Journal of Graph Algorithms and Applications
http://jgaa.info/ vol. 26, no. 3, pp. 307–334 (2022)
DOI: 10.7155/jgaa.00596

Planar L-Drawings of Bimodal Graphs

Patrizio Angelini[1] *Steven Chaplick*[2]
Sabine Cornelsen[3] *Giordano Da Lozzo*[4]

[1]John Cabot University, Rome, Italy
[2]Maastricht University, The Netherlands
[3]University of Konstanz, Germany
[4]Roma Tre University, Rome, Italy

Submitted:	October 2020	Reviewed:	May 2022	Revised:	May 2022
Accepted:	May 2022	Final:	June 2022	Published:	June 2022
Article type:	Regular paper	Communicated by:	D. Auber, P. Valtr		

Abstract. In a *planar L-drawing* of a directed graph (digraph) each edge e is represented as a polyline composed of a vertical segment starting at the tail of e and a horizontal segment ending at the head of e. Distinct edges may overlap, but not cross. Our main focus is on *bimodal graphs*, i.e., digraphs admitting a planar embedding in which the incoming and outgoing edges around each vertex are contiguous. We show that every plane bimodal graph without 2-cycles admits a planar L-drawing. This includes the class of upward-plane graphs. Bimodal graphs with 2-cycles admit a planar L-drawing if the underlying undirected graph with merged 2-cycles is a planar 3-tree. Finally, outerplanar digraphs admit a planar L-drawing – although they do not always have a bimodal embedding – but not necessarily with an outerplanar embedding.

1 Introduction

In an *L-drawing* of a directed graph (digraph), vertices are represented by points with distinct x- and y-coordinates, and each directed edge (u, v) is a polyline consisting of a vertical segment incident to the tail u and of a horizontal segment incident to the head v. This defines four *ports*, North, South, East, and West for each vertex. Two edges may overlap in a subsegment with end point at a common tail or head. An L-drawing is *planar* if no two edges cross (Fig. 1(c)). Non-planar L-drawings were first defined by Angelini et al. [3]. Chaplick et al. [13] showed that it is

Special Issue on the 28th Int. Symposium on Graph Drawing and Network Visualization, GD 2020

The work of Sabine Cornelsen was funded by the German Research Foundation DFG – Project-ID 50974019 – TRR 161 (B06). The work of Giordano Da Lozzo was partially supported by MIUR grant 20174LF3T8 *"AHeAD: efficient Algorithms for HArnessing networked Data"*.

E-mail addresses: pangelini@johncabot.edu (Patrizio Angelini) s.chaplick@maastrichtuniversity.nl (Steven Chaplick) sabine.cornelsen@uni-konstanz.de (Sabine Cornelsen) giordano.dalozzo@uniroma3.it (Giordano Da Lozzo)

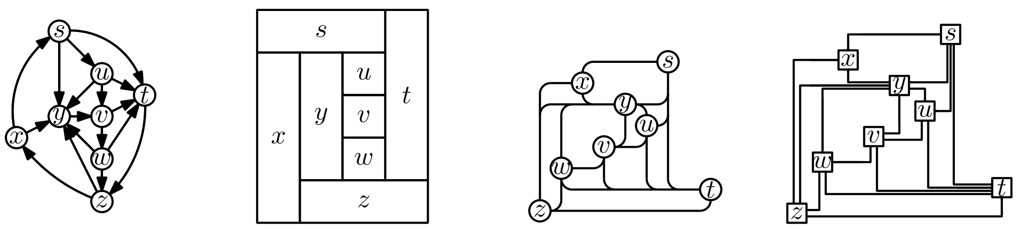

(a) Bimodal (b) Rectangular Dual (c) L-drawing (d) Kandinsky

Figure 1: Various representations of a bimodal irreducible triangulation.

NP-complete to decide whether a directed graph has a planar L-drawing if the embedding is not fixed. However it can be decided in linear time whether a planar st-graph has an *upward-planar L-drawing*, i.e., an L-drawing in which the vertical segment of each edge leaves its tail from the top.

A vertex v of a plane digraph G is *k-modal* ($\text{mod}(v) = k$) if in the cyclic sequence of edges around v there are exactly k pairs of consecutive edges that are neither both incoming nor both outgoing. A digraph G is *k-modal* if $\text{mod}(v) \leq k$ for every vertex v of G. The 2-modal graphs or 2-modal vertices are often referred to as *bimodal*, see Fig. 1(a). Any plane digraph admitting a planar L-drawing is clearly 4-modal. Upward-planar and level-planar drawings induce bimodal embeddings. While testing whether a graph has a bimodal embedding is possible in linear time, testing whether a graph has a 4-modal embedding [6] and testing whether a partial orientation of a plane graph can be extended to be bimodal [10] are NP-complete.

A *plane digraph* is a planar digraph with a fixed rotation system of the edges around each vertex and a fixed outer face. In an L-drawing of a plane digraph G, the clockwise cyclic order of the edges incident to each vertex and the outer face are the ones prescribed for G. In a planar L-drawing the edges attached to the same port of a vertex v are ordered as follows: There are first the edges bending to the left with increasing length of the segment incident to v and then those bending to the right with decreasing length of the segment incident to v.

This is analogous to the Kandinsky model [16], where vertices are drawn as squares of equal size on a grid and edges as orthogonal polylines on a finer grid (Fig. 1(d)). Bend-minimization in the Kandinsky model is NP-complete [11] even if the embedding is fixed and can be approximated within a factor of two [5]. Each undirected simple graph admits a Kandinsky drawing with one bend per edge [12]. The relationship between Kandinsky drawings and planar L-drawings was established in [13].

L-drawings of directed graphs can be considered as bend-optimal drawings, since one bend per edge is necessary in order to guarantee the property that edges must leave a vertex from the top or the bottom and enter it from the right or the left. Planar L-drawings can be also seen as a directed version of +-contact representations, where each vertex is drawn as a + and two vertices are adjacent if the respective +es touch. If the graph is bimodal, then the +es are ⊤s (including ⊢, ⊥, and ⊣). Undirected planar graphs always allow a ⊤-contact representation, which can be computed utilizing Schnyder woods [14].

Biedl and Mondal [9] showed that a +-contact representation can also be constructed from a rectangular dual (Fig. 1(b)). A plane graph with four vertices on the outer face has a rectangular dual if and only if it is an inner triangulation without separating triangles [21]. Bhasker and Sahni [7] gave the first linear-time algorithm for computing rectangular duals. Then, He [17] showed how to compute a rectangular dual from a regular edge labeling, while Kant and He [20]

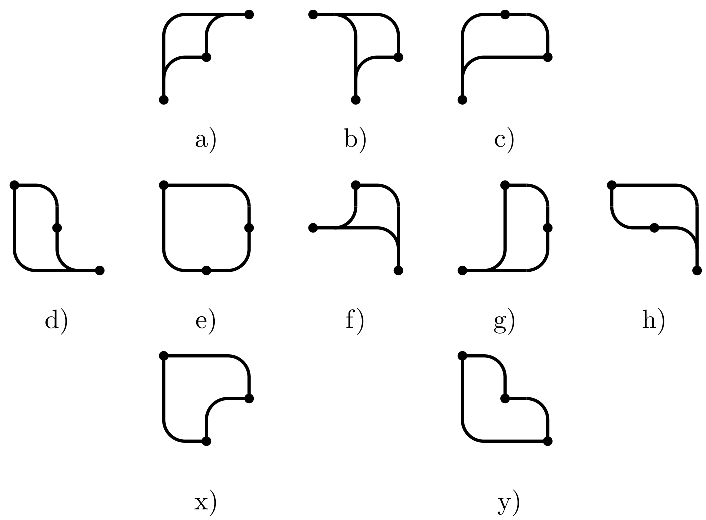

Figure 2: All planar L-drawings of a triangle up to symmetry.

gave two linear-time algorithms for computing regular edge labelings. Biedl and Derka [8] computed rectangular duals via (3,1)-canonical orderings.

Contribution: We show that every bimodal graph without 2-cycles admits a planar L-drawing respecting a given bimodal embedding (Sect. 4). This implies that every upward-planar graph admits a planar L-drawing respecting a given upward-planar embedding, which solves an open problem posed in [13]. The construction is based on rectangular duals. We also provide an iterative algorithm showing that any bimodal graph with 2-cycles admits a planar L-drawing if the underlying undirected graph with merged 2-cycles is a planar 3-tree (Sect. 5). Finally, we show that every outerplanar graph admits a planar L-drawing but not necessarily one where all vertices are incident to the outer face (Sect. 6). We conclude with open problems (Sect. 7).

2 Preliminaries

In this section, we provide further definitions and preliminary results.

2.1 L-Drawings

For each vertex we consider four *ports*, North, South, East, and West. An L-drawing implies a *port assignment*, i.e., an assignment of the edges to the ports of the end vertices such that the outgoing edges are assigned to the North and South port and the incoming edges are assigned to the East and West port. A port assignment for each edge e of a digraph G defines a pair $(\mathrm{out}(e),\mathrm{in}(e)) \in \{\mathrm{North},\mathrm{South}\} \times \{\mathrm{East},\mathrm{West}\}$. An L-drawing *realizes* a port assignment if each edge $e = (v, w)$ is incident to the out(e)-port of v and to the in(e)-port of w. A port assignment *admits* a planar L-drawing if there is a planar L-drawing that realizes it. Given a port assignment it can be tested in linear time whether it admits a planar L-drawing [13].

In this paper, we will distinguish between given L-drawings of a triangle.

Lemma 1 *Fig. 2 shows all planar L-drawings of a triangle up to symmetry.*

Proof: We first consider a triangle T that is not a directed cycle. Let t be the sink, s the source, and w the third vertex of T. We may assume that $\langle s, w, t \rangle$ is the counter-clockwise cyclic order of vertices around T and that the edge (s, w) uses the North port of s – the other cases being symmetric. We distinguish two cases.

1. w is to the right of s. In this case t cannot be below w: otherwise it is not possible to close T in counter-clockwise direction with only one bend per edge. For the x-coordinate of t there are three possibilities: t is to the right of w (a), between s and w (c), and to the left of s (b).

2. w is to the left of s. As in the former case, t cannot be above w in this case. If t is below s, its x-coordinate can be to the left of w (g), between w and s (e), or to the right of s (d). If t is vertically between w and s, its x-coordinate is between w and s (h) or to the left of w (f).

Consider now the case that triangle T is a directed cycle. Then no two edges of T can use the same port at any vertex. Thus, T is drawn as a 6-gon, which implies that the angular sum is 4π. Therefore, there is one $3\pi/2$ angle and five $\pi/2$ angles. The constructed drawings cover the cases in which the $3\pi/2$ angle is at a vertex (y) or at a bend (x), thus concluding the proof. □

Coordinates for the Vertices. Given a port assignment that admits a planar L-drawing, a planar L-drawing realizing it can be computed in linear time by the general compaction approach for orthogonal or Kandinsky drawings [15]. However, in this approach, the graph has to be first augmented such that each face has a rectangular shape. For L-drawings of plane triangulations it suffices to make sure that each edge has the right shape given by the port assignment. We will prove in Theorem 1 that this can be achieved using topological orderings only. For this, we first present an auxiliary lemma and give a further definition.

Lemma 2 *A drawing of a plane biconnected graph is planar if the ordering of the edges around each vertex is respected and the boundary of each face is drawn crossing-free.*

Proof: We exploit the same idea as in the proof of Tutte – see Item 9.1 on page 758 of [23]. Let G be a plane biconnected graph and let Γ be a drawing of G in which each face is drawn crossing-free. Assume two edges e_1 and e_2 cross. By assumption, e_1 and e_2 must belong to distinct faces, say f_1 and f_2, respectively. (We consider the boundary of a face not to be part of the face, i.e., faces are open regions). Then there must be a point $q \in f_1 \cap f_2$. But this is impossible: For each point p in the plane let $\delta(p)$ be the number of faces containing p. Since the inner faces are bounded, there must be a point q_o that is only contained in the outer face and thus, $\delta(q_o) = 1$. Consider a curve ℓ from q_o to q that does not contain vertices. Traversing ℓ, the count δ does not change if no edge is crossed. If we cross an edge then we leave a face and enter another face. Thus, the count will not change either, contradicting $\delta(q) \geq 2$. □

Consider two edges e_1 and e_2 in a planar L-drawing that are assigned to the same port of a vertex v, and let f be the face between e_1 and e_2. Then the bend at e_1 or e_2 must be concave in f. In this case, we say that e_1 and e_2 satisfy the *concave-bend condition*. A reader familiar with Kandinsky drawings [16] will notice that this condition is inspired by the so called bend-or-end property.

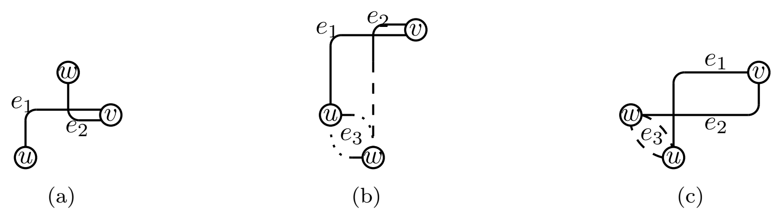

Figure 3: A triangular face cannot self-intersect if all edges have the correct shape.

Theorem 1 *Let $G = (V, E)$ be a plane triangulated graph with a port assignment that admits a planar L-drawing. Let X and Y be copies of G with a different orientation of the edges. More precisely, for each edge $e = (v, w) \in E$*

- *there is (v, w) in X if in(e) = West and (w, v) in X if in(e) = East.*
- *there is (v, w) in Y if out(e) = North and (w, v) in Y if out(e) = South.*

Let x and y be a topological ordering of X and Y, respectively. Drawing each vertex v at $(x(v), y(v))$ yields a planar L-drawing realizing the given port assignment.

Proof: By Lemma 2, it suffices to show that all faces are drawn in a planar way. Assume there are two edges e_1 and e_2 incident to the same face that cross. Since all faces are triangles, e_1 and e_2 are incident to the same vertex v. Assume without loss of generality that out(e_1) = North, in(e_1) = West, and that v is the head of e_1. We distinguish three ways e_2 could cross e_1 (see Fig. 3):

(a) v is the head of e_2, in(e_2) = West, out(e_2) = South, and e_2 is before e_1 in the clockwise order around v, or

(b) v is the head of e_2, in(e_2) = West, out(e_2) = North, and e_2 is after e_1 in the clockwise order around v, or

(c) v is the tail of e_2, out(e_2) = South, in(e_2) = East, the head w of e_2 is to the left, and above of the tail u of e_1.

Situation (a) violates the concave-bend condition. A crossing in Situation (b) implies that the edge e_3 closing the triangular face is either (u, w) with in(e_3) = West or (w, u) with in(e_3) = East. However, this port assignment would not be one that admits a planar L-drawing. Finally, a crossing in Situation (c) implies that the edge e_3 closing the triangular face is either (u, w) with in(e_3) = East and out(e_3) = North or (w, u) with in(e_3) = West and out(e_3) = South. Again, this port assignment would not be one that admits a planar L-drawing. □

Observe that in general it does not suffice to only consider a drawing respecting the port assignment of each edge in order to obtain a planar L-drawing even if the port assignment admits such a drawing, not even for a directed 4-cycle (Fig. 4(a)) or a 4-cycle consisting only of sinks and sources (Fig. 4(b)).

Observe that we can modify the edge lengths in a planar L-drawing independently in x- and y-directions in an arbitrary way, as long as we maintain the ordering of the vertices in x- and y-direction, respectively. This will still yield a planar L-drawing. This fact implies the following remark.

Remark 1 *Let G be a plane digraph with a triangular outer face, let Γ be a planar L-drawing of G, and let Γ_0 be a planar L-drawing of the outer face of G such that the edges on the outer face have the same port assignment in Γ and Γ_0. Then there exists a planar L-drawing of G with the same port assignment as in Γ in which the drawing of the outer face is Γ_0.*

2.2 Bimodal Graphs

We provide a small technical lemma stating that any given k-modal plane digraph can be "nicely" extended to a quadrangulated k-modal plane digraph. This is applied in the proof of Theorem 2 where we show that every plane bimodal graph without 2-cycles admits a planar L-drawing. A vertex v is a *sink-switch* or *source-switch* of a simple cycle C in a graph if and only if v is a sink or a source in C.

Lemma 3 *Let G be a plane digraph. We can augment G by adding edges to obtain a biconnected plane digraph with face-degree at most four such that each k-modal vertex remains k-modal if $k > 0$ and each 0-modal vertex becomes at most 2-modal. The construction can be done in linear time and introduces neither parallel edges nor 2-cycles. Moreover, each 4-cycle bounding a face consists of alternating source and sink switches.*

Proof: We perform the augmentation of G in three steps.

Making G connected: If G is not connected, let G_1 be a connected component of G such that the outer face f_o^1 of G_1 (considered as an open region) contains vertices of G. See Fig. 5(a). Pick a vertex v incident to f_o^1 with the property that v is an isolated vertex or v is the tail of an edge incident to f_o^1. Let f be the face of G that is contained in f_o^1 and that is incident to v. Pick a vertex w incident to f that is not in G_1 and such that w is isolated or the head of an edge incident to f. Add the edge (v, w).

Making G biconnected: If G contains a cut vertex v, let w_1 and w_2 be two consecutive neighbors of v in different biconnected components. Let f be the face between the edges connecting v to w_1 and w_2, respectively; see Fig. 5(b). If we can add the edge (w_1, w_2) or (w_2, w_1) such that the modalities of w_1 and w_2 do not increase, if they had been positive before, we do so. Otherwise the edges incident to f and w_1 and w_2 are either all outgoing edges of w_1 and w_2, respectively, or all incoming edges – we assume without loss of generality they are outgoing. Further, both w_1 and w_2 have at least two neighbors incident to f. Let $w_3 \neq v$ be a neighbor of w_2 such that the edge $\{w_2, w_3\}$ is incident to f. Add (w_1, w_3) to G.

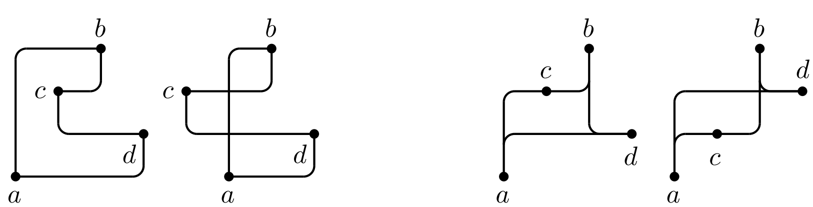

(a) directed 4-cycle (b) 4-cycle

Figure 4: Given a port assignment that admits a planar L-drawing, not every L-drawing that realizes it is also planar.

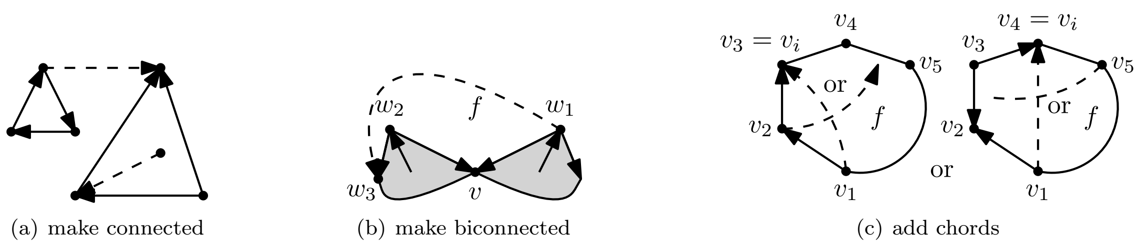

Figure 5: How to quadrangulate a digraph maintaining its modality

Making G quadrangulated: If G contains a face f of degree greater than four, let $v_1, \dots, v_k$ be the facial cycle of f such that v_1 is the tail of an edge incident to f; see Fig. 5(c). If the edge between v_3 and v_2 and the edge between v_3 and v_4 incident to f are both outgoing edges of v_3, let $i = 4$. Otherwise let $i = 3$. Now, if neither (v_1, v_i) nor (v_i, v_1) is present in G then add (v_1, v_i) to G. Otherwise we can add an edge between v_2 and either v_4 or v_5 if $i = 3$, or between v_5 and either v_3 or v_2 if $i = 4$.

Similarly to triangulating plane graphs (see, e.g., [18, Chapter 6]), this construction can be done in linear time. □

2.3 Irreducible Triangulations and Rectangular Duals

An *irreducible triangulation* is an internally triangulated graph without separating triangles, where the outer face has degree four (Fig. 1). A *rectangular tiling* of a rectangle R is a partition of R into a set of non-overlapping rectangles such that no four rectangles meet at the same point. A *rectangular dual* of a planar graph is a rectangular tiling such that there is a one-to-one correspondence between the inner rectangles and the vertices and there is an edge between two vertices if and only if the respective rectangles touch. We denote by R_v the rectangle representing the vertex v. Note that an irreducible triangulation always admits a rectangular dual, which can be computed in linear time [7, 8, 17, 20].

3 Perturbed Generalized L-Drawings

Motivated by the work of Biedl and Mondal [9], we will use rectangular duals in order to construct planar L-drawings. To this end, in Sect. 3.1, we introduce generalized L-drawings where segments are replaced by zig-zag-paths. We then essentially perturb the coordinate system such that the axes are the diagonals of the rectangles in the rectangular dual, i.e., we route the edges, such that each of their segments is parallel to these diagonals. See Sect. 3.2.

3.1 Generalized Planar L-Drawings

For two points p, q in the plane, we denote the line-segment $\{\lambda p + (1-\lambda)q;\ 0 \leq \lambda \leq 1\}$ connecting them by $\overline{pq}$. An *orthogonal polyline* $P = \langle p_1, \dots, p_n \rangle$ is a sequence of points s.t. $\overline{p_i p_{i+1}}$ is vertical or horizontal. For $1 \leq i \leq n-1$ and a point $p \in \overline{p_i p_{i+1}}$, the polyline $\langle p_1, \dots, p_i, p \rangle$ is a *prefix* of P and the polyline $\langle p, p_{i+1}, \dots, p_n \rangle$ is a *suffix* of P. Walking from p_1 to p_n, consider a *bend* p_i, $i = 2, \dots, n-1$. The rotation $\mathrm{rot}(p_i)$ is 1 if P has a left turn at p_i, -1 for a right turn, and 0 otherwise (when $\overline{p_{i-1}p_i}$ and $\overline{p_ip_{i+1}}$ are both vertical or horizontal). The *rotation* of P is $\operatorname{rot}(P) = \sum_{i=2}^{n-1} \operatorname{rot}(p_i)$.

In a *generalized planar L-drawing* of a digraph, vertices are still represented by points with distinct x- and y-coordinates, and edges by orthogonal polylines with the following three properties. (1) Each directed edge $e = (u, v)$ starts with a vertical segment incident to the tail u and ends with a horizontal segment incident to the head v. (2) The polylines representing two edges overlap in at most a common straight-line prefix or suffix, and they do not cross. This definition resembles the one of *angular drawings*, defined in the context of Windrose planarity [4], in which edges are represented by xy-monotone polylines with their endparts being straight-line segments with slope 1 or -1. However, in Windrose planarity the assignment of edges to ports is prescribed by the input.

In order to define the third property, let $\operatorname{init}(e)$ be the prefix of e overlapping with at least one other edge, let $\operatorname{final}(e)$ be the suffix of e overlapping with at least one other edge, and let $\operatorname{mid}(e)$ be the remaining individual part of e. Observe that the first and the last vertex of $\operatorname{init}(e)$, $\operatorname{final}(e)$, and $\operatorname{mid}(e)$ are end vertices of e, bends of e, or bends of some other edges. Now we define the third property: (3) For an edge e one of the following is true: (i) neither of the two end points of $\operatorname{mid}(e)$ is a bend of e and $\operatorname{rot}(e) = \pm 1$ or (ii) one of the two end points of $\operatorname{mid}(e)$, but not both, is a bend of e and $\operatorname{rot}(\operatorname{mid}(e)) = 0$. See Fig. 6. As a consequence of the flow model of Tamassia [22], we obtain the following lemma.

(a) not a generalized planar L-drawing (b) underlying orthogonal drawing

Figure 6: Cond. 3 of generalized planar L-drawings is fulfilled for all edges but for e_1 and e_2. The rotation of each edge is ± 1. However, $\operatorname{rot}(\operatorname{mid}(e_1)) = 2$ and both end vertices of $\operatorname{mid}(e_2)$ are bends of e_2.

Lemma 4 *A plane digraph admits a planar L-drawing if and only if it admits a generalized planar L-drawing with the same port assignment.*

Proof: A planar L-drawing is a planar generalized L-drawing. For the other direction, assume that a planar generalized L-drawing Γ of a digraph G is given. We consider Γ as an orthogonal drawing of a new graph G' – see Fig. 6(b). To this end, we replace every bend in Γ by a dummy vertex if it is contained in the polyline representing another edge. Consider now an edge of G that has more than one bend. This edge is decomposed in G' into an initial straight-line path P_{init}, an edge e_{mid}, and a final straight-line path P_{final}. The edge e_{mid} is represented by an orthogonal polyline with the same bends as $\operatorname{mid}(e)$. Using the flow model of Tamassia [22], we can remove all but $\operatorname{rot}(\operatorname{mid}(e))$ bends from e_{mid}. Now the cases are two: If none of the end vertices of e_{mid} is a bend of e then $\operatorname{rot}(e) = \pm 1$ and, thus, e ends up with exactly one bend. If the first or the last bend of e is an end vertex of e_{mid} then $\operatorname{rot}(e_{\text{mid}}) = 0$ and e_{mid} ends up with no bends. Thus, e has exactly one bend at exactly one end vertex of e_{mid}. □

3.2 Perturbed Generalized Planar L-drawing

Consider a rectangular dual for a directed irreducible triangulation G. We construct a drawing of G as follows. We place each vertex of G on the center of its rectangle. Each edge is routed as a *perturbed orthogonal polyline*, i.e., a polyline within the two rectangles corresponding to its two end vertices, such that each edge segment is parallel to one of the two diagonals of the rectangle containing it. See Fig. 7(a). This drawing is called a *perturbed generalized planar L-drawing* if and only if (1) each directed edge $e = (u, v)$ starts with a segment on the diagonal $\backslash_u$ of R_u from the upper left to the lower right corner and ends with a segment on the diagonal $/_v$ of R_v from the lower left to the upper right corner. Observe that a change of directions at the intersection of R_v and R_u is not considered a bend if the two incident segments in R_v and R_u are both parallel to $\backslash$ or to $/$. The definition of rotation and Conditions (2) and (3) are analogous to generalized planar L-drawings.

In a perturbed generalized planar L-drawing, the North port of a vertex is at the segment between the center and the upper left corner of the rectangle. The other ports are defined analogously. Since we can always approximate a segment with an orthogonal polyline (Figs. 7(b) to 7(d)), we obtain the following.

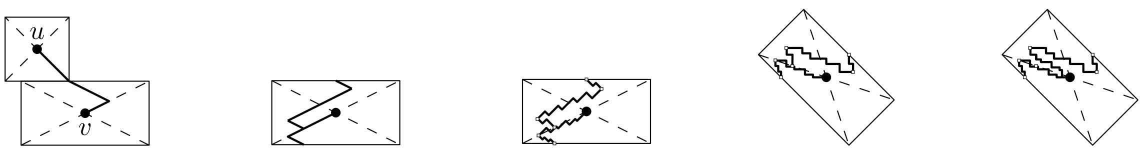

(a) (u, v) (b) perturbed (c) zig zag (d) rotate (e) split

Figure 7: (a) An edge in a perturbed generalized planar L-drawing. (b-e) From a perturbed generalized planar L-drawing to a generalized planar L-drawing.

Lemma 5 *If a directed irreducible triangulation has a perturbed generalized planar L-drawing, then it has a planar L-drawing with the same port assignment.*

Proof: Let a perturbed generalized planar L-drawing Γ of an irreducible triangulation G be given. By Lemma 4, it suffices to show that G has a generalized planar L-drawing. First, we construct a graph G' of maximum degree 4 by subdividing the edges at all bends and intersections with the boundary of a rectangle. Observe that now every edge lies in the interior of one rectangle. We approximate each edge e of G' with an orthogonal polyline rotated by 45°, arbitrarily close to it, in such a way that the following properties hold: No two polylines of two edges cross and the polyline of an edge is not self-intersecting. The rotation of any polyline is zero. The segments incident to the end vertices have both slope $-45°$ if e is parallel to $\backslash_v$ and slope 45° if e is parallel to $/_v$, where R_v is the rectangle containing e. See Fig. 7. Finally, rotate the drawing by 45° in clockwise direction. We obtain a drawing that fulfills all properties of a generalized planar L-drawing of G, except that edges might overlap in a prefix or a suffix with rotation 0 that might, however, not be a straight-line segment. We fix this as follows. Let v be a vertex and let $\langle e_1, \ldots, e_k \rangle$ be the sequence of edges assigned to a port of v in clockwise order. Assume without loss of generality that $e_1, \ldots, e_k$ are outgoing edges of v. Let $1 \leq m \leq k$ be such that $\text{init}(e_m)$ is longest. We redraw each edge e_i with $i \neq m$ so that the common part of e_i and e_m lies on the first segment and the lengths of $\text{init}(e_1), \ldots, \text{init}(e_{m-1})$ are still increasing and those of $\text{init}(e_{m+1}), \ldots, \text{init}(e_k)$ are still

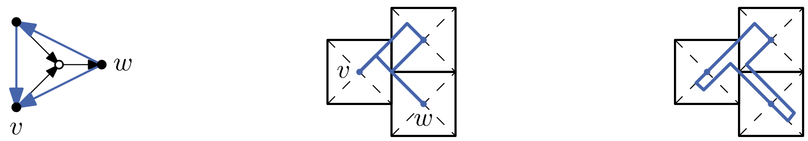

(a) pincers (b) bad pincers (c) alternative

Figure 8: (a) The (blue) outer edges incident to v and w, respectively, are pincers that are bad in the drawing of the (blue) outer triangle in (b) and not bad in (c).

decreasing. The rest of e_i is drawn arbitrarily close to e_m in such a way that the rotation of mid(e_i) is maintained. See Fig. 7(e). Observe that the bends of the original drawing still correspond to bends of the constructed drawing and have the same turns (left or right). Conditions (2) and (3) are still fulfilled. □

4 Bimodal Graphs without 2-Cycles

We study planar L-drawings of plane bimodal graphs. Our main contribution is to show that if the graph does not contain any 2-cycles, then it admits a planar L-drawing (Theorem 2). In Sect. 5, we also show that if there are 2-cycles, then there is a planar L-drawing if the underlying undirected graph after identifying parallel edges created by the 2-cycles is a planar 3-tree.

4.1 General Overview

Our approach is inspired by the work of Biedl and Mondal [9] to construct a +-contact representation for undirected graphs from a rectangular dual. We extend their technique in order to respect the given orientations of the edges.

The idea is to triangulate and decompose a given bimodal graph G into 4-connected components. Then, proceeding from the outermost to the innermost component, we construct a planar L-drawing of each component that respects a given shape of the outer face. This construction is based on rectangular duals and the corresponding perturbed generalized L-drawings.

In order to be able to reinsert the drawing of a component inside the respective triangular face in the drawing of its parent, we have to pay attention to the following situation. We call a pair of edges e_1, e_2 a *pincer* if e_1 and e_2 are on a triangle T, both are incoming or both outgoing edges of its common end vertex v (i.e., v is a sink- or a source switch of T), and there is another edge e of G incident to v in the interior of T but with the opposite direction. See Fig. 8. If the outer face of a 4-connected component contains a pincer, we have to make sure that e_1 and e_2 are not assigned to the same port of v in an ancestral component. In a partial perturbed (generalized) planar L-drawing of G, we call a pincer *bad* if e_1 and e_2 are assigned to the same port. Observe that in a bimodal graph, a pincer must be a source or a sink in an ancestral component. Moreover, in a 4-connected component at most one pair of incident edges of a vertex can be a pincer. The algorithm is described in detail in Sect. 4.2 and its correctness and run time are proved in Sect. 4.3. We summarize the result in the following theorem.

Theorem 2 *Every plane bimodal graph without 2-cycles admits a planar L-drawing. Moreover, such a drawing can be constructed in linear time.*

Theorem 2 yields the following implication, solving an open problem in [13].

Corollary 1 *Every upward-plane graph admits a planar L-drawing.*

4.2 Algorithm

In this subsection, we present the algorithm for the proof of Theorem 2. We start with a plane bimodal digraph D without 2-cycles. Triangulate D as follows: Add a new directed triangle in the outer face. Augment the graph by adding edges to obtain a plane bimodal graph in which each face has degree at most four as shown in Lemma 3. More precisely, now each non-triangular face is bounded by a 4-cycle consisting of alternating source and sink switches of the face. We finally insert a 4-modal vertex of degree 4 into each non-triangular face maintaining the 2-modality of the neighbors. Let G be the obtained triangulated graph without 2-cycles in which each vertex is 2-modal or an inner 4-modal vertex of degree four.

We construct a port assignment that admits a planar L-drawing of G as follows. Decompose G at separating triangles into 4-connected components. Proceeding from the outermost to the innermost components, we compute a port assignment for each 4-connected component H, avoiding bad pincers and such that the ports of the outer face of H are determined by the corresponding inner face of the parent component of H. We use an arbitrary shape for the outer face of G.

Let H be a 4-connected component of G and let Γ_0 be a planar L-drawing of the outer face of H without bad pincers of G. We now present an algorithm that constructs a planar L-drawing of H in which the drawing of the outer face is Γ_0 and no face contains bad pincers of G.

Port Assignment Algorithm

The aim of the algorithm is to compute a port assignment for the edges of H such that (i) there are no bad pincers and (ii) there exists a planar L-drawing realizing such an assignment. Note that the drawing Γ_0 already determines an assignment of the external edges to the ports of the external vertices. By Remark 1 any planar L-drawing with this given port assignment can be turned into one where the outer face has drawing Γ_0.

First, observe that H does not contain vertices on the outer face that are 4-modal in H: This is true since 4-modal vertices are inner vertices of degree four in the triangulated graph G and since G has no 2-cycles. Thus, a 4-modal vertex v has four distinct neighbors in G forming a cycle C enclosing v. If v was still 4-modal in H, then all four neighbors must still be present in H and, thus, v cannot be on the outer face of H.

Avoiding Bad Pincers. Next, we discuss the means that will allow us to avoid bad pincers. Let e_1 and e_2 be two edges with common end vertex v that are incident to an inner face f of H such that e_1, e_2 is a pincer of G. Note that the triangle bounding f is a separating triangle of G. We call f the *designated face* of v. Further, observe that v cannot be 4-modal in H, as otherwise it would be at least 6-modal in G. On the other hand, if v is 2-modal in H, then v was an inner 4-modal vertex of degree 4 in G, and e_1 and e_2 are two non-consecutive edges incident to v. It follows that H is a K_4 where the outer face is not a directed cycle; see Fig. 10(a). For any given drawing Γ_0 of the outer face (see Fig. 9 and Lemma 1 for the possible drawings of a triangle), the inner vertex can always be added such that no bad pincer is created. Therefore, in the following we can assume that v is 0-modal in H.

This implies that we only have to take care of pincers where the common end vertex is 0-modal in H. Since each 0-modal vertex was 2-modal in G, it has at most one designated face. In the

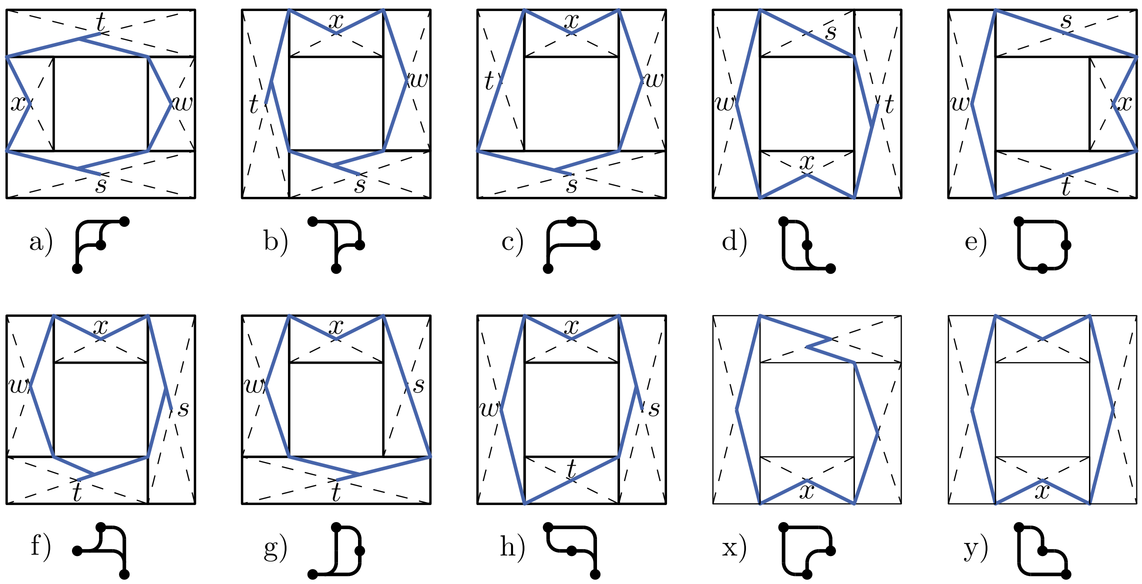

Figure 9: Realization in the rectangular dual for any kind of drawings of the outer face up to symmetries.

following, we assume that all 0-modal vertices are assigned a designated incident inner face where no 0° angle is allowed.

Constructing the Rectangular Dual. As an intermediate step towards a perturbed generalized planar L-drawing, we have to construct a rectangular dual of H. More precisely, we first augment H to an irreducible triangulation H_x for which a rectangular dual can be computed in linear time [7, 8, 17, 20]. In order to construct H_x, we proceed as follows. Let s, t, and w be the vertices on the outer face of H. Depending on the given drawing of the outer face, subdivide one of the edges of the outer face by a new vertex x according to the cases given in Fig. 9 – up to symmetries. Let f be the quadrangular inner face incident to x. Triangulate f by adding an edge e incident to x: Let y be the other end vertex of e. If y was 2-modal, we can orient e such that y is still 2-modal. If y was 0-modal and f was its designated face, then orient e such that y is now 2-modal. Otherwise, orient e such that y remains 0-modal. Observe that if y had degree 4 in the beginning, then it has degree 5 now.

The resulting graph H_x is triangulated, has no separating triangles, and its outer face is bounded by a quadrangle, hence it is an irreducible triangulation. Let R be a rectangular dual of H_x. Up to a possible rotation of a multiple of 90°, we can replace the four rectangles of R on the outer face with the configuration depicted in Fig. 9 that corresponds to the given drawing of the outer face. Let R_v be the rectangle of a vertex v in R.

Port Assignment. We now assign edges to the ports of the incident vertices. For the edges on the outer face the port assignment is given by Γ_0. Fig. 9 shows the assignments for the outer face.

Let v be a vertex of H_x. We define the *canonical assignment* of an edge incident to v to a port

(a) 2-modal pincer (b) virtual edges

Figure 10: (a) shows the only case (up to reversing directions) of a graph H in Sect. 4.2 with a pincer that is incident to a 2-modal vertex (the orientation of the undirected outer edge is irrelevant). (b) Avoiding bad pincers with virtual edges.

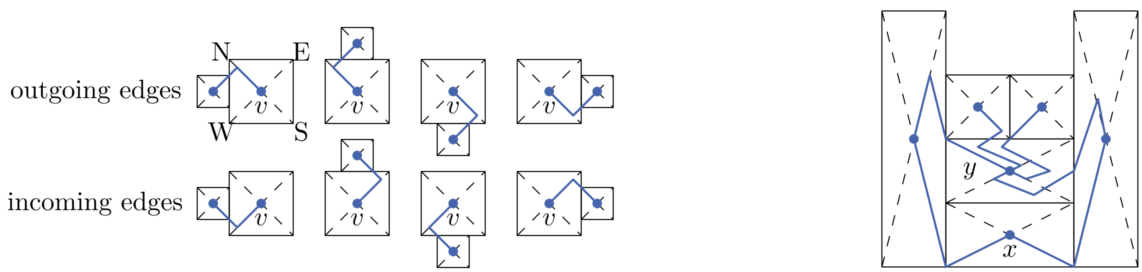

(a) canonical assignment (b) around y

Figure 11: Port assignment: (b) around neighbor y of outer subdivision vertex x.

around v as follows (see Fig. 11(a)). An outgoing edge (v, u) is assigned to the North port, if R_u is to the left or the top of R_v. Otherwise it is assigned to the South port. An incoming edge (u, v) is assigned to the West port, if R_u is to the left or the bottom of R_v. Otherwise it is assigned to the East port.

In the following we will compute a port assignment such that each edge is assigned in a canonical way to at least one of its end vertices and such that crossings between edges incident to the same vertex can be avoided within the rectangle of the common end vertex. We will exploit this property, alongside with the absence of 2-cycles, to prove that such an assignment determines a perturbed generalized planar L-drawing of the plane graph H_x. In the following, we separate the discussion according to the modality of the vertices.

0-Modal Vertices. We consider each 0-modal vertex v to be 2-modal by adding a *virtual edge* inside its designated face f. Namely, suppose v is a source, and let $e_1 = (v, w_1)$ and $e_2 = (v, w_2)$ be the two edges incident to v and to f. We add a virtual edge (w, v) connecting v to a new *virtual vertex* w, so that (w, v) appears between e_1 and e_2 in the circular order of the edges around v. Of course, there is not literally a rectangle R_w representing w, but for the assignments of the edges to the ports of v, we assume that R_w is the degenerate rectangle corresponding to the segment on the intersection of R_{w_1} and R_{w_2}. See Fig. 10(b).

2-Modal Vertices. Let now v be a 2-modal vertex. We discuss the cases where we have to deviate from the canonical assignment. We call a side s of a rectangle in the rectangular dual to be *mono-directed*, *bi-directed*, or *3-directed*, respectively, if there are 0, 1, or 2 changes of directions

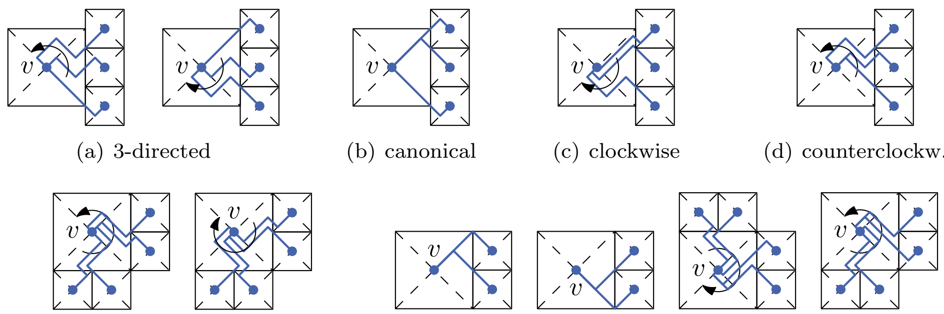

Figure 12: Assignment of ports when the direction of edges incident to one side of a rectangle changes a) twice b-e) once, or f) never.

of the edges across s. See Fig. 12. Observe that by 2-modality there cannot be more than two changes of directions.

Consider first the case that R_v has a side that is 3-directed, say its right side. See Fig. 12(a). If from top to bottom there are first outgoing edges, followed by incoming edges, and followed again by outgoing edges, then we assign from top to bottom first the North port, then the East port, and then the South port to the edges connecting v to vertices whose rectangles lie on the right of R_v (*counterclockwise switch*). Otherwise, we assign from top to bottom first the East port, then the South port, and then the West port (*clockwise switch*). All other edges are assigned in a canonical way to the ports of v; observe that there is no change of directions in the other sides of R_v.

Consider now the case that R_v has one side that is bi-directed, say again the right side of R_v. If the order from top to bottom is first incoming and then outgoing, then assign the edges incident to the right side of R_v in a canonical way (*canonical switch*, Fig. 12(b)). Otherwise *(unpleasant switch)*, we have two options: we either assign the outgoing edges to the North port and the incoming edges to the East port (*counter-clockwise switch*, Fig. 12(d)), or we assign the outgoing edges to the South port and the incoming edges to the West port (*clockwise switch*, Fig. 12(c)).

Observe that if there is an unpleasant switch on one side of R_v, then there cannot be a canonical switch on an adjacent side. Assume now that there are two adjacent sides s_1 and s_2 of R_v in this clockwise order around R_v with unpleasant switches. Then we consider both switches as counterclockwise or both as clockwise. Observe that, due to 2-modality, two opposite sides of R_v are neither both involved in unpleasant switches nor both in canonical switches.

Consider now the case that one side s of R_v is mono-directed, say again the right side of R_v. See Fig. 12(f). In most cases, we assign the edges incident to s in a canonical way. There are – up to symmetry – the following exceptions: The top side of R_v is involved in a clockwise switch and the edges at the right side are incoming edges. In that case we have a *clockwise switch* at s, i.e., the edges at the right side are assigned to the West port of v. The other possible exception occurs when the bottom side of R_v is involved in a counter-clockwise switch and the edges at the right side are outgoing edges. In that case we have a *counter-clockwise switch* at s, i.e., the edges at the right side are assigned to the North port of v.

In order to avoid switches at mono-directed sides, we handle these exceptions as follows: Let s_1, s, s_2 be three consecutive sides in this clockwise order around the rectangle R_v such that there is an unpleasant switch on side s; say s is the right side of R_v, s_1 is the top, and s_2 is the bottom,

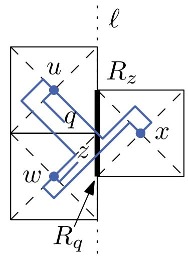

Figure 13: Extra Rule

and the edges on the right side are, from top to bottom, first outgoing and then incoming. By 2-modality, there cannot be a switch of directions on both, s_1 and s_2, i.e., s_1 contains no incoming edges or s_2 contains no outgoing edges. In the first case, we opt for a counterclockwise switch for s, otherwise, we opt for a clockwise switch.

There is one exception to the rule in the previous paragraph (which we refer to as Extra Rule): Let u and w be two adjacent 0-modal vertices with the same designated face f such that the two virtual end vertices are on a line ℓ. See Fig. 13. Let s_u be the side of R_u intersecting R_w and let s_w be the side of R_w intersecting R_u. Assume that u has an unpleasant switch at s_u and, consequently, w has an unpleasant switch at s_w. Let x be the third vertex on f and let s_x be the side of R_x intersecting R_u and R_w. Observe that $s_x \subset \ell$. Do the switch at s_u and s_w in clockwise direction if and only if the switch at s_x is in clockwise direction, otherwise in counterclockwise direction.

4-Modal Vertices. If v is an inner 4-modal vertex of degree 4, then each side of R_v is incident to exactly one rectangle, and we always use the canonical assignment. Suppose now that v is a 4-modal vertex of degree 5, which can only happen if v is the inner vertex y adjacent to the subdivision vertex x of an edge incident to the outer face. Note that we do not have to draw the edge between x and y. However, this case is still different from the previous one, since there are two rectangles incident to the same side s of R_y. If the switch at s is canonical, then there is no problem. Otherwise, we do the assignment as in Fig. 11(b).

Observe that we get one edge between y and a vertex on the outer face that is not assigned in a canonical way at y. But this edge is assigned in a canonical way at the vertex in the outer face. This completes the port assignments.

4.3 Correctness and Run Time

We prove that the constructed port assignment admits a perturbed generalized planar L-drawing and, thus, a planar L-drawing of H. The proof starts by showing that each edge is assigned in a canonical way at one end vertex at least (Lemma 6). Then we route the edges as indicated in the proof of Lemma 7. Finally, in Lemma 8, we discuss the run time. We conclude by presenting the proof of Theorem 2.

Lemma 6 *Every edge $e = (u, w)$ is assigned in a canonical way at u or w.*

Proof: Suppose, for contradiction, that there is an edge $e = (u, w)$ that is assigned in a canonical way neither at u nor at w. Assume, without loss of generality, that R_u is on top of R_w. This implies that e is assigned to the North port of u and to the West port of w. Moreover, the bottom side s_u of R_u is involved in a clockwise switch and the top side s_w of R_w in a counter-clockwise switch. See Fig. 14.

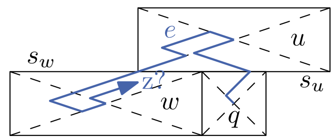

Figure 14: Edges are canonical at one end vertex at least.

First consider the case that neither s_u nor s_w is mono-directed. Then the only reason for these unpleasant switches is that 1. u has an incoming edge (q, u) such that R_q is incident to the bottom of R_u and is to the right of R_w, and 2. w has an outgoing edge (w, z) such that R_z is incident to the top of R_w and is to the right of R_u. This is not simultaneously possible if at least one among R_z or R_q is a real rectangle. In the case that R_z and R_q were both virtual, the designated face of u and w would be the same, R_z and R_q would be collinear and, thus, the Extra Rule would apply. However, this implies that the switches at u and w are both clockwise or both counterclockwise. See Fig. 13.

Consider now the case that one of the two sides s_u and s_w, say s_w, is mono-directed. By construction, a non-canonical switch at a mono-directed side can only happen if the vertex is one of the 0-modal vertices to which the Extra Rule was applied. See Fig. 15.

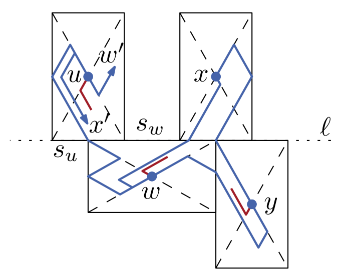

Figure 15: Illustration of the proof of Lemma 6 in the case of the Extra Rule.

We distinguish two cases based on whether s_u is bidirected or mono-directed. Suppose first that s_u is bidirected. As above there must still be a neighbor q of u such that R_q is to the right of R_w and below R_u, which however is not possible.

Assume now that also s_u is mono-directed, i.e., the Extra Rule was also applied to u and a vertex w', and the designated face f' of u and w' is as indicated by the (red) stub incident to u in Fig. 15. Let x' be the third vertex incident to f'. Observe that R_u and $R_{w'}$ must have a common corner that lies on a side of $R_{x'}$. However, this is only possible if the bottom right corner of R_u and the bottom left corner of $R_{w'}$ lie on the top side of $R_{x'}$, which implies that $x' = w$. But w is a 0-modal vertex and x' cannot be 0-modal. □

Lemma 7 *The constructed port assignment admits a planar L-drawing of H.*

Proof: In the following we show how to route the edges in order to obtain a perturbed generalized planar L-drawing with the given port assignment. By Lemma 5 this is sufficient to obtain a planar L-drawing with the same port assignment.

Recall that in a perturbed generalized planar L-drawing each edge is routed as a polyline lying in the two rectangles corresponding to its end vertices, composed of segments parallel to the respective diagonals and satisfying the following properties:

(1) each directed edge $e = (u, v)$ starts with a segment on the diagonal $\backslash_u$ of R_u from the upper left to the lower right corner and ends with a segment on the diagonal $/_v$ of R_v from the lower left to the upper right corner.

(2) The polylines representing two edges overlap in at most a common straight-line prefix or a suffix and they do not cross.

(3) For an edge e one of the following is true: (i) none of the two end vertices of $\text{mid}(e)$ is a bend of e and $\text{rot}(e) = \pm 1$ or (ii) one of the two end vertices of $\text{mid}(e)$, but not both, is a bend of e and $\text{rot}(\text{mid}(e)) = 0$.

Property 1 is already fulfilled by routing the initial and final segment of each edge along the diagonal prescribed by the constructed port assignment. Next, we show that the port assignment allows for a routing of the edges that also fulfills Properties 2 and 3.

Let e be an edge between v and w that is assigned in a canonical way at w. Let e be assigned to the port p of R_w. Let d be the corner of R_w at the end of the diagonal corresponding to the port p. We define how to route e distinguishing three main cases depending on the relationship of R_v and R_w with respect to d – see the columns of Fig. 16. Each case has two subcases depending on whether e is assigned in a canonical way at v or not – see the rows of Fig. 16. We subdivide the case where e is assigned in a canonical way at v into two additional subcases b.1 and b.2. Let u be the common neighbor of v and w that is incident to the side of w containing the corner d. Let d' be the common point of R_u, R_v, and R_w. The Subcase b.2 is the following: There is an edge connecting u to either v or w that is assigned in a canonical way at u and in a non-canonical way at the other end-vertex, whose drawing according to Case 2a passes through d'. Observe that in the last column of Row b.2 the corner d does not have to be on a side of R_u since R_u might be smaller.

For each of the cases, we draw e as sketched in the corresponding box in Fig. 16: The first and the last segment of an edge represents a straight line of an appropriate length, while every other segment represents a perturbed orthogonal polyline with rotation 0. If e is assigned to the same port as another edge e', we make sure that the routing respects the embedding by appropriately selecting the length of the first or last segment of e and e'. E.g., assume that e' follows e in counterclockwise order around vertex v such that both are assigned to the North port of v and assume that the first bend of both, e and e', is a right turn. Then the bend of e is closer to v than the bend of e'.

Observe that the edges are only routed within the rectangles of their end vertices. Thus, if there were crossings then they would involve edges incident to the same vertex, say v, and would lie within the rectangle R_v. However, the port assignment respects the planar embedding and, thus, ensures that the middle part of the edges can be routed such that no two edges cross. This guarantees that Property 2 is fulfilled.

It remains to show that also Property 3 is fulfilled. This is trivial for the cases in Row b.1, since in this case the edge is composed of two parts, each having rotation 0, which meet at a bend b. Thus, $\text{mid}(e)$ either starts at b and has rotation 0, or it contains b as an inner point and has rotation ± 1.

Analogously, for the other cases, it suffices to prove that $\text{mid}(e)$ contains two out of the three indicated bends as inner points – one with a left turn and one with a right turn. In fact, in this

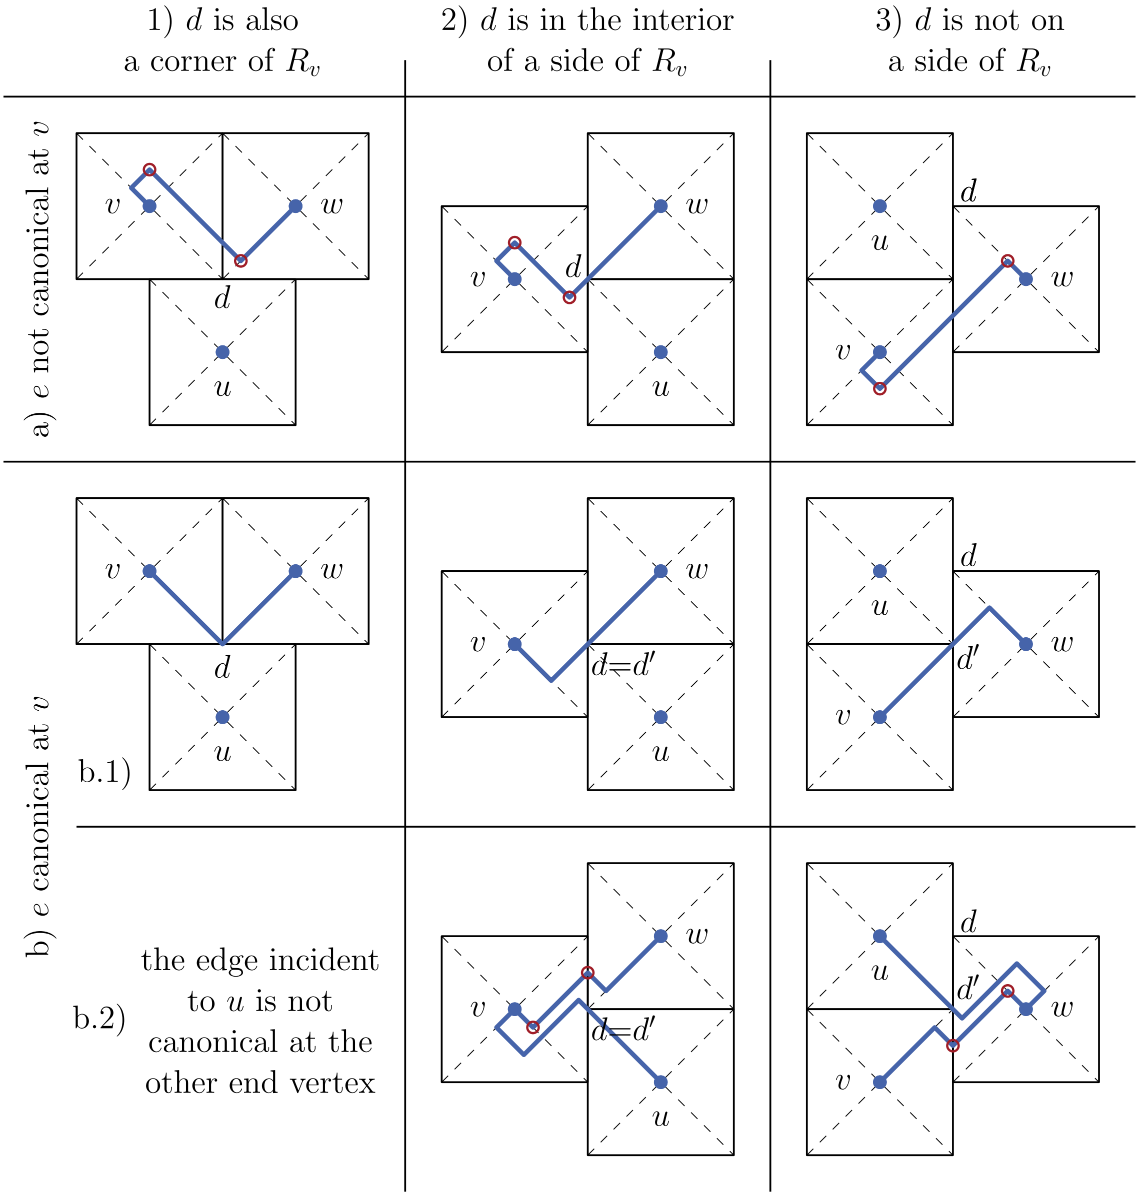

Figure 16: How to route the edge e between v and w. Point d is the corner at the end of the diagonal of R_w to which e is assigned.

case mid(e) has rotation 0 or ± 1 depending on whether the third bend is an end point of mid(e) or not.

To this end, we prove that each of the bends that are encircled red in Fig. 16 are inner points of mid(e), i.e., that they do not lie on any other edge. This is obvious, if the bend is not on a diagonal, since the end points of mid(e) lie on the diagonals of the rectangles of the end vertices of e. If the bend is on the diagonal, we prove that there are no edges leaving the diagonal after the bend. Note that if there was such an edge then it would be one that immediately follows or precedes e in the cyclic order around the respective end vertex.

b.2) We argue about the bend b on the diagonal of R_v in Column 2, the arguments for the bend on the diagonal of R_w in Column 3 is analogous. Observe that the edge e is immediately followed by the edge (u, v) in the cyclic order around v. Since (u, v) is assigned to a different port of v than e, the statement follows.

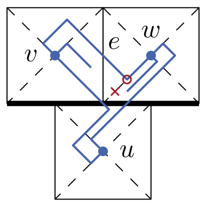

(a) Case 1a

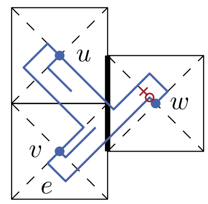

(b) Case 3a

Figure 17: The bend encircled red is an inner point of mid(e) in the case of the Extra Rule.

1a+3a) These two cases can only happen if we had applied the Extra Rule, as otherwise e would have been assigned in a canonical way at v. See Fig. 17. I.e., two among u, v, w are 0-modal vertices, their designated face is the one bounded by u, v, w, the respective virtual rectangles are collinear, and there is an unpleasant switch at the third vertex. Further, the non-virtual edge that immediately follows e (Case 3a) or precedes e (Case 1a) in the cyclic order around w is the edge (u, w). However, the port assignment in the Extra Rule guarantees that (u, w) is assigned to a different port of w than e.

This concludes the proof of the lemma. □

Lemma 8 *A planar L-drawing of H in which the drawing of the outer face is Γ_0 and no face contains bad pincers of G can be constructed in linear time.*

Proof: The construction guarantees a planar L-drawing of H. The port assignment is such that there are no bad pincers and for the outer face it is the same as in Γ_0. A rectangular dual can be constructed in linear time [7, 20]. The port assignment can also be done in linear time. Finally, the coordinates can be computed in linear time using topological ordering – see Theorem 1. □

Proof: [Proof of Theorem 2] Let D be a plane bimodal graph without 2-cycles. We can triangulate D in linear time, by adding a directed triangle in the outer face, by applying Lemma 3, and by inserting a vertex in each face of degree four. The decomposition of the resulting graph G into its 4-connected components can be performed in linear time [19]. Proceeding from the outermost to the innermost components, we compute a planar L-drawing for each 4-connected component H, avoiding bad pincers and such that the outer face Γ_0 is the one determined by the respective triangular face of the parent component of H. We can do so in linear time by Lemma 8. □

5 Bimodal Planar 3-Trees

A *planar 3-tree* is defined recursively: The complete graph K_4 on four vertices is a planar 3-tree. Adding a new vertex into an inner face f of a planar 3-tree and connecting it to the 3 vertices on the boundary of f yields again a planar 3-tree.

Theorem 3 *A bimodal graph has a planar L-drawing if the underlying undirected graph, after the identification of parallel edges due to 2-cycles, is a planar 3-tree.*

Before we prove the theorem, we collect some useful properties on 2-cycles in bimodal graphs and on pincers.[1]

Remark 2 *In a bimodal 3-connected graph*

1. *a vertex can be incident to at most two 2-cycles and*
2. *at most two edges of a triangle can be part of 2-cycles.*

Proof: There is no way of ordering the edges of three 2-cycles around a vertex in a bimodal way. Now consider a triangle and assume that each vertex of the triangle is incident to two 2-cycles. Observe that if a vertex is incident to two 2-cycles then the incoming incident edges and the outgoing incident edges, respectively, have to be consecutive. This is not possible simultaneously for all three vertices since a triangle has an odd length. □

Remark 3 *Let T be a triangle of a bimodal graph G and let v be a sink(source)-switch of T. If v is incident to an outgoing edge outside T, then v is not incident to a pincer on T.*

Proof: Otherwise v would also be incident to an outgoing edge inside T and, thus, v would not be bimodal. □

Proof: [of Theorem 3] Observe that there are no separating 2-cycles in the digraph since planar 3-trees are 3-connected. The idea of the construction is as follows. Although, by definition, the smallest planar 3-tree is a K_4, we start the induction with a graph containing three vertices. The underlying undirected graph induced by the three vertices contains a triangle plus some additional parallel edges. We find a planar L-drawing for that triangle with its 2-cycles and then keep inserting the vertices maintaining the invariant that there are no bad pincers. Observe that each inserted vertex has three adjacent vertices in the current digraph and up to five incident edges. We call a vertex an *in/out-vertex* of a face f if it is neither a source switch nor a sink switch of f.

Base Case. The input graph is a bimodal graph consisting of three vertices. Recall that by definition the edges form a triangle with possibly some 2-cycles. If the outer face is a 2-cycle, remove one of its two edges. Draw the triangle T bounding the outer face such that the respective vertices of the 2-cycle are extreme points of a diagonal of the bounding box of the drawing (Figs. 18(a) to 18(d)) and add the missing edge of the outer 2-cycle. If the outer 2-cycle had been between the source and the sink of T then – due to bimodality at the source and the sink – T is not involved in any further 2-cycles. Otherwise, by Remark 2.2, there can be at most one more 2-cycle. By Remark 3, bad pincers are not possible in these cases: The vertices incident to the external 2-cycle are already incident to two edges in opposite direction.

[1] Recall that (bad) pincers were defined in Sect. 4.1 and depicted in Fig. 8.

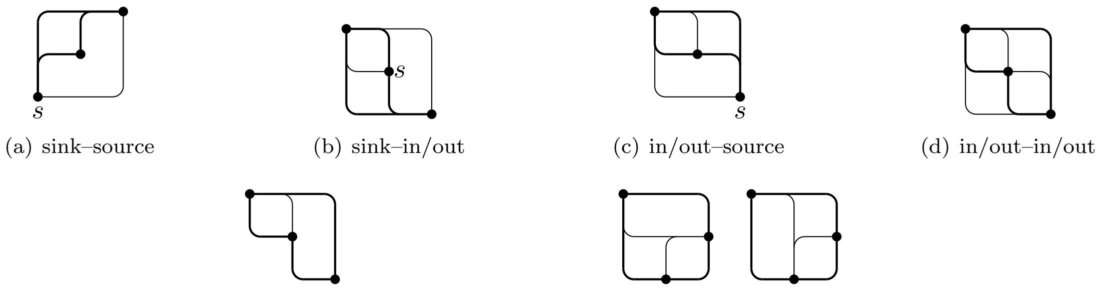

(a) sink–source (b) sink–in/out (c) in/out–source (d) in/out–in/out

(e) directed (f) not directed

Figure 18: Different ways of drawing a triangle T with possible 2-cycles. The edges of T are thicker. Top row: the outer face is bounded by a 2-cycle where the edge not in T connects the indicated vertices of T. Observe that in (d) only one of the edges inside T is possible. Bottom row: the outer face is bounded by T and T is directed or not.

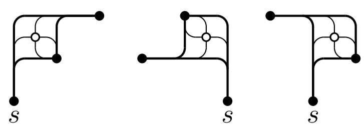

Figure 19: Different ways of drawing a triangle where the angle at both, the source and the sink, is 0° and how to add an additional internal vertex – where only one of the edges in the 2-cycle is present.

Consider now the case that the outer face is bounded by a triangle T. We say that a *triangle contains a 2-cycle* if one edge of the 2-cycle is an edge of the triangle and the other edge of the 2-cycle is in the interior of the triangle. Note that if a triangle contains two 2-cycles then, by bimodality, the common vertex of the two 2-cycles must be the source or the sink of the triangle. Thus, if the triangle is a directed cycle, it contains at most one 2-cycle.

Now, if T is a directed cycle, draw it as indicated in Fig. 18(e). Otherwise, the three vertices of T are a source, a sink, and an in/out vertex. Start with the drawing of T where no two edges are attached to the same port of a vertex. Add the 2-cycles. Thus, T with its 2-cycles is a subgraph of one of the two cases in Fig. 18(f). The choice of the initial drawing of T and Remark 3 imply that there are no bad pincers.

Inductive Case. Let T be a triangle bounding a face. We show how to insert a vertex v connected to the three vertices of T.

Assume first that T has a source s, a sink t and an in/out-vertex w. There are three cases: (1) The angles in T at s and t, respectively, are both 0° (Cases afb in Fig. 9 or Fig. 19, respectively), (2) the angle at either the source or the sink – say the source – is 0° (Cases cg and dh in Fig. 9 or Fig. 20, Columns a+b, respectively), or (3) T does not contain a 0° angle (Case e in Fig. 9 or Fig. 20, Column c, respectively).

In the first case – since by induction there are no bad pincers in T – the direction of the edges between v on one hand and the source and the sink on the other hand are fixed and it follows that – due to bimodality – v cannot be incident to a 2-cycle. By Remark 3, the pair of edges with

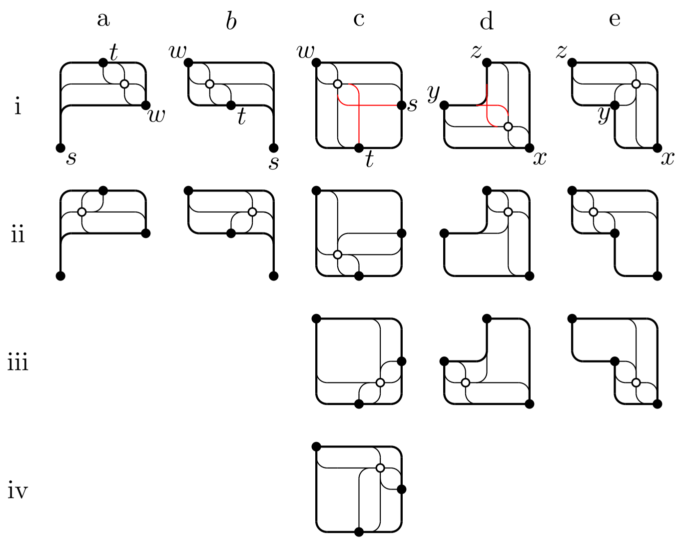

Figure 20: Different ways of inserting a vertex into a triangle. Not all edges are present simultaneously.

a 0° angle at the in/out-vertex w of T cannot be a pincer. Thus, we can add v and its incident edges as indicated in Fig. 19 – where only one of the edges in the 2-cycle is present.

In Cases 2 and 3, we have the property that if there is a 2-cycle between v and w then – due to bimodality of w – the order around w is fixed and must be such that on both sides an edge of T and an edge of the 2-cycle forms a 0° angle. Consider now Case 2 and assume w.l.o.g. that there are two edges of T that form a 0° angle at s. Up to symmetry there are two such drawings of T. See Fig. 20, Columns a+b. If t is incident to a 2-cycle with v then the order of the edges in the 2-cycle determines whether we are in Case (i) or (ii). By Remark 3, no two edges that are assigned to the same port are a pincer. If t is not incident to a 2-cycle with v, we can choose between Case (i) and (ii) and we do it such that we do not create a bad pincer incident to t. Observe that, due to bimodality, a vertex cannot be incident to two pincers whose triangles are internally disjoint.

Consider now Case 3 (Fig. 20, Column c). There can be at most two 2-cycles incident to v, and if so, fixing the ordering of the edges of one 2-cycle also fixes the ordering of the other (due to bimodality of v). Depending on this ordering, we can always choose one out of the Cases (ii) or (iv) such that no bad pincer occurs. If v is incident to at most one 2-cycle, the choice depends on pincers at v, s and t and can always be done without creating bad pincers: There are all four variants of pairs of 0° angles at s and t, so we can always choose one that does not contain a bad pincer.

If v is a source or a sink, we have to also make sure that there is no bad pincer incident to v. Observe that in this case v is not incident to any 2-cycles. It follows that T cannot contain edges incident to both s and t that are involved in pincers with edges incident to v. E.g., assume that v is a source. Then only (v, t) can be involved in a bad pincer: (v, s) and an edge on T incident to s do not have the same direction at s, and w is already incident to an edge in the opposite direction as (v, w). Among the four possibilities we can always choose one that has neither a bad pincer incident to v nor a bad pincer incident to t.

Assume now that T is a directed triangle (Cases x and y in Fig. 9 or Fig. 20, Columns d+e, respectively). Observe that due to bimodality v cannot be incident to two 2-cycles. By Remark 3, there cannot be a pincer that consists of an edge in T and an edge incident to v. For the same reason, if v is neither a source nor a sink, then there cannot be a pincer incident to v. It remains to check for bad pincers at v in the case were v is a source or a sink. Among the three possibilities in Column a or b, we can always choose one that has no bad pincer at v. □

6 Outerplanar Digraphs

Since there exist outerplanar digraphs that do not admit any bimodal embedding [6], we cannot exploit Theorem 2 to construct planar L-drawings for the graphs in this class. However, we are able to prove the following.

Theorem 4 *Every outerplanar graph admits a planar L-drawing.*

Proof: Put all vertices on a diagonal in the order in which they appear on the outer face – starting from an arbitrary vertex. The drawing of the edges is determined by the direction of the edges and the relative position of the end vertices, since an edge must start with a vertical segment and end with a horizontal segment. This implies that some edges are drawn above and some below the diagonal. By outerplanarity, there are no crossings. □

We remark that Theorem 4 provides an alternative proof to the one in [6] that any outerplanar digraph admits a 4-modal embedding. Observe that the planar L-drawings constructed in the proof of Theorem 4 are not necessarily outerplanar. In the following, we prove that this may be unavoidable.

Theorem 5 *Not every outerplanar graph admits an outerplanar L-drawing.*

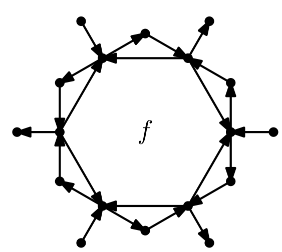

Proof: Consider the graph depicted above. It has a unique outerplanar embedding. Let f be the inner face of degree 6. Each vertex incident to f is 4-modal and is a source switch or a sink switch of f. Thus, the angle at each vertex is 0°. The angle at each bend is at most $3/2\pi$. Thus, the angular sum around f would imply $(2 \cdot \deg f - 2) \cdot \pi \leq 3/2 \cdot \deg f \cdot \pi$, which is not true for $\deg f = 6$. □

As established in Theorem 6 below, there are even 4-modal biconnected internally triangulated outerplane digraphs that do not admit an outerplanar L-drawing. The proof of this fact relies on the following lemma.

Lemma 9 *Every biconnected internally triangulated outerplanar digraph G with a 4-modal outerplanar embedding can be extended to an internally triangulated outerplanar digraph G' with a 4-modal outerplanar embedding in which all vertices of G are 4-modal.*

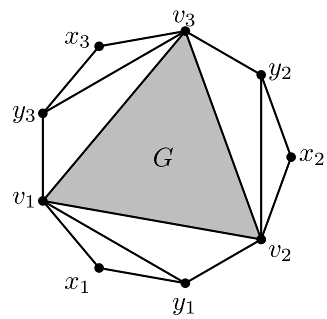

Figure 21: How to make the vertices of a biconnected internally triangulated outerplanar digraph 4-modal.

Proof: See Fig. 21 for an illustration. Let $v_1, \ldots, v_n$ be the vertices of G in the order in which they appear on the outer face and let $v_{n+1} = v_1$. For $i = 1, \ldots, n$ add a new vertex x_i with neighbor v_i and a vertex y_i with neighbors x_i, v_i, and v_{i+1}. Now each new vertex has degree at most three and thus, will be 2-modal, no matter how we orient the edges. Each vertex v_i, $i = 1, \ldots, n$ of G is incident to three new edges. These can be oriented such that each v_i, $i = 1, \ldots, n$ becomes 4-modal. □

Theorem 6 *Not every biconnected internally triangulated outerplanar digraph with a 4-modal outerplanar embedding has an outerplanar L-drawing.*

Proof: We consider the digraph $G = (V, E)$ in Fig. 22(a) augmented as described in Lemma 9. The resulting digraph G' is a biconnected internally triangulated outerplanar digraph, and its unique outerplanar embedding is 4-modal. It has 57 vertices. Each vertex of G is 4-modal in G'. We show that G' has no planar L-drawing with its outerplanar embedding. Fig. 22(b) shows a non-planar L-drawing of G where the cyclic order of the edges is as in the outerplanar embedding.

Consider now the following flow network $\mathcal{L}_G$ associated with a plane digraph $G = (V, E)$. Let F be the set of faces of G. $\mathcal{L}_G$ has node set $W := V \cup E \cup F$, arcs from vertices to incident faces, and from faces to incident edges, and supplies

- $b(v) = \frac{4-\text{mod}(v)}{2}$, $v \in V$,
- $b(f) = 2 + \#$source switches of f, f inner face,
- $b(f) = -2 + \#$source switches of f, f outer face,
- $b(e) = -1$, $e \in E$.

Based on the relationship to Kandinsky drawings, Chaplick et al. [13] showed that any planar L-drawing of a biconnected digraph G yields a flow ϕ in $\mathcal{L}_G$ as follows: Let $\alpha(v, f)$ be the angle in the face f at a vertex v. Then $\phi(v, f) = \alpha(v, f)/\pi$ if the two edges incident to both, v and f, are two outgoing or two incoming edges of v and $\phi(v, f) = \alpha(v, f)/\pi - 1/2$, otherwise. $\phi(f, e) = 1$ if and only if there is a convex bend in face f on edge e and 0 otherwise.

However, not every flow on $\mathcal{L}_G$ corresponds to a planar L-drawing. In effect, the angles and bends in the non-planar L-drawing in Fig. 22(b) also yield a flow ϕ' in $\mathcal{L}_G$.

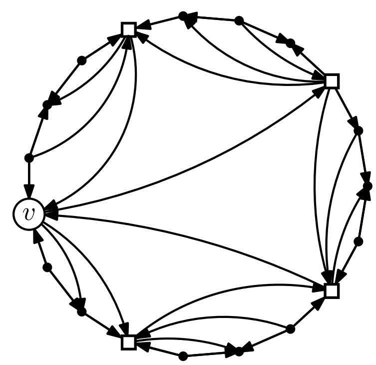

(a) graph

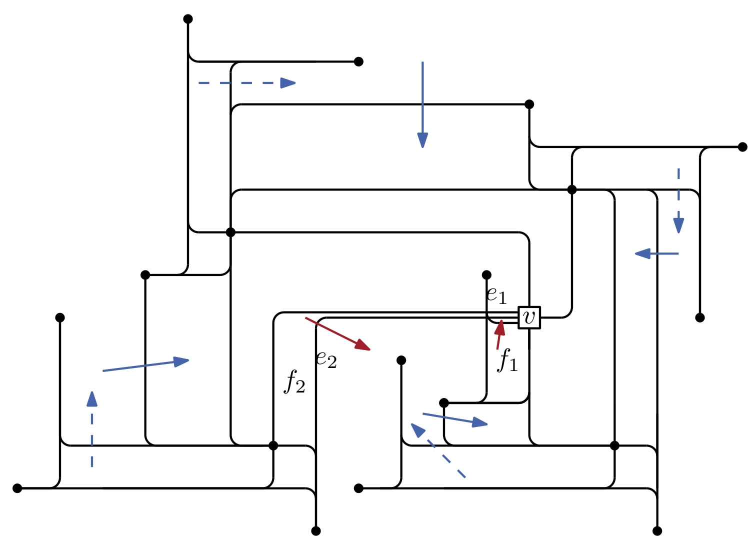

(b) non-planar L-drawing

Figure 22: A biconnected internally triangulated outerplanar digraph with a 4-modal outerplanar embedding that has no outerplanar L-drawing – in an extension that makes the present vertices 4-modal. b) Different from the drawing conventions of L-drawings, the edges at the West port of v are drawn slightly apart in order to make the embedding visible.

Assume now that there is a planar L-drawing Γ' of G' and consider the drawing Γ of G induced by Γ'. Let ϕ be the flow that corresponds to Γ. The difference between ϕ and ϕ' is a union C of directed cycles in the residual network $\mathcal{L}_{G,\phi'}$ of $\mathcal{L}_G$ with respect to the flow ϕ' (see e.g. [1, Theorem 3.7]). Since in G' all vertices of G are 4-modal, no arc from a vertex of G to an inner face of G can carry flow. Hence, C does not contain vertices of G. The direction of the arcs in $\mathcal{L}_{G,\phi'}$ is from a face f to an incident edge e with a concave bend in e and from e to the other incident face f'.

Since ϕ corresponds to a planar L-drawing it follows that at least one of the edges e_1 or e_2 must have a concave bend in the outer face f_o. Thus, C contains the arcs $(f_2, e_2), (e_2, f_o)$ or the arcs $(f_1, e_1), (e_1, f_o)$. This implies that C must also contain an arc from the outer face. But there are only three such arcs (solid blue arcs). Whenever such an arc is contained in C then also the respective dashed blue arc has to be contained in C. Since in a union of directed cycles the number of incoming and outgoing edges of each vertex must be the same, we have that C can neither contain the arc (e_2, f_o) nor (e_1, f_o). □

7 Open Problems

We conclude with several directions for future work.

- Are there bimodal graphs with 2-cycles that do not admit any planar L-drawing (with or without the given embedding)?
- What is the complexity of testing whether a 4-modal graph admits a planar L-drawing with a fixed embedding?
- In the directed Kandinsky model where edges leave a vertex to the top or the bottom and enter a vertex from the left or the right, for which k is there always a drawing with at most $1 + 2k$ bends per edge for any 4-modal graph? $k = 0$ does not suffice. What about $k = 1$?
- Can one efficiently test whether an outerplanar graph with a given 4-modal outerplanar embedding admits an outerplanar L-drawing?
- Can one efficiently test whether an upward-planar graph, with or without a given upward-planar embedding, admits an upward-planar L-drawing? Recently, some partial progress in this direction has been made [2].

Acknowledgments

We thank the anonymous reviewers for their useful comments that helped us to refine the presentation of the paper. We also thank Gregor Diatzko for suggesting improvements for the cases described in Fig. 9. He provided a complete implementation of the algorithm described in Sect. 4.2, which includes a new linear-time algorithm for the decomposition of a graph into 4-connected components. The code is available at https://github.com/Parranoh/ldrawing.

References

[1] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Pearson, 1993.

[2] P. Angelini, S. Chaplick, S. Cornelsen, and G. Da Lozzo. On upward-planar L-drawings of graphs. Technical Report arXiv:2205.05627, Cornell University Library, 2022. URL: http://arxiv.org/abs/2205.05627.

[3] P. Angelini, G. Da Lozzo, M. Di Bartolomeo, V. Di Donato, M. Patrignani, V. Roselli, and I. G. Tollis. Algorithms and bounds for L-drawings of directed graphs. *Int. J. Found. Comput. Sci.*, 29(4):461–480, 2018. doi:10.1142/S0129054118410010.

[4] P. Angelini, G. Da Lozzo, G. Di Battista, V. D. Donato, P. Kindermann, G. Rote, and I. Rutter. Windrose planarity: Embedding graphs with direction-constrained edges. *ACM Trans. Algorithms*, 14(4):54:1–54:24, 2018. doi:10.1145/3239561.

[5] W. Barth, P. Mutzel, and C. Yildiz. A new approximation algorithm for bend minimization in the Kandinsky model. In M. Kaufmann and D. Wagner, editors, *GD 2006*, volume 4372 of *LNCS*, pages 343–354. Springer, 2007. doi:10.1007/978-3-540-70904-6_33.

[6] J. J. Besa Vial, G. Da Lozzo, and M. T. Goodrich. Computing k-modal embeddings of planar digraphs. In M. A. Bender, O. Svensson, and G. Herman, editors, *ESA 2019*, volume 144 of *LIPIcs*, pages 19:1–19:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.ESA.2019.19.

[7] J. Bhasker and S. Sahni. A linear algorithm to find a rectangular dual of a planar triangulated graph. *Algorithmica*, 3:247–278, 1988. doi:10.1007/BF01762117.

[8] T. C. Biedl and M. Derka. The (3,1)-ordering for 4-connected planar triangulations. *JGAA*, 20(2):347–362, 2016. doi:10.7155/jgaa.00396.

[9] T. C. Biedl and D. Mondal. A note on plus-contacts, rectangular duals, and box-orthogonal drawings. Technical Report arXiv:1708.09560v1, Cornell University Library, 2017. URL: http://arxiv.org/abs/1708.09560v1.

[10] C. Binucci, W. Didimo, and M. Patrignani. Upward and quasi-upward planarity testing of embedded mixed graphs. *Theor. Comput. Sci.*, 526:75–89, 2014. doi:10.1016/j.tcs.2014.01.015.

[11] T. Bläsius, G. Brückner, and I. Rutter. Complexity of higher-degree orthogonal graph embedding in the Kandinsky model. In A. S. Schulz and D. Wagner, editors, *ESA 2014*, volume 8737 of *LNCS*, pages 161–172. Springer, 2014. doi:10.1007/978-3-662-44777-2_14.

[12] G. Brückner. Higher-degree orthogonal graph drawing with flexibility constraints. Bachelor thesis, Department of Informatics, Karlsruhe Institute of Technology, 2013. URL: https://i11www.iti.kit.edu/_media/teaching/theses/ba-brueckner-13.pdf.

[13] S. Chaplick, M. Chimani, S. Cornelsen, G. Da Lozzo, M. Nöllenburg, M. Patrignani, I. G. Tollis, and A. Wolff. Planar L-drawings of directed graphs. In F. Frati and K.-L. Ma, editors, *GD 2017*, volume 10692 of *LNCS*, pages 465–478. Springer, 2018. doi:10.1007/978-3-319-73915-1_36.

[14] H. de Fraysseix, P. Ossona de Mendez, and P. Rosenstiehl. On triangle contact graphs. *Combinatorics, Probability, and Computing*, 3(2):233–246, 1994. doi:10.1016/j.comgeo.2017.11.001.

[15] M. Eiglsperger and M. Kaufmann. Fast compaction for orthogonal drawings with vertices of prescribed size. In M. Jünger and P. Mutzel, editors, *GD 2001*, volume 2265 of *LNCS*, pages 124–138. Springer, 2002. doi:10.1007/3-540-45848-4_11.

[16] U. Fößmeier and M. Kaufmann. Drawing high degree graphs with low bend numbers. In F. J. Brandenburg, editor, *GD 1995*, volume 1027 of *LNCS*, pages 254–266. Springer, 1996. doi:10.1007/BFb0021809.

[17] X. He. On finding the rectangular duals of planar triangular graphs. *SIAM J. on Computing*, 22(6):1218–1226, 1993. doi:10.1137/0222072.

[18] G. Kant. Algorithms for drawing planar graphs, 1993.

[19] G. Kant. A more compact visibility representation. *Int. J. Comput. Geometry Appl.*, 7(3):197–210, 1997. doi:10.1142/S0218195997000132.

[20] G. Kant and X. He. Two algorithms for finding rectangular duals of planar graphs. In J. van Leeuwen, editor, *WG 1993*, volume 790 of *LNCS*, pages 396–410. Springer, 1994. doi:10.1007/3-540-57899-4_69.

[21] K. Koźmiński and E. Kinnen. Rectangular duals of planar graphs. *Networks*, 15(2):145–157, 1985. doi:10.1002/net.3230150202.

[22] R. Tamassia. On embedding a graph in the grid with the minimum number of bends. *SIAM J. on Computing*, 16:421–444, 1987. doi:10.1137/0216030.

[23] W. T. Tutte. How to draw a graph. *Proceedings of the London Mathematical Society*, s3-13(1):743–767, 1963. doi:10.1112/plms/s3-13.1.743.