Orthogonal Drawings of Cycles in 3D Space *
(Extended Abstract)

Giuseppe Di Battista[1], Giuseppe Liotta[2], Anna Lubiw[3], and Sue Whitesides[4]

[1] Dipartimento di Informatica ed Automazione, Università di Roma Tre, Roma, Italy.
gdb@dia.uniroma3.it

[2] Dipartimento di Ingegneria Elettronica e dell'Informazione,
Università di Perugia, Perugia, Italy.
liotta@diei.unipg.it

[3] Department of Computer Science, University of Waterloo,Waterloo, Canada.
alubiw@daisy.uwaterloo.ca

[4] School of Computer Science, McGill University,Montreal, Canada.
sue@cs.mcgill.ca

Abstract. Let C be a directed cycle, whose edges have each been assigned a desired direction in $3D$ (East, West, North, South, Up, or Down) but no length. We say that C is a shape cycle. We consider the following problem. Does there exist an orthogonal drawing Γ of C in $3D$ such that each edge of Γ respects the direction assigned to it and such that Γ does not intersect itself? If the answer is positive, we say that C is simple. This problem arises in the context of extending orthogonal graph drawing techniques and VLSI rectilinear layout techniques from $2D$ to $3D$. We give a combinatorial characterization of simple shape cycles that yields linear time recognition and drawing algorithms.

1 Introduction

The *topology-shape-metrics approach* [4] for constructing an orthogonal drawing of a planar graph in $2D$ consists of three main steps, called planarization, orthogonalization, and compaction. The planarization step determines an embedding, i.e., the face cycles, for the graph in the plane. The orthogonalization step determines an orthogonal representation of the input graph, i.e. a labeling for each edge (u, v) of the graph that defines the shape of (u, v) in the final drawing. For example, (u, v) could be labeled $NESNE$, which would say "starting from u first go North, then go East, etc." Finally, the compaction step computes the drawing, giving coordinates to vertices and bends while preserving the shape of the edges determined in the orthogonalization step.

* Research partially supported by operating grants from the Natural Sciences and Engineering Research Council (NSERC) of Canada, by the project "Algorithms for Large Data Sets: Science and Engineering" of the Italian Ministry of University and Scientific and Technological Research (MURST 40%), and by the project "Geometria Computazionale Robusta con Applicazioni alla Grafica ed a CAD" of the Italian National Research Council (CNR).

J. Marks (Ed.): GD 2000, LNCS 1984, pp. 272–283, 2001.

The topology-shape-metrics approach for $2D$ orthogonal drawings has been the subject of much literature. For each step of the approach, different optimization problems (for example minimizing the number of bends, minimizing the area, minimizing the maximum edge length) have been studied, and papers providing optimal algorithms and effective heuristics have been presented. An essential prerequisite of the topology-shape-metrics approach is a characterization of those graphs with edges labeled by orthogonal directions that can be drawn without crossings, while respecting the desired shapes for the edges. This problem has been studied in several papers, including [11,12]. The problem has been generalized to non-orthogonal polygons and to graphs in [6,9,13].

While the literature on $3D$ orthogonal drawings is quite rich (see, e.g. [1,7,8, 14,16]), the extension of the topology-shape-metrics approach to $3D$ has, as far as we know, not been previously explored. A major difficulty is that in $3D$, there is no counterpart to the $2D$ characterization of orthogonal representations. By studying orthogonal representations of cycles in $3D$, this paper represents a first step toward the goal of extending the topology-shape-metrics approach to $3D$.

A $3D$ *shape path* σ is an ordered sequence of labels for the edges of an oriented (graph theoretical) path P, where each label specifies a direction $East$, $West$, $North$, $South$, Up, or $Down$ for the corresponding edge, and σ contains at least one of each oppositely directed pair of direction labels. Similarly, a *shape cycle* σ is a circularly ordered sequence of direction labels for the edges of an oriented cycle C, where each label specifies a direction for the corresponding edge, and each of the six directions occurs at least once as a label. The *simplicity testing problem for* σ is to decide whether there exists an orthogonal drawing Γ of C so that Γ is *simple* (i.e., no two edges of Γ share any points except common endpoints) and satisfies the direction constraints on its edges as specified by σ. If so, then the shape cycle σ is said to be *simple*.

Not all shape cycles are simple. For example, consider the shape cycle given by the circular sequence of labels $ESUNDWUN$, where E stands for $East$, U stands for Up, and so on. This shape cycle has no simple orthogonal drawing, even though each direction label appears at least once (see Figure 1). By contrast, its subcycle $ESDWUN$ *is* simple.

Our main result is a combinatorial characterization of simple shape cycles that yields linear time testing and drawing algorithms.

2 Overview

In $2D$, simple shape cycles were characterized by Vijayan and Widgerson [12] in terms of editing operations on the sequence of labels in the shape cycle. If their editing operations are carried out until no further application is possible, the result is a unique, reduced form for the shape cycle. A $2D$ shape cycle is simple if and only if it can be edited to give the shape cycle for a rectangle, that is, a sequence of the four distinct labels $\{E, W, N, S\}$, with consecutive labels orthogonal. The editing operations arise in the context of repeatedly taking shortcuts at U turns in a rectilinear polygon, where the shortcuts do not intersect

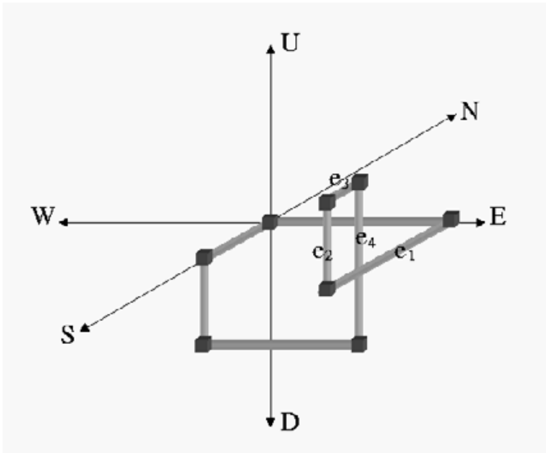

Fig. 1. In the shape cycle $\sigma = ESUNDWUN$, the labels assigned to the edges e_1, e_2, e_3, and e_4 define a "flat" . This shape cycle is *not* simple.

the boundary of the polygon. Their characterization can also be stated as follows: a $2D$ shape cycle is simple if and only if then number of right turns differs from the number of left turns by four. See also the paper by Tamassia [11].

Our recognition algorithm for simple shape cycles does *not* work by repeatedly applying editing operations to the shape cycle. Instead, it looks directly for a subsequence satisfying certain combinatorial properties, which were obtained by considering cycles of length six drawn on the edges of a cube. To state our characterization results precisely, we next introduce the concepts of *flat* and of *canonical sequence* for a shape cycle or path.

Let σ be a shape cycle or path. A *flat* of σ is a consecutive subsequence $F \subset \sigma$ that is maximal with respect to the property that any orthogonal drawing of F must consist of edges that lie on the same axis-aligned plane. For example, consider the shape cycle $\sigma = \sigma_1\sigma_2\ldots\sigma_8 = ESUNDWUN$, for which Figure 1 gives a (non-simple) drawing. Labels $\sigma_1 = E$ and $\sigma_2 = S$ must lie in an $EWNS$ flat, i.e., a flat whose labels belong to $\{E, W, N, S\}$. Since σ is a circular sequence and since, by definition, flats are maximal, label $\sigma_8 = N$ also belongs to the flat F_1 containing $\sigma_1 = E$ and $\sigma_2 = S$. Since an $EWNS$ flat cannot contain a U label, $F_1 = \sigma_8\sigma_1\sigma_2$. Shape σ contains three additional flats, namely $F_2 = \sigma_2\sigma_3\sigma_4\sigma_5 = SUND$; $F_3 = \sigma_5\sigma_6\sigma_7 = DWU$; and $F_4 = \sigma_7\sigma_8 = UN$. Note that each pair of consecutive flats F_iF_{i+1} share a *transition label* ($\sigma_2, \sigma_5, \sigma_7, \sigma_8$ in the example).

A not necessarily consecutive subsequence $\tau \subseteq \sigma$, where τ consists of k labels, is a *canonical sequence* if: (1) $1 \leq k \leq 6$; (2) the labels of τ are distinct; (3) no flat of σ contains more than three labels of τ; and (4) if a flat F of σ contains one or more labels of τ, then $\tau \cap F$ forms a consecutive subsequence of σ.

For example, the shape cycle $\sigma = ESUNDWUN$ of Figure 1 does not contain a canonical sequence of any length containing an S and a D: each of these labels occurs only once, in flat F_2, where S and D are not consecutive as elements of σ. Hence according to our characterization of simple shape cycles below, cycle σ is not simple.

Theorem 1. *A 3D shape cycle with at least two flats is simple if and only if it contains a canonical sequence of length six.*

In a companion paper [5], we introduced a simpler notion of canonical sequence in solving the *shape path reachability problem*, which is to determine, given a shape path σ and a point p in an octant, whether σ admits a simple orthogonal drawing that starts at the origin and ends at p. The main result of [5] can be summarized in concrete terms for a particular octant as follows.

Theorem 2. *[5] Let σ be a shape path, and let p be any point of the UNE octant. Then σ admits a simple orthogonal drawing that starts at the origin and ends at p if and only if σ contains a canonical sequence of length 3 containing the labels U, N, E in some order.*

At a very first glance, the necessity of the condition of Theorem 1 may appear to be an immediate consequence of Theorem 2: if σ admits a simple orthogonal drawing, then σ can be shown to split into the concatenation of two shape paths that reach opposite octants. However, the flats where two such paths join require special study, as they contain labels from each of the two paths. Hence the union of canonical sequences for each path need not yield a canonical sequence of length six for the cycle.

3 Preliminaries

We assume from now on that two adjacent labels of a shape path or cycle are neither identical, in which case they could be replaced by a single label, nor oppositely directed, in which case the shape could not be simple. Also, we omit some straightforward special case handling by considering here only shape cycles that contain at least four flats.

We regard $3D$ space as partitioned into eight open octants, eight open quadrants, six open (semi)axes directed away from the origin, and the origin itself. A triple XYZ of distinct unordered labels no two of which are opposite defines the *XYZ octant*. Similarly, a pair XY of distinct orthogonal labels defines the *XY quadrant* in $2D$ or $3D$, and a direction label X defines the X (semi)axis.

We sometimes use the term "shape" to refer either to a shape path or to a shape cycle. We sometimes say a shape is *drawable* if it is a simple shape.

To traverse a shape σ or a drawing $\Gamma(\sigma)$ of σ in the *positive sense* means to visit its labels or edges in the order specified by σ. If σ is a shape path, the starting point for a drawing $\Gamma(\sigma)$ of σ is regarded as the origin for that drawing. Visiting $\Gamma(\sigma)$ in the positive sense then orients each edge of $\Gamma(\sigma)$. An edge oriented in this way from u to v is denoted uv. It points in the direction specified by its associated label in σ.

Suppose that σ is a $3D$ shape path or cycle and that F is a flat of σ. Then *first* and *last* labels for F (also called its entry and exit labels) are determined by traversing the cyclic sequence σ in the positive sense. In a drawing of F, the *starting* or *entering point* for F is the vertex at the tail of the edge corresponding to the entry label of F.

A label Y of a shape σ is said to *occur between* two other labels X and Z if Y is met when traversing σ in the positive sense from X to Z.

Remark. Let $\phi()$ be a permutation of the six direction labels that maps opposite pairs of labels to possibly different opposite pairs (for example, ϕ might map N, S, E, W, U, D to E, W, N, S, D, U, respectively). Note that $\phi()$ defines a linear transformation of $3D$ space that determines a bijection between drawings of σ and drawings of $\phi(\sigma)$.

For concreteness, as in Theorem 2, we often state our results and proofs referring to some given octant, quadrant, or axis. However, the results can also be stated with respect to any other octant, quadrant, or axis since, by the Remark, they are preserved under the $\phi()$ transformation.

We sometimes specify the labels of a canonical sequence τ by using set notation. For example, $\{U, N, E, S\}$ might describe a canonical sequence whose directions labels are U, N, E, and S. In this notation, the order of the labels is not specified and is inherited from the shape σ once a particular subsequence τ has been chosen.

We sometimes distinguish the labels in a canonical sequence τ from the other labels of σ with special notation. To say that $\{\bar{U}, \bar{N}, \bar{E}, \bar{S}\}$ is a canonical sequence means not only that the canonical sequence contains a U, an S, an N, and an E direction label, but also that the U label in τ is a specific element $\bar{U}$ of σ, the S label is $\bar{S}$, and so on.

We say that the elements of σ that occur in a canonical sequence τ are *canonical labels*. It is useful to recall that a shape path is a sequence and that a shape cycle is a circular sequence. Thus a canonical sequence for a shape path is a sequence, and a canonical sequence for a shape cycle is a circular sequence.

4 Sufficiency

We now sketch a constructive proof that any shape cycle σ that contains a canonical sequence of length six admits a simple orthogonal drawing.

4.1 The Proof Technique

The intuition behind our construction of a drawing for a shape cycle is to imagine that it will be an elaboration of a cycle of six edges to be drawn along the edges of a box. A shape cycle of length six has one of two essentially different shapes, namely, a *chair* shape such as $UNDESW$, which has four flats, or a *skew* shape such as $UENDWS$, which has six flats. From a canonical sequence of length six we obtain (possibly after some modification) the six labels of a chair shape or a skew shape to follow the edges of a big box, with the remaining labels to be drawn as paths of short edges located near corners of the big box, serving to connect together the six long edges. The underlying chair or skew shape conveniently allows us to place the connecting paths of short segments in distinct octants by assigning long lengths to the canonical labels. In practice, the drawings we produce do not necessarily assign long lengths to the canonical labels, but this mental model gives the basic idea for the construction.

One difficulty is that, whereas σ does not contain any pairs of oppositely directed labels that are adjacent, a canonical sequence τ of σ may contain pairs of oppositely directed labels that are adjacent as elements of τ. Such a canonical sequence τ would then not provide the convenient underlying chair or skew shape for the construction. This motivates the following definition.

Definition 1. *A subsequence τ of a shape cycle σ is a* strong canonical cycle *if τ is a canonical cycle such that no two labels that are adjacent in τ are oppositely directed.*

The next lemma resolves the difficulty. Its proof is a technical case analysis in which new choices of canonical labels are substituted for old ones.

Lemma 1. *If a shape cycle σ contains a canonical sequence τ of length six, then σ contains a strong canonical sequence τ' of length six.*

Given a strong canonical sequence of length six (which can be found in linear time if one exists), we compute simple drawings for the connecting paths between the canonical labels, then assign lengths to the canonical labels so that these drawings remain in separate octants (this is made possible by the underlying chair or skew form of the canonical cycle). To ensure that this is the case, and to ensure that the cycle closes, we formulate and solve a system of linear inequalities expressing these constraints.

4.2 Constructing a Drawing

This subsection describes how to construct a drawing from a strong canonical sequence.

A drawing $\Gamma(\sigma)$ of a shape path σ is an *expanding drawing* if each segment travels one unit farther in its assigned direction than the extreme points, with respect to that direction, of the previous segments of $\Gamma(\sigma)$. A drawing $\Gamma(\sigma)$ of a shape path σ is a *doubly extensible drawing* if its first and last edges can be replaced by arbitrarily long edges without creating any intersections within the drawing of that shape path.

Lemma 2. *[5] Let σ be a shape path with n labels. Then σ admits an expanding drawing that can be computed in linear time on a real RAM. Also, if σ is such that either it consists of exactly two labels or it contains at least two flats, then σ has a doubly extensible drawing that can be computed in $O(n)$ time on a real RAM.*

We briefly review the proof, whose details are needed below for our cycle construction and its proof of correctness. The first part of the lemma follows from the algorithmic nature of the definition of expanding drawing. For the second part, note that if σ consists of exactly two labels, then it is clearly doubly extensible.

Now suppose σ has at least two flats. The subsequence strictly between the first and last elements thus contains a transition label. Place the tip of the first transition label at the origin. Working *backwards* through σ from this first transition label, create an expanding drawing for the initial subsequence of σ. Thus the transition label is the first label to be drawn, and has length 1. When eventually the first label of σ is reached, it can be drawn arbitrarily long.

To draw the remainder of σ, consider the label that immediately follows the first transition label. It must be drawn with its tail at the origin, and perpendicular to the plane of the previous flat. Working *forward* from this label, create an expanding drawing using the rule that when a new label is drawn, it extends farther by 1 in its direction than any previously drawn segment except the one that could be made arbitrarily long. Thus the first segment past the transition label is also assigned length 1, and when eventually the last segment of σ is drawn, it may be made arbitrarily long. This concludes the review of the proof.

Note that in the above proof sketch, the tip of the next-to-last segment of $\Gamma(\sigma)$, and hence the tail of its last segment, lies on the bounding box of the drawing of the remaining internal labels of σ.

Now we describe how to obtain a drawing for a shape cycle with a canonical subsequence τ. We assume, in accordance with Lemma 1, that τ has no adjacent, oppositely directed labels. Removal of τ from σ determines six *connecting* shape paths (some may be empty).

To each of these connecting paths, add back on the two elements of τ that bound it. Unless this path consists of just the two elements of τ, it must contain at least two flats. Otherwise, the two elements of τ, which are not adjacent in σ, would lie in the same flat of σ, contradicting the fact that τ is canonical.

The connecting parts of the doubly extensible drawings will be placed in separate octants. The segments of τ are precisely the end segments of these six doubly extensible drawings and can be drawn arbitrarily long. Their lengths will be chosen so long that they can connect the internal parts of the doubly extensible drawings isolated in distinct octants.

Make the drawings (and hence their bounding boxes) for the six connecting paths above. Some of these may be just points. Relative to a local origin of each drawing, we know the coordinates of all endpoints of the segments in that part of the drawing.

Now we determine lengths for the canonical segments and position the origins of the bounding boxes.

Look at the shape of τ. Since it is a strong canonical sequence, there are just two possibilities, the chair shape (e.g., $UNEDSW$) or the skew shape (e.g., $UNDESW$). Use this to determine octants for the placement of the bounding boxes containing the connecting drawings. Note that no two boxes are assigned to the same octant.

Let $l_E, l_N, \ldots$ denote the unknown lengths to be assigned to the canonical segments. A simple system of equations and inequalities must be satisfied by the unknowns for each oppositely directed pair of canonical segments: the total length of all segments directed E, say, must equal the total length of all segments directed W, and similarly for the other pairs. This will guarantee that the cycle closes.

Note that we have already determined the lengths of the segments that are not canonical, as well as the location of the endpoints of the canonical segments, in terms of the local coordinates of the boxes. Hence it easy to determine, for each local origin of a box, a system of three inequalities, one for each of the three orthogonal directions, that guarantees that the box stays strictly inside its assigned octant.

Satisfying these systems of inequalities implies that a corresponding system of inequalities on the lengths of the canonical segments must also be satisfied. This gives a lower bound on the length of each canonical segment of the form $l_E \geq c_E$ for some constant c_E, and so on.

To ensure that the cycle will close, we add to the system of inequalities on lengths three equations, one for each pair of opposite directions, as follows. The total length of segments directed E must equal the total length of segments directed W, and similarly for the other two pairs. The form for the E, W equation is either $l_E = l_W + c_{EW}$ or $l_W = l_E + c_{EW}$ for some positive constant c_{EW}, and similarly for the other two pairs.

Consider the constraints on the lengths of a particular oppositely directed pair, say on l_E and l_W. These constraints are

– for non-negative constant c_{EW} and for $l_E \geq l_W$, we have $l_E = l_W + c_{EW}$ (or if $l_E < l_W$, then $l_W = l_E + c_{EW}$;
– $l_E \geq c_E$;
– $l_W \geq c_W$.

These may be satisfied by assigning the value

$$l_W = max(c_W, c_E - c_{EW}) \text{ (or } l_E = max(c_E, c_W - c_{EW}) \text{ in case } l_E < l_W).$$

This determines the value of the length of the canonical segments directed E and W. The remaining lengths for the other directions may be determined similarly.

The lengths have now been chosen so that the path forms a closed cycle. To see that the cycle is simple, note that clearly, segments that are not canonical do not intersect each other. Hence it suffices to check that no canonical segment intersects another canonical segment or a non-canonical one (including ones in boxes not located at the endpoints of the canonical segment). This follows easily from the fact that the bounding boxes for the connecting paths are located in distinct octants.

4.3 Algorithmic Issues

To obtain a linear time algorithm for testing for the condition one must search for and produce, if one exists, a canonical sequence of length six in linear time. To do this, find, in linear time, a pair of parallel flats in σ. The proof of the necessity of the condition (see Section 5 for a sketch), reveals that if σ satisfies the condition, then it must contain one of a constant number of canonical sequences of special types defined by the relation of the labels in the canonical sequence to each other and to the two given parallel flats. Even though σ is a circular sequence, the fact that the pair of parallel flats can be chosen arbitrarily gives a starting label for σ, namely, the first label of one of these flats. Hence, it is not necessary to try each label of the entire sequence σ as a starting place when searching for a canonical sequence of one of the special types. Consequently, a linear time algorithm can be designed to check for the presence of one of these special canonical sequences. Given a strong canonical cycle, a simple orthogonal drawing for it can be constructed as described in the previous subsection. The computation of the coordinates of the endpoints of the segments of a drawing requires $O(n)$ time for the real RAM model of computation. Since the lengths of some segments might require $\Theta(lgn)$ bits to record, the running time becomes $O(nlgn)$ for a Turing machine model.

5 Necessity

Given a simple orthogonal drawing $\Gamma(\sigma)$ of a shape cycle σ, our goal is to show that σ contains a canonical sequence of length six. By slightly perturbing $\Gamma(\sigma)$ if necessary, we may assume without loss of generality that $\Gamma(\sigma)$ satisfies a *general position assumption*, namely, that no two vertices belonging to distinct flats of σ are drawn on the same axis-aligned plane. The lemmas and theorems that follow are based on this assumption.

5.1 The Proof Technique

The proof is based on the idea of cutting $\Gamma(\sigma)$ into two paths such that one reaches an octant and the other one goes back to the origin. We follow the two paths and look for canonical sequences on each path.

As mentioned in Section 2, a proof based on this approach does *not* follow easily from Theorem 2. It requires more elaborate machinery:

- We suitably choose the points a and b where we cut $\Gamma(\sigma)$ in order to define the two shape paths σ_{ab} and σ_{ba}.
- We find canonical sequences for $\sigma_{ab}\ell_{bb'}$ and $\sigma_{ba}\ell_{aa'}$ that may consist of three or four labels; $\ell_{bb'}$ is the label of σ after the last label of σ_{ab}, and $\ell_{aa'}$ is the label of σ after the last label of σ_{ba}.
- We use a certain necessary condition, together with Theorem 2 and the properties and lemmas of Subsection 5.2 to construct a canonical sequence of length six for σ.

5.2 Some Useful Properties

We now observe some basic properties of a canonical sequence τ of a shape cycle or path σ. These properties are useful and easy to prove. Unless specified otherwise, in this subsection σ is understood to denote either a shape path or a shape cycle. By the union $\tau_1 \cup \tau_2$ of two subsequences of σ, we mean the sequence whose elements are the elements of τ_1 and τ_2, ordered as in σ.

Property 1. If τ contains three labels XYZ such that they are consecutive on the same flat of σ, then X and Z define opposite directions.

Property 2. Let σ be a shape path. If we remove from τ its first or last label, the resulting sequence is still canonical for σ.

Property 3. If τ consists of three labels that define mutually orthogonal directions, then any subsequence of τ is a canonical sequence of σ .

The next property allows us to remove a label from a canonical sequence of length four; its proof is an immediate consequence of the definition of τ and of Property 1.

Property 4. If τ consists of four labels exactly two of which are oppositely directed, then a subsequence obtained by deleting from τ one of these two opposite labels is a canonical sequence of σ.

For example, if $\tau = \{\bar{D}, \bar{S}, \bar{W}, \bar{U}\}$ is a canonical sequence, then by Property 4, $\tau' = \{\bar{D}, \bar{S}, \bar{W}\}$ is also a canonical sequence.

The following lemmas allow us to merge two canonical sequences to obtain a new canonical sequence.

Lemma 3. *Let $\tau_1 \subset \sigma$ and $\tau_2 \subset \sigma$ be two canonical sequences such that (1) $\tau_1 \cap \tau_2 = \emptyset$, and (2) for all pairs of canonical labels X, Y such that $X \in \tau_1$ and $Y \in \tau_2$ there is no flat containing both X and Y. Then the sequence $\tau = \tau_1 \cup \tau_2$ is canonical for σ.*

Lemma 4. *Let σ have the form $\sigma = \sigma_2 X \sigma_1$ where X is a transition label for σ. Let $\tau_1 \subseteq X\sigma_1$ and $\tau_2 \subseteq \sigma_2 X$ be canonical sequences for σ such that (1) $\tau_1 \cap \tau_2 = X$ and (2) for all pairs of canonical labels $Y, Z \neq X$ such that $Y \in \tau_1$ and $Z \in \tau_2$, there is no flat containing both Y and Z. Then the sequence $\tau = \tau_1 \cup \tau_2$ is canonical for σ.*

Lemma 5. *Let σ be a shape cycle of the form $X\sigma_1 Y \sigma_2$, where X and Y are transition labels, and let τ_1 and τ_2 be canonical sequences for σ such that (1) $\tau_1 \subseteq X\sigma_1 Y$ and $\tau_2 \subseteq Y\sigma_2 X$, and (2) $\tau_1 \cap \tau_2 = X, Y$. Then $\tau_1 \cup \tau_2$ is a canonical sequence for σ.*

Next is a necessity result for $3D$ shape paths.

Lemma 6. *Let $\Gamma(\sigma)$ be a simple drawing of a shape path σ starting at the origin, and let uv be an edge of $\Gamma(\sigma)$. If u is in the DSW octant and v is in the DSE octant, then σ contains a canonical sequence $\tau = \{D, S, W, E\}$.*

5.3 Proof of Necessity

We sketch the proof in the case that shape cycle σ has at least four flats. Straightforward case analysis handles shape cycles with fewer than four flats.

Under the general position assumption, if σ has at least four flats, then it always has two flats F_a and F_b such that $\Gamma(F_a)$ and $\Gamma(F_b)$ lie on parallel planes. Let a be the starting point of F_a and let b be the starting point of F_b. Let aa' be the first edge of $\Gamma(F_a)$, and let bb' be the first edge of $\Gamma(F_b)$. Let $\ell_{aa'}$ be the direction label for aa' and let $\ell_{bb'}$ be the direction label for bb'. Observe that $\ell_{aa'}$ is a transition label shared by two flats of σ, the flat preceding F_a and flat F_a. Similarly, $\ell_{bb'}$ is a transition label shared by the flat preceding F_b and flat F_b. We denote with F_{a-} and F_{b-} the flats preceding F_a and F_b, respectively.

Observe that if the origin is chosen at a, then b is a point of an octant. We define two disjoint directed paths: $\Gamma(\sigma_{a'b})$ is the path from a' to b and $\Gamma(\sigma_{b'a})$ is the path from b' to a. We therefore have that $\sigma = \ell_{aa'}\sigma_{a'b}\ell_{bb'}\sigma_{b'a}$. We also have $\sigma_{ab} = \ell_{aa'}\sigma_{a'b}$ and $\sigma_{ba} = \ell_{bb'}\sigma_{b'a}$.

Suppose we locate the origin at a and let XYZ be the octant containing b. We say that a and a' are *equivalent with respect to* b if moving the origin from a to a' leaves b in the XYZ octant. A similar definition can be given for the relationship of b and b' with respect to a. Observe that if a and a' are not equivalent with respect to b, then when we locate the origin at b, we have that a' does not lie in the octant that contains a.

We consider four main cases, determined by whether or not a and a' are equivalent with respect to b, and by whether or not b and b' are equivalent with respect to a. For each case, we show how to choose a canonical sequence of length six for σ. This is done by using Properties 1, 3, 4, and Lemmas 3, 4, and 5 to perform merging operations on two canonical sequences, one defined in $\sigma_{ab}\ell_{bb'}$ and the other defined in $\sigma_{ba}\ell_{aa'}$. The canonical sequence of $\sigma_{ab}\ell_{bb'}$ ($\sigma_{ba}\ell_{aa'}$) can either consist of three labels if b and b' are equivalent with respect to a (a and a' are equivalent with respect to b), in which case Theorem 2 is used to define the canonical sequence; or it can consist of four labels if b and b' are not equivalent with respect to a (a and a' are not equivalent with respect to b), in which case Lemma 6 is used to define the canonical sequence. Since paths $\sigma_{ab}\ell_{bb'}$ and $\sigma_{ba}\ell_{aa'}$ are not disjoint, their canonical sequences need not be disjoint. However, the merging operations performed on these canonical sequences for paths produce a cyclic canonical sequence of length six for cycle σ.

We summarize the results of this section with the following theorem.

Theorem 3. *Let $\Gamma(\sigma)$ be a simple orthogonal drawing of a shape cycle σ. Then σ contains a canonical sequence of length six.*

6 Conclusion

This paper has characterized those shape cycles that admit a simple orthogonal drawing in $3D$. The characterization yields a linear time recognition algorithm, and a drawing algorithm that is linear in the real RAM model and $O(nlgn)$ in the Turing machine model. Interesting related problems that remain include: (1) characterizing simple shapes for graphs that are not just cycles, (2) minimizing the volume of bounding boxes of shape cycles that must be drawn with vertices at grid points (the coordinates of our drawing will be rational and can be scaled up to be integers; however, we have not attempted to minimize the volume of the drawing), and (3) extending the characterization of this paper to shape cycles with more than six directions and/or to dimension higher than three.

References

1. T. C. Biedl. Heuristics for *3d*-orthogonal graph drawings. In *Proc. 4th Twente Workshop on Graphs and Combinatorial Optimization*, pp. 41–44, 1995.
2. T. Biedl, T. Shermer, S. Wismath, and S. Whitesides. Orthogonal 3-D graph drawing. *J. Graph Algorithms and Applications*, 3(4):63–79, 1999.
3. R. F. Cohen, P. Eades, T. Lin and F. Ruskey. Three-dimensional graph drawing. *Algorithmica* , 17(2):199–208, 1997.
4. G. Di Battista, P. Eades, R. Tamassia, and I. Tollis. Graph Drawing. Prentice Hall, 1999.
5. G. Di Battista, G. Liotta, A. Lubiw, and S. Whitesides. Embedding problems for paths with direction constrained edges. In D.-Z. Du, P. Eades, V. Estivill-Castro, X. Lin, and A. Sharma, eds., *Computing and Combinatorics,* 6^{th} *Ann. Int. Conf., COCOON 2000*, Springer-Verlag LNCS vol. 1858, pp. 64-73, 2000.
6. G. Di Battista and L. Vismara. Angles of planar triangular graphs. *SIAM J. Discrete Math.*, 9(3):349–359, 1996.
7. P. Eades, C. Stirk, and S. Whitesides. The techniques of Komolgorov and Bardzin for three dimensional orthogonal graph drawings. *Inform. Process. Lett.*, 60:97–103, 1996.
8. P. Eades, A. Symvonis, and S. Whitesides. Three-dimensional orthogonal graph drawing algorithms. *Discrete Applied Math.*, vol. 103, pp. 55-87, 2000.
9. A. Garg. New results on drawing angle graphs. *Comput. Geom. Theory Appl.*, 9(1–2):43–82, 1998. Special Issue on Geometric Representations of Graphs, G. Di Battista and R. Tamassia, *eds.*.
10. A. Papakostas and I. Tollis. Algorithms for incremental orthogonal graph drawing in three dimensions. *J. Graph Algorithms and Appl.*, 3(4):81-115, 1999.
11. R. Tamassia. On embedding a graph in the grid with the minimum number of bends. *SIAM J. Comput.*, 16(3):421–444, 1987.
12. G. Vijayan and A. Wigderson. Rectilinear graphs and their embeddings. *SIAM J. Comput.*, 14:355–372, 1985.
13. V. Vijayan. Geometry of planar graphs with angles. In *Proc. 2nd Annu. ACM Sympos. Comput. Geom.*, pp. 116–124, 1986.
14. D. R. Wood. Two-bend three-dimensional orthogonal grid drawing of maximum degree five graphs. TR 98/03, School of Computer Science and Software Engineering, Monash University, 1998.
15. D. R. Wood. An algorithm for three-dimensional orthogonal graph drawing. In S. Whitesides, ed., *Graph Drawing (*6^{th} *Int. Symp., GD '98)*, Springer-Verlag, LNCS vol. 1547, pp. 332-346, 1998.
16. D. R. Wood. Three-Dimensional Orthogonal Graph Drawing. Ph.D. thesis, School of Computer Science and Software Engineering, Monash University, 2000.