Turn-Regularity and Optimal Area Drawings of Orthogonal Representations

Stina Bridgeman, Giuseppe Di Battista,
Walter Didimo, Giuseppe Liotta,
Roberto Tamassia, Luca Vismara

Department of Computer Science
Brown University
Providence, Rhode Island 02912

CS-99-04
March 1999

Turn-Regularity and Optimal Area Drawings of Orthogonal Representations*

Stina Bridgeman[†] *Giuseppe Di Battista*[‡] *Walter Didimo*[‡] *Giuseppe Liotta*[§]
Roberto Tamassia[†] *Luca Vismara*[†]

Abstract

Given an orthogonal representation H with n vertices and bends, we study the problem of computing a planar orthogonal drawing of H with small area. This problem has direct applications to the development of practical graph drawing techniques for information visualization and VLSI layout. In this paper, we introduce the concept of turn-regularity of an orthogonal representation H, provide combinatorial characterizations of it, and show that if H is turn-regular (i.e., all its faces are turn-regular), then a planar orthogonal drawing of H with minimum area can be computed in $O(n)$ time, and a planar orthogonal drawing of H with minimum area and minimum total edge length within that area can be computed in $O(n^{7/4} \log n)$ time. We also apply our theoretical results to the design and implementation of new practical heuristic methods for constructing planar orthogonal drawings. An experimental study conducted on a test suite of orthogonal representations of randomly generated biconnected 4-planar graphs shows that, on average, 95% of the faces are regular and that our heuristic drawing methods perform better than previous ones.

*Research supported in part by the National Science Foundation under grants CCR-9732327 and CDA-9703080, and by the U.S. Army Research Office under grant DAAH04-96-1-0013.

[†]Center for Geometric Computing, Department of Computer Science, Brown University, 115 Waterman Street, Providence, RI 02912–1910. {ssb,rt,lv}@cs.brown.edu

[‡]Dipartimento di Informatica e Automazione, Università degli Studi di Roma Tre, Via della Vasca Navale 79, 00146 Roma, Italy. {gdb,didimo}@dia.uniroma3.it

[§]Dipartimento di Ingegneria Elettronica e dell'Informazione, Università degli Studi di Perugia, Via Duranti 93, 06131 Perugia, Italy. liotta@istel.ing.unipg.it

1 Introduction

Orthogonal drawings are drawings of graphs in which every edge is represented by a chain of horizontal and vertical segments. An orthogonal representation is an equivalence class of orthogonal drawings that have the same "shape" (see Fig. 1). This class is formally described by specifying the bends along each edge and the angles between consecutive edges around each vertex. In this paper we consider planar orthogonal representations, that is, equivalence classes of orthogonal drawings for which at least one of the drawings is planar. Given a planar orthogonal representation H, the problem of finding a planar orthogonal grid drawing of H with small area is usually referred to as the *compaction* of H.

Orthogonal representations and planar orthogonal drawings have been extensively investigated (see, e.g., [1, 9, 11, 13, 15, 16, 17, 19, 31, 32, 33, 34]) because of their direct applications to the development of practical graph drawing techniques for information visualization [6]. In particular, it has been experimentally shown that drawing algorithms for general graphs based on the compaction of orthogonal representations with minimum number of bends perform better in practice than other known orthogonal drawing algorithms [7, 29]. Orthogonal representations and related concepts, such as slicing floorplans, are also widely used in VLSI layout compaction algorithms (see, e.g., [22, 25, 26, 30, 35]).

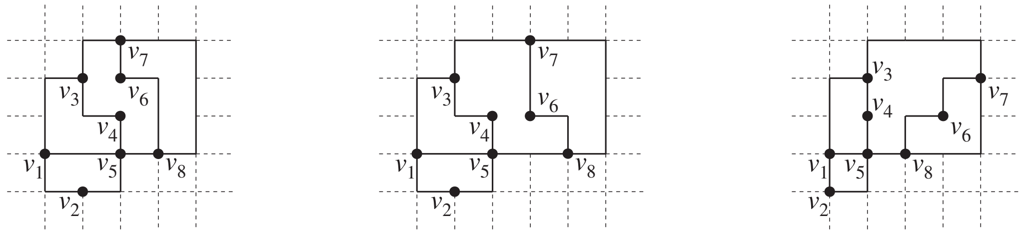

(a) (b) (c)

Figure 1: Three planar orthogonal drawings of a graph. Two drawings of the same orthogonal representation are shown in (a) and (b), while a drawing of a different orthogonal representation is shown in (c). The drawing in (a) has optimal area among all planar drawings of that orthogonal representation.

Despite the significant body of research on orthogonal representations, the development of effective compaction techniques remains a challenging task. It has been conjectured for a long time [34], and recently proved [28], that the optimal compaction of planar orthogonal representations, i.e., computing a minimum area planar orthogonal grid drawing of a given planar orthogonal representation, is NP-complete. The only class of planar orthogonal representations for which a polynomial-time optimal compaction algorithm is known is the trivial class of orthogonal representations whose faces are all rectangular [6].

From a practical perspective, the compaction algorithms used by current graph drawing libraries and systems, such as *AGD*[1], *GDToolkit*[2], and the *Graph Drawing Server*[3], are all variations of the compaction technique presented in [19, 31], which is based on the idea of splitting faces into

[1] `http://www.mpi-sb.mpg.de/AGD/`
[2] `http://www.dia.uniroma3.it/~gdt/`
[3] `http://www.cs.brown.edu/cgc/graphserver/`

rectangles. Since the splitting imposes unnecessary constraints on the geometry, the resulting drawings may have substantially suboptimal area.

The importance of compaction techniques for graph visualization applications is confirmed by a recent work of Klau and Mutzel [23, 24]. They consider the problem of assigning coordinates to vertices and edge bends of an orthogonal representation so that the total edge length is minimized. The problem is formulated as an integer linear program, whose practical performance is fairly good. Also, they show that the problem can be solved in polynomial time for those orthogonal representations in which there is only one possible relative position of any two vertices that results in a planar drawing; in this case, the inequalities of the corresponding ILP formulation form a totally unimodular matrix. The problem of minimizing the area of the drawing is not considered.

The main results of this paper can be summarized as follows.

- Given a planar orthogonal representation H, we define the concept of *turn-regularity* of a face of H, which is based on the structure of the sequence of left and right turns encountered when traversing the face. We show that the turn-regularity of a face can be tested in linear time.
- We relate turn-regularity to the concept of *switch-regularity* [8]. Namely, we characterize the turn-regularity of a face f in terms of the switch-regularity of two upward orientations of f.
- We introduce the concept of *orthogonal relation* between two vertices of H. This relation establishes the relative position of the two vertices in any planar orthogonal drawing of H. We show that an orthogonal relation is defined between every two vertices of H if and only if all the faces of H are turn-regular.
- We show that if H is turn-regular (i.e., all its faces are turn-regular), then any orthogonal drawing of H such that the orthogonal relations between every two vertices are satisfied is planar.
- We show that if H is turn-regular, then a planar orthogonal drawing of H with optimal area can be computed in $O(n)$ time and space, where n is the number of vertices and bends of H. Furthermore, a planar orthogonal drawing of H with optimal area and minimum total edge length within that area can be computed in $O(n^{7/4} \log n)$ time.
- We present the results of an experimental study on a test suite of planar orthogonal representations of randomly generated biconnected 4-planar graphs. The experiments show that the percentage of regular faces is quite high (95%). Motivated by this result, we have designed compaction heuristics based on the idea of "face regularization." Namely, we decompose non-regular faces into regular ones, and then perform an optimal compaction of the resulting planar orthogonal representation. We implemented our compaction algorithms and experimentally observed that the improvement in area is substantial when compared to the compaction algorithms available in state-of-the-art graph drawing libraries.

The paper is organized as follows. We recall some basic definitions, the notion of switch-regularity and its basic properties in Section 2. In Section 3, we define the orthogonal relations and the concept of turn-regularity, and relate the latter to switch-regularity. Two partial orientations of a turn-regular orthogonal representation and their properties are described in Section 4. In Section 5, we prove the existence of an orthogonal relation between every two vertices of a turn-regular orthogonal representation. The recognition algorithm and the two compaction algorithms are described in Section 6. In Section 7, we present the results of the experimental study. Section 8 contains the conclusions and some open problems. Finally, a table listing the various symbols used throughout the paper is given in the Appendix.

2 Preliminaries

2.1 Basic Definitions

We assume familiarity with graph terminology and basic properties of planar graphs (see, e.g., [5, 18]). The graphs we consider are assumed to be connected. For background on graph drawing, see [6].

A *drawing* of a graph G maps each vertex of G to a distinct point of the plane and each edge (u, v) of G to a simple Jordan curve with endpoints u and v. In an *orthogonal* drawing, each edge is represented as a polygonal chain of alternating horizontal and vertical segments. A drawing is *planar* if no two edges intersect, except, possibly, at common endpoints. A graph is planar if it has a planar drawing. A planar graph whose vertices have degree at most four is said to be 4-*planar*. Two planar drawings of a planar graph G are *equivalent* if, for each vertex v, they have the same clockwise circular sequence of edges incident on v. Hence, the planar drawings of G are partitioned into equivalence classes. Each of those classes is called an *embedding* of G. A planar drawing divides the plane into topologically connected regions, called *faces*. The *external* face is the unbounded region; all other faces are *internal*. Two equivalent planar drawings have the same faces. An *embedded planar* graph is a planar graph with a prescribed embedding.

An *st-digraph* is an acyclic digraph with a single source s (vertex with no incoming edges) and a single sink t (vertex with no outgoing edges). A *planar st-digraph* is an st-digraph that is planar and embedded with vertices s and t on the boundary of the external face. An important property of planar st-digraphs is that the incoming edges of each vertex v appear consecutively around v, as do the outgoing edges. Also, the boundary of each face f consists of two directed paths enclosing f, with common origin and destination.

Let G be an embedded 4-planar graph and let f be a face of G. In the following, we always traverse the boundary of f so that f is on the left, i.e., counterclockwise if f is internal and clockwise if f is external. The boundary of f consists of an alternating circular sequence of vertices and edges. Note that if G is not biconnected, there may be two occurrences of the same edge and multiple occurrences of the same vertex on the boundary of f. We denote by a_f the number of vertices (or edges) of f, each counted with its multiplicity.

Informally speaking, an *orthogonal representation* of an embedded 4-planar graph G describes an equivalence class of orthogonal drawings of G with "similar shape." It consists of a "decorated" version of G where each pair of consecutive edges around a vertex v is assigned an angle multiple of $\pi/2$ and each edge (u, v) is assigned a sequence of bends going from u to v, each a left or right turn. In this paper we consider *planar orthogonal representations*, that is, equivalence classes of orthogonal drawings for which at least one of the drawings is planar; in the rest of the paper, we omit the word planar when referring to orthogonal representations.

An orthogonal representation of G is formally defined as follows. Let v be a vertex of G. We assign an $\alpha \cdot \pi/2$ angle, $1 \leq \alpha \leq 4$, to each pair of consecutive edges around v (note that if v has degree one, the two consecutive edges around v coincide). We refer to these angles as *vertex-angles*. Let e be an edge of G with end-vertices u and v. We assign to e two sequences of $\pi/2$ and $-\pi/2$ angles; one contains the angles on the left of the bends along e when going from u to v, and the other the angles on the left of the bends along e when going from v to u. A $\pi/2$ angle corresponds to a left turn, while a $-\pi/2$ angle corresponds to a right turn. We refer to these angles as *bend-angles*. Note that one of the two sequences associated with e can be obtained from the other by reversing the order and changing the signs of its elements. The following properties must be satisfied by the above assignments:

Property 1 *The sum of the vertex-angles around each vertex is* 2π.

Property 2 *The sum of the vertex-angles minus the sum of the bend-angles along the boundary of each face* f *is*

$$\begin{cases} (2a_f - 4) \cdot \pi/2 & \textit{if } f \textit{ is an internal face,} \\ (2a_f + 4) \cdot \pi/2 & \textit{if } f \textit{ is the external face.} \end{cases}$$

Since each bend can be replaced by a dummy vertex of degree 2, in the rest of the paper we assume, for the sake of simplicity, that orthogonal representations have no bends. We also assume that different drawings of the same orthogonal representation are iso-oriented, i.e., each edge has the same direction and its end-vertices are in the same relative position.

2.2 Switch-Regularity

We recall some terminology and results from [2, 8, 10]. A drawing of a digraph is said to be *upward* if edges are mapped to curves monotonically increasing in a common direction, for instance the vertical one. A digraph is *upward planar* if it admits an upward planar drawing. As we are going to show in the next section, upward planar drawings and orthogonal representations are strictly related. We recall here some notations and results that will be useful in the rest of the paper.

A vertex v of an embedded planar digraph G is said to be *bimodal* if all the incoming/outgoing edges of v appear consecutively around v in the embedding. If all the vertices of G are bimodal then G and its embedding are called *bimodal*. Let f be a face of a bimodal embedded digraph G. A vertex v of f with incident edges e_1 and e_2 is a *switch* if e_1 and e_2 are both incoming or both outgoing edges (note that e_1 and e_2 may coincide if the digraph is not biconnected). In the former case v is a *sink switch* of f, in the latter a *source switch* of f. We denote by $2n_f$ the number of switches of f.

Assign S and L labels to the switches of each face f such that (see Fig. 2(a)): (i) each source or sink of G has exactly one L label; (ii) for each face f, the number of L-labeled switches assigned to f is equal to $n_f - 1$ if f is an internal face, and to $n_f + 1$ if f is the external face. The S-labeled (L-labeled) source switches are called s_S-switches (s_L-switches) and the S-labeled (L-labeled) sink switches are called t_S-switches (t_L-switches). The circular sequence of labels of f so obtained is a *labeling* of f and is denoted by σ_f. Also, S_{σ_f} (L_{σ_f}) denotes the number of S-labels (L-labels) of σ_f. A face f of G labeled in this manner is *upward consistent*.

Property 3 [2] *For each upward consistent face* f,

$$S_{\sigma_f} - L_{\sigma_f} = \begin{cases} 2 & \textit{if } f \textit{ is an internal face,} \\ -2 & \textit{if } f \textit{ is the external face.} \end{cases}$$

Theorem 1 [2] *A bimodal embedded digraph is upward planar if and only if all its faces have an upward consistent labeling.*

Let G be a bimodal embedded digraph such that all its faces have an upward consistent labeling. For each face of G, the S-label (L-label) assigned to a switch intuitively indicates that the angle formed by the two edges identifying the switch is smaller (larger) than π in an upward planar drawing of G. Any such drawing is said to *correspond* to the upward consistent labeling of the

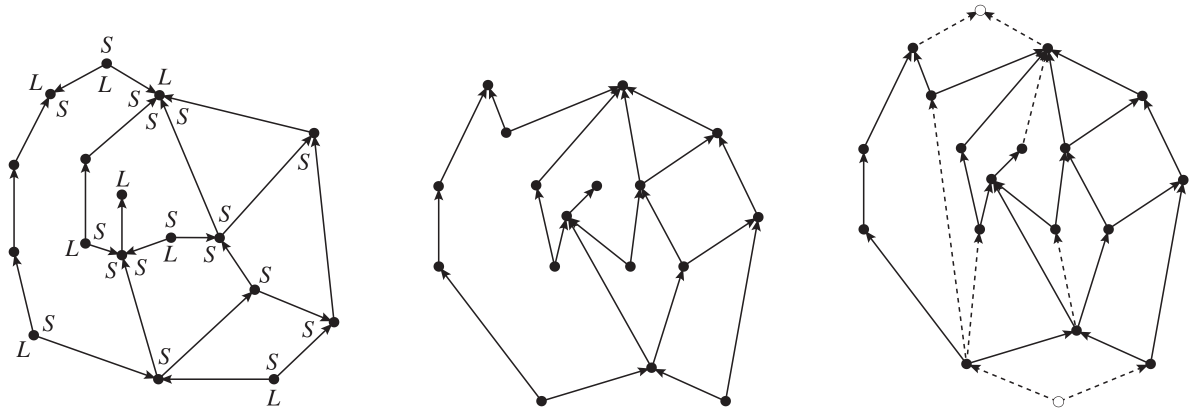

(a) (b) (c)

Figure 2: (a) A bimodal embedded planar digraph G with an upward consistent labeling of its faces. (b) An upward planar drawing of G corresponding to the upward consistent labeling in (a). (c) A complete saturator of G; s and t are represented as white circles, and saturating edges are represented as dashed segments.

faces of G. On the other hand, given an upward planar drawing of a bimodal embedded digraph G, an upward consistent labeling for each face of G can be obtained by simply checking whether the angle formed by each pair of edges identifying a switch is smaller or larger than π. Fig. 2(b) shows an upward planar drawing (corresponding to the upward consistent labeling) of the embedded planar digraph in Fig. 2(a).

Given an embedded upward planar digraph G, a *saturator* of G is a set of edges (each edge a *saturating edge*) plus two vertices s and t connected by edge (s,t). A saturating edge (u,v) is such that:

- Vertices u and v are switches of the same face, or u is a t_L-switch of the external face and $v = t$, or $u = s$ and v is an s_L-switch of the external face.
- If, $u, v \neq s, t$, either u is an s_S-switch and v is an s_L-switch or u is a t_L-switch and v is a t_S-switch. In the former case we say that *u saturates v* and in the latter case we say that *v saturates u*.
- The faces obtained with the insertion of a saturating edge are upward consistent.

A saturator of G is said to be *complete* if for every face f and for every switch u of f labeled L, u is an end-vertex of an edge of the saturator (see Fig. 2(c)). Clearly, adding to G a complete saturator yields a planar st-digraph.

Lemma 1 **[8]** *Every upward planar embedding admits a complete saturator.*

An upward planar embedding can have, in general, several complete saturators. The class of embedded upward planar digraphs for which there exists a unique complete saturator has been characterized in [8]. The characterization is based on a certain type of labeling. Namely, let G be an embedded upward planar digraph. An internal face f of G has a *switch-regular labeling* if

σ_f does not contain two distinct maximal subsequences σ_1 and σ_2 of S-labels such that $S_{\sigma_1} > 1$ and $S_{\sigma_2} > 1$. An external face f of G has a switch-regular labeling if σ_f does not contain two consecutive S-labels. (Note that a switch-regular labeling is called just regular labeling in [8].) A face of G with a switch-regular labeling is a *switch-regular* face. For example all faces of Fig. 2(a) are switch-regular. An embedded upward planar digraph is *switch-regular* if all its faces have a switch-regular labeling. The corresponding embedding is also called *switch-regular*.

Theorem 2 [**8**] *An upward planar embedding has a unique complete saturator if and only if it is switch-regular.*

3 Turn-Regularity and Switch-Regularity

3.1 Orthogonal Relations

Let G be an embedded 4-planar graph, H be an orthogonal representation of G, Γ be a planar drawing of H, and v be a vertex of G. We denote by $x(v)$ and $y(v)$ the x- and y-coordinates of the point representing v in Γ. We define four binary relations on the vertex set of G: for each pair $\{u, v\}$ of vertices of G, these relations determine the relative position of u and v in *all planar drawings* of H.

- $u <_x v$ if $x(u) < x(v)$ for all planar drawings of H; in this case, we say that u is *left* of v and that v is *right* of u.
- $u =_x v$ if $x(u) = x(v)$ for all planar drawings of H; in this case, we say that u is *x-aligned* with v.
- $u <_y v$ if $y(u) < y(v)$ for all planar drawings of H; in this case, we say that u is *below* v and that v is *above* u.
- $u =_y v$ if $y(u) = y(v)$ for all planar drawings of H; in this case, we say that u is *y-aligned* with v.

We refer to the first two binary relations as x-relations and to the second two binary relations as y-relations. As an example, in the orthogonal representation in Fig. 1 $v_2 <_x v_8$, $v_6 =_x v_7$, $v_2 <_y v_3$, and $v_1 =_y v_5$.

We define three new binary relations on the vertex set of G, obtained by combining an x-relation and a y-relation: $=_x \wedge <_y$, $<_x \wedge =_y$, and $<_x \wedge <_y$. These three binary relations together with the binary relations $<_x$ and $<_y$ are collectively referred to as *orthogonal relations*. As an example, in the orthogonal representation in Fig. 1 $v_5 =_x v_4 \wedge v_5 <_y v_4$, $v_1 <_x v_8 \wedge v_1 =_y v_8$, and $v_1 <_x v_7 \wedge v_1 <_y v_7$, while no orthogonal relation holds for $\{v_4, v_6\}$.

3.2 Turn-Regularity

To characterize those orthogonal representations that have an orthogonal relation for each pair of vertices, we introduce the notion of *turn-regularity*.

Let G be an embedded 4-planar graph, H be an orthogonal representation of G, and f be a face of G. For each occurrence of vertex v on the boundary of f, let *prev*(v) and *next*(v) be the edges preceding and following v, respectively, on the boundary of f. Note that *prev*(v) and *next*(v) may coincide if the graph is not biconnected. We associate with each occurrence of v one or two *corners*. Namely:

- If the angle internal to f between $prev(v)$ and $next(v)$ is $\pi/2$ in H, we associate with v one *convex* corner, and say that v corresponds to a *left turn*.
- If the angle internal to f between $prev(v)$ and $next(v)$ is π in H, we associate with v one *flat* corner, and say that v corresponds to a *flat turn*.
- If the angle internal to f between $prev(v)$ and $next(v)$ is $3\pi/2$ in H, we associate with v one *reflex* corner, and say that v corresponds to a *right turn*.
- If the angle internal to f between $prev(v)$ and $next(v)$ is 2π in H, we associate with v an ordered pair of *reflex* corners, and say that v corresponds to a *U-turn*.

Hence, a circular sequence of corners can be associated with the boundary of f. For each corner c of f, let $turn(c) = 1$ if c is convex, $turn(c) = 0$ if c is flat, and $turn(c) = -1$ if c is reflex. As an example, in Fig. 3 a convex corner is associated with v_1, a flat corner with v_2, a reflex corner with v_3, and an ordered pair of reflex corners with v_4. The grey portion of Fig. 3 (and of other figures of the paper) represents the rest of the graph.

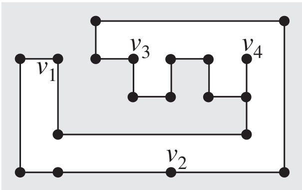

Figure 3: A face of an orthogonal representation. The grey portion represents the rest of the graph.

For each ordered pair $\{c_i, c_j\}$ of corners associated with vertices of f, let $rotation(c_i, c_j) = \sum_c turn(c)$ for all corners c along the boundary of f from c_i (included) to c_j (excluded). If c_i and c_j are associated with distinct vertices v_i and v_j, respectively, $rotation(c_i, c_j) \cdot \pi/2$ is the net angle turned between $prev(v_i)$ and $prev(v_j)$. As an example, in Fig. 3 let c_1, c_2, and c_3 be the corners associated with v_1, v_2, and v_3, respectively, and let $\{c'_4, c''_4\}$ be the ordered pair of corners associated with v_4; $rotation(c_1, c_2) = 3$, $rotation(c_3, c'_4) = 1$, $rotation(c_3, c''_4) = 0$, and $rotation(c_3, c_1) = -3$.

The following property is a direct consequence of the results in [31].

Property 4 *For each face f,*

$$rotation(c_i, c_i) = \begin{cases} 4 & \textit{if } f \textit{ is an internal face,} \\ -4 & \textit{if } f \textit{ is the external face.} \end{cases}$$

From the definition of $rotation(c_i, c_j)$ and from Property 4, Properties 5 and 6 easily follow.

Property 5 *For each ordered triplet of corners $\{c_i, c_j, c_k\}$ on the boundary of a face,*

$$rotation(c_i, c_j) = rotation(c_i, c_k) + rotation(c_k, c_j).$$

Property 6 *For each face f, $rotation(c_i, c_j) = 2$ if and only if*

$$rotation(c_j, c_i) = \begin{cases} 2 & \textit{if } f \textit{ is an internal face,} \\ -6 & \textit{if } f \textit{ is the external face.} \end{cases}$$

Two reflex corners c_i and c_j are called *kitty corners* if $rotation(c_i, c_j) = 2$ or $rotation(c_j, c_i) = 2$. In Fig. 1, the corners associated with vertices v_4 and v_6 are kitty corners. A face of an orthogonal representation is *turn-regular* if it has no kitty corners. An orthogonal representation is *turn-regular* if all its faces are turn-regular.

The reflex corners of a face of an orthogonal representation can be partitioned into four classes. Let f be a face of an orthogonal representation, and c be a reflex corner associated with vertex v of f.

- *NE-corners*: (*i*) if v is the left end-vertex of $prev(v)$ and the bottom end-vertex of $next(v)$; (*ii*) if v is the left end-vertex of both $prev(v)$ and $next(v)$, and c is the first corner associated with v; (*iii*) if v is the bottom end-vertex of both $prev(v)$ and $next(v)$, and c is the second corner associated with v.

- *NW-corners*: (*i*) if v is the bottom end-vertex of $prev(v)$ and the right end-vertex of $next(v)$; (*ii*) if v is the bottom end-vertex of both $prev(v)$ and $next(v)$, and c is the first corner associated with v; (*iii*) if v is the right end-vertex of both $prev(v)$ and $next(v)$, and c is the second corner associated with v.

- *SW-corners*: (*i*) if v is the right end-vertex of $prev(v)$ and the top end-vertex of $next(v)$; (*ii*) if v is the right end-vertex of both $prev(v)$ and $next(v)$, and c is the first corner associated with v; (*iii*) if v is the top end-vertex of both $prev(v)$ and $next(v)$, and c is the second corner associated with v.

- *SE-corners*: (*i*) if v is the top end-vertex of $prev(v)$ and the left end-vertex of $next(v)$; (*ii*) if v is the top end-vertex of both $prev(v)$ and $next(v)$, and c is the first corner associated with v; (*iii*) if v is the left end-vertex of both $prev(v)$ and $next(v)$, and c is the second corner associated with v.

As an example, the reflex corner associated with vertex v_3 in Fig. 3 is a SW-corner, and the two corners associated with vertex v_4 are a SE-corner (the first one) and a SW-corner (the second one). The following property is a direct consequence of the definition of kitty corners.

Property 7 *Two reflex corners are kitty corners only if they form either a SW-NE pair or a SE-NW pair.*

3.3 Turn-Regularity and Switch-Regularity

Let G be an embedded 4-planar graph, H be an orthogonal representation of G, and Γ be a planar drawing of H. Let Γ_r be an orientation of Γ such that all vertical segments are directed upward and all horizontal segments are directed rightward, and let Γ_ℓ be an orientation of Γ such that all vertical segments are directed upward and all horizontal segments are directed leftward. Observe that Γ_r is an upward planar drawing in the North-East direction and that Γ_ℓ is an upward planar drawing in the North-West direction. Γ_r and Γ_ℓ induce two orientations on H. We denote the oriented orthogonal representations by H_r and H_ℓ, respectively (see Fig. 4). In turn, H_r and H_ℓ induce two orientations on G. We denote the embedded 4-planar digraphs by G_r^H and G_ℓ^H, respectively. Observe that G_r^H and G_ℓ^H are embedded upward planar digraphs. Also, note that different orthogonal representations of G induce, in general, different orientations on G; since we work with a fixed orthogonal representation of a graph, we use G_r and G_ℓ in the rest of the paper, omitting the reference to H.

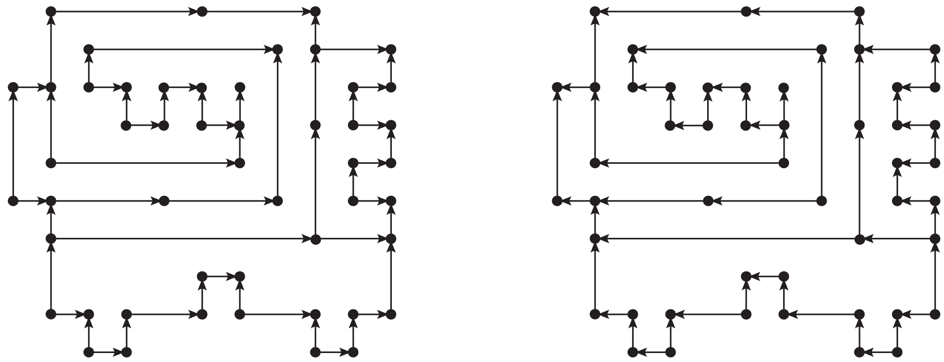

(a) (b)

Figure 4: Two orientations of the same orthogonal representation. (a) H_r. (b) H_ℓ.

To simplify the proofs of this section, we assume that the boundary of each face of H contains no vertices of degree one and no multiple occurrences of the same vertex. There is no loss of generality in this assumption, since a face f whose boundary contains vertices of degree one or multiple occurrences of the same vertex can be replaced, for the purpose of the proofs of this section, by an *expanded* face obtained from f as follows (see Fig. 5): (i) each degree one vertex v of f is replaced by a pair of vertices and by an edge connecting them, perpendicular to the edge incident with v in f; (ii) each vertex occurring k times on the boundary of f is replaced by k distinct vertices (note that $k \leq 4$ in an orthogonal representation); (iii) each edge occurring twice on the boundary of f is replaced by two edges having the same direction. Note that the proofs of this section consider a single face independently from the rest of H; thus, the expansion process is to be considered local to a single face of H. With the above assumption, each vertex v_i on the boundary of a face of H has exactly one associated corner, as described in Section 3.2. We denote with c_i the corner associated with v_i.

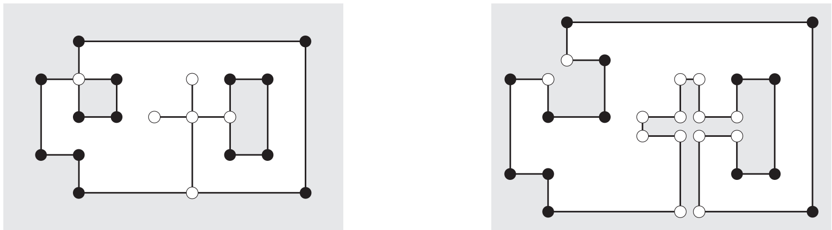

(a) (b)

Figure 5: (a) A face f of an orthogonal representation. The vertices affected by the expansion are represented as white circles. (b) The expanded face corresponding to f.

In order to show the connection between switch-regularity and turn-regularity, we first establish a connection between switches of G_r and G_ℓ and corners of H. We observe that since Γ_r (Γ_ℓ) is an upward planar drawing of G_r (G_ℓ), it induces an upward consistent labeling on G_r (G_ℓ). Let f be a face of G_r or G_ℓ; two switches v_i and v_j of f are *consecutive* if no vertex on the portion of the boundary of f from v_i to v_j is a switch of f.

Property 8 *Let v_i and v_j be two consecutive switches of a face of G_r or G_ℓ, and c_i and c_j be the associated corners. One of the following holds:*

- *If v_i and v_j are both L-labeled switches (an LL-transition), then $rotation(c_i, c_j) = -2$.*
- *If v_i is an L-labeled switch and v_j is an S-labeled switch (an LS-transition), then $rotation(c_i, c_j) = -1$.*
- *If v_i is an S-labeled switch and v_j is an L-labeled switch (an SL-transition), then $rotation(c_i, c_j) = 1$.*
- *If v_i and v_j are both S-labeled switches (an SS-transition), then $rotation(c_i, c_j) = 2$.*

Proof: We only consider G_r; the proof for G_ℓ is similar. Since v_i and v_j are consecutive switches, they are either connected by a single edge or by a "staircase." Fig. 6 shows the possible arrangements of non-flat corners between c_i and c_j for a face of G_r. Note that flat corners do not affect the value of *rotation*, and can safely be ignored. □

Let σ be a sequence of L-labels and S-labels. In the rest of the section, we denote the number of LL-transitions in σ as n^σ_{LL}, the number of LS-transitions in σ as n^σ_{LS}, the number of SL-transitions in σ as n^σ_{SL}, and the number of SS-transitions in σ as n^σ_{SS}.

Property 9 *Let σ be a sequence of L-labels and S-labels, and let $\rho = l_1 \sigma l_n$. We have:*

- *if $l_1 = L$ and $l_n = L$, then $n^\rho_{LS} = n^\rho_{SL}$;*
- *if $l_1 = L$ and $l_n = S$, then $n^\rho_{LS} = n^\rho_{SL} + 1$;*
- *if $l_1 = S$ and $l_n = L$, then $n^\rho_{SL} = n^\rho_{LS} + 1$;*
- *if $l_1 = S$ and $l_n = S$, then $n^\rho_{SL} = n^\rho_{LS}$.*

Lemma 2 *Let σ be a sequence of L-labels and S-labels, and let $\rho = L\sigma S$. Then, we have:*

- *$L_\rho = n^\rho_{LL} + n^\rho_{SL} + 1$;*
- *$S_\rho = n^\rho_{SS} + n^\rho_{SL} + 1$.*

Proof: We first prove that $L_\rho = n^\rho_{LL} + n^\rho_{SL} + 1$. For each sequence of L-labels, the number of labels is equal to the number of LL-transitions plus one. So the number of L-labels in ρ is equal to the number of LL-transitions in ρ plus the number of sequences of L-labels in ρ. In turn, since each sequence of L-labels in ρ, except the first one, is preceded by an S-label, the number of sequences of L-labels is equal to the number of SL-transitions in ρ plus one. Hence, the claim.

The proof that $S_\rho = n^\rho_{SS} + n^\rho_{SL} + 1$ is similar, once we observe that each sequence of S-labels in ρ, except the last one, is followed by an L-label. □

We now describe how to simplify portions of the boundary of a face of H by removing switches (and hence corners), without affecting the value of *rotation* for the remaining corners. In the rest of the section, we denote the label of a switch v_i by $label(v_i)$ and add an index to *rotation* to denote with respect to which face it is computed. Let $\{v_1, v_2, v_3, v_4\}$ be four consecutive switches of a face f of H. If either $label(v_2) = S$ and $label(v_3) = L$, or $label(v_2) = L$ and $label(v_3) = S$, we *collapse* the pair of switches $\{v_2, v_3\}$ by replacing the portion of the boundary of f between $prev(v_2)$ (included) and $next(v_3)$ (included) with a single edge having the same direction as $prev(v_2)$ and $next(v_3)$.

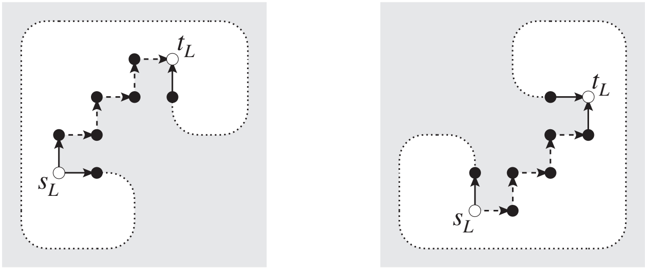

(a)

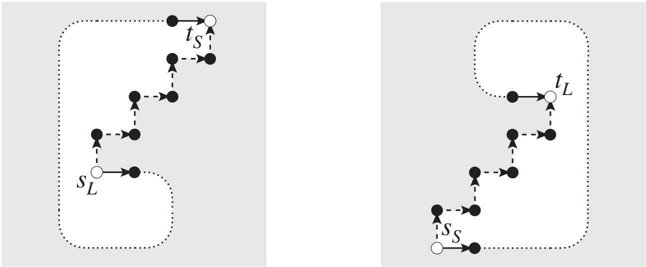

(b)

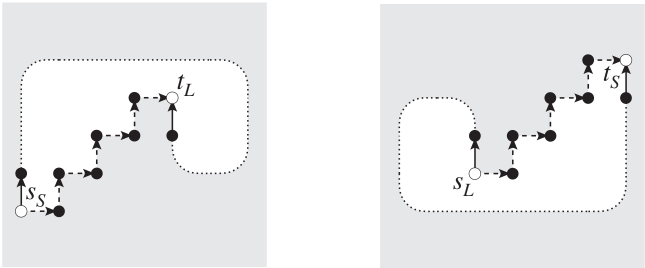

(c)

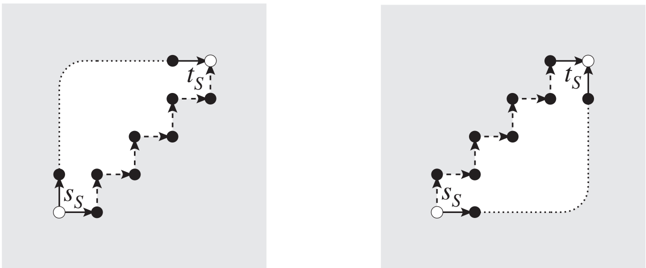

(d)

Figure 6: The shape of a face f of H_r between two consecutive switches. The "staircase" portion of the boundary of f (represented by dashed segments) may consist of a single edge. (a) LL-transitions. (b) LS-transitions. (c) SL-transitions. (d) SS-transitions.

Lemma 3 *Let f be a face of H and let $\{v_1, v_2, v_3, v_4\}$ be four consecutive switches of f in G_r or G_ℓ such that either label$(v_2) = S$ and label$(v_3) = L$, or label$(v_2) = L$ and label$(v_3) = S$. Let f' be the face obtained from f by collapsing the pair of switches $\{v_2, v_3\}$. Then, rotation$_f(c_1, c_4) =$ rotation$_{f'}(c_1, c_4)$.*

Proof: We prove the case $label(v_2) = S$ and $label(v_3) = L$ by case analysis. Let σ be the sequence of labels of v_1, v_2, v_3, and v_4 in f and σ' be the sequence of labels of v_1 and v_4 in f'. From Property 8 we have:

σ	$rotation_f(c_1, c_4)$	σ'	$rotation_{f'}(c_1, c_4)$
$S\ SL\ S$	$2+1-1=2$	SS	2
$S\ SL\ L$	$2+1-2=1$	SL	1
$L\ SL\ S$	$-1+1-1=-1$	LS	-1
$L\ SL\ L$	$-1+1-2=-2$	LL	-2

The proof of the case $label(v_2) = L$ and $label(v_3) = S$ is analogous. □

Lemma 4 *Let f be a face of H, $V = \{v_1, W, v_k\}$ be a sequence of k consecutive switches of f in G_r or G_ℓ, and σ be the sequence of labels corresponding to W. Let f' be the face obtained from f by applying the collapse operation to pairs of switches of V as many times as possible, and $V' = \{v_1, W', v_k\}$ be the resulting sequence of switches. We have that:*

- *$rotation_f(c_1, c_k) = rotation_{f'}(c_1, c_k)$;*
- *if $S_\sigma > L_\sigma$, then W' is a sequence of $S_\sigma - L_\sigma$ S-labeled switches;*
- *if $S_\sigma < L_\sigma$, then W' is a sequence of $L_\sigma - S_\sigma$ L-labeled switches;*
- *if $S_\sigma = L_\sigma$, then W' is empty.*

Proof: Clearly, $rotation_f(c_1, c_k) = rotation_{f'}(c_1, c_k)$ by Lemma 3. We now prove the case $S_\sigma > L_\sigma$; the other two cases are analogous.

If $L_\sigma = 0$, then σ consists entirely of S-labeled switches, and so the lemma is trivially true because no collapse operation can be applied, and hence $W' = W$. Otherwise, there is at least one subsequence $\{v_{i-1}, v_i, v_{i+1}, v_{i+2}\}$, $1 < i < k-1$ of V such that either $label(v_i) = S$ and $label(v_{i+1}) = L$, or $label(v_i) = L$ and $label(v_{i+1}) = S$. Hence, the collapse operation can be applied, and the process can be repeated as long as $L_\sigma > 0$. The first and last switches of V, v_1 and v_k, will never be eliminated, because they can only match v_{i-1} or v_{i+2} in the pattern of four consecutive switches. Their presence also guarantees that every SL or LS subsequence of σ can be matched to a pattern and thus eliminated. Since one L-labeled switch and one S-labeled switch are removed each time the collapse operation is applied, there will be $S_\sigma - L_\sigma$ S-labeled switches left when the process is complete. □

An example of the collapsing process is shown in Fig. 7. Let σ be the sequence of labels corresponding to a sequence W of consecutive switches of a face of G_r or G_ℓ, and let W' the sequence of switches obtained from W as described in Lemma 4. In the rest of the section, the sequence of labels σ' corresponding to W' is referred to as the *collapsed version* of σ.

Theorem 3 *An orthogonal representation H of an embedded 4-planar graph G is turn-regular if and only if the embedded upward planar digraphs G_r and G_ℓ are both switch-regular.*

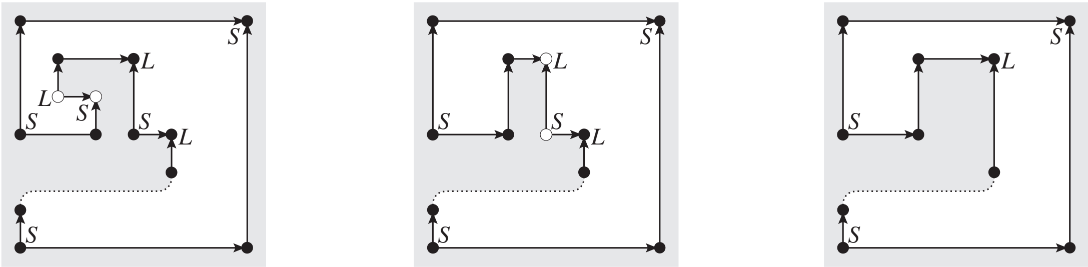

(a) $S\ SS(SL)LS\ L$ (b) $S\ SS(LS)\ L$ (c) $S\ SS\ L$

Figure 7: Collapsing a sequence of switches. In each step, the switches being collapsed are represented as white circles, and the SL or LS pair of labels being removed is shown in parentheses.

Proof: In the proof we adopt the regular expression formalism (see, e.g., [20]); in particular, we use the concatenation of strings and the Kleene star operator.

Only if. We prove the claim for G_r; the proof for G_ℓ is similar. Suppose, for a contradiction, that H is turn-regular and G_r is not switch-regular; then, there is a face f of G_r that does not have a switch-regular labeling.

We first consider the case in which f is an internal face: σ_f must contain two distinct maximal subsequences of S-labels with length greater than one. Thus, the labeling of f can be expressed as $\sigma_f = SS\sigma_1 SSS^* L\sigma_2$, where $\sigma_1 = L((SL)^*L^*)^*$ is any sequence with an initial L-label and with no two consecutive S-labels, and σ_2 is any (possibly empty) sequence of S-labels and L-labels such that σ_f satisfies Property 3. We show that there is an L-labeled switch v_j in σ_1 and an L-labeled switch v_k in $L\sigma_2$ such that c_j and c_k are kitty corners. This implies that f and thus H are not turn regular; a contradiction.

To simplify matters, we consider $\sigma'_f = SS\sigma'_1 SSS^* L\sigma'_2$ instead of σ_f, where σ'_1 and σ'_2 are the collapsed versions of σ_1 and σ_2. Note that every label in σ'_f is also in σ_f, and that the corresponding switch is associated with a corner of f. In the regular expression describing σ_1, each S-label is followed by an L-label; thus, the initial L-label of σ_1 guarantees that σ_1 contains at least one more L-label than S-label. Hence, by Lemma 4, $\sigma'_1 = LL^*$. Let $l \geq 1$ be the length of σ'_1 and $s \geq 2$ be the length of the SSS^* subsequence between σ'_1 and $L\sigma'_2$. The following table shows the general structure of σ'_f and the values of *rotation*, which can be easily verified by Property 8.

			σ'_1				SSS^*					σ'_2
v_i	v_1	v_2	v_3	v_4	$\cdots$	v_{2+l}	v_{3+l}	v_{4+l}	$\cdots$	v_{2+l+s}	v_{3+l+s}	$\cdots$
σ'_f	S	S	L	L	$\cdots$	L	S	S	$\cdots$	S	L	$\cdots$
$rotation(c_1, c_i)$	0	2	3	1	$\cdots$	$5-2l$	$4-2l$	$6-2l$	$\cdots$	$2(s+1)-2l$	$2s+3-2l$	$\cdots$

We consider two cases for the value of s, namely $2 \leq s \leq l+1$ and $s \geq l+2$, and prove the claim differently in the two cases.

If $2 \leq s \leq l+1$, then v_{4-s+l} is an L-labeled switch, since $3 \leq 4-s+l \leq 2+l$; v_{3+l+s} is also an L-labeled switch, and we have, by Property 5:

$$\begin{aligned} rotation(c_{4-s+l}, c_{3+l+s}) &= rotation(c_1, c_{3+l+s}) - rotation(c_1, c_{4-s+l}) \\ &= (2s+3-2l) - \{5 - 2[(4-s+l)-2]\} \\ &= 2 \end{aligned}$$

Hence, c_{4-s+l} and c_{3+l+s} are kitty corners.

If $s \geq l+2$, we consider the structure of the sequence $L\sigma_2 S$, where the final S is the label of v_1 (we recall that σ_f is a circular sequence). For ease of reference, we denote $L\sigma_2 S$ as ρ. By Properties 8 and 9, we have:

$$\begin{aligned} rotation(c_{3+l+s}, c_1) &= -2n^{\rho}_{LL} - n^{\rho}_{LS} + n^{\rho}_{SL} + 2n^{\rho}_{SS} \\ &= -2n^{\rho}_{LL} - (n^{\rho}_{SL}+1) + n^{\rho}_{SL} + 2n^{\rho}_{SS} \\ &= 2(n^{\rho}_{SS} - n^{\rho}_{LL}) - 1 \end{aligned}$$

On the other hand, by Properties 4 and 5, and by the above table, we have:

$$\begin{aligned} rotation(c_{3+l+s}, c_1) &= 4 - rotation(c_1, c_{3+l+s}) \\ &= 4 - (2s + 3 - 2l) \\ &\leq 4 - [2(l+2) + 3 - 2l] \\ &\leq -3 \end{aligned}$$

And thus, by combining the two results:

$$\begin{aligned} 2(n^{\rho}_{SS} - n^{\rho}_{LL}) - 1 &\leq -3 \\ n^{\rho}_{LL} - n^{\rho}_{SS} &\geq 1 \end{aligned}$$

Now, by Lemma 2, we have:

$$\begin{aligned} L_{\sigma_2} &= L_{\rho} - 1 = n^{\rho}_{LL} + n^{\rho}_{SL} \\ S_{\sigma_2} &= S_{\rho} - 1 = n^{\rho}_{SS} + n^{\rho}_{SL} \end{aligned}$$

And thus $L_{\sigma_2} > S_{\sigma_2}$, because

$$\begin{aligned} L_{\sigma_2} - S_{\sigma_2} &= n^{\rho}_{LL} - n^{\rho}_{SS} \\ &\geq 1 \end{aligned}$$

Let f' be the face obtained from f by repeatedly collapsing pairs of consecutive switches of $\{v_{3+l+s}, \ldots, v_1\}$. By Lemma 4 and the above discussion, σ'_2 has only L-labels; let v_k be the last (L-labeled) switch in σ'_2. Then, by Properties 4 and 5, and by Lemma 4, we have:

$$\begin{aligned} rotation(c_3, c_k) &= 4 - rotation_f(c_k, c_3) \\ &= 4 - rotation_{f'}(c_k, c_3) \end{aligned}$$

And since between v_k and v_3 in f' there are an LS-, an SS-, and an SL-transition, we have, by Property 8:

$$\begin{aligned} rotation(c_3, c_k) &= 4 - (-1 + 2 + 1) \\ &= 2 \end{aligned}$$

Hence, c_3 and c_k are kitty corners.

We now consider the case in which f is the external face: σ_f must contain two consecutive S-labels. Thus, the labeling of f can be expressed as $\sigma_f = SS\sigma_1$ where σ_1 is any (possibly empty) sequence of S-labels and L-labels such that σ_f satisfies Property 3. Again, to simplify matters, we consider $\sigma'_f = SS\sigma'_1$ instead of σ_f, where σ'_1 is the collapsed version of σ_1. Note that, by Property 3, $L_{\sigma_1} > S_{\sigma_1}$. Hence, by Lemma 4, σ'_1 has only L-labels; in particular, since also σ'_f must satisfy Property 3, $\sigma'_1 = LLLL$. The following table shows the structure of σ'_f and the values of *rotation*, which can be easily verified by Property 8.

v_i	v_1	v_2	v_3	v_4	v_5	v_6
σ'_f	S	S	L	L	L	L
$rotation(c_1, c_i)$	0	2	3	1	-1	-3

Then, by Property 5, we have:

$$\begin{aligned} rotation(c_3, c_6) &= rotation(c_1, c_6) - rotation(c_1, c_3) \\ &= -6 \end{aligned}$$

Hence, by Property 6, c_3 and c_6 are kitty corners.

If. Suppose, for a contradiction, that G_r and G_ℓ are both switch-regular and H is not turn-regular; then, there is a face f of H with a pair $\{c_j, c_k\}$ of kitty corners. Note that c_j and c_k are associated with a pair of L-labeled switches in either G_r (see Fig. 8(a)) or G_ℓ (see Fig. 8(b)).

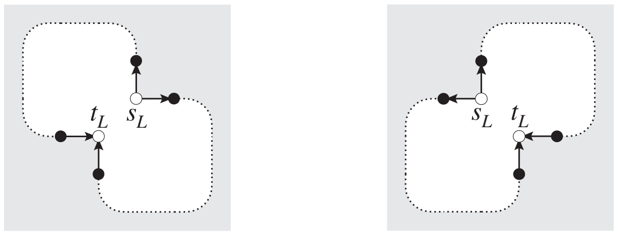

(a) (b)

Figure 8: The two possibilities for a pair of kitty corners: (a) a SW-NE pair of corners associated with two switches in G_r; (b) a SE-NW pair of corners associated with two switches in G_ℓ.

We assume that c_j and c_k are associated with the L-labeled switches v_j and v_k in G_r; the proof for G_ℓ is similar. The labeling of the face f of G_r can be expressed as $\sigma_f = L\sigma_1 L\sigma_2$, where σ_1 and σ_2 are any two sequences of S-labels and L-labels such that σ_f satisfies Property 3. The following table shows the general structure of σ_f.

		σ_1				σ_2					
v_i	v_j	v_{j+1}	$\cdots$	v_{k-1}	v_k	v_{k+1}	$\cdots$	v_n	v_1	$\cdots$	$v_j - 1$
σ_f	L		$\cdots$		L			$\cdots$			
$rotation(c_j, c_i)$	0		$\cdots$		2			$\cdots$			

We show that σ_1 contains at least two consecutive S-labeled switches and that, if f is an internal face, σ_2 also contains at least two consecutive S-labeled switches. Thus, f does not have a switch-regular labeling and G_r is not switch-regular; a contradiction. For ease of reference, we denote $L\sigma_1 L$ as ρ. By Property 9, $n^\rho_{LS} = n^\rho_{SL}$, and thus, by Property 8, the total contribution of the LS- and SL-transitions to $rotation(c_j, c_k)$ is zero. Since $rotation(c_j, c_k) = 2$, still from Property 8, it follows that $n^\rho_{SS} > n^\rho_{LL} \geq 0$. Hence, σ_1 contains at least two consecutive S-labeled switches. If f is an internal face, the same argument can be used to prove that σ_2 contains at least two consecutive S-labeled switches, because, by Property 6, also $rotation(c_k, c_j) = 2$. □

4 Orientations and Paths

Let G be an embedded 4-planar graph and H be a turn-regular orthogonal representation of G. As seen in Section 2.2, a complete saturator of an embedded upward planar digraph consists of two vertices s and t and a set of (directed) saturating edges. Fig. 9(a) shows a complete saturator of graph G_r corresponding to the oriented orthogonal representation H_r shown in Fig. 4(a). Fig. 9(b) shows a complete saturator of graph G_ℓ corresponding to the oriented orthogonal representation

H_ℓ shown in Fig. 4(b). In the rest of the paper we never consider the saturating edges of G_r and G_ℓ incident with s or t, even when not explicitly stated. Let f be an internal face of H; a maximal vertical or horizontal chain of f is said to be *unconstrained* if both its end-vertices correspond to a right turn of f. Note that an unconstrained maximal chain of f may consist of a single, degree one vertex.

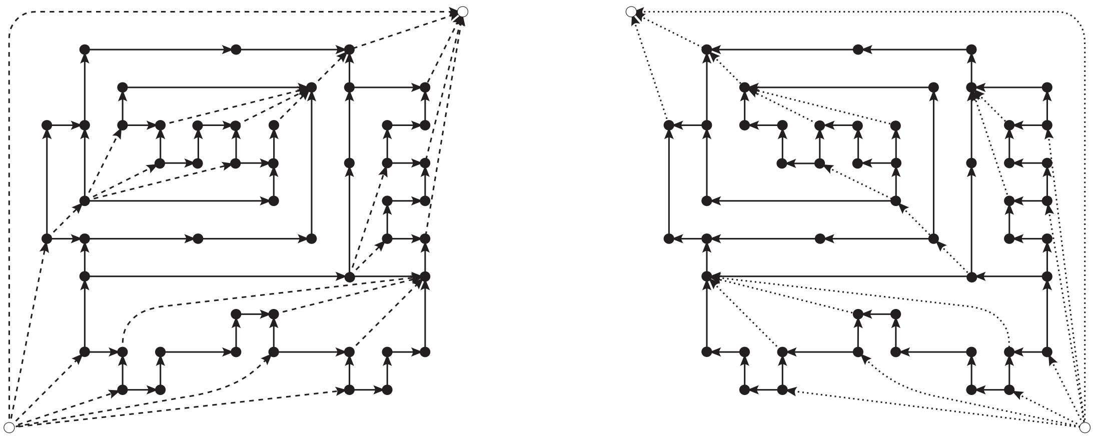

(a) (b)

Figure 9: (a) G_r (edges represented as solid segments) and a complete saturator (edges represented as dashed segments). (b) G_ℓ (edges represented as solid segments) and a complete saturator (edges represented as dotted segments).

We now construct two partially-directed graphs, one representing the "left" relation between maximal vertical chains of H, the other representing the "below" relation between maximal horizontal chains of H. The graph representing the "left" relation between maximal vertical chains of H is constructed as follows. We first augment H with the saturating edges of G_r and G_ℓ incident with an end-vertex of an unconstrained maximal vertical chain of H. We then orient the horizontal edges of H from left to right, reverse the orientation of the saturating edges of G_ℓ, and leave the vertical edges of H not oriented so that they can be traversed in both ways. We denote by H_x the resulting graph (see Fig. 10(a)). Similarly, the graph representing the "below" relation between maximal horizontal chains of H is constructed as follows. We first augment H with the saturating edges of G_r and G_ℓ incident with an end-vertex of an unconstrained maximal horizontal chain of H. We then orient the vertical edges of H from bottom to top and leave the horizontal edges of H not oriented so that they can be traversed in both ways. We denote by H_y the resulting graph (see Fig. 10(b)).

The following theorem shows how turn-regularity characterizes those orthogonal representations for which the "left" relation between maximal vertical chains and the "below" relation between maximal horizontal chains are uniquely determined.

Theorem 4 *Let H be an orthogonal representation. H_x and H_y are uniquely determined if and only if H is turn-regular.*

Proof: Easily follows from Theorem 3, the construction of H_x and H_y, and the definition of switch-regular embedded upward planar digraph. □

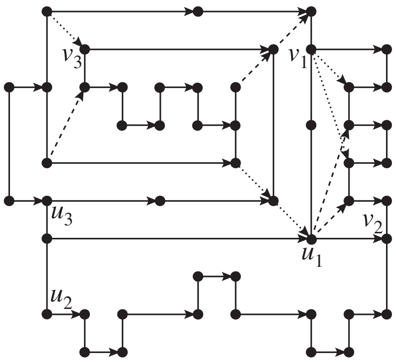

(a)

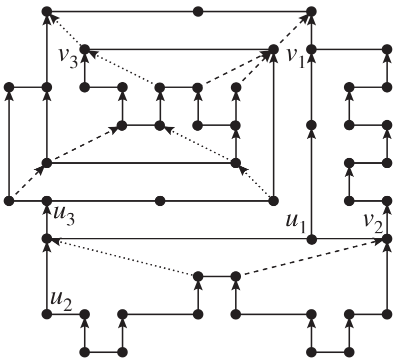

(b)

Figure 10: (a) H_x. (b) H_y. Both graphs are obtained using the complete saturators shown in Fig. 9.

Note that H_x and H_y are no longer orthogonal representations, and may, in general, be non-planar. From the definition of saturator, it follows that each saturating edge from G_r and G_ℓ used in the construction of H_x and H_y has both end-vertices on the same face of H. Two saturating edges in H_x or H_y are said to *cross* each other if their end-vertices appear alternately on the boundary of a common face of H. In the rest of the paper we refer to the maximal chains of non-oriented edges of H_x as maximal vertical chains of H_x, and denote by $mvc(v)$ the maximal vertical chain of H_x containing vertex v. Analogously, we refer to the maximal chains of non-oriented edges of H_y as maximal horizontal chains of H_y, and denote by $mhc(v)$ the maximal horizontal chain of H_y containing vertex v.

In the rest of this section, we present a series of technical lemmas that will be used in Section 5. In the proofs, we can assume the absence of degree one vertices; otherwise, the expansion mechanism described at the beginning of Section 3.3 can be applied.

Lemma 5 *Let H be a turn-regular orthogonal representation. Let (u, v) be a saturating edge from G_r not used in the construction of H_x (H_y); then, there exists a path from u to v in H_x (H_y). Similarly, let (u, v) be a saturating edge from G_ℓ not used in the construction of H_x (H_y); then, there exists a path from v to u in H_x (from u to v in H_y).*

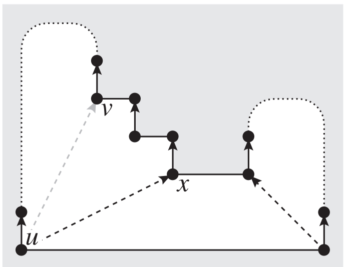

(a)

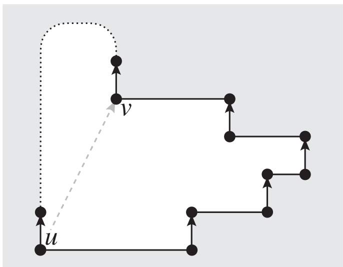

(b)

Figure 11: The two cases in the proof of Lemma 5.

Proof: We prove the claim for a saturating edge (u, v) from G_r not used in the construction of H_y; the proof for the other three cases is similar.

Let f be the face of H containing u and v. We consider in detail the case in which u is an s_S-switch and v an s_L-switch of f in G_r; the case in which u is a t_L-switch and v a t_S-switch of f in G_r is similar. Let $f(u, v)$ be the portion of f from u to v. Clearly, the last turn of $f(u, v)$ before v is a left turn, otherwise v would be the leftmost vertex of an unconstrained horizontal chain and (u, v) would be used in the construction of H_y. If $f(u, v)$ contains at least one unconstrained horizontal chain whose leftmost vertex is an s_L-switch of f in G_r, let *uhc* be the last of these chains and x be the leftmost vertex of *uhc* (see Fig. 11(a)). A saturating edge (u, x) from G_r is used in the construction of H_y, and the path between u and v in H_y consists of (u, x) and the subpath of $f(u, v)$ from x to v. Note that all the vertical edges of the subpath are traversed according to their direction in H_y (from bottom to top). If $f(u, v)$ does not contain such an unconstrained horizontal chain (see Fig. 11(b)), then the path between u and v in H_y is $f(u, v)$ itself. Note, again, that all the vertical edges of this path are traversed according to their direction in H_y (from bottom to top). If not, let x be the top end-vertex of the first vertical edge of $f(u, v)$ traversed from top to bottom; then, the corners associated with v and x would be kitty corners, and H would not be turn-regular. $\square$

Lemma 6 *Let H be a turn-regular orthogonal representation, and let (u, v) and (x, y) be two saturating edges in H_x or H_y crossing each other; then, (u, v) and (x, y) cannot be both from G_r or both from G_ℓ.*

Proof: We prove the claim for H_y; the proof for H_x is similar. In the proof, we denote the corner associated with vertex v as c_v. Suppose, for a contradiction, that (u, v) and (x, y) are both from G_r. We recall that u, v, x and y belong to the same face of H; four cases are possible:

1. u and x are both s_S-switches (see Fig. 12(a)). It follows that v and y are both t_L-switches. Let p be the portion of the boundary of f between u and x not containing v and y. Clearly, there must be at least one right turn in p; let w be the corresponding vertex. Then, the $\{c_w, c_y\}$ and $\{c_w, c_y\}$ are two pairs of kitty corners, and H is not turn-regular; a contradiction.

2. u is an s_S-switch and x is a t_L-switch (see Fig. 12(b)). It follows that v is an s_L-switch and y is a t_S-switch. Then, $\{c_x, c_v\}$ is a pair of kitty corners, and H is not turn-regular; a contradiction.

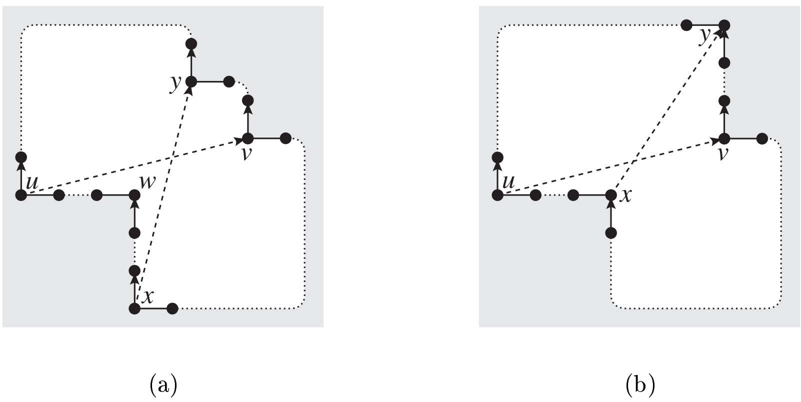

Figure 12: Impossible cases of crossing saturating edges in H_y.

3. u is a t_L-switch and x is an s_S-switch. Similar to Case 2.

4. u and x are both t_L-switches. Similar to Case 1.

The proof for (u, v) and (x, y) both from G_ℓ is analogous. □

Lemma 7 *Let H be a turn-regular orthogonal representation, and let (u, v) and (x, y) be two saturating edges in H_x crossing each other; then, either $mvc(u) = mvc(x)$ or $mvc(v) = mvc(y)$. Similarly, let (u, v) and (x, y) be two saturating edges in H_y crossing each other; then, either $mhc(u) = mhc(x)$ or $mhc(v) = mhc(y)$.*

Proof: We prove the claim for H_y; the proof for H_x is similar. In the proof, we denote the corner associated with vertex v as c_v. By Lemma 6, (u, v) and (x, y) cannot be both from G_r or G_ℓ. We assume that (u, v) is from G_r and (x, y) is from G_ℓ. We recall that u, v, x and y belong to the same face of H. The proof proceeds by case analysis:

1. u is an s_S-switch in G_r and x is an s_S-switch in G_ℓ (see Fig. 13(a)). It follows that v is a t_L-switch in G_r and y is a t_L-switch in G_ℓ. Let p be the portion of the boundary of f between u and x not containing v and y. We have that p contains no right turns; suppose the opposite, for a contradiction, and let w be the corresponding vertex. Either $\{c_w, c_v\}$ or $\{c_w, c_y\}$ is a pair of kitty corners, and H is not turn-regular; a contradiction. Hence, $mhc(u) = mhc(x)$.

2. u is an s_S-switch in G_r and x is a t_L-switch in G_ℓ (see Fig. 13(b)). It follows that v is an s_L-switch in G_r and y is a t_S-switch in G_ℓ. Note that, from the construction of H_y, $mhc(v)$ and $mhc(x)$ are unconstrained; let w and z be the other end-vertices of $mhc(v)$ and $mhc(x)$, respectively. This case is impossible, since $\{c_v, c_z\}$ and $\{c_w, c_x\}$ are two pairs of kitty corners and H is not turn-regular.

3. u is a t_L-switch in G_r and x is an s_S-switch in G_ℓ. Similar to Case 2.

4. u is a t_L-switch in G_r and x is a t_L-switch in G_ℓ. Similar to Case 1; hence, $mhc(v) = mhc(y)$.

The proof for (u, v) from G_ℓ and (x, y) from G_r is analogous. □

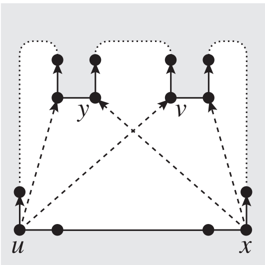

(a)

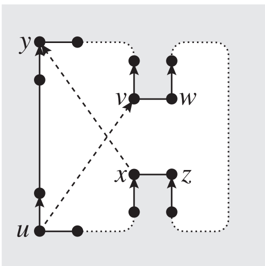

(b)

Figure 13: Crossing saturating edges of H_y.

Let H be an orthogonal representation and let f be a face of H. Two reflex corners of f are *mutually visible* if there exists a planar drawing of H such that the vertices associated with the two corners can be connected by a straight-line segment that does not cross any edge of f.

Lemma 8 *Let f be a face of an orthogonal representation, c_u be a NE-, or a NW-, or a SW-, or a SE-corner of f and c_v be a NW-, or a SW-, or a SE-, or a NE-corner of f, respectively. If c_u and c_v are mutually visible, then $rotation(c_u, c_v) = 1$.*

Proof: We prove the claim for a NE-corner c_u and a NW-corner c_v. The proofs for the other three cases are similar. Let H be an orthogonal representation and Γ be a planar drawing of H such that vertices u and v, associated with corners c_u and c_v, can be connected by a straight-line segment that does not cross any edge of f. Let Γ' be the drawing obtained from Γ as follows: if $y(u) = y(v)$, we add a horizontal segment between u and v; if $y(u) < y(v)$, we add a polyline from v to u that consists of a horizontal segment, a vertex v' corresponding to a right turn, a vertical segment, a vertex u' corresponding to a left turn, and a horizontal segment; if $y(u) > y(v)$, we add a polyline from v to u that consists of a horizontal segment, a vertex v' corresponding to a left turn, a vertical segment, a vertex u' corresponding to a right turn, and a horizontal segment. Let H' be the corresponding orthogonal representation. We denote by f' and f'' the two faces of H' replacing f, and let f' be the face above edge (u, v). Note that f' is an internal face of H', regardless of f being an internal or an external face of H, and that c_u and c_v are both convex corners in f'. Clearly, $rotation_{f'}(c_v, c_u) = turn_{f'}(c_v) = 1$, since $turn_{f'}(c_{u'})$ and $turn_{f'}(c_{v'})$ (if u' and v' exist) cancel each other out. Thus, by Properties 4 and 5, $rotation_{f'}(c_u, c_v) = 4 - rotation_{f'}(c_v, c_u) = 3$. Since c_u is a reflex corner in f, we have $rotation_f(c_u, c_v) = rotation_{f'}(c_u, c_v) - turn_{f'}(c_u) + turn_f(c_u) = 3 - 1 + (-1) = 1$ □

Lemma 9 *Let f be a face of a turn-regular orthogonal representation H. Let c_u be a NE-corner (SW-corner) of f, associated with vertex u, and let c_v be a NW-corner (SE-corner) of f, associated with vertex v, such that c_u and c_v are mutually visible; then, there exists a path from v to u (from u to v) in H_x. Similarly, let c_u be a NW-corner (SE-corner) of f and let c_v be a SW-corner (NE-corner) of f, such that c_u and c_v are mutually visible; then, there exists a path from v to u (from u to v) in H_y.*

Proof: We prove the claim for a NE-corner c_u and a NW-corner c_v. The proofs for the other three cases are similar. Let $f(u, v)$ be the portion of the boundary of f from u to v, and let $\{w_0 = u, w_1, \ldots, w_k, w_{k+1} = v\}$ be the sequence of vertices of $f(u, v)$. We assume that $f(u, v)$ does not contain vertices associated with flat corners, since their presence is irrelevant to this proof. Let c_i be the (only) corner associated with vertex w_i. From the definitions of NE- and NW-corners, it follows that $next(u)$ and $prev(v)$ are vertical edges; hence, $f(u, v)$ contains at least four vertices.

We first show that u cannot be followed by more than two consecutive left turns in $f(u, v)$. Suppose, for a contradiction, that vertices w_1, w_2, and w_3 all correspond to left turns (see Fig. 14(a)). Since v corresponds to a right turn, $w_3 \neq v$, and thus $k > 3$. We have $rotation(c_u, c_4) = 2$. Since, by Lemma 8, $rotation(c_u, c_v) = 1$, and since the only corners with a negative value of $turn$ (exactly -1) are reflex corners, it follows, by Property 5, that there is a vertex w_i (possibly coincident with w_4 itself) in $f(w_4, v)$, whose associated corner is reflex and such that $rotation(c_u, c_i) = 2$. Thus H is not turn-regular; a contradiction.

We now consider the possible cases for w_1, w_2, w_3, and w_4, discard the impossible ones, and prove the claim by induction on the number of vertices of $f(u, v)$. Vertex w_1 cannot correspond to a right turn, since otherwise c_1 and c_v would be kitty corners and H would not be turn-regular. Hence, w_1 corresponds to a left turn. If w_2 corresponds to a left turn as well, then, by the above discussion, w_3 must be a right turn. If $w_3 = v$ (see Fig. 14(b)), there clearly exists a path from v to u in H_x; this is the base case of the induction. If $w_3 \neq v$ (see Fig. 14(c)), then, since $turn(c_2) = 1$ and $turn(c_3) = -1$, we can remove w_2 and w_3 from the sequence of vertices of $f(u, v)$ without affecting the value of $rotation(c_u, c_v)$, and the claim is proved by the induction hypothesis. Note that this process is similar to the switch collapsing process described and used in Section 3.3. Analogously, if w_2 corresponds to a right turn (see Fig. 14(d)), then, since $turn(c_1) = 1$ and $turn(c_2) = -1$, we can remove w_1 and w_2 from the sequence of vertices of $f(u, v)$ without affecting the value of $rotation(c_u, c_v)$, and the claim is proved by the induction hypothesis. □

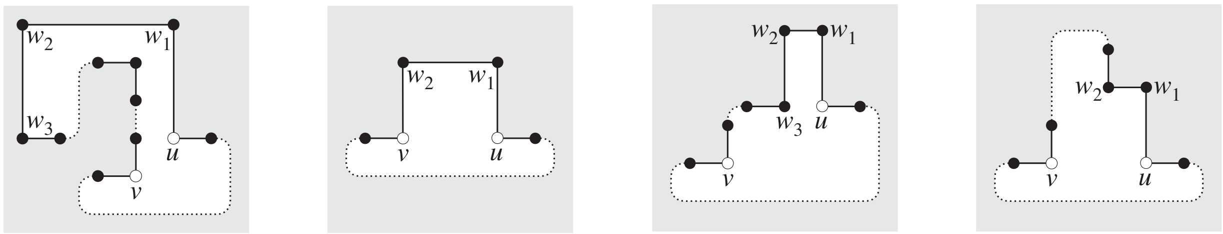

(a) (b) (c) (d)

Figure 14: Four cases in the proof of Lemma 9.

Lemma 10 *Let H be a turn-regular orthogonal representation, and let u and v be two vertices of H. If there is no path between u and v in H_x (H_y), then there is a path between u and v in H_y (H_x).*

Proof: We consider the case in which there is no path between u and v in H_x; the other case is similar. Let s_r and t_r be the source and the sink, respectively, of the complete saturator of G_r, and let s_ℓ and t_ℓ be the source and the sink, respectively, of the complete saturator of G_ℓ. In order to simplify the proof, instead of H_x and H_y, we consider the partially-directed graphs H'_x and H'_y constructed in the same way as H_x and H_y, but using all the saturating edges from G_r and G_ℓ (except (s_r, t_r) and (s_ℓ, t_ℓ)). Note that the existence of a path between u and v, $u, v \notin \{s_r, t_r, s_\ell, t_\ell\}$, in H'_x (H'_y) implies, by Lemma 5, the existence of a path between u and v, in H_x (H_y).

For each vertex w of H, we define four paths in H'_y, denoted $p_{\mathrm{NE}}(w)$, $p_{\mathrm{NW}}(w)$, $p_{\mathrm{SW}}(w)$, and $p_{\mathrm{SE}}(w)$. Informally speaking, they are paths in H'_y from w to the external face, going in the North-East, North-West, South-West, and South-East direction, respectively. In the rest of the proof, the subpath of path $p_{\mathrm{NE}}(w)$ ($p_{\mathrm{NW}}(w)$) from w to z is denoted by $p_{\mathrm{NE}}(w, z)$ ($p_{\mathrm{NW}}(w, z)$), and the subpath of path $p_{\mathrm{SW}}(w)$ ($p_{\mathrm{SE}}(w)$) from z to w is denoted by $p_{\mathrm{SW}}(z, w)$ ($p_{\mathrm{SE}}(z, w)$). The operator $+$ is used to denote the concatenation of vertices, edges, and subpaths of a given path. The formal definitions of all four paths are given for completeness, although they are very similar.

Path $p_{\mathrm{NE}}(w)$ is a path in H'_y from w to t_r, and is recursively defined as follows: (i) if $w = t_r$, then $p_{\mathrm{NE}}(w) = w$; (ii) otherwise, since the only unsaturated sink switch in G_r is t_r, there exists a vertex z such that (w, z) is either a saturating edge from G_r, or a (directed) vertical edge, or a horizontal edge with z as right end-vertex (tested in this order), and $p_{\mathrm{NE}}(w) = w + (w, z) + p_{\mathrm{NE}}(z)$. Note that $p_{\mathrm{NE}}(w)$ is non-decreasing in both the x- and y-coordinate. Similarly, path $p_{\mathrm{NW}}(w)$ is a path in H'_y from w to t_ℓ, and is recursively defined as follows: (i) if $w = t_\ell$, then $p_{\mathrm{NW}}(w) = w$; (ii) otherwise, since the only unsaturated sink switch in G_ℓ is t_ℓ, there exists a vertex z such that (w, z) is either a saturating edge from G_ℓ, or a (directed) vertical edge, or a horizontal edge with z as left end-vertex (tested in this order), and $p_{\mathrm{NW}}(w) = w + (w, z) + p_{\mathrm{NW}}(z)$. Note that $p_{\mathrm{NW}}(w)$ is non-increasing in the x-coordinate and non-decreasing in the y-coordinate.

Path $p_{\mathrm{SW}}(w)$ is a path in H'_y from s_r to w, and is recursively defined as follows: (i) if $w = s_r$, then $p_{\mathrm{SW}}(w) = w$; (ii) otherwise, since the only unsaturated source switch in G_r is s_r, there exists a vertex z such that (z, w) is either a saturating edge from G_r, or a (directed) vertical edge, or a horizontal edge with z as left end-vertex (tested in this order), and $p_{\mathrm{SW}}(w) = p_{\mathrm{SW}}(z) + (z, w) + w$. Note that $p_{\mathrm{NE}}(w)$ is non-decreasing in both the x- and y-coordinate. Similarly, path $p_{\mathrm{SE}}(w)$ is a path in H'_y from s_ℓ to w, and is recursively defined as follows: (i) if $w = s_\ell$, then $p_{\mathrm{SE}}(w) = w$; (ii) otherwise, since the only unsaturated source switch in G_ℓ is s_ℓ, there exists a vertex z such that (z, w) is either a saturating edge from G_ℓ, or a (directed) vertical edge, or a horizontal edge with z as right end-vertex (tested in this order), and $p_{\mathrm{SE}}(w) = p_{\mathrm{SE}}(z) + (z, w) + w$. Note that $p_{\mathrm{SE}}(w)$ is non-increasing in the x-coordinate and non-decreasing in the y-coordinate.

Let w_{NE} (w_{NW}) be the next-to-last vertex of $p_{\mathrm{NE}}(w)$ ($p_{\mathrm{NW}}(w)$), and let w_{SW} (w_{SE}) be the second vertex of $p_{\mathrm{SW}}(w)$ ($p_{\mathrm{SE}}(w)$). Note that w_{NE}, w_{NW}, w_{SW}, and w_{SE} are all vertices of the external face of H. For each vertex w of H, paths $p_{\mathrm{NE}}(w)$, $p_{\mathrm{NW}}(w)$, $p_{\mathrm{SW}}(w)$, and $p_{\mathrm{SE}}(w)$ define four regions in H'_y (see Fig. 15): $R_t(w)$ is the subgraph of H'_y with external face formed by $p_{\mathrm{NE}}(w)$, $p_{\mathrm{NW}}(w)$, and the portion of external face of H between w_{NE} and w_{NW}; $R_l(w)$ is the subgraph of H'_y with external face formed by $p_{\mathrm{NW}}(w)$, $p_{\mathrm{SW}}(w)$, and the portion of external face of H between w_{NW} and w_{SW}; $R_b(w)$ is the subgraph of H'_y with external face formed by $p_{\mathrm{SW}}(w)$, $p_{\mathrm{SE}}(w)$, and the portion of external face of H between w_{SW} and w_{SE}; $R_r(w)$ is the subgraph of H'_y with external face formed by $p_{\mathrm{SE}}(w)$, $p_{\mathrm{NE}}(w)$, and the portion of external face of H between w_{SE} and w_{NE}.

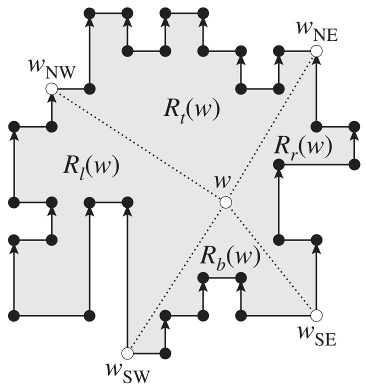

Figure 15: The four regions of vertex w.

We now show that if $u \in R_b(v)$, then there is a path from u to v in H'_y. We consider paths $p_{\mathrm{NE}}(u)$, $p_{\mathrm{NW}}(u)$, $p_{\mathrm{SW}}(v)$, and $p_{\mathrm{SE}}(v)$. By the definition of these paths, at least one of the following six cases applies: (i) $p_{\mathrm{NE}}(u)$ and $p_{\mathrm{SE}}(v)$ have a vertex q in common; then, $p_{\mathrm{NE}}(u, q) + p_{\mathrm{SE}}(q, v)$ is a path from u to v in H'_y. (ii) $p_{\mathrm{NE}}(u)$ and $p_{\mathrm{SE}}(v)$ cross, that is there is a saturating edge (w, x) from G_r in $p_{\mathrm{NE}}(u)$ and a saturating edge (y, z) from G_ℓ in $p_{\mathrm{SE}}(u)$ that cross each other; thus, by Lemma 7, either $mhc(w) = mhc(y)$ or $mhc(x) = mhc(z)$; w.l.o.g., assume that the former holds, and let $mhc(w, x)$ be the portion of the common maximal horizontal chain from w to x; then, $p_{\mathrm{NE}}(u, w) + mhc(w, x) + p_{\mathrm{SE}}(x, v)$ is a path from u to v in H'_y. (iii) $u_{\mathrm{NE}} \in R_b(v)$; note that u_{NE} is a t_L-switch of the external face of G_r, since it is adjacent to t_r, and that v_{SE} is an s_L-switch of the external face of G_ℓ, since it is adjacent to s_ℓ. Hence, u_{NE} is a SW-corner and v_{SE} is a NW-corner of the external face of H, and they are clearly mutually visible. By Lemma 9, there exists a path from u_{NE} to v_{SE} in H'_y. Let $f(u_{\mathrm{NE}}, v_{\mathrm{SE}})$ be this path (note that it is a portion of the external face of H); then, $p_{\mathrm{NE}}(u, u_{\mathrm{NE}}) + f(u_{\mathrm{NE}}, v_{\mathrm{SE}}) + p_{\mathrm{SE}}(v_{\mathrm{SE}}, v)$ is a path from u to v in H'_y. Cases $(iv$–$vi)$ are similar to Cases $(i$–$iii)$, and involve $p_{\mathrm{NW}}(u)$ and $p_{\mathrm{SW}}(v)$.

If $u \in R_t(v)$, then there is a path from v to u in H'_y. This can be proved in a similar way by considering paths $p_{\mathrm{SW}}(u)$, $p_{\mathrm{SE}}(u)$, $p_{\mathrm{NE}}(v)$, and $p_{\mathrm{NW}}(v)$.

Finally, with the same technique, it is possible to prove that if $u \in R_l(v)$ ($u \in R_r(v)$), then there is a path from u to v (from v to u) in H'_x. □

5 Turn-Regularity and Orthogonal Relations

In this section we use graphs H_x and H_y to characterize all possible orthogonal relations in a turn-regular orthogonal representation. This leads to a characterization of turn-regular orthogonal representations in terms of orthogonal relations. We denote by $u \to v$ a directed path from vertex u to vertex v in H_x containing at least a horizontal edge or in H_y containing at least a vertical edge, and by $u \not\to v$ the absence of such a path from vertex u to vertex v.

Lemma 11 *If H_x or H_y contains a saturating edge (u,v) from G_r, then $u <_x v$ and $u <_y v$. If H_x or H_y contains a saturating edge (u,v) from G_ℓ, then $v <_x u$ and $u <_y v$.*

Proof: We prove that, if H_x contains a saturating edge (u,v) from G_r, then u is left of and below v in any planar drawing of H; the proofs for a saturating edge (u,v) from G_ℓ, and for H_y instead of H_x are similar. In particular, we show this for a saturating edge from a t_L-switch to a t_S-switch; the proof for a saturating edge from an s_S-switch to an s_L-switch is similar.

Let f be a face of G_r, u be a t_L-switch of f, and v be a t_S-switch of f. We begin by showing that u is left of v in any planar drawing of H. Since u is a t_L-switch of f in G_r, it cannot be one of the rightmost vertices of f in any planar drawing of H. On the other hand, v is one of the rightmost vertices of f in any planar drawing of H. Suppose, for a contradiction, that there exists a planar drawing Γ of H in which v is not one of the rightmost vertices of f. Then there exists a maximal vertical chain *mvc* in f, containing at least one edge, whose vertices are right of v in Γ; let x and y be the bottommost and topmost vertices of *mvc*, respectively. Two cases are possible: either *mvc* precedes v in the traversal of f starting at u (see Fig. 16(a)), or *mvc* follows v (see Fig. 16(b)). In the first case, there must be at least one right turn in the path from y to v; in the second case, there must be at least two right turns in the path from v to x. In both cases, there exists a vertex w corresponding to one of these right turns such that its associated corner and the corner associated with u are kitty corners, thus violating the turn-regularity of H; a contradiction.

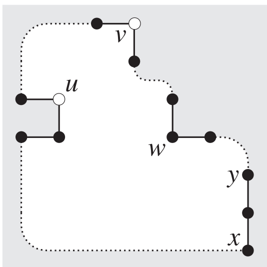

(a)

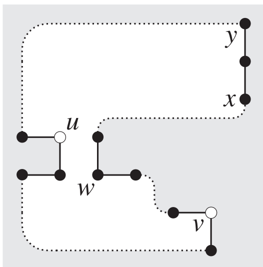

(b)

Figure 16: The two cases in the proof of Lemma 11.

Thus, we have proved that u is left of v in any planar drawing of H. The proof that u is below v in any planar drawing of H is similar, hence the thesis. □

Lemma 12 *Let Γ be a planar drawing of a turn-regular orthogonal representation H, and let mvc_1 and mvc_2 be two maximal vertical chains of Γ (possibly consisting of a single vertex) such that mvc_1 is to the left of mvc_2. If the endpoint of one chain can be connected to any point of the other by a horizontal segment that does not cross any other vertical chain of Γ, then there exists a path in H_x connecting any vertex of mvc_1 to any vertex of mvc_2.*

Proof: We consider only the case in which the bottom endpoint of mvc_1 can be connected to mvc_2; the other cases are similar. In this proof, we call two such chains *consecutive*. Clearly, if there is a path from some vertex of mvc_1 to some vertex of mvc_2 in H_x, there is a path from any vertex of mvc_1 to any vertex of mvc_2, because vertical edges of H_x are undirected. Let u be the bottom end-vertex of mvc_1; u and the object (either an edge or a vertex) on mvc_2 at the same y-coordinate of u belong to the same face f of H. Vertex u may correspond to a left turn, a right turn, or a U-turn of f; we consider the three cases separately.

If u corresponds to a left turn of f, then the first turn following u along the boundary of f corresponds to a vertex v of mvc_2, since otherwise the horizontal segment connecting u to mvc_2 would cross a vertical chain of Γ different from mvc_1 and mvc_2. The path from u to v in H_x is the portion of the boundary of f from u to v.

If u corresponds to a right turn of f, let c_u be the corner associated with u, and let vc_2 be the portion of mvc_2 that is part of the boundary of f. We first observe that vc_2 cannot be a single, degree one vertex, namely vertex v, since c_u and the second corner associated with v would be kitty corners, and H would not be turn-regular. Thus, let v be the bottom end-vertex of vc_2. Two cases are possible: (i) v corresponds to a right turn (see Fig. 17(a)) or a U-turn of f (see Fig. 17(b)). Then the corner c_v associated with v (the second one if v corresponds to a U-turn) is a NE-corner, and since c_u is a NW-corner, and c_u and c_v are mutually visible, there exists a path from u to v in H_x by Lemma 9. (ii) v corresponds to a left turn of f (see Fig. 17(c)). Then v is an s_S-switch in G_ℓ, and since u is an s_L-switch in G_ℓ, there exists a saturating edge (v, u) in G_ℓ (recall that, by Theorems 2 and 3, if H is turn regular, then G_ℓ has a unique complete saturator). If the saturating edge (v, u), with its orientation reversed, has been used in the construction of H_x, then the path from u to v in H_x is (v, u) itself; otherwise, such path exists by Lemma 5.

If u corresponds to a U-turn of f, the proof is analogous to that of the previous case, once we let c_u be the first corner associated with u, if the only edge incident with u is vertical, or the second corner associated with u, otherwise. □

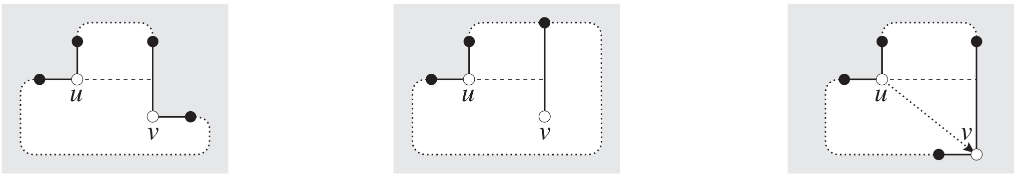

(a) (b) (c)

Figure 17: Three cases in the proof of Lemma 12. The horizontal segment connecting u to mvc_2 is represented as a dashed segment.

Lemma 13 *For each pair $\{u, v\}$ of vertices of H_x the following conditions hold:*

1. *$mvc(u) = mvc(v)$ if and only if $u =_x v$*
2. *$u \to v$ if and only if $u <_x v$*
3. *$v \to u$ if and only if $v <_x u$*
4. *$mvc(u) \neq mvc(v)$, $u \not\to v$, and $v \not\to u$ if and only if no x-relation can be established between u and v.*

Proof: *Only if.* We first prove that Conditions 1–4 are necessary.

Condition 1. If u and v belong to the same maximal vertical chain, they are clearly drawn with the same x-coordinate in any planar drawing of H.

Condition 2. We prove that $u <_x v$ by induction on the number of edges of the path. If $u \to v$ consists of only one edge $e = (u, v)$, then, since u and v do not belong to the same maximal vertical chain, there are three possible cases: (i) e is a horizontal edge of H; then, v is clearly right of u in any planar drawing of H. (ii) e is a saturating edge from G_r; then, by Lemma 11, u is left of v in any planar drawing of H; (iii) e is a saturating edge from G_ℓ with its orientation reversed; then, (v, u) is a saturating edge of G_ℓ and, again by Lemma 11, u is left of v in any planar drawing of H.

We now suppose that $u <_x v$ for each path $u \to v$ consisting of $k-1$ edges, $k > 2$, and prove the claim for a path $u \to v$ consisting of k edges. Let (w, v) be the last edge of of the path. Two cases are possible: (i) w belongs to the same maximal vertical chain of v; then, by Condition 1, we have $w =_x v$, and, by the inductive hypothesis, we have $u <_x w$, which implies $u <_x v$; (ii) w does not belong to the same maximal vertical chain of v; then, by the inductive hypothesis, we have $w <_x v$ and $u <_x w$, which implies $u <_x v$.

Condition 3. Analogous to the proof of Condition 2.

Condition 4. Let Γ be a planar orthogonal drawing of H. W.l.o.g., we assume that $x(u) \leq x(v)$ and $y(u) > y(v)$ in Γ. We show how to construct from Γ a different planar drawing of H such that $x(u) > x(v)$, thus showing that no x-relation can be established between u and v.

We consider the two vertical lines $x_\ell = x(u) - \frac{1}{2}$ and $x_r = x(v) + \frac{1}{2}$. We denote by Γ_c the portion of Γ between x_ℓ and x_r. We define a *moving-line* in Γ with the following properties:

- it is directed and orthogonal,
- it starts at the intersection u' between x_ℓ and a horizontal line y_u through u and ends at the intersection v' between x_r and a horizontal line y_v through v,
- it is entirely contained within the box B defined by x_ℓ, x_r, y_u, and y_v,
- it does not intersect any vertical edges of Γ_c, and
- its first and last segments are vertical.

Using compaction techniques developed for VLSI layout [12, 21, 25], the moving-line guarantees that it is possible to stretch and shift parts of Γ to obtain a planar orthogonal drawing such that $x(u) > x(v)$.

We construct a moving-line J such that all of its bends are displaced a half-unit from grid-points. This implies that the only intersections between J and Γ are between vertical segments of J and horizontal segments of Γ, and that the first and last segments of J are vertical, as required.

J is constructed as follows, where the possibility or impossibility of traveling in a certain direction is given by the above properties. Travel downward from u' until it is possible to travel right. Travel right as far as possible until either x_r is reached or a vertical segment of Γ is half-unit distant. If x_r is half-unit distant, travel downward to v', and the construction of J is complete. Otherwise, travel either downward or upward until it is again possible to travel right. Choose the upward direction only if it is not possible to travel right before reaching y_v when going down. If it is not possible to travel rightward by going either up or down (staying within B), there is no moving-line. Otherwise, continue the process until v' is reached. Observe that this method will construct an x-monotone moving-line if any moving-line exists, and that it will construct a y-monotone moving-line if there is any possible y-monotone moving-line.

If there is no moving-line, then there must be a vertical chain vc of Γ that intersects both y_u and y_v. Since both $mvc(u)$ and vc intersect y_u, either Lemma 12 applies and there is a path from u to any vertex of vc, or there is another vertical chain vc' between $mvc(u)$ and vc that also crosses y_u. The argument can then be applied to $mvc(u)$ and vc', and to vc' and vc, yielding a path from u to vc. A similar argument can be applied to vc and $mvc(v)$ because both chains intersect y_v, yielding a path from vc to v, and thus a path from u to v.

We now show that if J is not y-monotone, there is a path from u to v in H_x. Consider the sequence of vertical chains $C = \{vc_1, \ldots, vc_k\}$ that define J: $vc_1 = mvc(u)$, $vc_k = mvc(v)$, and the others are those that determine upward or downward turns of J, ordered from left to right. Let vc_i and vc_{i+1} be two (consecutive) elements of C. Then there is a horizontal segment connecting one endpoint of vc_i to vc_{i+1} that either crosses no other vertical chains of Γ or touches only one of their endpoints. By (possibly repeated) application of Lemma 12, there is a path from any vertex of vc_i to any vertex of vc_{i+1}. If J is not y-monotone, there is at least one chain in C that defines an upward turn in J; let vc_j be the first such chain. This chain vc_j must cross y_v, or else it would have been possible to make a downward turn instead. By repeatedly applying the previous argument, there is a path from $vc_1 = mvc(u)$ to vc_j. We now must show that there is a path from vc_j to $vc_k = mvc(v)$. Since both vc_j and vc_k intersect y_v, the argument used in the previous paragraph applies, resulting in a path from vc_j to vc_k, and thus a path from u to v.

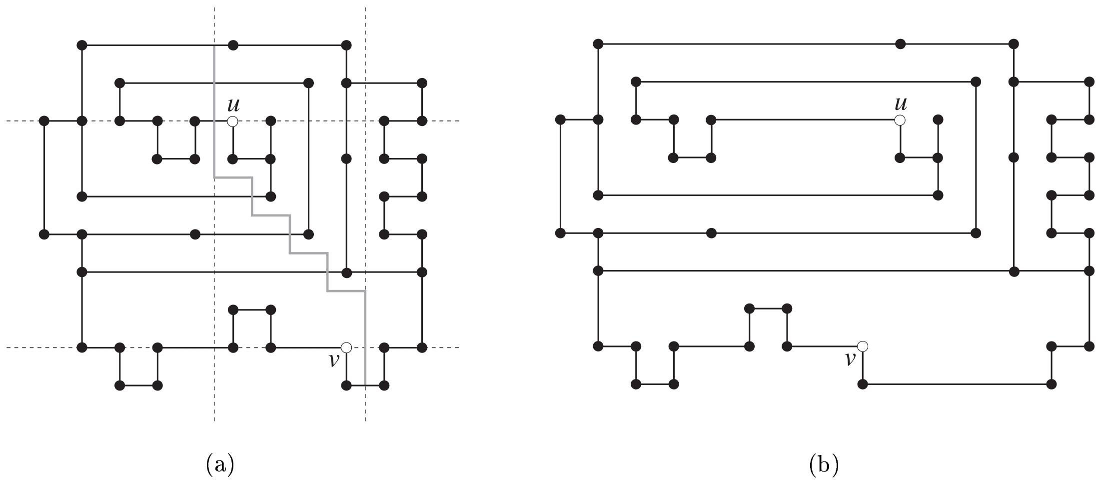

Figure 18: (a) A planar orthogonal drawing Γ with the extended moving-line for vertices u and v. (b) The stretched version of Γ.

As a result, there exists a moving-line J that is monotone in both x and y. This moving-line can then be used to stretch Γ horizontally so that $x(u) > x(v)$. Extend J by continuing its first vertical segment (i.e., the one that starts at u') in the positive y direction and its last vertical segment (i.e., the one that ends at v') in the negative y direction until they both intersect the external face of Γ (see Fig. 18(a)). Also, let $d_x = x(v) - x(u)$. Stretch Γ by increasing the x-coordinates of the vertices to the right of J by an amount $d > d_x$. As a result, a new planar orthogonal drawing is constructed in which $x(u) > x(v)$ (see Fig. 18(b)). This implies that there is no x-relation between u and v.

If. The sufficiency of Conditions 1–4 follows from the completeness and mutual exclusiveness of the four cases. □

As an example, we identify in the graph H_x shown in Fig. 10(a) three pairs of vertices corresponding to the various cases of Lemma 13: u_1 and v_1 belong to the same maximal vertical chain, there exists a directed path from u_2 to v_2, while there is neither a path between u_3 and v_3, nor they belong to the same maximal vertical chain.

The proof of the following lemma is similar to that of Lemma 13, and hence is omitted.

Lemma 14 *For each pair $\{u, v\}$ of vertices of H_y the following conditions hold:*

1. $mhc(u) = mhc(v)$ *if and only if* $u =_y v$
2. $u \to v$ *if and only if* $u <_y v$
3. $v \to u$ *if and only if* $v <_y u$
4. $mhc(u) \neq mhc(v)$, $u \not\to v$, *and* $v \not\to u$ *if and only if no y-relation can be established between u and v.*

As an example, we identify in the graph H_y shown in Fig. 10(b) three pairs of vertices corresponding to the various cases of Lemma 13: u_1 and v_2 belong to the same maximal horizontal chain, there exists a directed path from u_2 to v_3, while there is neither a directed path between u_3 and v_1, nor they belong to the same maximal vertical chain.

Lemma 15 *Let G be an embedded 4-planar graph and H be a turn-regular orthogonal representation of G. For each pair $\{u, v\}$ of vertices of G, exactly one orthogonal relation holds.*

Proof: By considering all possible combinations of the four (mutually exclusive) cases of Lemma 13 for H_x with the four (mutually exclusive) cases of Lemma 14 for H_y, we obtain the sixteen (mutually exclusive) cases shown in Table 1. For twelve of these cases, the orthogonal relations are simply obtained by taking the "logical and" of the corresponding relations given in Lemmas 13 and 14.

We now show that the remaining four cases are impossible. Clearly, if $mvc(u) = mvc(v)$ in H_x, then there is a path $u \to v$ or $v \to u$ in H_y, and if $mhc(u) = mhc(v)$ in H_y, then there is a path $u \to v$ or $v \to u$ in H_x. As for the "otherwise–otherwise" case, it is impossible by Lemma 10. □

We are interested in those orthogonal representations for which there is an orthogonal relation between each pair of vertices. As the following theorem shows, this class of orthogonal representations is characterized by turn-regularity.

Theorem 5 *An orthogonal representation H is turn-regular if and only if there is an orthogonal relation between every two vertices of H.*

		H_y			
		$mhc(u) = mhc(v)$	$u \to v$	$v \to u$	otherwise
H_x	$mvc(u) = mvc(v)$	impossible	$(u =_x v) \wedge (u <_y v)$ 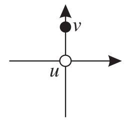	$(u =_x v) \wedge (v <_y u)$ 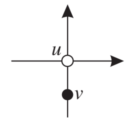	impossible
	$u \to v$	$(u <_x v) \wedge (u =_y v)$ 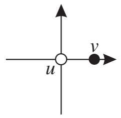	$(u <_x v) \wedge (u <_y v)$ 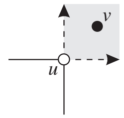	$(u <_x v) \wedge (v <_y u)$ 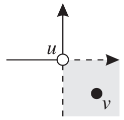	$u <_x v$ 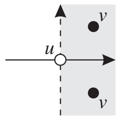
	$v \to u$	$(v <_x u) \wedge (u =_y v)$ 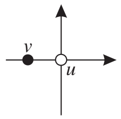	$(v <_x u) \wedge (u <_y v)$ 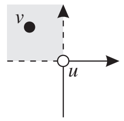	$(v <_x u) \wedge (v <_y u)$ 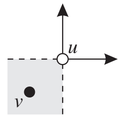	$v <_x u$ 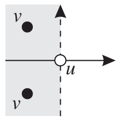
	otherwise	impossible	$u <_y v$ 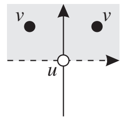	$v <_y u$ 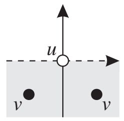	impossible

Table 1: Orthogonal relations for a pair $\{u, v\}$ of vertices in H.

Proof: *Only if.* It follows directly from Lemma 15.

If. Suppose, for a contradiction, that there is an orthogonal relation between every two vertices of H and that H is not turn-regular. Hence, there exists a face f of H with two vertices u and v, whose associated corners c_u and c_v are kitty corners. Assume, w.l.o.g., that the kitty corners form a SW-NE pair (see Fig. 8(a)). We generate two new orthogonal representations H_1 and H_2 as follows. H_1 is obtained by connecting u and v with a horizontal edge such that u is the left end-vertex of the new edge and v is the right end-vertex. H_2 is obtained by connecting u and v with a vertical edge such that u is the bottom end-vertex and v is the top end-vertex. Note that adding these edges is always possible because u is not the left or bottom end-vertex of an edge in H, and v is not the right or top end-vertex of an edge in H.

H_1 and H_2 are orthogonal representations, since they satisfy Properties 1 and 2. We show this for an internal face of H_1; the proofs for the external face of H_1 and for H_2 are similar. We assume that f contains no vertex of degree one and no multiple occurrences of the same vertex; if not, we can use the same expansion technique adopted for the proofs in Section 3.3. Note that the only vertices affected by the insertion of the new edge are u and v, and that the only face is f. Let f' and f'' be the two faces of H_1 replacing f, and let f' be the face below edge (u, v) and f'' be the face above. The angle at u in f' is equal to $\pi/2$, while the angle at u in f'' is equal to π. Similarly, the angle at v in f' is equal to π, while the angle at v in f'' is equal to $\pi/2$. Property 1 is clearly satisfied by u and v, since the angles at u and v in f are both $3\pi/2$. As for Property 2, we prove that it is satisfied by f'; the proof for f'' is analogous. Let p be the portion of f from u (excluded) to v (excluded), n_l be the number of left turns in p, n_f be the number of flat turns in p, and n_r be the number of right turns in p. Note that the vertices of f in p are also vertices of f', and that, in addition, f' contains u and v. Thus, Property 2 is satisfied by f' if

$$\begin{aligned}
n_l \cdot \pi/2 + n_f \cdot \pi + n_r \cdot 3\pi/2 + \pi/2 + \pi &= [2(n_l + n_f + n_r + 2) - 4] \cdot \pi/2 \\
(n_l + 2n_f + 3n_r + 1 + 2) \cdot \pi/2 &= (2n_l + 2n_f + 2n_r) \cdot \pi/2 \\
n_l + 3n_r + 3 &= 2n_l + 2n_r \\
n_l - n_r &= 3
\end{aligned}$$

And this is indeed the case. Let w be the first vertex of f after u, and let c_w be its associated corner. Since c_u and c_v are kitty corners in f, $rotation_f(c_u, c_v) = 2$, and thus $rotation_{f'}(c_w, c_v) = rotation_f(c_w, c_v) = 3$. That is, from the definition of *rotation*, the number of left turns in p minus the number of right turns in p is equal to three.

Since the edge connecting u and v in H_1 is horizontal, it follows that u must be left of and y-aligned with v in any planar drawing of H_1. Similarly, u must be below and x-aligned with v in any planar drawing of H_2. Also, note that a planar drawing of H can be obtained from any planar drawing of H_1 or H_2 by simply removing the extra edge. As a result, there is not an orthogonal relation between u and v; a contradiction. □

6 Turn-Regularity and Drawing Algorithms

In this section we first study the problem of efficiently checking whether an orthogonal representation is turn-regular. Then we show how an optimal area orthogonal drawing of a turn-regular orthogonal representation can be computed.

Theorem 6 *A turn-regular orthogonal representation with n vertices and bends can be recognized in $O(n)$ time and space.*

Proof: Let H be an orthogonal representation and let f be a face of H with n_f vertices. We show how to test whether f is turn-regular in $O(n_f)$ time and space. Since, for a planar graph, $\sum_f n_f = O(n)$, this proves the claim.

We first consider the case in which f is an internal face. We index the k reflex corners of f from c_1 to c_k, according to a counterclockwise visit of the boundary of f from an arbitrary vertex. We construct a circular list L of k elements, and we set its i-th element $L(i) = \mathit{rotation}(c_1, c_{i+1})$, where, by convention, $c_{k+1} = c_1$. Since $k = O(n_f)$, the size of L is linear with the number of vertices of f. Also, for each $1 \leq i \leq k$ we have $L(i) = O(n_f)$, because each vertex of f is associated with one or two corners, and for each corner c we have $-1 \leq \mathit{turn}(c) \leq 1$. Clearly, the construction of L requires linear time. By Property 6, in order to test the turn-regularity of f, we must verify whether there exist two indices $1 \leq i < j \leq k$ such that $L(j) - L(i) = \mathit{rotation}(c_1, c_{j+1}) - \mathit{rotation}(c_1, c_{i+1}) = \mathit{rotation}(c_{i+1}, c_{j+1}) = 2$.

Let min and max be the minimum and maximum values stored in L, respectively. We construct an array A of $\mathit{max} - \mathit{min} + 3$ boolean elements whose index is in the range $\mathit{min} - 2 \ \ldots \ \mathit{max}$. We denote the i-th element of A by $A(i)$ and we set it equal to true if there exists in L an element whose value is i, equal to false otherwise.

The algorithm to test whether f is turn-regular consists of the following test for each element j of L. Let ρ be the value of $L(j)$; if $A(\rho - 2)$ is true, then there exists an element $i < j$ of L such that $L(i) = \rho - 2$. Hence, $L(j) - L(i) = 2$, c_{i+1} and c_{j+1} are kitty corners, and f is not turn-regular. If for each element of L the result of the above test is false, then f is turn-regular.

We now consider the case in which f is the external face. By Property 6, in order to test the turn-regularity of f, we must verify whether there exist two indices $1 \leq i < j \leq k$ such that either $L(j) - L(i) = 2$ or $L(j) - L(i) = -6$. Thus, the array A consists of $\mathit{max} - \mathit{min} + 9$ boolean elements, its index is in the range $\mathit{min} - 2 \ \ldots \ \mathit{max} + 6$, and, in addition to $A(\rho - 2)$, we also test whether $A(\rho + 6)$ is true.

Since verifying the value of an element of A requires constant time, and the number of elements of L is $O(n_f)$, the overall procedure requires linear time. □

The optimal area drawings that we want to compute are planar. The next lemma guarantees the planarity of drawings that satisfy the orthogonal relations of a turn-regular orthogonal representation.

Theorem 7 *Let H be a turn-regular orthogonal representation, and let Γ be an orthogonal drawing of H such that, for each pair $\{u, v\}$ of vertices of H, the orthogonal relation between u and v is satisfied. Then Γ is planar.*

Proof: Since H is a turn-regular orthogonal representation, by Lemma 15 for each pair of vertices of H exactly one orthogonal relation is satisfied in all possible planar drawings of H. Suppose, for a contradiction, that there exists a drawing Γ of H such that:

- Γ is a non-planar drawing of H, and
- Γ satisfies all orthogonal relations defined by the turn-regularity of H.

Let (u, v) and (w, z) be two edges of Γ that cross each other. We assume that (u, v) is a vertical edge with u below v, and that (w, z) is a horizontal edge with w left of z. The proof for the case of overlapping vertices or edges is similar. Since Γ satisfies all orthogonal relations defined by the turn-regularity of H and, by Lemma 15, for each pair of vertices of H exactly one orthogonal relation is satisfied in all possible planar drawings of H, we conclude that the following orthogonal relations hold: (*i*) $u <_x z$, (*ii*) $w <_x u$, (*iii*) $z <_y v$, and (*iv*) $u <_y w$.

Since H is a planar orthogonal representation, there exists at least one planar drawing Γ' of H. Since in Γ' the pair of edges (u, v) and (w, z) do not cross, the four orthogonal relations given above cannot be simultaneously satisfied, contradicting the fact that in a turn-regular orthogonal representation exactly one orthogonal relation holds between any two vertices. □

We are now ready to present two different algorithms that compute optimal area drawings of turn-regular orthogonal representations. These algorithms are variations of the two compaction procedures described in [6]. For the first algorithm, we define two digraphs, denoted D_x and D_y. D_x is obtained from H_x by shrinking each maximal vertical chain to a single vertex, by removing possible multiple edges, and by adding a super-source and a super-sink (see Fig 19(a)). Thus, there is a one-to-one correspondence between maximal vertical chains of H_x and vertices of D_x, and a many-to-one correspondence between directed edge of H_x and edges of D_x. Note that in the shrinking process we "preserve the embedding," i.e., the circular ordering of the edges around each vertex v of D_x is induced by the circular ordering of the directed edges "around" the maximal vertical chain of H_x corresponding to v. D_y is obtained analogously from H_y by shrinking the maximal horizontal chains (see Fig 19(b)).

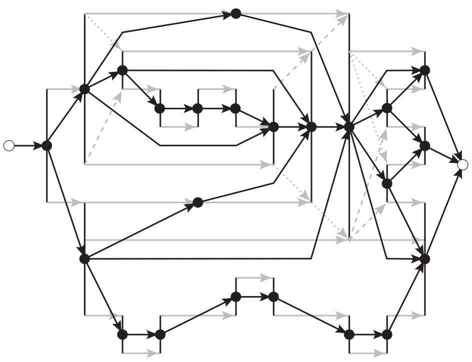

(a)

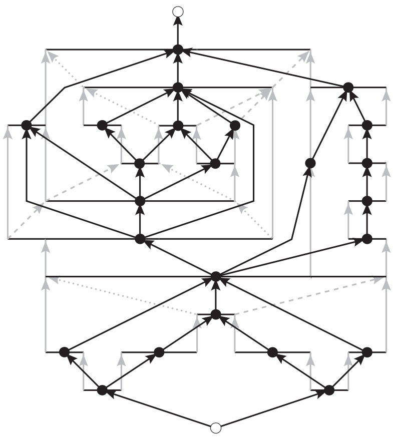

(b)

Figure 19: (a) D_x. (b) D_y. Both graphs are obtained from graphs H_x and H_y shown in Fig. 10.

Property 10 *D_x and D_y are planar st-digraphs.*

Proof: We first prove that D_x and D_y are planar. In particular, we prove the planarity of D_y; the proof for D_x is analogous. Let D be the graph obtained from D_y by removing the edges that correspond to saturating edges in H_y. Note that, if we ignore the direction of the edges, D can be thought of as the result of the shrinking process applied to H. Since H is an embedded planar graph, D is an embedded planar graph as well. In particular, there is a one-to-one correspondence between internal faces of H and faces of D, and if the end-vertices of a saturating edge e in H_y belong to face f of H, the end-vertices of the directed edge of D_y corresponding to e belong to the face of D corresponding to f. Let (u, v) and (w, z) be two saturating edges in H_y that cross each other. We recall that u, v, w, and z belong to the same face f of H. Thus, from the discussion above, we can concentrate our attention on f alone. We prove that the crossing "disappears" during the shrinking process through which D_y is obtained. By Lemma 7, either $mhc(u) = mhc(w)$ or

$mhc(v) = mhc(z)$. It follows that either u and w, or v and z are shrunk to the same vertex in D_y and the crossing "disappears" in the process.

It remains to prove that D_x and D_y are st-digraphs. The non-vertical edges of H_x are all directed from left to right, and the non-horizontal edges of H_y are all directed from bottom to top. Hence, by construction, D_x and D_y are acyclic. And still by construction, they have a single source and a single sink. □

Theorem 8 *Let H be a turn-regular orthogonal representation with n vertices and bends. A planar orthogonal drawing of H with optimal area can be constructed in $O(n)$ time and space.*

Proof: We recall that there is a one-to-one correspondence between maximal vertical chains of H_x (and hence of H) and vertices of D_x; similarly, there is a one-to-one correspondence between maximal horizontal chains of H_y (and hence of H) and vertices of D_y.

We compute the x-coordinates of the vertical segments representing the maximal vertical chains of H as follows. We assign unit weights to the edges of D_x and compute an optimal weighted topological numbering X of D_x (see page 89 of [6]). We then set the length of each horizontal directed edge (u, v) of H equal to $X(v') - X(u')$, where u' and v' are the vertices of D_x representing the maximal vertical chains of H containing u and v, respectively. In the same way, it is possible to compute the y-coordinates of the horizontal segments representing the maximal horizontal chains of H. Let Γ be the resulting drawing.

Lemma 13 shows that H_x represents all the x-relations between vertices of H. This information is represented by D_x, as well, since there is a many-to-one correspondence between directed edge of H_x and edges of D_x. A similar argument holds for D_y. Since the edges of D_x and D_y are assigned positive weights, and since X and Y are weighted topological numberings of D_x and D_y, the x-relations and y-relations between every two vertices of H are satisfied in Γ. Thus, by Theorem 7, Γ is planar.

Since the edges of D_x and D_y are assigned unit weights, and since the weighted topological numbering X and Y are both optimal, the width and height of Γ are both minimum. Hence, Γ has optimal area.

The time complexity of the algorithm is $O(n)$ since the number of vertices of D_x and of D_y are both not greater than the number of vertices of H, and computing an optimal weighted topological numbering of an n-vertex planar st-digraph requires $O(n)$ time and space. □

Theorem 9 *Let H be a turn-regular orthogonal representation with n vertices and bends. A planar orthogonal drawing of H with optimal area $\mathcal{A}$ and whose total edge length is optimal among all drawings with area $\mathcal{A}$ can be constructed in $O(n^{7/4} \log n)$ time and $O(n)$ space.*

Proof: To compute a planar orthogonal drawing with minimum area $\mathcal{A}$ and whose total edge length is minimum among all drawings with area $\mathcal{A}$, we use a flow technique similar to that described in [6]. Let H be a turn-regular orthogonal representation with n vertices. From Theorem 4, we know that the embedded planar st-digraphs D_x and D_y are uniquely determined. As seen in the proof of Theorem 8, a planar orthogonal drawing of H with optimal area can be found by computing an optimal weighted topological numbering on D_x and D_y, independently. This means that the width w and the height h of the drawing can be minimized independently.

We construct two flow networks N_x and N_y associated with D_x and D_y, respectively, and show that the cost of any integer feasible flow with value w on N_x plus the cost of any integer feasible flow with value h on N_y is equal to the total edge length of a planar orthogonal drawing of H with width w and height h. Consider the embedded planar st-digraph D_y, and consider the external face of D_y split into two regions, the "left" external face and the "right" external face. We construct N_y as follows (see Fig. 20(b)):

- For each internal face f of D_y, we consider a node v_f in N_y.
- For the external face of D_y, we consider two nodes in N_y, one for the "left" external face and one for the "right" external face.
- For each edge e of D_y, let f_l be the face to the left of e and f_r be the face to the right of e. We consider a dual arc from v_{f_l} to v_{f_r} in N_y. The upper capacity of this arc is set to $+\infty$ and the lower capacity is set to 1. Finally, the cost of the arc is set to 1 if e corresponds to an edge of H and to 0 if e corresponds to a saturating edge.

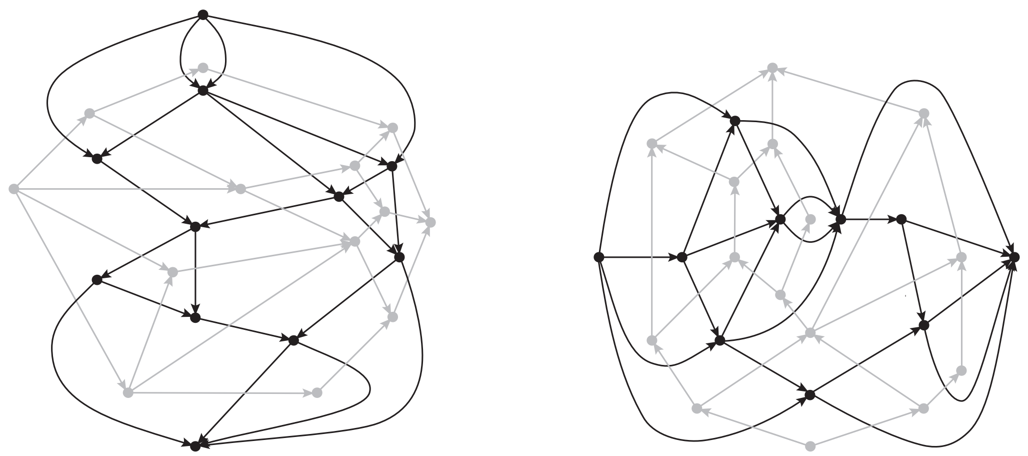

(a)

(b)

Figure 20: (a) N_x. (b) N_y. The corresponding planar st-digraphs D_x and D_y are represented in grey.

Since all arcs of N_y have an infinite upper capacity, there always exists an integer feasible flow with value h in N_y. Each unit of flow on an arc of N_y corresponds to a unit of length of the dual edge in D_y. In particular, the flow on each arc with cost 1 in N_y corresponds to the length of a vertical edge of H. Thus, computing a minimum cost flow with value h on N_y corresponds to minimizing the total length of the vertical edges of H in an orthogonal drawing with height h.

The construction of N_x is analogous to that of N_y, and computing a minimum cost flow with value w on N_x corresponds to minimizing the total length of the horizontal edges of H in an orthogonal drawing with width w.

In order to prove that the obtained drawing Γ is planar, we observe that D_x and D_y represent the x- and y-relations between every two vertices in any planar drawing of H. These relations are satisfied by the x- and y-coordinates obtained by computing the minimum cost flows on N_x and N_y, and thus, by Theorem 7, Γ is planar. In particular, we observe that:

- If there is a directed edge (u, v) in D_x (D_y), u will be right of (below) v in Γ, since the flow on the dual arc of (u, v) in N_x (N_y) is at least 1.

- The lengths of the edges of D_x and D_y are consistent. Recall that D_x and D_y are planar st-digraphs, and that the boundary of each face f of a planar st-digraph consists of two directed paths enclosing f, with common origin and destination. N_x and N_y are the dual planar st-digraphs of D_x and D_y; thus, for each face f of D_x (D_y), the incoming arcs of v_f in N_x (N_y) are duals of the edges of the "top" ("left") path of f, and the outgoing arcs of v_f in N_x (N_y) are duals of the edges of the "bottom" ("right") path of f. The consistency of the lengths of the edges of D_x and D_y then follows from the conservation property of the flow. And since there is a many-to-one correspondence between directed edge of H_x (H_y) and edges of D_x (D_y), also the lengths of the edges of Γ are consistent.

Hence, the minimum total edge length of a planar orthogonal drawing of H with optimal area $\mathcal{A} = w^* \cdot h^*$, is equal to the minimum cost of a flow with value w^* on N_x plus the minimum cost of a flow with value h^* on N_y. Each of the two minimum cost flows can be computed in $O(n^{7/4} \log n)$ time and $O(n)$ space as shown in [16]. □

We recall that the minimum number of bends for an orthogonal representation of a 4-planar graph with n vertices is $O(n)$ [3, 33]. The algorithm described in [31] produces such an orthogonal representation, and there exist various algorithms for producing an orthogonal representation with a sub-optimal, $O(n)$ number of bends (see, e.g, [4, 27, 32]).

7 Experiments

In this section, we present the results of an experimental study on a test suite of planar orthogonal representations of randomly generated biconnected 4-planar graphs. The analysis of the test suite has shown that the percentage of regular faces is quite high (95%). Motivated by this result, we have designed compaction heuristics based on the idea of "face regularization."

7.1 Compaction Heuristics

We have implemented a compaction algorithm for orthogonal representations based on the results described in the previous sections. Namely, let H be a given orthogonal representation of a 4-planar graph. The algorithm is as follows:

- H is first tested for turn-regularity, using the algorithm described in Theorem 6.
- If H is turn-regular, the algorithm computes an orthogonal drawing of H with optimal area and optimal total edge length within that area by applying the techniques in Theorem 9.
- If H contains some faces that are not turn-regular, an algorithm is applied to make these faces turn-regular. The algorithm adds dummy vertices and edges to H, creating a new orthogonal representation H' that is turn-regular. The techniques in Theorem 9 are then used to find a drawing Γ' of H' with optimal area and optimal total edge length within that area. Finally, the dummy vertices and edges are removed from Γ' to yield an orthogonal drawing Γ of H. In general the orthogonal drawing Γ does not have optimal area and total edge length.
- Two simple approaches are used to make non-regular faces turn-regular:
  1. The first approach is an improvement of the standard rectangularization method described in [6, 31]. When a dummy edge is inserted, a dummy vertex is added only if it really needed. Each non-regular face is divided into two or more smaller, rectangular faces.

2. The second approach recursively adds a straight edge (randomly chosen to be either horizontal or vertical) between each pair of kitty corners, until the face has been decomposed in smaller (but not necessarily rectangular) turn-regular faces. In general, this technique adds a much smaller number of dummy edges than the first approach.

In the following we call the two heuristic compaction algorithms derived from the two regularization approaches described above `Heur1` and `Heur2`, respectively. Our implementations of these algorithms use the *GDToolkit* library[4] and will be included in its next release.

7.2 Test Suite and Experimental Results

Heuristics `Heur1` and `Heur2` were tested on a set of 530 randomly generated biconnected 4-planar graphs with 10-3000 vertices. The results are compared to a third compaction heuristic, `StdComp`, in which all faces, both turn-regular and not, are decomposed into rectangles using the rectangularization method of `Heur1`.

The graphs in the test suite have been generated with a technique used in other experimental studies on orthogonal drawings [1]. Each biconnected 4-planar graph is generated from a cycle of three vertices by performing a random series of *InsertVertex* and *InsertEdge* operations. The *InsertVertex* operation subdivides an existing edge into two new edges separated by a new vertex. The *InsertEdge* operation inserts a new edge between two existing vertices on the same face. Any biconnected planar graph can be generated by a sequence of these two operations. Also, for each graph to be generated, the density of the graph (the number of edges divided by the number of vertices) is randomly chosen before the generation algorithm is run.

Our first experiment consisted in studying the percentage of non-regular faces in the graphs of our test suite. We have found that the percentage of non-regular faces decreases exponentially with the density of the graphs, i.e., with the ratio of the number of edges to the number of vertices (see Fig. 21).

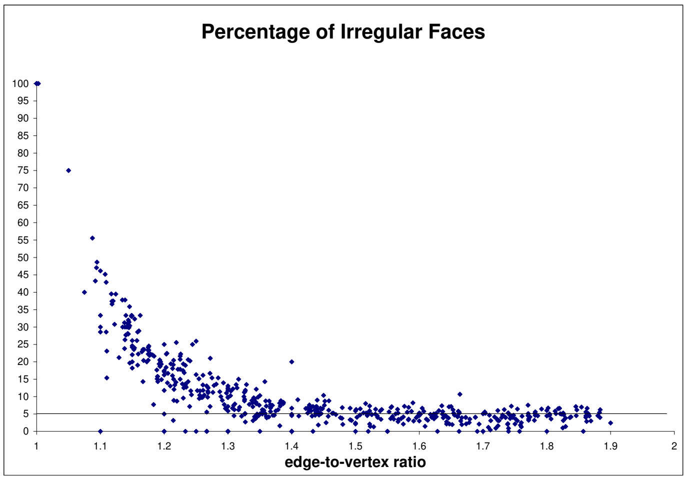

Figure 21: The percentage of non-regular faces with respect to the edge-to-vertex ratio.

We have then analyzed the results of the three heuristics. In particular, we have considered, for `Heur1` and `Heur2`, the improvement in the drawing area, total edge length, and maximum

[4] `http://www.dia.uniroma3.it/~gdt`

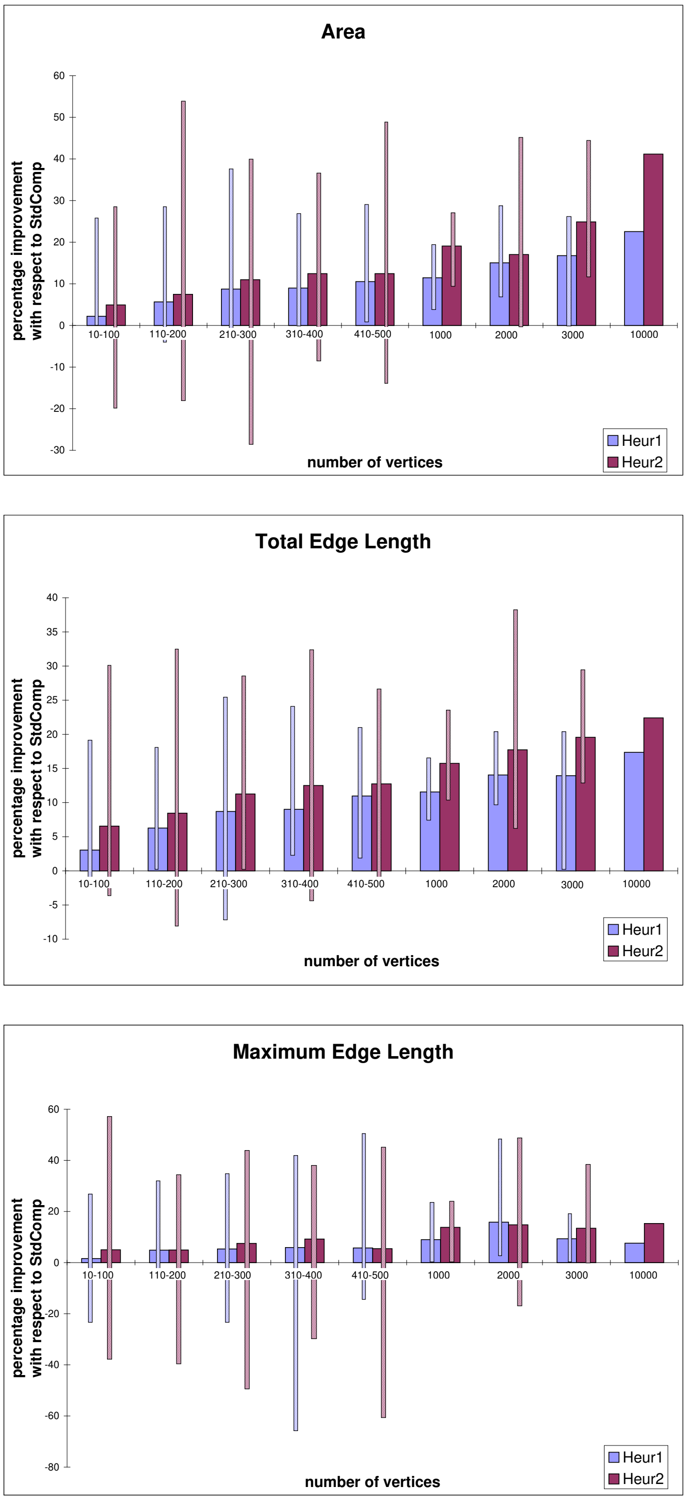

Figure 22: The average percentage improvement in area, total edge length, and maximum edge length of `Heur1` and `Heur2` with respect to `StdComp`. The narrow columns show the minimum and maximum improvement.

edge length with respect to `StdComp`. `Heur2` performs better than `Heur1` in most cases; also, the improvement in area and total edge length of `Heur2` with respect to `StdComp` increases with the number of vertices of the graph (see Fig. 22). The average improvement in area and total edge length is 19-25% for graphs with 3000 vertices, and there are some graphs in the test suite for which the improvement is more than 45%.

We have also run the three heuristics on a very large graph with $10,000$ vertices. The drawing computed by `Heur2` improves the area by 41% and the total edge length by 22.4% with respect to the drawing computed by `StdComp`.

8 Conclusions and Open Problems

We introduced the notion of turn-regularity which allows the characterization of a class of orthogonal representations that are optimally compactable in terms of area in polynomial-time. In particular, given a turn-regular orthogonal representation, we provided a linear-time algorithm to compute a planar drawing with minimum area, and a polynomial-time algorithm to compute a planar drawing with optimal area and minimum total edge length within that area of the given orthogonal representation.

This work bears some similarity to an ongoing project at the Max-Planck-Institute für Informatik [23, 24] in which the problem of minimizing the total edge length in planar drawings of orthogonal representations is considered by solving an associated ILP problem. Some overlaps are unavoidable since that research effort addresses a question in the same class of problems as the one that we study in this paper. However, the work differs significantly because the respective proof techniques are different, our approach is driven by a combinatorial analysis that does not lead to a formulation of the problem as an ILP problem, and the problem we study combines both area minimization and total edge length minimization.

We provided several implementations of heuristics for making orthogonal representations turn-regular and we used them in the compaction algorithm, instead of the standard rectangularization step. Experiments on a randomly generated test suite of biconnected 4-planar graphs show that the new compaction strategy works much better than the standard one, especially for very large graphs.

Some of the open problems related to this work are the following:

- To find other families of orthogonal representations for which an optimal area drawing can be computed in polynomial time.
- To investigate other effective heuristics for making an orthogonal representation turn-regular by adding a small number of edges.
- The problem of computing an orthogonal representation with the minimum number of bends has been extensively investigated in a variable embedding setting [9, 11, 14]. It would be interesting to study the compaction problem when it is possible to change the embedding of the input graph.

Acknowledgments

We would like to thank Maurizio Patrignani for useful discussions.

References

[1] P. Bertolazzi, G. Di Battista, and W. Didimo. Computing orthogonal drawings with the minimum number of bends. In F. Dehne, A. Rau-Chaplin, J.-R. Sack, and R. Tamassia, editors, *Algorithms and Data Structures (Proc. WADS '97)*, volume 1272 of *Lecture Notes Comput. Sci.*, pages 331–344. Springer-Verlag, 1997.

[2] P. Bertolazzi, G. Di Battista, G. Liotta, and C. Mannino. Upward drawings of triconnected digraphs. *Algorithmica*, 6(12):476–497, 1994.

[3] T. C. Biedl. New lower bounds for orthogonal graph drawings. In F. J. Brandenburg, editor, *Graph Drawing (Proc. GD '95)*, volume 1027 of *Lecture Notes Comput. Sci.*, pages 28–39. Springer-Verlag, 1996.

[4] T. C. Biedl and G. Kant. A better heuristic for orthogonal graph drawings. *Comput. Geom. Theory Appl.*, 9(3):159–180, 1998.

[5] J. A. Bondy and U. S. R. Murty. *Graph Theory with Applications.* North-Holland, Amsterdam, The Netherlands, 1976.

[6] G. Di Battista, P. Eades, R. Tamassia, and I. G. Tollis. *Graph Drawing.* Prentice Hall, Upper Saddle River, NJ, 1999.

[7] G. Di Battista, A. Garg, G. Liotta, R. Tamassia, E. Tassinari, and F. Vargiu. An experimental comparison of four graph drawing algorithms. *Comput. Geom. Theory Appl.*, 7:303–325, 1997.

[8] G. Di Battista and G. Liotta. Upward planarity checking: "Faces are more than polygons". In S. Whitesides, editor, *Graph Drawing (Proc. GD '98)*, volume 1547 of *Lecture Notes Comput. Sci.*, pages 72–86. Springer-Verlag, 1999.

[9] G. Di Battista, G. Liotta, and F. Vargiu. Spirality and optimal orthogonal drawings. *SIAM J. Comput.*, 27(6):1764–1811, 1998.

[10] G. Di Battista and R. Tamassia. Algorithms for plane representations of acyclic digraphs. *Theoret. Comput. Sci.*, 61:175–198, 1988.

[11] W. Didimo and G. Liotta. Computing orthogonal drawings in a variable embedding setting. In K. Chwa, editor, *Algorithms and Computation (Proc. ISAAC '98)*, volume 1533 of *Lecture Notes Comput. Sci.*, pages 79–88. Springer-Verlag, 1998.

[12] U. Fößmeier, C. Hess, and M. Kaufmann. On improving orthogonal drawings: The 4M-algorithm. In S. Whitesides, editor, *Graph Drawing (Proc. GD '98)*, volume 1547 of *Lecture Notes Comput. Sci.*, pages 125–137. Springer-Verlag, 1999.

[13] U. Fößmeier and M. Kaufmann. Drawing high degree graphs with low bend numbers. In F. J. Brandenburg, editor, *Graph Drawing (Proc. GD '95)*, volume 1027 of *Lecture Notes Comput. Sci.*, pages 254–266. Springer-Verlag, 1996.

[14] A. Garg and R. Tamassia. Planar drawings and angular resolution: Algorithms and bounds. In J. van Leeuwen, editor, *Algorithms (Proc. ESA '94)*, volume 855 of *Lecture Notes Comput. Sci.*, pages 12–23. Springer-Verlag, 1994.

[15] A. Garg and R. Tamassia. On the computational complexity of upward and rectilinear planarity testing. In R. Tamassia and I. G. Tollis, editors, *Graph Drawing (Proc. GD '94)*, volume 894 of *Lecture Notes Comput. Sci.*, pages 286–297. Springer-Verlag, 1995.

[16] A. Garg and R. Tamassia. A new minimum cost flow algorithm with applications to graph drawing. In S. C. North, editor, *Graph Drawing (Proc. GD '96)*, volume 1190 of *Lecture Notes Comput. Sci.*, pages 201–216. Springer-Verlag, 1997.

[17] N. Gelfand and R. Tamassia. Algorithmic patterns for orthogonal graph drawing. In S. Whitesides, editor, *Graph Drawing (Proc. GD '98)*, volume 1547 of *Lecture Notes Comput. Sci.*, pages 138–152. Springer-Verlag, 1999.

[18] F. Harary. *Graph Theory.* Addison-Wesley, Reading, MA, 1969.

[19] F. Hoffmann and K. Kriegel. Embedding rectilinear graphs in linear time. *Inform. Process. Lett.*, 29:75–79, 1988.

[20] J. E. Hopcroft and J. D. Ullman. *Introduction to Automata Theory, Languages, and Computation.* Addison-Wesley, Reading, MA, 1979.

[21] M. Y. Hsueh. *Symbolic layout and compaction of integrated circuits.* PhD thesis, Univ. California, Berkeley, CA, 1980.

[22] M. Y. Hsueh and D. O. Pederson. Computer-aided layout of lsi circuit building-blocks. In *Proc. IEEE Int. Symp. on Circuits and Systems*, 1979.

[23] G. W. Klau and P. Mutzel. Optimal compaction of orthogonal grid drawings. Technical Report MPI-I-98-1-031, Max–Planck–Institut für Informatik, Dec. 1998. `http://data.mpi-sb.mpg.de/internet/reports.nsf/NumberView/98-1-031`.

[24] G. W. Klau and P. Mutzel. Optimal compaction of orthogonal grid drawings. In G. P. Cornuejols, editor, *Integer Programming and Combinatorial Optimization (Proc. IPCO '99)*, Lecture Notes Comput. Sci. Springer-Verlag, to appear.

[25] T. Lengauer. *Combinatorial Algorithms for Integrated Circuit Layout.* Wiley-Teubner, 1990.

[26] R. H. J. M. Otten and J. G. van Wijk. Graph representations in interactive layout design. In *Proc. IEEE Internat. Sympos. on Circuits and Systems*, pages 914–918, 1978.

[27] A. Papakostas and I. G. Tollis. Algorithms for area-efficient orthogonal drawings. *Comput. Geom. Theory Appl.*, 9(1–2):83–110, 1998. Special Issue on Geometric Representations of Graphs, G. Di Battista and R. Tamassia, editors.

[28] M. Patrignani. On the complexity of orthogonal compaction. Technical Report RT-DIA-39-99, Dipartimento di Informatica e Automazione, Università di Roma Tre, Jan. 1999.

[29] J. M. Six, K. G. Kakoulis, and I. G. Tollis. Refinement of orthogonal graph drawings. In S. Whitesides, editor, *Graph Drawing (Proc. GD '98)*, volume 1547 of *Lecture Notes Comput. Sci.*, pages 302–315. Springer-Verlag, 1999.

[30] L. Stockmeyer. Optimal orientation of cells in slicing floorplan design. *Inform. Control*, 57:91–101, 1983.

[31] R. Tamassia. On embedding a graph in the grid with the minimum number of bends. *SIAM J. Comput.*, 16(3):421–444, 1987.

[32] R. Tamassia and I. G. Tollis. Planar grid embedding in linear time. *IEEE Trans. Circuits Syst.*, CAS-36(9):1230–1234, 1989.

[33] R. Tamassia, I. G. Tollis, and J. S. Vitter. Lower bounds for planar orthogonal drawings of graphs. *Inform. Process. Lett.*, 39:35–40, 1991.

[34] G. Vijayan and A. Wigderson. Rectilinear graphs and their embeddings. *SIAM J. Comput.*, 14:355–372, 1985.

[35] W. Wimer, I. Koren, and I. Cederbaum. Floorplans, planar graphs and layouts. *IEEE Trans. Circuits Syst.*, CAS-35:267–278, 1988.

A Appendix

Symbol	*Description*	*Section*
G	an embedded 4-planar graph	§2.1
H	an orthogonal representation of G	§2.1
Γ	a drawing of H	§3.1
Γ_r	an orientation of Γ with all vertical segments directed upward and all horizontal segments directed rightward	§3.3
Γ_ℓ	an orientation of Γ with all vertical segments directed upward and all horizontal segments directed leftward	§3.3
H_r	the orientation of H induced by Γ_r	§3.3
H_ℓ	the orientation of H induced by Γ_ℓ	§3.3
G_r	the orientation of G induced by H_r	§3.3
G_ℓ	the orientation of G induced by H_ℓ	§3.3
H_x	a partially-directed graph representing the "left" relation between maximal vertical chains of H	§4
H_y	a partially-directed graph representing the "below" relation between maximal horizontal chains of H	§4
D_x	a planar st-digraph obtained by shrinking the maximal vertical chains of H_x	§6
D_y	a planar st-digraph obtained by shrinking the maximal horizontal chains of H_y	§6
N_x	the dual planar st-digraph of D_x	§6
N_y	the dual planar st-digraph of D_y	§6

Table 2: The various symbols used throughout the paper to denote graphs, orthogonal representations, and drawings.