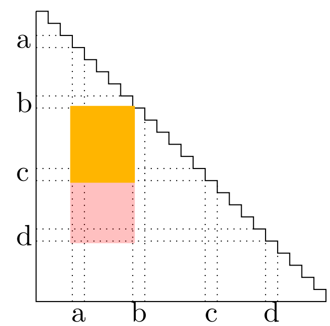

(d) Since $rect_M(abc)$ (in yellow) is contained in $rect_M(abd)$ (in red), we know $\alpha_M(abd) = -$ implies $\alpha_M(abc) = -$.

Figure 4.15: Case 1: $\alpha_M(bcd) = +$.

orientation sequence propagates all the way to the left, which implies that the sequence is monotone. Note that if $\alpha_M(abd) = -$ then $\alpha_M(abc) = -$. This is because from $rect_M(abd)$ containing only zero entries, we conclude that $rect_M(abc)$ consists of all zero entries. Therefore, the rightmost $+$ sign is either at position 2 or at position 3. We consider each case below.

Case 2.1. $\alpha_M(abd) = +$ and $\alpha_M(acd) = -$:
We have $rect_M(abd)$ contains a one entry whereas $rect_M(acd)$ consists of all zero entries. This implies that $rect_M(abc)$ contains a one entry. Thus $\alpha_M(abc) = +$, which concludes this case. See Figure 4.16.

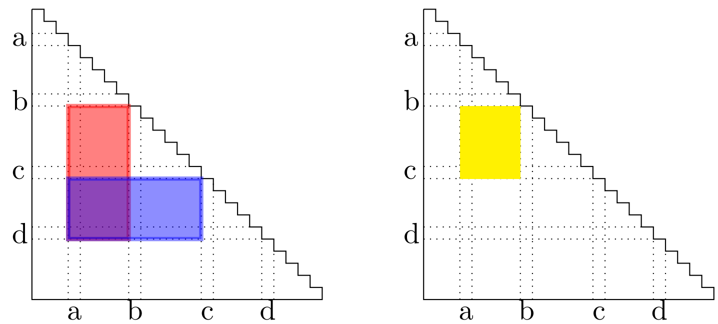

(a) From $\alpha_M(abd) = +$ and $\alpha_M(acd) = -$, we know $rect_M(abd)$ (in red) contains a one entry whereas $rect_M(acd)$ (in blue) consists of all zero entries.

(b) We conclude $rect_M(abc)$ (in yellow) contains a one entry, and hence $\alpha_M(abc) = +$.

Figure 4.16: Case 2.1: $\alpha_M(bcd) = -$, $\alpha_M(acd) = -$, and $\alpha_M(abd) = +$.

Case 2.2. $\alpha_M(acd) = +$:
We know $rect_M(acd)$ contains a one entry but $rect_M(bcd)$ consists of all zero entries. This implies that there is a one entry in the submatrix $M[[a..b);(c..d]]$. Since $M[[a..b);(c..d]]$ contains a one entry and $rect_M(bcd)$ consists of zero entries, by Lemma 5, we conclude that $rect_M(abc)$ contains a one entry. We know $rect_M(abd)$ covers $rect_M(abc)$ and hence it contains a one entry. Thus $\alpha_M(abc) = +$ and $\alpha_M(abd) = +$, which conclude this case. See Figure 4.17.

Therefore, α_M gives a monotone orientation sequence for each 4-tuple, and hence is a 3-signotope. Moreover, by the argument above, it is easy to see that it never gives the monotone sequence $(-,-,-,+)$ or $(+,-,-,-)$, and thus it is a restricted subclass of 3-signotopes.

□

This lemma together with Lemma 11 imply Lemma 10, which consequently concludes the proof of Theorem 2.

A 3-signotope is realizable as an ordered generalized configuration of points (by Theorem 7 of Felsner and Weil [32] and the concept of duality between pseudoline arrangements and generalized configurations of points). However, not all 3-signotopes are realizable as ordered point sets. Our orientation is a strict subclass of 3-signotopes. It excludes the monotone orientation sequences $(-,-,-,+)$ and $(+,-,-,-)$, which puts more constraints on the point set realizing it (that is, it forbids certain substructures in the realization). This may help prove realizability of this subclass.

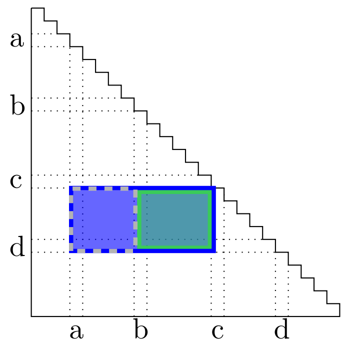

(a) From $\alpha_M(acd) = +$ and $\alpha_M(bcd) = -$, we know $rect_M(acd)$ (in blue) contains a one entry whereas $rect_M(bcd)$ (in green) consists of all zero entries. Since $rect_M(bcd)$ (in green) is contained in $rect_M(acd)$ (in blue), we conclude $M[[a..b);(c..d]]$ (dashed outline) contains a one entry.

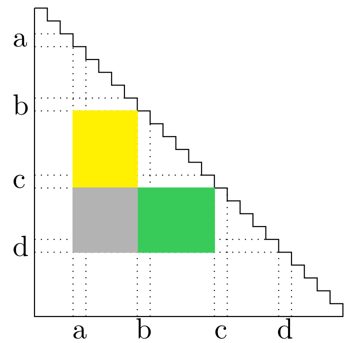

(b) The bar-property implies that if $M[[a..b);(c..d]]$ (in grey) contains a one entry and $rect_M(bcd)$ (in green) consists of all zero entries, then $rect_M(abc)$ (in yellow) contains a one entry.

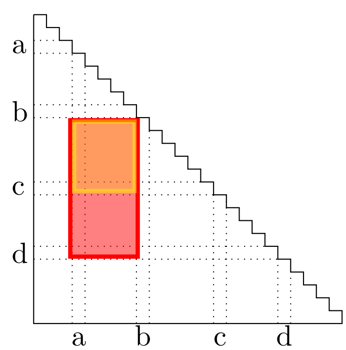

(c) Since $rect_M(abc)$ (in yellow) is contained in $rect_M(abd)$ (in red), we know $\alpha_M(abc) = +$ implies $\alpha_M(abd) = +$.

Figure 4.17: Case 2.2: $\alpha_M(bcd) = -$ and $\alpha_M(acd) = +$.

Chapter 5

Concluding remarks

Our approach may be interpreted in the context of oriented matroids. The same context has been used by Abello and Kumar [5] and later by O'Rourke and Streinu [50] in addressing visibility graphs of pseudo simple polygons. Both these results require more geometric information than the mere vertex visibility graphs in order to establish the triple orientations. In fact, they use additional combinatorial concepts to impute unique blocking vertices (articulation points in the terminology of O'Rourke and Streinu) to pairs of vertices that are not mutually visible, and consequently to determine a unique shortest path between two distinct vertices. The orientation of a triple is then based on whether all three vertices lie on the same shortest path or not. In our work, we restrict ourselves to terrain visibility graphs. We determine a triple orientation based on whether the associated half-strict rectangle contains a one entry or not. Our orientation simplifies the orientation defined by Abello and Kumar [7] and O'Rourke and Streinu [50], since it is established solely from the vertex visibility graph. For terrains, a shortest ordered path in the visibility graph determines the Euclidean shortest path in any realization of the graph. The term shortest path, in what follows, refers to the shortest ordered path in the visibility graph. It is easy to see that, for $a < b < c$, if the half-strict abc-rectangle contains a one entry, then b is not on any shortest path containing both a and c. This is because a one entry in the half-strict abc-rectangle indicates the existence of an edge $\{u, v\}$ that passes over b (where $a \leq u < b < v \leq c$). Let $\{u, v\}$ be the longest such edge. If b lies on a shortest path from a to c then $\{u, v\}$ is not $\{a, c\}$. Also, the shortest path must cross $\{u, v\}$, but then the X-property contradicts $\{u, v\}$ being the longest such edge. Similarly, if the half-strict abc-rectangle does not contain a one entry, then all shortest paths containing a and c, contain b as well. This implies that our orientation happens to be the same as the orientation defined in the previous results. However, we use specific graph theoretic properties of terrain visibility graphs, and consequently we show more about the structure of the underlying oriented matroid, which may help to prove realizability.

Since oriented matroids are closely related to Knuth's axioms [44] on three-point orientation predicates, we give a succinct description of these axioms and the relevant results here. Later we use Knuth's terminology, which may help us better distinguish between the results from Abello and Kumar [7], O'Rourke and Streinu [50], and our work.

The following axioms hold for counterclockwise relations between sets of up to five points in the Euclidean plane. (The primitive predicate pqr states that the circle through points (p, q, r) is traversed counterclockwise when we encounter the points in cyclic order p, q, r, p.):

Axiom 1 (cyclic symmetry) $pqr \Longrightarrow qrp$.

Axiom 2 (antisymmetry) $pqr \Longrightarrow \neg prq$.

Axiom 3 (nondegeneracy) $pqr \vee prq$.

Axiom 4 (interiority) $tqr \wedge ptr \wedge pqt \Longrightarrow pqr$.

Axiom 5 (transitivity) $tsp \wedge tsq \wedge tsr \wedge tpq \wedge tqr \Longrightarrow tpr$.

The ternary relations that satisfy Axioms $1-5$ are called CC systems (short for "counterclockwise systems"). The ternary relations that satisfy Axioms $1-3$ and Axiom 5 are pre-CC systems. *Uniform* oriented matroids correspond to Axiom 5. Axiom 4 is equivalent to a special class of uniform oriented matroids, called *acyclic*. In fact, Axiom 5 captures almost all the important properties of Axiom 4, and thus pre-CC systems are not much different from full CC systems. More precisely, a set of triples is a pre-CC system if and only if it can be obtained from a CC system by negating a subset of its points where negating a point complements the value of all triples that contain that point.

Abello and Kumar [7] show that their orientation forms a simplicial chirotope of rank 3, which equivalently represents a uniform rank 3 oriented matroid. It has been shown that simplicial rank 3 chirotopes are equivalent to Knuth's pre-CC systems [11, 44]. Abello and Kumar show that if the resulting chirotope is realizable then their necessary conditions would characterize visibility graphs of simple polygons (when the additional information of the blocking vertex assignment is given). However, since some pre-CC systems that satisfy their conditions are not realizable (that is, they do not arise from actual points in the plane), such conditions fail to characterize visibility graphs of straight-line simple polygons.

O'Rourke and Streinu [50] show that their orientation forms a CC system. CC systems are equivalent to uniform *acyclic* rank 3 oriented matroids (or identically *affine* rank 3 simplicial chirotopes). Uniform acyclic rank 3 oriented matroids are in turn equivalent to arrangements of pseudolines, by the Folkman-Lawrence representability theorem [33]. Shor [55] and Mnëv [47] show that not all arrangements of pseudolines are stretchable. In fact, Goodman and Pollack [39] show that almost all CC systems on n points are unrealizable, in the limit as $n \to \infty$. Even though O'Rourke and Streinu [50] suggest that their orientation forms a strict subclass of all CC systems, Streinu [58, 59] proves that their characterization (of a variant of pseudo visibility graphs containing more information than vertex pseudo visibility graphs) allows graphs that are not visibility graphs of straight-line polygons.

We study terrain visibility graphs, which induce stronger conditions. Namely, we know that terrain visibility graphs satisfy the X-property in addition to the properties of simple pseudo-polygon visibility graphs. Using the specific graph properties of terrain visibility graphs, we prove stronger results: our orientation is a strict subclass of 3-signotopes. By case analysis, it is easy to see that every 3-signotope is a CC system (and consequently a pre-CC system). However, we may have CC systems that are not 3-signotopes. For instance, let f be a *cyclic symmetric* and *antisymmetric* ternary relation on the set $\{a, b, c, d, e\}$ such that $f(abc) = -$, $f(abd) = +$, $f(abe) = -$, $f(acd) = -$, $f(ace) = -$, $f(ade) = +$, $f(bcd) = -$, $f(bce) = -$, $f(bde) = +$, and $f(cde) = +$, where $a < b < c < d < e$. Then f satisfies all the five axioms of CC systems, but does not satisfy the *monotonicity* property on the 4-tuples (a, b, c, d) and (a, b, d, e), and hence is not a 3-signotope. Figure 5.1 illustrates the relation between pre-CC systems, CC systems, and 3-signotopes.

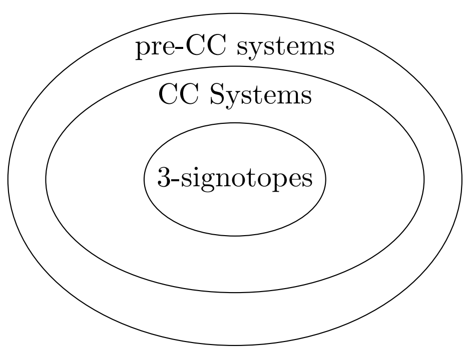

Figure 5.1: The relation between pre-CC systems, CC systems, and 3-signotopes.

Knuth's CC systems may be interpreted as characterizing the local sequences of generalized configurations of points. Likewise, Streinu [57] gives the necessary and sufficient properties of local sequences that arise from generalized configurations of points (where the resulting *allowable sequence*[2] may not include the identity permutation). These properties suggest that an orientation (clockwise or counterclockwise) of every triple that obeys a "generalized transitivity" law, is realizable as a generalized configuration of points *for some ordering of the point set* (or equivalently, it forms a uniform rank 3 acyclic oriented matroid). In this work, we study *ordered graphs* (that is, a graph with a total order over its vertices). We orient all triples such that they are *balanced*, and prove that our orientation obeys the transitivity law; and hence forms a partial order on the slope order requirements[3]. Felsner and Weil [32] show that the transitivity of the balanced orientations in an ordered set is identical to the *monotonicity* of all the 4-tuples (or in other words, to the orientation function being a 3-signotope). This suggests that our first proof for the acyclicity of the M-tableau relation digraph implies that our orientation function is a 3-signotope. We prove this fact directly in the second proof.

Abello et al. [4] show that persistent graphs, if representing visibility, guarantee a slope order on the lines connecting all pairs of terrain vertices (in a possible realization) that is consistent with the desired visibility graph, through an algorithmic approach of time complexity $O(n^5)$. In our work, not only do we prove this (by showing the acyclicity of the M-tableau relation digraph), but we also establish an orientation for *all* triples which gives a *certain* 3-signotope. Our proof techniques are much simpler, and clarify the implications of the X-property and the bar-property on the ordering of the slopes. Lastly, we may orient all triples in $O(n^3)$ time, which consequently gives an $O(n^3)$ time algorithm for constructing a slope ordering. The algorithm is as follows: (Remember that the orientation of a triple may be equivalently determined based on whether all three vertices lie on the same

[2] Suppose $\mathcal{P}$ is the set of numbered points $\{p_1, \ldots, p_n\}$ in the Euclidean plane. Let L be a directed line not orthogonal to any line determined by two members of $\mathcal{P}$, and project $\mathcal{P}$ orthogonally onto L. The projected points define a permutation of the indices $1, \ldots, n$. If L rotates counterclockwise, the permutation changes whenever L passes through a direction perpendicular to a line connecting two points of $\mathcal{P}$. (When L has turned through an angle π the resulting permutation of the indices will be the reverse of the initial permutation.) Such a periodic sequence of permutations is called an *allowable sequence*.

[3] The concept of the slope order may also be generalized to pseudolines. If a pseudoline arrangement is intersected with a vertical line such that all vertices of the arrangement lie to its right, then the order in which the pseudolines cross the vertical line (decreasing by the y-coordinates of the crossings) is the (increasing) slope order of the pseudolines [60].

shortest ordered path in the visibility graph or not.) First, we initialize all triple orientations negatively. Given a persistent graph $G = ([n], E)$, for every vertex i, we run BFS on the induced subgraph of G containing vertices $[i..n]$ such that the neighbours of vertices are visited in order. This way we obtain the ordered shortest path trees from the visibility graph. At the same time that we execute the BFS procedures, we positively orient the triples that occur on the same path in the BFS trees. Therefore we determine our triple orientation (that is, the function α_M) in $O(n^3)$ time in total. From the triple orientation α_M, we construct the digraph D_M (by Definition 5). The digraph has $O(n^2)$ vertices and $O(n^3)$ edges (which represent the inequality relations imposed by the orientation). We do a topological sort on the digraph in $O(n^3)$ time. As a result, we give an $O(n^3)$ time algorithm for constructing a slope ordering.

Finally, we would like to mention that the partial order resulting from our orientation may not encode all possible slope orderings. This is because we orient *all* triples, which may put more constraints on the tableau entries than are required in order to guarantee a balanced tableau with the desired skeleton. In general, a triple orientation of the point set realizing an induced subgraph may not be extensible to a triple orientation of the point set realizing the entire graph. For instance, in Figure 5.2, the induced subgraph on the vertex set $\{1, 2, 3, 4\}$ admits either orientation for the points $1, 3, 4$ for a realization. But the triple needs to be oriented positively in any realization of the entire graph. It is easy to see that the additional constraints that our

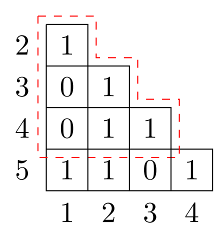

Figure 5.2: A triple orientation of a smaller graph may not be extensible to a triple orientation of a larger one: the smaller tableau admits either orientation for the points $1, 3, 4$ while the points need to be oriented positively in the larger tableau.

orientation defines make the triple orientation of smaller subgraphs extensible to the triple orientation of the entire graph. This suggests that, using our orientation, it may be possible to reconstruct the terrain by incrementally realizing the point set. In addition, our orientation forbids certain substructures in any realization (as it excludes $(-,-,-,+)$ and $(+,-,-,-)$ from the possible orientation sequences on the 4-tuples). This puts more constraints on the point set realizing it, which may help to prove realizability and perhaps to reconstruct the terrain.

Bibliography

[1] J. Abello. The majority rule and combinatorial geometry (via the symmetric group). *Annales Du Lamsade*, 3:1–13, October 2004.

[2] J. Abello and Ö. Egecioglu. Visibility graphs of staircase polygons with uniform step length. *International Journal of Computational Geometry and Applications*, 3(1):27–37, 1993.

[3] J. Abello, Ö. Egecioglu, and K. Kumar. Visibility graphs of staircase polygons and the weak Bruhat order II: From maximal chains to polygons. preprint.

[4] J. Abello, Ö. Egecioglu, and K. Kumar. Visibility graphs of staircase polygons and the weak Bruhat order I: From visibility graphs to maximal chains. *Discrete and Computational Geometry*, 14(3):331–358, 1995.

[5] J. Abello and K. Kumar. Visibility graphs and oriented matroids (extended abstract). In R. Tamassia and I. Tollis, editors, *Graph Drawing*, volume 894 of *Lecture Notes in Computer Science*, pages 147–158. Springer Berlin / Heidelberg, 1995.

[6] J. Abello and K. Kumar. Visibility graphs of 2-spiral polygons (extended abstract). In *Proceedings of the Second Latin American Symposium on Theoretical Informatics*, LATIN '95, pages 1–15, London, UK, 1995. Springer-Verlag.

[7] J. Abello and K. Kumar. Visibility graphs and oriented matroids. *Discrete and Computational Geometry*, 28:449–465, 2002. 10.1007/s00454-002-2881-6.

[8] J. Abello, K. Kumar, and Ö. Egecioglu. A combinatorial view of visibility graphs of simple polygons. In *Proceedings of the Fifth International Conference on Computing and Information*, ICCI '93, pages 87–92, Washington, DC, USA, 1993. IEEE Computer Society.

[9] J. Abello, H. Lin, and S. Pisupati. On visibility graphs of simple polygons. *Congressus Numerantium*, 90:119–128, 1992.

[10] T. Asano, T. Asano, L. Guibas, J. Hershberger, and H. Imai. Visibility of disjoint polygons. *Algorithmica*, 1:49–63, January 1986.

[11] P. Baier. NP-completeness of partial chirotope extendibility. *ArXiv Mathematics e-prints*, math/0504430, April 2005. arXiv:math/0504430v1 [math.CO].

[12] B. Ben-Moshe, M. Katz, and J. Mitchell. A constant-factor approximation algorithm for optimal terrain guarding. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, SODA '05, pages 515–524, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics.

[13] C. Berge. Some classes of perfect graphs. In *Six Papers in Graph Theory*, pages 1–21, Indian Statistical Institute, Calcutta, 1963.

[14] A. Björner, M. Las Vergnas, B. Sturmfels, N. White, and G. Ziegler. *Oriented matroids*, volume 46 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, 1993.

[15] J. Bokowski, J. Richter, and B. Sturmfels. Nonrealizability proofs in computational geometry. *Discrete and Computational Geometry*, 5:333–350, 1990. 10.1007/BF02187794.

[16] J. Bokowski and B. Sturmfels. On the coordinatization of oriented matroids. *Discrete and Computational Geometry*, 1:293–306, 1986. 10.1007/BF02187702.

[17] J. Bokowski and B. Sturmfels. Algebraic criteria for geometric realizability. In *Computational Synthetic Geometry*, volume 1355 of *Lecture Notes in Mathematics*, pages 61–86. Springer Berlin / Heidelberg, 1989. 10.1007/BFb0089257.

[18] J. Bokowski and B. Sturmfels. *Computational synthetic geometry.* Springer-Verlag, Berlin and New York, 1989.

[19] J. Bokowski and B. Sturmfels. An infinite family of minor-minimal nonrealizable 3-chirotopes. *Mathematische Zeitschrift*, 200:583–589, 1989. 10.1007/BF01160956.

[20] B. Chazelle. Triangulating a simple polygon in linear time. *Discrete and Computational Geometry*, 6:485–524, 1991.

[21] D. Z. Chen, V. Estivill-Castro, and J. Urrutia. Optimal guarding of polygons and monotone chains. In *Proceedings of the Seventh Canadian Conference on Computational Geometry*, pages 133–138, 1995.

[22] Seung-Hak Choi, Sung Yong Shin, and Kyung-Yong Chwa. Characterizing and recognizing the visibility graph of a funnel-shaped polygon. *Algorithmica*, 14:27–51, 1995.

[23] M. Chudnovsky, N. Robertson, P. Seymour, and R. Thomas. The strong perfect graph theorem. *Annals of Mathematics*, 164(1):51–229, 2006.

[24] P. Colley. Visibility graphs of uni-monotone polygons. Master's thesis, Department of Computer Science, University of Waterloo, Waterloo, Canada, 1991.

[25] P. Colley. Recognizing visibility graphs of unimonotone polygons. In *Proceedings of the Fourth Canadian Conference on Computational Geometry*, pages 29–34, 1992.

[26] P. Colley, A. Lubiw, and J. Spinrad. Visibility graphs of towers. *Computational Geometry: Theory and Applications*, 7:161–172, February 1997.

[27] C. Coullard and A. Lubiw. Distance visibility graphs. In *Proceedings of the seventh annual symposium on Computational geometry*, SCG '91, pages 289–296, New York, NY, USA, 1991. ACM.

[28] H. ElGindy. *Hierarchical decomposition of polygons with applications.* PhD thesis, School of Computer Science, McGill University, Montreal, Canada, 1985.

[29] H. Everett. *Visibility graph recognition.* PhD thesis, Department of Computer Science, University of Toronto, Toronto, Canada, 1990.

[30] H Everett and D.G Corneil. Recognizing visibility graphs of spiral polygons. *Journal of Algorithms*, 11(1):1 – 26, 1990.

[31] H. Everett, F. Hurtado, and M. Noy. Stabbing information of a simple polygon. In *Proceedings of the Eight Canadian Conference on Computational Geometry*, pages 74–79, 1996.

[32] S. Felsner and H. Weil. Sweeps, arrangements and signotopes. *Discrete Applied Mathematics*, 109:67–94, 2001.

[33] J. Folkman and J. Lawrence. Oriented matroids. *Journal of Combinatorial Theory, Series B*, 25(2):199 – 236, 1978.

[34] S. K. Ghosh. On recognizing and characterizing visibility graphs of simple polygons. In R. Karlsson and A. Lingas, editors, *SWAT 88*, volume 318 of *Lecture Notes in Computer Science*, pages 96–104. Springer Berlin / Heidelberg, 1988.

[35] S. K. Ghosh. On recognizing and characterizing visibility graphs of simple polygons. *Discrete and Computational Geometry*, 17(2):143–162, 1997.

[36] S. K. Ghosh. *Visibility Algorithms in the Plane*. Cambridge University Press, New York, NY, USA, 2007.

[37] S. K. Ghosh and P. P. Goswami. Unsolved problems in visibility graphs of points, segments and polygons. *CoRR*, abs/1012.5187, 2010.

[38] J. E. Goodman and R. Pollack. Semispaces of configurations, cell complexes of arrangements. *Journal of Combinatorial Theory, Series A*, 37(3):257 – 293, 1984.

[39] J. E. Goodman and R. Pollack. Upper bounds for configurations and polytopes in $\mathbf{R}^d$. *Discrete and Computational Geometry*, 1:219–227, 1986. 10.1007/BF02187696.

[40] J. E. Goodman and R. Pollack. Allowable sequences and order types in discrete and computational geometry. In J. Pach, editor, *New Trends in Discrete and Computational Geometry*, volume 10 of *Algorithms and Combinatorics*, pages 103–134. Springer-Verlag, Berlin, 1993.

[41] J. Hershberger. Finding the visibility graph of a simple polygon in time proportional to its size. In *Proceedings of the third annual symposium on Computational geometry*, SCG '87, pages 11–20, New York, NY, USA, 1987. ACM.

[42] L. Jackson and S. K. Wismath. Orthogonal polygon reconstruction from stabbing information. *Computational Geometry: Theory and Applications*, 23:69–83, July 2002.

[43] J. A. King. Approximation algorithms for guarding 1.5 dimensional terrains. Master's thesis, Department of Computer Science, The University of British Columbia, Vancouver, Canada, 2005.

[44] D. E. Knuth. *Axioms and hulls*, volume 606 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, 1992.

[45] Y. Lin and S. Skiena. Complexity aspects of visibility graphs. *International Journal of Computational Geometry and Applications*, 5(3):289–312, 1995.

[46] T. Lozano-Pérez and M. A. Wesley. An algorithm for planning collision-free paths among polyhedral obstacles. *Commun. ACM*, 22:560–570, October 1979.

[47] 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.

[48] J. O'Rourke. *Art gallery theorems and algorithms*, chapter 7. Oxford University Press, Inc., New York, NY, USA, 1987.

[49] J. O'Rourke. Computational geometry column 18. *SIGACT News*, 24:20–25, January 1993.

[50] J. O'Rourke and I. Streinu. Vertex-edge pseudo-visibility graphs: characterization and recognition. In *Proceedings of the thirteenth annual symposium on Computational geometry*, SCG '97, pages 119–128, New York, NY, USA, 1997. ACM.

[51] J. O'Rourke and I. Streinu. The vertex-edge visibility graph of a polygon. *Computational Geometry: Theory and Applications*, 10:105–120, May 1998.

[52] J. Richter-Gebert. *On the realizability problem of combinatorial geometries decision methods*. PhD thesis, Technische Hochschule Darmstadt, 1992.

[53] J. Richter-Gebert. Euclideaness and final polynomials in oriented matroid theory. *Combinatorica*, 13:259–268, 1993. 10.1007/BF01202352.

[54] J. Richter-Gebert and G. M. Ziegler. Oriented matroids. In J. E. Goodman and J. O'rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 6, pages 129–151. CRC Press, Boca Raton, second edition, 2004.

[55] 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.

[56] G. Srinivasaraghavan and A. Mukhopadhyay. A new necessary condition for the vertex visibility graphs of simple polygons. *Discrete and Computational Geometry*, 12:65–82, 1994. 10.1007/BF02574366.

[57] I. Streinu. Clusters of stars. In *Proceedings of the thirteenth annual symposium on Computational geometry*, SCG '97, pages 439–441, New York, NY, USA, 1997. ACM.

[58] I. Streinu. Stretchability of star-like pseudo-visibility graphs. In *Proceedings of the fifteenth annual symposium on Computational geometry*, SCG '99, pages 274–280, New York, NY, USA, 1999. ACM.

[59] I. Streinu. Non-stretchable pseudo-visibility graphs. *Computational Geometry: Theory and Applications*, 31:195–206, June 2005.

[60] I. Streinu, K. Bezdek, J. Pach, T. Dey, J. Chen, D. Kravets, N. Amato, and W. R. Franklin. Discrete and computational geometry. In K. H. Rosen, J. G. Michaels, J. L. Gross J. W. Grossman, and D. R. Shier, editors, *Handbook of Discrete and Combinatorial Mathematics*, chapter 13. CRC Press, 1999.

[61] R. Tamassia and I. Tollis. A unified approach to visibility representations of planar graphs. *Discrete and Computational Geometry*, 1:321–341, 1986. 10.1007/BF02187705.

[62] R. E. Tarjan and C. J. Van Wyk. An $O(n \log \log n)$-time algorithm for triangulating a simple polygon. *SIAM Journal on Computing*, 17:143–178, February 1988.

[63] N. Vijay. On testing the necessary conditions for visibility graphs of simple polygons. Research report, 1998.

[64] E. Welzl. Constructing the visibility graph for n-line segments in $O(n^2)$ time. *Information Processing Letters*, 20(4):167 – 171, 1985.

[65] S. K. Wismath. Characterizing bar line-of-sight graphs. In *Proceedings of the first annual symposium on Computational geometry*, SCG '85, pages 147–152, New York, NY, USA, 1985. ACM.