



A Topology-Shape-Metrics Framework for Ortho-Radial Graph Drawing

Lukas Barth¹ · Benjamin Niedermann² · Ignaz Rutter³ · Matthias Wolf¹ 

Received: 10 June 2021 / Revised: 14 October 2022 / Accepted: 6 November 2022

© The Author(s) 2023

Abstract

Orthogonal drawings, i.e., embeddings of graphs into grids, are a classic topic in Graph Drawing. Often the goal is to find a drawing that minimizes the number of bends on the edges. A key ingredient for bend minimization algorithms is the existence of an *orthogonal representation* that allows to describe such drawings purely combinatorially by only listing the angles between the edges around each vertex and the directions of bends on the edges, but neglecting any kind of geometric information such as vertex coordinates or edge lengths. In this work, we generalize this idea to *ortho-radial representations* of *ortho-radial drawings*, which are embeddings into an ortho-radial grid, whose gridlines are concentric circles around the origin and straight-line spokes emanating from the origin but excluding the origin itself. Unlike the orthogonal case, there

Editor in Charge: Csaba D. Tóth

This manuscript is based on the two conference papers (1) Barth, L., Niedermann, B., Rutter, I., Wolf, M.: Towards a topology-shape-metrics framework for ortho-radial drawings. In: 33rd International Symposium on Computational Geometry (Brisbane 2017). Leibniz Int. Proc. Inform., vol. 77, # 14. Leibniz-Zent. Inform., Wadern (2017); (2) Niedermann, B., Rutter, I., Wolf, M.: Efficient algorithms for ortho-radial graph drawing. In: 35th International Symposium on Computational Geometry (Portland 2019). Leibniz Int. Proc. Inform., vol. 129, # 53. Leibniz-Zent. Inform., Wadern (2019).

Lukas Barth
lukas.barth@kit.edu

Benjamin Niedermann
benjamin.niedermann@yworks.com

Ignaz Rutter
rutter@fim.uni-passau.de

Matthias Wolf
matthias.wolf@kit.edu

¹ Karlsruhe Institute of Technology, Karlsruhe, Germany

² yWorks GmbH, Tübingen, Germany

³ University of Passau, Passau, Germany

exist ortho-radial representations that do not admit a corresponding drawing, for example so-called strictly monotone cycles. An ortho-radial representation is called *valid* if it does not contain a strictly monotone cycle. Our first main result is that an ortho-radial representation admits a corresponding drawing if and only if it is valid. Previously such a characterization was only known for ortho-radial drawings of paths, cycles, and theta graphs (Hasheminezhad et al. in *Australas J Combin* 44:171–182, 2009), and in the special case of rectangular drawings of cubic graphs (Hasheminezhad et al. in *Comput Geom* 43(9):767–780, 2010), where the contour of each face is required to be a combinatorial rectangle. Additionally, we give a quadratic-time algorithm that tests for a given ortho-radial representation whether it is valid, and we show how to draw a valid ortho-radial representation in the same running time. Altogether, this reduces the problem of computing a minimum-bend ortho-radial drawing to the task of computing a valid ortho-radial representation with the minimum number of bends, and hence establishes an ortho-radial analogue of the topology-shape-metrics framework for planar orthogonal drawings by Tamassia (*SIAM J Comput* 16(3):421–444, 1987).

Keywords Graph drawing · Ortho-radial graph drawing · Ortho-radial representation · Topology-shape-metrics · Efficient algorithms

Mathematics Subject Classification 68R10

1 Introduction

Grid drawings of graphs embed graphs into grids such that vertices map to grid points and edges map to internally disjoint curves on the grid lines that connect their endpoints. Orthogonal grids, whose grid lines are horizontal and vertical lines, are popular and widely used in Graph Drawing. Among other applications, orthogonal graph drawings are used in VLSI design (e.g., [5, 35]), diagrams (e.g., [3, 14, 20, 38]), and network layouts (e.g., [24, 30]). They have been extensively studied with respect to their construction and properties (e.g., [1, 6, 7, 29, 33]). Moreover, they have been generalized to arbitrary planar graphs with degree higher than four (e.g., [8, 18, 32]).

Ortho-radial drawings are a generalization of orthogonal drawings to grids that are formed by concentric circles around the origin and straight-line spokes from the origin, but excluding the origin. Equivalently, they can be viewed as graphs drawn in an orthogonal fashion on the surface of a standing cylinder, see Fig. 1, or a sphere without the poles. Hence, they naturally bring orthogonal graph drawings to the third dimension.

Among other applications, ortho-radial drawings are used to visualize network maps; see Fig. 2¹. Especially, for metro systems of metropolitan areas they are highly suitable. Their inherent structure emphasizes the city center, the metro lines that run in circles as well as the metro lines that lead to suburban areas. While the automatic creation of metro maps has been extensively studied for other layout styles such as

¹ Note that ortho-radial drawings exclude the center of the grid, which is slightly different to the concentric circles maps by Maxwell J. Roberts.

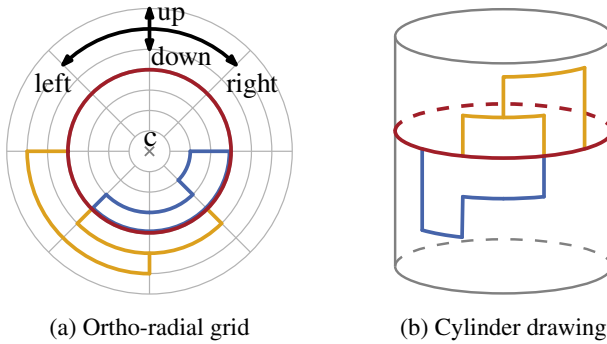


Fig. 1 An ortho-radial drawing of a graph on a grid (a) and its equivalent interpretation as an orthogonal drawing on a cylinder (b)

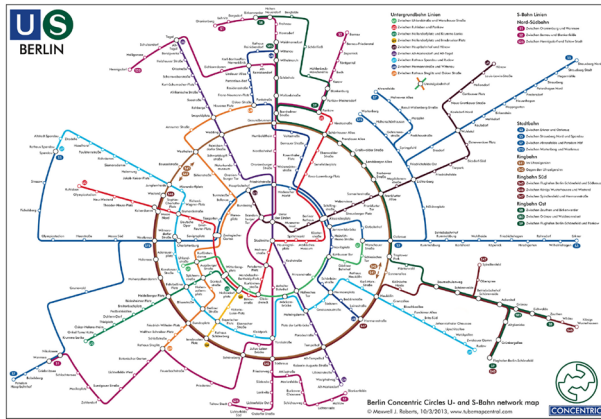


Fig. 2 Metro map of Berlin using an ortho-radial layout. Image copyright by Maxwell J. Roberts. Reproduced with permission

otilinear edge directions [23, 26, 28, 36] and curved metro lines [16], this is a new and wide research field for ortho-radial drawings [2, 17]. For a more detailed discussion on layouts of transit maps we refer to the survey by Wu et al. [37].

Adapting existing techniques and objectives from orthogonal graph drawings is a promising step to open up that field. One main objective in orthogonal graph drawing is to minimize the number of bends on the edges. The key ingredient of a large fraction of the algorithmic work on this problem is the *orthogonal representation*, introduced by Tamassia [31], which describes orthogonal drawings by listing (i) the angles formed by consecutive edges around each vertex, and (ii) the directions of bends along the edges. Such a representation is *valid* if (I) the angles around each vertex sum to 2π , and (II) the sum of the angles around each face with k vertices is $(k - 2) \cdot \pi$ for internal faces and $(k + 2) \cdot \pi$ for the outer face. The necessity of the first condition is obvious and the necessity of the latter follows from the sum of inner/outer angles of any polygon with k corners. It is thus clear that any orthogonal drawing yields a valid orthogonal representation, and Tamassia [31] showed that the converse holds true as well; for a valid orthogonal representation there exists a corresponding orthogonal drawing that

realizes this representation. Moreover, the proof is constructive and allows the efficient construction of such a drawing, a process that is referred to as *compaction*.

Altogether, this enables a three-step approach for computing orthogonal drawings, the so-called *Topology-Shape-Metrics Framework*, which works as follows. First, fix a *topology*, i.e., a combinatorial embedding of the graph in the plane (possibly planarizing it if it is non-planar); second, determine the *shape* of the drawing by constructing a valid orthogonal representation with few bends; and finally, compactify the orthogonal representation by assigning suitable vertex coordinates and edge lengths (*metrics*). As mentioned before, this reduces the problem of computing an orthogonal drawing of a planar graph with a fixed embedding to the purely combinatorial problem of finding a valid orthogonal representation, preferably with few bends. The task of actually creating a corresponding drawing in polynomial time is then taken over by the framework. It is this approach that is at the heart of a large body of literature on bend minimization algorithms for orthogonal drawings (e.g., [4, 9–13, 15]).

Contribution and Outline. In this paper we establish an analogous drawing framework for ortho-radial drawings. To this end, we introduce so-called ortho-radial representations, which give a combinatorial description of ortho-radial drawings, and therefore can be used to substitute orthogonal representations in the Topology-Shape-Metrics Framework.

More precisely, our contributions are as follows. We show that a natural generalization of the validity conditions (I) and (II) above is not sufficient, and introduce a third, less local condition that excludes so-called *strictly monotone cycles*, which do not admit an ortho-radial drawing. The third condition requires an additional edge, called the *reference edge*, from which we can derive whether an edge is part of a circle or a spoke of the ortho-radial grid. We prove that these three conditions together fully characterize ortho-radial drawings. Before that, characterizations for bend-free ortho-radial drawings were only known for paths, cycles and theta graphs [21]. Further, for the special case that each internal face is a combinatorial rectangle, a characterization for cubic graphs was known [22]. Here, an internal face is a combinatorial rectangle if it contains exactly four angles of $\pi/2$ and all other angles have π .

On the algorithmic side, we show that it can be tested in $\mathcal{O}(n^2)$ time whether an ortho-radial representation with a given reference edge is drawable. A corresponding drawing can be obtained in the same running time. While this does not yet directly allow us to compute ortho-radial drawings with few bends, our result paves the way for a purely combinatorial treatment of bend minimization in ortho-radial drawings, thus enabling the same type of tools that have proven highly successful in minimizing bends in orthogonal drawings. Recently, Niedermann and Rutter [27] presented such a tool based on an integer linear programming formulation showing that the topology-shape-metrics framework for ortho-radial drawings is capable of handling real-world networks such as metro systems.

We define basic terms and notations in Sect. 2. Afterwards, we formally introduce ortho-radial drawings and ortho-radial representations in Sect. 3, where we also establish basic properties that will be used throughout this paper. In Sect. 4 we introduce helpful tools for working with ortho-radial representations. Section 5 establishes basic properties of labelings that are used to describe ortho-radial representations. In Sects. 6 and 7 we prove that ortho-radial representations are drawable if and only if

they are valid. In Sect. 8 we give a validity test for ortho-radial representations that runs in $\mathcal{O}(n^2)$ time. Afterwards, in Sect. 9, we revisit the rectangulation procedure from Sect. 7 and show that using the techniques from Sect. 8 it can be implemented to run in $\mathcal{O}(n^2)$ time, improving over a naive application which would yield running time $\mathcal{O}(n^4)$. This enables a purely combinatorial treatment of ortho-radial drawings. In Sect. 10 we discuss how to include bends and we show that deciding whether a graph has a bend-free ortho-radial drawing is \mathcal{NP} -complete. We conclude with a summary and some open questions in Sect. 11.

2 Preliminaries

Let G be a plane graph with combinatorial embedding \mathcal{E} and outer face f_o . The embedding \mathcal{E} fixes for each vertex v of G the counterclockwise order of the edges incident to v around the vertex v . A *path* in G may contain vertices multiple times, and a *cycle* C may contain vertices multiple times but may not cross itself in the sense that the pairs of edges along which C enters and leaves a vertex v do not alternate in the cyclic order of edges around v in the embedding \mathcal{E} . We consider all paths and cycles to be directed. We represent a path P as the sequence $v_1 \dots v_k$ of its vertices in the order as they appear on P . Similarly, we represent a cycle as the sequence $v_1 \dots v_k$ of its vertices in the order as they appear on C , where v_1 is arbitrarily chosen. For any path $P = v_1 \dots v_k$ its *reverse* is $\overleftarrow{P} = v_k \dots v_1$. The concatenation of two paths P_1 and P_2 is written as $P_1 + P_2$. For two edges uv and wx on a path P the *subpath* from uv to wx is the unique path on P that starts with uv and ends with wx , and we denote it by $P[uv, wx]$. If P contains u (or x) only once, we may write u instead of uv (or x instead of wx). This can lead to abbreviations like $P[u, wx]$, $P[uv, x]$, and $P[u, x]$. In particular, if P is simple $P[u, x]$ denotes the subpath of P from u to x . For a cycle C , we similarly denote its reverse by \overrightarrow{C} , and for edges uv and wx on C the subpath of C from uv to wx in the direction of C is denoted by $C[uv, wx]$.

Moreover, a path or cycle is *simple* if it contains all vertices at most once. A *facial walk* C of a face f is a cycle in G that describes the boundary of f , i.e., the cycle C consists of edges of f and for any subpath uvw of C the edge uv precedes vw in the cyclic order of edges around v that is defined by \mathcal{E} . Any simple cycle C separates two sets of faces. One of these sets contains the outer face f_o , and we call these faces together with the vertices and edges incident to them the *exterior* of C . Conversely the faces of the other set and their incident vertices and edges form the *interior* of C . Note that C belongs to both its interior and its exterior. Unless specified explicitly, a simple cycle C is directed such that its interior lies to the right of C . Finally, a path P *respects* a cycle C if P lies in the exterior of C .

3 Ortho-Radial Drawings and Representations

Let $G = (V, E)$ be a planar, connected 4-graph with n vertices, where a graph is a 4-graph if it has maximum degree four. An *ortho-radial drawing* Δ of G is a plane drawing on an ortho-radial grid \mathcal{G} such that each vertex of G is a grid point of \mathcal{G} and

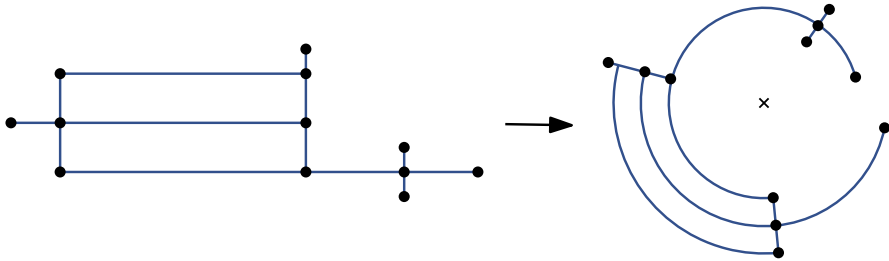


Fig. 3 The orthogonal drawing is transformed into an ortho-radial drawing by *bending* the horizontal edges into concentric circular arcs, while vertical edges become segments of rays that emanate from the center of the ortho-radial grid

each edge of G is a simple curve on \mathcal{G} . We observe that in any ortho-radial drawing there is an unbounded face f_o and a face f_c that contains the center of the grid; we call the former the *outer face* and the latter the *central face*; in our figures we mark the central face using a small “x”. All other faces are *regular*. We remark that f_c and f_o are not necessarily distinct. We further distinguish two types of simple cycles. If the central face lies in the interior of a simple cycle, the cycle is *essential* and otherwise *non-essential*.

In this paper, we assume that we are given G , a fixed combinatorial embedding \mathcal{E} of G and two (not necessarily distinct) faces f_c and f_o of \mathcal{E} . We seek an ortho-radial drawing Δ of G such that the combinatorial embedding of Δ is \mathcal{E} , the face f_c is the central face of Δ and f_o is the outer face of Δ . We call the tuple $I = (G, \mathcal{E}, f_c, f_o)$ an *instance* of ortho-radial drawing and Δ a drawing of I .

We observe that the definition of ortho-radial drawings allows edges to have bends, i.e., an edge may consist of a sequence of straight-line segments and circular arcs. In this paper, we focus on ortho-radial drawings without bends; we call such drawings *bend-free*. Hence, each edge is either part of a radial ray or of a concentric circle of \mathcal{G} . This is not a restriction as any ortho-radial drawing can be turned into a bend-free drawing by replacing bends with subdivision vertices; see Sect. 10 for a detailed discussion.

In a bend-free ortho-radial drawing of G each edge has a *geometric direction* in the sense that is drawn either clockwise, counterclockwise, towards the center, or away from the center. Hence, using the metaphor of a cylinder, the edges point *right*, *left*, *down*, or *up*, respectively. Moreover, *horizontal edges* point left or right, while *vertical edges* point up or down; see Fig. 1.

We further observe that if the central and outer face are identical then an ortho-radial drawing can be interpreted as a distorted orthogonal drawing, in which the horizontal edges are bent to circular arcs, while the vertical edges remain straight segments; see Fig. 3 for an example. Hence, utilizing the framework for orthogonal drawings by Tamassia [31], this allows us to easily create an ortho-radial drawing of an instance I for the case that the central face f_c and the outer face f_o are the same. Therefore, we assume $f_c \neq f_o$ in the remainder of this work, which changes the problem of finding an ortho-radial drawing of I substantially.

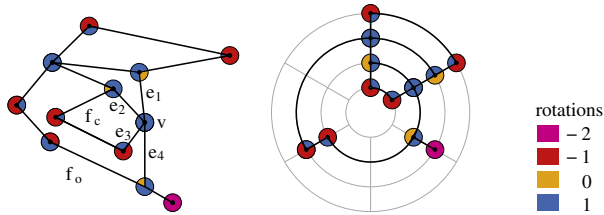


Fig. 4 A combinatorial embedding (left) and an ortho-radial drawing (right) of a graph. For each combinatorial angle a rotation is given. Refer to the online version of this article for colored figures

We first introduce concepts that help us to combinatorially describe the ortho-radial drawing Δ . Let v be a vertex of G and let $\mathcal{E}(v)$ be the counterclockwise order of the edges in \mathcal{E} around v . A *combinatorial angle* at v is a pair of edges (e_1, e_2) that are both incident to v and such that e_1 immediately precedes e_2 in $\mathcal{E}(v)$; see Fig. 4. An *angle assignment* Γ of an instance $I = (G, \mathcal{E}, f_c, f_o)$ assigns to each combinatorial angle (e_1, e_2) of \mathcal{E} a *rotation* $\text{rot}(e_1, e_2) \in \{-2, -1, 0, 1\}$. For an ortho-radial drawing Δ of I we can derive an angle assignment that defines $\text{rot}(e_1, e_2) = 2 - 2\alpha/\pi$ for each angle (e_1, e_2) at v , where α is the counterclockwise geometric angle between e_1 and e_2 in Δ . Hence, the rotation of a combinatorial angle counts the number of right turns that are taken when going from e_1 to e_2 via v , where negative numbers correspond to left turns; see Fig. 4. In particular, when traversing the boundary of a face such that the face lies to the right, then we perform two left turns at a degree-1 vertex. Hence, in case that $e_1 = e_2$ we derive $\text{rot}(e_1, e_2) = -2$ from Δ , i.e., v contributes two left turns. But conversely, we cannot derive an ortho-radial drawing from every angle assignment.

For a face f of \mathcal{E} with facial walk $v_1 \dots v_k$ around f (where f is oriented in clockwise order) we define $\text{rot}(f) = \sum_{i=1}^k \text{rot}(v_{i-1}v_i, v_i v_{i+1})$, where we define $v_0 := v_k$ and $v_{k+1} := v_1$. Every angle assignment Γ that is derived from a bend-free ortho-radial drawing is locally consistent in the following sense [21].

Definition 3.1 An angle assignment is *locally consistent* if it satisfies the following two conditions.

- (i) For each vertex, the sum of the rotations around v is $2(\text{deg}(v) - 2)$.
- (ii) For each face f , we have

$$\text{rot}(f) = \begin{cases} 4, & f \text{ is a regular face,} \\ 0, & f \text{ is the outer or the central face but not both,} \\ -4, & f \text{ is both the outer and the central face.} \end{cases}$$

Unlike for orthogonal representations, conditions (i) and (ii) do not guarantee that for a locally consistent angle assignment of I there is an ortho-radial drawing of I having the same angles; see Fig. 5. In this paper, we introduce a third more global condition that characterizes all locally consistent angle assignments of I that can be drawn. To that end, we first introduce basic concepts on rotations and directions in locally consistent angle assignments in Sect. 3.1, which leads us to ortho-radial representations. Afterwards, we define this global condition in Sect. 3.2.

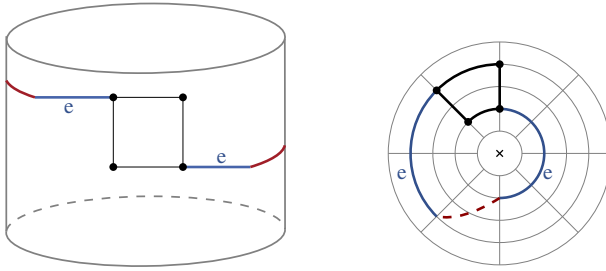


Fig. 5 Illustration of an ortho-radial representation where the angles around vertices sum up to 2π , and also the sum of angles for each face is as expected for an ortho-radial drawing. However, the graph does not have an ortho-radial drawing without bends as illustrated on a cylinder (left) and an ortho-radial grid (right)

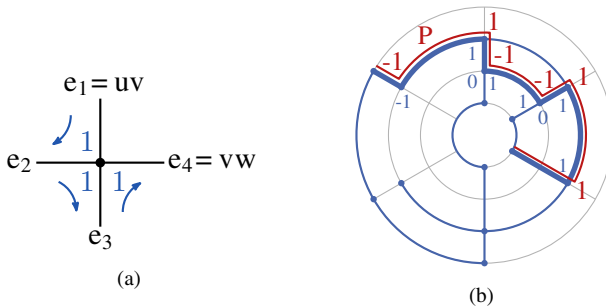


Fig. 6 Generalization of rotations. **a** The rotation of the two edges $e_1 = uv$ and $e_2 = vw$ is $\text{rot}(uvw) = \sum_{i=1}^3 1 - 2(4 - 2) = -1$. **b** The rotation of the path P is $\text{rot}(P) = -1 + 1 - 1 - 1 + 1 + 1 = 0$

3.1 Ortho-Radial Representations

We transfer two basic properties of ortho-radial drawings to locally consistent angle assignments. First, the rotations of all cycles are either 0 or 4. Second, fixing the geometric direction of a single edge e^* , fixes the geometric directions of all edges. We call e^* a *reference edge* and assume that it points to the right and lies on the outer face of \mathcal{E} .

Definition 3.2 An *ortho-radial representation* is a locally consistent angle assignment with a reference edge e^* on the outer face such that the outer face lies to the left of e^* .

For two edges $e = uv$ and $e' = vw$ we define the rotation between them as $\text{rot}(uvw) = \sum_{i=1}^{k-1} \text{rot}(e_i, e_{i+1}) - 2(k - 2)$, where $e = e_1, \dots, e_k = e'$ are the edges that are incident to v and lie between e and e' in counterclockwise order; see Fig. 6a.

The rotation of a path $P = v_1 \dots v_k$ is the sum of the rotations at its internal vertices, that is $\text{rot}(P) = \sum_{i=2}^{k-1} \text{rot}(v_{i-1}v_i v_{i+1})$; see Fig. 6(b).

Observation 3.3 Let P be a path with start vertex s and end vertex t .

- It is $\text{rot}(\overline{P}) = -\text{rot}(P)$.
- For every edge e on P it is $\text{rot}(P) = \text{rot}(P[s, e]) + \text{rot}(P[e, t])$.

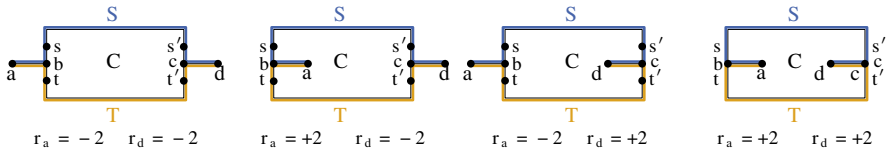


Fig. 7 Illustration for proof of Lemma 3.5. Cycle C with two additional edges ab and cd . If a lies in the exterior of C then $r_a = +2$ and otherwise $r_a = -2$. Similarly, if d lies in the exterior of C then $r_d = +2$ and $r_d = -2$ otherwise

Similarly, for a cycle $C = v_1 \dots v_k$, its rotation is the sum of the rotations at all its vertices (where we define $v_0 = v_k$ and $v_{k+1} = v_1$), i.e., $\text{rot}(C) = \sum_{i=1}^k \text{rot}(v_{i-1}v_i v_{i+1})$. We observe that the rotation of a face f is equal to the rotation of the cycle that we obtain from the facial walk around f .

Lemma 3.4 *Let C be a simple cycle in an ortho-radial representation Γ . Then, $\text{rot}(C) = 0$ if C is essential and $\text{rot}(C) = 4$ if C is non-essential.*

Proof Let H be the sub-graph of G that is contained in the interior of C ; we note that C belongs to H . Let n_H , m_H , and f_H denote the number of vertices, edges and faces of H , respectively. Let Γ' be the angle assignment that coincides with Γ on all angles in the interior of C and has $\text{rot}_{\Gamma'}(uvw) = \text{rot}_{\Gamma}(uvw)$ for any three consecutive vertices along C . By construction, Γ' satisfies property (i) and therefore the sum R of all rotations is $R = \sum_{v \in V} 2(\text{deg}(v) - 2) = 4(m_H - n_H) = 4(f_H - 2)$ by Euler’s formula.

If C is essential, then $f_H - 2$ interior faces contribute a rotation of 4 each and the central face contributes a rotation of 0, leaving a rotation of 0 for the outer face. If C is non-essential, then the $f_H - 1$ interior faces contribute a rotation of 4 each, leaving a rotation of -4 for the outer face. Since C is oriented clockwise while the outer face is oriented counterclockwise, $\text{rot}(C) = 4$ follows. \square

The next lemma relates the rotations of two paths S and T that use the same edges except on a cycle C ; see Fig. 7.

Lemma 3.5 *Let C be a cycle and let ab and cd be two edges (with $b \neq c$) such that b and c lie on C , but a and d do not. Further, let $S = ab + C[b, c] + cd$ and $T = ab + \overline{C}[b, c] + cd$. Then*

$$\text{rot}(S) - \text{rot}(T) = \text{rot}(C) + r_a + r_d,$$

where for $z \in \{a, d\}$ we define $r_z = +2$ if z lies in the interior of C and $r_z = -2$ if z lies in the exterior of C .

Proof Let t and s be the vertices immediately before and after b on C . We observe that a lies in the interior of C if and only if ab lies locally to the right of the path sbt . Considering all six cases how the edges ab , sb , and tb can be arranged (see Fig. 8), we obtain $\text{rot}(abs) - \text{rot}(abt) = \text{rot}(tbs) + r_a$. Similarly, we define s' and t' as the vertices immediately before and after c on C . Considering all cases as above we get

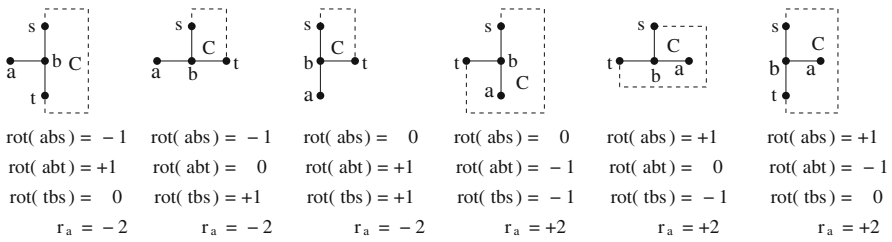


Fig. 8 Illustration for proof of Lemma 3.5. In all six cases it holds $\text{rot}(abs) - \text{rot}(abt) = \text{rot}(tbs) + r_a$

$\text{rot}(s'cd) - \text{rot}(t'cd) = \text{rot}(s'ct') + r_d$. Splitting S and T into three parts (see Fig. 7), we have

$$\begin{aligned} \text{rot}(S) - \text{rot}(T) &= \text{rot}(abs) + \text{rot}(C[b, c]) + \text{rot}(s'cd) - \text{rot}(abt) \\ &\quad - \text{rot}(\overline{C}[b, c]) - \text{rot}(t'cd). \end{aligned}$$

Combining the rotations at b and c using the observations from above, we get

$$\begin{aligned} \text{rot}(S) - \text{rot}(T) &= \text{rot}(tbs) + r_a + \text{rot}(C[b, c]) + \text{rot}(C[c, b]) + \text{rot}(s'ct') + r_d \\ &= \text{rot}(C) + r_a + r_d. \end{aligned}$$

□

For two edges e and e' let P be an arbitrary simple path that starts at the source or target of e and ends at the source or target of e' , and that neither contains e nor e' . We call P a *reference path* from e to e' . We define the *combinatorial direction* of $e' = xy$ with respect to $e = uv$ and P as

$$\text{dir}(e, P, e') = \begin{cases} \text{rot}(e + P + e') & P \text{ starts at } v \text{ and ends at } x, \\ \text{rot}(\overline{e} + P + e') + 2 & P \text{ starts at } u \text{ and ends at } x, \\ \text{rot}(e + P + \overline{e}') - 2 & P \text{ starts at } v \text{ and ends at } y, \\ \text{rot}(\overline{e} + P + \overline{e}') & P \text{ starts at } u \text{ and ends at } y. \end{cases}$$

With the fixed direction of the reference edge e^* , it is natural to determine the direction of any other edge e by considering the direction of any reference path from e^* to e . In order to get consistent results, any two reference paths P and Q from e^* to e must induce the same direction of e , which means that $\text{dir}(e^*, P, e)$ and $\text{dir}(e^*, Q, e)$ may only differ by a multiple of 4. In the following lemma we show that this is indeed the case.

Lemma 3.6 *Let e and e' be two edges of an ortho-radial representation Γ , and let P and Q be two reference paths from e to e' .*

- It holds $\text{dir}(e, P, e') \equiv \text{dir}(e, Q, e') \pmod{4}$.
- It holds $\text{dir}(e, P, e') = \text{dir}(e, Q, e')$, if there are two essential cycles C and C' such that

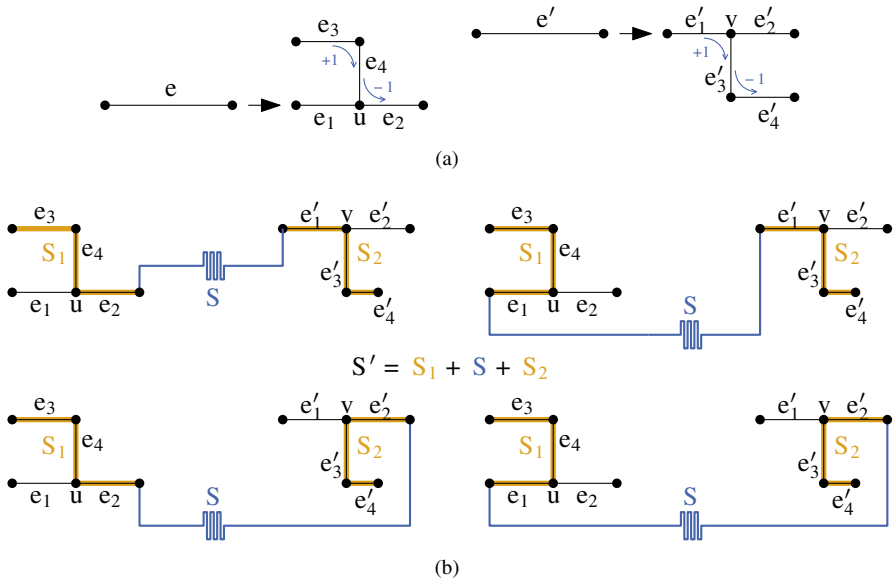


Fig. 9 Illustration of the proof for Lemma 3.6. **a** The edges e and e' are replaced by the depicted construction to reduce the number of cases to be considered. **b** Four cases are considered how the edges e and e' can be connected by the path S . The paths S_1 and S_2 are defined depending on the particular case

- (a) C' lies in the interior of C ,
- (b) e lies on C and e' lies on C' , and
- (c) P and Q lie in the interior of C and in the exterior of C' .

Proof First, we define a construction that helps us to reduce the number of cases to be considered. We subdivide e by a vertex u into two edges e_1 and e_2 ; see Fig. 9a. Further, we add a path consisting of two edges e_3 and e_4 such that the target of e_4 is u . We define that $\text{rot}(e_3e_4) = 1$ and $\text{rot}(e_4e_2) = -1$. Similarly, we subdivide e' by a vertex v into two edges e'_1 and e'_2 . Further, we add a path consisting of two edges e'_3 and e'_4 such that the source of e'_3 is v . We define that $\text{rot}(e'_1e'_3) = 1$ and $\text{rot}(e'_3e'_4) = -1$. Let S be a reference path from e to e' ; see Fig. 9b. Let S_1 be the path that starts at the source of e_3 and ends at the starting point of S only using edges from $\{e_1, e_2, e_3, e_4\}$. Further, let S_2 be the path that starts at the end point of S and ends at the target of e'_4 only using edges from $\{e'_1, e'_2, e'_3, e'_4\}$. The *extension* S' of S is the path $S_1 + S + S_2$. The following claim shows that we can consider S' instead of S such that it is sufficient to consider the rotation of S' instead of the direction $\text{dir}(e, S, e')$, which distinguishes four cases.

Claim 3.7 $\text{dir}(e, S, e') = \text{rot}(S')$.

The detailed proof of Claim 3.7 is found at the end of this proof. In the following let P' and Q' be the extensions of P and Q , respectively. We show that $\text{rot}(P') \equiv \text{rot}(Q') \pmod{4}$. Moreover, for the case that e coincides with the reference edge e^* and e' lies on an essential cycle that is respected by P and Q , we show that $\text{rot}(P') = \text{rot}(Q')$.

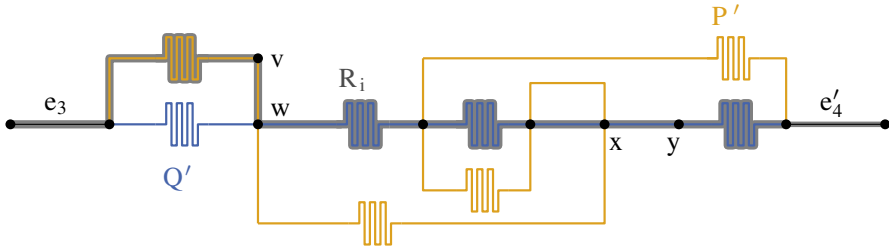


Fig. 10 Illustration for proof of Lemma 3.6. The path R_i (gray) consists of a prefix of the path P' (orange) and a suffix of the path Q' (blue)

Altogether, due to Claim 3.7 this proves Lemma 3.6. We show $\text{rot}(P') = \text{rot}(Q')$ by converting P' into Q' successively. More precisely, we construct paths $R_1 \dots R_k$ such that R_i consists of a prefix of P' followed by a suffix of Q' such that with increasing i the used prefix of P' becomes longer, while the used suffix of Q' becomes shorter. In particular, we have $R_1 = Q'$ and $R_k = P'$. We show that $\text{rot}(R_i) \equiv \text{rot}(R_{i+1}) \pmod{4}$. If $R_i \neq P'$, we construct R_{i+1} from R_i as follows; see Fig. 10.

There is a first edge vw on P' such that the following edge does not lie on R_i . Let x be the first vertex on P' after w that lies on R_i and let y be the vertex on R_i that follows x immediately. As both P' and R_i end at the same edge, these vertices always exist. We define $R_{i+1} = P'[e_3, x] + R_i[x, e_4]$. We observe that $R_i[x, e_4] = Q'[x, e_4]$, as $R_i[w, e_4] = Q'[w, e_4]$ and x occurs on Q' after w . Further, R_{i+1} is a path as we can argue as follows. We can decompose R_{i+1} into three paths: $R_{i+1}^1 = P'[e_3, w] = R_i[e_3, w]$, $R_{i+1}^2 = P'[w, x]$ and $R_{i+1}^3 = Q'[x, e_4] = R_i[x, e_4]$. The paths R_{i+1}^1 and R_{i+1}^3 do not intersect as both also belong to R_i , which is a path by induction. The paths R_{i+1}^1 and R_{i+1}^2 do not intersect (except at their common vertex w), because both belong to P' . The paths R_{i+1}^2 and R_{i+1}^3 do not intersect (except at their common vertex x), because by the definition of x no vertex of P' between w and x lies on R_i .

Next, we show that $\text{rot}(R_i) \equiv \text{rot}(R_{i+1}) \pmod{4}$. To that end consider the (initially undirected) cycle C_i that consists of the two paths $R_i[w, x]$ and $R_{i+1}[w, x]$. We orient C_i such that the interior of the cycle locally lies to the right of its edges. By the definition of R_{i+1} we obtain $\text{rot}(R_i[e_3, vw]) = \text{rot}(R_{i+1}[e_3, vw])$ and $\text{rot}(R_i[xy, e_4]) = \text{rot}(R_{i+1}[xy, e_4])$, as these subpaths of R_i and R_{i+1} coincide, respectively. Hence, it remains to show that $\text{rot}(R_i[vw, xy]) \equiv \text{rot}(R_{i+1}[vw, xy]) \pmod{4}$ and $\text{rot}(R_i[vw, xy]) = \text{rot}(R_{i+1}[vw, xy])$ in the special case that $e = e^*$ and e' lies on an essential cycle that is respected by P and Q .

In general we can describe the obtained situation as follows. We are given a cycle and two edges ab and cd (with $b \neq c$) such that b and c lie on that cycle, but a and d not; see also Fig. 7. Hence, we can apply Lemma 3.5.

We distinguish two cases: if the interior of C_i lies locally to the right of R_i we define $S = vw + R_i[w, x] + xy$ and $T = vw + R_{i+1}[w, x] + xy$, and otherwise $S = vw + R_{i+1}[w, x] + xy$ and $T = vw + R_i[w, x] + xy$. We only consider the first case, as the other case is symmetric. By Lemma 3.5 we obtain $\text{rot}(S) - \text{rot}(T) = \text{rot}(C_i) + r_v + r_y$. As $\text{rot}(C_i) \equiv 0 \pmod{4}$ and $r_v, r_y \equiv 2 \pmod{4}$ we obtain

$\text{rot}(S) \equiv \text{rot}(T) \pmod{4}$ and with this $\text{rot}(R_i[vw, xy]) \equiv \text{rot}(R_{i+1}[vw, xy]) \pmod{4}$. Altogether, in the general case we obtain $\text{rot}(R_i) \equiv \text{rot}(R_{i+1})$.

Finally, we prove the second statement of the lemma. Hence, there are two essential cycles C and C' such that

- C' lies in the interior of C ,
- e lies on C and e' lies on C' , and
- P and Q lie in the interior of C and in the exterior of C' .

In particular, the paths P and Q respect C' as they lie in the exterior of C' . We show that $\text{rot}(S) = \text{rot}(T)$. First, we observe that $R_i[e, vw]$ respects the cycle C_i by the simplicity of R_i and R_{i+1} . In particular, vw lies in the exterior of C_i so that $r_v = -2$. We distinguish the two cases whether C_i is essential or non-essential. If C_i is a non-essential cycle, the edge xy is also contained in the exterior of C_i as P and Q end on the essential cycle C' , which both respect. Hence, we obtain $r_y = -2$. Thus, we get by Lemmas 3.4 and 3.5 that $\text{rot}(S) - \text{rot}(T) = \text{rot}(C_i) + r_v + r_y = 4 - 2 - 2 = 0$. If C_i is an essential cycle, then the cycle C' is contained in the interior of C_i as both contain the central face, and C_i is composed by parts of paths that respect C' . Consequently, the edge xy lies in the interior of C_i so that we obtain $r_y = 2$. By Lemmas 3.4 and 3.5 we get $\text{rot}(S) - \text{rot}(T) = \text{rot}(C_i) + r_v + r_y = 0 - 2 + 2 = 0$. It remains to prove Claim 3.7. □

Proof of Claim 3.7 We distinguish the four cases given by the definition of $\text{dir}(e, S, e')$; see Fig. 9b. If S starts at the target of e and ends at the source of e' , we obtain $\text{dir}(e, S, e') = \text{rot}(e + S + e') = \text{rot}(e_2 + S + e'_1)$ as subdividing e and e' transfers the directions of e and e' to e_2 and e'_1 , respectively. Hence, we obtain

$$\begin{aligned} \text{rot}(S') &= \text{rot}(S_1 + S + S_2) \\ &= \text{rot}(e_3 + e_4 + e_2 + S + e'_1 + e'_3 + e'_4) \\ &= \text{rot}(e_3 + e_4 + e_2) + \text{rot}(e_2 + S + e'_1) + \text{rot}(e'_1 + e'_3 + e'_4) \\ &= 0 + \text{dir}(e, S, e') + 0 = \text{dir}(e, S, e'). \end{aligned}$$

If S starts at the source of e and ends at the source of e' , we obtain $\text{dir}(e, S, e') = \text{rot}(\overline{e} + S + e') + 2 = \text{rot}(\overline{e}_1 + S + e'_1) + 2$ and with this we obtain

$$\begin{aligned} \text{rot}(S') &= \text{rot}(S_1 + S + S_2) \\ &= \text{rot}(e_3 + e_4 + \overline{e}_1 + S + e'_1 + e'_3 + e'_4) \\ &= \text{rot}(e_3 + e_4 + \overline{e}_1) + \text{rot}(\overline{e}_1 + S + e'_1) + \text{rot}(e'_1 + e'_3 + e'_4) \\ &= 2 + \text{dir}(e, S, e') - 2 + 0 = \text{dir}(e, S, e'). \end{aligned}$$

If S starts at the target of e and ends at the target of e' , we obtain $\text{dir}(e, S, e') = \text{rot}(e + S + \overline{e'}) - 2 = \text{rot}(e_2 + S + \overline{e}'_2) - 2$ and with this we obtain

$$\begin{aligned} \text{rot}(S') &= \text{rot}(S_1 + S + S_2) \\ &= \text{rot}(e_3 + e_4 + e_2 + S + \overline{e}'_2 + e'_3 + e'_4) \end{aligned}$$

$$\begin{aligned}
 &= \text{rot}(e_3 + e_4 + e_2) + \text{rot}(e_2 + S + \overline{e_2}) + \text{rot}(\overline{e_2} + e'_3 + e'_4) \\
 &= 0 + \text{dir}(e, S, e') + 2 - 2 = \text{dir}(e, S, e').
 \end{aligned}$$

If S starts at the source of e and ends at the target of e' , we obtain $\text{dir}(e, S, e') = \text{rot}(\overline{e} + S + \overline{e'}) = \text{rot}(\overline{e_1} + S + \overline{e'_2})$ and with this we obtain

$$\begin{aligned}
 \text{rot}(S') &= \text{rot}(S_1 + S + S_2) \\
 &= \text{rot}(e_3 + e_4 + \overline{e_1} + S + \overline{e'_2} + e'_3 + e'_4) \\
 &= \text{rot}(e_3 + e_4 + \overline{e_1}) + \text{rot}(\overline{e_1} + S + \overline{e'_2}) + \text{rot}(\overline{e'_2} + e'_3 + e'_4) \\
 &= 2 + \text{dir}(e, S, e') - 2 = \text{dir}(e, S, e').
 \end{aligned}$$

Altogether, this shows the claim $\text{dir}(e, S, e') = \text{rot}(S')$. □

Corollary 3.8 *If e is the reference edge e^* and e' lies on an essential cycle that is respected by P and Q , then $\text{dir}(e, P, e') = \text{dir}(e, Q, e')$.*

Proof The statement directly follows from the second statement of Lemma 3.6 by assuming that C is the outermost essential cycle of Γ and C' is the cycle containing e' . □

Using this result, the geometric directions of all edges of a given ortho-radial representation Γ can be determined as follows. Let P be any reference path from the reference edge e^* to any edge e , the edge e points right, down, left, and up if $\text{dir}(e^*, P, e)$ is congruent to 0, 1, 2, and 3, respectively. Edges that point up or down are *vertical* and edges that point left or right are *horizontal*. Lemma 3.6 ensures that the result is independent of the choice of the reference path. In fact, Lemma 3.6 even gives a stronger result as we can infer the geometric direction of one edge from the geometric direction of another edge locally without having to resort to paths to the reference edge. We often implicitly make use of this observation in our proofs.

3.2 Drawable Ortho-Radial Representations

In this section we introduce concepts that help us to characterize the ortho-radial representations that have an ortho-radial drawing. To that end, consider an arbitrary bend-free ortho-radial drawing Δ of a plane 4-graph G . As we assume throughout this work that the outer and central face are not the same, there is an essential cycle C that lies on the outer face f of Δ . Let e be a horizontal edge of C that points to the right and that lies on the outermost circle of the ortho-radial grid among all such edges, and let R_Δ be the set of all edges e' with $\text{dir}(e, P, e') = 0$ for any path P on f ; see Fig. 11. We observe that R_Δ is independent of the choice of e , because if there are multiple choices for e , then all of these edges are contained in R_Δ . We call the edges in R_Δ the *outlying edges* of Δ . Note that not all outlying edges lie on the outermost circle of Δ ; consider for example the edge e' in Fig. 11. Without loss of generality, we require that the reference edge of the according ortho-radial representation Γ stems from R_Δ ; as R_Δ is not empty and all of the edges in R_Δ are possible candidates for

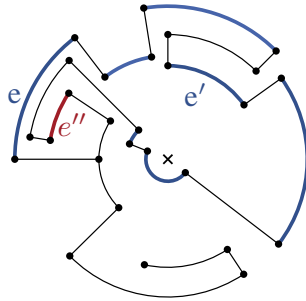


Fig. 11 The outer face of an ortho-radial drawing. All outlying edges are marked blue

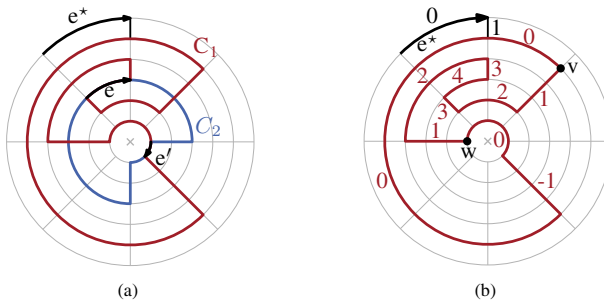


Fig. 12 a Two cycles C_1 and C_2 may have both common edges with different labels ($\ell_{C_1}(e) = 4 \neq 0 = \ell_{C_2}(e)$) and ones with equal labels ($\ell_{C_1}(e') = \ell_{C_2}(e') = 0$). **b** All labels of $C_1[v, w]$ are positive, implying that C_1 goes down. Note that not all edges of $C_1[v, w]$ point downwards

being a reference edge, we can always change the reference edge of Γ to one of the edges in R_Δ .

An ortho-radial representation Γ of a graph G with reference edge e^* is *drawable* if there exists a bend-free ortho-radial drawing Δ of G embedded as specified by Γ such that the corresponding angles in Δ and Γ are equal and the edge e^* is an outlying edge, i.e., $e^* \in R_\Delta$. Note that it is not required that e^* lies on the outermost circle of Δ . For example in Fig. 11 the edge e' is a suitable reference edge, but there is no drawing where e' lies on the outermost circle. Rather, the requirement is that every path along the outer face from e' to an edge on the outermost circle has rotation 0. Unlike for orthogonal representations, conditions (i) and (ii) do not guarantee that the ortho-radial representation is drawable; see Fig. 5. Therefore, we introduce a third condition, which is formulated in terms of labelings of essential cycles.

Let e be an edge on an essential cycle C in G and let P be a reference path from the reference edge e^* to e that respects C . We define the *label* of e on C as $\ell_C(e) = \text{dir}(e^*, P, e)$. By Corollary 3.8 the label $\ell_C(e)$ of e does not depend on the choice of P . However, if an edge e is contained in two essential cycles C_1 and C_2 , then their labels may generally differ, i.e., $\ell_{C_1}(e) \neq \ell_{C_2}(e)$. In fact, Fig. 12 shows that two cycles C_1, C_2 may share two edges e, e' such that $\ell_{C_1}(e) \neq \ell_{C_2}(e)$ and $\ell_{C_1}(e') = \ell_{C_2}(e')$.

We call the set of all labels of an essential cycle its *labeling*.

We call an essential cycle *monotone* if either all its labels are non-negative or all its labels are non-positive. A monotone cycle is a *decreasing* cycle if it has at least one strictly positive label, and it is an *increasing* cycle if it has at least one strictly negative label. We also refer to increasing and decreasing cycles as *strictly monotone*. An ortho-radial representation is *valid* if it contains no strictly monotone cycles. The validity of an ortho-radial representation ensures that on each essential cycle with at least one non-zero label there is at least one edge pointing up and one pointing down. We note that whether an essential cycle is monotone crucially depends on the reference edge. In Fig. 11 the cycle bounding the outer face is not monotone if e or e' is the reference edge, but it is strictly monotone if e'' is the reference edge.

A main goal of this paper is to show that a graph with a given ortho-radial representation can be drawn if and only if the representation is valid. Further, we show how to test validity of a given representation and how to obtain a bend-free ortho-radial drawing from a valid ortho-radial representation in quadratic time. Altogether, this yields our main results:

Theorem 3.9 *An ortho-radial representation is drawable if and only if it is valid.*

Theorem 3.10 *Given an ortho-radial representation Γ , it can be determined in $\mathcal{O}(n^2)$ time whether Γ is valid. In the negative case a strictly monotone cycle can be computed in $\mathcal{O}(n^2)$ time.*

Theorem 3.11 *Given a valid ortho-radial representation, a corresponding drawing can be constructed in $\mathcal{O}(n^2)$ time.*

We prove the three theorems in the given order. Sections 4–6 deal with Theorem 3.9. In Sect. 7 we prove Theorem 3.10. In particular, together with the proof of Theorem 3.9 this already leads to a version of Theorem 3.11, but with a running time of $\mathcal{O}(n^4)$. In Sect. 9 we show how to achieve $\mathcal{O}(n^2)$ running time proving Theorem 3.11.

4 Transformations of Ortho-Radial Representations

In this section we introduce helpful tools that we use throughout this work. In the remainder of this work we assume that we are given an ortho-radial representation Γ with reference edge e^* .

Since the reference edge lies on an essential cycle by definition, we can compute the labelings of essential cycles via the rotation of paths as shown in the following lemma. This simplifies the arguments of our proofs.

Lemma 4.1 *For every edge e on an essential cycle C there is a reference path P from e^* to e such that P respects C , starts at the target of e^* and ends on the source of e . Moreover, $\ell_C(e) = \text{rot}(e^* + P + e)$.*

Proof Let Q be a reference path from the reference edge $e^* = rs$ to $e = tu$ respecting C . We construct the desired reference path P as follows. Let C' be the outermost essential cycle, i.e., C' is the unique essential cycle such that every edge of C' bounds the outer face. If C and C' have a common vertex, we define v to be the first common

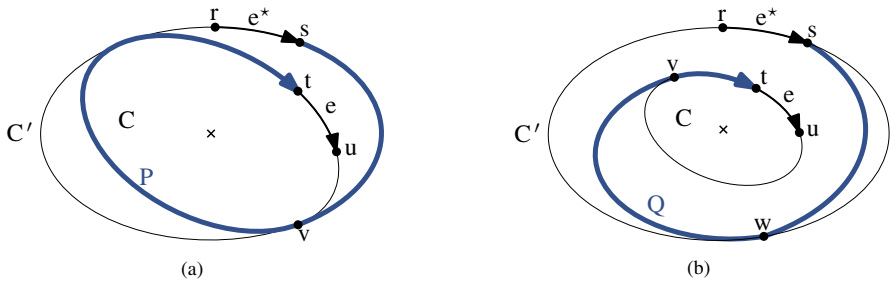


Fig. 13 Illustration of proof for Lemma 4.1. The outermost cycle C' contains the essential cycle C . There is always a reference path from s on C' to t on C respecting C . **a** The cycles C and C' have vertices in common. **b** The cycles C and C' are disjoint

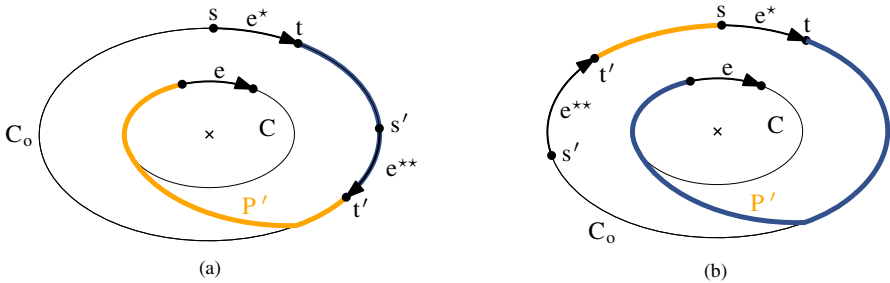


Fig. 14 Illustration of proof for Lemma 4.2. The outermost cycle C_o contains the essential cycle C . Further, there is a reference path P from e^* to an edge e on C . **a** The path P contains the edge e^{**} . **b** The path P does not contain the edge e^{**}

vertex on C' after s ; see Fig. 13a. We set $P = C'[s, v] + C[v, t]$. By the choice of v this concatenation is a path. Moreover, it is a reference path from e^* to e that respects C . If C and C' are disjoint, let v be the first vertex of Q lying on C and let w be the last vertex of Q before v that lies on C' ; see Fig. 13b. We set $P = C'[s, w] + Q[w, v] + C[v, t]$. We observe that the concatenation P is a path by the choice of v and w , and that P is a reference path from e^* to e respecting C .

By Corollary 3.8 we have $\ell_C(e) = \text{dir}(e^*, P, e)$ since P is a reference path from e^* to e respecting C . Further, since P starts at the target s of the reference edge e^* and ends at the source t of the edge e , we can express the direction of P as $\text{dir}(e^*, P, e) = \text{rot}(e^* + P + e)$. \square

The next lemma shows how we can change the reference edge e^* on the unique essential cycle C_o whose edges are all incident to the outer face; Fig. 14 illustrates the lemma.

Lemma 4.2 *Let Γ be an ortho-radial representation and let e^* be the reference edge of Γ . Further, let C_o be the essential cycle that lies on the outer face and let e^{**} be an edge on C_o such that $\text{rot}(C_o[e^*, e^{**}]) = 0$. For every edge e on an essential cycle C of Γ it holds $\ell_C(e) = \bar{\ell}_C(e)$, where $\bar{\ell}_C$ is the labeling of C with respect to e^{**} . In particular, Γ with reference edge e^* is valid if and only if Γ with reference edge e^{**} is valid.*

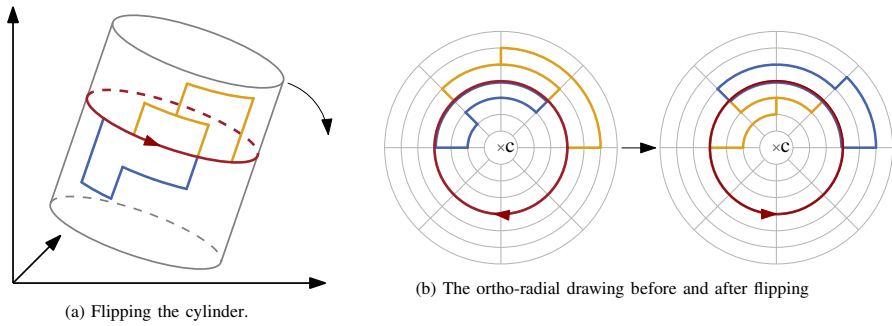


Fig. 15 Illustration of flipping the cylinder

Proof Let e be an edge of an essential cycle C in Γ and let P be a reference path from e^* to e . By Lemma 4.1 we assume that P starts at the target of e^* and ends at the source of e . Without loss of generality we assume that only a prefix of P is part of C_o . Further, let $e^* = st$ and $e^{**} = s't'$.

If P contains e^{**} , then

$$\ell_C(e) = \text{rot}(e^* + P + e) = \text{rot}(e^* + P[t, e^{**}]) + \text{rot}(e^{**} + P' + e),$$

where P' is the suffix of P that starts at t' ; see Fig. 14a. By assumption $\text{rot}(e^* + P[t, e^{**}]) = 0$ and $\bar{\ell}_C(e) = \text{rot}(e^{**} + P' + e)$.

If P does not contain e^{**} , then $C_o[t', e^*] + P$ is a reference path of e with respect to e^{**} , which contains e^* ; see Fig. 14b. Swapping e^* and e^{**} in the argument of the previous case yields the claim. Note that $\text{rot}(C_o[e^{**}, e^*]) = 0$ as $\text{rot}(C_o) = 0$. \square

In our arguments we frequently exploit certain symmetries. For an ortho-radial representation Γ we introduce two new ortho-radial representations, its *flip* $\text{flip}(\Gamma)$ and its *mirror* $\text{mirror}(\Gamma)$. Geometrically, viewed as a drawing on a cylinder, a flip corresponds to rotating the cylinder by an angle of π around a line perpendicular to the axis of the cylinder so that is upside down, see Fig. 15, whereas mirroring corresponds to mirroring it at a plane that is parallel to the axis of the cylinder; see Fig. 16. Intuitively, the first transformation exchanges left/right and top/bottom, and thus preserves monotonicity of cycles, while the second transformation exchanges left/right but not top/bottom, and thus maps increasing cycles to decreasing ones and vice versa. This intuition indeed holds with the correct definitions of $\text{flip}(\Gamma)$ and $\text{mirror}(\Gamma)$, but due to the non-locality of the validity condition for ortho-radial representations and the dependence on a reference edge this requires some care. The following two lemmas formalize flipped and mirrored ortho-radial representations. We denote the reverse of a directed edge e by \bar{e} .

To define $\text{flip}(\Gamma)$ it is necessary that the cycle bounding the central face f_c contains an edge e^{**} with $\ell_{f_c}(e^{**}) = 0$. In particular, such an edge exists if the cycle bounding the central face is not strictly monotone. We choose \bar{e}^{**} as the reference edge of $\text{flip}(\Gamma)$. The central face of Γ becomes the outer face of $\text{flip}(\Gamma)$ and the outer face of Γ becomes

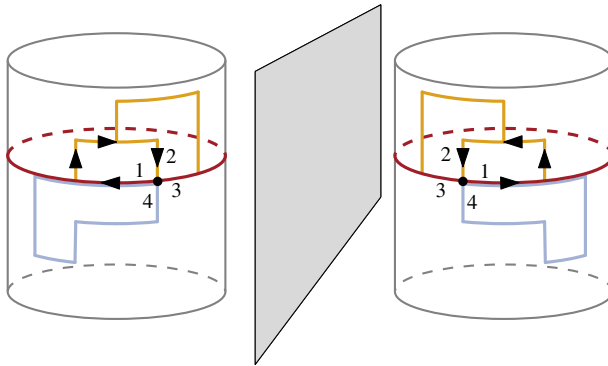


Fig. 16 Mirroring the cylinder

the central face of $\text{flip}(\Gamma)$. All other information of Γ is transferred to $\text{flip}(\Gamma)$ without modification.

Lemma 4.3 (Flipping) *Let Γ be an ortho-radial representation with outer face f_o and central face f_c . If the cycle bounding the central face is not strictly monotone, then $\text{flip}(\Gamma)$ is an ortho-radial representation such that*

- (i) \bar{f}_c is the outer face of $\text{flip}(\Gamma)$ and \bar{f}_o is the central face of $\text{flip}(\Gamma)$,
- (ii) $\ell_{\bar{C}}^f(\bar{e}) = \ell_C(e)$ for all essential cycles C and edges e on C , where ℓ^f is the labeling in $\bar{\Gamma}$.

In particular, increasing and decreasing cycles of Γ correspond to increasing and decreasing cycles of $\text{flip}(\Gamma)$, respectively.

Proof As the local structure is unchanged, $\text{flip}(\Gamma)$ is an ortho-radial representation. The essential cycles in Γ bijectively correspond to the essential cycles in $\text{flip}(\Gamma)$ by reversing the direction of the cycles. That is, any essential cycle C in Γ corresponds to the cycle \bar{C} in $\text{flip}(\Gamma)$. Note that the reversal is necessary since we always consider essential cycles to be directed such that the center lies in its interior, which is defined as the area locally to the right of the cycle.

Consider any essential cycle C in Γ . We denote the labeling of C in Γ by ℓ_C and the labeling of \bar{C} in $\text{flip}(\Gamma)$ by $\ell_{\bar{C}}^f$. We show that for any edge e on C , $\ell_C(e) = \ell_{\bar{C}}^f(\bar{e})$, which in particular implies that any monotone cycle in Γ corresponds to a monotone cycle in $\bar{\Gamma}$ and vice versa. By Lemma 4.1 there is a reference path P from the target of e^* to the source of the edge e respecting C . Similarly, there is a path Q from the target of e to the source of e^{**} that lies in the interior of C . The path \bar{Q} is a reference path from e^{**} to \bar{e} respecting \bar{C} in $\text{flip}(\Gamma)$.

Assume for now that $P + e + Q$ is simple. We shall see at the end how the proof can be extended if this is not the case. By the choice of e^{**} , we have

$$0 = \ell_{f_c}(e^{**}) = \text{rot}(e^* + P + e + Q + e^{**}) = \text{rot}(e^* + P + e) + \text{rot}(e + Q + e^{**}). \tag{1}$$

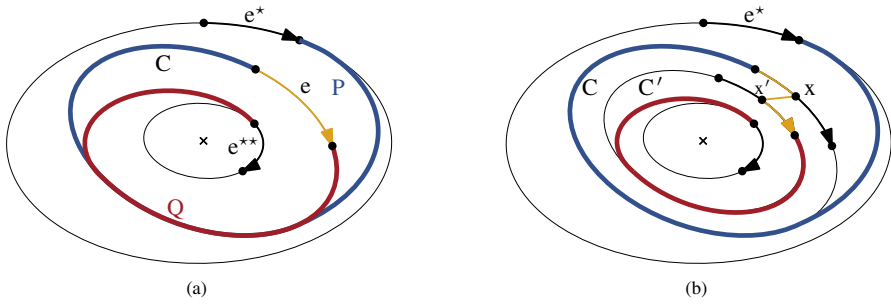


Fig. 17 Illustration of proof for Lemma 4.3. **a** The path P is a reference path from e^* to e in Γ respecting C and the path Q is a reference path from e^{**} to \bar{e} in $\text{flip}(\Gamma)$ respecting \bar{C} . **b** The edge e is subdivided by a vertex x . Afterwards, G is cut at C such that the interior and the exterior get their own copies of C . The copies are connected by an edge between x and x'

Hence, $\text{rot}(e^* + P + e) = -\text{rot}(e + Q + e^{**}) = \text{rot}(e^{**} + \bar{Q} + \bar{e})$ and in total

$$\ell_C^f(\bar{e}) = \text{rot}(e^{**} + \bar{Q} + \bar{e}) = \text{rot}(e^* + P + e) = \ell_C(e). \tag{2}$$

Thus, any monotone cycle in Γ corresponds to a monotone cycle in $\text{flip}(\Gamma)$ and vice versa.

If $P + e + Q$ is not simple, we make it simple by cutting G at C such that the interior and the exterior of C get their own copies of C ; see Fig. 17. We connect the two parts by an edge between two new vertices x and x' on the two copies of e , which we denote by vw in the exterior part and $v'w'$ in the interior part. The new edge is placed perpendicular to these copies. The path $P + vx'w' + Q$ is simple and its rotation is 0. Hence, the argument above implies $\ell_C^f(\bar{e}) = \ell_C(e)$. \square

We define $\text{mirror}(\Gamma)$ as follows. We reverse the direction of all faces and reverse the order of the edges around each vertex. The outer and central face are equal to those in Γ (except for the directions) and the reference edge is \bar{e}^* . By this definition (e_1, e_2) is a combinatorial angle in Γ if and only if (\bar{e}_2, \bar{e}_1) is a combinatorial angle in $\text{mirror}(\Gamma)$. We define $\text{rot}_{\text{mirror}(\Gamma)}(\bar{e}_2, \bar{e}_1) = \text{rot}_\Gamma(e_1, e_2)$, where the subscript indicates the ortho-radial representation that defines the rotation. Thus, edges that point left in Γ point right in $\text{mirror}(\Gamma)$ and vice versa, but the edges that point up (down) in Γ also point up (down) in $\text{mirror}(\Gamma)$.

Lemma 4.4 (Mirroring) *Let Γ be an ortho-radial representation with outer face f_o and central face f_c . Then, $\text{mirror}(\Gamma)$ is an ortho-radial representation such that*

- (i) \bar{f}_o is the outer face of $\text{mirror}(\Gamma)$ and \bar{f}_c is the central face of $\text{mirror}(\Gamma)$,
- (ii) $\ell_C^m(\bar{e}) = -\ell_C(e)$ for all essential cycles C and edges e on C , where ℓ^m is the labeling in $\text{mirror}(\Gamma)$.

In particular, increasing and decreasing cycles of Γ correspond to decreasing and increasing cycles of $\text{mirror}(\Gamma)$, respectively.

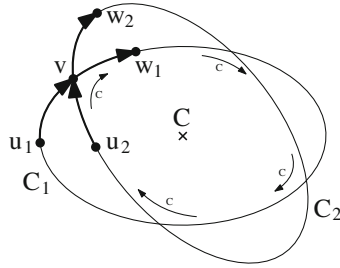


Fig. 18 Illustration of proof for Lemma 5.2. The essential cycles C_1 and C_2 intersect each other having a common vertex v . The essential cycle C bounds the central face of H

Proof Note that the construction of $\text{mirror}(\Gamma)$ satisfies the conditions for ortho-radial representations. Let $e = tu$ be an edge on an essential cycle C and let P be a reference path from $e^* = rs$ to e that respects C ; by Lemma 4.1 we assume that P starts at s and ends at t . After mirroring, P still is a reference path from e^* to e , but its rotation in $\text{mirror}(\Gamma)$ may be different from its rotation in Γ . As above, to distinguish the directions and rotations of paths in $\text{mirror}(\Gamma)$ from the ones in Γ , we include $\text{mirror}(\Gamma)$ and Γ as subscripts to rot and dir .

As P starts at s and ends at t we have by definition

$$\begin{aligned} \text{dir}_\Gamma(e^*, P, e) &= \text{rot}_\Gamma(e^* + P + e), \\ \text{dir}_{\text{mirror}(\Gamma)}(\overline{e^*}, P, \overline{e}) &= \text{rot}_{\text{mirror}(\Gamma)}(e^* + P + e). \end{aligned}$$

As for any path Q we have $\text{rot}_\Gamma(Q) = -\text{rot}_{\text{mirror}(\Gamma)}(Q)$, we obtain $\text{dir}_{\text{mirror}(\Gamma)}(e^*, P, \overline{e}) = -\text{dir}_\Gamma(e^*, P, e)$. By the definition of labels as directions of reference paths we directly obtain that $\ell_C^m(\overline{e}) = -\ell_C(e)$. In particular, if C is increasing (decreasing) in Γ , then \overline{C} is decreasing (increasing) in $\text{mirror}(\Gamma)$. \square

5 Properties of Labelings

In this section we study the properties of labelings in more detail to derive useful tools for proving Theorem 3.9. Throughout this section, we are given an instance $(G, \mathcal{E}, f_c, f_o)$ with an ortho-radial representation Γ and a reference edge e^* . The following observation follows immediately from the definition of labels and the fact that the rotation of any essential cycle is 0.

Observation 5.1 *Let C be an essential cycle. Then, for any two edges e and e' on C , it holds that $\text{rot}(C[e, e']) = \ell_C(e') - \ell_C(e)$.*

The rest of this section is devoted to understanding the relationship between labelings of essential cycles that share vertices or edges. The following technical lemma is a key tool in this respect; see also Fig. 18.

Lemma 5.2 *Let C_1 and C_2 be two essential cycles and let $H = C_1 + C_2$ be the subgraph of G formed by C_1 and C_2 . Let v be a common vertex of C_1 and C_2 that is incident*

to the central face of H . For $i = 1, 2$, let further u_i and w_i be the vertices preceding and succeeding v on C_i , respectively. Then $\ell_{C_1}(vw_1) = \ell_{C_1}(u_1v) + \text{rot}(u_1vw_1) = \ell_{C_2}(u_2v) + \text{rot}(u_2vw_1)$. Moreover, if $w_1 = w_2$, then $\ell_{C_1}(vw_1) = \ell_{C_2}(vw_2)$.

Proof By Observation 5.1 it holds that $\ell_{C_1}(vw_1) = \ell_{C_1}(u_1v) + \text{rot}(u_1vw_1)$.

Let C be the cycle that bounds the central face of H . First assume that the edge vw_1 is incident to the central face of H . Every reference path that respects C_1 also respects C . Hence,

$$\ell_C(vw_1) = \ell_{C_1}(vw_1).$$

Therefore, our goal is to show $\ell_C(vw_1) = \ell_{C_2}(u_2v) + \text{rot}(u_2vw_1)$.

Let P be a reference path from the reference edge rs to u_2v respecting C_2 . By Lemma 4.1 we assume that P starts at the target of the reference edge. Further, we assume that P contains either u_2 or v but not both.

Let x be the first vertex of P that lies on C_2 . If $x \neq v$, then $Q = P[s, x] + C_2[x, u_2]$ is a reference path for u_2v that respects C_2 and does not contain v . Then, $Q' = Q + u_2v$ is a reference path for vw_1 respecting C . Hence,

$$\begin{aligned} \ell_C(vw_1) &= \text{rot}(e^* + Q' + vw_1) = \text{rot}(e^* + Q + u_2v) + \text{rot}(u_2vw_1) \\ &= \ell_{C_2}(u_2v) + \text{rot}(u_2vw_1). \end{aligned}$$

Now assume that $x = v$. Let $Q = P[s, v]$ and let y be the vertex before v on Q . The vertex y lies strictly in the exterior of C_2 and thus not on C . In particular, $y \neq w_1$.

By local consistency we get

$$\text{rot}(yvu_2) + \text{rot}(u_2vw_1) + \text{rot}(w_1vy) = 2.$$

Since Q respects C_2 and thus also C , we find

$$\begin{aligned} \ell_{C_2}(u_2v) &= \text{rot}(e^* + Q + vu_2) - 2 = \text{rot}(e^* + Q) + \text{rot}(yvu_2) - 2, \\ \ell_C(vw_1) &= \text{rot}(e^* + Q + vw_1) = \text{rot}(e^* + Q) + \text{rot}(yvw_1). \end{aligned}$$

Subtracting the first equation from the second we get

$$\begin{aligned} \ell_C(vw_1) - \ell_{C_2}(u_2v) &= \text{rot}(yvw_1) - \text{rot}(yvu_2) + 2 \\ &= -\text{rot}(w_1vy) - \text{rot}(yvu_2) + 2 \\ &= \text{rot}(u_2vw_1), \end{aligned}$$

where the last step uses the equation obtained by local consistency.

If vw_1 does not lie on C , then the edge vw_2 does. By swapping the roles of C_1 and C_2 and using the same argument as above, we obtain

$$\ell_{C_1}(u_1v) + \text{rot}(u_1vw_2) = \ell_{C_2}(u_2v) + \text{rot}(u_2vw_2).$$

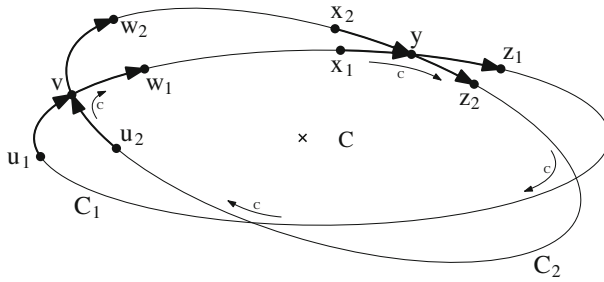


Fig. 19 Illustration of proof for Proposition 5.4. All edges of C_2 are labeled with 0. In this situation there are edges on C_1 with labels -1 and 1 . Hence, C_2 is neither increasing nor decreasing

Since vw_1 lies locally to the left of both u_1vw_2 and u_2vw_2 , it is $\text{rot}(u_ivw_1) = \text{rot}(u_ivw_2) - \alpha$ for $i = 1, 2$ and the same constant α , which is either 1 or 2. Hence, we get

$$\begin{aligned} \ell_{C_1}(vw_1) &= \ell_{C_1}(u_1v) + \text{rot}(u_1vw_1) = \ell_{C_1}(u_1v) + \text{rot}(u_1vw_2) - \alpha \\ &= \ell_{C_2}(u_2v) + \text{rot}(u_2vw_2) - \alpha = \ell_{C_2}(u_2v) + \text{rot}(u_2vw_1). \end{aligned}$$

Finally, if $w_1 = w_2$, i.e., vw_1 lies on both C_1 and C_2 , then $\ell_{C_1}(vw_1) = \ell_{C_2}(u_2v) + \text{rot}(u_2vw_1) = \ell_{C_2}(u_2v) + \text{rot}(u_2vw_2) = \ell_{C_2}(vw_2)$. □

Corollary 5.3 *Let C_1 and C_2 be two essential cycles, let $H = C_1 + C_2$ be the subgraph of G formed by C_1 and C_2 , and let e be an edge that lies on both C_1 and C_2 and that is incident to the central face of H . Then $\ell_{C_1}(e) = \ell_{C_2}(e)$.*

This allows us to prove an important criterion to exclude strictly monotone cycles. We call an essential cycle C *horizontal* if $\ell_C(e) = 0$ for all edge e of C . We show that a strictly monotone cycle and a horizontal cycle cannot share vertices.

Proposition 5.4 *Let C_2 be a horizontal cycle and let C_1 be an essential cycle that shares at least one vertex with C_2 . Then C_1 is not strictly monotone.*

Proof The situation is illustrated in Fig. 19.

If the two cycles are equal, the claim clearly holds. Otherwise, we show that one can find two edges on C_1 whose labels have opposite signs.

Let v be a shared vertex of C_1 and C_2 that is incident to the central face f of $H = C_1 + C_2$ and such that the edge vw_1 on C is not incident to f . For $i = 1, 2$ denote by u_i and w_i the vertex preceding and succeeding v on C_i , respectively. By Lemma 5.2 it is

$$\ell_{C_1}(vw_1) = \ell_{C_2}(u_2v) + \text{rot}(u_2vw_1) = \text{rot}(u_2vw_1), \tag{3}$$

where the second equality follows from the assumption that $\ell_{C_2}(u_2v) = 0$.

Let y be the first common vertex of C_1 and C_2 on the central face f after v . That is, $f[v, y]$ is a part of one of the cycles C_1 and C_2 , and it intersects the other cycle only

at v and y . For $i = 1, 2$, we denote by x_i and z_i the vertices preceding and succeeding y on C_i . Again by Lemma 5.2 (this time swapping the roles of C_1 and C_2), we have

$$0 = \ell_{C_2}(yz_2) = \ell_{C_1}(x_1y) + \text{rot}(x_1yz_2). \tag{4}$$

Overall, we have $\ell_{C_1}(vw_1) = \text{rot}(u_2vw_1)$ and $\ell_{C_1}(x_1y) = -\text{rot}(x_1yz_2)$.

By construction vw_1 and x_1y lie on the same side of C_2 . Hence, u_2vw_1 and x_1yz_2 both make a right turn if vw_1 and x_1y lie in the interior of C_2 and a left turn otherwise. Thus, it is $\text{rot}(u_2vw_1) = \text{rot}(x_1yz_2) \neq 0$, and therefore $\ell_{C_1}(vw_1)$ and $\ell_{C_1}(x_1y)$ have opposite signs. Hence, C_1 is not strictly monotone. \square

In many cases we cannot assume that one of two essential cycles sharing a vertex is horizontal. However, we can still draw conclusions about their intersection behavior from their labelings and find conditions under which shared edges have the same label on both cycles.

Intuitively, positive labels can often be interpreted as going downwards and negative labels as going upwards. In Fig. 12b all edges of $C_1[v, w]$ have positive labels and in total the distance from the center decreases along this path, i.e., the distance of v from the center is greater than the distance of w from the center. Yet, the edges on $C_1[v, w]$ point in all possible directions—even upwards. One can still interpret a maximal path with positive labels as leading downwards with the caveat that this is a property of the whole path and does not impose any restriction on the directions of the individual edges.

Using this intuition, we expect that a path P going down cannot intersect a path Q going up if P starts below Q . In Lemma 5.5, we show that this assumption is correct if we restrict ourselves to intersections on the central face.

Lemma 5.5 *Let C_1 and C_2 be two simple, essential cycles in G sharing at least one vertex. Let $H = C_1 + C_2$, and denote the central face of H by f . Let v be a vertex that is shared by C_1 and C_2 that is incident to f and, for $i = 1, 2$, let u_i and w_i be the vertices preceding and succeeding v on C_i , respectively.*

- (i) *If $\ell_{C_1}(u_1v) \geq 0$ and $\ell_{C_2}(u_2v) \leq 0$, then u_2v lies in the interior of C_1 .*
- (ii) *If $\ell_{C_1}(vw_1) \geq 0$ and $\ell_{C_2}(vw_2) \leq 0$, then vw_2 lies in the exterior of C_1 .*

Proof The second case follows from the first by taking the mirror representation; this reverses the order on the cycles and changes the sign of each label by Lemma 4.4, but does not change the notion of interior and exterior. It therefore suffices to consider the first case.

Since the central face f lies in the interior of both C_1 and C_2 and v lies on the boundary of f , one of the edges vw_1 and vw_2 is incident to f . We denote this edge by vx and it is either $x = w_1$ (as in Fig. 20a) or $x = w_2$. By Lemma 5.2, we have $\ell_{C_1}(u_1v) + \text{rot}(u_1vx) = \ell_{C_2}(u_2v) + \text{rot}(u_2vx)$. Applying $\ell_{C_1}(u_1v) \geq 0$ and $\ell_{C_2}(u_2v) \leq 0$, we obtain $\text{rot}(u_1vx) \leq \text{rot}(u_2vx)$. Therefore, u_2v lies to the right of or on u_1vx and thus in the interior of C_1 . \square

The next lemma is a direct consequence of Lemma 5.5 when applied to decreasing and increasing cycles.

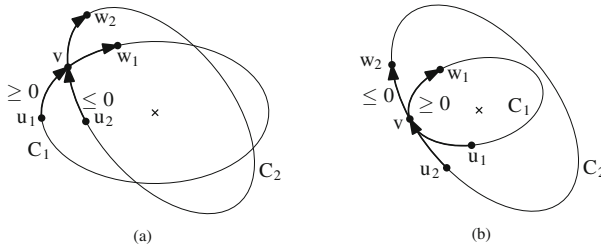


Fig. 20 Possible intersection of two cycles C and C' at v . **a** The labels of the incoming edges satisfy $\ell_{C_1}(u_1v) \geq 0$ and $\ell_{C_2}(u_2v) \leq 0$. The edges vw_1 and vw_2 could be exchanged. **b** The labels of the outgoing edges satisfy $\ell_{C_1}(vw_1) \geq 0$ and $\ell_{C_2}(vw_2) \leq 0$. The edges u_1v and u_2v could be exchanged

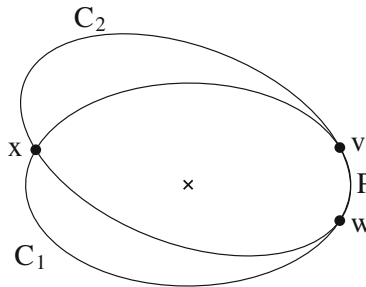


Fig. 21 Illustration of proof for Lemma 5.6. The cycle C_1 is increasing and the cycle C_2 is decreasing. The edge of C_2 ending at v strictly lies in the exterior of C_1 and the edge of C_2 starting at w strictly lies in the interior of C_1 . Hence, the edge of C_2 ending at x strictly lies in the interior of C_1 , which contradicts Lemma 5.5

Lemma 5.6 *A decreasing and an increasing cycle do not have any common vertex.*

Proof For an illustration see Fig. 21. Let C_1 be an increasing and C_2 a decreasing cycle. Assume that they have a common vertex. But then there also is a common vertex on the central face f of the subgraph $C_1 + C_2$. Consider any maximal common path P of C_1 and C_2 on f . We denote the start vertex of P by v and the end vertex by w . Note that v may equal w . By Lemma 5.5 the edge to v on the decreasing cycle C_2 lies strictly in the exterior of C_1 , where the strictness follows from the maximality of P . Similarly, the edge from w on C_2 lies strictly in the interior of C_1 . Hence, $C_2[w, v]$ crosses C_1 . Let x be the first intersection of C_1 and C_2 on f after w . Then the edge to x on C_2 lies strictly in the interior of C_1 , contradicting Lemma 5.5. \square

For the correctness proof in Sect. 7, a crucial insight is that for essential cycles using an edge that is part of a regular face, we can find an alternative cycle without this edge in a way that preserves labels on the common subpath.

Lemma 5.7 *If an edge e belongs to both a simple essential cycle C and a regular face f , then there exists a simple essential cycle C' that can be decomposed into two paths P and Q such that*

- (i) P or \bar{P} lies on f ,
- (ii) $Q = C \cap C'$, and

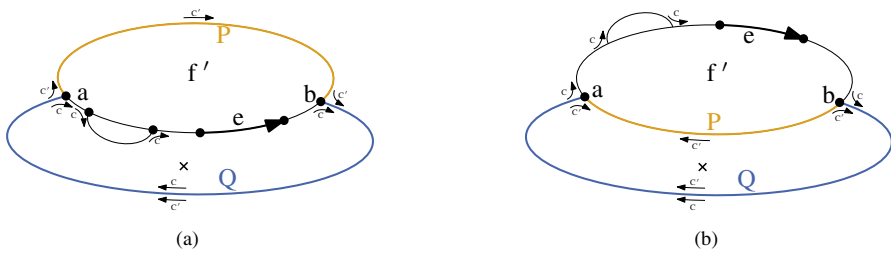


Fig. 22 The edge e cannot lie on both the outer and the central face. In both cases C' can be subdivided in two paths P and Q on C and f , respectively. Here, these paths are separated by the vertices a and b . **a** The cycle bounding the outer face is C' . The edge e does not lie on the outer face, and hence the cycle bounding this face is defined as C' . **b** The cycle bounding the central face is C'

(iii) $\ell_C(e) = \ell_{C'}(e)$ for all edges e on Q .

Proof Consider the graph $H = C + f$ composed of the essential cycle C and the regular face f . Since e is incident to f , the edge e cannot lie on both the outer and the central face in H . If e does not lie on the outer face, we define C' as the cycle bounding the outer face but directed such that it contains the center in its interior; see Fig. 22a. Otherwise, C' denotes the cycle bounding the central face; see Fig. 22b.

Since C lies in the exterior of f , the intersection of C with C' forms one contiguous path Q . Setting $P = C' - Q$ yields a path that lies completely on f (it is possible though that P and f are directed differently).

By the construction of C' the edges of Q are incident to the central face of $C + C'$. Then Lemma 5.2 implies that $\ell_C(e) = \ell_{C'}(e)$ for all edges e of Q . \square

The last lemma of this section shows that we can replace single edges of an essential cycle C with complex paths without changing the labels of the remaining edges on C .

Lemma 5.8 *Let C be an essential cycle in an ortho-radial representation Γ and let uv be an edge of C . Consider an ortho-radial representation Γ' that is created by replacing uv with a path P such that the interior vertices of P do not belong to Γ , i.e., they are newly inserted vertices in Γ' . If the cycle $C' = C[v, u] + P$ is essential, then the labels of C and C' coincide on $C[v, u]$ ($= C'[v, u]$).*

Proof For an illustration of the proof see Fig. 23. Let $e = xy$ be an arbitrary edge on $C[v, u]$ and let R be a reference path from the reference edge e^* to e that respects C . We first construct a new reference path R' that does not contain uv as follows. Let w be the first vertex of R that lies on C . If $C[w, x]$ does not contain e , we define $R' = R - R[w, x] + C[w, x]$ and otherwise $R' = R - R[w, x] + \overline{C}[w, y]$. We observe that R' is again a reference path of e that respects C . Further, it can be partitioned into a prefix that only consists of edges that do not belong to C and a suffix that only consists of edges that belong to $C[v, u]$.

We now show that R' is a reference path of e in Γ' that respects C' . As R' does not use uv , it is still contained in Γ' and hence it is a reference path of e . So assume that R' does not respect C' . Hence, R' and C' have a vertex z in common such that

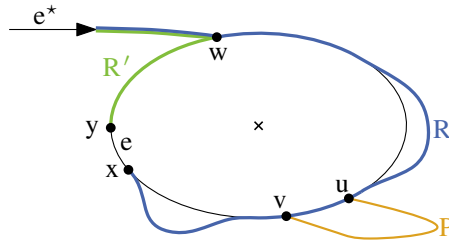


Fig. 23 Illustration of the proof for Lemma 5.8. The reference path R' (green) for the edge xy is constructed based on the reference path R . Replacing the edge uv with a path P , does not impact R' , which implies that the label of xy remains the same

the outgoing edge e' of R' at z strictly lies in the interior of C' . As R' respects C , this vertex lies on P . It cannot be an intermediate vertex of P , because these are newly inserted in Γ' . Hence, z is either u or v and thus part of $C[v, u]$. In particular, it occurs on R' after w , which implies that e' belongs to C . This contradicts that e' strictly lies in the interior of C' . Altogether, this shows that R' is a reference path of e both in Γ and Γ' such that C and C' are respected. Consequently, $\ell_C(e) = \ell_{C'}(e)$. \square

6 Characterization of Rectangular Ortho-Radial Representations

Throughout this section, assume that $I = (G, \mathcal{E}, f_c, f_o)$ is an instance with an ortho-radial representation Γ and a reference edge e^* . We prove Theorem 3.9 for the case that Γ is *rectangular*. In a *rectangular ortho-radial representation* the central face and the outer face are horizontal cycles, and every regular face is a rectangle, i.e., it has exactly four right turns, but no left turns.

We first observe that a bend-free ortho-radial drawing Δ can be described by an angle assignment together with the lengths of its vertical edges and the angles of the circular arcs representing the horizontal edges; we call the angles of the circular arcs *central angles*. We define two flow networks that assign consistent lengths and central angles to the vertical edges and horizontal edges, respectively. These networks are straightforward adaptations of the networks used for drawing rectangular graphs in the plane [34]. In the following, *vertex* and *edge* refer to the vertices and edges of G , whereas *node* and *arc* are used for the flow networks.

The network $N_{\text{ver}} = (F_{\text{ver}}, A_{\text{ver}})$ contains one node for each face of G except for the central and the outer face. All nodes have a demand of 0. For each vertical edge e in G , which we assume to be directed upwards, there is an arc $a_e = fg$ in N_{ver} , where f is the face to the left of e and g the one to its right. The flow on fg has the lower bound $l(fg) = 1$ and upper bound $u(fg) = \infty$. An example of this flow network is shown in Fig. 24a.

Similarly, the network $N_{\text{hor}} = (F_{\text{hor}}, A_{\text{hor}})$ assigns the central angles of the horizontal edges. There is a node for each face of G , and an arc $a_e = fg$ for every horizontal edge e in G such that f lies locally below e and g lies locally above e . Additionally, N_{hor} includes one arc from the outer to the central face. Again, all edges require a minimum flow of 1 and have infinite capacity. The demand of all nodes is 0.

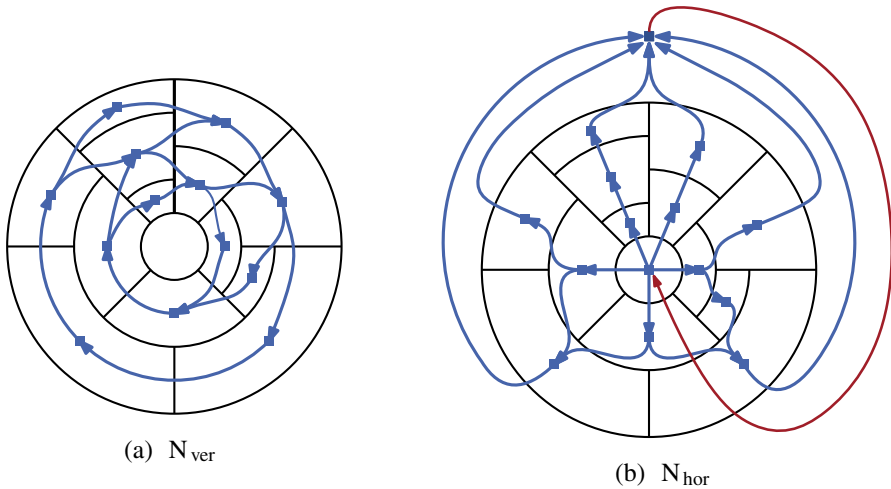


Fig. 24 Flow networks N_{ver} and N_{hor} (blue arcs) for an example graph G (black edges)

Figure 24b shows an example of such a flow network. Valid flows in these two flow networks then yield an ortho-radial drawing of Γ .

Lemma 6.1 *A pair of valid flows in N_{hor} and N_{ver} corresponds to a bend-free ortho-radial drawing of Γ and vice versa.*

Proof Given a bend-free ortho-radial drawing Δ of Γ , we can extract flows in the two networks. For each vertical edge e we set the flow $\varphi_{\text{ver}}(a_e)$ of the corresponding arc a_e to l_e/l_{min} , where l_e is the length of e in Δ and l_{min} is the length of the shortest edge in Δ . With the scaling, we ensure that the flow of each arc is at least 1. Similarly, for the horizontal edges we assign to each arc a_e of each horizontal edge e the flow $\varphi_{\text{hor}}(a_e) = \alpha_e/\alpha_{\text{min}}$, where α_e is the central angle of e in Δ and α_{min} is the smallest central angle of any horizontal edge in Δ . Again, the scaling ensures that each arc has flow at least 1. Since the opposing vertical sides of the regular faces have the same lengths and the opposing horizontal sides of the regular faces have the same central angles, the flow is preserved at all nodes.

Conversely, given feasible flows in the flow networks, we construct a corresponding drawing by interpreting the flow values in φ_{ver} as edge lengths, and the flow values in φ_{hor} as central angles. The flow conservation ensures that opposing sides have the same lengths (vertical edges) and same central angles (horizontal edges), respectively. To obtain a drawing, we normalize the central angles so that the total flow on the back edge of N_{hor} is 2π . To prove that this works, we reduce it to the orthogonal case by cutting performing a straight-line cut from the central face to the outer face.

First observe that if there exists a path P of only vertical edges that runs from a vertex incident to the central face to a vertex incident to the outer face, then we can split this path into a left copy P_L and a right copy P_R as in Fig. 25a. Then the resulting flow corresponds to an orthogonal drawing (on the standard grid). However, since for each edge e of G its copies on P_L and P_R have the same lengths, this orthogonal drawing can be merged into an ortho-radial drawing.

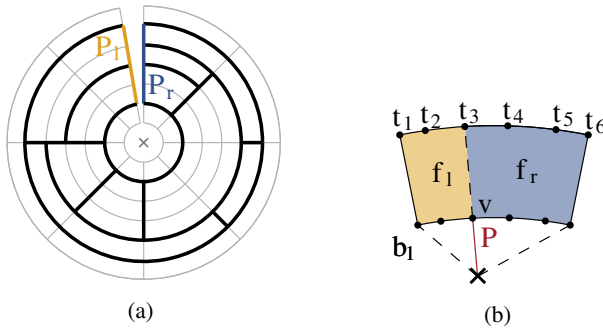


Fig. 25 Illustration of proof of Lemma 6.1. **a** The ortho-radial representation is cut along the vertical path that goes from the center to the outer face. **b** The path P ends at a face at vertex v . The face is divided into two sub-faces

We now show that, if G does not contain such a path, then we can augment G so it contains one. Let P be the maximum length path of only upward edges that starts at a vertex incident to the central face. Assume P ends at a vertex v and that has no edge pointing upward; see Fig. 25b. Since Γ is rectangular, the vertex v then has left and right-pointing edges, and we denote by f the face between them. By assumption f is not the outer face and it is therefore regular. Let B be the right-ward path that bounds the bottom side of f ; we denote the first vertex of B by b_1 . We observe that v lies on B . Let $\phi_B = \sum_{e \in B[b_1, v]} \varphi_{\text{hor}}(a_e)$ denote the sum of the flow values of N_{hor} on the arcs that correspond to edges of $B[b_1, v]$, where a_e denotes the arc of e in N_{hor} . Consider now the right-ward path $T = t_1, \dots, t_k$ that bounds the top side of f . Let i be the smallest index such that $\phi_T = \sum_{e \in T[t_1, t_i]} \varphi_{\text{hor}}(a_e) \geq \phi_B$. If $\phi_B = \phi_T$, then we add to Γ the vertical edge vt_i , thereby splitting f into a left f_l and a right face f_r . We then reconnect the incoming arcs of f before v to f_l and the incoming arcs after v to f_r . Similarly, we reconnect the outgoing arcs of f before t_i to f_l and the outgoing arcs of f after t_i to f_r . Observe that, by construction, we have flow conservation at f_l and f_r and we have extended the maximum length of an outward path from the central face.

Now assume $\phi_B < \phi_T$. In that case, we add a subdivision vertex t' on the edge $t_{i-1}t_i$. We then add the edge vt' , thereby again splitting the face f into a left face f_l and a right face f_r . As before, we reconnect the incoming arcs before v to f_l and the incoming arcs after v to f_r . The outgoing edges of f that correspond to an edge of $T[t_1, t_{i-1}]$ we simply reconnect to f_l and similarly, we reconnect the outgoing edges of f that correspond to an edge of $T[t_i, t_k]$ to f_r . The flow on the outgoing edge fg that corresponds to $t_{i-1}t_i$, we split on the two edges f_lg and f_rg in such a way that the total outflow of f_l is ϕ_B . Observe that this yields again a valid flow in the augmented flow network. Thus also in this case, we have extended the maximum length of an outward path from the central face. Observe further that a drawing of the modified ortho-radial representation Γ' contains a drawing of Γ . It follows that Γ is drawable. \square

Using this correspondence of drawings and feasible flows, we show the characterization of rectangular graphs.

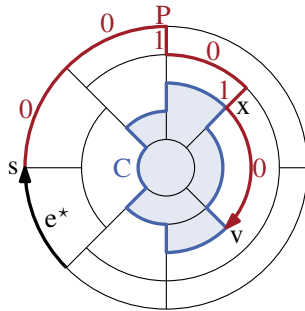


Fig. 26 The path P from s to v —constructed backwards by going only up or left—does not intersect the interior of C . The rotations of the edges on P relative to e^* are 0 or 1

Theorem 6.2 *Let Γ be a rectangular ortho-radial representation and let $N_{\text{hor}} = (F_{\text{hor}}, A_{\text{hor}})$ and $N_{\text{ver}} = (F_{\text{ver}}, A_{\text{ver}})$ be the flow networks as defined above. The following statements are equivalent:*

- (i) Γ is drawable.
- (ii) Γ is valid.
- (iii) For every subset $S \subseteq F_{\text{ver}}$ such that there is an arc from $F_{\text{ver}} \setminus S$ to S in N_{ver} , there is also an arc from S to $F_{\text{ver}} \setminus S$.

Proof “(i) \Rightarrow (ii)”: Let Δ be a bend-free ortho-radial drawing of G preserving the embedding described by Γ and let C be an essential cycle. Our goal is to show that C is not strictly monotone. To this end, we construct a path P from the reference edge of Γ to a vertex on C such that either the labeling of C induced by P attains both positive and negative values or it is 0 everywhere.

In Δ either all vertices of C lie on the same concentric circle, or there is a maximal subpath Q of C whose vertices all have maximum distance to the center of the ortho-radial grid among all vertices of C . In the first case, we may choose the endpoint v of the path P arbitrarily, whereas in the second case we select the first vertex of Q as v ; for an example see Fig. 26.

We construct the path P backwards (i.e., the construction yields \overline{P}) as follows: Starting at v we choose the edge going upwards from v , if it exists, or the one leading left. Since all faces of Γ are rectangles, at least one of these always exists. This procedure is repeated until the target s of the reference edge is reached.

To show that this algorithm terminates, we assume that this was not the case. As G is finite, there must be a first time a vertex w is visited twice. Hence, there is a cycle C' in Δ containing w that contains only edges going left or up. As all drawable essential cycles with edges leading upwards must also have edges that go down [21], all edges of C' are horizontal. By construction, there is no edge incident to a vertex of C' that leads upwards. The only cycle with this property, however, is the one enclosing the outer face because G is connected. But this cycle contains the reference edge, and therefore the algorithm halts.

This not only shows that the construction of P ends, but also that P is a path (i.e., the construction does not visit a vertex twice). Thus, P is a reference path from the

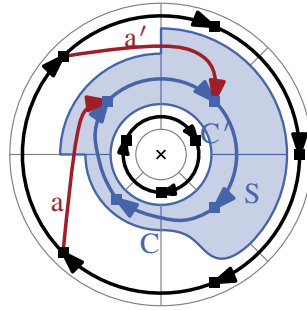


Fig. 27 A set S of nodes in a graph G such that $N_{\text{ver}}[S]$ has no outgoing but two incoming arcs a and a' . The set of faces \mathcal{S} corresponding to the nodes in S are shaded with blue. The outermost boundary of \mathcal{S} forms an increasing cycle C . The edges on this cycle with label -1 are exactly those that are crossed by a or a' . All other edges on C are labeled with 0 . Note that the edge on C at the bottom is curved because G does not admit an ortho-radial drawing

reference edge e^* to the edge vv' , where v' is the vertex following v on C . Further, P respects C as \bar{P} starts at the outermost circle of the ortho-radial grid that is used by C and as by construction all edges of \bar{P} point left or upwards.

By the construction of P , the label of vv' induced by P is 0 . If all edges of C are horizontal, this implies $\ell_C(e) = 0$ for all edges e of C , which shows that C is not strictly monotone. Otherwise, we claim that the edges $e_- = uv$ and $e_+ = wx$ directly before and after Q on C have labels -1 and $+1$, respectively. Since all edges on Q are horizontal and e_- goes down, we have $\text{rot}(C[v, x]) = 1$ and therefore $\ell_C(e_+) = 1$. Similarly, $\text{rot}(uvv') = 1$ implies that $\ell_C(uv) = \ell_C(vv') - \text{rot}(uvv') = -1$.

“(ii) \Rightarrow (iii)”: Instead of proving this implication directly, we show the contrapositive. That is, we assume that there is a set $S \subsetneq F_{\text{ver}}$ of nodes in N_{ver} such that S has no outgoing but at least one incoming arc. From this assumption we derive that Γ is not valid, as we find a strictly monotone cycle.

Let $N_{\text{ver}}[S]$ denote the node-induced subgraph of N_{ver} induced by the set S . Without loss of generality, S can be chosen such that $N_{\text{ver}}[S]$ is weakly connected, i.e., the underlying undirected graph is connected. If N_{ver} is not weakly connected, at least one weakly-connected component of $N_{\text{ver}}[S]$ possesses an incoming arc but no outgoing arc, and we can work with this component instead.

As each node of S corresponds to a face of G , S can also be considered as a collection of faces of G . To distinguish the two interpretations of S , we refer to this collection of faces by \mathcal{S} . Our goal is to show that the innermost or the outermost boundary of \mathcal{S} forms a strictly monotone cycle in Γ .

Figure 27 shows an example of such a set S of nodes. Here, the arcs a and a' lead from a node outside of S to one in S . These arcs cross edges on the outer boundary of \mathcal{S} , which point upwards.

Let \mathcal{F} be the set of faces of \mathcal{E} (including the central and outer face). Let Z be a connected component of $\mathcal{F} \setminus \mathcal{S}$ such that there exists an arc from Z to S in N_{ver} and let C be the cycle in G that separates Z from \mathcal{S} . If C were non-essential, then $\text{rot}(C) = 4$ and C would therefore contain an upward and a downward edge. One of these edges

would correspond to an incoming arc of S and the other edge to an outgoing arc of S , contradicting the choice of S . Thus, C is essential.

As usual we consider C in clockwise direction. We may assume without loss of generality that C contains S in its interior; otherwise, we consider $\text{flip}(\Gamma)$ and \bar{C} . Note that for each edge of C the face locally to the right belongs to S whereas the face locally to the left does not. Hence, upward edges of C correspond to incoming arcs of S and downward edges to outgoing arcs. Since there is an arc from Z to S but not vice versa, the cycle C contains at least one upward but no downward edge. Hence, there is some integer k such that all labels of C belong to $L_k = \{4k, 4k + 1, 4k + 2\}$. Since the numbers in L_k are either all non-negative (if $k \geq 0$) or all negative (if $k < 0$), the cycle C is monotone. Moreover, C is not horizontal because it has an upward edge.

“(iii) \Rightarrow (i)”: By Lemma 6.1 the existence of a drawing is equivalent to the existence of feasible flows in N_{hor} and N_{ver} . If a flow network N contains for each arc a a cycle C_a that contains a , then routing one unit flow along each of these cycles C_a and adding all flows gives a circulation in N where at least one unit flows along each arc. Hence, it suffices to prove that in N_{hor} and in N_{ver} each arc is contained in a cycle.

Note that N_{hor} without the arc from the outer face g to the central face f is a directed acyclic graph with f as its only source and g as its only sink. For each arc $a \neq gf$ in N_{hor} there is a directed path P_a from f to g via a . Adding the arc gf , we obtain the cycle $C_a = P_a + gf$.

For N_{ver} we consider an arc $a = fg$, and we define the set S_g of all nodes h for which there exists a directed path from g to h in N_{ver} . By definition, there is no arc from a vertex in S_g to a vertex not in S_g . As N_{ver} satisfies (iii), S_g does not have any incoming arcs either. Hence, $f \in S_g$ and there is a directed path P_a from g to f . Then $C_a = P_a + fg$ is the desired cycle. \square

By [21] the angle assignment defined by an ortho-radial drawing of a graph is locally consistent (in the sense of Definition 3.1). Therefore, Theorem 6.2 implies the characterization of ortho-radial drawings for rectangular graphs.

Corollary 6.3 (Theorem 3.9 for Rectangular Ortho-Radial Representations) *A rectangular ortho-radial representation is drawable if and only if it is valid.*

We note that we can construct the flows in N_{hor} and N_{ver} using standard techniques based on flows in planar graphs with multiple sinks and sources [25]. With this a drawing can be computed in $\mathcal{O}(n^{\frac{3}{2}})$ time.

7 Drawable Representations of Planar 4-Graphs

In the previous section we proved that a rectangular ortho-radial representation is drawable if and only if it is valid. We extend this result to general ortho-radial representations by reduction to the rectangular case. In Sect. 7.1 we present a procedure that augments a given instance such that all faces become rectangles. For readability we defer some of the proofs to Sect. 7.2. In Sect. 7.3 we use the rectangulation procedure and Corollary 6.3 to show Theorem 3.9. We remark that all our proofs are constructive,

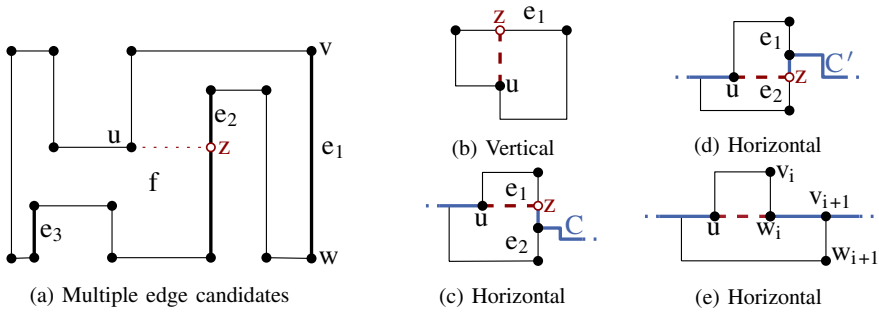


Fig. 28 Examples of augmentations. **a** The candidate edges of u are e_1, e_2 and e_3 . **b** Insertion of vertical edge uz . **c** $\Gamma_{e_1}^u$ contains a decreasing cycle. **d** $\Gamma_{e_2}^u$ is valid. **e** Insertion of horizontal edge uw_i because there is a horizontal path from w_i to u

but make use of tests whether certain modified ortho-radial representations are valid. We develop an efficient testing algorithm for this in Sect. 8.

7.1 Rectangulation Procedure

Throughout this section, we are given an instance $I = (G, \mathcal{E}, f_c, f_o)$ with a valid ortho-radial representation Γ and a reference edge e^* . The core of the argument is a *rectangulation procedure* that successively augments G with new vertices and edges to a graph G' along with a valid rectangular ortho-radial representation Γ' . Then, Γ' is drawable by Corollary 6.3, and removing the augmented parts yields a drawing of Γ .

The rectangulation procedure works by augmenting non-rectangular faces one by one, thereby successively removing concave angles at the vertices until all faces are rectangles. Traversing the boundary of a face in clockwise direction yields a sequence of left and right turns, where a degree-1 vertex contributes two left turns. Note that concave angles correspond exactly to left turns in this sequence. Consider a face f with a left turn (i.e., a concave angle) at u such that the following two turns when walking along f (in clockwise direction) are right turns; see Fig. 28. We call u a *port* of f . We define a set of *candidate edges* (or just *candidates*) that contains precisely those edges vw of f , for which $\text{rot}(f[u, vw]) = 2$; see Fig. 28a. We treat this set as a sequence, where the edges appear in the same order as in f , beginning with the first candidate after u . The *augmentation* Γ_{vw}^u with respect to a candidate edge vw is obtained by splitting the edge vw into the edges vz and zw , where z is a new vertex, and adding the edge uz in the interior of f such that the angle formed by zu and the edge following u on f is $\pi/2$. The direction of the new edge uz in Γ_{vw}^u is the same for all candidate edges. If this direction is vertical, we call u a *vertical port* and otherwise a *horizontal port*. We note that any vertex with a concave angle in a face becomes a port during the augmentation process. For regular faces Tamassia [31] shows that they always contain a port. Moreover, the following observation can be proven analogously.

Observation 7.1 *If u is a port of a face f and vw is a candidate edge for u , then Γ_{vw}^u is an ortho-radial representation.*

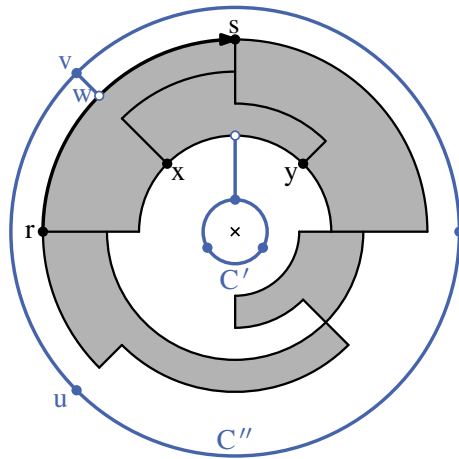


Fig. 29 The outer and the central face are rectangulated by adding cycles of length 3. The cycle C' is connected to an arbitrary edge xy that has label 0 and C'' is connected to a new vertex on the old reference edge rs . The edge uv is selected as the new reference edge

However, an augmentation Γ_{vw}^u is not necessarily valid. We prove that we can always find an augmentation that is valid. The crucial ingredient is the following proposition.

Proposition 7.2 *Let G be a planar 4-graph with valid ortho-radial representation Γ , let f be a regular face of G and let u be a port of f with candidate edges $e_1 = v_1w_1, \dots, e_k = v_kw_k$. Then the following facts hold:*

1. *If u is a vertical port, then $\Gamma_{e_1}^u$ is a valid ortho-radial representation; see Fig. 28b.*
2. *If u is a horizontal port, then $\Gamma_{e_1}^u$ does not contain an increasing cycle and $\Gamma_{e_k}^u$ does not contain a decreasing cycle; see Fig. 28c, d.*
3. *Let P_i be the maximal path in Γ that contains the vertex w_i of the candidate edge $e_i = v_iw_i$ and that consists of only horizontal edges. If u is a horizontal port and $\Gamma_{e_i}^u$ contains a decreasing cycle and $\Gamma_{e_{i+1}}^u$ contains an increasing cycle, then u is an endpoint of P_i and adding the horizontal edge uz to the other endpoint z of P_i yields a horizontal cycle; see Fig. 28e. In particular, $\Gamma + uz$ is valid.*

To increase the readability we split the proof of Proposition 7.2 into the separate Lemmas 7.3–7.7, which we defer to Sect. 7.2.

We are now ready describe the rectangulation procedure. Let G be a planar 4-graph with valid ortho-radial representation Γ . Without loss of generality, we can assume that G is connected, otherwise we can treat the connected components separately. We further insert triangles in both the central and outer face and suitably connect these to the original graph; see Fig. 29. Namely, for the central face g we identify an edge e on the simple cycle C bounding g such that $\ell_C(e) = 0$. Since Γ is valid and C is an essential cycle, such an edge exists. We then insert a new cycle C' of length 3 inside g and connect one of its vertices to a new vertex on e . The new cycle C' now forms the boundary of the central face. Analogously, we insert into the outer face a cycle C_o of length 3 which contains G and is connected to the reference edge e^* . We choose

an arbitrary edge e^{**} on C_o as new reference edge. We observe that there is a path P from e^{**} to e^* with rotation 0. Hence, each reference path from e^* to an essential cycle in Γ can be extended by P such that the new path is a reference path with respect to e^{**} and has the same rotation.

After this preprocessing any face f that is not a rectangle is regular, and it therefore contains a port u . If any of the candidate augmentations $\Gamma_{e_i}^u$ is valid, then $\Gamma_{e_i}^u$ has fewer concave corners than Γ , and we continue the augmentation procedure with $\Gamma_{e_i}^u$. On the other hand, if none of these augmentations is valid, then each $\Gamma_{e_i}^u$ contains a strictly monotone cycle. Let i be the smallest index such that $\Gamma_{e_i}^u$ contains an increasing cycle and note that such an index i exists and that $i > 1$ by property 2 of Proposition 7.2. Then, by definition of i , $\Gamma_{e_{i-1}}^u$ contains a decreasing cycle and $\Gamma_{e_i}^u$ contains an increasing cycle. But then property 3 of Proposition 7.2 guarantees the existence of a vertex z such that $\Gamma + uz$ is valid and has fewer concave corners than Γ . Using this procedure we can iteratively augment Γ to a rectangular ortho-radial representation that contains a subdivision of the ortho-radial representation Γ .

7.2 Proof of Proposition 7.2

Throughout this section, we assume that we are in the situation described by Proposition 7.2. That is, G is a planar 4-graph with valid ortho-radial representation Γ , and u is a port of a regular face f with candidate edges $e_1 = v_1w_1, \dots, e_k = v_kw_k$. After possibly replacing Γ with $\text{flip}(\Gamma)$, we may assume that the edge uz resulting from an augmentation with a candidate is directed to the right or up. By Lemma 4.3 there is a one-to-one correspondence between increasing (decreasing) cycles in Γ and increasing (decreasing) cycles in $\text{flip}(\Gamma)$.

Lemma 7.3 *If u is a vertical port, then $\Gamma_{e_1}^u$ is a valid ortho-radial representation.*

Proof Assume for the sake of contradiction that $\Gamma_{e_1}^u$ contains a strictly monotone cycle C . As Γ is valid, C must contain the new edge uz in either direction (i.e., uz or zu). Let f' be the new rectangular face of $G + uz$ containing u, v_1 and z , and consider the subgraph $H = C + f'$ of $G + uz$. According to Lemma 5.7 there exists an essential cycle C' that does not contain uz . Moreover, C' can be decomposed into paths P and Q such that P lies on f' and Q is a part of C ; see Fig. 30.

The goal is to show that C' is increasing or decreasing. We present a proof only for the case that C is an increasing cycle. The proof for decreasing cycles can be obtained by flipping all inequalities.

For each edge e on Q the labels $\ell_C(e)$ and $\ell_{C'}(e)$ are equal by Lemma 5.7, and hence $\ell_{C'}(e) \leq 0$. For an edge e of P , there are two possible cases: e either lies on the side of f' parallel to uz or on one of the two other sides. In the first case, the label of e is equal to the label $\ell_C(uz)$ ($\ell_C(zu)$ if C contains zu instead of uz). In particular the label is negative.

In the second case, we first note that $\ell_{C'}(e)$ is even, since e points left or right. Assume that $\ell_{C'}(e)$ was positive and therefore at least 2. Then, let e' be the first edge on C' after e that points to a different direction. Such an edge exists, since otherwise C' would be an essential cycle whose edges all point to the right, but they are not

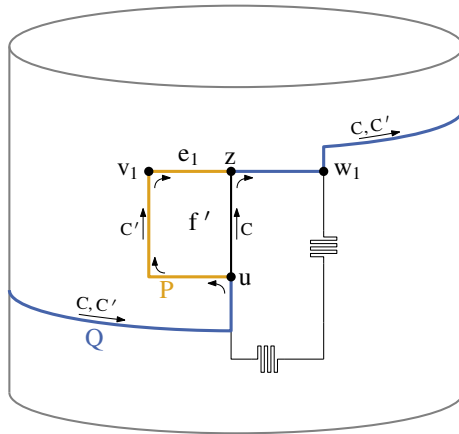


Fig. 30 Illustration of proof for Lemma 7.3. In this illustration it is assumed that inserting the edge uz into a valid ortho-radial representation creates an increasing cycle C that uses uz . However, then there is cycle C' not using uz that is also increasing

labeled with 0. This edge e' lies on Q or is parallel to uz . Hence, the argument above implies that $l_{C'}(e') \leq 0$. However, $l_{C'}(e')$ differs from $l_{C'}(e)$ by at most 1, which requires $l_{C'}(e) \geq 1$. Therefore, $l_{C'}(e)$ cannot be positive.

We conclude that all edges of C' have a non-positive label. If all labels were 0, C would not be an increasing cycle by Proposition 5.4. Thus, there exists an edge on C' with a negative label and C' is an increasing cycle in Γ . But as Γ is valid, such a cycle does not exist, and therefore C does not exist either. Hence, $\Gamma_{e_1}^u$ is valid. \square

Lemma 7.4 *If u is a horizontal port, then $\Gamma_{e_1}^u$ contains no increasing cycle.*

Proof Let f' be the new rectangular face of $\Gamma_{e_1}^u$ containing u, v_1 and z , and assume for the sake of contradiction that there is an increasing cycle C in $\Gamma_{e_1}^u$. This cycle must use either uz or zu . Similar to the proof of Lemma 7.3, we find an increasing cycle C' in Γ , contradicting the validity of Γ .

Applying Lemma 5.7 to C and f' yields an essential cycle C' without uz and zu that can be decomposed into a path P on f' and a path $Q \subseteq C \setminus f'$ such that all edges of Q have non-positive labels. We show in the following that the edges of P also have non-positive labels.

If C contains uz , there are three possibilities for an edge e of P , which are illustrated in Fig. 31: The edge e lies on the left side of f' and points up, e is parallel to uz , or $e = v_1z$. In the first case $l_{C'}(e) = l_C(uz) - 1 < 0$ and in the second case $l_{C'}(e) = l_C(uz) \leq 0$. If $e = v_1z$, C cannot contain zv_1 and therefore $zw_1 \in C$. Then, $l_{C'}(e) = l_C(zw_1) < 0$. In all three cases the label of e is at most 0.

If C contains zu , the label of zu has to leave a remainder of 2 when it is divided by 4 since zu points to the left. As the label is also at most 0, we conclude $l_C(zu) \leq -2$. The edges of P lie either on the left, top or right of f' . Therefore, the label of any edge e on P differs by at most 1 from $l_C(zu)$, and thus we get $l_{C'}(e) \leq 0$.

Summarizing the results above, we see that all edges on C' are labeled with non-positive numbers. The case that all labels of C' are equal to 0 can be excluded, since C

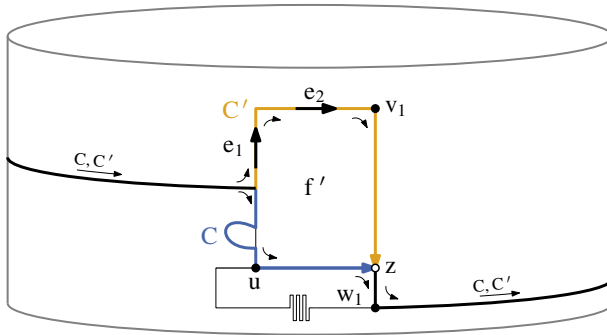


Fig. 31 The increasing cycle C contains uz . There are three possibilities for edges on C' that lie not on C : They lie on the left side of f' (like e_1), on the top (like e_2), or on the right side formed by only the edge v_1z

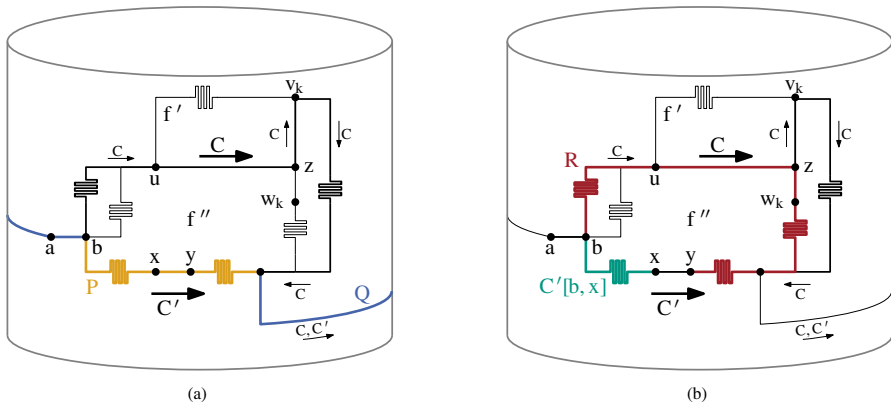


Fig. 32 The situation in the proof of Lemma 7.5. The cycle C is decreasing and it is assumed that $\ell_{C'}(xy) < 0$. **a** The decomposition of C' into P and Q . **b** The reference paths R and $C'[b, x]$ from ab to xy

would not be an increasing cycle by Proposition 5.4. Hence, C' is an increasing cycle, which was already present in Γ , contradicting the validity of Γ . \square

Lemma 7.5 *If u is a horizontal port, then $\Gamma_{e_k}^u$ contains no decreasing cycle.*

Proof Let uz be the new edge inserted in $\Gamma_{e_k}^u$. In $\Gamma_{e_k}^u$, the face f is split in two parts. Let f' be the face containing v_k and f'' the one containing w_k . Assume for the sake of contradiction that there is a decreasing cycle C in $\Gamma_{e_k}^u$. Then, either uz or zu lies on C . By Lemma 5.7 there exists an essential cycle C' that can be decomposed into a path P on f'' and $Q = C \cap C'$; see Fig. 32a. For all edges e of Q , we have $\ell_C(e) = \ell_{C'}(e) \geq 0$ by Lemma 5.7. Since C' is already present in Γ and Γ is valid, C' cannot be decreasing. Moreover, since C and C' intersect, C' cannot be horizontal by Proposition 5.4. Therefore, C' must contain an edge xy with $\ell_{C'}(xy) < 0$, which hence has to lie on P .

Our goal is to show that there must be a candidate on f after y and in particular after the last candidate e_k —a contradiction. The following claim gives a sufficient condition for the existence of such a candidate.

Claim 7.6 If $\text{rot}(f[u, yx]) \leq 2$, then there is a candidate on $f[yx, u]$.

To prove Claim 7.6, we determine for each edge e on f the value $r(e) := \text{rot}(f[u, e])$. By assumption, it is $r(yx) \leq 2$. For the last edge e_{last} on $f[yx, u]$ it is $r(e_{\text{last}}) = \text{rot}(f) - \text{rot}(tuv) = 5$, where t and v are the preceding and succeeding vertices of u on f , respectively. Here, we use that f is a regular face (i.e., $\text{rot}(f) = 4$) and $\text{rot}(tuv) = -1$ since u is a port.

Note that for two consecutive edges e, e' on the boundary of f , it is $r(e') \leq r(e) + 1$. Therefore, there exists an edge e that lies between yx and e_{last} on the boundary of f that satisfies $r(e) = 2$. Hence, e is a candidate that lies after yx on the boundary of f . This finishes the proof of Claim 7.6.

To finish the proof of the lemma it hence suffices to show that $\text{rot}(f[u, yx]) \leq 2$. As e_k is a candidate, we have $\text{rot}(f[u, zw_k]) = 2$ and therefore

$$\text{rot}(f[u, yx]) = \text{rot}(f[u, zw_k]) + \text{rot}(f[zw_k, yx]) = \text{rot}(f[zw_k, yx]) + 2.$$

Thus, it suffices to show $\text{rot}(f[zw_k, yx]) \leq 0$.

We present a detailed argument for the case that C uses uz as illustrated in Fig. 32. At the end of the proof, we briefly outline how the argument can be adapted if C uses zu .

If C uses uz , then P is directed such that f'' lies to the left of P . Thus, C' lies in the interior of C . Let now ab be the last edge of Q , and let R be the path defined by $C[b, uz] + f[zw_k, y]$; see Fig. 32b. Both R and $C'[b, x]$ are reference paths from ab to xy that lie in the interior of C and in the exterior of C' . Applying the second statement of Lemma 3.6 hence gives

$$\text{dir}(ab, R, xy) = \text{dir}(ab, C'[b, x], xy). \tag{5}$$

The direction along R is defined as

$$\begin{aligned} \text{dir}(ab, R, xy) &= \text{rot}(ab + R + yx) - 2 \\ &= \text{rot}(ab + C[b, uz] + f[zw_k, y] + yx) - 2 \\ &= \text{rot}(C[ab, uz] + f[zw_k, yx]) - 2 \\ &= \text{rot}(C[ab, uz]) + \text{rot}(uzw_k) + \text{rot}(f[zw_k, yx]) - 2 \\ &= \ell_C(uz) - \ell_C(ab) + \text{rot}(f[zw_k, yx]) - 1. \end{aligned}$$

The last step uses $\text{rot}(C[ab, uz]) = \ell_C(uz) - \ell_C(ab)$ and $\text{rot}(uzw_k) = 1$.

The rotation along $C'[b, x]$ is defined as

$$\begin{aligned} \text{dir}(ab, C'[b, x], xy) &= \text{rot}(ab + C'[b, x] + xy) \\ &= \text{rot}(C'[ab, xy]) \\ &= \ell_{C'}(xy) - \ell_{C'}(ab) \\ &= \ell_{C'}(xy) - \ell_C(ab). \end{aligned}$$

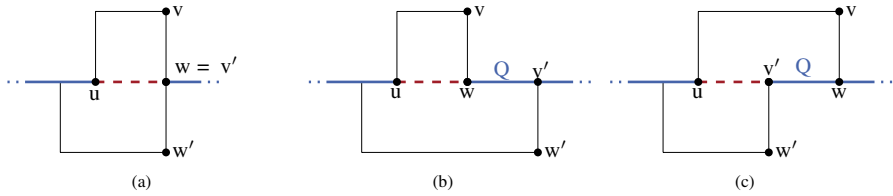


Fig. 33 Three possibilities how the path between w and v' can look like: **a** $w = v'$, **b** all edges point right, and **c** all edges point left. In the first two cases the edge uw is inserted and in **c** uv' is added

The last step uses that $\ell_C(ab) = \ell_{C'}(ab)$ by Lemma 5.7. Altogether, we obtain

$$\ell_C(uz) - \ell_C(ab) + \text{rot}(f[zw_k, yx]) - 1 = \ell_{C'}(xy) - \ell_C(ab),$$

which can be rearranged to

$$\text{rot}(f[zw_k, yx]) = \ell_{C'}(xy) - \ell_C(uz) + 1.$$

With $\ell_{C'}(xy) \leq -1$ and $\ell_C(uz) \geq 0$ we obtain $\text{rot}(f[zw_k, yx]) \leq 0$. This completes the proof for the case that uz lies on C .

If zu lies on C , we consider the flipped representation $\text{flip}(\Gamma_{e_k}^u)$. In $\text{flip}(\Gamma_{e_k}^u)$ the cycle \bar{C} is decreasing and contains the edge uz . The cycle \bar{C}' is not decreasing and contains the edge yx with label $\ell_{\bar{C}'}^f(yx) = \ell_{C'}(xy) < 0$. Moreover, the cycle \bar{C} contains \bar{C}' in its interior. Thus, the argument above can be applied to \bar{C} , \bar{C}' , and yx instead of C , C' , and xy . □

Lemma 7.7 *Let P_i be the maximal path that contains the vertex w_i of the candidate edge $e_i = v_i w_i$ and that consists of only horizontal edges. If u is a horizontal port and $\Gamma_{e_i}^u$ contains a decreasing cycle and $\Gamma_{e_{i+1}}^u$ contains an increasing cycle, then u is an endpoint of P_i and adding the horizontal edge uz to the other endpoint z of P_i yields a horizontal cycle. In particular, $\Gamma + uz$ is valid.*

Proof Figure 33 illustrates three possibilities of the path. Let z_i be the new vertex inserted in $\Gamma_{e_i}^u$ and z_{i+1} the one in $\Gamma_{e_{i+1}}^u$. Since both uz_i and uz_{i+1} point to the right, there is no augmentation of Γ containing both edges. We compare $\Gamma_{e_i}^u$ and $\Gamma_{e_{i+1}}^u$ by the following construction (see also Fig. 34), which models all important aspects of both representations: Starting from Γ we insert new vertices z_i on e_i and z_{i+1} on e_{i+1} . We connect u and z_i by a path of length 2 that points to the right and denote its internal vertex by x . Furthermore, a path of length 2 from x via a new vertex y to z_{i+1} is added. The edge xy points down and yz_{i+1} to the right. In the resulting ortho-radial representation $\tilde{\Gamma}$ the edge uz_i in $\Gamma_{e_i}^u$ is modeled by the path uxz_i . Similarly, the edge uz_{i+1} in $\Gamma_{e_{i+1}}^u$ is modeled by the path $uxyz_{i+1}$ in $\tilde{\Gamma}$.

Take any decreasing cycle in $\Gamma_{e_i}^u$. As Γ is valid, this cycle must contain either uz_i or $z_i u$. We obtain a cycle C_1 in $\tilde{\Gamma}$ by replacing uz_i with uxz_i (or $z_i u$ with $z_i x u$). Note that ux and xz_i have the same label as uz_i , and the labels of all other edges on the cycles stay the same by Lemma 5.8. Therefore, C_1 is a decreasing cycle.

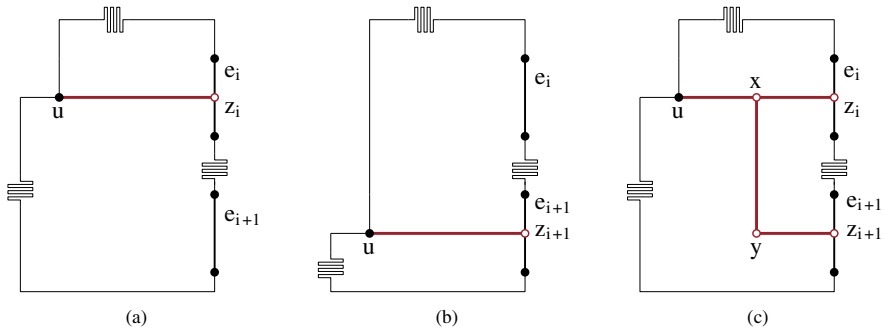


Fig. 34 Illustration of proof for Lemma 7.7. **a** $\Gamma_{e_i}^u$ is obtained by inserting the edge uz_i into Γ . **b** $\Gamma_{e_{i+1}}^u$ is obtained by inserting the edge uz_{i+1} into Γ . **c** $\tilde{\Gamma}$ is obtained by combining $\Gamma_{e_i}^u$ and $\Gamma_{e_{i+1}}^u$

Similarly, there exists an increasing cycle in $\Gamma_{e_{i+1}}^u$, which contains uz_{i+1} or $z_{i+1}u$. Replacing uz_{i+1} with $uxyz_{i+1}$ (or $z_{i+1}u$ with $z_{i+1}yxu$) we get a cycle C_2 in $\tilde{\Gamma}$. By Lemma 5.8 the labels of C_2 are non-positive outside of $uxyz_{i+1}$ (or $z_{i+1}yxu$) and there is an edge with negative label. Consequently, the only edge of C_2 that may have a positive label is the edge e between x and y . Since C_1 and C_2 intersect, Lemma 5.6 implies that C_2 is not increasing. Thus, the label of e is positive. If $e = yx$, then e points upwards. Hence, $\ell_{C_2}(e) \equiv 3 \pmod 4$, which implies that $\ell_{C_2}(e) \geq 3$. Thus, its succeeding edge has a positive label as well contradicting that e is the only edge on C_2 with a positive label. Hence, the cycle C_2 contains the edge xy and consequently the edge ux .

Using this construction we show that one endpoint of P_i is u and the other is another vertex of f . To that end, we prove the following claims.

Claim 7.8 The cycles C_1 and C_2 both contain the edge ux and it is $\ell_{C_1}(ux) = \ell_{C_2}(ux) = 0$.

Claim 7.9 The vertices w_i and v_{i+1} have a degree of at least 2. Further, C_1 contains w_i and C_2 contains v_{i+1} .

Claim 7.10 The edges of $f[w_i, v_{i+1}]$ are part of P_i . In particular, C_2 contains w_i or C_1 contains v_{i+1} .

Claim 7.11 The right endpoint of P_i is u and the left endpoint z lies on f .

Using these claims we prove the lemma as follows. Due to Claim 7.11 we can insert the horizontal edge uz into Γ obtaining a horizontal cycle $P_i + uz$. The resulting ortho-radial representation is valid, because any strictly monotone cycle necessarily contains uz and hence shares a vertex with a horizontal cycle, contradicting Lemma 5.4. This finishes the proof. In the following we prove the claims.

Claim 7.8. Let $H = C_1 + C_2$ be the graph formed by the two cycles C_1 and C_2 and let \tilde{f} be its central face. We first show that ux is incident to \tilde{f} . If all edges of \tilde{f} belong to one of the two cycles, say C_1 , then $\tilde{f} = C_1$ since C_1 is simple. Thus, ux is incident to \tilde{f} . Otherwise, there is an edge e of $C_2 \setminus C_1$ on \tilde{f} whose source v lies on C_1 ; see

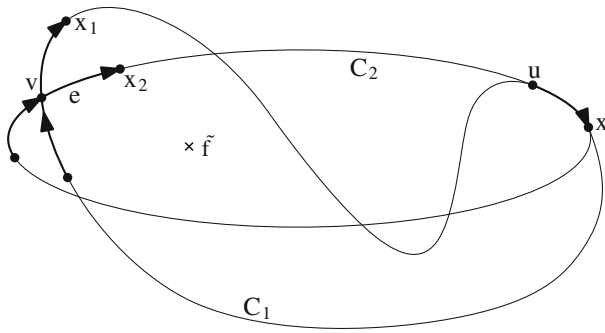


Fig. 35 Illustration of proof for Claim 7.8 of Lemma 7.7

Fig. 35. Let x_1 and x_2 be the succeeding vertices of v on C_1 and C_2 , respectively. As C_1 is a decreasing cycle, we have $\ell_{C_1}(vx_1) \geq 0$. If $\ell_{C_2}(vx_2) \leq 0$, then by Lemma 5.5 the edge $e = vx_2$ lies in the exterior of C_1 , which contradicts the choice of e . Hence, we have $\ell_{C_2}(vx_2) > 0$, which implies that $e = xy$. Further, the cyclic order of the vertex x implies that ux is the predecessor of xy on the central face.

Since C_2 contains ux , the central face \hat{f} lies to the right of ux . Since C_1 is also directed such that \hat{f} lies to the right of C_1 , it cannot contain xu , but it contains ux . By Lemma 5.2 we further obtain that $\ell_{C_1}(ux) = \ell_{C_2}(ux) = 0$.

Claim 7.9. We consider the setting in $\tilde{\Gamma}$; see Fig. 34c. From Claim 7.8 we obtain $\ell_{C_1}(xz_i) = 0$. Consequently, since C_1 is a decreasing cycle and since $\text{rot}(xz_i w_i) = 1$ and $\text{rot}(xz_i v_i) = -1$, the cycle contains the edge $z_i w_i$ but not the edge $z_i v_i$. In particular, this implies that w_i has a degree of at least 2, as otherwise C_1 would not be simple. Similarly, from Claim 7.8 we obtain $\ell_{C_1}(yz_{i+1}) = 0$. Consequently, since C_2 has only non-positive labels except on xy and since $\text{rot}(xz_{i+1} v_{i+1}) = -1$ and $\text{rot}(xz_{i+1} w_{i+1}) = 1$, the cycle contains the edge $z_{i+1} v_{i+1}$ but not the edge $z_{i+1} w_{i+1}$. In particular, this implies that v_{i+1} has a degree of at least 2, as otherwise C_2 would not be simple.

Claim 7.10. Let e_s be the direct successor of e_i and let e_p be the direct predecessor of e_{i+1} on the boundary of f . In order to show the claim, we do a case distinction on the rotations $\text{rot}(e_i, e_s)$ and $\text{rot}(e_p, e_{i+1})$. From Claim 7.9 it follows that $\text{rot}(e_i, e_s) \in \{-1, 0, 1\}$ and $\text{rot}(e_p, e_{i+1}) \in \{-1, 0, 1\}$. Further, it cannot be both $\text{rot}(e_i, e_f) = 1$ and $\text{rot}(e_p, e_{i+1}) = 1$. Otherwise, since $\text{rot}(f[u, e_i]) = 2$ and $\text{rot}(f[u, e_{i+1}]) = 2$, it would hold $\text{rot}(f[u, e_s]) = 3$ and $\text{rot}(f[u, e_p]) = 1$. In that case, there would be a further candidate edge e with $\text{rot}(f[u, e]) = 2$ between e_i and e_{i+1} . Thus, we obtain $\text{rot}(e_i, e_s) \in \{-1, 0\}$ or $\text{rot}(e_p, e_{i+1}) \in \{-1, 0\}$.

First, assume that $\text{rot}(e_i, e_s) = 0$. We obtain $e_s = e_{i+1}$ and hence $e_p = e_i$. Otherwise, e_s would be a candidate edge between e_i and e_{i+1} . Consequently, $f[w_i, v_{i+1}]$, which consist of the single vertex $v = w_i = v_{i+1}$, trivially belongs to P_i . Moreover, by Claim 7.8 the cycle C_2 then also contains w_i and C_1 contains v_{i+1} . The very same argument can be applied for $\text{rot}(e_p, e_{i+1}) = 0$, as we also obtain $e_s = e_{i+1}$ and $e_p = e_i$. In particular, we obtain $\text{rot}(e_i, e_s) = 0$ if and only if $\text{rot}(e_p, e_{i+1}) = 0$.

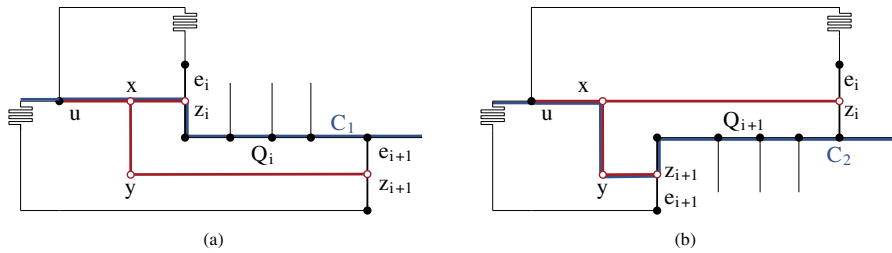


Fig. 36 Illustration of proof for Claim 7.10 of Lemma 7.7

Now, assume that $\text{rot}(e_i, e_s) = -1$ or $\text{rot}(e_p, e_{i+1}) = -1$. We first consider the case that $\text{rot}(e_i, e_s) = -1$; see Fig. 36a. We observe that C_1 contains e_s as it also contains the vertex w_i by Claim 7.8. Let Q_i be the maximal horizontal path on f that starts at w_i . We observe that Q_i is a prefix of P_i , and the cycle C_1 shares at least e_s with Q_i . Since by Claim 7.8 the cycle C_1 has label 0 on ux , it also has label 0 on any of its edges that lie on Q_i . In fact, this implies that C_1 completely contains the path Q_i . Otherwise, it would contain an edge that does not belong Q_i , but starts at an intermediate vertex of Q_i . By the choice of Q_i such an edge has a negative label since its preceding edge has label 0, which contradicts that C_1 is decreasing. For the same reason, the edge e that succeeds Q_i on f has a negative label. Hence, the edge e is the candidate e_{i+1} . Consequently, C_1 contains v_{i+1} and $f[w_i, v_{i+1}] = Q_i$, which shows Claim 7.10. For the case that $\text{rot}(e_p, e_{i+1}) = -1$ we use similar arguments; see Fig. 36b. Let Q_{i+1} be the maximal horizontal path on f that starts at v_{i+1} when going along f counterclockwise. The first edge of Q_{i+1} is e_p , which also belongs to C_2 . Further, C_2 has label 0 on any of its edges that belong to Q_{i+1} . Consequently, C_2 contains Q_{i+1} completely, as otherwise it would contain an edge that points downwards, contradicting that C_2 has only non-positive labels except for the edge xy . Hence, the edge of f that is incident with the endpoint of Q_{i+1} and does not belong to Q_{i+1} is the candidate edge e_i , which concludes the proof.

Claim 7.11. We first show that u is the right endpoint of P_i . To that end, let $H = C_1 + C_2$ be the graph formed by the two cycles C_1 and C_2 and let \tilde{f} be its central face. From the proof of Claim 7.8 we know that ux is incident to the central face. Further, from Claims 7.9 and 7.10 it follows that C_1 and C_2 have the vertex w_i or the vertex v_{i+1} in common; see Figs. 36a and 36b, respectively. Let $v = w_i$ in the former case and let $v = v_{i+1}$ in the latter case.

We first show that $C_1[v, x] = C_2[v, x]$. Assume that this is not the case. Consider the maximal common prefix of $C_1[v, x]$ and $C_2[v, x]$, and let w be the endpoint of that prefix. As v is incident to \tilde{f} , this also applies to w . Further, by Lemma 5.5 the outgoing edge of w on C_2 lies in the exterior of C_1 . Let t be the first intersection of C_1 and C_2 on \tilde{f} after w . Then the edge to t on C_2 lies strictly in the exterior of C_1 contradicting Lemma 5.5.

Hence, we have $C_1[v, x] = C_2[v, x]$; we denote that path by P . Since ux is incident to \tilde{f} , all edges of P are incident to \tilde{f} as well. Hence, by Corollary 5.3 all edges of P have the same label on C_1 and C_2 . Since C_1 is decreasing, i.e., its labels are non-negative, and the labels on C_2 are non-positive on $C_2[z_{i+1}, u]$, the labels of the edges

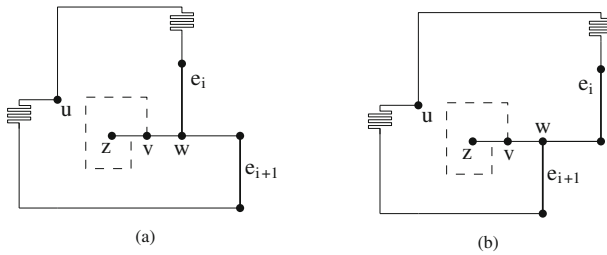


Fig. 37 Illustration of proof for Claim 7.11 of Lemma 7.7

of P are 0. Therefore, these edges are horizontal. As $f[w_i, v_{i+1}]$ also belongs to P_i , we conclude that u is the right endpoint of P_i .

Finally, we argue that z belongs to the face f . As by Claim 7.10, $f[w_i, v_{i+1}]$ belongs to P_i , the vertex z belongs to a prefix Q of P_i that does not contain any edge of $f[w_i, v_{i+1}]$ and ends either at w_i or at v_{i+1} ; see Figs. 37a and 37b, respectively. If Q is empty, w_i or v_{i+1} is z , respectively, which shows the claim. So assume that Q is not empty. Let vw be the last edge of Q , i.e., depending on the considered case we either have $w = w_i$ or $w = v_{i+1}$. By the local ordering of the incident edges of w , the path Q lies in the interior of f . Further, vw is incident to f on both sides, as otherwise the two other incident edges of w could not be incident to f as well. Hence, we obtain $\text{rot}(f[u, vw]) = 3$ and $\text{rot}(f[u, v]) = 0$. Thus, the path $f[wv, vw]$ of f consists only of horizontal edges, as otherwise there would be an edge e on that path with $\text{rot}(f[u, e]) = 2$, which contradicts the choice of e_i and e_{i+1} . Altogether, the path $f[wv, vw]$ is $\bar{Q} + Q$, which proves that z belongs to f . \square

7.3 Proof of the Main Theorem

We are now ready to prove the characterization of drawable ortho-radial representations.

Theorem 7.12 *An ortho-radial representation is drawable if and only if it is valid.*

For an illustration of the proof see Fig. 38. If Γ is a valid ortho-radial representation, then using the rectangulation procedure of Sect. 7.1 we obtain a valid rectangular ortho-radial representation Γ' that contains a subdivision of Γ . Then Γ' has a drawing Δ' by Theorem 6.3, and a drawing Δ of Γ is obtained from Δ' by removing edges that are in Γ' but not in Γ , and by undoing subdivisions. It remains to show that the reference edge e^* of Γ is an outlying edge in Δ , which then implies that Γ is drawable.

Let C_o be the essential cycle that is part of the outer face of Γ and let e_o be an outermost horizontal edge of C_o in Δ , i.e., an horizontal edge of C_o with maximum distance to the center; see Fig. 38a. To show that e^* is an outlying edge in Δ , we show $\text{rot}(C_o[e_o, e^*]) = 0$.

Let C_T be the essential cycle that forms the outer face of Γ' . Further, let l be the ray that emanates from the center of Δ' and goes through the source vertex v of e_o . Let p_1, \dots, p_k be the intersection points of l with Δ' from v on. We observe that $p_1 = v$ and p_k lies on C_T . For each intersection point p_i with $1 \leq i \leq k$ we insert a vertex v_i

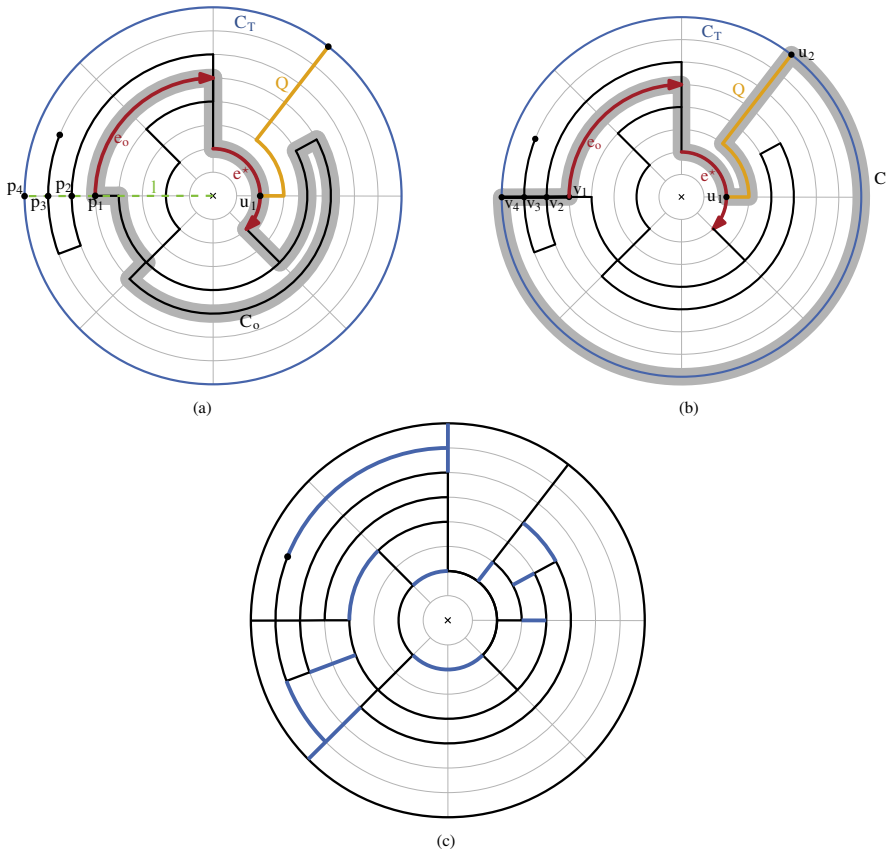


Fig. 38 Illustration of proof for Theorem 3.9

at p_i if there is no vertex so far. Further, we add vertical edges to Δ' such that the new drawing contains the path $P = v_1 \dots v_k$. We denote the resulting drawing by Δ'' ; see Fig. 38b.

By the construction of the rectangulation the drawing Δ'' further contains a path Q that connects a subdivision vertex u_1 on e^* with a vertex u_2 on C_T . We observe that the cycle $C = C_0[v_1, u_1] + Q + T[u_2, v_k] + \overline{P}$ is essential. Moreover, by construction the path $C[u_1, v_1] = Q + C_T[u_2, v_k] + \overline{P}$ has rotation 2. Since C has a left bend at both v_1 and u_1 it follows that $C_0[v_1, u_1]$ has rotation 0. Consequently, e^* is an outlying edge in Δ . Altogether, we obtain that Γ is drawable.

Conversely, assume that Γ is drawable. Hence, Γ has a drawing Δ in which the reference edge e^* of Γ is an outlying edge. By Lemma 4.2 we can assume without loss of generality that e^* lies on the outermost circle that is used by an essential cycle of Δ . By [21] the representation Γ satisfies conditions (i) and (ii) of Definition 3.1. To prove that Γ is valid, we show how to reduce the general case to the more restricted one, where all faces are rectangles. By Corollary 6.3 the existence of a drawing and the validity of the ortho-radial representation are equivalent.

Given the drawing Δ , we augment it such that all faces are rectangles. This rectangulation is similar to the one described in Sect. 7.1 but works with a drawing and not only with a representation. We first insert the missing parts of the innermost and outermost circle that are used by Δ such that the outer and the central face are already rectangles. For each left turn on a face f at a vertex u , we then cast a ray from v in f in the direction in which the incoming edge of u points. This ray intersects another edge in Δ . Say the first intersection occurs at the point p . Either there already is a vertex z drawn at p or p lies on an edge. In the latter case, we insert a new vertex, which we call z , at p . We then insert the edge uz in G and update Δ and Γ accordingly.

Repeating this step for all left turns, we obtain a drawing Δ' and an ortho-radial representation Γ' of the augmented graph G' ; see Fig. 38c for an example of Δ' . As the labelings of essential cycles are unchanged by the addition of edges elsewhere in the graph, any increasing or decreasing cycle in Γ would also appear in Γ' . But by Corollary 6.3 Γ' is valid, and hence neither Γ nor Γ' contain increasing or decreasing cycles. Thus, Γ is valid. \square

We remark that the proof of Theorem 3.9 is in fact constructive and it can easily be implemented in polynomial time, provided that we can check in polynomial time whether a given ortho-radial representation is valid. We develop such an algorithm with running time $\mathcal{O}(n^2)$ in Sect. 8. In Sect. 9 we show how compute a rectangulation of Γ in time $\mathcal{O}(n^2)$.

8 Validity Testing

In this section, we show how to check whether a given ortho-radial representation is valid in polynomial time, which yields the following statement.

Theorem 8.1 *Given an ortho-radial representation Γ , it can be determined in $\mathcal{O}(n^2)$ time whether Γ is valid. In the negative case a strictly monotone cycle can be computed in $\mathcal{O}(n^2)$ time.*

The two conditions for ortho-radial representations are local and checking them can easily be done in linear time. Throughout this section, we therefore assume that we are given an instance $I = (G, \mathcal{E}, f_c, f_o)$ with an ortho-radial representation Γ and a reference edge e^* . The condition for validity however references all essential cycles of which there may be exponentially many. We present an algorithm that checks whether Γ contains a strictly monotone cycle and computes such a cycle if one exists. The main difficulty is that the labels on a decreasing cycle C depend on a reference path P that runs from e^* to C and respects C . However, we know neither the path P nor the cycle C in advance, and choosing a specific cycle C may rule out certain paths P and vice versa.

We only describe how to search for decreasing cycles; increasing cycles can be found by searching for decreasing cycles in the mirrored representation by Lemma 4.4. A decreasing cycle C is *outermost* if it is not contained in the interior of any other decreasing cycle. Clearly, if Γ contains a decreasing cycle, then it also has an outermost one. We first show that in this case this cycle is uniquely determined.

Lemma 8.2 *If Γ contains a decreasing cycle, there is a unique outermost decreasing cycle.*

Proof Assume that Γ has two outermost decreasing cycles C_1 and C_2 , i.e., C_1 does not lie in the interior of C_2 and vice versa. Let C be the cycle bounding the outer face of the subgraph $H = C_1 + C_2$ that is formed by the two decreasing cycles. By construction, C_1 and C_2 lie in the interior of C , and we claim that C is a decreasing cycle contradicting that C_1 and C_2 are outermost. To that end, we show that $\ell_C(e) = \ell_{C_1}(e)$ for any edge e that belongs to both C and C_1 , and $\ell_C(e) = \ell_{C_2}(e)$ for any edge e that belongs to both C and C_2 . Hence, all edges of C have a non-negative label since C_1 and C_2 are decreasing. By Proposition 5.4 there is at least one label of C that is positive, and hence C is a decreasing cycle.

It remains to show that $\ell_C(e) = \ell_{C_1}(e)$ for any edge e that belongs to both C and C_1 ; the case that e belongs to both C and C_2 can be handled analogously. Let Γ_H be the ortho-radial representation Γ restricted to H . We flip the cylinder to exchange the outer face with the central face and vice versa. More precisely, Lemma 4.3 implies that the reverse edge \bar{e} of e lies on the central face of the flipped representation $\text{flip}(\Gamma_H)$ of Γ_H . Further, it proves that $\ell_C^f(\bar{e}) = \ell_C(e)$ and $\ell_{C_1}^f(\bar{e}) = \ell_{C_1}(e)$, where ℓ^f is the labeling in Γ_H . Hence, by Corollary 5.3 we obtain $\ell_C^f(\bar{e}) = \ell_{C_1}^f(\bar{e})$. Flipping back the cylinder, again by Lemma 4.3 we obtain $\ell_C(e) = \ell_{C_1}(e)$. \square

The core of our algorithm is an adapted left-first DFS. Given a directed edge e it determines the outermost decreasing cycle C in Γ such that C contains e in the given direction and e has the smallest label among all edges on C , if such a cycle exists. By running this test for each directed edge of G as the start edge, we find a decreasing cycle if one exists.

More precisely, the DFS visits each vertex at most once and it maintains for each visited vertex v a reference edge $\text{ref}(v)$, the edge of the search tree via which v was visited. Whenever it has a choice which vertex to visit next, it picks the first outgoing edge in clockwise direction after the reference edge that leads to an unvisited vertex. In addition to that, we employ a filter that ignores certain outgoing edges during the search. To that end, we define for all outgoing edges e incident to a visited vertex v a search label $\tilde{\ell}(e)$ by setting $\tilde{\ell}(e) = \tilde{\ell}(\text{ref}(v)) + \text{rot}(\text{ref}(v) + e)$ for each outgoing edge e of v . In our search we ignore edges with negative search labels. For a given directed edge vw in G we initialize the search by setting $\text{ref}(w) = vw$, $\tilde{\ell}(vw) = 0$ and then start searching from w .

Let T denote the directed search tree with root w constructed by the DFS in this fashion. If T contains v , then this determines a candidate cycle C containing the edge vw . If C is a decreasing cycle, which we can easily check by determining a reference path, we report it. Otherwise, we show that there is no outermost decreasing cycle C such that vw lies on C and has the smallest label among all edges on C .

It is necessary to check that C is essential and decreasing. For example the cycle in Fig. 39 is found by the search and though it is essential, it is non-decreasing. This is caused by the fact that the label of vw is actually -4 on this cycle but the search assumes it to be 0.

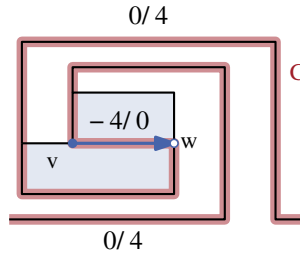


Fig. 39 The search from vw finds the non-decreasing cycle C . Edges are labeled $\ell_C(e)/\tilde{\ell}(e)$

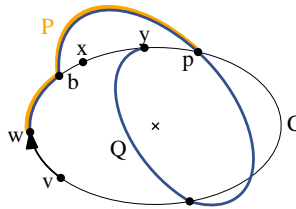


Fig. 40 Path Q and its prefix P that leaves C once and ends at a vertex p of C

Lemma 8.3 Assume that Γ contains a decreasing cycle. Let C be the outermost decreasing cycle of Γ and let vw be an edge on C with the minimum label, i.e., $\ell_C(vw) \leq \ell_C(e)$ for all edges e of C . Then the left-first DFS from vw finds C .

Proof Assume for the sake of contradiction that the search does not find C . Let T be the tree formed by the edges visited by the search. Since the search does not find C by assumption, a part of $C[w, v]$ does not belong to T . Let xy be the first edge on $C[w, v]$ that is not visited, i.e., $C[w, x]$ is a figure part of T but $xy \notin T$. There are two possible reasons for this. Either $\tilde{\ell}(xy) < 0$ or y has already been visited before via another figure from w with $Q \neq C[w, y]$.

The case $\tilde{\ell}(xy) < 0$ can be excluded as follows. By the construction of the labels $\tilde{\ell}$, for any path P from w to a vertex z in T and any edge e' incident to z we have $\tilde{\ell}(e') = \text{rot}(vw + P + e')$. In particular, $\tilde{\ell}(xy) = \text{rot}(C[vw, xy]) = \ell_C(xy) - \ell_C(vw) \geq 0$ since the rotation can be rewritten as a label difference (see Observation 5.1) and vw has the smallest label on C .

Hence, T contains a path Q from w to x that was found by the search before and Q does not completely lie on C . There is a prefix of Q (possibly of length 0) lying on C followed by a subpath not on C until the first vertex p of Q that again belongs to C ; see Fig. 40. We set $P = Q[w, p]$ and denote the vertex where P leaves C by b . By construction the edge vw lies on $C[p, b]$. The subgraph $H = P + C$ that is formed by the decreasing cycle C and the path P consists of the three internally vertex-disjoint paths $P[b, p]$, $C[b, p]$ and $\bar{C}[b, p]$ between b and p . Since edges that are further left are preferred during the search, the clockwise order of these paths around b and p is fixed as in Fig. 40. In H there are three faces, bounded by C , $\bar{C}[b, p] + \bar{P}[p, b]$ and $P[b, p] + \bar{C}[p, b]$, respectively. Since C is an essential cycle and it bounds a face in H , it bounds the central face and one of the two other faces is the outer face. These



(a) The edge vw lies on the outer face of H (b) The edge vw does not lie on the outer face of H

Fig. 41 The two possible embeddings of the subgraph formed by the decreasing cycle C and the path P , which was found by the search

two possibilities are shown in Fig. 41. We denote the cycle bounding the outer face but in which the edges are directed such that the outer face lies locally to the left by C' . That is, the boundary of the outer face is $\overline{C'}$. We distinguish cases based on which of the two possible cycles constitutes $\overline{C'}$.

If $\overline{C'} = \overline{C}[b, p] + \overline{P}[p, b]$ forms the outer face of H , vw lies on C' as illustrated in Fig. 41a and we show that C' is a decreasing cycle, which contradicts the assumption that C is the outermost decreasing cycle. Since P is simple and lies in the exterior of C , the path P is contained in C' , which means $C'[w, p] = P$. The other part of C' is formed by $C[p, w]$. Since C forms the central face of H , the labels of the edges on $C[p, w]$ are the same for C and C' by Proposition 5.2. In particular, $\ell_C(vw) = \ell_{C'}(vw)$ and all the labels of edges on $C[p, w]$ are non-negative because C is decreasing. The label of any edge e on both C' and P is $\ell_{C'}(e) = \ell_C(vw) + \text{rot}(vw + P[w, e]) = \ell_C(vw) + \ell(e) \geq 0$. Thus, the labeling of C' is non-negative. Further, not all labels of C' are 0 since otherwise C would not be a decreasing cycle by Proposition 5.4. Hence, C' is decreasing and contains C in its interior, a contradiction.

If $\overline{C'} = \overline{C}[p, b] + P[b, p]$, the edge vw does not lie on C' ; see Fig. 41b. We show that C' is a decreasing cycle containing C in its interior, again contradicting the choice of C . As above, Proposition 5.2 implies that the common edges of C and C' have the same labels on both cycles. It remains to show that all edges xy on $\overline{P}[p, b]$ have non-negative labels. To establish this we use paths to the edge that follows b on C . This edge bc has the same label on both cycles and thus provides a handle on $\ell_{C'}(xy)$. We make use of the following equations, which follow immediately from the definition of the (search) labels.

$$\begin{aligned} \ell_{C'}(bc) &= \ell_{C'}(xy) + \text{rot}(\overline{P}[xy, db]) + \text{rot}(dbc) \\ &= \ell_{C'}(xy) - \text{rot}(P[bd, yx]) - \text{rot}(cbd) \\ \ell_C(bc) &= \ell_C(vw) + \text{rot}(C[vw, ab]) + \text{rot}(abc) \\ \tilde{\ell}(yx) &= \text{rot}(C[vw, ab]) + \text{rot}(abd) + \text{rot}(P[bd, yx]). \end{aligned}$$

Since $\ell_C(bc) = \ell_{C'}(bc)$ and $\text{rot}(abd) = -\text{rot}(dba)$, we thus get

$$\begin{aligned} \ell_{C'}(xy) &= \ell_C(vw) + \text{rot}(C[vw, ab]) + \text{rot}(abc) + \text{rot}(P[bd, yx]) + \text{rot}(cbd) \\ &= \ell_C(vw) + \tilde{\ell}(yx) + \text{rot}(dba) + \text{rot}(abc) + \text{rot}(cbd). \end{aligned}$$

Since $\ell_C(vw) \geq 0$ and $\tilde{\ell}(yx) \geq 0$ (as yx was not filtered out), it follows that $\ell_{C'}(xy) \geq \text{rot}(abc) + \text{rot}(dba) + \text{rot}(cbd) = 2$ as this is the sum of clockwise rotations around a degree-3 vertex. Hence, C' is decreasing and contains C in its interior, a contradiction. Since both embeddings of H lead to a contradiction, we obtain a contradiction to our initial assumption that the search fails to find C . \square

The left-first DFS clearly runs in $\mathcal{O}(n)$ time. In order to guarantee that the search finds a decreasing cycle if one exists, we run it for each of the $\mathcal{O}(n)$ directed edges of G . Since some edge must have the lowest label on the outermost decreasing cycle, Lemma 8.3 guarantees that we eventually find a decreasing cycle if one exists. Increasing cycles can be found by searching for decreasing cycles in the mirror representation $\text{mirror}(\Gamma)$ (Lemma 4.4). Altogether, this proves Theorem 3.10.

9 Efficient Rectangulation Procedure

Let G be a planar 4-graph with valid ortho-radial representation Γ . By Theorem 3.9 Γ is drawable. The proof of Theorem 3.9 is constructive and shows how to augment Γ to a valid rectangular ortho-radial representation Γ^* . Then a drawing Δ^* of Γ^* can be computed by determining flows in two flow networks by Theorem 6.1. A drawing Δ of Γ can be finally obtained by undoing augmentation steps. Hence, it remains to show how a valid ortho-radial representation can be rectangulated efficiently.

To follow the construction of the proof of Theorem 3.9, one needs to determine efficiently whether an augmentation yields a valid ortho-radial representation Γ_{vw}^u . We call such an augmentation a *valid augmentation*. Since each valid augmentation reduces the number of concave angles, we obtain a rectangulation after $\mathcal{O}(n)$ valid augmentations. Moreover, there are $\mathcal{O}(n)$ candidates for each augmentation, each of which can be tested for validity (and increasing/decreasing cycles can be detected) in $\mathcal{O}(n^2)$ time by Theorem 3.10. Thus, the rectangulation algorithm can be implemented to run in $\mathcal{O}(n^4)$ time.

In the remainder of this section we present an improvement to $\mathcal{O}(n^2)$ time, which is achieved in three steps. First, we show that, due to the nature of augmentations, each validity test can be done in $\mathcal{O}(n)$ time. This reduces the running time of the augmentation procedure to $\mathcal{O}(n^3)$; see Sect. 9.1. Second, we show how to find a valid augmentation for a given port u using only $\mathcal{O}(\log n)$ validity tests, thus improving the running time to $\mathcal{O}(n^2 \log n)$; see Sect. 9.2. Finally, we design an algorithm that can be used to post-process a valid augmentation and which reduces the number of validity tests to $\mathcal{O}(n)$ and the running time to $\mathcal{O}(n^2)$ in total; see Sect. 9.3. Altogether, this proves our third main result.

Theorem 9.1 *Given a valid ortho-radial representation, a corresponding drawing can be constructed in $\mathcal{O}(n^2)$ time.*

9.1 1st Improvement—Faster Validity Test

The general test for strictly monotone cycles performs one left-first DFS per edge and runs in $\mathcal{O}(n^2)$ time. However, we can exploit the special structure of the augmentation

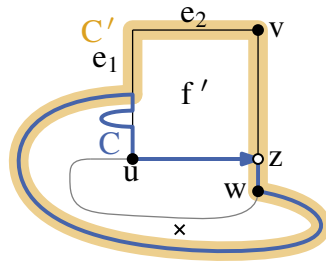


Fig. 42 A decreasing cycle C that uses uz and an essential cycle C' derived from C

to reduce the running time to $\mathcal{O}(n)$. For the proof we restrict ourselves to the case that the inserted edge uz points to the right. The case that it points left can be handled by flipping the representation using Lemma 4.3.

The key result is that in any decreasing cycle of an augmentation the new edge uz has the minimum label. Thus, performing only one left-first DFS starting at uz is sufficient. For increasing cycles the arguments do not hold, but in a second step we show that the test for increasing cycles can be replaced by a simple test for horizontal paths.

Recall that the augmentations Γ_{vw}^u that are tested during the rectangulation are built by adding one edge uz to a valid representation Γ . Hence, any strictly monotone cycle in Γ_{vw}^u contains the edge uz .

We first show that the new edge uz has label 0 on any decreasing cycle in the augmentation Γ_{vw}^u if vw is the first candidate. We extend this result afterwards to augmentations to all candidates. Since the label of edges on decreasing cycles is non-negative, this implies in particular that the label of uz is minimum, which is sufficient for the left-first DFS to succeed (see Lemma 8.3).

Lemma 9.2 *Let Γ be a valid ortho-radial representation and let u be a horizontal port on f with first candidate vw . If Γ_{vw}^u contains a decreasing cycle C , then C contains uz in this direction and $\ell_C(uz) = 0$.*

Proof We first consider the case that C uses uz (and not zu) and assume for the sake of contradiction that $\ell_C(uz) \neq 0$; see Fig. 42. Since uz points right, $\ell_C(uz)$ is divisible by 4. Together with $\ell_C(uz) \geq 0$ because C is decreasing, we obtain $\ell_C(uz) \geq 4$. By Lemma 5.7 there is an essential cycle C' without uz in the subgraph H that is formed by the new rectangular face f' and C . We show that C' is a decreasing cycle. We observe that each edge of C' either lies on f' or is an edge of C . For any edge e of $C' \cap C$ Lemma 5.7 gives $\ell_{C'}(e) = \ell_C(e) \geq 0$. Since f' is rectangular, the labels of edges in $C' \cap f'$ differ by at most 1 from $\ell_C(uz)$. By assumption it is $\ell_C(uz) \geq 4$ and therefore $\ell_{C'}(e) \geq 3$ for all edges e in $C' \cap f'$. Hence, C' is a decreasing cycle in G contradicting the validity of Γ .

If zu lies on C , it is $\ell_C(zu) \geq 2$ and a similar argument yields a decreasing cycle in Γ . □

While the same statement does not generally hold for all candidates, it does hold if the first candidate creates a decreasing cycle.

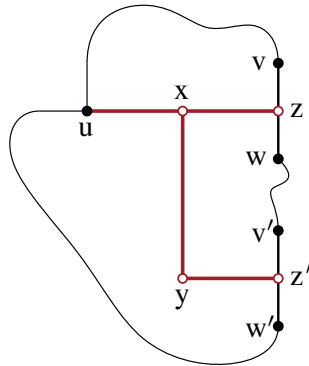


Fig. 43 The structure used to simulate the simultaneous insertion of uz to vw and uz' to $v'w'$

Lemma 9.3 *Let Γ be a valid ortho-radial representation and let u be a horizontal port on f with first candidate vw . Further, let $v'w'$ be another candidate and denote the edge inserted in $\Gamma_{v'w'}^u$ by uz' . If Γ_{vw}^u contains a decreasing cycle, any decreasing cycle C' in $\Gamma_{v'w'}^u$ uses uz' in this direction and $\ell_{C'}(uz') = 0$.*

Proof In order to simulate the simultaneous insertion of two new edges to both vw and $v'w'$ we use the structure from the proof of Lemma 7.7; see Fig. 43. We denote the resulting augmented representation by $\tilde{\Gamma}$. There is a one-to-one correspondence between decreasing cycles in Γ_{vw}^u and decreasing cycles in $\tilde{\Gamma}$ containing uxz . Let C be a decreasing cycle in $\tilde{\Gamma}$ containing uxz . By Lemma 9.2 the cycle C contains uxz in this direction, and we have $\ell_C(ux) = 0$.

Similarly, for any decreasing cycle in $\Gamma_{v'w'}^u$, there is a cycle in $\tilde{\Gamma}$ where uz' ($z'u$) is replaced by the path $uxyz'$ ($z'yxu$). Let C' be the cycle in $\tilde{\Gamma}$ that corresponds to the decreasing cycle C' in $\Gamma_{v'w'}^u$.

Suppose for now that C' uses uz' in this direction, which means that \tilde{C}' uses ux . In particular, since $\ell_{\tilde{C}'}(ux) = \ell_{C'}(uz')$, it holds $\ell_{\tilde{C}'}(xy) = \ell_{\tilde{C}'}(ux) + 1 \geq 1$. Let \tilde{f} be the central face of $H = \tilde{C} + \tilde{C}'$. We note that \tilde{f} is a decreasing cycle by Lemma 5.2 and Proposition 5.4. Since Γ is valid, \tilde{f} is not exclusively formed by edges of G . Thus, by the construction of H , the path $uxyz'$ lies on \tilde{f} . Lemma 5.2 therefore implies that $\ell_{\tilde{C}'}(ux) = \ell_C(ux) = 0$, where the last equality follows from Lemma 9.2.

Above we assumed that C' uses uz' in this direction. This is in fact the only possibility. Assume for the sake of contradiction that C' contains $z'u$ and hence $xu \in \tilde{C}'$. As above we can argue that the central face \tilde{f} of $\tilde{C} + \tilde{C}'$ is not exclusively formed by edges of G and that the path $z'yxz$ lies on \tilde{f} . By Lemma 5.2 we have $\ell_{\tilde{C}'}(yx) = \ell_C(xz) - \text{rot}(yxz) = -1$. Hence, we obtain $\ell_{C'}(z'u) = \ell_{\tilde{C}'}(xu) = \ell_{\tilde{C}'}(yx) + \text{rot}(yxu) = -2$, which contradicts that C' is a decreasing cycle. \square

Altogether, we can efficiently test which of the candidates e_1, \dots, e_k produce decreasing cycles as follows. By Lemma 9.2, if the first candidate is not valid, then $\Gamma_{e_1}^u$ has a decreasing cycle that contains the new edge uz with label 0, which is hence the minimum label for all edges on the cycle. This can be tested in $\mathcal{O}(n)$ time by

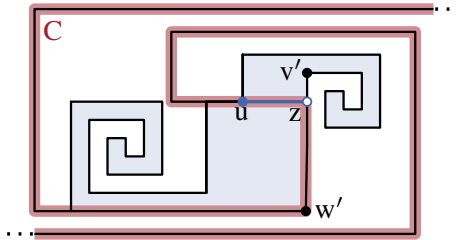


Fig. 44 Here, the insertion of the edge uz to the last candidate $v'w'$ introduces an increasing cycle C with $\ell_C(uz) = -4$

Lemma 8.3. Fact 2. of Proposition 7.2 guarantees that we either find a valid augmentation or a decreasing cycle. In the former case we are done, in the second case Lemma 9.3 allows us to similarly restrict the labels of uz to 0 for the remaining candidate edges, thus allowing us to detect decreasing cycles in $\Gamma_{e_i}^u$ in $\mathcal{O}(n)$ time for $i = 2, \dots, k$.

It is tempting to use the mirror symmetry (Lemma 4.4) to exchange increasing and decreasing cycles to deal with increasing cycles in an analogous fashion. However, this fails as mirroring invalidates the property that u is followed by two right turns in clockwise direction. For example, in Fig. 44 inserting the edge to the last candidate introduces an increasing cycle C with $\ell_C(uz) = -4$. We therefore give a direct algorithm for detecting increasing cycles in this case.

Let $e_i = v_i w_i$ and $e_{i+1} = v_{i+1} w_{i+1}$ be two consecutive candidates for u such that $\Gamma_{e_i}^u$ contains a decreasing cycle but $\Gamma_{e_{i+1}}^u$ does not. If $\Gamma_{e_{i+1}}^u$ contains an increasing cycle, then by Fact 3. of Proposition 7.2 the vertices w_i, v_{i+1} and u lie on a horizontal path that starts at a vertex z incident to f and ends at u . The presence of such a horizontal path P can clearly be checked in linear time, thus allowing us to also detect increasing cycles provided that the previous candidate produced a decreasing cycle. If P exists, we insert the edge uz . By Proposition 5.4 this does not produce strictly monotone cycles. Otherwise, if P does not exist, the augmentation $\Gamma_{e_{i+1}}^u$ is valid. In both cases we have resolved the horizontal port u successfully.

Summarizing, the overall algorithm for augmenting from a horizontal port u now works as follows. By exploiting Lemmas 9.2 and 9.3, we test the candidates in the order as they appear on f until we find the first candidate e for which Γ_e^u does not contain a decreasing cycle. Using Fact 3. of Proposition 7.2 we either find that Γ_e^u is valid, or we find a horizontal path as described above. In both cases this allows us to determine an edge whose insertion does not introduce a strictly monotone cycle. Since in each test for a decreasing cycle the edge uz can be restricted to have label 0, each of the tests takes linear time. This improves the running time of the rectangulation algorithm to $\mathcal{O}(n^3)$.

9.2 2nd Improvement—Fewer Validity Tests

Instead of linearly searching for a suitable candidate for u , we can employ a binary search on the candidates, which reduces the number of validity tests for u from linear to

logarithmic. To do this efficiently, we first compute the list of all candidates e_1, \dots, e_k for u in time linear in the size of f . Next, we test if the augmentation $\Gamma_{e_1}^u$ is valid. If it is, we are done.

Otherwise, we start the binary search on the list e_1, \dots, e_k , where k is the number of candidates for u . The search maintains a sublist e_i, \dots, e_j of consecutive candidates such that $\Gamma_{e_i}^u$ contains a decreasing cycle and $\Gamma_{e_j}^u$ does not. Note that this invariant holds in the beginning, because we explicitly test for a decreasing cycle in $\Gamma_{e_1}^u$ and there is no decreasing cycle in $\Gamma_{e_k}^u$ by Fact 2. of Proposition 7.2. If the list consists of only two consecutive candidates, i.e., $j = i + 1$, we stop. Otherwise, we set $m = \lfloor (i + j)/2 \rfloor$ and test if $\Gamma_{e_m}^u$ contains a decreasing cycle. If it does, we recurse on e_m, \dots, e_j and otherwise on e_i, \dots, e_m . As the invariant is preserved, we end up with two consecutive candidates e_i and e_{i+1} such that $\Gamma_{e_i}^u$ contains a decreasing cycle and $\Gamma_{e_{i+1}}^u$ does not. In this situation Fact 3. of Proposition 7.2 guarantees that we find a valid augmentation. Clearly this only requires $\mathcal{O}(\log n)$ validity tests in total. Further, as argued in the previous subsection each of these tests can be performed in linear time. Altogether, we obtain the following lemma.

Lemma 9.4 *Using binary search we find a valid augmentation for u in $\mathcal{O}(n \log n)$ time.*

Since there are $\mathcal{O}(n)$ ports to remove, we obtain that any planar 4-graph with valid ortho-radial representation can be rectangulated in $\mathcal{O}(n^2 \log n)$ time.

Theorem 9.5 *Given a valid ortho-radial representation Γ , a corresponding rectangulation can be computed in $\mathcal{O}(n^2 \log n)$ time.*

9.3 3rd Improvement—Linear Number of Validity Tests

In this section we describe an improvement of our algorithm that reduces the total number of validity tests to $\mathcal{O}(n)$ such that the running time of our algorithm becomes $\mathcal{O}(n^2)$. The improvement adapts the augmentation step for horizontal ports of the rectangulation procedure. Let u be a horizontal port of a face f and let e_1, \dots, e_k be its candidate edges. The adapted augmentation step resolves u in two steps.

Step 1. We do a linear scan on e_1, \dots, e_k to search for the first candidate e_m of u that gives rise to a valid augmentation with an additional edge uz . Recall that z either subdivides e_m or uz completes a horizontal cycle that contains the source of e_m . We note that we apply m validity tests in this step. In case that $m < 4$, the augmentation step is stopped. Otherwise we continue with the following step.

Step 2. In the following let Γ_0 be the valid ortho-radial representation that we obtain after Step 1. We observe that the inserted edge uz splits the face f into two smaller faces. In Step 2.1 we partition the face that contains the edges e_1, \dots, e_{m-1} by inserting two paths T and B of constant size that start at subdivision vertices on uz and end at subdivision vertices on e_1 and e_{m-1} , respectively; see Fig. 45. These two paths separate two sub-faces f_1 and f_2 of f that contain the candidate edges e_1, \dots, e_{m-2} , which are all but a constant number of edges for which we have performed validity tests in Step 1. In Step 2.2. we rectangulate f_1 without performing any validity test and in Step 2.3 we rectangulate f_2 performing a number of validity tests that is proportional to the size of f_2 . In the following we describe the three steps in greater detail.

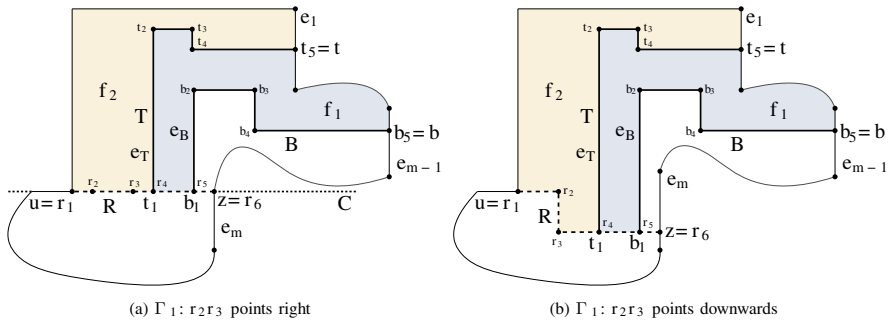


Fig. 45 Illustration of Step 1, which inserts R, T and B into Γ_0 . Depending on whether uz lies on a horizontal cycle C in Γ_0 , the edge r_2r_3 points **a** to the right or **b** downwards

Step 2.1. We adapt Γ_0 as follows. We introduce a path $B = b_1 \dots b_5$ in f that connects a subdivision vertex b_1 on uz with a subdivision vertex on e_{m-1} ; see Fig. 45. The edge b_1b_2 points upwards, the edge b_3b_4 points downwards, and the other two edges point to the right. The face that lies to the left of B is the face f' we seek to rectangulate. Similarly, we introduce a path $T = t_1 \dots t_5$ in f' that connects a subdivision vertex t_1 on ub_1 with a subdivision vertex on e_1 . The edge t_1t_2 points upwards, the edge t_3t_4 points downwards, and the other two edges point to the right. Finally, we subdivide the edge ut_1 by two additional vertices. Altogether, the edge uz has been replaced by a path $R = r_1 \dots r_6$ with $r_1 = u, r_4 = t_1, r_5 = b_1$ and $r_6 = z$. As we have obtained the edges of R by subdividing uz , they all point to the right. In the case that uz does not lie on a horizontal cycle, we orient r_2r_3 such that it points downwards.

We denote the resulting ortho-radial representation by Γ_1 . Further, let f_1 be the face that lies to the right of T , and let f_2 be the face that lies to the left of T .

Step 2.2. We iteratively resolve the ports in f_1 until the face is rectangulated. For each port u' of f_1 we augment the ortho-radial representation as follows. If u' is a vertical port, we augment with respect to the first candidate edge and if u' is a horizontal port we augment with respect to the last candidate edge of u' . The procedure stops when f_1 is completely rectangulated. We denote the resulting ortho-radial representation by Γ_2 . We emphasize that this step does not execute any validity test.

Step 2.3. Starting with Γ_2 , we rectangulate the face f_2 , which has a constant number of ports, by iteratively applying the original augmentation step. We denote the resulting ortho-radial representation by Γ_3 .

In the following we show that the modified rectangulation procedure runs in $\mathcal{O}(n^2)$ time and yields a valid ortho-radial representation.

Lemma 9.6 *The modified rectangulation procedure produces a valid, rectangulated ortho-radial representation.*

Since the proof is rather technical, we defer it to Sect. 9.6. In the following we argue the running time. To that end, we first prove that the output rectangulation has linear size.

Lemma 9.7 *The rectangulated ortho-radial representation produced by the modified rectangulation procedure has size $\mathcal{O}(n)$.*

Proof For the proof we define the potential function

$$\Phi = 3 \cdot \text{horizontal corners of } \Gamma + \text{vertical corners of } \Gamma,$$

where a *horizontal* (*vertical*) corner is a concave corner that becomes a horizontal (vertical) port of a face during the rectangulation procedure. At the beginning of the rectangulation procedure it holds $\Phi \leq 4n$, because each vertex with degree greater than 1 can be either a horizontal or a vertical corner, but not both. Further, each vertex with degree 1 is both a horizontal and vertical corner. We show that for each augmentation step of the rectangulation procedure the potential Φ decreases by some value $\Delta\Phi \geq 1$ and that the number ΔV of inserted vertices is proportional to $\Delta\Phi$. Since the rectangulation procedure terminates with $\Phi = 0$, Lemma 9.7 follows.

In case that the augmentation step handles a vertical port, this vertical corner is resolved but no new corner is created. Hence, Φ decreases by $\Delta\Phi = 1$. Moreover, $\Delta V = 1$. For the case that the augmentation step handles a horizontal port, we distinguish two sub-cases. If $m < 4$, this horizontal corner is resolved but no new corner is created. Hence, Φ decreases by $\Delta\Phi = 3$. Moreover, $\Delta V \leq 1$ (in case that the augmentation step closes a horizontal cycle we have $\Delta V = 0$). Now consider the case that $m \geq 4$. Let k_1 and k_2 be the number of the horizontal and vertical corners, respectively, that are resolved during the rectangulation of f_1 and f_2 in Step 2.2 and Step 2.3, excluding those that lie on R , T and B .

Due to the insertion of the vertices r_3 and b_4 the potential increases by at most 2. Further, by resolving the horizontal corner u the potential decreases by 3. Further, rectangulating f_1 and f_2 decreases the potential by $3k_1 + k_2$. Altogether, we obtain $\Delta\Phi \geq (3 - 2) + 3k_1 + k_2 \geq 1 + k_1 + k_2$. Moreover, in Step 2.1 we add at most 13 vertices. In Step 2.2 we add $k_1 + k_2$ vertices for the corners not on T and B , and we add three vertices for t_4 , b_2 , and b_3 . In Step 2.3 we add at most three vertices for t_2 , t_3 , and possibly r_2 . In total, we add $\Delta V \leq 19 + k_1 + k_2$ vertices.

In all cases $\Delta V \leq 19\Delta\Phi$. Altogether, we obtain that the rectangulation procedure terminates (since Φ decreases in every rectangulation step) and the resulting rectangulation has $\mathcal{O}(n)$ vertices, and therefore $\mathcal{O}(n)$ edges. \square

Lemma 9.8 *The modified rectangulation procedure applies $\mathcal{O}(n)$ validity tests and runs in $\mathcal{O}(n^2)$ time.*

Proof We now show that by replacing the original augmentation step with this adaptation, the number of validity tests is linear and the rectangulation procedure runs in $\mathcal{O}(n^2)$ time. We use a charging argument that assigns to each vertex and to each edge of the output rectangulation a constant number of the validity tests that have been applied during the rectangulation procedure. Further, we distribute the total running time such that running time linear in the output size is assigned to each vertex and each edge. Hence, by Lemma 9.7 the rectangulation procedure applies $\mathcal{O}(n)$ validity tests and runs in $\mathcal{O}(n^2)$ time.

Let u be a port that is considered in a rectangulation step. If u is a vertical port, we determine its first candidate edge in $\mathcal{O}(n)$ time, which we charge on u . Further, we do not apply a validity test.

If u is a horizontal port, we determine its candidate edges in $\mathcal{O}(n)$ time, which we charge on u . Further, we apply m validity tests in Step 1. If $m < 4$, only a constant number of validity tests is applied, which we charge on u . Otherwise, we charge the validity tests of e_1, e_{m-1} and e_m on u . We observe that after the augmentation step, the horizontal port u is resolved and no further validity tests can be charged on u . Further, we charge the validity test of e_i with $2 \leq i \leq m - 2$ on e_i . Hence, since by construction the candidates e_2, \dots, e_{m-2} belong to rectangles after the augmentation step, each edge can only be charged twice (once from each side).

Step 2.1 has constant running time and applies no validity tests. We charge the running time on u . Step 2.2 requires no validity tests and resolves each concave corner v on f_1 in $\mathcal{O}(n)$ time, which we charge on v . Afterwards, the face f_1 is rectangulated and v cannot be charged again. In Step 2.3 we need $\mathcal{O}(n)$ time for each concave corner v to identify the candidate edges, which we charge on v . Further, when applying a validity test in this step, we charge it on the corresponding candidate edge. As f_2 has only a constant number of concave corners, each edge of f_2 is charged with at most a constant number of validity tests. Further, since f_2 is rectangulated afterwards its edges cannot be charged again (from the side of f_2).

Hence, over all rectangulation steps we obtain $\mathcal{O}(n)$ validity tests and running time $\mathcal{O}(n^2)$. \square

Altogether, we obtain that any planar 4-graph with valid ortho-radial representation can be rectangulated in $\mathcal{O}(n^2)$ time. Using Corollary 6.3 this further implies that a corresponding bend-free ortho-radial drawing can be computed in $\mathcal{O}(n^2)$ time.

Theorem 9.9 *Given a valid ortho-radial representation Γ , a corresponding rectangulation can be computed in $\mathcal{O}(n^2)$ time.*

9.3.1 Proof of Lemma 9.6

We prove Lemma 9.6 by showing that each rectangulation step yields a valid ortho-radial representation. By Proposition 7.2, Step 1 yields a valid ortho-radial representation Γ_0 . If Step 2 is not considered, Γ_0 is the output of the rectangulation step. So assume that Step 2 is executed. We use the same notation as in the description of the algorithm. In particular, Step 2.1, Step 2.2 and Step 2.3 produce the ortho-radial representations Γ_1, Γ_2 and Γ_3 from the ortho-radial representation of the preceding step. In the following we consider the sub-steps of Step 2 separately and show that Γ_1, Γ_2 and Γ_3 are valid.

Correctness of Step 2.1. In order to show the correctness of the first step we successively add the paths R, T and B to Γ_0 and prove the validity of each created ortho-radial representation. To that end, let $\Gamma_R = \Gamma_0 - uz + R$, $\Gamma_T = \Gamma_R + T$ and $\Gamma_B = \Gamma_T + B = \Gamma_1$.

Lemma 9.10 *The ortho-radial representation Γ_R is valid.*

Proof Let C be an essential cycle in Γ_R . If C does not contain R , it is contained in Γ_0 . Since Γ_0 is valid, C is not strictly monotone. It remains to consider the case that C uses R .

If the edge r_2r_3 of R points to the right, R lies on a horizontal cycle by construction. By Proposition 5.4 C is not strictly monotone. If r_2r_3 points downwards, the cycle C corresponds to an essential cycle $C' = C[z, u] + uz$ in Γ_0 . By Lemma 5.8 the labels of C and C' coincide on $C[z, u]$. By construction C' is not horizontal. Hence, it contains edges e and e' with $\ell_{C'}(e) = 1$ and $\ell_{C'}(e') = -1$. Since uz is a horizontal edge, we have $e \neq uz \neq e'$. Therefore, $\ell_C(e) = 1$ and $\ell_C(e') = -1$, that is, C is not strictly monotone. \square

Next, we prove that Γ_T and Γ_B are valid. To that end, we introduce the following definition. A *cascading cycle* is a non-monotone essential cycle that can be partitioned into two paths P and Q such that the labels on P are -1 and the labels on Q are non-negative. We further require that the edges incident to the internal vertices of P either all lie in the interior of C or they all lie in the exterior of C . In the first case we call C an *outer* cascading cycle and in the second case an *inner* cascading cycle. The path P is the *negative path* of the cycle. We first prove the following two general lemmata on cascading cycles. The first lemma shows that a cascading cycle cannot be *crossed* by an increasing cycle.

Lemma 9.11 *Let C_1 be a cascading cycle and C_2 an increasing cycle. Either C_1 lies in the interior of C_2 or vice versa.*

Proof We assume without loss of generality that C_1 is an outer cascading cycle. The case that it is an inner cascading cycle can be handled by flipping the cylinder, which exchanges the exterior and interior of essential cycles but keeps the labels.

Let g be the central face of the subgraph formed by the cycles C_1 and C_2 . Assume that g is neither C_1 nor C_2 . Hence, there exist (not necessarily distinct) vertices v and v' on g such that $g[v, v']$ belongs to both C_1 and C_2 , the edge that precedes v on g does not belong to C_2 , and the edge that succeeds v' on g does not belong to C_1 ; see Fig. 46.

For $i = 1, 2$ let u_i and w_i be the vertices of C_i that precede and succeed v , respectively. Further, let u'_i and w'_i be the vertices of C_i that precede and succeed v' , respectively. Since u_2v strictly lies in the exterior of g , we obtain $\text{rot}(u_2vw_1) < \text{rot}(u_1vw_1)$. Hence, Lemma 5.2 gives $\ell_{C_1}(u_1v) < \ell_{C_2}(u_2v)$. Further, by definition of C_1 and C_2 it holds $\ell_{C_1}(u_1v) \geq -1$ and $\ell_{C_2}(u_2v) \leq 0$. Thus, we obtain $\ell_{C_1}(u_1v) = -1$ and $\ell_{C_2}(u_2v) = 0$. Since C_1 is an outer cascading cycle this also implies that v is the endpoint of the negative path P of C_1 .

Since $v'w'_1$ strictly lies in the exterior of g , we obtain $\text{rot}(u'_2v'w'_2) > \text{rot}(u'_2v'w'_1)$. Further, by the definition of labels we obtain $\ell_{C_2}(u'_2v') = \ell_{C_2}(v'w'_2) - \text{rot}(u'_2v'w'_2)$. Applying this in Lemma 5.2 gives $\ell_{C_1}(v'w'_1) = \ell_{C_2}(v'w'_2) - \text{rot}(u'_2v'w'_2) + \text{rot}(u'_2v'w'_1)$. Thus, we obtain $\ell_{C_1}(v'w'_1) < \ell_{C_2}(v'w'_2)$. Analogously to the arguments above, $\ell_{C_2}(v'w'_2) = 0$ and $\ell_{C_1}(v'w'_1) = -1$. Thus, v' lies on P and it is not the endpoint of P . It follows $v \neq v'$ and $C_1 - P$ is contained in $g[v, v']$. Hence, there is an edge e on $g[v, v']$ with label $\ell_{C_1}(e) > 0$. By Corollary 5.3 $\ell_{C_2}(e) = \ell_{C_1}(e)$, which contradicts that C_2 is an increasing cycle. \square

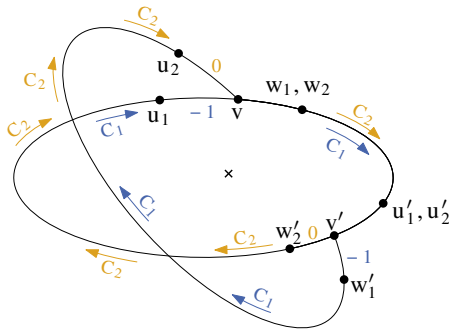


Fig. 46 Illustration of proof for Lemma 9.11. It is assumed that the central face g is neither the outer cascading cycle C_1 nor the increasing cycle C_1 . Further, C_1 and C_2 have the edges $g[v, v']$ in common. It is proven that any edge of $C_1[v, v']$ has label 0 and any edge of $C_1[v', v]$ has label -1 , which contradicts that C_1 is a cascading cycle

We use the next lemma to show that the introduced paths T and B do not impact the validity of the ortho-radial representation.

Lemma 9.12 *Let C be a cascading cycle in an ortho-radial representation Γ and let P be a sub-path of C such that*

- (i) *the intermediate vertices of P have degree 2 in the graph,*
- (ii) *P contains the negative path of C , and*
- (iii) *P contains an edge e with label $\ell_C(e) > 0$.*

If Γ without P is valid, then Γ is valid.

Proof First assume that Γ contains a decreasing cycle C' , which implies that P is contained in C' (in either direction). Let $H = C + C'$ be the common sub-graph of C and C' , and let g be the central face of H . We distinguish the following two cases.

Case 1, P is part of g . First assume that C and C' use P in opposite directions. Since the central face locally lies to the right of any essential cycle, this implies that the central face lies to the left and right of P . Consequently, the central face is not simple, which contradicts that H is biconnected. So assume that C and C' use P in the same direction. By Proposition 5.2 it holds $\ell_C(e') = \ell_{C'}(e') = -1$ for any edge e' on the negative path of C . Thus, C' is not a decreasing cycle.

Case 2, P is not part of g . Let C_g be the essential cycle formed by g . Since C_g consists of edges of C' and C , the corresponding labels of C' and C also apply on C_g by Proposition 5.2. Further, since C' is a decreasing cycle and P is the only part of C that has a negative label on C , the cycle C_g only has non-negative labels. Since P does not lie on C_g but on C , C' has at least one vertex with C_g in common. By Proposition 5.4 the cycle C_g is not horizontal. Altogether, C_g is a decreasing cycle that also exists in $\Gamma - P$, which contradicts its validity.

Finally, assume that Γ contains an increasing cycle C' , which implies that P is contained in C' in either direction. Lemma 9.11 implies that the central face g of the subgraph formed by the two essential cycles C and C' is either C or C' . In particular, P or \bar{P} lies on g . Hence, both C and C' use P in the same direction as otherwise g

would lie in the exterior of one of these cycles. But this would contradict that they are essential. Hence, they both contain P in this direction and P also lies on g . By Proposition 5.2 both cycles have the same labels on P . Since P contains the edge e with $\ell_C(e) > 0$ and thus $\ell_{C'}(e) > 0$, the cycle C' is not an increasing cycle. \square

We construct a cascading cycle C_T in Γ_T as follows. Let C be the outermost decreasing cycle in $\Gamma_{e_1}^u$ and let ut_5 be the newly inserted edge in $\Gamma_{e_1}^u$. We replace ut_5 by $R[u, r_4] + T$ obtaining the cycle C_T , which is well-defined because C uses ut_5 in that direction by Lemma 9.2.

Lemma 9.13 *C_T is a cascading cycle with negative path t_1t_2 no matter whether r_2r_3 points to the right or downwards.*

Proof Let C be the outermost decreasing cycle in $\Gamma_{e_1}^u$. By Lemma 5.8 the labels of C and C_T coincide on $C[t_5, u]$ ($= C_T[t_5, u]$). Hence, since C is a decreasing cycle all labels on $C_T[t_5, u]$ are non-negative. Further, by Lemma 9.2 the edge ut_5 has label 0 on C . If r_2r_3 points to the right the sequence of the labels on $R[u, r_4] + T$ is therefore 0, 0, 0, -1, 0, 1, 0. If r_2r_3 points downwards the sequence is 0, 1, 0, -1, 0, 1, 0. In both cases C_T is a cascading cycle. \square

Applying Lemma 9.11 to the situation of C_T proves that Γ_T does not contain any increasing cycles. Together with Lemma 9.12 this yields that Γ_T is valid. We analogously prove the validity of Γ_B as for Γ_T . Let C be the outermost decreasing cycle in $\Gamma_{e_{m-1}}^u$ and let ub_5 be the newly inserted edge in $\Gamma_{e_{m-1}}^u$. We replace ub_5 by $R[u, r_5] + B$ obtaining the cycle C_B , which is well-defined because C uses ub_5 in that direction by Lemma 9.2.

Lemma 9.14 *C_B is a cascading cycles no matter whether r_2r_3 points to the right or downwards. In particular, b_1b_2 is the negative path of C_B .*

Proof Let C be the outermost decreasing cycle in $\Gamma_{e_{m-1}}^u$. By Lemma 5.8 the labels of C and C_B coincide on $C[b_5, u]$ ($= C_B[b_5, u]$). Hence, since C is a decreasing cycle all labels on $C_B[b_5, u]$ are non-negative. Further, by Lemma 9.3 the edge ub_5 has label 0 on C . If r_2r_3 points to the right the sequence of the labels on $R[u, r_4] + B$ is therefore 0, 0, 0, 0, -1, 0, 1, 0. If r_2r_3 points downwards the sequence is 0, 1, 0, 0, -1, 0, 1, 0. In both cases C_B is a cascading cycle. \square

Applying Lemma 9.11 to the situation of C_T proves that Γ_B does not contain any increasing cycles. Together with Lemma 9.12 this yields that Γ_B is valid. The following lemma summarizes the result.

Lemma 9.15 *The ortho-radial representation $\Gamma_B = \Gamma_1$ is valid.*

Correctness of Step 2.2. By Lemma 9.15 the ortho-radial representation Γ_1 of Step 1 is valid. We now prove that Γ_2 is a valid ortho-radial representation. We use the same notation as in the description of the algorithm.

Starting with the valid ortho-radial representation Γ_1 , the procedure iteratively resolves ports in the face f_1 , which locally lies to the right of T . In case that we resolve

a vertical port u' in a representation Π , the resulting ortho-radial representation $\Pi_{e'_1}^{u'}$ is valid by Fact 1. of Proposition 7.2, where e'_1 is the first candidate of u' . So assume that u' is a horizontal port. In that case we take $\Pi_{e'_1}^{u'}$ for the next iteration, where e'_1 is the last candidate of u' . We observe that the augmentation of f_1 may subdivide edges on the negative paths of C_T and C_B , but the added edges lie in the interior of C_T and the exterior of C_B . Hence, C_T remains an outer cascading cycle and C_B an inner cascading cycle.

Lemma 9.16 *The ortho-radial representation $\Pi_{e'_1}^{u'}$ is valid.*

Proof Assume that $\Pi_{e'_1}^{u'}$ is not valid. Hence, there is a strictly monotone cycle C that uses $e = u'z'$, where z' is the vertex subdividing e'_1 . Since e'_1 is the last candidate of u' , the cycle C is increasing by Fact 2. of Proposition 7.2. By construction e strictly lies in the interior of C_T and the exterior of C_B . This implies that C lies in the interior of C_T and the exterior of C_B by Lemma 9.11. In other words, C is contained in the subgraph H formed by the intersection of the interior of C_T and the exterior of C_B . As $R[r_1, r_4]$ belongs to both C_T and C_B , it is incident to the outer and the central face of H . Hence, removing $R[r_1, r_4]$ leaves a subgraph without essential cycles. Thus, the essential cycle C includes $R[r_1, r_4]$.

By Proposition 5.2 the labels of C and C_T are the same on $R[r_1, r_4]$. If r_2r_3 points downwards, its label is 1, which contradicts that C is increasing. If otherwise r_2r_3 points right, it lies on a horizontal cycle. But then C is not increasing by Proposition 5.4. \square

Altogether, applying the lemma inductively on the inserted edges, we obtain that Γ_2 is valid.

Correctness of Step 2.3. As we only apply the first phase of the augmentation step on f_2 , the resulting ortho-radial representation Γ_3 is also valid due to the correctness of the first phase. This concludes the correctness proof of the second phase.

10 Bend Minimization

We have considered bend-free ortho-radial drawings so far. In this section we shortly describe how to extend our results to ortho-radial drawings with bends. We model bends by subdividing each edge of G with K degree-2 vertices. Let S denote the set of all these subdivision vertices. A vertex $v \in S$ is a *bend* in an ortho-radial representation if for its two incident edges e_1 and e_2 it holds $\text{rot}(e_1, e_2) \neq 0$. A valid ortho-radial representation of I is *bend-minimal* if there is no other valid ortho-radial representation of I that has fewer bends. Niedermann and Rutter [27] showed that $K = 2n - 4$ subdivision vertices per edge are sufficient for creating bend-minimal valid ortho-radial representations. They propose an integer linear programming formulation for creating valid bend-minimal ortho-radial representations assuming a pre-defined embedding of the graph.

If we do not assume a pre-defined embedding of the graph, finding a bend-minimal drawing becomes \mathcal{NP} -hard even for the orthogonal case. Garg and Tamassia [19]

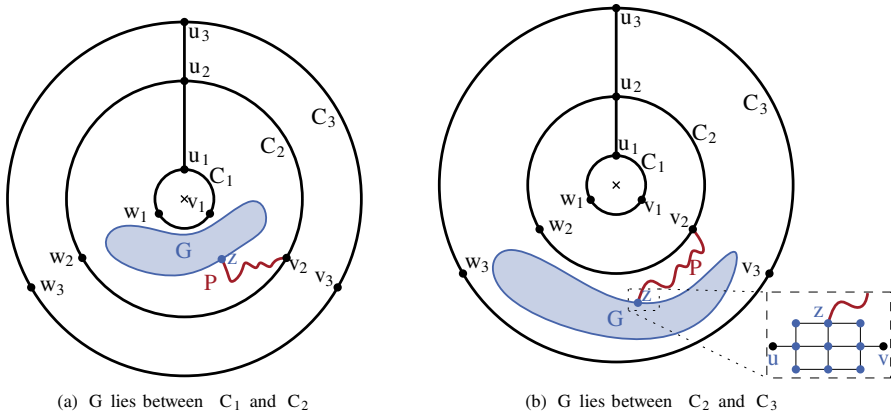


Fig. 47 Possible embeddings of G' : In both cases G contains no essential cycles. The roles of C_1 and C_3 can be exchanged

showed that the problem **ORTHOGONAL EMBEDDABILITY** to decide whether a given planar 4-graph admits an orthogonal drawing without bends is \mathcal{NP} -complete. In the remaining part of this section we study the analogous problem **ORTHO-RADIAL EMBEDDABILITY** for ortho-radial drawings and prove that it is \mathcal{NP} -complete as well. We say a graph G admits an *ortho-radial* (or *orthogonal*) *embedding* if there is an embedding of G such that G can be drawn ortho-radially (or orthogonally) without bends.

We give a reduction from **ORTHOGONAL EMBEDDABILITY**. To do so, we note that the reduction by Garg and Tamassia [19] actually produces instances $G = (V, E)$ with a fixed edge $e \in E$ such that it is \mathcal{NP} -complete to decide whether G has an orthogonal embedding where e is incident to the outer face.

Given such a graph, we build a structure around G that yields a graph G' such that in any ortho-radial embedding of G' the induced representation Γ of G does not contain any essential cycles. In other words, Γ is actually an orthogonal representation of G . Hence, an ortho-radial embedding of G' can only exist if G admits an orthogonal embedding. We may assume without loss of generality that G is connected as otherwise, we handle each component separately.

The construction of G' from G is based on the fact that there is only one way to ortho-radially draw a triangle C , i.e., a cycle of length 3, without bends: as an essential cycle on one circle of the grid. We build a graph H consisting of three triangles called C_1, C_2 and C_3 and denote the vertices on C_i by u_i, v_i and w_i . Furthermore, H contains the edges u_1u_2 and u_2u_3 . In Fig. 47 H is formed by the black edges. To connect H and G , we replace the special edge $e = uv$ of G by a 3×3 -grid and connect one of the degree-3 vertices z of that grid by a path P to v_2 , where we choose the length of the path equal to the number of edges in G . The resulting graph is G' ; see Fig. 47. The reduction clearly runs in polynomial time. Moreover, if G' admits an ortho-radial embedding, then the triangles of H must be drawn as essential cycles, and therefore G must be contained in one of the two regular faces of H , and can hence not contain any essential cycles. We therefore find an orthogonal embedding of G . Conversely, if we

have an orthogonal embedding of G with e on the outer face, then it can be inserted into a face of the drawing of H and the path P can be drawn as it is sufficiently long. We therefore obtain an ortho-radial embedding of G' . This proves the following theorem.

Theorem 10.1 ORTHO-RADIAL EMBEDDABILITY is \mathcal{NP} -complete.

11 Conclusion

In this paper we considered orthogonal drawings of graphs on cylinders. Our main result is a characterization of the plane 4-graphs that can be drawn bend-free on a cylinder in terms of a combinatorial description of such drawings. These ortho-radial representations determine all angles in the drawing without fixing any lengths, and thus are a natural extension of Tamassia's orthogonal representations. However, compared to those, the proof that every valid ortho-radial representation has a corresponding drawing is significantly more involved. The reason for this is the more global nature of the additional property required to deal with the cyclic dimension of the cylinder.

Our ortho-radial representations establish the existence of an ortho-radial TSM framework in the sense that they are a combinatorial description of the graph serving as interface between the “Shape” and “Metrics” step.

For rectangular plane 4-graphs, we gave an algorithm producing a drawing from a valid ortho-radial representation. Our proof reducing the drawing of general plane 4-graphs with a valid ortho-radial representation to the case of rectangular plane 4-graphs is constructive. We have described an algorithm that checks the validity of an ortho-radial representation in $\mathcal{O}(n^2)$ time. In the positive case, we can also produce a corresponding drawing in the same running time, whereas in the negative case we find a strictly monotone cycle. These algorithms correspond to the “Metrics” step in a TSM framework for ortho-radial drawings. It is an interesting question whether the running time can be improved further.

While bend-minimal orthogonal representations can be created in polynomial time, it remains an open question whether the same can be achieved for bend-minimal valid ortho-radial representations.

Finally, we want to emphasize that we deem the generalization of ortho-radial drawings from the cylinder to the torus or even more complex surfaces an interesting and promising research question. It is far from clear how to transfer our results to the torus as its two cyclic dimensions lead to different types of essential cycles.

Acknowledgements Lukas Barth was funded by the German Research Foundation (DFG) as part of the Research Training Group GRK 2153: Energy Status Data—Informatics Methods for its Collection, Analysis and Exploitation. Matthias Wolf was funded by the Helmholtz Association Program Storage and Cross-linked Infrastructures, Topic 6 Superconductivity, Networks and System Integration and by the German Research Foundation (DFG) as part of the Research Training Group GRK 2153: Energy Status Data—Informatics Methods for its Collection, Analysis and Exploitation.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence,

and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Alam, M.J., Kobourov, S.G., Mondal, D.: Orthogonal layout with optimal face complexity. *Comput. Geom.* **63**, 40–52 (2017)
2. Bast, H., Brosi, P., Storandt, S.: Metro maps on flexible base grids. In: 17th International Symposium on Spatial and Temporal Databases (SSTD'21), pp. 12–22. ACM, New York (2021)
3. Batini, C., Nardelli, E., Tamassia, R.: A layout algorithm for data flow diagrams. *IEEE Trans. Softw. Eng.* **SE-12**(4), 538–546 (1986)
4. Bertolazzi, P., di Battista, G., Didimo, W.: Computing orthogonal drawings with the minimum number of bends. *IEEE Trans. Comput.* **49**(8), 826–840 (2000)
5. Bhatt, S.N., Leighton, F.T.: A framework for solving VLSI graph layout problems. *J. Comput. Syst. Sci.* **28**(2), 300–343 (1984)
6. Biedl, T.C.: New lower bounds for orthogonal graph drawings. In: Brandenburg, F.J. (ed.) *Graph Drawing (GD'95)*. Lecture Notes in Computer Science, vol. 1027, pp. 28–39. Springer, Berlin (1996)
7. Biedl, T., Kant, G.: A better heuristic for orthogonal graph drawings. *Comput. Geom.* **9**(3), 159–180 (1998)
8. Biedl, T.C., Madden, B.P., Tollis, I.G.: The three-phase method: a unified approach to orthogonal graph drawing. In: DiBattista, G. (ed.) *Graph Drawing (GD'97)*. Lecture Notes in Computer Science, vol. 1353, pp. 391–402. Springer, Berlin (1997)
9. Bläsius, T., Rutter, I., Wagner, D.: Optimal orthogonal graph drawing with convex bend costs. *ACM Trans. Algorithms* **12**(3), Art. No. 33 (2016)
10. Bläsius, T., Lehmann, S., Rutter, I.: Orthogonal graph drawing with inflexible edges. *Comput. Geom.* **55**, 26–40 (2016)
11. Chang, Y.-J., Yen, H.-C.: On bend-minimized orthogonal drawings of planar 3-graphs. In: Aronov, B., Katz, M.J. (eds.) *Computational Geometry (SoCG'17)*. Leibniz International Proceedings in Informatics (LIPIcs), vol. 77, Art. No. 29. Schloss Dagstuhl—Leibniz-Zentrum für Informatik, Wadern (2017)
12. Cornelsen, S., Karrenbauer, A.: Accelerated bend minimization. In: van Kreveld, M., Speckmann, B. (eds.) *Graph Drawing (GD'12)*. Lecture Notes of Computer Science, vol. 7034, pp. 111–122. Springer, Heidelberg (2012)
13. Eiglsperger, M., Kaufmann, M., Siebenhaller, M.: A topology-shape-metrics approach for the automatic layout of UML class diagrams. In: *Proceedings of the 2003 ACM Symposium on Software Visualization (SoftVis'03)*, pp. 189–ff. ACM, New York (2003)
14. Eiglsperger, M., Gutwenger, C., Kaufmann, M., Kupke, J., Jünger, M., Leipert, S., Klein, K., Mutzel, P., Siebenhaller, M.: Automatic layout of UML class diagrams in orthogonal style. *Inf. Vis.* **3**(3), 189–208 (2004)
15. Felsner, S., Kaufmann, M., Valtr, P.: Bend-optimal orthogonal graph drawing in the general position model. *Comput. Geom.* **47**(3, Part B), 460–468 (2014)
16. Fink, M., Haverkort, H., Nöllenburg, M., Roberts, M., Schuhmann, J., Wolff, A.: Drawing metro maps using Bézier curves. In: Didimo, W., Patrignani, M. (eds.) *Graph Drawing (GD'13)*. Lecture Notes in Computer Science, vol. 7704, pp. 463–474. Springer, Heidelberg (2013)
17. Fink, M., Lechner, M., Wolff, A.: Concentric metro maps. In: Roberts, M.J., Rodgers, P. (eds.) *Abstracts of the Schematic Mapping Workshop 2014* (2014)
18. Fößmeier, U., Kaufmann, M.: Drawing high degree graphs with low bend numbers. In: *International Symposium on Graph Drawing (Passau 1995)*. Lecture Notes in Computer Science, vol. 1027, pp. 254–266. Springer, Berlin (1996)
19. Garg, A., Tamassia, R.: On the computational complexity of upward and rectilinear planarity testing. *SIAM J. Comput.* **31**(2), 601–625 (2001)

20. Gutwenger, C., Jünger, M., Klein, K., Kupke, J., Leipert, S., Mutzel, P.: A new approach for visualizing UML class diagrams. In: 1st ACM Symposium on Software Visualization (San Diego 2003), pp. 179–188. ACM, New York (2003)
21. Hasheminezhad, M., Hashemi, S.M., Tahmabasi, M.: Ortho-radial drawings of graphs. *Australas. J. Combin.* **44**, 171–182 (2009)
22. Hasheminezhad, M., Hashemi, S.M., McKay, B.D., Tahmasbi, M.: Rectangular-radial drawings of cubic plane graphs. *Comput. Geom.* **43**(9), 767–780 (2010)
23. Hong, S.-H., Merrick, D., do Nascimento, H.A.D.: Automatic visualisation of metro maps. *J. Vis. Lang. Comput.* **17**(3), 203–224 (2006)
24. Kieffer, S., Dwyer, T., Marriott, K., Wybrow, M.: Hola: human-like orthogonal network layout. *IEEE Trans. Visual. Comput. Graph.* **22**(1), 349–358 (2016)
25. Miller, G.L., Naor, J.: Flow in planar graphs with multiple sources and sinks. *SIAM J. Comput.* **24**(5), 1002–1017 (1995)
26. Nickel, S., Nöllenburg, M.: Towards data-driven multilinear metro maps. In: *Diagrammatic Representation and Inference. Lecture Notes in Artificial Intelligence*, vol. 12169, pp. 153–161. Springer, Cham (2020)
27. Niedermann, B., Rutter, I.: An integer-linear program for bend-minimization in ortho-radial drawings. In: 28th International Symposium on Graph Drawing and Network Visualization (Vancouver 2020). *Lecture Notes in Computer Science*, vol. 12590, pp. 235–249. Springer, Cham (2020)
28. Nöllenburg, M., Wolff, A.: Drawing and labeling high-quality metro maps by mixed-integer programming. *IEEE Trans. Vis. Comput. Graph.* **17**(5), 626–641 (2011)
29. Papakostas, A., Tollis, I.G.: Algorithms for area-efficient orthogonal drawings. *Comput. Geom.* **9**(1–2), 83–110 (1998)
30. Rüegg, U., Kieffer, S., Dwyer, T., Marriott, K., Wybrow, M.: Stress-minimizing orthogonal layout of data flow diagrams with ports. In: 22nd International Symposium on Graph Drawing (Würzburg 2014). *Lecture Notes in Computer Science*, vol. 8871, pp. 319–330. Springer, Heidelberg (2014)
31. Tamassia, R.: On embedding a graph in the grid with the minimum number of bends. *SIAM J. Comput.* **16**(3), 421–444 (1987)
32. Tamassia, R., Di Battista, G., Batini, C.: Automatic graph drawing and readability of diagrams. *IEEE Trans. Syst. Man Cybern.* **18**(1), 61–79 (1988)
33. Tamassia, R., Tollis, I.G., Vitter, J.S.: Lower bounds for planar orthogonal drawings of graphs. *Inf. Process. Lett.* **39**(1), 35–40 (1991)
34. Tollis, I.G., Di Battista, G., Eades, P., Tamassia, R.: *Graph Drawing: Algorithms for the Visualization of Graphs*. Prentice Hall, Upper Saddle River (1999)
35. Valiant, L.G.: Universality considerations in VLSI circuits. *IEEE Trans. Comput.* **30**(2), 135–140 (1981)
36. Wang, Y.-S., Chi, M.-T.: Focus+context metro maps. *Trans. Vis. Comput. Graph.* **17**(12), 2528–2535 (2011)
37. Wu, H.-Y., Niedermann, B., Takahashi, S., Roberts, M.J., Nöllenburg, M.: A survey on transit map layout—from design, machine, and human perspectives. *Comput. Graph. Forum* **39**(3), 619–646 (2020)
38. Wybrow, M., Marriott, K., Stuckey, P.J.: Orthogonal connector routing. In: 17th International Symposium on Graph Drawing (Chicago 2009). *Lecture Notes in Computer Science*, vol. 5849, pp. 219–231. Springer, Berlin (2010)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.