On the Curve Complexity of Upward Planar Drawings *

Franz J. Brandenburg

University of Passau,
94030 Passau, Germany
Email: brandenb@informatik.uni-passau.de

Abstract

We consider directed graphs with an upward planar drawing on the plane, the sphere, the standing and the rolling cylinders. In general, the drawings allow complex curves for the edges with many zig-zags and windings around the cylinder and the sphere.

The drawings are simplified to polyline drawings with geodesics as straight segments and vertices and bends at grid points. On the standing cylinder the drawings have at most two bends per edge and no windings of edges around the cylinder. On the rolling cylinder edges may have one winding and five bends, and there are graphs where edges must wind. The drawings have a discrete description of linear size. The simplifications can be computed efficiently in $\mathcal{O}(\tau\, n^3)$ time, where τ is the cost of computing the point of intersection of a curve and a horizontal line through a vertex. The time complexity does not depend on the description complexity of the drawing and its curves, but only on $\mathcal{O}(n^3)$ sample points.

1 Introduction

Graph drawing is mostly concerned with the problem to map a graph in the plane. The objective are nice drawings which shall be constructed by efficient algorithms. Such a map assigns the vertices of a graph to distinct points. The edges are simple Jordan curves between the endpoints. Hence, graph drawing realizes a transformation from a topological into a geometrical structure.

The task of drawing a graph consists of two phases: placement and routing. Often both phases are merged, particularly, for straight-line drawings. These drawings are completely described by the placement of the vertices. Moreover, there are polyline drawings with straight-line segments, which are used e.g. in the hierarchical approach for directed graphs or in orthogonal graph drawings, see [12, 22]. These drawings have a discrete description, which is given by the points for the vertices and the bends. More complex curves are used in a postprocessing phase, where bends of polylines are smoothed by splines [15] or spiral segments [2].

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

The full generality of Jordan curves is used in the definition of planar graphs and plane drawings. A drawing$\Gamma(\ G)$ of a graph G is plane if it is one-to-one on the vertices and the curves of distinct edges are disjoint except at common endpoints. What does this really mean? How complex are the curves? This question has rarely been addressed, since mostly the curves are simplified. For effective computations the curves must have a finite description, which is given by a complete listing, a program or a Turing machine. However, the full description may be very long and may exceed the size of the graph by any (super-exponential) bound. Such a complexity would be out of scope for graph drawing.

For planar (undirected) graphs it is well-known that all curves can be straightened. This was proved by Steinitz and Rademacher [30], Wagner [32] and Stein [29] using the one-to-one correspondence between (triconnected) planar graphs and the mesh of convex polyhedra. An alternative proof by Fáry [18] uses induction and is based on the fact that the outer face of a triangulated planar graph can be drawn as a straight-line triangle. This proof can be implemented by an $\mathcal{O}(n^3)$ algorithm, since it uses separating triangles. These investigations were finalized by de Fraysseix, Pach and Pollack [11] and by Schnyder [28] who showed that every planar graph has a straight-line grid drawing of $\mathcal{O}(n^2)$ area, which can be computed in $\mathcal{O}(n)$ time.

Directed graphs are most commonly drawn as hierarchies using the approach introduced by Sugiyama et al. [31]. This drawing style produces polyline drawings and transforms the edge direction into a geometric direction: all edges point upward if the graph is acyclic. If graphs have cycles then there is a unidirectional representation on the rolling cylinder. Roll the cylinder and so follow the direction of the edges. This drawing style is a particularity of the rolling cylinder.

The combination of upward and planar leads to upward planarity. A graph is upward planar or an *up*-graph if it can be drawn in the plane such that the edges are monotonically increasing in y-direction and there are no edge crossings. *up*-graphs were studied intensively; for a comprehensive study see [12]. They were characterized as the spanning subgraphs of planar *st*-graphs [13, 23], which are directed acyclic graphs with a single source s and a single sink t and an (s, t) edge.

Concerning the curve complexity *up*-graphs are simple, too, since they admit straight-line upward planar drawings. This can be proved by Fáry's [18] inductive technique. On the other hand, straight-line upward planar drawings may require an area of exponential size [14], whereas there are upward polyline drawings on quadratic area with at most two bends per edge [13].

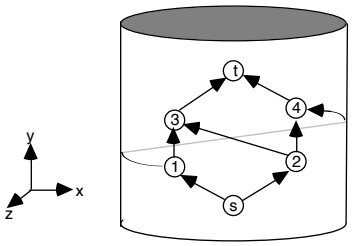

Figure 1: *st*-$K_{2,2}$ on the standing cylinder

Drawing graphs has also been considered on other surfaces, such as the sphere [21] and the torus [27], and there are 3D approaches, see [22]. In particular, upward planarity has been generalized to the sphere and the standing cylinder using a solid in $\mathcal{R}^3$ and a fixed embedding of the respective surface which is used as the canvas for the drawing. The respective classes of upward planar graphs were studied separately by Hashemi et al. [16, 20, 21] and by Hansen [19] and Limaye et al. [24, 25].

For upward planarity the surfaces of genus zero no longer coincide: the plane is weaker than the sphere or the standing cylinder. A counterexample is the *st*-$K_{2,2}$ graph consisting of the complete bipartite $K_{2,2}$ with a source and a sink, see Figure 1 and [13, 23]. The (s,t)-edge makes the distinction in the characterization. The upward planar graphs on the sphere [20] and on the standing cylinder [24] were characterized as spanning subgraphs of acyclic planar graphs with a single source and a single sink. This implies that the respective classes of graphs coincide.

Recently, Auer et al. [1] introduced a universal approach towards upward planar drawings defining the upward direction by a vector field on an arbitrary surface. This approach gives raise to a general classification of upward planarity including spherical, cylindrical, rotational and toroidal. These come from the sphere, the standing and rolling cylinders and the torus with the homogeneous field on the respective surfaces. Here, graphs with a planar upward drawing on a standing and a rolling cylinder were characterized in terms of upward drawability in the plane with a radial and a cyclic field for the upward direction, and the coincidence of upward planarity on the sphere and the standing cylinder was formally established.

Upward planar drawings are related to level planar graphs, where each vertex is assigned to a fixed level and the edge direction is given by the numbering of the levels. Level planarity is an important topic in graph drawing [22]. In the plane the levels are horizontal lines. On the standing cylinder and the sphere the levels are latitudes. This is equivalent to concentric circles in the plane [2, 5], which is the view on a standing cylinder or a sphere from below. On the rolling cylinder the levels are horizontal lines, which form a recurrent hierarchy [31], where there are k levels which are numbered modulo k. Such drawings were recently investigated in depth in [3, 4, 6, 7].

In this work we focus on the simplification of upward planar drawings on the standing and rolling cylinders. We call them *sup*- and *rup*-drawings of *sup*- and *rup*-graphs. Our goal are polyline drawings with geodesics as straight segments and vertices and bends at grid points. Then the drawings are completely specified by the positions of the vertices and the edge bends.

Initially, the edges are arbitrary Jordan curves, which are monotonically increasing and disjoint except at common endpoints. The curves may be highly complex and zig-zag and frequently wind around the cylinders. There may even be infinitely many zigzags, e.g. driven by the function $f(t) = t \sin \frac{1}{t}$ with $t \to 0$ if we accept this as a description. Infinitely many windings with an ultimate convergence towards an endpoint are excluded by the continuity of a Jordan curve. However, the number of windings can be arbitrarily large. For example, a curve may form a spiral with successive windings around the cylinder at distance 2^{-i} for $i = 1, \ldots, N$, and N is e.g. the value of Ackermann's function $A(n,n)$, where n is the size of the graph.

We jump over this complexity and show that every *sup*- and *rup*-graph has a polyline drawing with geodesics as straight segments. The edges have only a few bends per edge, and the vertices and the bends are placed on a grid of quadratic size. Edges do not wind around the standing cylinder and wind at most once around the rolling cylinder. Here single windings of some edges are unavoidable. However, it is unknown whether or not there are straight *sup*- and *rup*-drawings for the respective graphs.

To cope with the complexity of the description of Jordan curves we suppose that it takes time τ to compute the coordinates of the intersection of an edge and a horizontal line through a vertex. τ may be huge. In the previous example with windings described by Ackermann's function one must evaluate the Ackermann function. However, τ is a constant in the resulting drawings. On the standing cylinder and the sphere the polyline drawing can be computed in $\mathcal{O}(\tau\, n^2)$ time, and it takes $\mathcal{O}(\tau\, n^3)$ time on the rolling cylinder. This is a significant reduction of the description complexity of upward planar drawings. It may be quite easy to detect the intersections of curves with the horizontal lines through the vertices; however it may be very complex to trace the edges, particularly, if they zig-zag and frequently wind around the cylinder. Thus we close the gap between drawings with arbitrary Jordan curves and polyline drawings, which are the standard in graph drawing, and we capture Jordan curves computationally using a value τ for the computation of an intersection of curves in a neighborhood of vertices.

In the next section we introduce the basic notions. Then we study upward planar drawings first on the standing and thereafter on the rolling cylinder, and we conclude with some open problems.

2 Upward Drawings

We assume familiarity with the fundamental concepts on graphs and graph drawing, see [12, 22].

A *graph* $G = (V, E)$ is a simple directed graph with a finite set of vertices V and a finite set of directed edges E. An edge $e = (u, v)$ is directed from u to v. G has n vertices and m edges. Moreover, we assume that the underlying undirected graph is connected; disconnected components can be drawn on top of each other on the standing cylinder and side-by-side on the rolling cylinder.

We consider well-known two-dimensional surfaces which are given by a fixed embedding in $\mathcal{R}^3$, where x is horizontal, y is vertical, and z is directed towards the spectator according to the right-hand rule. The *plane* is the surface with the Cartesian coordinates and $z = 0$. Since scaling is allowed the area of the plane may be restricted to a range from -1 to +1 in each dimension. The *sphere* is defined by

$\{(x,y,z) \mid x^2+y^2+z^2=1\}$, the *standing cylinder* by $\{(x,y,z) \mid x^2+z^2=1 \text{ and } -1 \leq y \leq 1\}$ and the *rolling cylinder* by $\{(x,y,z) \mid y^2+z^2=1 \text{ and } -1 \leq x \leq 1\}$.

A planar drawing Γ maps each vertex of G to a distinct point on a surface $\mathcal{S}$ and each directed edge $e=(u,v)$ is mapped to a simple Jordan curve $J:[0,1] \to \mathcal{S}$ from the endpoint$\Gamma(u)$ to the endpoint$\Gamma(v)$ such that the curve is continuous and no two edges have a common point except a common endpoint.

When it is clear from the context, we say that a vertex v is *placed* at$\Gamma(v)$ and we do not distinguish between an edge e and its curve$\Gamma(e)$. Additionally, $\Gamma(G)$ stands for the set of points of the drawing.

A drawing is *upward planar* if the curves of the edges are monotonically increasing in y-direction and the edges are disjoint except for common endpoints. In the plane, on the sphere and on the standing cylinder the y-direction is the y-dimension, i.e., $a<b$ implies $y(J(a)) < y(J(b))$, where $y(J(a))$ is the value of the y-coordinate at $J(a)$. On the rolling cylinder the y-direction is taken in counterclockwise order in the (z,y) plane. This means up for $z>0$ and down for $z<0$. For the spectator it is upward on the front side, and then over the upper side and down on the backside. All edges are unidirectional if the cylinder is rolled clockwise. Upward is formally described by the angle of the tangent of the curve at a point in (z,y) dimension and the horizontal line through the point. It is readily seen that upward is equivalent to increasing in y-dimension in the two-cell approach where the surface is unrolled, or in the fundamental polygon approach where opposite sides of a square are identified [1, 26, 27].

An *upward planar graph* G on a surface $\mathcal{S}$ admits a upward planar drawing on $\mathcal{S}$. Let $\boldsymbol{UP}, \boldsymbol{SUP}$ and $\boldsymbol{RUP}$ be the classes of graphs with an upward planar drawing on the plane, the standing and the rolling cylinders. We call these graphs *up-*, *sup-* and *rup-*graphs, respectively, see [1, 9].

It is well-known that the plane, the sphere, the standing and the rolling cylinders coincide on planarity - the undirected view. The picture changes for upward planarity - the directed view. Here, the plane is weaker than the sphere and the standing cylinder, since the *st*-$K_{2,2}$ graph from Figure 1 is non-upward planar [13, 23]. Moreover, the sphere and the standing cylinder are equivalent and are weaker than the rolling cylinder as first established by Auer et al. [1] and deepened in [9]. They are dominated by surfaces of higher genus, such as the torus. These facts are summarized to:

Proposition 1.

$$\boldsymbol{UP} \subsetneq \boldsymbol{SUP} \subsetneq \boldsymbol{RUP}.$$

3 The Standing Cylinder and the Sphere

Our topic is the curve complexity of upward planar drawings. We transform drawings with arbitrary Jordan curves for the edges into polyline drawings with few bends and windings around the cylinder. Hence, curves with a high description complexity and even infinitely many zig-zags and many windings are simplified significantly. An edge zig-zags if it has many bends and it winds around the standing cylinder or the sphere if it intersects a vertical line or a longitude at least twice. The resulting upward planar drawing is completely specified by $\mathcal{O}(n)$ points on the standing cylinder or the sphere. Hence, the data for the description of the drawing is reduced from (even super-exponential) high volumes to a linear size.

A drawing$\Gamma(G)$ on the standing cylinder is seen as a level drawing (the sphere can be treated in a similar way). Then upward planarity turns into radial level planarity as studied by Bachmaier et al. [2, 5]. Each level defines a circular list of *observable points* $OBS(v)$, which are the intersection points of the edges and the horizontal level $L(v)$ of a vertex v. $OBS(v)$ together with the vertices on the level are sorted counterclockwise from v and has at most $n+m$ entries.

We suppose that each observable point can be computed in time τ, whereas vertices can be accessed directly in$\Gamma(G)$ in $\mathcal{O}(1)$ time. There may be a huge gap between the complete description of a drawing $\Gamma(G)$ and the data stored by the sets of observable points, which consists of only a few sampling points per edge. The gap between the original drawing and this data is bridged by τ. In the resulting polyline drawings τ is in $\mathcal{O}(1)$, since each edge consists of finitely many straight segments.

The simplification of upward planar drawings on the standing cylinder to polyline drawings was first stated in [1]. Here, we improve these results and also consider windings, bends, a grid, and the time complexity. These parameters were not addressed before together with planar upward drawings.

Lemma 1. *Let $\Gamma(G)$ be an upward planar drawing of a graph G on the standing cylinder (or the sphere). Then there is a planar upward polyline drawing $\tilde{\Gamma}(G)$ with at most n-2 bends per edge and no windings of edges around the cylinder.*

$\tilde{\Gamma}(G)$ can be computed in time $\mathcal{O}(\tau\, n^2)$, where τ is the time to compute the (x,z)-coordinates of the intersection of an edge and the latitude of a vertex.

Proof. The *sup*-drawing$\Gamma(G)$ is cut horizontally at the y-coordinates of the vertices. Each cut defines a circle of points on the perimeter of the cylinder. The points are the vertices on the cut and the intersection points for each proper crossing of an edge and the cut.

A *disk* is a piece of$\Gamma(G)$ between two adjacent cuts. On its borders it has a lower and an upper circle of points. There are no vertices in its interior and segments of edges are routed planar upward between the points on the lower and upper circles.

We process the disks iteratively from bottom to top. For a disk choose an edge segment connecting two points on the opposite borders for an alignment, say p_1 on the lower and q_1 on the upper border. Rotate the upper ring and all rings above such that p_1 and q_1 have the same (x,z)-coordinates. To save computation time these rotations are performed in a lazy fashion and are first stored as offsets. In a final bottom to top sweep the rotations are summed up over all disks. In a rotation the edge segments are stretched or contracted like rubber bands.

From (p_1,q_1) process the edge segments between the opposite borders in counterclockwise order and replace each edge segment from a point p on the lower circle to a point q on the upper circle by the geodesic from p to q. Here the counterclockwise order of the geodesics must be preserved, which is obtained from the embedding induced by the drawing. The geodesic from p to q is not necessarily the shortest curve between these points, since (p_1,q_1) must not be intersected. Hence, the cyclic order of all (segments of the) curves is preserved, which implies that there are no crossings. So the curve of an edge is a polyline of geodesics with at most n-2 bends. Let$\Gamma'(G)$ be the so obtained drawing of G. $\Gamma'(G)$ has a discrete description with $\mathcal{O}(n^2)$ points for the vertices and the bends.

Let's consider windings. An edge cannot wind around a disk, since the first segment is vertical. Windings of edges can be avoided by an appropriate choice of the points for the alignment. Suppose that these points are placed on the front line of the cylinder with $(x, z) = (0, 1)$. This line will not be crossed by any edge.

We proceed from bottom to top in$\Gamma'(G)$. Suppose that every edge segment on a disk has unit length. Then we compute longest paths from appropriate starting points $t_1, \ldots, t_r$ for some $r \geq 1$. The paths are ordered monotonically and each path ends at a sink of G. They are routed on the front line. The first path starts at a vertex on the lower border of the bottommost disk, and the last path ends at a vertex on the upper border of the topmost disk. Consider the i-th path p_i from t_i to t'_i. If there are vertices on the front line or edge segments crossing the front line above t'_i, then get the first such item. Suppose that an edge segment $s = (p_s, q_s)$ crosses the front line next above t'_i, and let $s = (q_s, q_s)$ if it were a vertex. Then choose q_s for the alignment, place it on the front line, and use it as the starting point of the i+1-st path. Since q_s is taken as the first point above the i-th path there is no other edge passing the gap between t'_i and q_s on the front line. By the planarity edges cannot cross paths. Hence, no edge crosses the front line, which excludes windings. Let $\tilde{\Gamma}(G)$ be the so obtained drawing. $\tilde{\Gamma}(G)$ is a polyline drawing without edge windings.

Concerning the time complexity there are at most n-1 disks, and each circle has at most $\mathcal{O}(n)$ points, since each upward curve intersects at most once. The coordinates of each such point are computed in time τ. Each of the at most $\mathcal{O}(n^2)$ many segments of geodesics can then be computed in $\mathcal{O}(1)$ time. Using lazy evaluation all rotations are performed in $\mathcal{O}(n^2)$ time. Hence, the computation of$\Gamma'(G)$ takes $\mathcal{O}(\tau\, n^2)$ time.

The computation of $\tilde{\Gamma}(G)$ from$\Gamma'(G)$ can be done in linear time in the number of edge segments. The coordinates of their endpoints were computed before such that it takes $\mathcal{O}(1)$ per segment, and each segment must be considered at most once for the longest paths computations. □

In a post-processing phase we reduce the number of bends from $\tilde{\Gamma}(G)$. We consider two approaches. The first goes via the visibility representation of *up*-graphs in the plane by Di Battista and Tamassia [13], where the vertices are represented by horizontal segments and the edges by straight vertical lines. This representation can be transformed into a monotone grid drawing with at most two bends per edge. The used coordinates are integral and vertices and bends are placed on a grid of $\mathcal{O}(n^2)$ size.

However, on the standing cylinder edges can be routed over the backside. These edges have an extra bend at a vertical cut line along the backside. Such edges may have up to five bends. We do not pursue this approach here because it is outperformed by the *linear segments model*, where each edge is represented by a polyline with at most three straight segments. The approach by Brandes and Köpf [10] realizes this model and operates even on arbitrary directed graphs in the plane. It uses a thinning technique and extracts edges or segments of edges until all (dummy) vertices have degree at most two. Then only a collection of paths remains, which are straightened vertically and are ordered left to right and compacted by a longest path heuristic. In addition, there is a balancing phase in order to center the vertices over their neighbors. This is skipped here. The paths determine the vertical grid coordinates of the vertices and the horizontal grid coordinate is given by the leveling.

The approach was extended to radial drawings with concentric circles for the levels by Bachmaier [2]. The running time of the algorithms is linear in the number of proper segments between adjacent levels. This number may be quadratic in the size of the graph. An improvement in the running time was obtained by Eiglsperger et al. [17] using edge bundling. We specialize and adapt the approach to upward planar graphs.

Lemma 2. *Let $\tilde{\Gamma}(G)$ be a polyline sup-drawing without edge windings and bends only at the latitudes of vertices. Then there is a polyline drawing $\hat{\Gamma}(G)$ on the standing cylinder with no windings and at most two bends per edge, and all vertices and all bends are placed at grid points of a grid of size $\mathcal{O}(n^2)$.*

$\hat{\Gamma}(G)$ can be computed from $\tilde{\Gamma}(G)$ in $\mathcal{O}(n)$ time.

Proof. Consider $\tilde{\Gamma}(G)$ as constructed in the proof of Lemma 1 with disks and horizontal latitudes $L(v)$ through the vertices. Then the front line can be transformed into a vertical *cut line* C. If edges are completely routed on C they are kept. If an edge e enters C from the left (right) and meets C at latitude $L(u)$ it remains on C up to the latitude of its end node $L(v)$. Then move the segments between $L(u)$ and $L(v) - 1$ slightly to the left (right), such that no other edge or vertex is touched or crossed. This preserves the property that no edge crosses C and prevents edge windings.

Project a grid on the standing cylinder, such that the horizontal grid lines are the latitudes of the vertices. There are at most m vertical grid lines, which are set by the subsequent procedure. The grid points are regarded as integer coordinates. Take C as the first vertical grid line and fix the vertices and edges on C as they are given after the slight move.

First, remove all edge segments entering C from either side. In particular, this removes the outer segments from $L(v) - 1$ to $L(v)$ that were slightly moved previously. Thereafter, we pursue the thinning technique from [10]. An edge e of $\tilde{\Gamma}(G)$ is *short*, if it spans just one disk and is *long*, otherwise. A long edge consists of two outer segments on its top- and bottommost disks and of a *stick* over the intermediate disks. A stick may degenerate to a single point.

The approach aligns the sticks as straight vertical lines on the grid and attempts to obtain even more straight edges and save bends. If a vertex u has outgoing outer segments of long and short edges, then temporarily remove the short edges and finally keep only the left median of the outer segments of the long edges. If u has only outgoing short edges the left median edge is kept. Accordingly, keep the left median incoming outer segment at v or the left median incoming short edge, if there are no outer segments entering v. All other outer segments and short edges are temporarily removed. Thereafter all endpoints of edges, outer segments or sticks have degree at most two. Together with the kept edge segments they form a set of paths P. Moreover, for every vertex v of G there is exactly one path which contains v. In addition P contains sticks.

First, align C vertically as the first (leftmost) vertical grid line. Then order the paths from left to right starting at C. A path p is at least one grid unit to the left of a path p' if at some latitude p is to the left of p' in a counterclockwise (or left to right) traversal

from C. This defines a partial order on P. The paths on C are the minimal and the paths immediately to the left of C are the maximal elements in this partial order. Sort P by a longest path heuristic (or by topological sorting). Assign the number of this sorting as the horizontal grid coordinate to each path p and each vertex v on p. So the paths and vertices on C are on the first vertical grid line. In consequence, all sticks are vertically aligned on the grid. Hence, long edges have a bend at most between an outer segment and the stick. The vertical grid coordinate for the vertices and the endpoints of the sticks is taken from the latitudes. The previously removed segments are reinserted and are drawn as straight lines on the grid, which is folded on the standing cylinder. The drawing is planar since the partial order of the paths preserves the order of incoming and outgoing edges (edge segments) at each vertex. Our simplification of the algorithm of Brandes and Köpf [10] needs only linear time, since each edge is partitioned into at most three pieces, and the removal of edge segments, the computation of the medians and the subsequent construction of the path can be done in linear time in the number of pieces. The longest path heuristic is a modified topological sorting, which computes the distance of a path from C in linear time. Topological sorting produces wider drawings also in linear time. The underlying grid has at most n horizontal lines and at most m vertical lines. Hence, the grid size is at most quadratic in the size of the given graph. □

For the standing cylinder and the sphere we can summarize:

Theorem 1. *For every upward planar drawing $\Gamma(G)$ of a graph G on the standing cylinder or the sphere there is an upward planar drawing $\hat{\Gamma}(G)$ of G*

- *with polylines and geodesics as straight segments*
- *no edge windings around the cylinder (or sphere)*
- *at most two bends per edge*
- *a placement of all vertices and bends on a grid of quadratic size, and*
- *which can be computed in time $\mathcal{O}(\tau\, n^2)$, where τ is the time to compute an intersection of an edge and a horizontal line.*

4 The Rolling Cylinder

For the rolling cylinder we apply a similar technique and transform an upward planar drawing into a recurrent hierarchy with horizontal levels at the vertices. Recurrent hierarchies were introduced by Sugiyama et al. [31] in their pioneering work on drawing hierarchies and were recently studied in depth by Bachmaier et al. [3, 6]. A k-level recurrent hierarchy is a directed graph together with an assignment of the vertices to levels $L_0, \ldots, L_{k-1}$, which are ordered modulo k. An edge (u, v) from vertex u on level L_i to vertex v on level L_j is directed upward in the cyclic order of the levels. In particular, there are edges from vertices on high levels to vertices on low levels, which first meet or cross level L_0. Recurrent hierarchies are drawn in 3D on the rolling cylinder or in 2D, where either the levels are rays from a common center and the edges are routed as poly-spiral curves or L_0 is duplicated and its copy appears at the top of the common hierarchical drawing in the plane with horizontal levels. These drawing styles are equivalent in the sense that there are geometric transformations from one drawing into the other, as elaborated in [3].

The complexity of Jordan curves in upward planar drawings on the rolling cylinder comes from zig-zags and vertical windings around the cylinder. Zig-zags are dealt with as above and are replaced by geodesics between points on adjacent levels. To simplify the drawings we proceed in three steps. First, we smooth Jordan curves to polylines with at most $\mathcal{O}(n)$ many windings and bends. Then redundant windings are removed and finally the number of bends is reduced to at most five per edge and the vertices are placed on a grid of size $\mathcal{O}(n^2)$.

In addition we reduce the computational effort and achieve time bounds that are independent of the description complexity of the drawing and can be given in terms of the size of the graph. For our constructions a viewer of the drawing must not (be able to) trace the Jordan curves and follow the zig-zags and too many windings. It suffices to compute the sets of *observable points* of the vertices $OBS(v)$. Again an observable point p is the intersection point of an edge e and a horizontal line $L(v)$ through a vertex v. Its computation takes τ units of time. However, the sets $OBS(v)$ are different from those on the standing cylinder. Each vertex v first considers the intersection points of some edges immediately to its left and right. So it determines its *passing* edges and whether these edges wind frequently and purely. Each passing edge is observed for at most two rounds; multiple observations of edges from several vertices don't matter. A more frequently winding edge is observed piecewise by several vertices or it disappears from the screen for a while, since there are pure windings and the edge can be cut short. The short cut is done just below the level of the *winding successor*, which is determined by observing the bundles of edges with pure windings.

Let's make these ideas precise.

Let$\Gamma(\ G)$ be an upward planar drawing on a rolling cylinder. The *level* $L(v)$ of a vertex v is given by the horizontal line through its (y, z)-coordinates. These levels $L_0, \ldots, L_{k-1}$ are ordered counterclockwise in the (y, z)-plane such that each level has a successor in up direction and L_0 is the successor of L_{k-1}.

An edge $e = (u, v)$ of$\Gamma(\ G)$ is *winding* if its curve intersects $L(u)$ at some point p. It is *right-winding* (*left-winding*), if p is to the right (left) of u. A right-winding (left-winding) edge also intersects the level of v to the left (right) of v.

Consider edges in a *rup*-drawing. Clearly, an edge is self-intersecting if it is left- and right-winding. This follows from the mean-value theorem using the representation on the fundamental polygon, where the upper and lower sides are identified. By the same reasoning, left-winding and right-winding edges cannot interleave. A bundle of winding edges must wind in the same direction. If e is left-winding and e' is right-winding, and e' is to the left of e at some level, then e' is completely to the left of e except at a common endpoint, otherwise they would cross. However, a vertex may have outgoing left- and right-winding edges, or incoming right- and left-winding edges. Then all outgoing left-winding edges are to the left of all outgoing right-winding edges, and all incoming left-winding edges are to the right of all incoming right-winding edges.

It is a crucial fact that a winding together with a horizontal connection partitions the given drawing into a left and a right part. In the plane this is a path from the bottom to the top of the drawing and

a connection of the ends in the outer face.

Lemma 3. *Let $\Gamma(G)$ be an upward planar drawing on a rolling cylinder. A curve C along edges of G from a point $p = (x, y, z)$ to a point $p' = (x', y, z)$ together with the horizontal line L between p and p' and such that C does not cross L partitions $\Gamma(G)$ and G into a left and a right part G_l and G_r, such that each edge between G_l and G_r must meet a point on $C + L$.*

In particular, if C is a segment of an edge winding once around the cylinder, then each edge between G_l and G_r intersects the horizontal line L.

Proof. $C + L$ is a closed curve on the rolling cylinder. Now planarity and the Jordan's curve theorem enforces a partition of$\Gamma(\ G)$ into an inner and outer part, which induces a partition of G into G_l and G_r, respectively. By the planarity the edges between vertices of G_l and G_r must cross $C + L$ and hence must pass through points of $C + L$. □

For convenience we shall assume henceforth that there are only right-winding edges; left-winding is symmetric, where left and right must be interchanged.

Moreover, the out-going edges of a vertex v are regarded as the first edges intersecting the level $L(v)$ immediately to the right of v and passing v to the right. The out-going edges are taken from left to right in clockwise order. Accordingly, the in-coming edges of a vertex v' in counter-clockwise order are the last edges passing v' to its left. Each vertex has at least an out-going or an in-coming edge, since the underlying graph is connected.

Definition 1. *Let $\Gamma(G)$ be an upward planar drawing of a graph G on the rolling cylinder without left-winding edges.*

For a vertex v let $s_r(v) = e_1, \ldots, e_r$ be the maximal sequence of edges which intersect $L(v)$ immediately to the right of v and such that each edge occurs at most twice and there is no other vertex on $L(v)$ between v and the rightmost intersection point from e_r.

$s_r(v)$ is the right-signature *of v, where the edges are taken as symbols and $s_r(v)$ as a string. The out-going edges of v are the prefix of this string.*

The left-signature *consisting of the sequence of edges $s_l(v)$ entering v is defined symmetrically. Here, an edge may intersect or meet the level of v at most twice to the left of v, and the in-coming edges are the suffix of $s_l(v)$.*

The asymmetry of left- and right-signatures comes from the restriction to right-winding edges. For left-winding it is reversed. The pair $s_l(v)$ and $s_r(v)$ is characteristic for the vertex v, since v has in- or out-going edges, which appear as suffix and prefix of the signatures, respectively. However, there may be several vertices with the same right (or left) signature. If $s_r(v) = s_r(v')$ then v' is a source and lies to the right of the out-going edges of v and on a level before v (or vice-versa). Moreover, windings of edges and pure windings can easily be described and detected using the signatures. An edge (u, v) is winding if it repeats in $s_r(u)$ and $s_l(v)$.

If there are no edges with more than two windings, then $s_r(v)$ is the sequences of all edges crossing the level $L(v)$ from left to right to the right of v, provided there is no other vertex on $L(v)$. This sequence has a common substring with $s_l(v')$, where v' defines the next level after $L(v)$ in the clockwise order of the levels.

If there are frequently (right-) winding edges, then $s_r(v)$ takes only the first two windings to the right into account. Further windings are not observed by v. They are either observed by other vertices or there are pure windings of a bundle of edges β. A bundle of *purely winding* edges is easily detected, since $s_r(v)$ is of the form $\beta\,\beta$, where the outgoing edges of v are a prefix of β. β occurs as a maximal substring of $s_l(v')$ for some vertices v', from which we choose the *winding successor* as the leftmost one. If an edge from $s_r(v)$ reaches its endpoint during the next two windings or if new edges start at a vertex then the sequence of edges does not repeat. In that case v is static and observes its passing edges for less than two rounds and then passes control to another vertex.

Definition 2. *Let $\Gamma(G)$ be an upward planar drawing on the rolling cylinder.*

A vertex v is an out-winder *if its right-signature $s_r(v)$ is of the form $\beta\,\beta$ for a non-empty sequence of edges β; otherwise, v is* static. *v* is an in-winder if $s_l(v)$ is of the form $\gamma\,\gamma$ for a nonempty sequence of edges γ.

Let v be an out-winder with $s_r(v) = \beta\,\beta$. A vertex v' is the winding successor of v if v' is an out-winder with $s_l = \alpha\beta\gamma$, where α and γ are substrings of β and α is of minimal length.

Lemma 4. *Let $\Gamma(G)$ be a rup-drawing. Every* out-winder *v has a unique* winding successor.

Proof. At every vertex v there is a change in the passing edges just below and above $L(v)$, since G is connected.

Let v be an out-winder with $s_r(v) = \beta\,\beta$. Then the edges of β have at least two windings around the cylinder immediately to the right of v. Let e be the first edge of β. Then the first winding of e together with the horizontal line from the first to the second intersection of e and $L(v)$ is a closed curve. By Lemma 3 the graph is partitioned and all edges from one part to the other pass the horizontal line. These are the remaining edges of β.

Consider the intersection points of $L(v)$ to the right of v. There are further intersections of the edges from β and $L(v)$ after more windings. Consider the last complete winding of the edges of β on $L(v)$, and thereafter, consider the edges of β from left to right and follow them for at most one round, until there is a change for the first time. This change occurs at a vertex v', where some edges of β end or new edges start. v' is the leftmost such vertex with respect to winding edges. Then the left-signature of v' is $s_l(v') = \beta''\,\beta\,\beta'\,\mathrm{in}(v'')$, where $\beta = \beta'\,\mathrm{in}(v')\,\beta''$, since the edges from β make at least two windings. Moreover, the length of β'' is minimal for all vertices with β as a substring of their left-signature, since v' is reached first. $s_l(v')$ is of the form $\gamma\,\gamma$ with $\gamma = \beta''\beta'\,\mathrm{in}(v'')$. Hence, v' is an in-winder and β is a substring of $s_l(v')$. The uniqueness of v' follows from the minimality of the length of β''. □

Example 1. Suppose that vertex v is the source of an edge a and v' is the winding successor such that c ends at v', and there is a bundle $abcd$ of frequently winding edges. Then $s_r(v) = a\,b\,c\,d\,a\,b\,c\,d$ and $s_l(v') = d\,a\,b\,c\,d\,a\,b\,c$. The signature or string of edges intersecting $L(v)$ to the right of v and up to v' is of the form $abcd(abcd)^+ab[c]$, where "+" means at least one repetition and brackets mean at most once.

Later on there is a short cut from level $L(v')$-1 to level $L(v')$, where the leftmost intersection points of the edges in the bundle $dabc$ on level $L(v')$-1 are directly connected with the intersection points of $dabc$ immediately to the left of v' on level $L(v')$ such that c enters vertex v'.

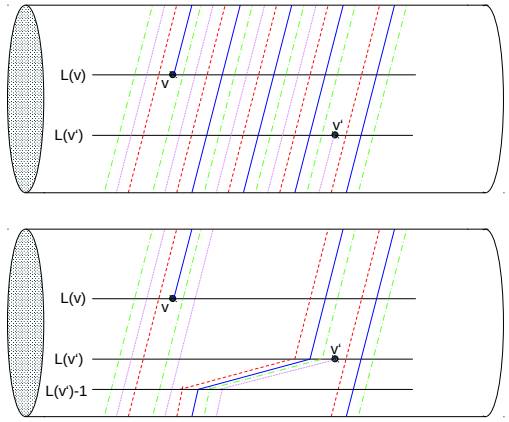

Figure 2: A bundle of winding edges and its short cut

This situation is illustrated in Figure 2.

Let$\Gamma(\ G)$ be an upward planar drawing on a rolling cylinder, and treat$\Gamma(\ G)$ as a recurrent hierarchy with levels $L(v)$, which are ordered from L_0 to L_{k-1}. We smooth the edges from$\Gamma(\ G)$.

If v is static then some of the edges passing or leaving v reach their endpoint during the next two windings. Then v observes the edges of $s_r(v)$ for at most two rounds on the cylinder. Each edge segment from level $L(v)$ to the next level is replaced by the geodesic between the points of the intersection of the curve $J(e)$ with the levels.

If v is an out-winder with a frequently winding bundle of edges β consider the *winding successor* v'. Now v observes the edges from β from its level $L(v)$ to level $L(v')$-1 and replaces the curves level by level by geodesics. Between the levels $L(v')$-1 and $L(v')$ there is a short cut.

Lemma 5. *Let $\Gamma(G)$ be an upward planar drawing on a rolling cylinder. Then there is an upward planar polyline drawing $\Gamma'(G)$ with at most $\mathcal{O}(n)$ many windings per edge.*

$\Gamma'(G)$ can be computed from $\Gamma(G)$ in time $\mathcal{O}(\tau\, n^3)$, where τ is the time to compute an observable point.

Proof. If v is static it observes the edges on its level and the edges of $s_r(v)$ for up to two rounds on the cylinder. Every edge segment from a level L_i to the next level is replaced by the geodesic between the points of the intersection of the curve $J(e)$ with the levels. This preserves the left-to-right order at the vertices and thus planarity.

If v is an out-winder with $s_r(v) = \beta\,\beta$ it observes the edges of $\beta = e_1, \ldots, e_q$ from level $L(v)$ to level $L(v')$ - 1, where v' is the winding successor of v. As above the curve of each edge e_i is replaced by segments of geodesics between adjacent levels. All edges e_i with $i = 1, \ldots, q$ wind at least $r \geq 2$ times around the cylinder, they do it in a bundle, and some end at v'. e_1 is the first outgoing edge of v or the first edge passing v to the right if v has no outgoing edges. Between $L(v')$-1 and $L(v)$ the curves are replaced by geodesics from the intersection points of the first winding after v on $L(v')$-1 to the intersection points on $L(v')$, when the edges pass v' immediately to the left or enter v'. See Figure 2 for an illustration.

All edges remain by the smoothing and the short cut, so the graph is unchanged. The transformation preserves the order of all segments of edges on all levels. Since all edges between adjacent levels are replaced by geodesics this excludes crossings. Also the upward direction on the rolling cylinder is preserved. Hence,$\Gamma\ '(G)$ is upward planar.

All vertices of$\Gamma\ '(G)$ are static and edges can wind at most $2n$ times around the cylinder.

Concerning the time complexity every vertex observes edges for at most two windings. Hence, it controls at most $\mathcal{O}(n^2)$ many observable points. In total, $\mathcal{O}(n^3)$ many observable points must be taken into account, each of which can be computed in time τ. □

Next we transform the polyline drawing with at most $\mathcal{O}(n)$ windings into one with at most one winding per edge. If an edge winds at least twice then the subgraphs between successive windings are upward planar in the plane. Their compression and contraction eliminates an edge winding. If there is a bundle of winding edges, all windings are reduced simultaneously. Then the complexity is only linear in the number of edge segments of such polyline drawings.

Lemma 6. *For every upward planar polyline drawing of a graph $\Gamma'(G)$ there is an upward planar polyline drawing $\tilde{\Gamma}(G)$ with at most one winding per edge.*

$\tilde{\Gamma}(G)$ can be computed in $\mathcal{O}(n^3)$ time if $\Gamma'(G)$ has at most $\mathcal{O}(n)$ windings per edge.

Proof. Again consider the drawing$\Gamma\ '(G)$ as a recurrent hierarchy with horizontal levels $L(v)$ for the vertices and geodesics as segments of edges between adjacent levels, as obtained in Lemma 5.

First we consider the removal of a single winding. The extension towards many windings follows by induction. For the time bound it needs some adjustments.

Let $e_1 = (u, v)$ be the leftmost edge with at least two windings. Assume that e_1 is right-winding (otherwise proceed in the opposite direction from v). If e_1 makes exactly two windings, lift v by some small amount such that $L(u) \neq L(v)$ and planarity is preserved. Thereafter e_1 has more than two windings. Also suppose that there is no other vertex w on $L(u)$; otherwise move w slightly up or down. For convenience, rotate the cylinder such that $L(u)$ is at the bottom line $(0, -1)$ in the (y, z) plane and unroll it such that there are no vertices and levels on the backside. So $L(u)$ reappears at the top, where it is called the top line $\hat{L}(u)$. (In fact this is the fundamental polygon approach). Let $L(u')$ be the first level following $L(u)$.

Let $p_1, \ldots, p_r$ be the leftmost points of intersection of the edges $e_1, \ldots, e_r$ with the top line $\hat{L}(u)$, where p_r is the point of intersection of e_1 in the second winding. These points reappear on the bottom line. Let $q_1, \ldots, q_r$ be the points of intersection of the edges $e_1, \ldots, e_r$ with $L(u')$ between the second and third intersection of e_1. These are the endpoints of the segments of the edges from p_i to q_i for $i = 1, \ldots, r$ on $L(u')$. See Figure 3 for an illustration.

Consider the stripe of the drawing between the top and bottom lines and with the first winding of e_1 to the left and the second winding to the right. This subgraph is upward planar, since the bottom and top lines serve as s and t and e_1 as the s-t edge.

Introduce a new level L' just below $L(u')$. Now compress the stripe and contract it towards the bottom line such that it fits into the trapezoid with the bottom line $L(u)$, the top line L' and the segments of e_1 from the first and second windings to the left and right.

Let $p'_1, \ldots, p'_r$ be the points of intersection of the edges $e_1, \ldots, e_r$ with L' after this compression and

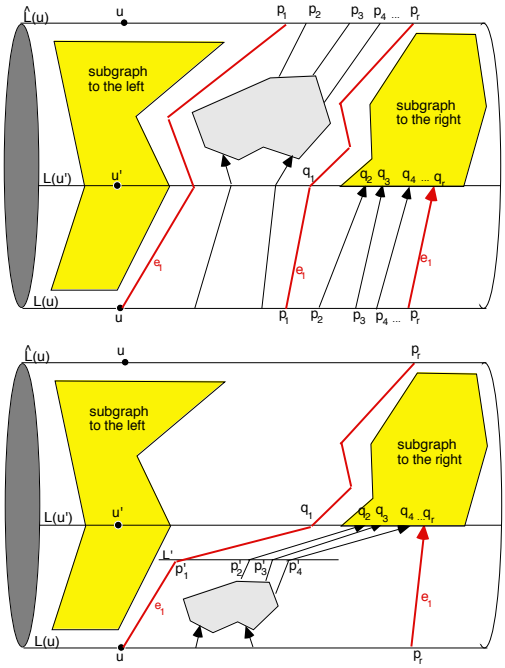

Figure 3: Reducing edge windings

contraction. As the main step reroute the edges e_i from p'_i on level L' to q_i on level $L(u')$ and replace it by the geodesic (p'_i, q_i) for $i = 0, \ldots, r-1$. All other segments of edges are kept.

This transformation leaves the subgraphs to the left of the first and to the right of the second windings of e untouched by Lemma 3. It preserves the upward planarity of the compressed subgraph in the stripe. The rerouting of the edges between the levels L' and $L(u')$ preserves the given left-to-right order and thus planarity, and the upwardness of the edge segments. It decreases the number of windings of e_1 and does dot increase the windings of other edges. By induction, we obtain the drawing $\tilde{\Gamma}(G)$, where no edge winds more than once.

However, the previous transformation takes $\mathcal{O}(n^2)$ time, since up to $\mathcal{O}(n^2)$ segments and points are involved, and there are $\mathcal{O}(n)$ edges with $\mathcal{O}(n)$ windings. To achieve a cubic time bound consider the most frequently winding edge $e = (u, v)$ with k windings. Starting from $L(u)$ compress the subgraphs between the first k-1 windings and stack them such that they fit between $L(u)$ and L' as above, and then reroute the edge segments from L' to the original points at $L(u')$. The cost of this transformation is linear in the number of segments. This procedure is repeated until all windings are treated. Maximal edge windings do not interact. A winding belongs to at most two maximal windings. Hence, the subgraph and the segments between two windings is treated at most twice. Since there are at most $\mathcal{O}(n^3)$ segments the total time is cubic, too. □

Edge windings cannot be eliminated on the rolling cylinder, as is was doable on the standing cylinder. On the standing cylinder horizontal rotations (of disks) and the upward direction are orthogonal and independent. On the rolling cylinder they are

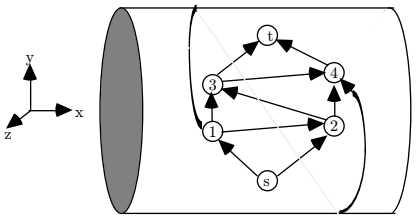

Figure 4: An unavoidable edge winding

unidirectional. An extension of the st-$K_{2,2}$ serves as a counterexample.

Lemma 7. *There are rup-graphs which need winding edges on the rolling cylinder.*

Proof. Consider the graph in Figure 4. For upward planarity the vertices s and t must be on the outer face on the cylinder. Thus the embedding is unique up to reflection. Then the edge $(1, 4)$ must wind around the cylinder, since the paths $(1, 2, 4)$ and $(1, 3, 4)$ are upward. □

In our final step we reduce the number of bends to at most five per edge. Therefore, we adapt the technique of Di Battista and Tamassia [13] from visibility representations and monotone grid drawings.

Lemma 8. *Let $\tilde{\Gamma}(G)$ be an upward planar polyline drawing with at most one winding per edge. Then there is an upward planar drawing $\hat{\Gamma}(G)$ with at most one winding and at most five bends per edge. Moreover, the vertices and bends can be placed on a grid of size $\mathcal{O}(n^2)$. The computation of $\hat{\Gamma}(G)$ from $\tilde{\Gamma}(G)$ takes $\mathcal{O}(n)$ time.*

Proof. Assume that there is no vertex with coordinates $(x, 0, 1)$. Cut $\tilde{\Gamma}(G)$ at the horizontal front line L with $(y, z) = (0, 1)$ and place a dummy vertex at each intersection point of an edge e with L. Duplicate L into a low and a high copy L_l and L_h and add a source s just below L_l and a sink t just above L_h. The so obtained graph G' is upward planar and has a monotonic grid drawing with at most two bends per edge [13] on a grid of quadratic size. The algorithm VISIBILITY-DRAW from [13] places all vertices from L_l one level above s. It can be modified to place all vertices on L_h one level below t. The algorithm GRID-DRAW from [13] contracts the horizontal segments for the vertices to grid points at the cost of two bends per edge. It can be modified to place the vertices of L_l and L_h at points with the same x-coordinates. Both algorithms operate in linear time.

Finally, remove s and t and embed the underlying grid and G' on the rolling cylinder by the identification of L_l and L_h. Then the edges crossing L have up to five bends, two before and after L and one at L. □

We summarize:

Theorem 2. *For every upward planar drawing $\Gamma(G)$ of a graph G on the rolling cylinder there is an upward planar drawing $\hat{\Gamma}(G)$*

- *with polylines and geodesics as straight segments*
- *at most one winding per edge*
- *at most five bends per edge*

- *a placement of all vertices and bends on a grid of quadratic size, and*
- *which can be computed in time $\mathcal{O}(\tau\, n^3)$, where τ is the time to compute an intersection of an edge and a horizontal line.*

5 Conclusion and Open Problems

We have simplified planar upward drawings on the standing and rolling cylinders to polyline drawings with at most one winding and only a few bends per edge and with vertices and bends placed on a grid.

It is unknown whether or not there are straight planar upward drawings on cylinders. If bends are unavoidable, what is the least number of bends in total and per edge? The approach of Bachmaier et al. [3] promises only two bends per edge also in the rolling cylinder. However, it needs edges which do not wind.

Finally, do similar simplifications of drawings hold for more general classes of graphs, such as quasi-upward planar graphs [8] or graphs with an upward planar drawing on the torus [1]?

6 Acknowledgement

I wish to thank C. Auer, C. Bachmaier and A. Gleißner for inspiring discussions and useful hints, and in particular, the observation that the front line in the proof of Lemma 1 is not a path, and an anonymous referee the useful comments, particularly on winding edges on the rolling cylinder.

References

[1] C. Auer, C. Bachmaier, F. J. Brandenburg, and A. Gleißner. Classification of plane upward embedding. In M. van Kreveld and B. Speckmann, editors, *GD 2011*, volume 7034 of *LNCS*, pages 415–426, 2011.

[2] C. Bachmaier. A radial adaption of the sugiyama framework for visualizing hierarchical information. *IEEE Trans. Vis. Comput. Graphics*, 13(3):583–594, 2007.

[3] C. Bachmaier, F. J. Brandenburg, W. Brunner, and R. Fülöp. Coordinate assignment for cyclic level graphs. In H. Q. Ngo, editor, *COCOON 2009*, volume 5609 of *LNCS*, pages 66–75, 2009.

[4] C. Bachmaier, F. J. Brandenburg, W. Brunner, and F. Hübner. Global k-level crossing reduction. *J. Graph Algorithms Appl.*, 2011. (to appear). A short version has appeared in M.S. Rahman and S. Fujita (editors), WALCOM 2010, volume 5942 of LNCS, pages 70–81, 2010.

[5] C. Bachmaier, F.-J. Brandenburg, and M. Forster. Radial level planarity testing and embedding in linear time. *J. Graph Algorithms Appl.*, 9(1):53–97, 2005.

[6] C. Bachmaier and W. Brunner. Linear time planarity testing and embedding of strongly connected cyclic level graphs. In D. Halperin and K. Mehlhorn, editors, *ESA 2008*, volume 5193 of *LNCS*, pages 136–147, 2008.

[7] C. Bachmaier, W. Brunner, and C. König. Cyclic level planarity testing and embedding. In S.-H. Hong, T. Nishizeki, and W. Quan, editors, *Graph Drawing*, volume 5942 of *LNCS*, pages 50–61, 2007.

[8] P. Bertolazzi, G. Di Battista, and W. Didimo. Quasi-upward panarity. *Algorithmica*, 32(3):474–506, 2002.

[9] F. J. Brandenburg. Upward planar drawings on the standing and rolling cylinders. In *preparation*, 2011.

[10] U. Brandes and B. Köpf. Fast and simple horizontal cordinate assignment. In P. Mutzel, M. Jünger, and S. Leipert, editors, *Graph Drawing*, volume 2265 of *LNCS*, pages 31–44, 2002.

[11] H. de Fraysseix, J. Pach, and R. Pollack. How to draw a planar graph on a grid. *Combinatorica*, 10:41–51, 1990.

[12] G. Di Battista, P. Eades, R. Tamassia, and I. G. Tollis. *Graph Drawing: Algorithms for the Visualization of Graphs*. Prentice Hall, 1999.

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

[14] G. Di Battista, R. Tamassia, and I. G. Tollis. Area requirement and symmetry display of planar upward drawings. *Discrete Comput. Geom.*, 7:381–401, 1992.

[15] D. P. Dobkin, E. R. Gansner, E. Koutsofios, and S. C. North. Implementing a general-purpose edge router. In G. Di Battista, editor, *Graph Drawing*, volume 1353 of *LNCS*, pages 262–271, 1997.

[16] A. Dolati and S. M. Hashemi. On the sphericity testing of single source digraphs. *Discrete Math.*, 308(11):2175–2181, 2008.

[17] M. Eiglsperger, M. Siebenhaller, and M. Kaufmann. An efficient implementation of sugiyama's algorithm for layered graph drawing. *J. Graph Algorithms Appl.*, 9(3):305–325, 2005.

[18] I. Fáry. On straight line representation of planar graphs. *Acta Sci. Math. Szeged*, 11:229–233, 1948.

[19] K. A. Hansen. Constant width planar computation characterizes ACC^0. *Theor. Comput. Sci.*, 39(1):79–92, 2006.

[20] S. M. Hashemi. Digraph embedding. *Discrete Math.*, 233(1–3):321–328, 2001.

[21] S. M. Hashemi, I. Rival, and A. Kisielewicz. The complexity of upward drawings on spheres. *Order*, 14:327–363, 1998.

[22] M. Kaufmann and D. Wagner. *Drawing Graphs*, volume 2025 of *LNCS*. Springer, 2001.

[23] D. Kelly. Fundamentals of planar ordered sets. *Discrete Math.*, 63:197–216, 1987.

[24] N. Limaye, M. Mahajan, and J. M. N. Sarma. Evaluating monotone circuits on cylinders, planes and tori. In B. Durand and W. Thomas, editors, *STACS 2006*, volume 3884 of *LNCS*, pages 660–671, 2006.

[25] N. Limaye, M. Mahajan, and J. M. N. Sarma. Upper bounds for monotone planar circuit value and variants. *Computational Complexity*, 18(3):377–412, 2009.

[26] B. Mohar and P. Rosenstiel. Tessellation and visibility representations of maps on the torus. *Discrete Comput. Geom.*, 19:249–263, 1998.

[27] B. Mohar and C. Thomassen. *Graphs on Surfaces*. John Hopkins University Press, 2001.

[28] W. Schnyder. Embedding planar graphs on the grid. In *Proc 1st Annual ACM-SIAM Sympos. on Discrete Algorithms (SODA), San Francisco*, pages 138–147, 1990.

[29] S. Stein. Convex maps. *Proc. Amer. Math. Soc.*, 2:464–466, 1951.

[30] E. Steinitz and H. Rademacher. *Vorlesungen über die Theorie der Polyeder*. Julius Springer, Berlin, 1934.

[31] K. Sugiyama, S. Tagawa, and M. Toda. Methods for visual understanding of hierarchical system structures. *IEEE Trans. Systems, Man, and Cybernetics*, 11(2):109–125, 1981.

[32] K. Wagner. Bemerkungen zum Vierfarbenproblem. *Jahresbericht Deutsche Math.-Vereinigung*, 46:26–32, 1936.