Karlsruhe Institute of Technology

Upward-Rightward-Prescribed Planarity

Bachelor Thesis of

Vera Chekan

At the Department of Informatics
Institute of Theoretical Informatics

Reviewers:	Prof. Dr. Dorothea Wagner
	Prof. Dr. Peter Sanders
Advisors:	Guido Brückner, M.Sc.

Time Period: 1st June 2019 – 30th September 2019

Statement of Authorship

I hereby declare that this document has been composed by myself and describes my own work, unless otherwise acknowledged in the text.

Karlsruhe, September 29, 2019

Abstract

We introduce and investigate a new planarity variant for directed graphs called upward-rightward-prescribed planarity. An upward-rightward-prescribed graph is a directed graph in which every edge is assigned either a u- or an r-label. Such a graph is called upward-rightward-prescribed planar if there exists a planar drawing in which each u-labeled edge is drawn y-monotonic and each r-labeled edge is drawn x-monotonic. This planarity variant is strongly related to upward planarity and windrose planarity.

We show that testing graphs for upward-rightward-prescribed planarity is $\mathcal{NP}$-hard even for single-source graphs and consider a restricted setting where the subgraph induced by the u-labeled edges is a biconnected spanning subgraph and the embedding of this subgraph is fixed.

For this restriction, we provide a linear-time algorithm that decides if an upward-rightward-prescribed drawing exists and in the positive case, constructs such a drawing. This algorithm reduces upward-rightward-prescribed-planarity testing to windrose-planarity testing. We show that there exist upward-rightward-prescribed planar graphs that admit no straight-line drawing. The relation to windrose planarity implies that every upward-rightward-prescribed planar graph has an (upward-rightward-prescribed planar) drawing on a grid of polynomial size with at most three bends per edge.

Deutsche Zusammenfassung

Wir stellen eine neue Planaritätsvariante für gerichtete Graphen vor, welche als aufwärts-rechts bezeichnet wird, und untersuchen diese. In einem aufwärts-rechts Graphen wird jeder Kante entweder eine u- oder eine r-Markierung zugewiesen. Ein solcher Graph ist aufwärts-rechts-planar, falls er so gezeichnet werden kann, dass jede Kante mit einer u-Markierung monoton entlang der y-Achse und jede Kante mit einer r-Markierung monoton entlang der x-Achse verläuft. Diese Variante der Planarität ist eng mit Aufwärtsplanarität und Windrosenplanarität verwandt.

Wir zeigen, dass das Testen eines Graphen, selbst mit nur einer Quelle, auf Aufwärts-rechtsplanarität im Allgemeinen $\mathcal{NP}$-schwer ist, und betrachten daher nur eine eingeschränkte Variante des Problems, in welcher der durch die u-markierten Kanten induzierte Subgraph aufspannend und zweifach zusammenhängend ist und dessen Einbettung fest ist.

Für diese eingeschränkte Variante geben wir einen Linearzeit-Algorithmus an, welcher entscheidet, ob für einen gegebenen Graphen eine aufwärts-rechts-planare Zeichnung existiert und gegebenenfalls eine solche konstruiert. Der Algorithmus reduziert den Test auf Aufwärts-Rechtsplanarität auf den Test auf Windrosenplanarität. Wir zeigen, dass es Graphen gibt, welche zwar aufwärts-rechts-planar sind, jedoch nicht geradlinig gezeichnet werden können. Aus der Beziehung zu Windrosenplanarität schließen wir auf die Existenz von Zeichnungen auf einem polynomiellen Gitter mit höchstens drei Knicken pro Kante.

Contents

1. Introduction

Many problems deal with drawings for which the relative or absolute positions of vertices are prescribed. Upward planarity is a property of digraphs that have a planar drawing in which every edge is represented by a curve that monotonically increases in the vertical direction. This planarity variant has attracted considerable attention. Garg and Tamassia have proven that upward-planarity testing is $\mathcal{NP}$-hard [GT95]. Bertolazzi and DiBattista have found a polynomial algorithm that solves the problem for a fixed combinatorial embedding [BDB91]. Bertolazzi et al. have also found a linear-time algorithm that decides if a single-source digraph admits an upward drawing and, if so, constructs one [BDBMT98]. In this thesis, we consider a planarity variant that generalizes the concept of upward planarity.

Level planarity is a similar concept. The difference to upward planarity is that the y-coordinates of vertices are fixed. It has been shown that level planarity can be tested in linear time [JLM98].

In Upward-Rightward Planarity the question is whether there exists a planar drawing of a directed graph in which every edge monotonically increases in the horizontal or in the vertical direction or in simple words, no edge points to the southwest. It has been shown that every directed planar graph has an upward-rightward drawing with straight-line edges and such a drawing can be computed in linear time and polynomial area [DGDK$^+$14]. In this thesis, we consider the setting where we additionally prescribe for each edge if it has to increase in the horizontal or in the vertical direction.

In Bi-Monotonicity the coordinates of every vertex of an undirected graph are fixed and the question is whether there exists a planar drawing that respects this prescription and in which every edge is both x- and y-monotone. It has been proven that testing graphs for bi-monotonicity is $\mathcal{NP}$-hard [KR17].

Another related concept is HV-rectilinear planarity. An HV-restricted planar graph is a planar undirected graph G with vertex-degree at most four and such that each edge is labeled either H (horizontal) or V (vertical). The HV-Rectilinear Planarity Testing problem asks whether G admits a planar drawing where every edge labeled V is drawn as a vertical segment and every edge labeled H is drawn as a horizontal segment. It has been proven that HV-rectilinear-planarity testing is $\mathcal{NP}$-complete even for graphs having vertex-degree at most three. It has also been shown that HV-rectilinear planarity can be tested in polynomial time for partial 2-trees of maximum degree four [DLP14].

And finally, Windrose Planarity: an instance of this problem consists of a planar graph and labeling that, for every edge $\{u, v\}$, prescribes in which quadrant around u the vertex v has to be drawn. A windrose drawing is a planar drawing that respects this prescription and in which every edge is monotone in both the horizontal and the vertical direction. Windrose-planarity testing is also $\mathcal{NP}$-hard but becomes efficiently solvable for a fixed embedding [ALB$^+$18]. In this work, we provide a simple criterion of windrose planarity of graphs with a fixed embedding and further apply it to prove the existence of certain drawings.

In this thesis, we consider a new planarity variant called upward-rightward-prescribed planarity. The concept was suggested by Angelini et al. as an open question in [ALB$^+$18]. In Upward-Rightward-Prescribed Planarity every edge of a directed graph is labeled either u (upward) or r (rightward) and we ask whether there exists a planar drawing that respects this labeling, that is, every u-edge is drawn y-monotonic and every r-edge is drawn x-monotonic. Such drawings represent two partial orders on the same set of vertices whereas upward or level planarity only deals with one partial order at a time. Although upward-rightward-prescribed planarity is related to the aforementioned planarity variants, the concept is new: in the Bi-Monotonicity problem setting, the vertices are totally ordered relative to both the x- and the y-axis, windrose planarity only deals with a local relationship between adjacent vertices and in HV-rectilinear and upward-rightward planarity, the direction of the monotonicity is not fixed.

Because Upward-Rightward-Prescribed Planarity is a generalization of Upward Planarity, it is only worth considering a restricted problem setting: the general setting is $\mathcal{NP}$-hard. In this thesis, the combinatorial embedding is fixed and additionally, for every vertex the left-to-right order of the incoming u-edges and the left-to-right order of the outgoing u-edges are prescribed; this embedding is called an upward embedding. Furthermore, we only consider graphs whose upward or rightward part is a spanning subgraph and this subgraph is biconnected.

This thesis is structured as follows. First, in Chapter 2, we introduce definitions and related results that we will work with. Then, in Chapter 3, we provide a simple criterion of windrose planarity of graphs with fixed embedding that will be used throughout the thesis. After that, in Chapters 4 and 5, we begin to study the problem by looking into an arbitrary upward-rightward-prescribed drawing and making important observations that result in necessary conditions for upward-rightward-prescribed planarity. In Chapter 6, we construct a 2-Sat instance that represents this set of necessary conditions. Next, in Chapter 7, we show that the satisfiability of this instance is sufficient for the existence of an upward-rightward-prescribed drawing. To prove this, we transform our instance into an instance of Windrose Planarity, apply the criterion provided in Chapter 3 to show that there exists a windrose planar drawing. Using the algorithm by Angelini et al., we obtain a 3-bend drawing of the original graph. In this chapter, we also show that not every upward-rightward-prescribed planar graph admits a straight-line drawing. In the aforementioned chapters, we only consider the instances of the problem with a special type of embedding called *linearized embedding*: for every vertex, the clockwise order of the edges in each quadrant is fixed. Chapter 8 provides the generalization of the approach so that the problem can be solved if the linearized embedding is not fixed but the above-mentioned upward embedding of the upward part of the graph is still prescribed. This is the main result of this work. Additionally, in Chapter 9, we prove that upward-rightward-prescribed-planarity testing is $\mathcal{NP}$-hard even for single-source graphs.

2. Preliminaries

This chapter provides definitions and related results that will be used within the next chapters.

2.1 Embeddings of Planar Graphs

A *drawing* of graph fixes the exact coordinates of each vertex and the exact polyline for each edge: in particular, every edge consists of a finite number of straight-line segments. A drawing is *planar* if edges intersect only at their endpoints. A planar drawing of a graph divides the plane into regions called *faces*. The unbounded region is called the *outer face*. For each planar drawing there exists a unique *combinatorial embedding* that defines the clockwise order of edges incident to the same vertex. A combinatorial embedding defines uniquely the faces of the drawing but it does not define the outer face. Indeed, for a fixed combinatorial embedding every face can be chosen to be the outer face.

2.2 Upward Planarity

A directed graph (also called a *digraph*) $G = (V, E)$ is *upward planar* if there exists a planar drawing of G in which each edge monotonically increases in the vertical direction; such a drawing is an *upward drawing*. Garg and Tamassia have shown that upward-planarity testing is $\mathcal{NP}$-hard [GT95]. However, Bertolazzi and DiBattista have shown that for a fixed combinatorial embedding, the problem can be solved in polynomial time [BDB91]. Within the work, we will use their results, so now we introduce the definitions and the main results provided in this paper.

A connected graph is *biconnected* if the connectivity of the graph is maintained after the removing an arbitrary vertex. In the context of upward planarity, we only consider biconnected graphs. A vertex in a digraph is a *source* (*sink*) if it only has incident outgoing (incoming) edges. A vertex v incident to a face f is a *face-source* (*face-sink*) of f if both edges incident to v and lying on the border of f are outgoing (incoming). A vertex is a *switch* of a face if it is a face-source or a face-sink of this face.

Given an upward drawing Γ of a digraph G, for each vertex v, we can identify two (possibly empty) linear lists of consecutive incoming and outgoing edges incident to v by visiting these edges from left to right. Two upward drawings are *equivalent* if, for every vertex v, they define the same two linear lists. A class of equivalent upward drawings is an *upward* *embedding*. Further on, writing about *fixed upward embedding* we mean that for every vertex the aforementioned lists are fixed.

Let v be a source or a sink of G. Let l and r be two edges incident to v and neighboring in its adjacency list. A face f incident to l and r is *identified* by l and r if f lies to the left of l and to the right of r. Notice that left and right of edges are well-defined for directed graphs. For such a vertex, only one of the lists associated with it is not empty. We say that v is *assigned* to a face f of Γ if f is the face *identified* by the first and the last edges of this list and we also say that the angle between these edges is *large*. We say that v is *not assigned* to all the other faces and the corresponding angle is called *small*.

An *upward face* of an upward drawing is a face f in which for every switch is prescribed if it is assigned to this face. Recall that a switch can only be assigned to a face if it is also a sink or a source of the entire graph.

Observation 2.1. *All the upward drawings of the same upward embedding have the same set of upward faces. Moreover, they all have the same external upward face [BDB91].*

Lemma 2.2. *Let Φ be an upward embedding of G, let h be its external upward face, and let f be an internal upward face. Denote with $2n_h$ and $2n_f$ the number of switches of h and f, respectively. The number of sources and sinks of G assigned to f in Φ is equal to $n_f - 1$. The number of sources and sinks of G assigned to face h is equal to $n_h + 1$ [BDB91].*

A planar embedding of an acyclic digraph is a *candidate upward embedding* if, for each vertex v, all the outgoing (incoming) edges appear consecutively in the adjacency list of v.

Given a candidate upward embedding, an outer face, and an assignment of sources and sinks to faces, i.e., an angle assignment for switches of faces, this assignment is *upward-consistent* if it satisfies:

1. Each source or sink is assigned to exactly one face.
2. For each face, the condition of Lemma 2.2 is satisfied.

Lemma 2.3. *Given a candidate upward embedding Ψ and a candidate external upward face h of Ψ, the embedding Φ obtained from the assignment A of sources and sinks to faces is an upward embedding of G with external face h if and only if A is an upward-consistent assignment [BDB91].*

This theorem provides a simple way to test if an edge (u, v) can be inserted into a fixed upward embedding of a graph so that the arising embedding is also upward: firstly, test if both the incoming and outgoing edges incident to u and v are consecutive and secondly, check if both faces arising by this insertion satisfy the condition stated in Lemma 2.2. Notice that a fixed upward embedding prescribes the assignment of switches to faces and hereby, we only apply this approach if the underlying angle assignment has to be maintained throughout the insertion of the edge.

2.3 Upward-Rightward-Prescribed Planarity

Let $G = (V, E = E_r \dot\cup E_u)$ be a digraph. An *upward-rightward-prescribed drawing* of G is a planar drawing in which every edge $e \in E_u$ (called a *u-edge*) monotonically increases in the vertical direction and every edge $e \in E_r$ (called an *r-edge*) monotonically increases in the horizontal direction. If G admits such a drawing, it is called *upward-rightward-prescribed planar*.

We emphasize that the partition of E into E_r and E_u is a part of an instance. If this partition is not fixed, it is about an instance of UPWARD-RIGHTWARD PLANARITY [DGDK+14].

This problem setting is a generalization of the UPWARD PLANARITY problem: an instance $G' = (V', E')$ of UPWARD PLANARITY is equivalent to the instance $(V', (E_u = E')\dot{\cup}(E_r = \emptyset))$ of UPWARD-RIGHTWARD-PRESCRIBED PLANARITY. This immediately implies the $\mathcal{NP}$-hardness of UPWARD-RIGHTWARD-PRESCRIBED PLANARITY.

Theorem 2.4. *Upward-rightward-prescribed-planarity testing is $\mathcal{NP}$-hard.*

For this reason, we are only interested in a restricted problem setting. Firstly, in this thesis, we only consider graphs for which the subgraph G_u induced by E_u is a biconnected spanning subgraph of G, which means in particular that every vertex has an incident u-edge. Secondly, the algorithm provided in this thesis works with a fixed combinatorial embedding of G and a fixed upward embedding (a fixed angle assignment) of G_u. The algorithm replaces every r-edge by a set of u-edges. The above restrictions allow to apply the criterion stated in Lemma 2.3 to test whether these u-edges can be embedded into G_u. First, we consider a stricter embedding called "linearized embedding" introduced in the next paragraph and then in Chapter 8, we generalize the approach to solve the problem for the above-mentioned variant of an embedding.

Consider a planar drawing of a digraph and let v be a vertex in it. A *linearized order* of edges incident to v is the clockwise order in which edges appear around v starting at the ray that exits v to the left (see Figure 2.1). A *linearized embedding* is an assignment of the linearized order of incident edges to each vertex. As already mentioned, before Chapter 8, we consider instances of UPWARD-RIGHTWARD-PRESCRIBED PLANARITY with a fixed linearized embedding $\mathcal{E}$.

Notice that, amongst other things, $\mathcal{E}$ uniquely defines the upward embedding of G_u as follows. If we consider only u-edges, then for an arbitrary upward-rightward-prescribed drawing, in the linearized order around every vertex, first the outgoing u-edges appear in the left-to-right-order and then the incoming u-edges appear in the right-to-left-order. Hereby, the upward embedding of G_u can be extracted from $\mathcal{E}$ in linear time.

Let Γ be an arbitrary upward-rightward-prescribed drawing of $(G, \mathcal{E})$, the omission of E_r yields an upward drawing of G_u. This results in the following lemma:

Lemma 2.5. *Upward planarity of $\mathcal{E}_u := \mathcal{E}_{|G_u}$ is a necessary condition for the upward-rightward-prescribed planarity of $(G, \mathcal{E})$.*

From now on, we assume that $\mathcal{E}_u$ is upward. Notice that this is not a restriction because, given $\mathcal{E}$, it is possible to extract $\mathcal{E}_u$ and test the condition in Lemma 2.3 in linear time [BDB91].

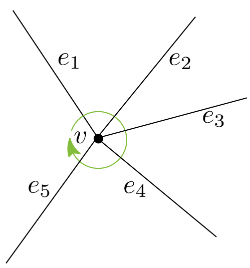

Figure 2.1: The linearized order of edges around v is $[e_1, e_2, e_3, e_4, e_5]$.

2.4 Windrose Planarity

Windrose Planarity is a related problem. An instance of this problem consists of a graph and labeling that, for every edge $\{u, v\}$, prescribes the quadrant around u in which v lies. The quadrants are denoted by NW (North-West), NE (North-East), SW (South-West), and SE (South-East). The question is whether there exists a planar drawing that respects this prescription and in which every edge is represented by a curve that is both x- and y-monotone.

This problem was introduced by Angelini et al. in [ALB+18]. They have shown that the problem is $\mathcal{NP}$-hard in its general form but can be solved in polynomial time for a fixed combinatorial embedding. They have also proven that every windrose planar graph has a 1-bend drawing.

In this thesis, we reformulate the problem setting: we consider directed graphs that arise if we direct every edge from the South to the North. This is not a restriction: if both end-vertices of an edge are prescribed to lie to the North (South) of each other, then there exists no windrose planar drawing and this can be detected in linear time. In Chapter 3, we will provide and use a simple criterion of windrose planarity of graphs with fixed upward embedding.

An instance of the adapted problem setting contains two types of edges: the ones that monotonically increase to the north-east (NE-edges) and the ones that monotonically increase to the north-west (NW-edges). A polyline is called an $NE(NW)$-line if it monotonically increases in the vertical direction and monotonically increases (decreases) in the horizontal direction. Hereby, we look for a drawing in which every $NE(NW)$-edge is drawn as an $NE(NW)$-line.

2.5 2-Sat

In Chapter 6, we represent constraints as a 2-Sat instance.

Let W be a set of boolean variables. The set of *literals* over W is defined as $L = \{w, \neg w \mid w \in W\}$. A *formula* over W is a disjunction of two literals from L. A 2-Sat instance is denoted by (V, F) where F is a set of formulas. Instance (V, F) is a satisfiable exactly if there exists an assignment of boolean values to variables so that in every formula at least one literal is assigned a $true$-value. There exist linear-time algorithms solving the 2-Sat problem, for example [APT79].

2.6 Main Approach

The approach that we follow in this work:

1. The criterion of upward planarity provided in [BDB91] yields a simple algorithm to test if certain sets of edges can be inserted into a fixed upward embedding so that the arising embedding is also upward.
2. We will consider an arbitrary upward-rightward-prescribed drawing and detect how an r-edge can be drawn in such a drawing.
3. We will also find dependencies between representations of different edges in G (for example, incident to the same vertex) and represent these dependencies as a 2-Sat instance. If no solution exists, then the graph is not upward-rightward-prescribed planar. Otherwise, a solution for this instance provides more information about how exactly the edges can look like in the desired drawing.

4. After that, this solution will be transformed into an instance of Windrose Planarity.

5. We will show that a windrose drawing of this graph indeed exists and it yields an upward-rightward-prescribed drawing of the initial graph.

6. As a result:

    a) If no upward-rightward prescribed drawing exists, then we find this out (the corresponding 2-Sat instance is unsolvable).

    b) If such a drawing exists, then we provide one.

Further on, we provide illustrations. Unless specified otherwise, r-edges are drawn red and u-edges are drawn black.

3. Windrose Planarity for a Fixed Embedding

Before we begin considering the Upward-Rightward-Prescribed Planarity problem, we first provide a simple criterion of windrose planarity of graphs with fixed upward embedding. We recall that an upward embedding determines the left-to-right orders of incoming and outgoing edges incident to the same vertex. This criterion can be derived from the original paper [ALB$^+$18] but we still provide our constructive proof as we follow another approach that can be interesting on its own. Within the work, we will deal a lot with Windrose Planarity and this criterion will be widely used to prove the existence of particular drawings.

As mentioned earlier in preliminaries, we reformulate the original setting of the Windrose Planarity problem so that it better suits our purposes. Instead of considering undirected graphs as in the original paper, we assign the direction "from South to North" to every edge and then, concentrate on directed graphs. This is not a restriction: if both end-vertices of an edge are prescribed to lie to the North (South) of each other, then there is no windrose planar drawing and this can be detected in linear time.

Recall that an $NE(NW)$-line is a polyline that monotonocally increases in the vertical direction and monotonically increases (decreses) in the horizontal direction. So we consider the following setting of the Windrose Planarity problem: let $G = (V, E_{NW} \dot\cup E_{NE})$ be a digraph. We look for a planar drawing of G in which:

1. Every edge $e \in E_{NE}$ is represented by an NE-line. These edges are called *NE-edges*.
2. Every edge $e \in E_{NW}$ is represented by an NW-line. These edges are called *NW-edges*.

Every edge has to increase in y-direction and hereby, a windrose drawing of G is also an upward drawing of G. Hence, the upward planarity of G is a necessary condition for the windrose planarity of G. From now on, we are only interested in the restriction of the problem to fixed upward embeddings. Notice that given an assignment of left-to-right orders of incoming and outgoing edges incident to the same vertex, it is possible to test whether it yields an upward embedding by checking the condition in Lemma 2.2. This checking is possible in linear time. So from now on, we assume that the embedding is upward.

The following observation was made in [ALB$^+$18]:

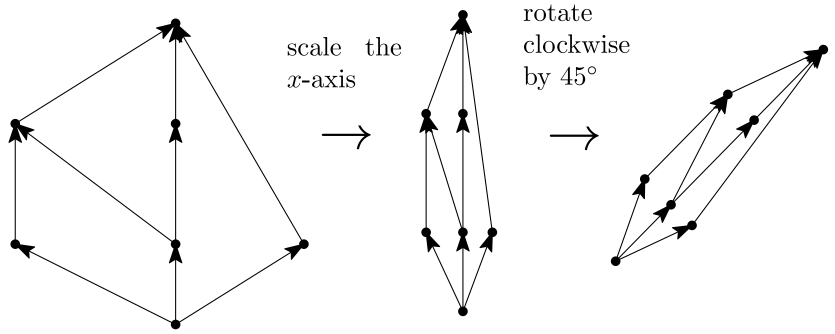

Figure 3.1: The construction of an upward-rightward straight-line drawing of an upward planar digraph. On the left: Γ, in the middle: Γ', on the right: Γ^*.

Observation 3.1. *For an upward embedding $\mathcal{E}$, there exists a straight-line upward drawing in which every edge has a positive slope (such a drawing is called upward-rightward).*

Proof. We briefly describe the main idea of the construction. According to [DBT88], there exists a straight-line upward drawing Γ of $\mathcal{E}$. Transform Γ into a new drawing Γ' of G in which the absolute value of each slope is larger than 1 by scaling the x-axis. Finally, rotate Γ' clockwise by 45°. Let Γ^* be the arising drawing. This drawing is also upward and each edge is represented by a straight-line segment having a positive slope in it. Hereby, Γ^* is the desired drawing (see Figure 3.1). □

Observation 3.2. *A drawing that satisfies the conditions stated in Observation 3.1 can be easily adapted so that it still has these properties and furthermore, no two vertices are drawn with the same x-coordinate.*

The following observation was also made in [ALB$^+$18] (we reformulated it so that it suits our problem setting):

Observation 3.3. *In any embedding corresponding to a windrose-planar drawing of a digraph G, for each vertex v of G, the edges incident to v appear around v in this clockwise order starting at 0°: first the incoming NW-edges, then the incoming NE-edges, then the outgoing NW-edges, and finally, the outgoing NE-edges.*

As a result, for every vertex in a windrose embedding, its adjacency list can be partitioned into four lists of consecutive edges that contain only incoming NW-edges, incoming NE-edges, outgoing NW-edges, and outgoing NE-edges respectively. This property can be checked in linear time. For a fixed upward embedding, this condition is also sufficient for the existence of a windrose drawing and the remainder of the chapter deals with proving it. Further, we only consider an embedding with this property and show how to create a windrose drawing of it.

In this chapter, we consider drawings of digraphs. If the drawing is clear from the context, then we sometimes write "edges (vertices)" instead of "points (lines) which represent these edges (vertices)".

In order to produce a windrose drawing of an embedding, we start with a certain drawing and change it iteratively. Each iteration corresponds to an NW-edge. In the iteration that corresponds to an NW-edge e, we *repair* e (afterward, e is drawn with an NW-line) and maintain the following invariants:

1. The drawing is planar.

2. The drawing is upward.
3. Every edge is drawn as a polyline.
4. Every NE-edge is drawn as an NE-line.
5. Every NW-edge whose iteration has already been carried out is drawn as an NW-line.
6. No two vertices are drawn with the same x-coordinate.

A drawing that has properties stated in Observation 3.2 is used as an initial drawing as it satisfies these invariants: the fifth invariant is satisfied because at the beginning no NW-edge has been repaired yet, other invariants are clear.

3.1 One iteration

In this section, we describe the transformation which is carried out during one iteration.

Let Γ be the actual drawing and let $e = (v, w)$ be the NW-edge that is repaired in the current iteration. First, remove $\Gamma(e)$ from the drawing.

The drawing Γ is upward and therefore, $\Gamma(v)_y < \Gamma(w)_y$. Define a polyline g_e as follows:

$$g_e(y) = \begin{cases} (\Gamma(w)_x, y) & y \geq \Gamma(w)_y \\ (\Gamma(v)_x, y) & y \leq \Gamma(v)_y \\ (\Gamma(e)(y), y) & \text{otherwise} \end{cases}$$

This polyline is well-defined because $\Gamma(e)$ is y-monotone. It is also continuous (endpoints of $\Gamma(e)$ are $\Gamma(v)$ and $\Gamma(w)$) and y-monotone.

Choose $\alpha_1 \in (0, \pi/2)$ so that the ray starting at $\Gamma(v)$ and pointing upwards with counterclockwise angle α_1 to the vertical line $x = x_v$ doesn't cross $\Gamma(e)$ and lies to the left of it. Choose $\alpha_2 \in (0, \pi/2)$ so that the ray starting at $\Gamma(w)$ and pointing downwards with clockwise angle α_2 to the vertical line $x = x_w$ doesn't cross $\Gamma(e)$ and lies to the right of it. Such angles α_1 and α_2 exist because $\Gamma(e)$ is a polyline (see Figure 3.2).

Set $\alpha = max\{\alpha_1, \alpha_2\}$. Set for a small μ:

$$\Delta = \max\{\frac{\Gamma(w)_y - \Gamma(v)_y}{\tan(\alpha)}, \max_{(x,y)\in c_e} x - \min_{(x,y)\in c_e} x\} + \mu \tag{3.1}$$

Shift the drawing strictly to the right of g_e **and** the point $\Gamma(v)$ on Δ to the right; in particular, the vertex w stays represented by the same point. The motivation behind the shifting is to guarantee that the straight-line segment connecting v and w has a negative

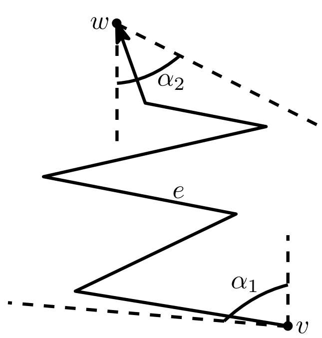

Figure 3.2: The choice of Δ

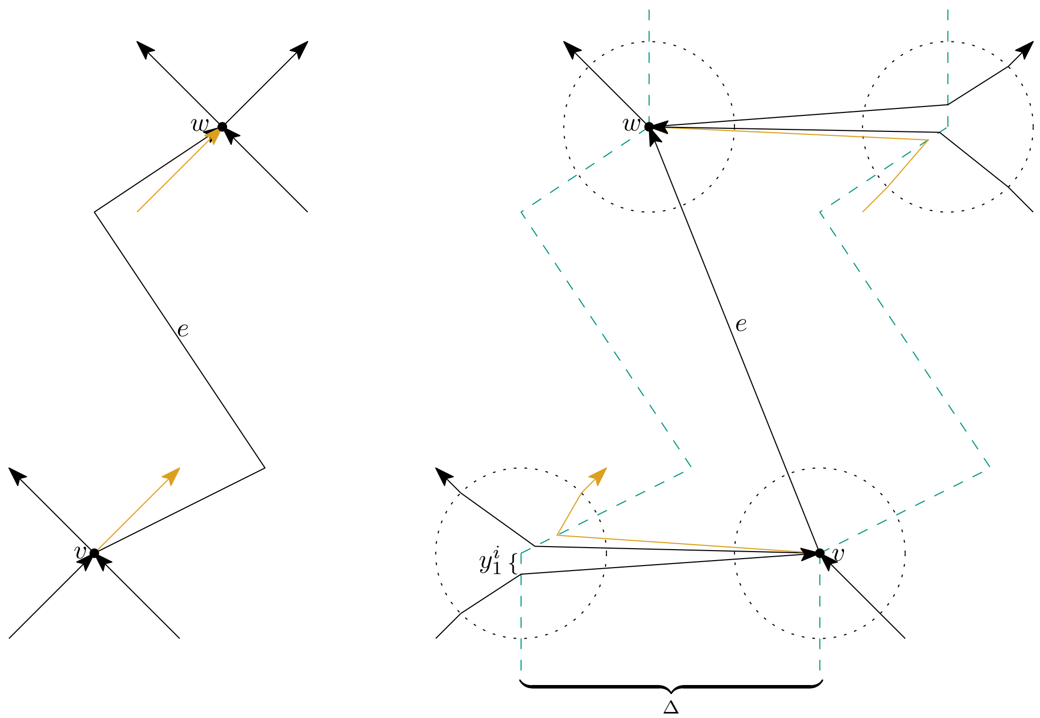

Figure 3.3: The drawings Γ (on the left) and Γ^* (on the right) in the neighborhood of (v, w). The only edges whose x-monotonicity can be destroyed by this transfomation are drawn orange.

slope and hence, it is an NW-line. Further on, we write p^Δ for the point with coordinates $(p_x + \Delta, p_y)$ and we write g_e^Δ for the polyline which arises from g_e by shifting it on Δ to the right. Let Γ′ be the actual drawing.

As a result of this shifting, some edges can be *torn*, which means the polyline representing such an edge can now be discontinuous. In the next subsection, we describe how to restore the continuity of these edges. Notice that an edge could have been torn at several points; these points are considered independently from each other.

3.1.1 Repairing Torn Edges

An edge f can only be torn if $\Gamma(f)$ crosses g_e at some point p. This point p satisfies $p_y \leq \Gamma(v)$ or $p_y \geq \Gamma(w)$; otherwise, the edges e and f would cross at an inner point of e in Γ which is prevented by the planarity of Γ. Hereby, p lies outside of $\Gamma(e)$ or it coincides with either $\Gamma(v)$ or $\Gamma(w)$. We distinguish between these three cases.

3.1.1.1 Edges Torn at $\Gamma(v)$

First, we describe how to repair edges torn at $\Gamma(v)$.

Further, $U_\epsilon(A)$ denotes the ϵ-neighborhood of the point A. Choose a small ϵ_v, so that in Γ:

1. The neighborhood $U_{\epsilon_v}(\Gamma(v)) =: U_v$ contains no other vertices.
2. Only edges incident to v have common points with U_v.
3. Each edge is represented by a straight-line segment in U_v.

Note that such a neighbourhood exists as every edge is a polyline. Remove the drawing from U_v to the left of g_e.

Every edge (v_e, w_e) is a polyline in Γ. Further, the straight-line segment incident to $\Gamma(v)$ or $\Gamma(w)$ is called the *first* or the *last* straight-line segment of (v_e, w_e) respectively.

Let H_v denote the border of U_v. Let $[i_1, \ldots, i_k, o_1, \ldots, o_t]$ be the clockwise order of edges incident to v that cross H_v to the left of g_e in Γ so that $\{o_1, \ldots, o_t\}$ are outgoing edges and the remaining edges are incoming. As these edges lie to the left of g_e, their last straight-line segments have not been shifted whereas the vertex v was shifted and hence, these edges are exactly the edges torn at $\Gamma(v)$.

Choose $y_1^i < \cdots < y_k^i < \Gamma(v)_y < y_1^o < \cdots < y_t^o$ close to $\Gamma(v)_y$ so that:

1. The horizontal line $y = y_1^i$ lies above the intersection of $\Gamma(i_k)$ and $H(v)$.
2. The horizontal line $y = y_1^o$ lies under the intersection of $\Gamma(o_1)$ and $H(v)$.
3. The horizontal line $y = y_1^o$ lies under the intersection of $\Gamma(o_t)$ and $H(v)$.
4. The straight-line segment $\overline{\Gamma(v)^\Delta g_e(y_t^o)}$ lies under the straight-line segment $\overline{\Gamma(v)^\Delta \Gamma(w)}$ (this straight-line segment is going to represent (v, w) in the arising drawing).

We clarify the choice of this values three paragraphs later.

For $j \in \{1, \ldots, k\}$, let p_j^i be the intersection of $H(v)$ and $\Gamma(i_j)$. Draw the straight-line segments $\overline{\Gamma(v)^\Delta g_e(y_j^i)}$ and $\overline{g_e(y_j^i) p_j^i}$.

For $j \in \{1, \ldots, t\}$, let p_j^obe the intersection of $H(v)$ and $\Gamma(o_j)$. Draw the straight-line segments $\overline{\Gamma(v)^\Delta g_e(y_j^o)}$ and $\overline{g_e(y_j^o) p_j^o}$.

As a result, the gaps at $\Gamma(v)$ that arose by the shifting are now eliminated.

The first three conditions for the choice of the values y_j^i and y_j^o guarantee that the y-monotonicity of the torn edges is maintained. The last one preserves the combinatorial embedding.

It should be noted that simply connecting v^Δ with every p_j^o (p_j^i) by a straight-line segment would possibly change the combinatorial embedding of the drawing. For this reason, we are forced to produce an additional bend at the point $g_e(y_j^i)$ $(g_e(y_j^o))$.

3.1.1.2 Edges Torn at $\Gamma(w)$

A symmetrical transformation is carried out with edges torn at $\Gamma(w)$. The main difference is that in this case the torn edges lie in Γ to the right of g_e. The remainder is analogous (see Figure 3.3).

3.1.1.3 Edges Torn above $\Gamma(w)$ or under $\Gamma(v)$

Now we describe how to restore the continuity of edges that cross g_e under $\Gamma(v)$ or above $\Gamma(w)$ in Γ. Since Γ satisfies the sixth invariant, there are no vertices lying on g_e other than v and w and hereby, the aforementioned edges are torn at an internal point. Let P be the set of such points. For every $p \in P$, let f_p be the edge to which p belongs; choose a small axis-parallel rectangle R_p with the midpoint p of height h_p and width b_p so that:

1. $\Gamma(f_p)$ is represented by a straight-line segment in R_p
2. R_p contains no vertices and no other edges in Γ.

Set

$$\beta = \min_{p \in P} b_p$$

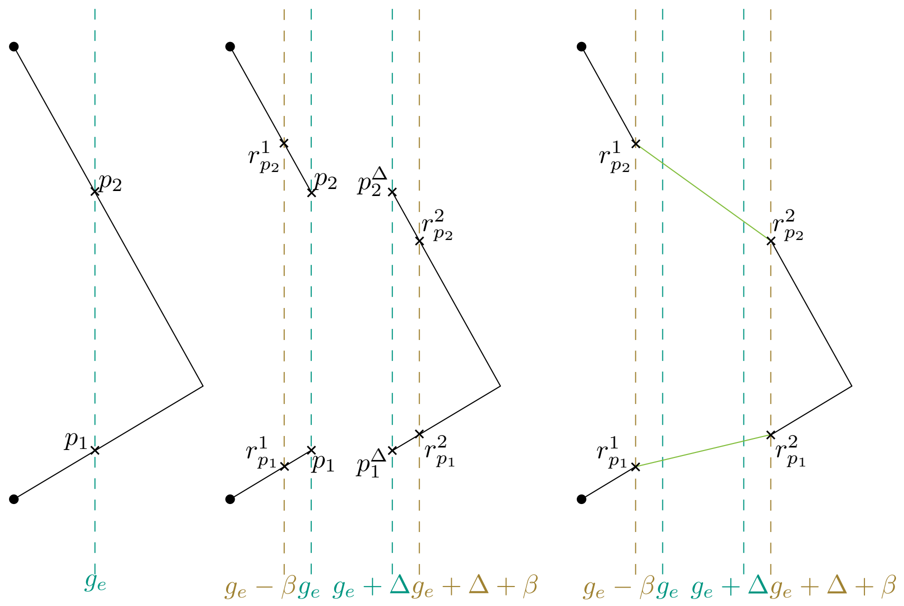

Figure 3.4: Repairing an edge under $\Gamma(v)$ or above $\Gamma(w)$. On the left: Γ; in the middle: the edge is torn after the shifting; on the right: the continuity of the edge is restored.

Remove the current drawing from

$$([\Gamma(w)_x - \beta, \Gamma(w)_x] \cup [\Gamma(w)^\Delta_x, \Gamma(w)^\Delta_x + \beta]) \times ((-\infty, \Gamma(v)_y - \epsilon_v] \cup [\Gamma(w)_y + \epsilon_w, \infty))$$

For $p \in P$ let $(r^1_p, p), (p^\Delta, r^2_p)$ be the straight line-segments that have just been removed. For every $p \in P$, connect r^1_p, r^2_p by a straight-line segment (see Figure 3.4).

In simple words, for every gap in P, we connect two points, at which the corresponding edge is torn, by a straight-line segment. But first, we remove the drawing from thin strips to the left of g_e and to the right of g_e^Δ so that the arising straight-line segments have are not horizontal.

The last thing to carry out is to draw (v, w) with a straight-line segment $\overline{\Gamma(v)^\Delta\Gamma(w)}$. By the choice of Δ, this line does not cross any other edge. Let Γ^* be the drawing after this iteration. The remainder of the chapter deals with proving the correctness of the algorithm.

3.2 Correctness

Here, we show that Γ^* satisfies the aforementioned invariants:

1. The drawing is planar.
2. The drawing is upward.
3. Every edge is drawn as a polyline.
4. Every NE-edge is drawn as an NE-line.
5. Every NW-edge whose iteration has already been carried out is drawn as an NW-line.
6. No two vertices are drawn with the same x-coordinate.

“The drawing is planar” because every step of the iteration respects the planarity.

“The drawing is upward” as the drawing Γ is upward and every step of the iteration replaces a straight-line segment with a y-monotone polyline.

Every edge in Γ is a polyline and hence, there is only a finite number of crossings of g_e with edges. As a result of the iteration, the total number of straight-line segments in the drawing increases by at most two per such a crossing and hence, stays finite. Hereby, the third invariant is satisfied.

The edges that have not been torn satisfy the fourth and the fifth invariants as these edges either have not been changed at all or have been shifted on Δ to the right and hence, their monotonicity has not been changed. It is also clear that the edges incident neither to v nor w also satisfy the fourth and the fifth invariants: the iteration has only replaced some straight-line segments with positive or negative slopes by straight-line segments with positive or negative slopes respectively.

Now, consider the edges incident to v or w in more detail. If an edge has not been torn on $\Gamma(v)$ or $\Gamma(w)$ and was only shifted, then the above argument can be applied. Consider the edges that have been torn on $\Gamma(v)$. These are the edges whose representation in Γ crossed H_v to the left of g_e. Distinguish the cases:

1. If f is an incoming edge, then the last straight-line segment of $\Gamma(f)$ has a positive slope because it lies to the left of a vertical part of g_e. A part of this segment was replaced with an x-monotone polyline that consists of two straight-line segments, both having positive slopes.
2. If f is an outgoing edge, then it is an NW-edge as it lies to the left of NW-edge e:
   a) If the first straight-line segment of $\Gamma(f)$ has a positive slope, then a part of this segment was replaced with an x-monotone polyline that consists of two straight-line segments, both having positive slopes and the x-monotonicity (if it existed) is maintained.
   b) If the first straight-line segment of $\Gamma(f)$ has a positive slope, then $\Gamma(f)$ is not an NW-line. Hence, f is an NW-edge that has not been repaired yet because Γ satisfies the fifth invariant.

Hence, the edges torn at $\Gamma(v)$ also satisfy the fourth and the fifth invariant. An analogous argument is applied to edges torn at $\Gamma(w)$.

The sixth invariant is also satisfied: let $r \neq s$ be two arbitrary vertices. Γ satisfies this invariant and hence, $\Gamma(r) \neq \Gamma(s)$. Distinguish two cases depending on the positions of these vertices in Γ:

1. If both vertices lie on the same side of g_e, then both vertices are shifted on Δ or both are not. As a result, the x-coordinates stay distinct.
2. If one of the vertices lies to the left of g_e and the other one to the right of g_e, then the second component of the maximization in Equation 3.1 guarantees that such two vertices are not drawn with the same x-coordinate.

As a result, Γ^* satisfies all invariants.

Hereby, the drawing arising after carrying out the iterations corresponding to every NW-edge is a windrose planar drawing of the digraph.

Theorem 3.4. *Let G be an instance of* Windrose Planarity *and let $\mathcal{E}$ be an candidate upward embedding of G. There exists a windrose planar drawing of $\mathcal{E}$ if and only if $\mathcal{E}$ is upward and $(G, \mathcal{E})$ satisfies the condition stated in Observation 3.3.*

4. Face-Local Observations

In the following two chapters, we consider an arbitrary upward-rightward-prescribed drawing and make some observations concerning it. There are two types of observations we make: vertex-local ones and face-local ones. Vertex-local observations describe the interaction of edges incident to the same vertex. Face-local observations tell how an r-edge can be drawn depending on the face of G_u in which it lies.

First, we provide these observations in natural language and then, in Chapter 6, we will represent them more formally as a 2-SAT instance. If this instance is not satisfiable, then no upward-rightward-prescribed drawing exists. In other words, the satisfiability of these constraints is a necessary condition for the existence of an upward-rightward-prescribed drawing. Later, in Chapter 7, we will show that this condition is also sufficient.

In Chapter 7, we will transform an instance of the UPWARD-RIGHTWARD-PRESCRIBED PLANARITY problem into an instance of WINDROSE PLANARITY. For this purpose, we will replace r-edges by gadgets comprised of u-edges. In this chapter, we provide the intuition behind the choice of the gadget-replacement model.

First, we concentrate on instances of UPWARD-RIGHTWARD-PRESCRIBED PLANARITY with a fixed linearized embedding. In Chapter 8, we provide the generalization of the approach to a fixed upward embedding of G_u.

Consider an arbitrary upward-rightward-prescribed drawing Γ of a digraph G. Recall that every edge is drawn as a polyline. We assume that none of the r-edges contains a horizontal straight-line segment: if a drawing includes a horizontal segment $\overline{AB}$, it can be replaced by two straight-line segments $\overline{AC}$, $\overline{CB}$ where $C = ((A_x + B_x)/2, A_y + \epsilon)$ for a small ϵ. The arising drawing ist still an upward-rightward-prescribed planar drawing of G.

For every r-edge e, $\Gamma(e)$ has a finite number of local extrema. This means $\Gamma(e)$ is a concatenation of a finite number of alternating NE- and NW-lines as shown in Figure 4.1.

The remainder of the chapter deals with proving that there is an upward-rightward-prescribed drawing of G in which every edge is drawn with at most three local extrema.

Let $\tilde{e}$ be an r-edge so that $\Gamma(\tilde{e})$ has at least four local extrema; if such an edge does not exist, then Γ is the desired drawing.

For an r-edge $e = (v_e, w_e)$, let $[p^e_2, \ldots, p^e_{t_e-1}]$ be the left-to-right-order of local extrema of $\Gamma(e)$ in open interval $(\Gamma(v_e)_x, \Gamma(w_e)_x)$. For each $i \in \{2, \ldots, t_e - 1\}$, we identify p^e_i with a new vertex v^e_i. As a result, e is partitioned. Set $v^e_1 := v_e$ and $v^e_{t_e} := w_e$. Let g_e denote the

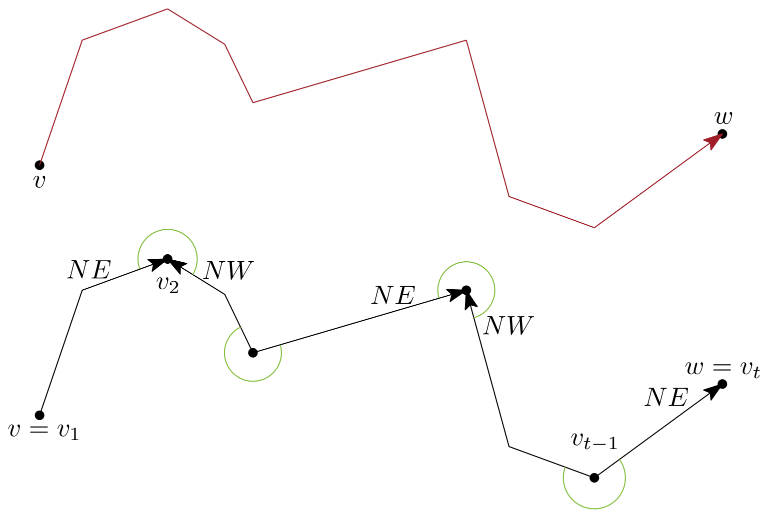

Figure 4.1: Replacing an r-edge with a concatenation of alternating NW- and NE-edges.

set of directed edges $\{\{v_i^e, v_{i+1}^e\} \mid i \in \{1, \dots, t_e - 1\}\}$ so that every edge is directed "from the local minimum to the local maximum". Then, the drawing Γ is an upward drawing of a digraph $\tilde{G} = G + \{g_e \mid e \in E_r\}$. Let $\tilde{\mathcal{E}}$ denote the corresponding upward embedding.

For readability purposes, v_i denotes $v_i^{\tilde{e}}$ and t denotes $t_{\tilde{e}}$. Let f_1, f_2 be the faces of G incident to $g = g_{\tilde{e}}$ in $\tilde{G}$. Observe that for two switches v_i, v_{i+1} ($i \in \{2, \dots, t-2\}$), one is assigned to f_1 and another one is assigned to f_2 (see Figure 4.2).

Transform $\tilde{G}$ and $\tilde{\mathcal{E}}$ into a new digraph G' and a candidate upward embedding $\mathcal{E}'$ of G' as follows:

1. Set $\tilde{G} := G$, $\tilde{\mathcal{E}} := \mathcal{E}$.
2. Remove edges $e_1 := \{v_1, v_2\}, \{v_2, v_3\}, e_2 := \{v_3, v_4\}$.
3. Insert edge $e^* := \{v_1, v_4\}$ which is directed as follows:
    a) If $\Gamma(v_1)$ is a local minimum of $c_{\tilde{e}}$, then e^* is directed from v_1 to v_4.
    b) Otherwise, e^* is directed from v_4 to v_1.
4. Replace e_1 (e_2) by an edge e^* in the left-to-right order of outgoing or incoming edges incident to v_1 (v_4). Hereby, the angle assignment is adopted from $\tilde{\mathcal{E}}$.
5. Remove vertices v_2, v_3.

In simple words, we replaced the path $[v_1, v_2, v_3, v_4]$ by edge $e^* == \{v_1, v_4\}$ and therefore, eliminated two local extrema of $\tilde{e}$ that corresponded to vertices v_2 and v_3. First, we show that $\mathcal{E}'$ is an upward embedding of G'.

Embedding $\tilde{\mathcal{E}}$ is a candidate upward embedding: the direction of e^* is chosen with regard to the directions of e_1 and e_2 so that an incoming edge is replaced by an incoming edge and an outgoing edge is replaced by an outgoing edge. It is also clear that every sink and every source stays assigned to exactly one face because the angle assignment is adopted from $\mathcal{E}$.

The last thing to consider is the condition of Lemma 2.2. The only faces to check are the faces f_1', f_2' of $\mathcal{E}'$ incident to e^* (see Figure 4.2) as other faces are also upward faces of $\mathcal{E}$. Face f_1' (f_2') arises from upward face f_1 (f_2) by removing two switches and one large angle and hence, it also satisfies the condition stated in this lemma. As a result, $\mathcal{E}'$ is an upward embedding of G'.

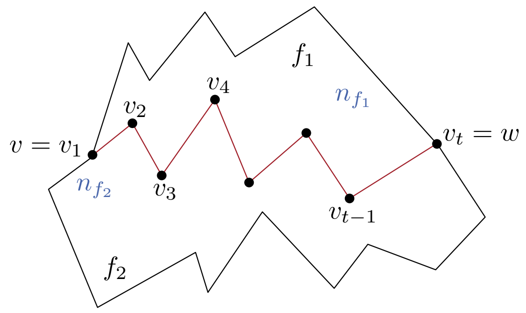

Figure 4.2: Redrawing an r-edge $\tilde{e}$ so that it has two local extrema less.

So far, we have only shown that G' is upward. This doesn't guarantee that the edges corresponding to the r-edges of G can be drawn "rightward". To show this, we construct an instance of WINDROSE PLANARITY. This is conducted by replacing every edge of G' with at most two labeled edges as follows.

For every u-edge $e = (x, y)$ of G, let Q_x (Q_y) be the quadrant around $\Gamma(x)$ ($\Gamma(y)$) in which e lies:

1. If $Q_x = NE$ and $Q_y = SW$, then replace e with an NE-edge (x, y).
2. If $Q_x = NW$ and $Q_y = SE$, then replace e with an NW-edge (x, y).
3. If $Q_x = NW$ and $Q_y = SW$, then add a new vertex v_e and replace e with an NW-edge (x, v_e) and an NE-edge (v_e, y).
4. If $Q_x = NE$ and $Q_y = SE$, then add a new vertex v_e and replace e with an NE-edge (x, v_e) and an NW-edge (v_e, y).

Other constellations are not possible because $\Gamma_{|G_u}$ is upward.

For every r-edge e of G, for every $i, j \in \{1, \ldots, t_e\}$ so that $(v_i^e, v_j^e) \in E_{G'}$:

1. If $i < j$, then replace (v_i^e, v_j^e) with a NE-edge (v_i^e, v_{i+1}^e).
2. If $i > j$, then replace (v_i^e, v_j^e) with a NW-edge (v_i^e, v_{i+1}^e).

For every edge (v_i, v_j) of G':

1. $\Gamma(v_i)$ is a local minimum of $\Gamma(e)$: replace (v_i, v_j) with an NW-edge (v_i, v_j)
2. $\Gamma(v_i)$ is a local maximum of $\Gamma(e)$: replace (v_i, v_j) with an NW-edge (v_j, v_i)

Let $G^w = (V^w, E_r^w \dot\cup E_u^w)$ be the arising instance of the WINDROSE PLANARITY problem. The digraph G^w is upward as it arises by partitioning some edges of the upward digraph G'. By the construction, the above transformation yields the labeling of edges which satisfies the condition stated in Theorem 3.4. As a result, there exists a windrose drawing Γ^w in which every edge of G^w is drawn with at most two straight-line segments [ALB$^+$18]. Observe that by the construction of G^w and Γ^w:

1. Every u-edge of G corresponds to a concatenation of at most two polylines that monotonically increase in the vertical direction in Γ^w. Hereby, this concatenation is also a y-monotone polyline.
2. Every r-edge e other than $\tilde{e}$ corresponds to $t_e - 1$ alternating (as shown in Figure 4.1) NW- and NE-lines, which means its representation monotonically increases in the horizontal direction and has t_e local extrema.

3. The r-edge $\tilde{e}$ corresponds to $t-3$ alternating NW- and NE-lines, which means its representation also monotonically increases in the horizontal direction and has $t-2$ local extrema.

It's noteworthy that this transformation into an instance of Windrose Planarity is only carried out for this proof; the transformation into an instance of Windrose Planarity, which yields an upward-rightward-prescribed drawing of the initial graph, is first made in Chapter 7.

Now, we have shown that there exists an upward-rightward-prescribed drawing of G in which $\tilde{e}$ has two local extrema less than in Γ and all other r-edges have the same number of local extrema as they have in Γ.

This proves the following lemma.

Lemma 4.1. *Let Γ be an upward-rightward-prescribed drawing G and let $\tilde{e}$ be an arbitrary r-edge (if it exists). For every r-edge e let u_e^Γ denote the number of local extrema of $\Gamma(e)$. If $u_{\tilde{e}}^\Gamma \geq 4$, then there exists an upward-rightward-prescribed drawing Γ' of G such that:*

1. *for every $e \in E_r \setminus \{\tilde{e}\}$: $u_e^{\Gamma'} = u_e^\Gamma$*
2. *$u_{\tilde{e}}^{\Gamma'} = u_{\tilde{e}}^\Gamma - 2$*

Repeat applying this lemma until the representation of every r-edge has at most three local extrema.

Lemma 4.2. *Let Γ be an upward-rightward-prescribed drawing G. For every r-edge e let u_e^Γ denote the number of local extrema of $\Gamma(e)$. Then, there exists an upward-rightward-prescribed drawing Γ' of G such that: for every $e \in E_r$, $u_e^{\Gamma'} \leq 3$.*

The above lemma results in an important lemma about the set of gadgets we have to consider.

Further, *gadget* stands for a set of u-edges. For a gadget g, we call $G_u + g$ upward *with respect to* an upward embedding $\mathcal{E}_u$ of G_u if there exists an upward embedding $\mathcal{E}_u'$ of $G_u + g$ whose restriction to G_u coincides with $\mathcal{E}_u$.

Lemma 4.3. *Let $G = (V, E_u \dot\cup E_r)$ be an upward-rightward-prescribed planar digraph with underlying upward embedding $\mathcal{E}_u$ of G_u. Then, for every r-edge $e = (v, w) \in E_r$, one of the following digraphs is upward with respect to $\mathcal{E}_u$:*

1. *$G + (v, w)$*
2. *$G + (w, v)$*
3. *$G + (v, q_e) + (w, q_e)$, where v_e is a new vertex and a large angle at q_e clockwise between (v, q_e) and (w, q_e) is prescribed*
4. *$G + (q_e, v) + (q_e, w)$, where v_e is a new vertex and a large angle at q_e clockwise between (q_e, w) and (q_e, v) is prescribed*

As a result, we only have to consider four types of gadgets to replace an r-edge. The numeration used in the above lemma is further used to distinguish between different types of gadgets (see Figure 4.3). At this point, we are not interested in labels (NE or NW) which are denoted blue in the figure; they will be introduced and used in Chapter 7.

Further, writing about one of the digraphs listed in Lemma 4.3, we implicitly mean an upward embedding that respects $\mathcal{E}_u$. A gadget g corresponding to (v, w) is called *valid* if $G + g$ is upward.

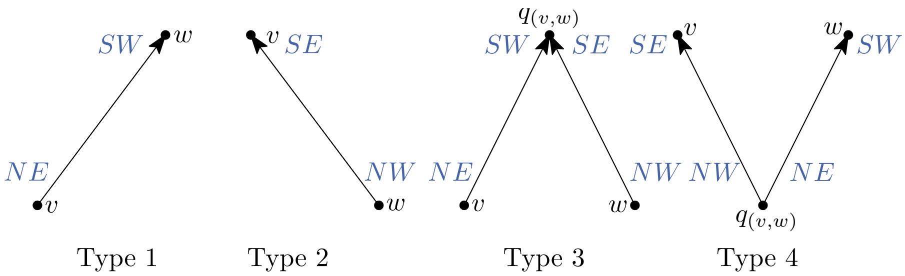

Figure 4.3: Gadgets of all four types for an r-edge (v, w)

As mentioned earlier, in Chapter 7, we will show that the satisfiability of constraints based on the face- and vertex-local observations provided in this chapter and the next one is a criterion of upward-rightward-prescribed planarity of graphs with fixed linearized embedding. To conclude this chapter, we look at only so-called “face-local” constraints, which guarantee that for every r-edge the condition in Lemma 4.3 is fulfilled, and justify that their satisfiability is not sufficient for the existence of an upward-rightward-prescribed drawing.

Consider graph G in Figure 4.4. Notice that for each of three r-edges, the gadget of the first type is valid and hereby, each r-edge can individually be embedded into the upward part G_u. However, the r-edges form a cycle and hence, the graph is not upward-rightward-prescribed planar even though that face-local constraints are satisfiable. Thus, we indeed have to consider the so-called “vertex-local” observations provided in the next chapter.

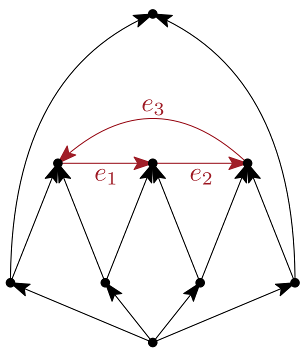

Figure 4.4: Although every r-edge can be induvidually inserted into the upward part G_u, graph G is not upward-rightward-prescribed planar, as the r-edges form a cycle.

5. Vertex-Local Observations

In this chapter, we continue considering an arbitrary upward-rightward-prescribed drawing. The observations we make here are called vertex-local as we concentrate on edges incident to the same vertex and see how they are located around this vertex.

Let Γ be an arbitrary upward-rightward-prescribed drawing of a digraph G. Every edge is drawn as a polyline. In the previous chapter, we have shown that Γ can be easily adapted so that no edge includes a horizontal straight-line segment so further, we assume that every straight-line segment has a non-zero slope. For a similar reason, we can also assume that none of the edges includes a vertical segment. Hereby, every vertex v has a neighborhood $U_\epsilon(\Gamma(v))$ in which every edge is xy-monotone and hence, the assignment of edges to the quadrants is well-defined (notice that such an assignment is specific to the drawing and can be different for the same digraph).

Here, we only consider edges incident to v; for this reason, writing about "edges lying in the X-quadrant" or simply "X-edges", we mean edges whose representation in $U_\epsilon(\Gamma(v))$ lies in the X-quadrant around v (for $X \in \{NW, NE, SW, SE\}$).

Observation 5.1. *Let Γ be an arbitrary upward-rightward-prescribed drawing of a digraph G and $v \in V_G$ be a vertex. Consider the location of edges incident to v around this vertex:*

1. *Every outgoing u-edge lies in one of the North-quadrants.*
2. *Every incoming u-edge lies in one of the South-quadrants.*
3. *Every outgoing r-edge lies in one of the East-quadrants.*
4. *Every incoming r-edge lies in one of the West-quadrants.*

Recall that we consider instances of Upward-Rightward-Prescribed Planarity with a fixed linearized embedding. Let l_v be the linearized order of edges incident to a vertex v. The NW-edges are first to appear in l_v, followed by the NE-, then the SE- and finally, the SW-edges. This yields a natural order $[NW, NE, SE, SW]$ of the quadrants. Further, writing about the quadrants "following" or "preceding" other quadrants, we refer to this order.

Observation 5.2. *Let $G = (V, E)$ be a planar digraph, let Γ be a planar drawing of G and let $\mathcal{E}$ be the linearized embedding defined by Γ. For an arbitrary vertex v of G, let l be the linearized order of edges incident to v in $\mathcal{E}$. The list l can be partitioned into four*

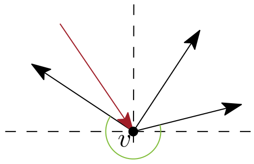

Figure 5.1: There can exist a large angle around a vertex v, although v is neither a source nor a sink. The large angle is depicted green.

(possibly empty) lists $l = l_{NW} \cdot l_{NE} \cdot l_{SE} \cdot l_{SW}$ so that l_X only contains edges that lie in the X-quadrant around v (for $X \in \{NW, NE, SE, SW\}$). In particular, an edge e_1 that precedes and edge e_2 in l does not lie in a quadrant that follows the quadrant in which e_2 lies.

We call an angle around a vertex v *large* if both North- or both South-quadrants lie entirely within this angle. Notice that there can be a large angle around a vertex that is neither a source nor a sink (see Figure 5.1). Hereby, this definition is used in a different way than within the Upward Planarity problem setting. Later, we will see why we still use the same term in different contexts. In the following figures, large angles are depicted green.

6. 2-Sat Representation

In Chapters 4 and 5, we have made important face- and vertex-local observations stated in Lemma 4.2 and Observation 5.2 about an arbitrary upward-rightward-prescribed drawing of a digraph and the assignment of edges to the quadrants in this drawing. We look for a quadrant assignment so that the conditions stated in these observations are satisfied. As every upward-rightward-prescribed drawing satisfies these conditions, if there is no such assignment, then there is also no upward-rightward-prescribed drawing. Later on, in Chapter 7, we will also show the reverse: if there exists such an assignment, then it is possible to find an upward-rightward-prescribed drawing.

Further, for every edge and both of its end-vertices, we represent possible quadrants as boolean variables and represent the dependencies stated in the previous two chapters as an instance of the 2Sat problem. Hence, the satisfiability of the 2Sat-instance constructed in this chapter is necessary for the existence of an upward-rightward-prescribed drawing.

6.1 Trivial Constraints

For each edge $e = (v, w)$, we define boolean variables $v_e^{NW}, v_e^{NE}, v_e^{SE}, v_e^{SW}, w_e^{NW}, w_e^{NE}, w_e^{SE}, w_e^{SW}$. The values assigned to these variables are interpreted as follows: for example, if v_e^{NE} is assigned a *true*-value, then, in a small neighbourhood of v, e lies in the NE-quadrant around v.

First, set the constraints which guarantee that per end-vertex, every edge is assigned exactly one quadrant and this quadrant complies with the type (u or r) of the edge (see Observation 5.1).

For every r-edge $e = (v, w)$ set:

1. $\neg v_e^{NW}$, $\neg v_e^{SW}$ and $v_e^{NE} \nleftrightarrow v_e^{SE}$
2. $\neg w_e^{NE}$, $\neg w_e^{SE}$ and $w_e^{NW} \nleftrightarrow w_e^{SW}$

Similarly, for every u-edge $e = (v, w)$ set:

1. $\neg v_e^{SW}$, $\neg v_e^{SE}$ and $v_e^{NE} \nleftrightarrow v_e^{NW}$
2. $\neg w_e^{NE}$, $\neg w_e^{NW}$ and $w_e^{SE} \nleftrightarrow w_e^{SW}$

6.2 Linearized Embedding

Let v be a vertex and $[e_1, \ldots, e_t]$ be the linearized order of edges incident to v. For every $i \in \{1, \ldots, t-1\}$, set the following constraints (see Observation 5.2):

- $v_{e_i}^{NE} \to \neg v_{e_{i+1}}^{NW}$
- $v_{e_i}^{SE} \to \neg v_{e_{i+1}}^{NW}$, $v_{e_i}^{SE} \to \neg v_{e_{i+1}}^{NE}$
- $v_{e_i}^{SW} \to \neg v_{e_{i+1}}^{NW}$, $v_{e_i}^{SW} \to \neg v_{e_{i+1}}^{NE}$, $v_{e_i}^{SW} \to \neg v_{e_{i+1}}^{SE}$ (or shorter: $v_{e_i}^{SW} \to v_{e_{i+1}}^{SW}$)

Notice that this representation of the dependencies is only possible as the linearized embedding is prescribed.

6.3 Face-Local Constraints

The constraints set in this section represent the property stated in Lemma 4.3.

We recall that a linearized embedding of G prescribes in particular an upward embedding $\mathcal{E}_u$ of G_u and consequentially, also an angle assignment of G_u.

Further on, $\mathcal{H}^i_{(v,w)}$ denotes the gadget of type i ($i \in \{1, 2, 3, 4\}$, see Figure 4.3) that corresponds to an r-edge (v, w).

Let $e = (v, w)$ be an r-edge and let f be the face of G_u into which e is to be embedded. Let $S_e = \{i \in \{1, 2, 3, 4\} \mid G_u + \mathcal{H}^i_e$ is upward with respect to $\mathcal{E}\}$ denote the set of valid gadgets corresponding to the r-edge e. Observe that the linearized order of edges around v and w prescribes the angle assignment of $G_u + \mathcal{H}^i_e$ in a unique way. The set S_e can be easily obtained in linear time: for every gadget, count the number of switches and large angles on the border of faces arising after the insertion of the gadget and check if the condition in Lemma 2.2 is satisfied. Now we describe that this is also feasible in constant time once a linear-time precalculation is done.

6.3.1 Obtaining S_e in $\mathcal{O}(1)$

Here we provide a data structure that allows the computation of S_e in $\mathcal{O}(1)$. The precalculation is done in linear time.

Let f be a face of G_u and let $V_f = [v_1, \ldots, v_{t_f}]$ be the clockwise order of vertices on the border of f. Two linear lists $l^f_s = [s^f_1, \ldots, s^f_{t_f}]$ and $l^f_b = [b^f_1, \ldots, b^f_{t_f}]$ are associated with f. The value s^f_i is the number of vertices in $\{v_1, \ldots, v_i\}$ that are switches of f and the value b^f_i is the number of vertices in $\{v_1, \ldots, v_i\}$ which are sinks or sources of G_u assigned to f.

For a set A, indicator function 1_A is defined as follows:

$$1_A(x) = \begin{cases} 1 & x \in A \\ 0 & otherwise \end{cases}$$

Let S_f and B_f be the set of switches on the border of f and the set of switches assigned to f respectively. Then, the lists l^f_s and l^f_b correspond to the prefix sum (relative to order V_f of vertices) of indicator functions 1_{S_f} and 1_{B_f} respectively. This function can be evaluated in $\mathcal{O}(1)$ because an upward embedding is given. Hence, both lists can be calculated in $\mathcal{O}(t_f)$.

The total running time of the precalculation is $\mathcal{O}(\sum_{f \in \mathcal{F}} t_f) = O(|V| + |E|)$ where $\mathcal{F}$ denotes the set of faces of $\mathcal{E}_u$.

Let (v, w) be an r-edge which is inserted into f. Consider a gadget g which replaces (v, w). Let f_1 and f_2 be the faces arising after the insertion of g. Our aim is to test whether $G_u + g$ is upward with respect to $\mathcal{E}_u$. As previously mentioned, the linearized embedding determines the only angle assignment that comes into question.

Prefix sums make range queries possible, which means for two vertices on the border of the same face, the number of switches or large angles between them can be acquired in constant time. As a result, for a fixed angle assignment of $G_u + g$, the number of large angles and switches on the borders of f_1 and f_2 can also be obtained in constant time. Lemma 2.2 allows to check whether this angle assignment yields an upward embedding or not in $\mathcal{O}(1)$.

Altogether, there are four different gadgets for e and hence, there is only a constant number of constellations to check. As a result, S_e can be obtained in constant time.

In Chapter 8, we omit the prescription of the linearized embedding but we again consider only one angle assignment per gadget and hereby, this data structure can still be used to compute S_e in constant time.

6.3.2 Constraints for $\mathbf{S_e}$

Here, we show how the face-local constraints concerning an r-edge e look like depending on S_e. The whole case distinction is too massive to be provided here, so we only consider two examples in detail.

First, we consider $S_e = \{1\}$: in this case the gadget is fixed and it must be guaranteed that e lies in the NE-quadrant around v and in the SW-quadrant around w. Hence, the corresponding constraints are v_e^{NE}, w_e^{SW}.

Now we consider $S_e = \{1, 2, 3\}$: in this case, the gadget $\mathcal{H}_e^4$ is prohibited and it has to be guaranteed that the pair of quadrants around its end-vertices which corresponds to this gadget is ruled out: $\neg(v_e^{SE} \wedge w_e^{SW}) = \neg v_e^{SE} \vee \neg w_e^{SW} = v_e^{NE} \vee w_e^{NW}$.

The remaining cases are considered in a similar way. Depending on the content of S_e the constraints are set as follows:

1. $S_e = \emptyset$: in this case, there exists no upward-rightward-prescribed drawing, so we can simply stop or add an unsatisfiable set of constraints, for example $x, \neg x$
2. $S_e = \{1\}$: v_e^{NE}, w_e^{SW}
3. $S_e = \{2\}$: v_e^{SE}, w_e^{NW}
4. $S_e = \{3\}$: v_e^{NE}, w_e^{NW}
5. $S_e = \{4\}$: v_e^{SE}, w_e^{SW}
6. $S_e = \{1, 2\}$: $v_e^{NE} \leftrightarrow w_e^{SW}$
7. $S_e = \{1, 3\}$: v_e^{NE}
8. $S_e = \{1, 4\}$: w_e^{SW}
9. $S_e = \{2, 3\}$: w_e^{NW}
10. $S_e = \{2, 4\}$: v_e^{SE}
11. $S_e = \{3, 4\}$: $v_e^{NE} \leftrightarrow w_e^{NW}$
12. $S_e = \{2, 3, 4\}$: $v_e^{SE} \vee w_e^{NW}$
13. $S_e = \{1, 3, 4\}$: $v_e^{NE} \vee w_e^{SW}$

14. $S_e = \{1, 2, 4\}$: $v_e^{SE} \vee w_e^{SW}$
15. $S_e = \{1, 2, 3\}$: $v_e^{NE} \vee w_e^{NW}$
16. $S_e = \{1, 2, 3, 4\}$: no additional constraints are needed because every gadget is admissible.

The 2-Sat instance constructed as described in this chapter is denoted by $\mathcal{I}_G^{\mathcal{E}}$.

7. From Upward-Rightward-Prescribed Planarity to Windrose Planarity

In the previous chapter, we provided the construction of 2-SAT instance $\mathcal{I}_G^{\mathcal{E}}$ whose satisfiability is necessary for the existence of an upward-rightward-prescribed drawing of digraph G with underlying linearized embedding $\mathcal{E}$. In this chapter, we prove that this condition is also sufficient.

First, given a solution $\mathcal{S}$ for this instance, we describe the construction of the corresponding instance of WINDROSE PLANARITY. After that, we show that this instance indeed has a windrose drawing and finally, see that this drawing immediately yields an upward-rightward-prescribed drawing of the initial graph.

The main idea of the construction is the following: $\mathcal{S}$ prescribes the assignment of edges to the quadrants. This assignment provides the main structure of the corresponding windrose labeling. But notice that we can not simply use $\mathcal{S}$ as the labeling as there can still be contradictions: for example, there can exist such an u-edge $e = (v, w)$ that $\mathcal{S}$ assigns e to the NW-quadrant around v and the SW-quadrant around w (see Figure 7.1 on the left). In this form, it would be a no-instance of the WINDROSE PLANARITY problem so we have to partition contradictive u-edges (see Figure 7.1 on the right). Another thing to accomplish is the replacement of the r-edges by gadgets listed in Lemma 4.3 in accordance with $\mathcal{S}$.

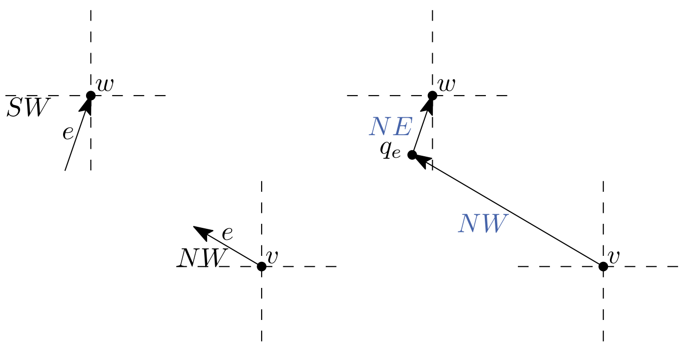

Figure 7.1: The quadrant assignment $\mathcal{S}$ can be inconsistent in terms of the Windrose Planarity. This problem can be resolved by partitioning inconsistent labeled edges. The labeling of the arising edges is depicted blue.

We define an instance $\hat{G}$ of Windrose Planarity that consists of a digraph, labeling and an embedding as follows. Initially, the set of vertices is V_G and the set of edges is empty.

For every u-edge $e = (v, w)$, the labeling of the corresponding edges in $\hat{G}$ depends on the quadrants around v and w assigned to e by $\mathcal{S}$ as follows:

1. If e is assigned to the NW-quadrant around v and the SE-quadrant around w, then add an NW-edge (v, w).
2. If e is assigned to the NE-quadrant around v and the SW-quadrant around w, then add an NE-edge (v, w).
3. If e is assigned to the NW-quadrant around v and the SW-quadrant around w, then add a new vertex q_e, an NW-edge $e_1 = (v, q_e)$, and an NE-edge $e_2 = (q_e, w)$ (see Figure 7.1).
4. If e is assigned to the NE-quadrant around v and the SE-quadrant around w, then add a new vertex q_e, an NE-edge $e_1 = (v, q_e)$, and an NW-edge $e_2 = (q_e, w)$.

For every r-edge $e = (v, w)$, the gadget choice is determined by $\mathcal{S}$ as follows:

1. If e is assigned to the NE-quadrant around v and the SW-quadrant around w, then add an NE-edge (v, w) ($\mathcal{H}_e^1$).
2. If e is assigned to the SE-quadrant around v and the NW-quadrant around w, then add an NW-edge (w, v) ($\mathcal{H}_e^2$).
3. If e is assigned to the NE-quadrant around v and the NW-quadrant around w, then add a new vertex q_e, an NE-edge $e_1 = (v, q_e)$ and an NW-edge $e_2 = (w, q_e)$; the left-to-right order of incoming edges incident to q_e is $[e_1, e_2]$ ($\mathcal{H}_e^3$).
4. If e is assigned to the SE-quadrant around v and the SW-quadrant around w, then add a new vertex q_e, an NW-edge $e_1 = (q_e, v)$, and an NE-edge $e_2 = (q_e, w)$; the left-to-right order of incoming edges incident to q_e is $[e_1, e_2]$ ($\mathcal{H}_e^4$).

The constraints set according to Section 6.1 guarantee that there are no possible constellations other than the ones listed above and hence, every edge $e \in E_G$ is represented by a gadget in $\hat{G}$.

The constraints set according to Section 6.2 guarantee that the incoming and outgoing edges incident to the same vertex of $\hat{G}$ are consecutive. The left-to-right order of edges incident to $v \in V_G$ in $\hat{G}$ is extracted from the linearized order l_v of edges around v in G and the quadrant assignment $\mathcal{S}$ as follows. The edges of G assigned to the North-quadrants around v correspond to outgoing edges in $\hat{G}$ and the order in which they appear in l_v corresponds to the left-to-right order in $\hat{G}$. Similarly, the edges of G assigned to the South-quadrants around v correspond to incoming edges in $\hat{G}$ and the order in which they appear in l_v corresponds to the right-to-left order in $\hat{G}$. These left-to-right orders uniquely define the angle assignment $\mathcal{A}$ of $\hat{G}$. Notice that every sink and every source is assigned to exactly one face by $\mathcal{A}$ (to be exact, to the face identified by the leftmost and the rightmost edge incident to this sink or source).

7.1 Existence of a Windrose Drawing of $\hat{\mathrm{G}}$

This section deals with proving that $\hat{G}$ satisfies the conditions stated in Theorem 3.4 and hence, has a windrose drawing.

The constraints set according to Section 6.2 guarantee that the order of edges around every vertex $v \in V_G$ in $\hat{G}$ satisfies the condition stated in the Observation 3.3. Obviously, every vertex q_e also satisfies it. The only thing still to prove is that the embedding of $\hat{G}$ is upward. For this purpose, we show that $\mathcal{A}$ satisfies the condition stated in Lemma 2.2.

Observe how $\hat{G}_u$ (the part of $\hat{G}$ that arose from G_u) was created:

1. The constraints set according to Section 6.1 guarantee that every outgoing (incoming) u-edge in G is assigned to one of the North (South)-quadrants in $\mathcal{S}$ and hence, is transformed into an outgoing (incoming) edge in $\hat{G}$.
2. Some of these edges were partitioned.

Let $\mathcal{A}_u$ denote the restriction of $\mathcal{A}$ to $\hat{G}_u$. The digraph $\hat{G}_u$ arises from G_u by partitioning some edges and hence, $\hat{G}_u$ is upward. Observe, that the angle assignment of G_u induced by $\mathcal{E}_u$ coincides with $\mathcal{A}_u$; this allows us to write about inserting gadgets into G_u and $\hat{G}_u$ interchangeably. Hereby, if we show that the upward planarity of G_u is maintained throughout the insertion of the gadgets chosen according to $\mathcal{S}$, then $\hat{G}$ is also upward planar.

Further on, writing about gadgets we implicitly mean the gadgets representing r-edges of G in $\hat{G}$ and hence, chosen in accordance with $\mathcal{S}$. In particular, every gadget is of one of the four types listed in Lemma 4.3. Recall that face-local constraints set according to Section 6.3 guarantee that every gadget is valid and hence, can individually be inserted into G_u so that the embedding stays upward. We still have to prove that all gadgets can be simultaneously inserted into G_u so that the embedding also stays upward. First, we show that two different gadgets can be simultaneously inserted into G_u and then, prove by induction on the number of gadgets that this also holds for all gadgets.

If two gadgets g_1, g_2 are inserted into different faces of G_u, then after insertion of both of them, the embedding stays upward as every upward face of the arising graph is an upward face of $G_u + g_1$ or $G_u + g_2$.

For a subgraph H of $\hat{G}$, let $\mathcal{A}_H$ denote the restriction of $\mathcal{A}$ to H. For readability purposes, $\mathcal{A}_g$ stands for $\mathcal{A}_{G_u+g}$ and $\mathcal{A}_{g_1,g_2}$ stands for $\mathcal{A}_{G_u+g_1+g_2}$ where g, g_1, g_2 are gadgets.

Now, we show that the insertion of two gadgets into the same face of G_u also maintains the upward planarity. Every gadget corresponds to an r-edge and there are no multi-edges in G. Hence, two different gadgets share at most one vertex. We distinguish two cases, that is, whether or not two gadgets share a vertex.

7.1.1 Insertion of Two Gadgets with Disjoint Sets of Vertices

First, we consider the simpler case: two gadgets have no common vertices.

Lemma 7.1. *Let f be a face of G_u and let $[u_1, u_k, v_t, v_1]$ be pairwise distinct vertices on the border of f in the clockwise order. Let g_1, g_2 be the gadgets with the disjoint sets of vertices so that:*

1. *The end-vertices of g_1 are u_1 and u_k.*
2. *The end-vertices of g_2 are v_1 and v_t.*
3. *Gadgets g_1, g_2 are inserted into f.*

Then $\mathcal{A}_{g_1,g_2}$ satisfies the condition stated in Lemma 2.2.

Proof. First, assume that f is an inner face. Consider angle assignment $\mathcal{A}_{g_1,g_2}$: notice that its restrictions to G_u, $G_u + g_1$ and $G_u + g_2$ are $\mathcal{A}_u$, $\mathcal{A}_{g_1}$ and $\mathcal{A}_{g_2}$ respectively. We recall that face-local constraints guarantee that g_1 and g_2 are valid and hence, $\mathcal{A}_{g_1}$ and $\mathcal{A}_{g_2}$ are upward-consistent. Now, we show that $\mathcal{A}_{g_1,g_2}$ satisfies the condition stated in Lemma 2.2.

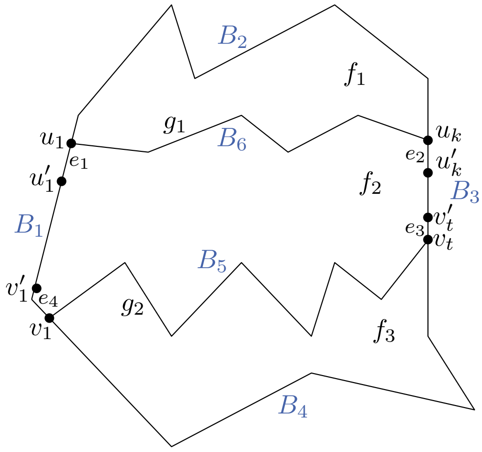

Figure 7.2: Simultaneous insertion of two gadgets that have no common vertices.

Let f_1, f_2, f_3 be the faces arising from f by insertion of g_1 and g_2 (see Figure 7.2).

Angle assignment $\mathcal{A}_{g_1,g_2}$ arises from quadrant assignment $\mathcal{S}$ by extracting the lists of incoming and outgoing edges associated with every vertex. Hence, each source and each sink is assigned to exactly one face. The capacity condition stated in Lemma 2.2 holds for upward faces other than f_1, f_2, f_3 as these are also upward faces of G_u. Face f_1 (f_3) satisfies this condition because it is also an upward face of $G_u + g_1$ ($G_u + g_2$).

Finally, consider face f_2. Let e_1 (e_2, e_3, e_4) be edges incident to u_1 (u_k, v_t, v_1) and lying on the border of f_2 in $G_u + g_1 + g_2$ and on the border of f in G. Let u_1' (u_k', v_t', v_1') be the vertices that arise by partitioning the edges e_1, e_2, e_3, e_4. This partitioning is made to guarantee that u_1' (u_k', v_t', v_1') are not switches and we do not count any switch twice later on. Notice that partitioning the edges changes nothing on upward planarity: for readability purposes and to avoid the definition of new graphs, we assume that these vertices are already contained in G_u (again see Figure 7.2).

The main idea is to apply Lemma 2.2 to different faces of different graphs to show that f_2 also satisfies the condition stated in this lemma. For this purpose, we partition the borders of f_1, f_2, f_3 into the sets of consecutive vertices as follows. Writing about the "set of vertices between v and w", we mean that in particular, the set contains v and w.

1. Let B_1 be the set of vertices incident to f clockwise between v_1' and u_1'.
2. Let B_2 be the set of vertices incident to f clockwise between u_1 and u_k.
3. Let B_3 be the set of vertices incident to f clockwise between u_k' and v_t'.
4. Let B_4 be the set of vertices incident to f clockwise between v_t and v_1.
5. Let B_5 be the set of vertices incident to f_2 clockwise between v_t and v_1.
6. Let B_6 be the set of vertices incident to f_2 clockwise between u_1 and u_k.

Let $f_{1,2}$ ($f_{2,3}$) be the face of $G_u + g_2$ ($G_u + g_1$) which corresponds to the union of the faces f_1 and f_2 (f_2 and f_3).

Let T_1 (T_2, T_3, T_4) be the number of switches in B_1 (B_2, B_3, B_4) assigned to f in $\mathcal{A}_u$. Let T_5 (T_6) be the number of switches in B_5 (B_6) assigned to $f_{1,2}$ ($f_{2,3}$) in $\mathcal{A}_{g_2}$ ($\mathcal{A}_{g_1}$).

Let U_1 (U_2, U_3, U_4) be the number of switches of f in B_1 (B_2, B_3, B_4). Let U_5 (U_6) be the number of switches of f_2 in B_5 (B_6).

Observe that by the construction of these angle assignments:

1. T_1 (T_3) is also the number of switches in B_1 (B_3) assigned to f_2 in $\mathcal{A}_{g_1,g_2}$.
2. T_1 (T_3) is also the number of switches in B_1 (B_3) assigned to $f_{1,2}$ in $\mathcal{A}_{g_2}$.
3. T_1 (T_3) is also the number of switches in B_1 (B_3) assigned to $f_{2,3}$ in $\mathcal{A}_{g_1}$.
4. T_5 (T_6) is also the number of switches in B_5 (B_6) assigned to f_2 in $\mathcal{A}_{g_1,g_2}$.

Similar observations can be made concerning U_i-s.

The number of switches on the border of f is $U_1+U_2+U_3+U_4$ and the number of switches assigned to f in $\mathcal{A}_u$ is $T_1+T_2+T_3+T_4$. The condition stated in Lemma 2.2 for $\mathcal{A}_u$ and f:

$$(U_1 + U_2 + U_3 + U_4)/2 - 1 = T_1 + T_2 + T_3 + T_4 \tag{7.1}$$

The same condition for $\mathcal{A}_{g_1}$ and $f_{2,3}$:

$$(U_1 + U_6 + U_3 + U_4)/2 - 1 = T_1 + T_6 + T_3 + T_4 \tag{7.2}$$

Same for $\mathcal{A}_{g_2}$ and $f_{1,2}$:

$$(U_1 + U_2 + U_3 + U_5)/2 - 1 = T_1 + T_2 + T_3 + T_5 \tag{7.3}$$

Subtracting (7.3) from (7.1) yields:

$$(U_2 - U_6)/2 = T_2 - T_6 \tag{7.4}$$

Subtracting (7.2) from (7.1) yields:

$$(U_4 - U_5)/2 = T_4 - T_5 \tag{7.5}$$

Now consider face f_2. Equations (7.1), (7.4) and (7.5) yield:

$$\begin{aligned} T_1 + T_6 + T_3 + T_5 &= T_1 + T_2 - (U_2 - U_6)/2 + T_3 + T_4 - (U_4 - U_5)/2 \\ &= (U_1 + U_2 + U_3 + U_4)/2 - 1 - U_2/2 + U_6/2 - U_4/2 + U_5/2 \\ &= (U_1 + U_6 + U_3 + U_5)/2 - 1 \end{aligned} \tag{7.6}$$

The number of switches on the border of f_2 is $U_1 + U_6 + U_3 + U_5$ and the number of switches assigned to f_2 in $\mathcal{A}_{g_1,g_2}$ is $T_1 + T_6 + T_3 + T_5$. Therefore, the face f_2 also satisfies the required condition:

$$(U_1 + U_6 + U_3 + U_5)/2 - 1 = T_1 + T_6 + T_3 + T_5 \tag{7.7}$$

If f is the outer face, the proof is analogous. The only difference is that -1 is replaced by $+1$ on the left side of every equation (see Lemma 2.2).

As a result, the angle assignment $\mathcal{A}_{g_1,g_2}$ is upward-consistent.

□

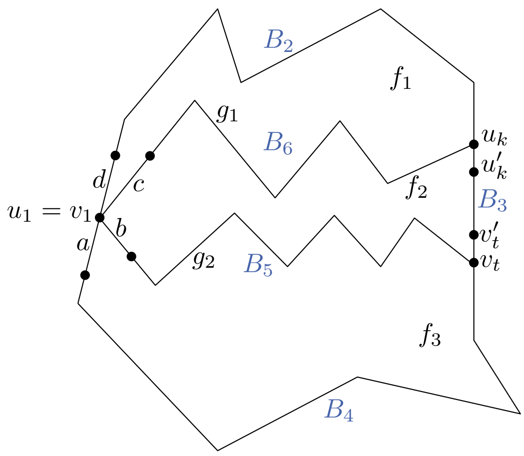

Figure 7.3: The simultaneous insertion of two gadgets that share a vertex.

7.1.2 Insertion of Two Gadgets that Share a Vertex

Here we consider a situation similar to the one from the previous subsection with the only difference that $v_1 = u_1$ (two gadgets share a vertex) and briefly justify that in this case, $\mathcal{A}_{g_1,g_2}$ is also upward-consistent.

The structure of the proof stays the same but in order to set up the equations similar to (7.1), (7.2) and (7.3) we now have to distinguish if v_1 is a switch of each f_1, f_2 and f_3 and if it is assigned to one of these faces. The proof is carried out in the same way for all these cases so we only consider one of them as an example.

The edges a, b, c, d incident to v_1 are defined as shown in Figure 7.3. As an aforementioned example, we consider the constellation where a, b, c are incoming edges and d is an outgoing edge regarding v_1. Hence, the angles around v_1 incident to f, $f_{1,2}$ and $f_{2,3}$, f_2 are small in $\mathcal{A}_u$, $\mathcal{A}_{g_2}$, $\mathcal{A}_{g_1}$ and $\mathcal{A}_{g_1,g_2}$ respectively. It is the same argument as the one used before for justifying that faces other than f_2 satisfy the condition in Lemma 2.2. Now, similarly to the previous subsection, we set up the system of equations to show that f_2 also satisfies this condition.

The B_i-s are defined in a similar way as in the previous chapter with the following difference. Vertex v is excluded from areas B_2, B_4, B_5, B_6: it is considered individually in this subsection. The U_i-s and T_i-s are also defined in the same way as in the previous subsection according to the new B_i-s.

We assume that f is an inner face. For the outer face the proof is analogous.

Vertex v_1 is not a switch of f, so the application of Lemma 2.2 on the face f in $\mathcal{A}_u$ yields:

$$T_2 + T_3 + T_4 = (U_2 + U_3 + U_4)/2 - 1 \tag{7.8}$$

Same for $f_{1,2}$ in $\mathcal{A}_{g_2}$:

$$T_2 + T_3 + T_5 = (U_2 + U_3 + U_5)/2 - 1 \tag{7.9}$$

Notice that v_1 is a switch of $f_{2,3}$ with a small angle. The condition in Lemma 2.2 for $f_{2,3}$ in $\mathcal{A}_{g_1}$:

$$T_3 + T_4 + T_6 = (U_3 + U_4 + U_6 + 1)/2 - 1 \tag{7.10}$$

Subtracting (7.8) from (7.9) yields:

$$T_5 - T_4 = (U_5 - U_4)/2 \tag{7.11}$$

Subtracting (7.8) from (7.10) yields:

$$T_6 - T_2 = (U_6 + 1 - U_2)/2 \tag{7.12}$$

Now consider the face f_2; similarly to the previous subsection, equations (7.8), (7.11) and (7.12) yield:

$$\begin{aligned} T_3 + T_5 + T_6 &= T_3 + (T_4 + (U_5 - U_4)/2) + (T_2 + (U_6 + 1 - U_2)/2) \\ &= (T_2 + T_3 + T_4) + (U_5 - U_4)/2 + (U_6 + 1 - U_2)/2 \\ &= (U_2 + U_3 + U_4)/2 - 1 + (U_5 - U_4)/2 + (U_6 + 1 - U_2)/2 \\ &= (U_3 + U_5 + U_6)/2 - 1/2 = (U_3 + U_5 + U_6 + 1)/2 - 1 \end{aligned} \tag{7.13}$$

Notice that v_1 is a switch of f_2 with a small angle and therefore, the face f_2 also satisfies the required condition:

$$T_3 + T_5 + T_6 = (U_3 + U_5 + U_6 + 1)/2 - 1 \tag{7.14}$$

As a result, the angle assignment $\mathcal{A}_{g_1,g_2}$ is upward-consistent.

To consider the remaining constellations, equations (7.8), (7.9) and (7.10) are adjusted so that they comply with the directions of edges a, b, c, d (for example, $+1$ on the right side of equation (7.10) corresponds to the fact that v_1 is a switch of $f_{2,3}$). Some constellations are symmetrical so the whole proof can be realized as the distinction of 7 cases. As previously stated, the remaining cases are omitted here as they do not have much sense by themselves.

This results in the following lemma:

Lemma 7.2. *Let $[u_1 = v_1, u_k, v_t]$ be pairwise distinct vertices on the border of f in the clockwise order. Let g_1, g_2 be two gadgets so that:*

1. *The end-vertices of g_1 are u_1 and u_t.*
2. *The end-vertices of g_2 are v_1 and v_t.*
3. *Gadgets g_1, g_2 are inserted into f.*

Then $\mathcal{A}_{g_1,g_2}$ satisfies the condition stated in Lemma 2.2.

Lemmas 7.1 and 7.2 result in the following lemma.

Lemma 7.3. *Let g_1, g_2 be different gadgets chosen according to $\mathcal{S}$. Then, $\mathcal{A}_{g_1,g_2}$ is upward-consistent.*

7.1.3 Upward Planarity of $\hat{\mathbf{G}}$

In this subsection, we prove that the upward planarity of G_u is maintained throughout the insertion of the gadgets chosen in accordance with quadrant assignment $\mathcal{S}$.

First, we recap the notation that we will work with here. The graph $G = (V, E_u \dot\cup E_r)$ denotes an instance of the UPWARD-RIGHTWARD-PRESCRIBED PLANARITY problem. The linearized embedding $\mathcal{E}$ has the following property: the corresponding 2-SAT instance $\mathcal{I}_G^{\mathcal{E}}$ (created according to Chapter 6) is satisfiable and there exists a solution $\mathcal{S}$ for it. This solution is interpreted as an assignment of edges to the quadrants. And $\mathcal{A}$ denotes the angle assignment induced by $\mathcal{S}$. Finally, for $e \in E_r$, $g(e)$ denotes the gadget chosen to replace e according to $\mathcal{S}$.

Lemma 7.4. *Let $e = (v_e, w_e) \in E_r$ be an r-edge. Set $E'_u := E_u \dot\cup g(e)$ and $E'_r := E_r \setminus \{e\}$. Let $(G' = (V', E'_u \dot\cup E'_r), \mathcal{E}')$ be an instance of* Upward-Rightward-Prescribed Planarity *where the linearized order of edges around a vertex v in $\mathcal{E}'$ is defined as follows:*

1. *If $v \in v_G$, then the linearized order of edges around v is adopted from $\mathcal{E}$.*
2. *If $g(e) = \mathcal{H}_e^3$ and $v = q_e$, then the linearized order of edges around this vertex is set to $[(w_e, q_e), (v_e, q_e)]$ (see Figure 4.3).*
3. *If $g(e) = \mathcal{H}_e^4$ and $v = q_e$, then the linearized order of edges around this vertex is set to $[(q_e, v_e), (q_e, w_e)]$ (see Figure 4.3).*

Let $\mathcal{S}'$ be the quadrant assignment of G' defined as follows:

1. *Every edge $e' \in E \setminus \{e\}$ is assigned in $\mathcal{S}'$ to the same quadrants around its end-vertices as in $\mathcal{S}$.*
2. *The edges that belong to $g(e)$ are assigned to the quadrants as depicted blue in Figure 4.3.*

Then, G'_u is upward planar and $\mathcal{S}'$ is a solution for $\mathcal{I}_{G'}^{\mathcal{E}'}$.

Proof. Face-local constraints set according to Section 6.3 guarantee that for every r-edge e, $G_u + g(e) = G'_u$ is upward planar.

Next, we show that the constraints set according to Chapter 6 with regard to $(G', \mathcal{E}')$ are satisfied by $\mathcal{S}'$. The following argument is based on the fact that $\mathcal{S}$ satisfies these constraints with regard to G.

The quadrants to which the gagdet-edges of $g(e)$ are assigned in $\mathcal{S}'$ comply with the direction of these edges and hence, the trivial constraints set according to Section 6.1 are satisfied by $\mathcal{S}'$.

Quadrant assignment $\mathcal{S}'$ also satisfies the constraints from Section 6.2: in the linearized order of edges around v_e (w_e), the edge e was replaced by a gadget-edge of $g(e)$ assigned to the same quadrant to which e is assigned in $\mathcal{S}$ and the quadrant assignment of the remaining edges is adopted from $\mathcal{S}$.

Face-local constraints set with regard to $(G', \mathcal{E}')$ have to guarantee that for an arbitrary r-edge $e' \in E'_r = E_r \setminus \{e\}$, the insertion of the gadget $g'(e')$ chosen according to $\mathcal{S}'$ into G'_u maintains the upward planarity. Lemma 7.3 says that $\mathcal{A}_{g(e), g(e')}$ induces an upward embedding of $G_u + g(e) + g(e') = G'_u + g(e')$. Observe that the quadrant assignment of e' is the same in $\mathcal{S}'$ and $\mathcal{S}$ and hence, $g(e') = g'(e')$. Hereby, the face-local constraints are also satisfied by $\mathcal{S}'$ and finally, $\mathcal{S}'$ is a solution for $\mathcal{I}_{G'}^{\mathcal{E}'}$.

□

Lemma 7.5. *Let $R \subseteq E_r$ be a set of r-edges. Set $E_r^* := E_r - R$ and $E_u^* := E_u \dot\cup g(R)$. Let $(G^* = (V^*, E_u^* \dot\cup E_r^*), \mathcal{E}^*)$ be an instance of* Upward-Rightward-Prescribed Planarity *where the linearized order of edges around a vertex v in $\mathcal{E}^*$ is defined as follows:*

1. *If $v \in V_G$, then the linearized order of edges around v is adopted $\mathcal{E}$.*
2. *If $v = q_e$ for $e = (v_e, w_e) \in R$ so that $g(e) = \mathcal{H}_e^3$, then the linearized order of edges around this vertex is set to $[(w_e, q_e), (v_e, q_e)]$ (see Figure 4.3).*
3. *If $v = q_e$ for $e = (v_e, w_e) \in R$ so that $g(e) = \mathcal{H}_e^4$, then the linearized order of edges around this vertex is set to $[(q_e, v_e), (q_e, w_e)]$ (see Figure 4.3).*

Let $\mathcal{S}^*$ be the quadrant assignment of G^* defined as follows:

1. *Every edge $e \in E \setminus R$ is assigned in $\mathcal{S}^*$ to the same quadrants of its end-vertices as in $\mathcal{S}$.*
2. *For every $e \in R$, edges of $g(e)$ are assigned the quadrants as depicted blue in Figure 4.3.*

Then, G_u^* is upward planar and $\mathcal{S}^*$ is a solution for $\mathcal{I}_{G^*}^{\mathcal{E}^*}$.

Proof. We carry out a proof by induction on $|R| = k$.

Base case $k = 0$: $R = \emptyset$, $E_u^* = E_u$, $E_r^* = E_r$, $\mathcal{S} = \mathcal{S}^*$ and hence, $G = G^*$. Hereby, the claim is true.

Induction step $k \curvearrowright k + 1$:

1. There exist $S \subseteq E_r$ and $e \in E_r$ so that $R = S \,\dot\cup\, \{e\}$.
2. Set $G'_u := G_u + g(S)$. Clearly, $|S| = k - 1$. Induction hypothesis yields: the digraph G'_u is upward and the quadrant assignment $\mathcal{S}'$ adopted from $\mathcal{S}$ satisfies the constraints created for the instance G' of Upward-Rightward-Prescribed Planarity with "upward part" induced by G'_u and "rightward part" $E_r \setminus S = E_r^* \,\dot\cup\, \{e\}$.
3. Observe that e is assigned to the same quadrants according to $\mathcal{S}'$ and $\mathcal{S}$. Thus, the gadget $g(e)$ chosen in accordance with $\mathcal{S}'$ is also chosen in accordance with $\mathcal{S}'$.
4. According to Lemma 7.4, $G'_u + g(e) = G_u^*$ is upward and the assignment $\mathcal{S}^*$ adopted from $\mathcal{S}'$ (and by the construction of $\mathcal{S}, \mathcal{S}'$, and $\mathcal{S}^*$, also adopted from $\mathcal{S}$) satisfies the constraints created for the instance of Upward-Rightward-Prescribed Planarity with "upward part" E_u^* and "rightward part" $(E_r^* \,\dot\cup\, \{e\}) \setminus \{e\} = E_r^*$. This instance is exactly $(G^*, \mathcal{E}^*)$.

□

For $R = E_r$: $G_u^* = G_u + g(R) = G_u + g(E_r)$ is upward planar. Recall that $\hat{G}_u$ arises from G_u^* by partitioning edges and hence, $\hat{G}_u$ is also upward planar. This concludes the proof that $\hat{G}$ admits a windrose planar drawing.

7.2 Upward-Rightward-Prescribed Drawing of G

Now, we can apply the linear-time algorithm provided in [ALB+18] to create a 1-bend windrose drawing Γ_w of $\hat{G}$ on a polynomial grid. This drawing is transformed into an upward-rightward-prescribed drawing of G as follows.

Draw every partitioned u-edge (v, w) as the concatenation of polylines representing $(v, q_{(v,w)})$ and $(q_{(v,w)}, w)$; both polylines monotonically increase in the vertical direction and hence, also their concatenation does. Every unpartitioned u-edge stays drawn with the same polyline as in Γ_w. As a result, all u-edges are represented by polylines that monotonically increase in the vertical direction.

Draw every r-edge (v, w) as the drawing of the corresponding gadget $g(e)$:

1. If $g(e) = \mathcal{H}_e^1$, then $g(e) = (v, w) = e$ is represented by a NE-line and hence, e monotonically increases in x-direction.
2. If $g(e) = \mathcal{H}_e^2$, then $g(e) = (w, v) = e^{-1}$ is represented by a NW-line and hence, e monotonically increases in x-direction.
3. If $g(e) = \mathcal{H}_e^3$, then $g(e) = (v, q_e) + (w, q_e)$ and (v, q_e) is represented by a NE-line, (w, q_e) is represented by a NW-line; as a result, e is represented by a polyline that monotonically increases in the horizontal direction.

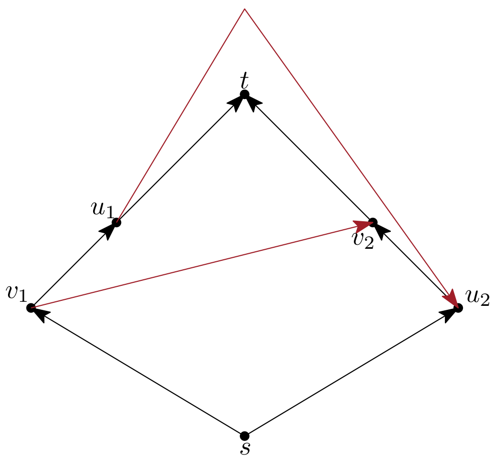

Figure 7.4: An upward-rightward-prescribed planar digraph that admits no straight-line upward-rightward-prescribed drawing.

4. If $g(e) = \mathcal{H}_e^4$, then $g(e) = (q_e, v) + (q_e, w)$ and (q_e, v) is drawn as a NW-line, (q_e, w) is drawn as a NE-line; as a result, e is represented by a polyline that monotonically increases in the horizontal direction.

The drawing Γ_w is planar and hence, the arising drawing is also planar, which means we have found an upward-rightward-prescribed drawing of G. Recall that every edge $e \in E$ is represented by at most two edges in $\hat{G}$ and Γ_w is a 1-bend drawing. Hereby, the arising upward-rightward-prescribed drawing of G is a 3-bend drawing. In the next section, we show that not every upward-rightward-prescribed admits a straight-line drawing.

7.3 Straight-Line Drawings

DiBattista and Tamassia have shown that every upward embedding admits a straight-line upward planar drawing [DBT88]. The situation with Upward-Rightward-Prescribed Planarity is different. To prove this, we provide an upward-rightward-prescribed planar digraph that has no straight-line drawing in Figure 7.4.

Clearly, the provided drawing is upward-rightward-prescribed planar. Now, we briefly justify that this digraph has no straight-line drawing. Consider an arbitrary upward-rightward-prescribed drawing Γ of this digraph. The restriction of the graph to the upward paths $[s, v_1, u_1, t]$ and $[s, u_2, v_2, t]$ separates the plane into two faces: the inner one and the outer one. Exactly one of the edges (u_1, u_2) and (v_1, v_2) is drawn in the outer face as otherwise, these two edges would intersect. Without loss of generality, assume that this edge is (u_1, u_2). The following applies: $\Gamma(s)_y < \Gamma(u_1)_y < \Gamma(t)_y$ and $\Gamma(s)_y < \Gamma(u_2)_y < \Gamma(t)_y$. Hereby, this edge runs under s or above t and has at least one bend. As a result, every upward-rightward-prescribed drawing of G is not straight-line.

Hereby, there exist upward-rightward-prescribed planar graphs that admit no straight-line drawing. We have already shown that every upward-rightward-prescribed planar graph has a 3-bend drawing. The question, whether every graph of this type has a 1-bend or a 2-bend drawing remains open.

7.4 Running Time

Here we justify that the algorithm described in the last two chapters can be implemented in linear time.

1. The construction of $\mathcal{I}_G^{\mathcal{E}}$:

   a) Per edge, there are $\mathcal{O}(1)$ trivial constraints. Hereby, there are $\mathcal{O}(|E|)$ trivial constraints.

   b) Per vertex v, $\mathcal{O}(|N(v)|+1)$ vertex-local constraints are set where $|N(v)|$ denotes the number of edges incident to v. Altogether, there are $\mathcal{O}(|E|+|V|)$ vertex-local constraints.

   c) For the subsequent calcucation of the valid gadgets, the precalculation is done in $\mathcal{O}(|E|+|V|)$ (see Subsection 6.3.1). After that, face-local constraints can be set in $\mathcal{O}(1)$ per edge.

   d) Per edge, there are $\mathcal{O}(1)$ face-local constraints: altogether there are $\mathcal{O}(|E|)$ face-local constraints.

2. Finding a solution $\mathcal{S}$ for $\mathcal{I}_G^{\mathcal{E}}$ if any exists: the size of $\mathcal{I}_G^{\mathcal{E}}$ is in $\mathcal{O}(|E|+|V|)$ and hence, $\mathcal{S}$ can be computed in $\mathcal{O}(|E|+|V|)$ [APT79].

3. The construction of $\hat{G}$ can also be implemented in $\mathcal{O}(|E|+|V|)$:

   a) Every edge $e \in E$ corresponds to at most two edges in $\hat{G}$.

   b) The number of vertices in $\hat{G}$ is also in $\mathcal{O}(|E|+|V|)$: these are the vertices of G and at most one new vertex per edge.

4. A windrose drawing of Γ_w of $\hat{G}$ can be calculated in linear time [ALB$^+$18] (notice that an upward-consistent angle assignment of $\hat{G}$ does not have to be found as it is prescribed by $\mathcal{S}$).

5. Finally, Γ_w is transformed into an upward-rightward-prescribed drawing of G: per edge, at most two polylines are concatenated. This can also be implemented in $\mathcal{O}(|E|)$.

Altogether, for a fixed linearized embedding, the linear-time algorithm decides if there exists an upward-rightward-prescribed drawing and in the positive case, finds one.

8. Instances of Upward-Rightward-Prescribed Planarity without a Fixed Linearized Embedding

In the previous chapters, we have considered instances of UPWARD-RIGHTWARD-PRESCRIBED PLANARITY with a fixed linearized embedding, which means for every vertex v, the clockwise order in which edges appear around v starting at the ray that exits v to the left was fixed. Recall that the reason to prescribe such an embedding was that otherwise, for a fixed quadrant assignment around a vertex, different angle assignments around this vertex are still possible. Hereby, it was unclear how to set vertex-local constraints. In order to clarify the problem once again, we consider a simple example. Let v be a vertex which only has an incoming u-edge e_u and an outgoing r-edge e_r (see Figure 8.1). Consider a constellation in which e_r lies in the SE-quadrant. The "influence" of e_r on the quadrant in which e_u lies depends on the angle between e_u and e_r (large or small). The corresponding SAT-constraint would inevitably contain the third boolean variable that somehow represents the position of the large angle. Such a constraint can not be represented within a 2-SAT instance. The representation in HORN-SAT would also not succeed, as it is not possible to represent "exclusive or (XOR)" in it: this operator is needed to assign every edge to exactly one quadrant per end-vertex. So the difficulty without a fixed linearized embedding was that it is unclear how to set the vertex-local constraints.

The main idea of the approach we have previously followed was based on these facts:

1. For a fixed linearized embedding, every r-edge e and every gadget $\mathcal{H}_e^i$, there exists at most one angle assignment of $G_u+\mathcal{H}_e^i$ that respects the underlying upward embedding of G_u.
2. For a fixed quadrant assignment, a linearized embedding uniquely defines the angle assignment.

As a result, for a gadget $g(e)$ chosen in accordance with a solution $\mathcal{S}$ for $\mathcal{I}_G^{\mathcal{E}}$, the digraph $G + g(e)$ is upward.

If we omit the linearized embedding, two problems arise:

1. For a fixed quadrant assignment, the number of possible angle assignments can become exponential in the number of vertices.

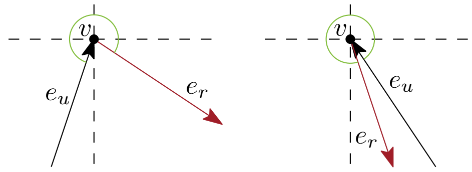

Figure 8.1: On the left, e_r does not influence the position of e_u: the edge can still lie either in the SW- or in the SE-quadrant around v. However, on the right, e_u is forced to also lie in the SE-quadrant around v. The problem was previously avoided because of a fixed linearized embedding, as these two constellations induce distinct linearized embeddings.

2. For a fixed quadrant assignment, it is not clear if there exists an angle assignment so that for every gadget g chosen according to $\mathcal{S}$, the graph $G_u + g$ is upward.

In this chapter, we provide the adaptation of the previous approach to cope with the generalized problem setting. From now on, the embedding consists of a combinatorial embedding of G with a fixed outer face and an upward embedding $\mathcal{E}_u$ of G_u; in particular, an upward-consistent angle assignment of G_u is still prescribed.

Similarly to the previous chapters, we first make face- and vertex-local observations and sketch their correctness. Afterward, we represent the corresponding necessary conditions as a 2-Sat instance. We will also show that the satisfiability of the new 2-Sat instance is again sufficient for the existence of an upward-rightward-prescribed drawing.

In the face-local part, we will widely use the observation that in certain constellations, certain pairs of gadgets are "equivalent", which means either both of them are valid or both of them are not. This will allow us to consider only particular combinations of gadgets and angle assignments so that the angle assignment is later extracted from the quadrant assignment in a unique way.

Vertex-local constraints are also less obvious in this setting: we have to distinguish between sources, sinks, and inner vertices of G_u. Further on, for a sink or a source v of G_u, l_v and r_v denote the leftmost and the rightmost u-edge incident to v respectively.

First, we summarize the observations and constraints that can be undertaken from the previous chapters.

1. The constraints set according to Section 6.1 are adopted. They guarantee that each edge incident to v is assigned to exactly one quadrant around v and this quadrant complies with the type of the edge.
2. The observation stated in Lemma 4.3 determines the set of gadgets that come into question. Recall that this observation was made for an arbitrary upward-rightward-prescribed drawing and hence, is valid regardless of whether a linearized embedding is fixed or not. It is worth noting that unlike the previous chapters, for a gadget g, different angle assignments of $G + g$ that respect $\mathcal{E}_u$ are possible and we have to permit only the upward-consistent ones.

8.1 Vertex-Local: Inner vertices of $\mathbf{G_u}$

Similarly to the previous approach, vertex-local constraints are set to guarantee that the circular adjacency list of every vertex can be partitioned into $l_{NW} \cdot l_{NE} \cdot l_{SE} \cdot l_{SW}$, where l_X only contains edges assigned to the X-quadrant around this vertex.

Here we show that the inner vertices of G_u are simple to treat, which means it is intuitively clear how to guarantee the above condition for such a vertex. Let v be an inner vertex of G_u. Fix an arbitrary outgoing (incoming) u-edge o (i). The left and the right side of the directed path $[i, o]$ are well-defined. The main idea is that the linearized order of edges around v to the right of $[i, o]$ is fixed. For the edges on the left side, the situation is symmetrical.

First, consider the clockwise order $[f_1 = o, f_2, \dots f_t = i]$ of the edges to the right of $[i, o]$; the order of quadrants in which these edges lie has to be $[NW, NE, SE, SW]$ (or a sublist of it). Consequentially, we set for every $j \in \{1, \dots t-1\}$:

1. $v_{f_j}^{NE} \to \neg v_{f_{j+1}}^{NW}$
2. $v_{f_j}^{SE} \to \neg v_{f_{j+1}}^{NW}$, $v_{f_j}^{SE} \to \neg v_{f_{j+1}}^{NE}$
3. $v_{f_j}^{SW} \to \neg v_{f_{j+1}}^{NW}$, $v_{f_j}^{SW} \to \neg v_{f_{j+1}}^{NE}$, $v_{f_j}^{SW} \to \neg v_{f_{j+1}}^{SE}$ (or simpler: $v_{f_j}^{SW} \to v_{f_{j+1}}^{SW}$)

In a similar way, consider the clockwise order $[e_1 = i, e_2, \dots, e_k = o]$ of the edges to the left of $[i, o]$; the order of quadrants in which these edges appear has to be $[SE, SW, NW, NE]$ (or, again, a sublist of it). Analogously, we set for every $j \in \{1, \dots k-1\}$:

1. $v_{e_j}^{SW} \to \neg v_{e_{j+1}}^{SE}$
2. $v_{e_j}^{NW} \to \neg v_{e_{j+1}}^{SE}$, $v_{e_j}^{NW} \to \neg v_{e_{j+1}}^{SW}$
3. $v_{e_j}^{NE} \to \neg v_{e_{j+1}}^{SE}$, $v_{e_j}^{NE} \to \neg v_{e_{j+1}}^{SW}$, $v_{e_j}^{NE} \to \neg v_{e_{j+1}}^{NW}$ (or simpler: $v_{e_j}^{NE} \to v_{e_{j+1}}^{NE}$)

Notice that i and o are considered as edges both on the left and on the right side; this performs the interaction between both sides. It is also worth mentioning that all angles around such a vertex are small and hence, the angle assignment around such a vertex is unique.

8.2 Face-Local Observations

In the previous subsection, we have described how to guarantee that the quadrant assignment around an inner vertex of G_u is valid.

From now on, we are mostly interested in sinks and sources of G_u. Further on, we show that in certain constellations certain pairs of gadgets are *equivalent*, which means one gadget can be replaced by another one and vice versa. Gadgets are still identified according to the numeration used in Lemma 4.3 (see Figure 4.3). We recall that $\mathcal{H}_e^i$ denotes the gadget of type i that corresponds to an r-edge e.

As already stated, for gadget $\mathcal{H}_e^i$, the angle assignment of $G_u + \mathcal{H}_e^i$ that respects $\mathcal{E}_u$ is not necessarily unique and we have to define the angle assignments that we write about more concretely. Let a and b be two edges incident to a vertex v and neighboring in its adjacency list. To avoid ambiguities, from now on writing about a large angle between a and b we implicitly means that a lies to the left of b. "Gadget $\mathcal{H}_e^i$ is valid with a large (small) angle between the edges a and b" means that there exists an upward embedding of $G_u + \mathcal{H}_e^i$ that respects $\mathcal{E}_u$ and in which the angle between a and b is large (small).

8.2.1 A Sink of $\mathbf{G_u}$

First, consider a sink v of G_u. Let f be the face of G_u to which v is assigned in $\mathcal{E}_u$.

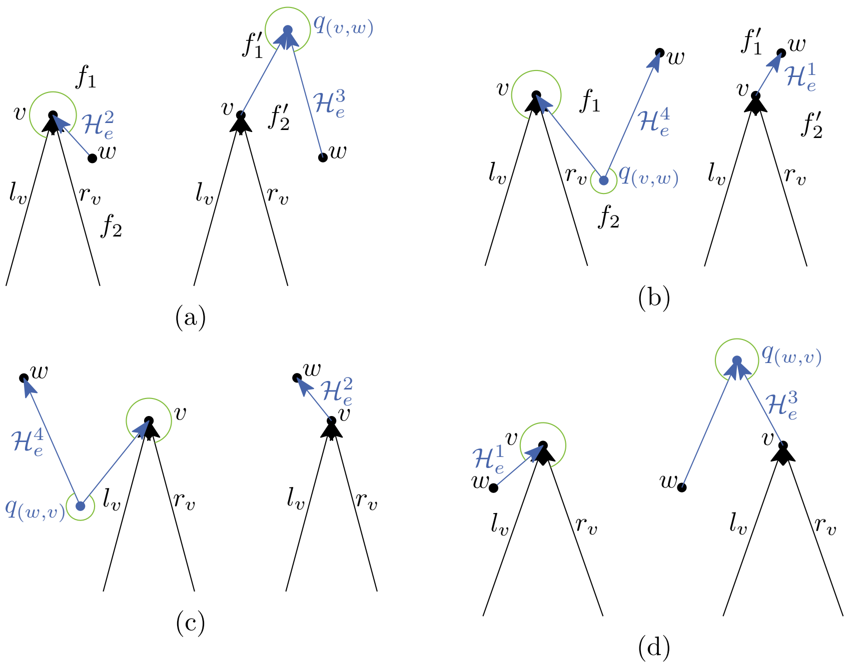

Figure 8.2: Four pairs of equivalent gadgets incident to a sink v of G_u.

8.2.1.1 Outgoing Edges

Consider an r-edge $e = (v, w)$ that has to be embedded into f. Gadget $\mathcal{H}_e^2$ is valid with a large angle between l_v and (w, v) exactly if the gadget $\mathcal{H}_e^3$ is valid. To verify this, we compare faces f_1, f_2 with faces f_1', f_2' that arise from f through the insertion of $\mathcal{H}_e^2$ and $\mathcal{H}_e^3$ respectively (see Figure 8.2 (a)):

1. Upward face f_1 contains the same number of switches and large angles as f_1'.
2. Upward face f_2 contains the same number of switches and large angles as f_2'.

Consequentially, upward face f_1 (f_2) satisfies the condition of Lemma 3.3 exactly if f_1' (f_2') does. Hereby, gadgets $\mathcal{H}_e^2$ and $\mathcal{H}_e^3$ are equivalent in this constellation (a large angle between l_v and (w, v)). We will indeed use it: setting the face-local constraints and trying to insert gadget $\mathcal{H}_e^2$, we will only consider angle assignments of $G_u + \mathcal{H}_e^2$ with a large angle between (w, v) and r as otherwise, $\mathcal{H}_e^3$ can be used instead.

Similarly, gadget $\mathcal{H}_e^4$ with a large angle between l_v and $(q_{(v,w)}, v)$ is equivalent to gadget $\mathcal{H}_e^1$. Faces f_1, f_2, f_1', f_2' are defined analogously (see Figure 8.2 (b)):

1. Upward face f_1 contains two more switches and one more large angle than f_1'.
2. Upward face f_2 contains two more switches and one more large angle than f_2'.

Hereby, f_1 (f_2) satisfies the condition of Lemma 3.3 exactly if f_1' (f_2') does and the gadgets are equivalent in this constellation.

8.2.1.2 Incoming Edges

Let $e = (w, v)$ be an r-edge. We provide two pairs of equivalent gadgets; the proof is completely analogous to the previous ones and we omit it here.

1. Gadget $\mathcal{H}_e^4$ with a large angle between $(q_{(w,v)}, v)$ and r_v is equivalent to gadget $\mathcal{H}_e^2$ (see Figure 8.2 (c)).
2. Gadget $\mathcal{H}_e^1$ with a large angle between (w, v) and r_v is equivalent to gadget $\mathcal{H}_e^3$ (see Figure 8.2 (d)).

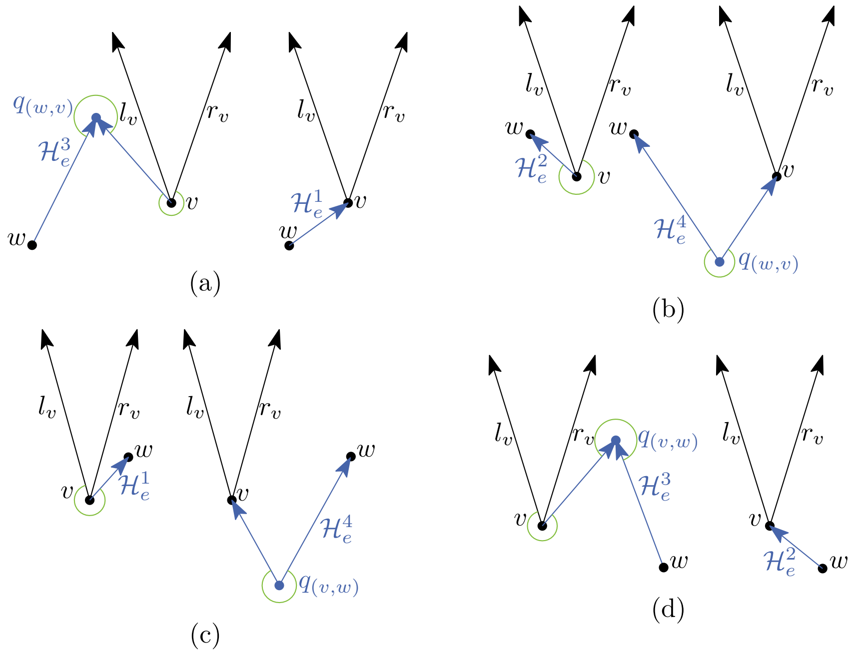

Figure 8.3: Four pairs of equivalent gadgets incident to a source of G_u.

8.2.2 Source of G_u

Here we provide the pairs of gadgets that are equivalent for the r-edges incident to a source v of G_u. The proof is again similar to the previous ones and we also omit it here.

8.2.2.1 Incoming Edges

For an incoming r-edge $e = (w, v)$:

1. Gadget $\mathcal{H}_e^3$ with a large angle between $(v, q_{(w,v)})$ and r_v is equivalent to gadget $\mathcal{H}_e^1$ (see Figure 8.3 (a)).
2. Gadget $\mathcal{H}_e^2$ with a large angle between (v, w) and r_v is equivalent to gadget $\mathcal{H}_e^4$ (see Figure 8.3 (b)).

8.2.2.2 Outgoing Edges

For an outgoing r-edge $e = (v, w)$:

1. Gadget $\mathcal{H}_e^1$ with a large angle between l_v and (v, w) is equivalent to gadget $\mathcal{H}_e^4$ (see Figure 8.3 (c)).
2. Gadget $\mathcal{H}_e^3$ with a large angle between l_v and $(v, q_{(v,w)})$) is equivalent to gadget $\mathcal{H}_e^2$ (see Figure 8.3 (d)).

8.2.3 Special Upward-Rightward-Prescribed Drawings

In Lemma 4.2, we have proven that for every upward-rightward-prescribed embedding, there exists a drawing in which every r-edge is drawn with at most three local extrema. Now we combine this result with just made observations to concentrate on a certain type of upward-rightward-prescribed drawings later on.

Let Γ be an upward-rightward-prescribed drawing of a digraph G which satisfies the condition stated in Lemma 4.2. This property of Γ allows to identify a gadget $\mathcal{H}_e^i$ with every r-edge e. It is also possible to extract the linearized embedding $\mathcal{E}_\Gamma$ and the quadrant assignment $\mathcal{S}_\Gamma$ from this drawing. Quadrant assignment $\mathcal{S}_\Gamma$ is thus a solution for $\mathcal{I}_G^{\mathcal{E}_\Gamma}$.

Now consider an arbitrary sink v of G_u. Let E_{SE} (see Figure 8.4 (a)) denote the set of r-edges e incident to v such that:

1. The edge e is embedded into the face f of G_u to which v is assigned in $\mathcal{E}_u$.
2. The edge e lies in the SE-quadrant around v.
3. There is a large angle between l_v and e (in simple words, e lies to the right of l_v).

Notice that every edge e in E_{SE} is identified with a gadget $\mathcal{H}_e^2$ or $\mathcal{H}_e^4$. Replace every such gadget by an equivalent gadget in accordance with 8.2.1.1. Observe that after this transformation, every edge $e \in E_{SE}$ is moved into the NE-quadrant. The arising quadrant assignment $\mathcal{S}'$ is also a solution for $\mathcal{I}_G^{\mathcal{E}_\Gamma}$:

1. Clearly, trivial constraints stay satisfied.
2. The equivalence of gadgets proven in 8.2.1.1 guarantees that face-local constraints stay satisfied.
3. It is also clear that vertex-local constraints around every vertex other than v stay satisfied.
4. Let $l_{NW} \cdot l_{NE} \cdot l_{SE} \cdot l_{SW}$ be the linearized order of edges around v in $\mathcal{E}_\Gamma$ so that l_X only contains edges lying in the X-quadrant v according to $\mathcal{S}$. Observe: $l_{SE} = E_{SE} \cdot l'_{SE}$. The reason is the following: there are only incoming u- and outgoing r-edges in the SE-quadrant around v and every edge in E_{SE} lies right of the rightmost u-edge r_v. The above-described transformation moves E_{SE} into the NE-quadrant so that the linearized order of edges in the resulting embedding is $(l'_{NW} := l_{NW})\cdot(l'_{NE} := l_{NE}\cdot E_{SE})\cdot l'_{SE}\cdot(l'_{SW} := l_{SW})$. Notice that l'_X only contains edges lying in the X-quadrant around v according to $\mathcal{S}'$, which means the edges around v appear in quadrants in the valid order and vertex-local constraints are satisfied.

Hereby, there exists an upward-rightward-prescribed drawing of G with underlying quadrant assignment $\mathcal{S}'$. This transformation is carried out for every sink of G_u.

For symmetry reasons, for every sink v of G_u it is possible to replace the gadgets corresponding to E_{SW} (defined analogously; see Figure 8.4 (b)) by equivalent gadgets described in 8.2.1.2. The proof is similar except for the following particularity. Let $\mathcal{E}'_\Gamma$ be the linearized

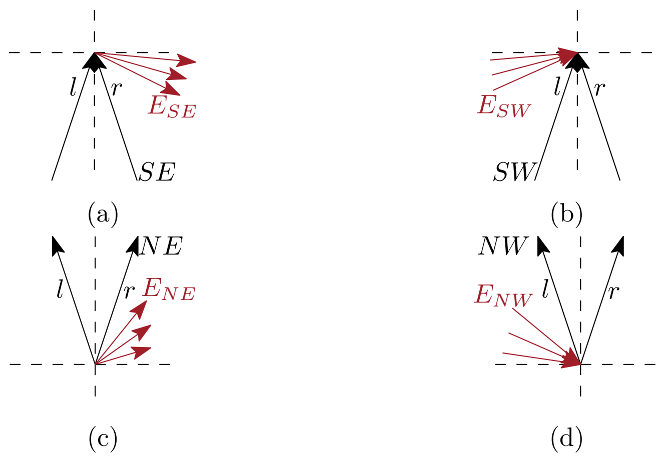

Figure 8.4: The red edges can be redrawn.

embedding of G arising after this transformation and let t_v be the linearized order of edges around v in $\mathcal{E}_\Gamma$. Indeed, the linearized embedding is changed by this transformation because the edges in E_{SW} are moved from the end of t_v (the SW-quadrant) to its beginning (the NW-quadrant). Now, the same argument as before can be applied to show that the arising quadrant assignment is a solution for $\mathcal{I}_G^{\mathcal{E}'_\Gamma}$. The details are omitted here.

Finally, in the same way, it can be proven that for every source v of G_u, it is possible to carry out the following transformations so that there still exists an upward-rightward-prescribed drawing:

1. Replace gadgets corresponding to E_{NW} (see Figure 8.4 (c)) by equivalent gadgets described in 8.2.2.1.
2. Replace gadgets corresponding to E_{NE} (see Figure 8.4 (d)) by equivalent gadgets described in 8.2.2.2.

Each of the aforementioned replacements yields a new quadrant assignment for which there exists an upward-rightward-prescribed drawing of G. Let Γ' be an upward-rightward-prescribed drawing that corresponds to the quadrant assignment that arises after carrying out all the previously mentioned transformations. This drawing satisfies the following conditions:

1. For every sink v of G_u:
   a) For every incoming r-edge e lying in the SW-quadrant around v, there is a large angle between l_v and e.
   b) For every outgoing r-edge e lying in the SE-quadrant around v, there is a large angle between e and r_v.
2. For every source v of G_u:
   a) For every incoming r-edge e lying in the NW-quadrant around v, there is a large angle between l_v and e.
   b) For every outgoing r-edge e lying in the NE-quadrant around v, there is a large angle between e and r_v.

We call such a drawing a *special* upward-rightward-prescribed drawing. Here we have shown that every upward-rightward-prescribed embedding admits a special upward-rightward-prescribed drawing. Consequentially, we can restrict our constraints so that we only look for such a drawing.

8.3 Face-Local Constraints

In this section, we describe how to apply the observations made in the previous section to set face-local constraints. The approach is similar to Section 6.3. The main difference is the following: for a gadget g, an angle assignment of $G + g$ that respects $\mathcal{E}_u$ is possibly not unique, but we consider only the one that corresponds to a special upward-rightward-prescribed drawing. Later on, this guarantees that the angle assignment is extracted from the quadrant assignment in a unique way.

Let $e = (v, w)$ be an r-edge. The set S_e of valid gadgets is defined similarly to Section 6.3 with the following additional limitation: the set S_e contains gadget $\mathcal{H}_e^i$ only if there exists an upward embedding of $G_u + \mathcal{H}_e^i$ that respects $\mathcal{E}_u$ and satisfies the following restrictions:

1. If v is a sink in G_u:

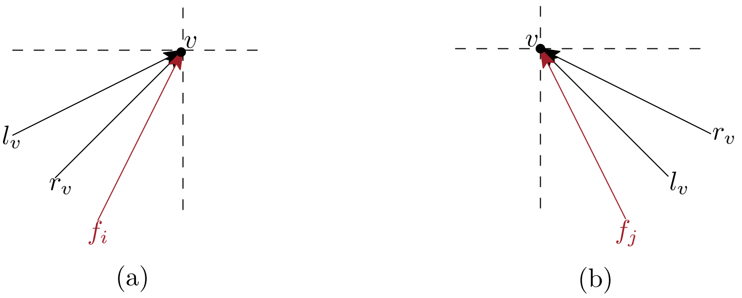

Figure 8.5: In a special upward-rightward-prescribed drawing, there can not exist both r-edges in the SE- and SW-quadrant within the large angle between l and r.

   a) If $i = 2$, there is no large angle between l_v and (w, v).
   b) If $i = 4$, there is no large angle between l_v and $(q_{(v,w)}, v)$.

2. If v is a source in G_u:
   a) If $i = 1$, there is no large angle between l_v and (v, w).
   b) If $i = 3$, there is no large angle between l_v and $(v, q_{(v,w)})$.
3. If w is a sink in G_u:
   a) If $i = 1$, there is no large angle between (v, w) and r_w.
   b) If $i = 4$, there is no large angle between $(q_{(v,w)}, w)$ and r_w.
4. If w is a source in G_u:
   a) If $i = 2$, there is no large angle between (w, v) and r_w .
   b) If $i = 3$, there is no large angle between $(w, q_{(v,w)})$ and r_w.

Notice that S_e can be obtained in linear time (see Subsection 6.3.1) as we still consider only one angle assignment per gadget. After that, face-local constraints are set in accordance with the prescription provided in Subsection 6.3.2.

8.4 Vertex-Local Constraints for a Sink of $\mathbf{G_u}$

In this section, we provide the vertex-local constraints that guarantee that the quadrants, in which the edges incident to the same vertex lie, appear in a valid order. Again, let Γ be an arbitrary special upward-rightward-prescribed drawing and let v be a sink of G_u.

First, consider the edges "under the peak": let $[e_1 = l_v, \ldots, e_k = r_v]$ be the counterclockwise order of edges between l_v and r_v. Because of the large angle between l_v and r_v, these edges lie in the South-quadrants of v and the above order coincides with the left-to-right order. Hence, in this list the SW-edges appear first and then, the SE-edges. Consequentially, we set:

1. For each $i \in \{1, \ldots, k\}$: $v^{SE}_{e_i} \vee v^{SW}_{e_i}$.
2. For each $i \in \{1, \ldots, k-1\}$: $v^{SE}_{e_i} \rightarrow v^{SE}_{e_{i+1}}$.

Now, consider the edges "above the peak": let $F = [f_1, \ldots, f_t]$ be the clockwise order of edges to the left of l_v and to the right of r_v (exclusive l_v and r_v). Recall that Γ is special: for an edge f_i in the SW-quadrant, there is a large angle between l_v and f_i and hence, the existence of such an edge implies that all the u-edges incident to v lie in the SW-quadrant (see Figure 8.5 (a)). Similarly, the existence of an edge f_j in the SE-quadrant implies that all the u-edges incident to v lie in the SE-quadrant (see Figure 8.5 (b)). Hereby, both the SE- and the SW-edges can not exist in F at the same time.

Hence, the order of the quadrants in which these edges appear is either $[NW, NE, SW]$ or $[SE, NW, NE]$ (or a sublist of one them). Thus, if we guarantee that there are no SW- and SE-edges at the same time and that the quadrants in which the edges in F lie appear in the order $[SE, NW, NE, SW]$, then the valid order of quadrants is ensured.

For this purpose, we set the following constraints:

1. For all $i \in \{1, \dots t\}$: $v_{f_i}^{SE} \to v_{l_v}^{SE}$, $v_{f_i}^{SW} \to v_{r_v}^{SW}$.
2. For all $i \in \{1, \dots t-1\}$: $v_{f_i}^{NW} \to \neg v_{f_{i+1}}^{SE}$.
3. For all $i \in \{1, \dots t-1\}$: $v_{f_i}^{NE} \to \neg v_{f_{i+1}}^{NW}$.
4. For all $i \in \{1, \dots t-1\}$: $v_{f_i}^{NE} \to \neg v_{f_{i+1}}^{SE}$.
5. For all $i \in \{1, \dots t-1\}$: $v_{f_i}^{SW} \to v_{f_{i+1}}^{SW}$.

8.5 Vertex-Local Constraints for a Source of $\mathbf{G_u}$

Here we only provide the vertex-local constraints for a source v of G_u without a justification: for symmetry reasons, it is analogous to the previous section. Let $[e_1 = l_v, e_2, \dots, e_k = r_v]$ the clockwise order of edges between l_v and r_v. Set:

1. For all $i \in \{1, \dots, k\}$: $v_{e_i}^{NE} \vee v_{e_i}^{NW}$.
2. For all $i \in \{1, \dots, k-1\}$: $v_{e_i}^{NE} \to v_{e_{i+1}}^{NE}$.

Let $F = [f_1, \dots, f_t]$ be the counterclockwise order of edges to the left of l_v and to the right of r_v (again, exclusive l_v and r_v). In this constellation, two possible orders in which quadrants of edges in F appear are $[NE, SW, SE]$ and $[SW, SE, NW]$. Hence, we have to guarantee that there are no NE- and NW-edges in F at the same time and that the order of quadrants in which the edges in F lie coincides with $[NE, SW, SE, NW]$ or with a sublist of it. For this reason, set:

1. For all $i \in \{1, \dots, k\}$: $v_{f_i}^{NE} \to v_{l_v}^{NE}, v_{f_i}^{NW} \to v_{r_v}^{NW}$.
2. For all $i \in \{1, \dots, k-1\}$: $v_{f_i}^{SW} \to \neg v_{f_{i+1}}^{NE}$.
3. For all $i \in \{1, \dots, k-1\}$: $v_{f_i}^{SE} \to \neg v_{f_{i+1}}^{SW}$.
4. For all $i \in \{1, \dots, k-1\}$: $v_{f_i}^{SE} \to \neg v_{f_{i+1}}^{NE}$.
5. For all $i \in \{1, \dots, k-1\}$: $v_{f_i}^{NW} \to v_{f_{i+1}}^{NW}$.

8.6 Sufficiency

The proof, that the satisfiability of these constraints is sufficient for the existence of an upward-rightward-prescribed drawing, is now analogous to the one in Chapter 7: solution $\mathcal{S}$ for these constraints is again interpreted as a quadrant assignment. By the construction of S_e, corresponding angle assignment $\mathcal{A}$ is extracted in a unique way from $\mathcal{S}$. We again consider the corresponding instance $\hat{G}$ of Windrose Planarity (compare with Chapter 7).

Two properties of $\mathcal{S}$ and $\mathcal{A}$ were used in Chapter 7 to show the existence of the windrose drawing of $\hat{G}$.

First, to show that $\hat{G}$ is upward we used that:

1. Quadrant assignment $\mathcal{S}$ yields a candidate upward embedding of $\hat{G}$ (the incoming and outgoing edges are both consecutive around every vertex): the above-described vertex-local constraints guarantee that edges assigned to North- and South-quadrants are consecutive around every vertex.

2. Every gadget chosen according to $\mathcal{S}$ is valid with the angle assignment $\mathcal{A}$: face-local constraints are set so that the angle-assignment $\mathcal{A}$ is extracted from $\mathcal{S}$ in a unique way and for every gadget g chosen according to $\mathcal{S}$, $\mathcal{A}_g$ is the angle-assignment for which $G_u + g$ is upward-consistent.

As a result, the same argument as in Chapter 7 can be applied to show that $\hat{G}$ is upward planar.

Further, vertex-local constraints guarantee that for every vertex v, the circular adjacency list of v can be partitioned into $l_{NW} \cdot l_{NE} \cdot l_{SE} \cdot l_{SW}$ so that l_X only contains edges assigned to the X-quadrant around v (for every $X \in \{NW, NE, SE, SW\}$). By the construction of $\mathcal{A}$ (see Chapter 7), for every vertex v in $\hat{G}$, in the left-to-right order of incoming (outgoing) edges the NE- and then the NW-edges (first the NW- and then the NE-edges) appear. Hereby, $\hat{G}$ satisfies the condition stated in Theorem 3.3

These two properties of $\hat{G}$ are sufficient for the existence of a windrose drawing of $\hat{G}$ and hence, also for the existence of an upward-rightward-prescribed drawing of G as shown in Chapter 7.

The last thing to evaluate is the running time of the algorithm sketched in this chapter. This algorithm differs from the one, that was introduced in the previous chapters, in the definition of the constraints. Their number remains linear, they can also be obtained in linear time and hereby, the algorithm provided in this chapter can also be implemented in linear time.

9. NP-hardness of Single-Source Upward-Rightward-Prescribed Planarity

Upward Planarity is an $\mathcal{NP}$-hard problem [GT95]. In several works, the single-source restriction of the problem has been studied. First, Hutton and Lubiw found a polynomial algorithm to test if a single-source digraph has an upward drawing [HL96]. Later, Bertolazzi et al. improved this result by providing a linear-time algorithm that solves this problem [BDBMT98]. For this reason, it is natural to consider the single-source special case of Upward-Rightward-Prescribed Planarity. In this chapter, we prove the $\mathcal{NP}$-hardness of Single-Source Upward-Rightward-Prescribed Planarity.

Theorem 9.1. *For a single-source digraph, it is $\mathcal{NP}$-hard to decide if it admits an upward-rightward-prescribed drawing.*

Proof. To show this, we provide the reduction of Upward Planarity to Single-Source Upward-Rightward-Prescribed Planarity. Let $G = (V, E)$ be an instance of Upward Planarity.

If G has no edges, then $g(G) := (V'_G, E'_u \,\dot{\cup}\, E'_r)$ where $V_{G'} = \{s, t\}$, $E_{G'} = \{(s, t)\}$ and $E'_r = \emptyset$. Graph $g(G)$ is a trivial yes-instance with the unique source s.

If G has no sources, then it necessarily contains a cycle and hence, G is not upward planar. Hereby, we set $g(G) := (V'_G, E'_u \dot{\cup} E'_r)$ where $V_{G'} = \{s, v_1, v_2, v_3\}$, $E'_u = \{(s, v_1), (v_1, v_2), (v_2, v_3), (v_3, v_1)\}$, $E'_r = \emptyset$. Graph $g(G)$ a trivial no-instance with the unique source s.

Otherwise, let $\{s^1, s^2, \ldots, s^t\}$ be the set of sources of G ($t > 0$). This set can be obtained in linear time. For $i \in \{2, \ldots, t\}$, let $\{(s^i, v^i_1), \ldots, (s^i, v^i_{k^i})\}$ be the set of the edges incident to s^i.

Graph G is transformed into an instance $g(G)$ of Upward-Rightward-Prescribed Planarity as follows:

1. For every $i \in \{2, \ldots, t\}$ and for every $j \in \{1, \ldots, k^i\}$:
   a) Replace (s^i, v^i_j) by u-edges (s^i, x^i_j), (x^i_j, y^i_j), and (y^i_j, v^i_j). Hence, edge (s^i, v^i_j) is partitioned.
   b) For each y^i_j, insert r-edge (y^i_j, s^i).

2. The remaining edges of G, that is, the edges that are incident to none of $\{s^2, \dots, s^t\}$, stay unchanged. To be exact, these edges are u-edges in $g(G)$.

Notice that s^1 is the unique source of $g(G)$.

Now we prove that g is a reduction of Upward Planarity to Single-Source Upward-Rightward-Prescribed Planarity.

"$\Rightarrow$" Let $g(G) = G'$ be upward-rightward-prescribed planar. Then, there exists an upward-rightward-prescribed drawing Γ of G'. We can easily construct an upward drawing Γ' of G as follows:

1. Set $\Gamma' := \Gamma$.
2. The drawing of every r-edge is removed.
3. Every edge that is incident to one of $\{s^2, \dots, s^t\}$ is drawn as the concatenation of the edges arisen by partitioning. The drawing of each of these three edges increases monotonically in the vertical direction and hence, their concatenation also does.

Hereby, G is upward planar.

"$\Leftarrow$" Let G be upward planar. In Chapter 3, we have seen that every upward planar digraph has a straight-line upward-rightward drawing (see Figure 3.1). Analogously, every upward planar digraph has a straight-line *upward-leftward* drawing in which every edge monotonically increases in the vertical direction and monotonically decreases in the horizontal direction. Let Γ be such a drawing of G.

Then, for every $i \in \{2, \dots, t\}$, there exists a small neighborhood $U_{\epsilon_i}(\Gamma(s^i)) =: U_i$ that contains no vertices other than s^i and that only contains edges incident to s^i. For every $j \in \{1, \dots, k^i\}$ we transform the drawing as follows.

1. Every vertex y^i_j is placed at some point (other than s^i) on the edge (s^i, v^i_j) in U_i.
2. Every vertex y^i_j is placed in the same way with the additional restriction: y^i_j lies above x^i_j. As a result, $(s^i, x^i_j), (x^i_j, y^i_j), (y^i_j, v^i_j)$ are drawn upward.
3. Finally, draw the r-edge (y^i_j, s^i) with a 1-bend polyline running along (s^i, v^i_j). The drawing Γ is upward-leftward and hereby, this curve monotonically increases in the horizontal direction (see Figure 9.1).

By the construction, the arising drawing is an upward-rightward-prescribed drawing of $g(G)$ and hence, $g(G)$ is upward-rightward-prescribed planar.

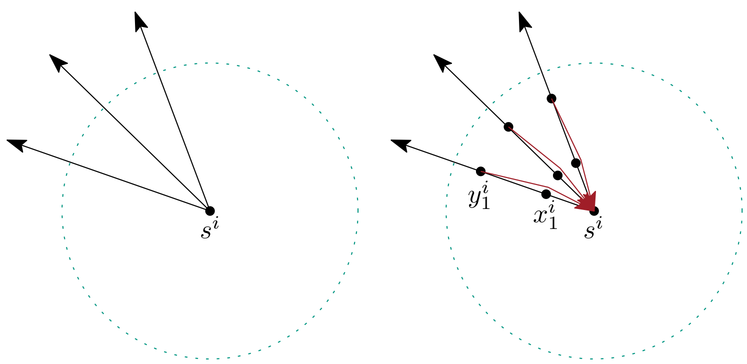

Figure 9.1: On the left: upward drawing of G in a small neighborhood of s^i; on the right: upward-rightward-prescribed drawing of $g(G)$ in the same neighborhood.

Finally, $g(G)$ can be constructed in linear time and hence, g is a polynomial transformation from Upward Planarity to Single-Source Upward-Rightward-Prescribed Planarity. Upward Planarity is $\mathcal{NP}$-hard, which implies that Single-Source Upward-Rightward-Prescribed Planarity is $\mathcal{NP}$-hard as well. $\square$

10. Open questions

In this thesis, we introduced a new planarity variant called upward-rightward-prescribed planarity. Upward-rightward-prescribed-planarity testing is $\mathcal{NP}$-hard, so we concentrated on a restricted problem setting. We considered only the graphs where the upward part G_u is a biconnected spanning subgraph of G. For such a graph, we also assumed that a combinatorial embedding of G and an upward embedding of G_u are fixed. For this special case, we provided a linear-time algorithm that reduces upward-rightward-prescribed-planarity testing to windrose-planarity testing. This relation to windrose planarity implies that every upward-rightward-prescribed planar graph admits a 3-bend (upward-rightward-prescribed planar) drawing on a polynomial grid. We have also provided a graph that admits no straight-line drawing. As a result, the following question arises:

Question 10.1. *Does every upward-rightward-prescribed planar graph have a 2-bend, or even a 1-bend drawing?*

If it can be shown that every windrose planar graph has a straight-line drawing, then by the construction in Chapter 7, every upward-rightward-prescribed planar digraph indeed admits a 1-bend drawing. Another possible approach can be to consider only the instances of windrose planarity arising in Chapter 7: the intuition behind this is that in such instances almost every edge is already partitioned. We could additionally partition the remaining edges and look for a straight-line windrose drawing of the arising instance. Maybe, this construction could allow avoiding additional bends.

So far, we considered the digraphs where G_u is a biconnected spanning subgraph of G. For symmetry reasons, we can treat the graphs where G_r is a biconnected spanning subgraph analogously. In Subsection 6.3, for r-edge e we tested whether $G_u + \mathcal{H}_e^i$ is upward planar or not. We recall that the criterion of upward planarity provided in [BDB91] only works for biconnected graphs. Thus, if an end-vertex of e has no incident u-edges or G_u is not biconnected, then the aforementioned graph is not necessarily biconnected and the criterion can not be applied.

Question 10.2. *Which approach can be followed to consider the digraphs where neither G_u nor G_r is a spanning subgraph of G?*

We have only considered digraphs with a fixed combinatorial embedding of G and a fixed angle assignment of G_u. The next natural generalization is to omit such an angle assignment. For example, maybe, it is possible to construct a flow network (similar to [BDB91]) to

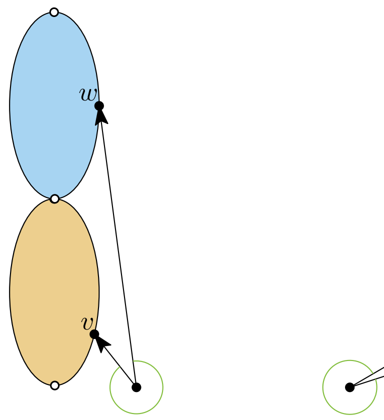

Figure 10.1: The ellipses represent two R-nodes of an UP-tree. Large angles are depicted green. In the left configuration, the gadget of the fourth type is embeddable whereas in the right configuration, it is not embeddable and hence, the gadget choice depends on the concrete configuration of the UP-tree.

find a valid set of gadgets. For this approach, we do not understand which gadgets come into question or if the embeddability of individual gadgets is sufficient for the upward and windrose planarity of the arising graph $\hat{G}$.

Question 10.3. *Is there any way to efficiently solve the problem for a fixed combinatorial embedding?*

Besides the results provided in this thesis, we tried to consider instances of Upward-Rightward-Prescribed Planarity without a fixed combinatorial embedding in which every vertex, again, has an incident u-edge but G_u is now a single-source digraph. The main idea was to follow the approach of Angelini et al. in [ADBF+12] and replace an SPQR-tree with a UP-tree introduced by Brückner et al. in [BHR19]. We encountered some difficulties with this approach. Firstly, the gadget choice depends on flips in the concrete configuration of the UP-tree (see Figure 10.1). We emphasize that the position of the big angle in a gadget is crucial: if we reflect a gadget of the third or the fourth type (see Figure 4.3) and thus, change the position of the big angle around the gadget-vertex, then the arising gadget corresponds to an r-edge (w, v) instead of (v, w). Secondly, it is also unclear how to guarantee that the order of edges incident to the same vertex is "valid" regarding windrose planarity; in particular, this becomes difficult for the P-nodes of the UP-tree at which the edges can be arbitrarily permuted. We have not found any way to represent these constraints so that their satisfiability can be efficiently tested.

Question 10.4. *Is there any way to use UP-trees to solve the* Upward-Rightward-Prescribed Planarity *problem for digraphs where G_u is a single-source spanning subgraph of G?*

11. Conclusion

In this thesis, we introduced and considered the generalization of the UPWARD PLANARITY problem called UPWARD-RIGHTWARD-PRESCRIBED PLANARITY. Upward-planarity testing is $\mathcal{NP}$-hard which implies that UPWARD-RIGHTWARD-PRESCRIBED PLANARITY is $\mathcal{NP}$-hard as well, so we concentrated on special cases.

We considered the graphs where the upward part G_u is a biconnected spanning subgraph. For such a graph G, we assumed that a combinatorial embedding of G and an angle assignment of G_u are fixed.

For this restriction first, we have provided a linear-time algorithm for instances with a fixed linearized embedding, that is for every vertex v, we additionally prescribed the clockwise order in which edges appear around v starting at the ray that exits v to the left. For this purpose, we have considered an arbitrary upward-rightward-prescribed drawing and found a relation between the quadrants around vertices in which the edges lie: first, we considered how an r-edge can be represented in such a drawing and then, we examined the edges incident to the same vertex. We represented this relation as a set of constraints in 2-SAT: the satisfiability of these constraints is necessary for the existence of an upward-rightward-prescribed drawing. After that, we have also shown that this satisfiability is also sufficient: given a solution for these constraints, we provided a transformation of the initial graph into an equivalent instance of WINDROSE PLANARITY that has a windrose drawing. This drawing yields an upward-rightward-prescribed drawing of the initial graph.

The relation to windrose planarity resulted in the existence of 3-bend upward-rightward-prescribed drawings on a polynomial grid. Then, we provided a graph that admits no straight-line drawing.

Afterward, we provided the generalization of the approach to provide a linear-time algorithm to solve the problem if the linearized embedding is not fixed.

Finally, we have shown UPWARD-RIGHTWARD-PRESCRIBED PLANARITY remains $\mathcal{NP}$-hard for single-source digraphs even though UPWARD PLANARITY can be polynomially solved for this type of graphs. In conclusion, we have proposed a number of open questions.

Bibliography

[ADBF+12] Patrizio Angelini, Giuseppe Di Battista, Fabrizio Frati, Maurizio Patrignani, and Ignaz Rutter. Testing the Simultaneous Embeddability of two Graphs whose Intersection is a Biconnected or a Connected Graph. *Journal of Discrete Algorithms*, 14:150–172, 2012. Special issue of IWOCA'10.

[ALB+18] Patrizio Angelini, Giordano Da Lozzo, Giuseppe Di Battista, Valentino Di Donato, Philipp Kindermann, Günter Rote, and Ignaz Rutter. Windrose planarity: Embedding graphs with direction-constrained edges. *ACM Transactions on Algorithms (TALG)*, 14(4):54, 2018.

[APT79] Bengt Aspvall, Michael F. Plass, and Robert Endre Tarjan. A linear-time algorithm for testing the truth of certain quantified boolean formulas. *Information Processing Letters*, 8(3):121 – 123, 1979.

[BDB91] Paola Bertolazzi and Giuseppe Di Battista. On Upward Drawing Testing of Triconnected Digraphs (extended abstract). In *Proceedings of the 7th Annual ACM Symposium on Computational Geometry*, pages 272–280. ACM Press, 1991.

[BDBMT98] Paola Bertolazzi, Giuseppe Di Battista, Carlo Mannino, and Roberto Tamassia. Optimal upward planarity testing of single-source digraphs. *SIAM Journal on Computing*, 27(1):132–169, 1998.

[BHR19] Guido Brückner, Markus Himmel, and Ignaz Rutter. An spqr-tree-like embedding representation for upward planarity. *arXiv preprint arXiv:1908.00352*, 2019.

[DBT88] Giuseppe Di Battista and Roberto Tamassia. Algorithms for Plane Representations of Acyclic Digraphs. *Theoretical Computer Science*, 61(2-3):175–198, 1988.

[DGDK+14] Emilio Di Giacomo, Walter Didimo, Michael Kaufmann, Giuseppe Liotta, and Fabrizio Montecchiani. Upward-rightward planar drawings. In *IISA 2014, The 5th International Conference on Information, Intelligence, Systems and Applications*, pages 145–150. IEEE, 2014.

[DLP14] Walter Didimo, Giuseppe Liotta, and Maurizio Patrignani. On the complexity of hv-rectilinear planarity testing. In *International Symposium on Graph Drawing*, pages 343–354. Springer, 2014.

[GT95] Ashim Garg and Roberto Tamassia. On the Computational Complexity of Upward and Rectilinear Planarity Testing. In *DIAMCS International Workshop*, volume 894 of *Lecture Notes in Computer Science*, pages 286–297. Springer, January 1995.

[HL96] Michael D Hutton and Anna Lubiw. Upward planar drawing of single-source acyclic digraphs. *SIAM Journal on Computing*, 25(2):291–311, 1996.

[JLM98] Michael Jünger, Sebastian Leipert, and Petra Mutzel. Level planarity testing in linear time. In *International Symposium on Graph Drawing*, pages 224–237. Springer, 1998.

[KR17] Boris Klemz and Günter Rote. Ordered level planarity, geodesic planarity and bi-monotonicity. In *International Symposium on Graph Drawing and Network Visualization*, pages 440–453. Springer, 2017.