The Duals of Upward Planar Graphs on Cylinders

Christopher Auer, Christian Bachmaier,
Franz J. Brandenburg, Andreas Gleißner, and Kathrin Hanauer

Department of Informatics and Mathematics, University of Passau
{auerc,bachmaier,brandenb,gleissner,hanauer}@fim.uni-passau.de

UNIVERSITÄT PASSAU
Fakultät für Informatik und Mathematik

Technical Report, Number MIP-1204
Department of Informatics and Mathematics
University of Passau, Germany
September 2011

The Duals of Upward Planar Graphs on Cylinders*

Technical Report MIP-1204
May, 2012

Christopher Auer, Christian Bachmaier, Franz J. Brandenburg, Andreas Gleißner, and Kathrin Hanauer

University of Passau, Germany,
{auerc|bachmaier|brandenb|gleissner|hanauer}@fim.uni-passau.de

Abstract. We consider directed planar graphs with an upward planar drawing on the rolling and standing cylinders. These classes extend the upward planar graphs in the plane. Here, we address the dual graphs. Our main result is a combinatorial characterization of these sets of upward planar graphs. It basically shows that the roles of the standing and the rolling cylinders are interchanged for their duals.

1 Introduction

Directed graphs are used as a model for structural relations where the edges express dependencies. Such graphs are often acyclic and are drawn as hierarchies using the framework introduced by Sugiyama et al. [20]. This drawing style transforms the edge direction into a geometric direction: all edges point upward. If only plane drawings are allowed, one obtains *upward planar graphs*, for short **UP**. These graphs can be drawn in the plane such that the edge curves are monotonically increasing in y-direction and do not cross. Hence, **UP** graphs respect the unidirectional flow of information as well as planarity.

There are some fundamental differences between upward planar and undirected planar graphs. For instance, there are several linear time planarity tests [16], whereas the recognition problem for **UP** is $\mathcal{NP}$-complete [12]. The difference between planarity and upward planarity becomes even more apparent when different types of surfaces are studied: For instance, it is known that every graph embeddable on the plane is also embeddable on any surface of genus 0, e. g., the sphere and the cylinder, and vice versa. However, there are graphs with an upward embedding on the sphere with edge curves increasing from the south to the north pole, which are not upward planar [15]. The situation becomes even more challenging if upward embeddability is extended to other surfaces even if these are of genus 0.

* Supported by the Deutsche Forschungsgemeinschaft (DFG), grant Br835/15-2.

Upward planarity on surfaces other than the plane generally considers embeddings of graphs on a fixed surface in $\mathbb{R}^3$ such that the curves of the edges are monotonically increasing in y-direction. Examples for such surfaces are the standing [6, 13, 18, 19, 21] and the rolling cylinder [6], the sphere and the truncated sphere [9, 11, 14, 15], and the lying and standing tori [8, 10]. We generalized upward planarity to arbitrary two-dimensional manifolds endowed with a vector field which prescribes the direction of the edges [1]. We also studied upward planarity on standing and rolling cylinders, where the former plays an important role for radial drawings [2] and the latter in the context of recurrent hierarchies [3]. We showed that upward planar drawings on the rolling cylinder can be simplified to polyline drawings, where each edge needs only finitely many bends and at most one winding around the cylinder [6]. The same holds for the standing cylinder, where all windings can be eliminated [6]. In accordance to [1], we use the fundamental polygon to define the plane, the standing and the rolling cylinders. The *plane* is identified with $I \times I$, where I is the open interval from -1 to $+1$, i. e., $I \times I$ is the (interior of the) square with side length two. The *rolling* (*standing*) *cylinder* is obtained by identifying the bottom and top (left and right) sides. By identifying the boundaries of I, we obtain $I_\circ$. Then, the standing and the rolling cylinder are defined by $I_\circ \times I$ and $I \times I_\circ$, respectively. Let **RUP** be the set of graphs which can be drawn on the rolling cylinder such that the edge curves do not cross and are monotonically increasing in y-direction. If the edge curves are permitted to be non-decreasing in y-direction, horizontal lines are allowed. Since the top and bottom sides of the fundamental polygon are identified, "upward" means that edge curves wind around the cylinder all in the same direction. Specifically, **RUP** allows for cycles. Accordingly, let **SUP** denote the class of graphs with a planar drawing on the standing cylinder and increasing curves for the edges and let **wSUP** be the corresponding class of graphs with non-decreasing curves. The novelty of **wSUP** graphs are cycles with horizontal curves, whereas **SUP** graphs are acyclic, i. e., $\mathbf{SUP} \subsetneq \mathbf{wSUP}$. In [1] we established that a graph is in **SUP** if and only if it is upward planar on the sphere. These *spherical* graphs were studied in [9, 11, 14, 15]. Finally, let **UP** be the class of upward planar graphs (in the plane) [7, 17]. Note that for **UP** and **RUP** graphs non-decreasing curves can be replaced by increasing ones and the corresponding classes coincide [1].

Upward planar graphs in the plane and on the sphere or on the standing cylinder were characterized by using *acyclic dipoles*. An acyclic dipole is a directed acyclic graph with a single source s and a single sink t. More specifically, a graph G is **SUP**/spherical if and only if it is a spanning subgraph of a planar acyclic dipole [13, 15, 18]. The idea behind acyclic dipoles is that s corresponds to the south and t to the north pole of the sphere. Moreover, a graph G is in **UP** if and only if the dipole has in addition the (s, t) edge [7, 17].

In contrast, there is no related characterization of **RUP** graphs. Acyclic dipoles cannot be used since **RUP** graphs may have cycles winding around the rolling cylinder. However, the idea behind dipoles can be applied indirectly to

RUP graphs, namely, to their duals. For this, we generalize acyclic dipoles to *dipoles* which may also contain cycles.

Section 2 provides the necessary definitions. We develop our new characterization of **RUP** and **SUP** graphs in terms of their duals in Sect. 3. In Sect. 4 we obtain related results for **wSUP** graphs.

2 Preliminaries

The graphs in this work are connected, planar (unless stated otherwise), directed multigraphs $G = (V, E)$ with non-empty sets of vertices V and edges E, where pairs of vertices may be connected by multiple edges. G can be drawn in the plane such that the vertices are mapped to distinct points and the edges to non-intersecting Jordan curves. Then, G has a planar *drawing*. It implies an *embedding* of G, which defines (cyclic) orderings of incident edges at the vertices. In the following, we only deal with embedded graphs and all paths and cycles are simple.

A *face* f of G is defined by a (underlying undirected) circle $C = (v_1, e_1, v_2, e_2, \ldots, v_{k-1}, e_{k-1}, v_k = v_1)$ such that $e_i \in E$ is the direct successor of $e_{i-1} \in E$ according to the cyclic ordering at v_i. The edges/vertices of C are said to be the *boundary* of f and C is a *clockwise traversal* of f. Accordingly, the *counterclockwise traversal* of f is obtained by choosing the predecessor edge at each vertex in the circle. The embedding defines a unique *(directed) dual* graph $G^* = (F, E^*)$, whose vertex set is the set of faces F of G [4]. Let $f \in F$ be a face of G and $e = (u, v) \in E$ be part of its boundary. If the counterclockwise traversal of f passes e in its direction, we say that *f is to the left of e*. If the same holds for e and another face g in clockwise direction, then *g is to the right of e*. For each edge $e \in E$ there is an edge in E^* from the face to the left of e to the face right of e. This definition establishes a bijection between E and E^*. Whenever necessary, we refer to G as the *primal* of G^*. By vertex we mean an element of V, whereas the vertices F of G^* are called faces.

Note that G^* is connected and the dual of G^* is isomorphic to the *converse* G^{-1} of G where all edges are reversed, since G is connected. Hence, an embedding of G implies an embedding of G^*, and vice versa. G and G^{-1} share many properties, see Proposition 1.

An embedding of a graph is an X *embedding* with $X \in \{\mathbf{RUP}, \mathbf{SUP}, \mathbf{wSUP}, \mathbf{UP}\}$ if it is obtained from an X drawing. For every graph in class X, we assume that a corresponding X embedding is given. Given an embedded graph G, a face *f is to the left of a face g* if there is a path $f \rightsquigarrow g$ in G^*. Note that a face can simultaneously lie to the left and to the right of another face, and “to the left” does not directly correspond to the geometric left-to-right relation in a drawing. A cycle in a **RUP** embedding winds exactly once around the cylinder [6]. We say that a face *$f \in F$ lies left (right) of a cycle C* if there is another face $g \in F$ such that f is to the left (right) of g and each path $f \rightsquigarrow g$ in the dual contains at least one edge of C. Each edge/face of f's boundary is then also said to *lie to the left (right) of C*.

Next we introduce graphs which represent the high-level structure of a given graph and which inherit its embedding. Let the equivalence class $[v]$ denote the set of vertices of the strongly connected component containing the vertex $v \in V$ and let $\mathbb{V}$ be the set of strongly connected components of G. The *component graph* $\mathbb{G} = (\mathbb{V}, \mathbb{E})$ of G contains an edge $([v], [w]) \in \mathbb{E}$ for each original edge $(v, w) \in E$ with $[v] \neq [w]$. $\mathbb{G}$ is an acyclic multigraph which inherits the embedding of G. A component $\gamma \in \mathbb{V}$ is a *compound*, if it contains more than one vertex or consists of a single vertex with a loop. Its induced subgraph is denoted by $G_\gamma \subseteq G$. For the sake of convenience, we identify G_γ with γ and call both compound. The set of all compounds is denoted by $\mathbb{V}_C$. Each component $[v]$ that is *not a compound* consists of a single vertex v and is called *trivial component*. A trivial component which is a source (sink) in $\mathbb{G}$ is called *source (sink) terminal* and the set of all terminals is denoted by $\mathbb{T} \subseteq \mathbb{V}$. Based on the component graph, we define the *compound graph* $\overline{\mathbb{G}} = (\mathbb{V}_C \cup \mathbb{T}, \overline{\mathbb{E}})$, whose vertices are the compounds and terminals. Let $u, v \in \mathbb{V}_C \cup \mathbb{T}$ be two vertices of the compound graph. There is an edge $(u, v) \in \overline{\mathbb{E}}$ if there is a path $u \rightsquigarrow v$ in $\mathbb{G}$ which internally visits only trivial components. Note that $\overline{\mathbb{G}}$ is a simple graph. Each edge $\tau \in \overline{\mathbb{E}}$ corresponds to a set of paths in $\mathbb{G}$. Denote by $\mathbb{G}_\tau$ the subgraph of $\mathbb{G}$ which is induced by the set of paths belonging to edge τ. We call τ and its induced graph $\mathbb{G}_\tau$ *transit*. See Fig. 1 for an example, where the fundamental polygon of the rolling cylinder is represented by rectangles with identified bottom and top sides. Based on these definitions, we are now able to define dipoles.

Definition 1. *A graph is a* dipole *if it has exactly one source s and one sink t and its compound graph is a path from s to t.*

Note that similar to the definition of *st*-graphs [7, 17], a dipole is not necessarily planar.

Lemma 1. *Let $G = (V, E)$ be a graph with a source s and a sink t. Then, G is a dipole if and only if every path $s \rightsquigarrow t$ contains at least one vertex of each compound and for every vertex $v \in V$ there are paths $s \rightsquigarrow v$ and $v \rightsquigarrow t$.*

Proof. "$\Rightarrow$": Since G is a dipole, its compound graph $\overline{\mathbb{G}} = (\mathbb{V}_C \cup \mathbb{T}, \overline{\mathbb{E}})$ is a path $p = (v_1, v_2, \ldots, v_k)$ with $s = v_1$ and $v_k = t$ and $(v_i, v_{i+1}) \in \overline{\mathbb{E}}$ for $1 \leq i < k$. (u, v) is an edge in $\mathbb{E}$ if and only if there is a path $u \rightsquigarrow v$ in the component graph $\mathbb{G}$ which internally visits only trivial components. Any path in $s \rightsquigarrow t$ which does not visit all the compounds of G in exactly the order as given by p would imply an edge $(v_i, v_j) \in \overline{\mathbb{E}}$ with $j \neq i + 1$ and, hence, the compound graph would be no path.

Note that there is at least one path $p = s \rightsquigarrow t$ in G. Since p contains at least one vertex of each compound γ and γ is a set of strongly connected vertices, there are also paths $s \rightsquigarrow v$ and $v \rightsquigarrow t$ for each vertex v in compound γ. Hence, what is left to show is that there also paths $s \rightsquigarrow \hat{v}$ and $\hat{v} \rightsquigarrow t$ for each trivial component $\hat{v} \neq s, t$. Assume for contradiction that there is no path $s \rightsquigarrow \hat{v}$. Let $\hat{V} \subseteq V$ be the set of vertices $u \in \hat{V}$ for which there is a path $u \rightsquigarrow \hat{v}$. No vertex u of a compound can be in $\hat{V}$ since, otherwise, there would be a path $s \rightsquigarrow u$ and,

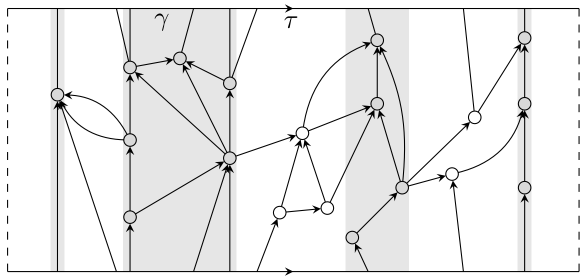

(a) Graph $G \in \mathbf{RUP}$

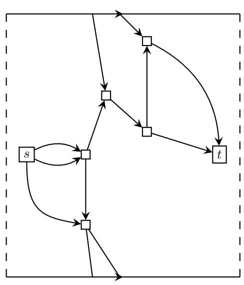

(b) The dual of the second compound γ of G

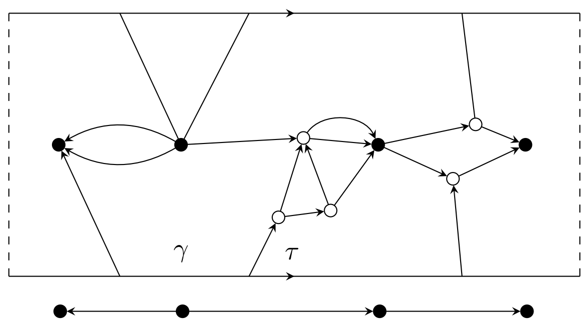

(c) The component graph $\mathbb{G}$ and the compound graph $\overline{\mathbb{G}}$ of G

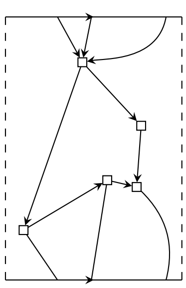

(d) The dual of the second transit τ of G

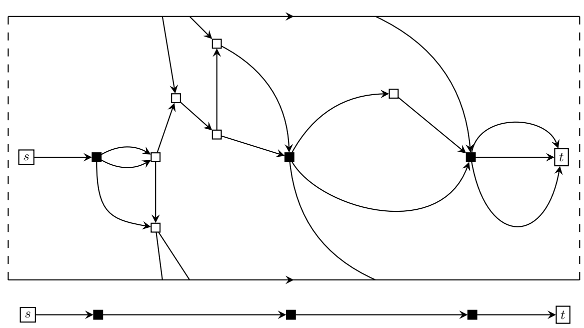

(e) The component graph $\mathbb{G}^*$ and the compound graph $\overline{\mathbb{G}^*}$ of the dual G^* with $s, t \in \mathbb{T}$

Fig. 1. A **RUP** example

therefore, $s \rightsquigarrow \hat{v}$ since $u \rightsquigarrow \hat{v}$. Hence, the subgraph induced by $\hat{V}$ is an acyclic subgraph of G which does not contain s. Therefore, there must be a source $\hat{s} \in \hat{V}$ with $\hat{s} \neq s$; a contradiction. By the same reasoning it can be shown that there is a path from every vertex to the sink t.

"$\Leftarrow$": Since $s \rightsquigarrow v$ and $v \rightsquigarrow t$ for every $v \in V$, there is a path $s \rightsquigarrow t$. Let $p = (v_1, v_2, \ldots, v_k)$ with $s = v_1$ and $v_k = t$ and $(v_i, v_{i+1}) \in \overline{\mathbb{E}}$ for $1 \leq i < k$ be a path in the compound graph $\overline{\mathbb{G}}$. Note that p is a Hamiltonian path since every path visits each compound by assumption. Suppose now that the compound graph $\overline{\mathbb{G}}$ is not a path $s \rightsquigarrow t$. Since $\overline{\mathbb{G}}$ is acyclic, there is a (transitive) edge $(v_i, v_j) \in \overline{\mathbb{E}}$ with $1 < i + 1 < j \leq k$. However, then there exists a path $v_i \rightsquigarrow v_j$ and, hence, also a path $s \rightsquigarrow t$ which does not visit any vertex from compounds $v_{i+1}, \ldots, v_{j-1}$; a contradiction. □

Proposition 1. *A graph G is (i) acyclic, (ii) strongly connected, (iii) upward planar, or (iv) a dipole if and only if the same holds for its converse G^{-1}.*

Thus, in the subsequent statements on the relationship between a graph G and its dual G^*, the roles of G and G^* are interchangeable.

Lemma 2. *A graph G is acyclic if and only if its dual G^* is strongly connected.*

The proof is deduced from the one for polynomial solvability of the feedback arc set problem on planar graphs as given in [4].

Proof. "$\Rightarrow$": Consider an embedding of $G = (V, E)$ and suppose that the dual $G^* = (F, E^*)$ is not strongly connected. Then, G^* has a dicut $(X, F \setminus X)$, i. e., a subset $X \subsetneq F$, such that there are no edges in G^* with source in $F \setminus X$ and target in X. Let L be the set of edges of $(X, F \setminus X)$, i. e., those that are directed from X to $F \setminus X$. Since G^* is connected, this set is not empty. Let F_L^* be the set of faces left and right of the edges in L. Consider any traversal of the boundary C of a face $f \in F_L^*$ starting with the edge $l_1 = (u, v) \in L$ with $u \in X$ and $v \in F \setminus X$. In order to close C, at least one edge $l_2 \in L$ is traversed against its direction. Note that l_1 and l_2 may coincide. Recall that the vertices of G are the faces of G^*. Hence, l_1 and l_2 imply an incoming edge to f and an outgoing edge from f with their other endpoints in the opposite faces of l_1 and l_2. Due to the definition of F_L^*, both opposite faces are in F_L^*. Consequently, the subgraph induced by F_L^* has neither sources nor sinks and thus contains a cycle, which is a contradiction.

"$\Leftarrow$": Suppose that G has a cycle. Then, the cycle separates the inner and the outer faces of G, which leads to a dicut in G^* and, hence, G^* cannot be strongly connected; a contradiction. □

We need some additional notation. For $\epsilon > 0$ and an arbitrary metric d on a surface $\mathbb{S}$, the *ϵ-environment* of a point set $P \subseteq \mathbb{S}$ is the union of all open balls with radius ϵ around points in P, i. e.,

$$\bigcup_{p \in P} \{q \in \mathbb{S} : d(p, q) < \epsilon\} . \tag{1}$$

For $x, y \in \mathbb{R}$, let $x \bmod y$ be the non-negative remainder $x - y\lfloor \frac{x}{y} \rfloor$ of dividing a real x by y. By *rotation around the rolling cylinder* we mean the transformation

$$T : I \times I_\circ \to I \times I_\circ : (x, y) \mapsto (x, (y + \Delta + 1) \bmod 2 - 1) \tag{2}$$

for some $\Delta \in \mathbb{R}$, which also can be considered as a translation in y-direction by Δ. Note that rotating a **RUP** drawing does neither affect its upwardness nor its implied embedding.

3 RUP and SUP Graphs and their Duals

We consider **RUP** graphs, i. e., upward planar graphs on the rolling cylinder, and characterize them in terms of their duals. Our main result is:

Theorem 1. *A graph G is a **RUP** graph if and only if G is a spanning subgraph of a planar graph H without sources or sinks whose dual H^* is a dipole.*

The theorem is proved by a series of lemmata which are also of interest in their own. For our first observation, consider the **RUP** drawing of graph G in Fig. 1(a), where all vertices within a compound are drawn on a shaded background. The component graph $\mathbb{G}$ of G is displayed in Fig. 1(c) along with its compound graph $\overline{\mathbb{G}}$ below, where the compounds are shaded black. Note that $\overline{\mathbb{G}}$ has the structure of an (undirected) path. Due to Lemma 2, each transit, i. e., edge in $\overline{\mathbb{G}}$, becomes a compound and each compound, i. e., vertex in $\overline{\mathbb{G}}$, becomes a transit in the dual G^* of G. Hence, the path-like structure of $\overline{\mathbb{G}}$ must carry over to the compound graph $\overline{\mathbb{G}^*}$ of G^*. Moreover, since all cycles in the **RUP** drawing have the same orientation, i. e., they all wind around the cylinder in the same direction, the transits in G^* point into the same direction. Also note that G contains neither sources nor sinks, i. e., both the left and right border of the drawing are directed cycles C_l and C_r, respectively. Hence, in the dual G^* of G, the face to the left of C_l is a source s and the face to the right of C_r is a sink t. All these observations together indicate that the compound graph of G^* is a path $s \rightsquigarrow t$, i. e., G^* is a dipole. Indeed, this can be seen for the example in Fig. 1(e), where the component graph of G^* and its compound graph are depicted.

Lemma 3. *The dual G^* of a **RUP** graph G without sources and sinks is a dipole.*

Proof. First, we show that G^* contains at least one source and one sink. Let C_l be a leftmost cycle according to the embedding of G, i. e., there is no vertex to the left of C_l which belongs to another cycle. Since G contains neither sources nor sinks, it contains at least one directed cycle. Hence, C_l exists. We now show that there is not only no cycle to the left of C_l but also no vertex at all. Otherwise, assume for contradiction that there is a non-empty subgraph $G_l \subseteq G$ induced by the vertices lying to the left of C_l. Since C_l is the leftmost cycle, G_l does not contain any cycle and, hence, it must either contain a source or a sink. This is a contradiction to our assumption. Consequently, the face s bounded by C_l to the left is a source in G^*. By a similar reasoning, it can be shown that there exists a sink t in G^*.

In the following, we use Lemma 1 to show that $G^* = (F, E^*)$ is a dipole. Let $s \in F$ be a source bounded by a leftmost cycle C_l as defined in the previous paragraph. We show that there is path from s to every other face $f \in F$. Let F_s be the set of faces reachable from s. Assume for contradiction that $F_s \subsetneq F$. F is partitioned into F_s and $X = F \setminus F_s$. By the definition of F_s, any edge connecting a vertex in F_s with a vertex in X must point from X to F_s. This set of edges, denoted by $E' \subsetneq E^*$, is not empty since G^* is connected. On the rolling cylinder, the edges E' correspond to a cycle that winds around the cylinder in the opposite orientation of C_l, which contradicts our assumption of a **RUP** embedding. The situation is illustrated in Fig. 2, where the shaded area covers the faces F_s. Analogously, it can be shown that there is a path from every face to sink t.

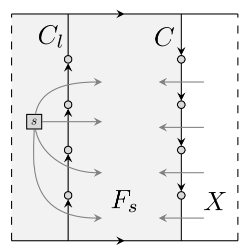

Fig. 2. Two cycles winding around the cylinder with opposite orientations

Since every face is reachable from s, there is a path $s \rightsquigarrow t$ in G^*. It remains to show that every path $s \rightsquigarrow t$ contains a vertex of each compound in the compound graph $\overline{\mathbb{G}^*} = (\mathbb{F}_C \cup \{s, t\}, \overline{\mathbb{E}})$, where $\mathbb{F}_C$ is the set of compounds in $\mathbb{G}^*$. If G is strongly connected, then G^* is acyclic and connected. Then, $\overline{\mathbb{G}^*}$ contains no compounds at all and we are done. Hence, we assume that $\overline{\mathbb{G}^*}$ contains at least one compound γ. Compound γ contains at least one cycle C that winds exactly once around the cylinder. In the embedding of G^* on the rolling cylinder, C divides the cylinder into two regions, where one contains s and the other t. Hence, each path $s \rightsquigarrow t$ must contain a vertex of C and, thus, a vertex of γ. $\square$

For the following lemma, there is a physical interpretation: Consider an upward drawing of a planar acyclic dipole on the standing cylinder and suppose that an electric current flows from the bottom to the top of the cylinder in direction of the edges. This current induces a magnetic field wrapping around the standing cylinder. Intuitively, by Lemma 4, we can show that a dipole's dual is upward planar with respect to the induced magnetic field, i. e., its embedding is a **RUP** embedding.

Lemma 4. *The embedding of a strongly connected graph G is a* ***RUP*** *embedding if and only if its dual G^* is an acyclic dipole.*

Proof. "$\Rightarrow$": Follows directly from Lemmata 2 and 3.

"$\Leftarrow$": Since G^* is acyclic, G is strongly connected by Lemma 2 and Proposition 1. To show that G is **RUP** embedded we inductively construct a **RUP** drawing of G on the fundamental polygon of the rolling cylinder $I \times I_\circ$ such that all edge curves are upward and the embedding of G^* is preserved.

Let $f_1, f_2, \ldots, f_k \in F$ be the faces of G in a topological ordering of G^*, i. e., f_1 is the single source and f_k is the single sink of G^*. Let G_i $(1 \leq i \leq k)$ be the embedded subgraph of G induced by the faces $f_1, \ldots, f_i$, i. e., G_i contains exactly those edges and vertices bounding the faces $f_1, \ldots, f_i$. Then, $G_k = G$.

The basic idea of the inductive proof is to add new edges to G_i such that f_{i+1} is enclosed and lies to the left of all newly added edges. To assure a plane drawing, the x-coordinates of the newly added vertices are strictly greater than all x-coordinates of all previously added vertices. In the following let $x_1, x_2, \ldots, x_k \in I$ be a sequence of strictly increasing x-coordinates, i. e., $x_i < x_{i+1}$ for all $1 \leq i < k$.

As induction invariant, the following conditions hold for the drawing Γ_i of each G_i. Γ_i is a **RUP** drawing which respects the embedding of G and lies in $[x_1, x_i] \times I_\circ$. Additionally, the dual G_i^* of each G_i is a planar, acyclic dipole. Especially, the right border of Γ_i is a directed cycle C_r and all faces $f_1, \ldots, f_i$ are on the left of C_r.

For $i = 1$, G_1 consists of a single cycle with $\mathrm{d}^+(f_1)$ edges, since f_1 is a source in G^* with out-degree $\mathrm{d}^+(f_1)$. All vertices of the cycle are assigned the x-coordinate x_1 and y-coordinates according to the cyclic order of the edges around f_1, see Fig. 3(a) for an example. The drawing of G_1 guarantees the induction invariants.

For $i = k - 1$, the faces $f_1, \ldots, f_{k-1}$ lie to the left of cycle C_r, which is the right border of Γ_{k-1}. Then, the right face of C_r is f_k. Thus, Γ_{k-1} is a drawing of G and we are done.

Now assume that $1 < i < k-1$. In the embedding of G^* and, consequently, in the embedding of G_{i+1}^*, all incoming edges are consecutive in the cyclic order of the edges around f_{i+1}, and accordingly for the outgoing edges. This follows from the fact that G^* is an embedded planar acyclic dipole and, hence, its embedding is quasi upward planar [5]. Denote by $e_1^-, \ldots, e_p^-$ and $e_1^+, \ldots, e_q^+$ the incoming and outgoing edges of f_{i+1}, respectively, ordered according to the embedding of G^*. Note that f_{i+1} has at least one incoming edge as otherwise it would be a source different from f_1, which contradicts the assumption of G^* being a dipole. Analogously, f_{i+1} has at least one outgoing edge. Due to the topological ordering of the faces, all faces that have an outgoing edge to f_{i+1} are already present in the drawing of G_i. Additionally, all these edges e_j^- are part of the cycle C_r at the right border of Γ_i (black solid vertices in Fig. 3(b)). Otherwise, there is an edge e_j^- $(1 \leq j \leq p)$ such that one of its end points lies to the left of C_r. However, then either e_j^- cannot be an edge bounding f_{i+1} or G_i does not respect the embedding of G. Moreover, the edges $e_1^-, \ldots, e_p^-$ correspond to a directed path $p^- = (v_1^-, \ldots, v_{p+1}^-)$ in G_i, which is part of C_r (see again Fig. 3(b)), since

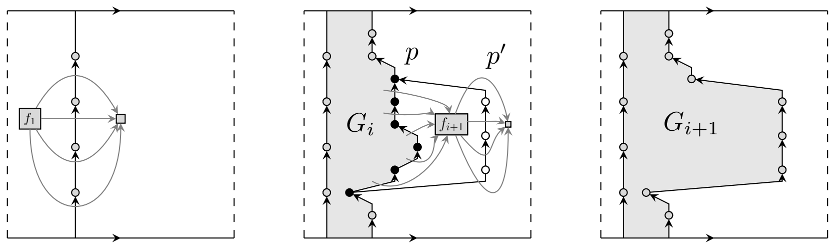

(a) Base case (b) Induction step (c) After the induction step

Fig. 3. Inductive construction of a **RUP** drawing from its dual

Γ_i respects the embedding of G_i^*. Let $a \in I_\circ$ ($b \in I_\circ$) be the y-coordinate of v_1^- (v_{p+1}^-). Note that v_{p+1}^- must lie "above" v_1^- as p must be drawn upward. W.l.o.g., we assume that $a < b$.[1] Accordingly, the edges $e_1^+, \dots e_q^+$ correspond to a path $p^+ = (v_1^+, \dots, v_{q+1}^+)$ in G_{i+1}. Note that $v_1^+ = v_1^-$ and $v_{q+1}^+ = v_{p+1}^-$ since p^+ and p^- together bound face f_{i+1}. Assign to each vertex $v_2^+, \dots, v_q^+$ the x-coordinate x_{i+1}. For every vertex v_j^+ choose a y-coordinate y_j^+ such that $a < y_j^+ < b$ for all $2 \leq j \leq q$ and $y_j^+ < y_{j+1}^+$ for $2 \leq j < q$. Now the edges of p^+ can be drawn upward as straight lines with a single bend at the first and the last edge of the path, see Figs. 3(b) and (c). For the position of the bend in the edges e_1^+ and e_q^+ choose as y-coordinate some value in the intervals (a, y_2^+) and (y_q^+, b), respectively, such that the straight lines from the end points on p^- cause no crossing with any other edge from p^-. Note that this is always possible due to the construction of the drawing. For the x-coordinate of the bends choose x_{i+1}. If p^+ consists of a single edge, this edge has two bends. The resulting drawing Γ_{i+1} of G_{i+1} is a **RUP** drawing respecting the embedding of G and it lies within $[x_1, x_{i+1}] \times I_\circ$. In Γ_{i+1} there is a newly formed cycle C'_r containing p^+ on the right border of the drawing such that all faces $f_1, \dots, f_{i+1}$ lie to the left of C'_r, see Fig. 3(c). Thus, the dual G_{i+1}^* is again a planar, acyclic dipole. $\square$

Since each **SUP** graph is a subgraph of a planar, acyclic dipole [15], Lemma 4 implies:

Corollary 1. *The dual G^* of a strongly connected **RUP** graph G is in **SUP**.*

Consider again the component graph $\mathbb{G}$ and its compound graph $\overline{\mathbb{G}}$ in Fig. 1(c) of the **RUP** graph G in Fig. 1(a). In the dual G^* of G, compounds and transits of G swap their roles, i.e., compounds become transits and vice versa, cf. Fig. 1(e). As a compound of G is a strongly connected **RUP** graph, its dual is an acyclic dipole by Lemma 4. For instance, consider the second compound γ in Fig. 1(a),

[1] If $b > a$, rotate the drawing around the cylinder until $a < b$.

i. e., the vertices on the second shaded area labeled with γ. Its dual is indeed an acyclic dipole as depicted in Fig. 1(b). For the transits, the same holds but with swapped roles, i. e., the dual of a transit is a strongly connected **RUP** graph. As an example, the dual of the second transit τ in Fig. 1(a) is shown in Fig. 1(d) and it is indeed a strongly connected **RUP** graph. The following lemma subsumes these observations.

Lemma 5. *Let G be a **RUP** graph without sources and sinks and $\overline{\mathbb{G}} = (\mathbb{V}_C, \overline{\mathbb{E}})$ be its compound graph. Then,*

*(i) the dual of each compound $\gamma \in \mathbb{V}_C$ is a planar, acyclic dipole and, thus, it is in **SUP**.*
*(ii) each transit $\tau \in \overline{\mathbb{E}}$ is a planar, acyclic dipole and, thus, its dual is a strongly connected **RUP** graph.*

Proof. *(i)*: The induced subgraph G_γ of a compound γ is strongly connected and has a **RUP** embedding. According to Lemma 4, its dual G^*_γ is a planar acyclic dipole and, hence, its embedding is a **SUP** embedding.

(ii): The graph G_τ induced by a transit τ is subgraph of $\mathbb{G}$ induced by paths from a source s to a sink t via trivial components only. Note that s may be a compound or a trivial component and the same holds for t. Since all other vertices of G_τ are trivial and G contains neither sources nor sinks, s is the only source and t the only sink. Furthermore, G_τ is acyclic since it is a subgraph of the component graph $\mathbb{G}$. Hence, G_τ is an embedded, acyclic dipole. By Lemma 4 and Proposition 1 we can conclude that G^*_τ is strongly connected and a **RUP** graph. □

By Lemma 3 we have seen that the dual of a **RUP** graph that contains neither sources nor sinks is a dipole. Also the converse holds:

Lemma 6. *A graph G without sources and sinks is a **RUP** graph if its dual G^* is a dipole.*

Consider again the example **RUP** graph in Fig. 1(a) and the compound graph $\overline{\mathbb{G}^*}$ of its dual G^*. Since G^* is a dipole, $\overline{\mathbb{G}^*}$ is a path $p = (s, \tau^*_1, \gamma^*_1, \tau^*_2, \gamma^*_2, \dots, \tau^*_4, t)$ consisting of compounds γ^*_i, transits τ^*_j, and two terminals s and t. Note that each element on p corresponds to a subgraph in the primal G, i. e., for each γ^*_i there is a transit τ_i in G and for each τ^*_j there is a compound γ_j in G. In the proof of Lemma 6, we construct a **RUP** drawing of G by subsequently processing the elements of p. We start with transit τ^*_1, whose induced subgraph in G^* is an acyclic dipole, and obtain a **RUP** drawing of γ_1 which respects the given embedding by Lemma 4. Then we proceed with γ^*_1, a compound in G^*, for which we obtain a **SUP** drawing of τ_1 which respects the given embedding by Lemma 4. However, this **SUP** drawing is upward only with respect to the x-direction, i. e., from left to right. We transform this drawing, while preserving its embedding, such that it is also upward in y-direction. The so obtained drawing of τ_1 is then attached to the right border of the drawing of γ_1. Then, the drawing of γ_2 is attached to the right side of τ_1 and so forth until we reach t. Note that since all transits τ_j^* point into the same direction in $\overline{\mathbb{G}^*}$, i. e., from s to t, all cycles of the compounds in G have the same orientation in the obtained drawing, i. e., they all wind around the cylinder in the same direction.

Before we can prove Lemma 6, we need a supporting lemma.

Lemma 7. *Let Γ be a* ***SUP*** *drawing of* ***SUP*** *graph G. Then, there is a* ***RUP*** *drawing of G which implies the same embedding as Γ.*

Proof. Consider a **SUP** drawing Γ of $G = (V, E)$ on the standing cylinder. If we negate the x-coordinate of each point in Γ and swap the axes of the fundamental polygon, we yield a drawing of G on the rolling cylinder where all edge curves are increasing monotonically in x-direction.[2] We assume that all curves are differentiable. Otherwise, (the finitely many [6]) critical points must be excluded from the following reasoning. The idea is to shear the drawing vertically such that the curves are increasing additionally in y-direction. Then, we have a **RUP** drawing.

Edge curves are usually represented by continuous maps of the interval $[0, 1]$ to points of the surface. In the following, we take a different approach and use real-valued partial functions as they allow for a simpler mathematical treatment. Note that no curve has multiple points with the same x-coordinate. Thus, we can represent each curve of an edge e by a differentiable and partial function f_e from I to $\mathbb{R}$, whose domain is a closed interval, which we denote by $\mathrm{dom}(f_e)$. Then, let the drawing of e be the point set

$$\{(x, y) \in \mathrm{dom}(f_e) \times I_\circ : y = f_e(x) \bmod 2 - 1\} . \tag{3}$$

At a first glance, this definition may seem odd, but mapping the image of f_e via $(\cdot \bmod 2) - 1$ to y-coordinates on the rolling cylinder allows differentiable real-valued functions to represent all curves increasing in x-direction, even if they wind multiple times around the cylinder. See Fig. 4 for an example.

Let a define the least gradient of all functions representing an edge of G, i. e.,[3]

$$a = \min_{e \in E} \min_{x \in \mathrm{dom}(f_e)} f_e'(x) . \tag{4}$$

If $f_e' > 0$ for all functions f_e representing an edge curve, the curves are increasing monotonically in y-direction and we are done. Otherwise, assume for the rest of the proof that there is at least one function f_e with a non-positive gradient at some point, i. e., $a \leq 0$.

Now we are ready to define the shearing S by (for simplicity, we omit the application of $(\cdot \bmod 2) - 1$ to the y-coordinate)

$$S : I \times I_\circ \to I \times I_\circ : (x, y) \mapsto (x, y + (1 - a) \cdot x) . \tag{5}$$

[2] Swapping the coordinates alone would alter the embedding as it reverses the cyclic order of incident edges. Negation cancels this effect.

[3] If there are critical points, then choose $a = \min_{e \in E} \inf_{x \in \mathrm{dom}(f_e)} f_e'(x)$.

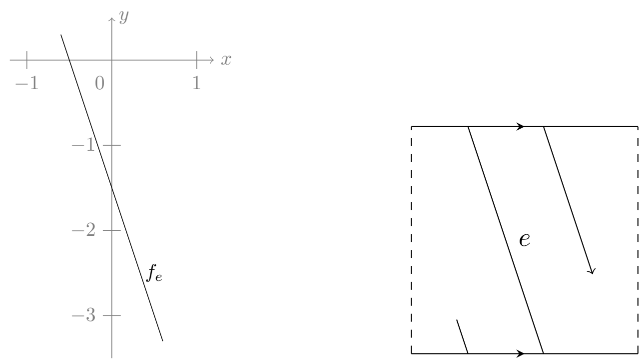

(a) The function f_e representing edge e

(b) The edge curve of e

Fig. 4. Edge curve represented by the partial function $f_e : I \times \mathbb{R} : x \mapsto -3(x + \frac{1}{2})$ with $\mathrm{dom}(f_e) = [-0.6, 0.6]$

Denote by $S[\Gamma]$ the image of Γ under S. Observe that S preserves the embedding of G, i. e., $S[\Gamma]$ is plane and implies the same embedding as Γ. It remains to show that $S[\Gamma]$ is a **RUP** drawing. Let f represent an edge curve in Γ. Then the function g representing the corresponding edge curve in the transformed drawing $S[\Gamma]$ is

$$g : \mathrm{dom}(f) \to \mathbb{R} : x \mapsto f(x) + (1 - a) \cdot x\,. \tag{6}$$

We derive

$$\underset{x \in \mathrm{dom}(f)}{\forall} g'(x) = f'(x) + 1 - a > 0 \tag{7}$$

since $f'(x) \geq a$ by the definition of a. $\square$

We are now able to prove Lemma 6.

Proof. If G consists of a single compound, then it is strongly connected as it contains neither sources nor sinks. Thus, G^* is an acyclic dipole, the embedding of G is a **RUP** embedding according to Lemma 4, and we are done. In the following we assume that G contains at least two compounds.

Let $s \in F$ and $t \in F$ be the source and the sink of $G^* = (F, E^*)$, respectively. We show that each cycle C in G separates the faces s and t, i. e., s lies to the left and t to the right of C or vice versa. Assume for contradiction that the left side of C neither contains s nor t. Let $f \in F$ be a face to the left of C. By the definition of the dual graph and as C is a directed cycle and f lies to its left, we can conclude that there is no path from any face to the right of C to f in G^*. Especially, there is no path from $s \rightsquigarrow f$. However, by Lemma 1, there must be a path $s \rightsquigarrow f$; a contradiction. Hence, both s and t must be to the left of C. By a similar reasoning, there is no path $f \rightsquigarrow t$; again a contradiction.

Let C and C' be two cycles in G belonging to different compounds. We show that C and C' have the same orientation, i. e., either s is to the left or to the right of both of them. As C and C' belong to different compounds, they are (vertex) disjoint. As both C and C' separate s from t, they determine three disjoint, non-empty regions in the embedding of G. One region containing face s, one containing t, and a third region containing the faces $\hat{F} \subsetneq F$. If C and C' have opposite orientations, then we obtain the same situation as in the proof of Lemma 3 and as displayed in Fig. 2, where t is situated within region X. Let $f \in \hat{F}$. Then, either there is no path $s \rightsquigarrow f$ or no path $f \rightsquigarrow t$, which contradicts the assumption that G^* is a dipole (Lemma 1). Hence, C and C' have the same orientation and, consequently, so have all cycles of G. From now on, we assume w. l. o. g. that s is to the left and t is to the right of all cycles in G.

By the reasoning in the previous paragraph, we can also conclude that the compounds of G properly nest, i. e., there is an ordering $\gamma_1, \gamma_2, \ldots, \gamma_k$ of the compounds $\mathbb{V}_C$ of G with the following properties. The region to the left of any cycle in compound γ_i ($1 < i < k$) contains all vertices of compounds $\gamma_1, \ldots, \gamma_{i-1}$ and the region to the right of any cycle in compound γ_i contains all vertices of compounds $\gamma_{i+1}, \ldots, \gamma_k$. Compound γ_1 is the leftmost compound in the sense that no compound is to its left side and all other compounds are to its right side. In the same sense, γ_k is the rightmost compound.

In the following, consider a drawing of G in the plane which respects the given embedding. Fig. 5 shows the principle structure of such a drawing. The compounds are displayed as rings which are shaded gray and the arrows at the rings' borders indicate the direction of the compounds' cycles. Face s is situated in the middle and lies left to all compounds $\gamma_1, \ldots, \gamma_k$. Face t is the outer face to the right of all compounds. Let γ_i and γ_j be two compounds of G with $1 \leq i < j \leq k$ such that $j - k > 1$, i. e., in the ordering of the compounds, there is at least one compound between γ_i and γ_j. We now show that there is no transit between γ_i and γ_j, i. e., neither $(\gamma_i, \gamma_j) \in \overline{\mathbb{E}}$ nor $(\gamma_j, \gamma_i) \in \overline{\mathbb{E}}$. Assume for contradiction that a transit $\hat{\tau} = (\gamma_i, \gamma_j) \in \overline{\mathbb{E}}$ exists (for the converse edge, the following reasoning proceeds analogously). Then, there is a path p from a vertex in γ_i to a vertex in γ_j which visits internally only trivial components. However, there is at least one compound γ_l between γ_i and γ_j ($i < l < j$). In the plane drawing of G, this implies that p must visit at least one vertex of γ_l before it can reach γ_j, which is a contradiction since p must be a path via trivial components only. For instance, in Fig. 5, the path of transit $\hat{\tau} = (\gamma_1, \gamma_3) \in \overline{\mathbb{E}}$ must contain at least one vertex of γ_2 due to planarity.

Further, since G is connected, there must be a transit between directly adjacent compounds γ_i and γ_{i+1}, i. e., for all i with $1 \leq i < k$, either $(\gamma_i, \gamma_{i+1}) \in \overline{\mathbb{E}}$ or $(\gamma_{i+1}, \gamma_i) \in \overline{\mathbb{E}}$. In the following, let $\gamma_1, \tau_1, \gamma_2, \ldots, \tau_{k-1}, \gamma_k$ be the sequence of compounds and transits in G such that τ_i is the transit connecting compounds γ_i and γ_{i+1}. Analogously, let $\tau_1^*, \gamma_1^*, \tau_2^*, \ldots, \gamma_{k-1}^*, \tau_k^*$ be the sequence of compounds and transits in G^* in order of the path from the source to the sink in G^*.

By Lemma 4 we know that each compound in G has a **RUP** embedding. Each transit is an acyclic, planar dipole by Lemma 4, which also has a **RUP**

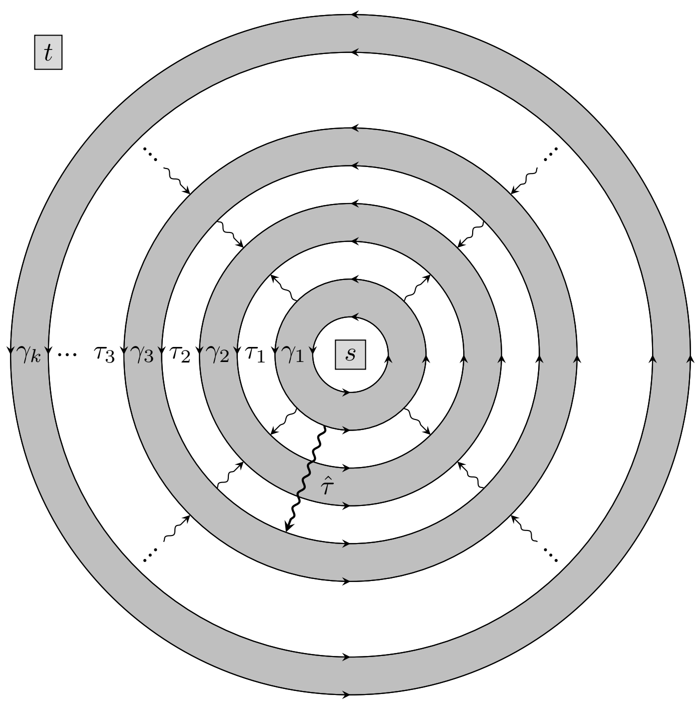

Fig. 5. Plane drawing of a graph G whose dual is a dipole

embedding as shown by Lemma 7. We conclude the proof by showing that the **RUP** embeddings of the individual compounds and transits can be merged into a single consistent **RUP** embedding of the original graph G.

We construct a **RUP** drawing of G by subsequently processing the elements in the order $\gamma_1, \tau_1, \gamma_2, \tau_2, \ldots, \gamma_k$. For an example, see Fig. 6. As scaling a drawing in x-direction does not impair its upwardness in y-direction [1], we do not have to bother with the width of the rolling cylinder.

We start with transit τ_1^*, which is an acyclic dipole, and obtain a **RUP** drawing $\mathit{\Gamma}_{\gamma_1}$ of γ_1, which respects the given embedding by Lemma 4, see Fig. 6(b). Note that G_{γ_1} has a rightmost cycle C_1 defined by its embedding. Then, we proceed with γ_1^*, a compound in G^*, for which we obtain a **SUP** drawing $\mathit{\Gamma}_{\tau_1}$ of $\mathbb{G}_{\tau_1}$ which respects the given embedding by Lemma 4. First assume that τ_1 is directed from left to right, i. e., $\tau_1 = (\gamma_1, \gamma_2)$. We shear $\mathit{\Gamma}_{\tau_1}$ as shown in Lemma 7 such that it becomes a **RUP** drawing and place it to the right of $\mathit{\Gamma}_{\gamma_1}$. However, $\mathbb{G}_{\tau_1}$ is not a subgraph of G but of its component graph. $\mathbb{G}_{\tau_1}$ consists of a source s, a sink t, and other vertices, which are neither a source nor a sink. For the latter there is a one-to-one correspondence to the vertices of G, as they represent trivial components of size 1. So we simply identify each component with the single vertex it contains. Thus, we reinterpret the drawing of $\mathbb{G}_{\tau_1}$ effectively as a drawing of τ_1. s is the compound γ_1 and its incident edges correspond to edges in the original graph, which are incident to vertices in γ_1, more precisely to those of C_1. Therefore, we remove s and all points of its incident edge curves from $\mathit{\Gamma}_{\tau_1}$ within a rectangular ϵ-environment (for a suitably small ϵ under the maximum metric). This results in edge curves starting in "cutting points" rather than in s (Fig. 6(d)). We rotate $\mathit{\Gamma}_{\tau_1}$ around the rolling cylinder such that the y-coordinate of the cutting points is greater than the y-coordinates of any of the vertices in γ_1. Let (u, v) be the edge in G corresponding to that edge in $\mathbb{G}_{\tau_1}$ whose edge curve has the leftmost cutting point. Recall that $u \in C_1$ and $v \in G_{\tau_1}$. Next rotate the drawing of γ_1 such that u is the topmost vertex, but has a smaller y-coordinate than the cutting points. Since both the embedding of C_1 implied by $\mathit{\Gamma}_{\gamma_1}$ and the embedding of $\mathbb{G}_{\tau_1}$ implied by $\mathit{\Gamma}_{\tau_1}$ obey the initial planar embedding of G, the order of the cutting points from right to left corresponds to the order of the vertices in C_1 from bottom to top. Hence, we can connect the vertices of C_1 with edge curves increasing monotonically in y-direction to the respective cutting points without introducing crossings.

The resulting drawing $\mathit{\Gamma}'$, see Fig. 6(e), forms the basis for the next step, where we obtain (again as in Lemma 4) a **RUP** drawing $\mathit{\Gamma}_{\gamma_2}$ of graph G_{γ_2} induced by the compound γ_2, and place it to the right of $\mathit{\Gamma}'$. In a similar way, we remove t from the drawing and reconnect the resulting cutting points to the respective vertices in the leftmost cycle of G_{γ_2}.

If, contrary to our aforementioned assumption, a transit is directed from right to left, we proceed similarly except that we switch the roles of s and t and rotate the cutting points around t to the bottom rather than the top.

Analogously, we proceed with $\tau_2, \gamma_3, \tau_3, \gamma_4, \ldots$ until we have processed all components, resulting in a **RUP** drawing of G. □

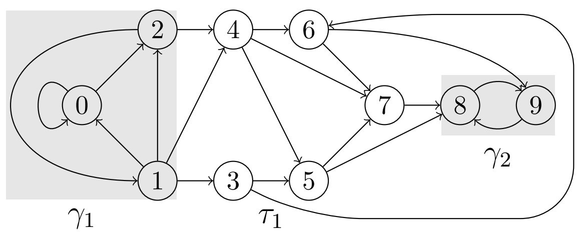

(a) Graph G without sources and sinks, whose dual is an acyclic dipole. The shaded subgraphs are compounds.

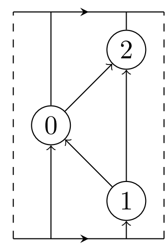

(b) A **RUP** embedding of γ_1

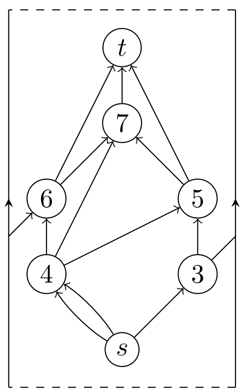

(c) A **SUP** embedding of τ_1

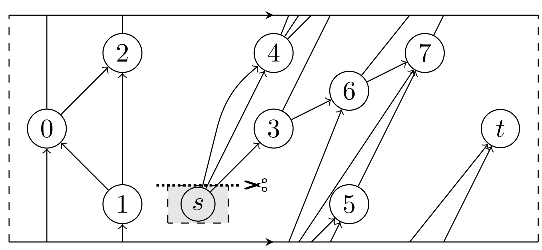

(d) Γ_{τ_1} has been sheared to become a **RUP** drawing and placed to the right of Γ_{γ_1}.

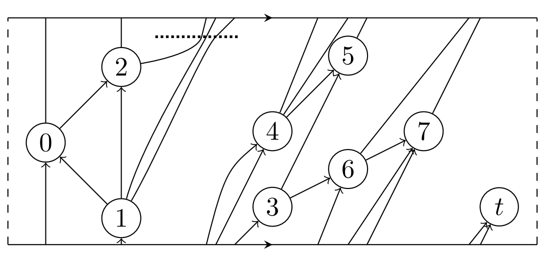

(e) Intermediate result Γ' after the drawings of γ_1 and τ_1 have been merged.

Fig. 6. Construction of a **RUP** drawing

Lemmata 3 and 6 both require that the graph at hand contains neither sources nor sinks. At a first glance, this requirement seems to be a strong limitation. However, in the following lemma we show that each **RUP** graph can be augmented by edges such that all sources and sinks vanish while still preserving **RUP** embeddability.

Lemma 8. *A **RUP** graph G is a spanning subgraph of a **RUP** graph H without sources and sinks.*

Proof. Consider an upward drawing of G on the rolling cylinder. We iteratively add edges until all vertices have both ingoing and outgoing edges. Let t be a sink of G. Shoot a ray from the position of t in upward direction and determine where it first meets some point p of the drawing. If p belongs to a vertex v, we can introduce a geodesic edge, i. e., a straight line on the fundamental polygon, from t to v. Note that $v = t$ if no other vertex or edge has a point with the x-coordinate of p, such that the ray wraps exactly once around the cylinder. If p belongs to an edge (u, v), proceed as follows. The drawing of (u, v) has an ϵ-environment which contains no other point of the drawing except within the ϵ-environment of u and v. Thus, we can route a new edge (t, v) in upward direction and without introducing crossings, which goes first from t towards p on the ray, then runs alongside (u, v) such that it finally meets v. Analogously, we add incoming edges to the sources of G. □

The proof of Theorem 1 is now complete. The only-if direction follows from Lemmata 8 and 3 and the if direction is a consequence of Lemma 6 and the fact that every subgraph of a **RUP** graph is a **RUP** graph.

4 wSUP Graphs and their Duals

We now turn to spherical graphs and upward planar embeddings on the standing cylinder. These graphs were characterized as spanning subgraphs of planar, acyclic dipoles [13,15,18]. We already provided a new characterization for **SUP** in terms of dual graphs in Lemma 4 in combination with Proposition 1. Now we consider graphs which have a weakly upward planar drawing on the standing cylinder. These graphs have not been characterized before.

For a start, consider an upward drawing of a **wSUP** graph. If there are cycles, they must wind around the cylinder horizontally, which leads us to the following observation.

Lemma 9. *Let G be a graph in **wSUP**. Then, all cycles of G are (vertex) disjoint.*

Proof. Consider an arbitrary, weakly upward planar drawing of G on the standing cylinder. If G has cycles, then all cycles must wind around the cylinder exactly once and at the same ordinate, i. e., horizontally. Suppose that G has two non-disjoint cycles C_1, C_2. Then, C_1 and C_2 share a common vertex v. Let y_1 and y_2 be the ordinates of C_1 and C_2, respectively. As v is part of both, $y_1 = y_2$.

Thus, C_1 and C_2 interfere and the drawing is not planar; a contradiction. Note that vertex disjointness implies edge disjointness. $\square$

For the characterization of **wSUP** graphs, we use supergraphs which may have an extra source or sink and extend techniques for **SUP** graphs from [15].

Lemma 10. *A graph G is a **wSUP** graph if and only if it has a **wSUP** supergraph $H \supseteq G$ with one source and one sink.*

Proof. "$\Rightarrow$": Let G be a **wSUP** graph. Consider the component graph $\mathbb{G}$ of G. According to Lemma 9, cycles in G are disjoint. Each such cycle becomes a compound in $\mathbb{G}$ and there are no other compounds. The compounds subdivide $\mathbb{G}$ into *sections*, i. e., the union of transits that have their source in the lower compound or sink in the upper compound or both. As G is planar, no edge can span multiple sections. A strongly connected component contains neither sources nor sinks, hence, it suffices to focus on sections.

Let $\mathbb{G}_\sigma$ be an intermediate section, i. e., a section bounded by a lower compound s and an upper compound t. As $\mathbb{G}_\sigma \in$ **SUP**, it has a planar, acyclic dipole $\mathbb{H}_\sigma$ as supergraph. The construction of $\mathbb{H}_\sigma$ given in [15] allows for $\mathbb{H}_\sigma$ to have a **wSUP** embedding which follows the embedding of $\mathbb{G}_\sigma$ and s is its single source and t its single sink. If $\mathbb{G}_\sigma$ is extremal, i. e., it is the lowermost or the uppermost section (or both) and not bounded by a compound on the lower or upper end, then a source or sink, respectively, is chosen according to the construction given in [15]. Thus, we obtain a planar, acyclic dipole for every section of $\mathbb{G}$. Now expand the compounds to horizontal cycles again by adjusting the embedding accordingly. If, during the construction, incoming or outgoing edges were added to a compound, add them to the vertices of the compound such that the cyclic ordering in a re-contraction of the vertices remains the same. This does not introduce any new cycles. In the end, every intermediate section is a **wSUP** graph without sources and sinks and the two extremal sections are **wSUP** graphs with exactly one source or sink, respectively. If there is no uppermost (lowermost) extremal section, simply add a vertex and an edge from (to) one of the vertices of the uppermost (lowermost) compound.

Let H be the graph constructed as described. Then, H is a **wSUP** graph and a supergraph of G with exactly one source and sink. Note that due to the elimination of sources and sinks within the sections, the component graph of H only consists of strongly connected components and transits and, therefore, H is a dipole.

"$\Leftarrow$": Follows immediately, since any subgraph of a **wSUP** graph is in **wSUP**. $\square$

We are now able to give a first characterization of **wSUP** graphs.

Theorem 2. *A graph G is a **wSUP** graph if and only if it has a supergraph $H \supseteq G$ such that H is a planar dipole whose cycles are (vertex) disjoint.*

Proof. "$\Rightarrow$": Let G be a **wSUP** graph. Then, the supergraph H constructed according to Lemma 10 has exactly one source s and one sink t. Additionally, H is in **wSUP** and a dipole. By Lemma 9 all cycles of H are disjoint.

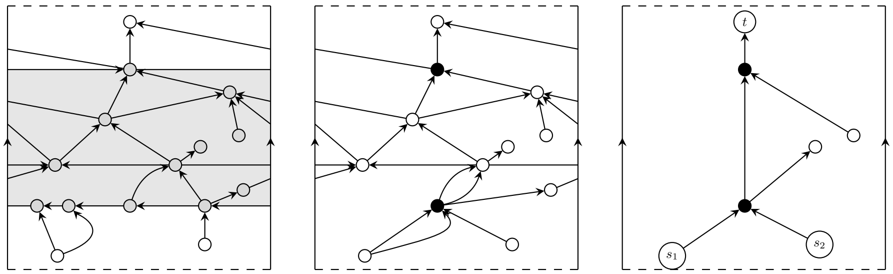

(a) Graph $G \in$ **wSUP**

(b) The component graph $\mathbb{G}$ of G

(c) The compound graph $\overline{\mathbb{G}}$ of G

Fig. 7. A **wSUP** example

"$\Leftarrow$": Let H be a planar dipole whose cycles are disjoint. Consider the component graph $\mathbb{H}$ of H. As H is a planar dipole with source s and sink t, the compound graph of H is a path $s \rightsquigarrow t$. Each transit τ is a planar, acyclic dipole and, therefore, has a **SUP** embedding. We can construct a drawing of H on the standing cylinder by assembling the compounds and transits of H on on top of each other (with respect to the upward direction) according to their appearance in a traversal of the compound graph from s to t. Every transit is drawn according to its **SUP** embedding. Every compound of H consists of a single cycle only, so it can be drawn horizontally, i. e., it winds around the cylinder exactly once and at the same ordinate. This leads to a weakly upward planar drawing, so H is a **wSUP** graph. Now G is in **wSUP** as it is a subgraph of the **wSUP** graph H. □

Next, we turn to the duals of **wSUP** graphs. Recall from Lemma 4 in combination with Proposition 1 that a graph with one source and one sink is in **SUP** if and only if its dual is a strongly connected **RUP** graph. Introducing vertex disjoint cycles, the characterization via dual graphs now reads as follows.

Theorem 3. *A graph G with exactly one source and sink is a **wSUP** graph if and only if its dual G^* is a **RUP** graph that has no trivial strongly connected components.*

Proof. "$\Rightarrow$": Let G be a **wSUP** graph with exactly one source and one sink. Then, G is a dipole and Theorem 1 together with Proposition 1 implies that G^* is in **RUP**.

If G is acyclic, then G^* is strongly connected by Lemma 4 and Proposition 1. Hence, G^* has no trivial strongly connected components. Otherwise, G may contain only simple, edge disjoint ("horizontal") cycles, which wind around the cylinder at the same ordinate. Therefore, a cycle cannot "split", i. e., in a traversal of the cycle, the successor vertex is unambiguous. These cycles are the compounds in the component graph $\mathbb{G}$.

Consider the dual graph G^* of G. By Lemma 5, the dual of each transit of G is in **RUP** and is strongly connected. According to the definition, a transit contains at least one edge. Therefore, the dual of a transit is a compound. The dual of every cycle C consists of simple edges from the faces to the left of C to the faces to its right, which are also part of the transits' duals. Since the cycles are disjoint, no faces are enclosed by their edges. Thus, all vertices of G^* are part of a compound and G^* has no trivial strongly connected components.

"$\Leftarrow$": Let G^* be a **RUP** embedded graph without any trivial strongly connected components. If G^* is strongly connected, then G is an acyclic dipole by Lemma 4 and Proposition 1 and, thus, a **SUP** graph [15].

Otherwise, consider the component graph of G^*. According to Lemma 4 and Proposition 1, the primal of each strongly connected component is an acyclic dipole. Since G^* has no trivial strongly connected components, all transits consist of paths of length 1. Hence, the primal of each transit is a simple, disjoint cycle which winds around the cylinder exactly once and G does not contain any other cycles. G^* has neither sources nor sinks, since they would be trivial strongly connected components. Thus, by Lemma 3 and Proposition 1, G is a planar dipole with disjoint cycles only. Theorem 2 now implies that G is a **wSUP** graph with only one source and one sink. □

From Theorem 3 and Lemma 10 we directly obtain the following corollary, which concludes our second characterization of **wSUP** graphs.

Corollary 2. *Every **wSUP** graph G has a **wSUP** supergraph H whose dual H^* is a **RUP** graph without trivial strongly connected components.*

5 Summary

We have shown that a directed graph has a planar upward drawing on the rolling cylinder if and only if it is a spanning subgraph of a planar graph without sources and sinks whose dual is a dipole. This result completes the known characterizations of planar upward drawings in the plane [7,17] and on the sphere [9,11,14,15]. Every **SUP** graph is a spanning subgraph of a planar, acyclic dipole and every **UP** graph is a spanning subgraph of a planar, acyclic dipole with an *st*-edge. Moreover, a graph has a weakly upward drawing on the standing cylinder if and only if it is a subgraph of a planar dipole with disjoint cycles.

Concerning dual graphs, the duals of the acyclic components of **RUP** graphs are in **RUP** and the duals of the strongly connected components are in **SUP**. In particular, the dual of a strongly connected **RUP** graph is in **SUP**. Every **wSUP** graph has a planar supergraph whose dual is a **RUP** graph without trivial strongly connected components.

Future work is to investigate whether the characterization by means of dual graphs leads to new insights on the upward embeddability on other surfaces, e. g., the torus. Also, the duals of quasi-upward planar graphs [5] shall be considered.

References

1. Auer, C., Bachmaier, C., Brandenburg, F.J., Gleißner, A.: Classification of planar upward embedding. In: van Kreveld, M., Speckmann, B. (eds.) GD 2011. LNCS, vol. 7034, pp. 415–426. Springer (2012)
2. Bachmaier, C.: A radial adaption of the Sugiyama framework for visualizing hierarchical information. IEEE Trans. Vis. Comput. Graphics 13(3), 583–594 (2007)
3. Bachmaier, C., Brandenburg, F.J., Brunner, W., Fülöp, R.: Drawing recurrent hierarchies. J. Graph Alg. App. 16(2), 151–198 (2012)
4. Bang-Jensen, J., Gutin, G.: Digraphs: Theory, Algorithms and Applications. Springer, 1st edn. (2000)
5. Bertolazzi, P., Di Battista, G., Didimo, W.: Quasi-upward planarity. Algorithmica 32(3), 474–506 (2002)
6. Brandenburg, F.J.: On the curve complexity of upward planar drawings. In: Mestre, J. (ed.) Computing: The Australasian Theory Symposium, CATS 2012. Conferences in Research and Practice in Information Technology (CRPIT), vol. 128, pp. 27–36. Australian Computer Society (ACS) (2012)
7. Di Battista, G., Tamassia, R.: Algorithms for plane representations of acyclic digraphs. Theor. Comput. Sci. 61(2–3), 175–198 (1988)
8. Dolati, A.: Digraph embedding on t_h. In: Cologne-Twente Workshop on Graphs and Combinatorial Optimization, CTW 2008. pp. 11–14 (2008)
9. Dolati, A., Hashemi, S.M.: On the sphericity testing of single source digraphs. Discrete Math. 308(11), 2175–2181 (2008)
10. Dolati, A., Hashemi, S.M., Kosravani, M.: On the upward embedding on the torus. Rocky Mt. J. Math. 38(1), 107–121 (2008)
11. Foldes, S., Rival, I., Urrutia, J.: Light sources, obstructions and spherical orders. Discrete Math. 102(1), 13–23 (1992)
12. Garg, A., Tamassia, R.: On the computational complexity of upward and rectilinear planarity testing. SIAM J. Comput. 31(2), 601–625 (2001)
13. Hansen, K.A.: Constant width planar computation characterizes ACC^0. Theor. Comput. Sci. 39(1), 79–92 (2006)
14. Hashemi, S.M., Rival, I., Kisielewicz, A.: The complexity of upward drawings on spheres. Order 14, 327–363 (1998)
15. Hashemi, S.M.: Digraph embedding. Discrete Math. 233(1–3), 321–328 (2001)
16. Kaufmann, M., Wagner, D.: Drawing Graphs, LNCS, vol. 2025. Springer (2001)
17. Kelly, D.: Fundamentals of planar ordered sets. Discrete Math. 63, 197–216 (1987)
18. Limaye, N., Mahajan, M., Sarma, J.M.N.: Evaluating monotone circuits on cylinders, planes and tori. In: Durand, B., Thomas, W. (eds.) STACS 2006. LNCS, vol. 3884, pp. 660–671. Springer (2006)
19. Limaye, N., Mahajan, M., Sarma, J.M.N.: Upper bounds for monotone planar circuit value and variants. Comput. Complex. 18(3), 377–412 (2009)
20. Sugiyama, K., Tagawa, S., Toda, M.: Methods for visual understanding of hierarchical system structures. IEEE Trans. Syst., Man, Cybern. 11(2), 109–125 (1981)
21. Thomassen, C.: Planar acyclic oriented graphs. Order 5(1), 349–361 (1989)