Canonical Ordering for Triangulations on the Cylinder, with Applications to Periodic Straight-line Drawings

Luca Castelli Aleardi, Olivier Devillers, Eric Fusy

Canonical Ordering for Triangulations on the Cylinder, with Applications to Periodic Straight-line Drawings[*]

Luca Castelli Aleardi[†]   Olivier Devillers[‡]   Éric Fusy[§]

Abstract

We extend the notion of canonical orderings to cylindric triangulations. This allows us to extend the incremental straight-line drawing algorithm of de Fraysseix, Pach and Pollack to this setting. Our algorithm yields in linear time a crossing-free straight-line drawing of a cylindric triangulation G with n vertices on a regular grid $\mathbb{Z}/w\mathbb{Z} \times [0..h]$, with $w \leq 2n$ and $h \leq n(2d+1)$, where d is the (graph-) distance between the two boundaries. As a by-product, we can also obtain in linear time a crossing-free straight-line drawing of a toroidal triangulation with n vertices on a periodic regular grid $\mathbb{Z}/w\mathbb{Z} \times \mathbb{Z}/h\mathbb{Z}$, with $w \leq 2n$ and $h \leq 1 + n(2c+1)$, where c is the length of a shortest non-contractible cycle. Since $c \leq \sqrt{2n}$, the grid area is $O(n^{5/2})$. Our algorithms apply to any triangulation (whether on the cylinder or on the torus) that have no loops nor multiple edges in the periodic representation.

1   Introduction

The problem of efficiently computing straight-line drawings of planar graphs has attracted a lot of attention over the last two decades. Two combinatorial concepts for planar triangulations turn out to be the basis of many classical straight-line drawing algorithms: the *canonical ordering* (a special ordering of the vertices obtained by a shelling procedure) and the closely related *Schnyder wood* (a partition of the inner edges of a triangulation into 3 spanning trees with specific incidence conditions). Algorithms based on canonical ordering [7, 10] are typically incremental, adding vertices one by one while keeping the drawing planar. Algorithms based on Schnyder woods [14] are more global, the (barycentric) coordinates of each vertex have a clear combinatorial meaning (typically the number of faces in certain regions associated to the vertex). Algorithms of both types make it possible to draw in linear time a planar triangulation with n vertices on a grid of size $O(n \times n)$. They can also both be extended to obtain (weakly) convex drawings of 3-connected maps on a grid of

size $O(n \times n)$. The problem of obtaining planar drawings of higher genus graphs has been addressed less frequently [11, 9, 12, 13, 4, 6, 15], from both the theoretical and algorithmic point of view. Recently some methods for the straight-line planar drawing of genus g graphs with polynomial grid area (of size $O(n^3)$, in the worst case) have been described in [4, 6] (to apply these methods the graph needs to be unfolded planarly along a *cut-graph*). However, it does not yield (at least easily) periodic representations: for example, in the case of a torus, the boundary vertices (on the boundary of the rectangular polygon) might not be aligned, so that the drawing does not give rise to a periodic tiling.

Our main contribution is to generalize the notion of canonical ordering and the incremental straight-line drawing algorithm of de Fraysseix, Pach and Pollack [7] (shortly called FPP algorithm thereafter) to triangulations on the cylinder. For any triangulation G of the cylinder, our algorithm yields in linear time a crossing-free straight-line drawing of G on a regular grid (on the flat cylinder) of the form $\mathbb{Z}/w\mathbb{Z} \times [0..h]$, with $w \leq 2n$ and $h \leq n(2d+1)$, where n is the number of vertices of G and d is the (graph-) distance between the two boundaries of G. By a reduction to the cylindric case (the reduction is done with the help of a so-called *tambourine* [2]), we also get a drawing algorithm on the torus. Precisely, for any toroidal triangulation G, we can obtain in linear time a crossing-free straight-line drawing of G on a regular grid (on the flat torus) $\mathbb{Z}/w\mathbb{Z} \times \mathbb{Z}/h\mathbb{Z}$, with $w \leq 2n$ and $h \leq 1 + n(2c-1)$, where n is the number of vertices of G and c is the length of a shortest non-contractible cycle. Since $c \leq (2n)^{1/2}$ as shown in [1], we have $h \leq (2n)^{3/2}$, so that the grid area is $O(n^{5/2})$.

For the toroidal case we mention that a notion of canonical ordering has been introduced in [5] (this actually works in any genus and yields an efficient encoding procedure) but we do not use it here. We also mention that, independently, an elegant periodic straight-line drawing algorithm for toroidal triangulations has been very recently described in [8], based on so-called *toroidal Schnyder woods* and face-counting operations; in their case the area of the periodic grid is $O(n^4)$.

2 Preliminaries

Graphs embedded on surfaces.   A *map* of genus g is a connected graph G embedded on the compact orientable surface S of genus g, such that all components of $S \backslash G$ are topological disks, which are called the *faces* of the map. The map is called *planar* for $g = 0$ (embedding on the sphere) and *toroidal* for $g = 1$ (embedding on the torus). The *dual* of a map G is the map G^* representing the adjacencies of the faces of G, i.e., there is a vertex v_f of G^* in each face f of G, and each edge e of G gives rise to an edge $e^* = \{v_f, v_{f'}\}$ in G^*, where f and f' are the faces on each side of e. A *cylindric map* is a planar map with two marked faces B_1 and B_2 whose boundaries $C(B_1)$ and $C(B_2)$ are simple cycles (possibly $C(B_1)$ and $C(B_2)$ share vertices and edges). The faces B_1 and B_2 are called the *boundary-faces*. Boundary vertices and edges are those belonging to $C(B_1)$ (black circles in Fig. 1) or $C(B_2)$ (gray circles in Fig. 1); the other ones are called *inner* vertices (white circles in Fig. 1) and edges.

Periodic drawings.   Here we consider the problem of drawing a cylindric triangulation on the flat cylinder and drawing a toroidal triangulation on the flat torus. For $w > 0$ and $h > 0$,

Figure 1: A cylindric triangulation with boundary faces $B_1 = \{a, b, c, d\}$ and $B_2 = \{h, i, j, k\}$. Left: annular representation. Right: x-periodic representation.

the *flat cylinder* of width w and height h is the rectangle $[0, w] \times [0, h]$ where the vertical sides are identified. A point on this cylinder is located by two coordinates $x \in \mathbb{R}/w\mathbb{Z}$ and $y \in [0, h]$. The *flat torus* of width w and height h is the rectangle $[0, w] \times [0, h]$ where both pairs of opposite sides are identified. A point on this torus is located by two coordinates $x \in \mathbb{R}/w\mathbb{Z}$ and $y \in \mathbb{R}/h\mathbb{Z}$. Assume from now on that w and h are positive integers. For a cylindric triangulation G, a *periodic straight-line drawing* of G of width w and height h is a crossing-free straight-line drawing (edges are drawn as segments, two edges can meet only at common end-points) of G on the flat cylinder of width w and height h, such that the vertex-coordinates are in $\mathbb{Z}/w\mathbb{Z} \times [0..h]$ (i.e., are integers). Similarly for a toroidal triangulation G, a *periodic straight-line drawing* of G of width w and height h is a crossing-free straight-line drawing (edges are drawn as segments, two edges can meet only at common end-points) of G on the flat torus of width w and height h, such that the vertex-coordinates are in $\mathbb{Z}/w\mathbb{Z} \times \mathbb{Z}/h\mathbb{Z}$ (i.e., are integers).

3 Periodic drawings of cylindric triangulations

A *cylindric triangulation* is a cylindric map with no loops nor multiple edges and such that all non-boundary faces are triangles (in Section 5 we will allow such triangulations to have 1-cycles or 2-cycles separating the two boundary-faces). We introduce at first a notion of canonical ordering for cylindric triangulations:

Definition 1. *Let G be a cylindric triangulation with boundary-faces B_1 and B_2, and such that the cycle $C(B_1)$ has no chords (i.e., there is no edge that is not on $C(B_1)$ and has both ends on $C(B_1)$). An ordering $\pi = \{v_1, v_2, \ldots, v_n\}$ of the vertices of $G\backslash C(B_1)$ is called a (cylindric) canonical ordering if it satisfies:*

- *For each $k \geq 0$ the map G_k induced by $C(B_1)$ and by the vertices $\{v_1, \ldots, v_k\}$ is a cylindric triangulation.*

- *The other boundary-face B' of G_k (the one different from B_1), whose contour is denoted C_k, contains B_2 (to clarify the meaning of "B' contains B_2", remember that a cylindric map is a special kind of map on the sphere; since G_k is a submap of G, any face of G_k is made of a set of faces of G).*

- *The vertex v_k is on C_k, and its neighbours in G_{k-1} are consecutive on C_{k-1}.*
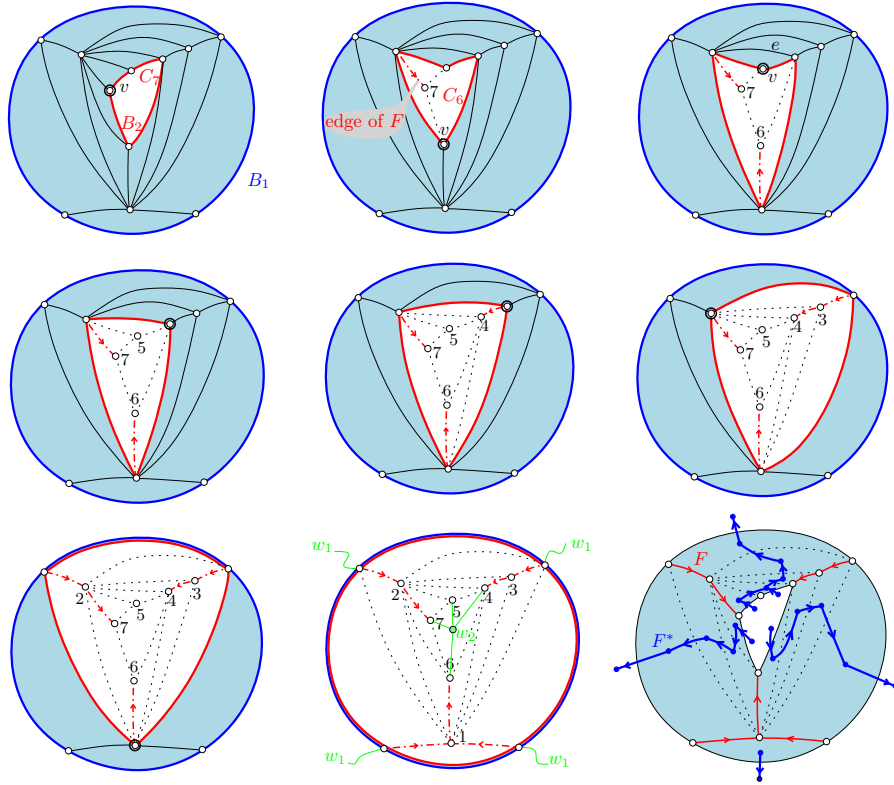
3

Figure 2: Shelling procedure to compute a canonical ordering of a given cylindric triangulation. The underlying forest is computed on the fly; the last drawing shows the underlying forest superimposed with the dual forest. The graph is the one of Fig. 1.

The notion of canonical ordering makes it possible to construct a cylindric triangulation G incrementally, starting from $G_0 = C(B_1)$ and adding one vertex (and incident edges) at each step. This is similar to canonical orderings for planar triangulations, as introduced by de Fraysseix, Pach and Pollack [7] (the main difference is that for a planar triangulation one starts with G_0 being an edge, called the *base-edge*, whereas here one starts with G_0 being a cycle, seen as a cylindric map without non-boundary faces).

Shelling procedure. We now describe a shelling procedure to compute a canonical ordering of a cylindric triangulation G with boundary-faces B_1, B_2. At each step the graph formed by the remaining vertices is a cylindric triangulation, one boundary face remains B_1 all the way, while the other boundary-face (initially B_2) has its contour, denoted by C_k, getting closer to $C(B_1)$. A vertex $v \in C_k$ is *free* if v is incident to no chord of C_k and if $v \notin C(B_1)$ (see Fig. 2 top left). The shelling procedure goes as follows (n is the number of vertices in $G \backslash C(B_1)$): for k from n to 1, choose a free vertex v on C_k, assign $v_k \leftarrow v$, and then delete v together with all its incident edges. The existence of a free vertex at each step follows from the same argument as in the planar case [3]. Indeed, if there is no chord incident to C_k, then any vertex $v \in C_k \backslash C(B_1)$ is free, while if there is a chord e for C_k, then the set of chords incident to C_k forms a system of archs (relative to C_k). If we look at a chord $e = \{u, v\}$ that is "innermost" (i.e., seen as an arch, no other arch is nested inside), then the path between u and v on C_k contains at least one vertex, which has to be free (see Fig. 2 top right).

4

Figure 3: One step of the incremental drawing algorithm. Two vertical strips of width 1 (each one along a path in the dual forest) are inserted in order to make the slopes of e_ℓ and e_r smaller than 1 in absolute value. Then the new vertex and its edges connected to the upper boundary can be drawn in a planar way.

Underlying forest and dual forest. Given a cylindric triangulation G (with boundary faces B_1 and B_2) endowed with a canonical ordering π, define the *underlying forest F* for π as the oriented subgraph of G where each vertex $v \in C(B_2)$ has outdegree 0, and where each $v \notin C(B_2)$ has exactly one outgoing edge, which is connected to the adjacent vertex of v of largest label in π. The forest F can be computed on the fly during the shelling procedure: when treating a free vertex v_k, for each neighbour v of v_k such that $v \notin C_k$, add the edge $\{v, v_k\}$ to F, and orient it from v to v_k. Since the edges are oriented in increasing labels, F is an oriented forest; it spans all vertices of $G \backslash C(B_2)$ and has its roots on $C(B_2)$. The *augmented map $\widehat{G}$* is obtained from G by adding a vertex w_1 inside B_1, a vertex w_2 inside B_2, and connecting all vertices around B_1 to w_1 and all vertices around B_2 to w_2 (thus triangulating the interiors of B_1 and B_2, see Fig. 2 bottom middle). Define $\widehat{F}$ as F plus all edges incident to w_1 and all edges incident to w_2. Define the *dual forest F^** for π as the graph formed by the vertices of $\widehat{G}^*$ (the dual of $\widehat{G}$) and by the edges of $\widehat{G}^*$ that are dual to edges not in $\widehat{F}$. Since $\widehat{F}$ is a spanning connected subgraph of $\widehat{G}$, F^* is a spanning forest of $\widehat{G}^*$. Precisely, each of the trees (connected components) of F^* is rooted at a vertex "in front of" each edge of B_1, and the edges of the tree can be oriented toward this root-vertex (see Fig. 2 bottom right).

Drawing algorithm. Given a cylindric triangulation G such that $C(B_1)$ has no chord, we first compute a canonical ordering of G, and then draw G in an incremental way. We start with a cylinder of width $2|C(B_1)|$ and height 0 (i.e., a circle of length $2|C(B_1)|$) and draw the vertices of $C(B_1)$ equally spaced on the circle (space 2 between two consecutive vertices). Then the strategy for each $k \geq 1$ is to compute the drawing of G_k out of the drawing of G_{k-1} by first stretching the cylinder (increasing the width by 2) and then placing the vertex v_k and its incident edges (in G_k) in a planar way. Define the *x-span* of an edge e in the cylindric drawing as the number of columns $[i, i+1] \times [0, +\infty]$ that meet the interior of e (we have no need for a more complicated definition since, in our drawings, a column will never meet an edge more than once).

Consider the dual forest F^* for the canonical ordering restricted to G_{k-1}. Note that the set of vertices of C_{k-1} that are neighbours of v_k form a path on C_{k-1}. Traversing this path
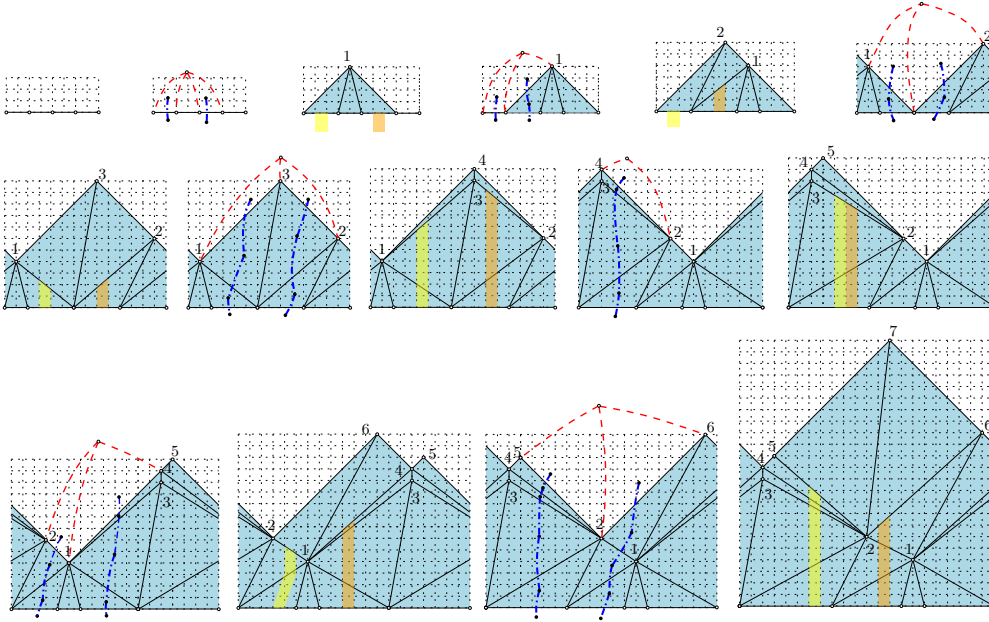
5

Figure 4: Algorithm (Prop. 1) to compute an x-periodic drawing of a cylindric triangulation (no chordal edges incident to B_1). The vertices are treated in increasing label (the canonical ordering is the one computed in Fig. 2).

γ with the boundary-face of G_{k-1} (the one different from B_1) to the left, let e_ℓ be the first edge of γ and e_r be the last edge of γ. Let also a_k be the starting vertex and let b_k be the ending vertex of γ. Let P_ℓ (resp. P_r) be the path in F^* from e_ℓ^* (resp. e_r^*) to the root in its connected component (which is a vertex "in front of" an edge of B_1)). We stretch the cylinder by inserting a vertical strip of length 1 along P_ℓ and another along P_r (see Fig. 3). This comes down to increasing by 1 the x-span of each edge of G_k dual to an edge in P_ℓ, and then increasing by 1 the x-span of each edge dual to an edge in P_r (note that P_ℓ and P_r are not necessarily disjoint, in which case the x-span of an edge dual to an edge in $P_\ell \cap P_r$ is increased by 2). After these stretching operations, whose effect is to make the slopes of e_ℓ and e_r strictly smaller than 1 in absolute value, we insert (as in the planar case) the vertex v_k at the intersection of the ray of slope 1 starting from a_k and the ray of slope -1 starting from b_k, and we connect v_k to all vertices of γ by segments[1]. These two rays actually intersects at a grid point since the Manhattan distance between any two vertices on C_{k-1} is even. Fig. 4 shows the execution of the algorithm on the example of Fig. 1.

Proposition 1. *For each cylindric triangulation G with no chord incident to $C(B_1)$, one can compute in linear time a crossing-free straight-line drawing of G on an x-periodic regular grid $\mathbb{Z}/w\mathbb{Z} \times [0..h]$ where —with n the number of vertices of G and d the (graph-)distance between the two boundaries— $w = 2n$ and $h \leq n(2d+1)$, such that: every $v \in C(B_1)$ has*

---

[1] In the FPP algorithm for planar triangulations, the step to make the (absolute value of) slopes of e_ℓ and e_r smaller than 1 is formulated as a shift of certain subgraphs described in terms of the underlying forest F. The extension of this formulation to the cylinder would be quite cumbersome. We find the alternative formulation with strip insertions more convenient for the cylinder (it also gives rise to a very easy linear-time implementation).

$y(v) = 0$ *(so every edge in $C(B_1)$ has slope 0), and every edge belonging to $C(B_2) \backslash C(B_1)$ has slope ± 1.*

Proof. The fact that the drawing remains crossing-free relies on the slope-property for the upper boundary and on the following inductive property, which is easily shown to be maintained at each step k from 1 to n:

Pl: for each edge e on C_k (the upper boundary of G_k), let P_e be the path in F^* from e^* to the root, let E_e be the set of edges dual to edges in P_e, and let δ_e be any nonnegative integer. Then the drawing remains planar when successively increasing by δ_e the x-span of all edges of E_e, for all $e \in C_k$.

We now prove the bounds on the grid-size. If $|C(B_1)| = t$ then the initial cylinder is $2t \times 0$; and at each vertex insertion, the grid-width grows by 2. Hence $w = 2n$. In addition, due to the slope conditions (slopes of boundary-edges are at most 1 in absolute value), the y-span (vertical span) of every edge e is not larger than the current width at the time when e is inserted in the drawing. Hence, if we denote by v the vertex of $C(B_2)$ that is closest (at distance d) to $C(B_1)$, then the ordinate of v is at most $d \cdot (2n)$. And due to the slope conditions, the vertical span of $C(B_2)$ is at most $w/2 \leq n$. Hence the grid-height is at most $n(2d + 1)$. The linear-time complexity is shown next. $\qquad \square$

Linear-time implementation. An important remark is that, instead of computing the x-coordinates and y-coordinates of vertices in the drawing, one can compute the y-coordinates of vertices and the x-span of edges (as well as the knowledge of which extremity is the left-end vertex and which extremity is the right-end vertex). In a first pass (for k from 1 to n) one computes the y-coordinate of vertices and the x-span r_e of each edge $e \in G$ at the time $t = k$ when it appears on G_k (as well one gets to know which extremity of e is the left-end vertex). Afterwards if $e \notin F$, the x-span of e might further increase due to insertion of new vertices. Precisely, let v_j be a vertex inserted afterwards (i.e., $j > t$), and (with the notations a_j, b_j of the drawing algorithm description) let $\epsilon_\ell = \{v_j, a_j\}$ and $\epsilon_r = \{v_j, b_j\}$. Note that e is stretched due to the insertion of the strip along P_ℓ iff ϵ_ℓ^* is in the subtree T_e of F^* formed by the edges descending from e^*. Similarly e is stretched due to the insertion of the strip along P_r iff ϵ_r^* is in T_e. To state it more clearly, each edge in T_e is responsible for an increase (by 1) of the x-span of e. Hence the total x-span of each edge $e \in G$ is given by $r_e + s_e$, where $s_e = 0$ if $e \in F$, and, if $e \notin F$, s_e is the number of edges in T_e. Since all quantities s_e can easily be computed in linear time, this gives a linear-time implementation.

Allowing for chordal edges at $\mathbf{C(B_1)}$. We finally explain how to draw a cylindric triangulation when allowing for chords incident to $C_0 = C(B_1)$; it is good to view B_2 as the top boundary-face and B_1 as the bottom-boundary face (and imagine a standing cylinder). For each chordal edge e for the cycle C_0, the *component under e* is the face-connected part of G that lies below e; such a component is a quasi-triangulation (polygonal outer face, triangular inner faces) rooted at the edge e. A chordal edge e of C_0 is *maximal* if the component Q_e under e is not strictly included in the component under another chordal edge for C_0. The *size* of such an edge e is defined as $|e| = 2|V(Q_e)| - 4$. (the size $|e|$ is actually

Figure 5: Drawing a cylindric triangulation with chords at B_1. To make enough space to place the component under e_2, one takes 8 (instead of 2) as the initial x-span of e_2.

the width of the drawing of Q_e using the FPP algorithm). If we delete the component under each maximal chordal edge (i.e., delete everything from the component except for the chordal edge itself) we get a new bottom cycle C_0' that is chordless, so we can draw the reduced cylindric triangulation G' using the algorithm of Proposition 1. As we have seen in Section 3 (implementation paragraph), for each edge e of C_0', the initial x-stretch is $r_e = 2$ and then the further increase s_e of the x-stretch equals the number of edges descending from e^* in the dual forest F^*. Note that we have actually some freedom to choose the initial x-stretch r_e of each edge $e \in C_0'$ (just it has to be a positive even number since at each step of the incremental algorithm the vertices of the current upper boundary have to be at even Manhattan distance). If $e \in C_0'$ is on C_0 we take $r_e = 2$. If $e \in C_0'$ is not on C_0 (i.e., e was a maximal chordal edge for C_0), we take for r_e the minimal even positive number such that $r_e + s_e \geq |e|$, i.e., $r_e = 2 \cdot \max(1, \lceil (|e| - s_e)/2 \rceil)$. Hence, at the end of the execution of the drawing of G', the length $\ell_e = r_e + s_e$ of each maximal chord e satisfies $\ell_e \geq |e|$. Then for each maximal chord e of C_0, we draw the component Q_e under e using the FPP algorithm. This drawing has width $|e|$, with e as horizontal bottom edge of length $|e|$ and with the other outer edges of slopes ± 1. We shift the left-extremity of e so that the drawing of Q_e gets width $\ell(e)$, then we rotate the drawing of Q_e by 180 degrees and plug it into the drawing of G' (see Fig. 5). The overall drawing of G is clearly planar. We obtain:

Theorem 1. *For each cylindric triangulation G, one can compute in linear time a crossing-free straight-line drawing of G on an x-periodic regular grid $\mathbb{Z}/w\mathbb{Z} \times [0, h]$, where —with n the number of vertices and d the (graph-) distance between the two boundaries— $w \leq 2n$ and $h \leq n(2d+1)$. The drawing is x-monotone (the intersection with any vertical line is an interval) and the slopes of boundary-edges are at most 1 in absolute value.*

4 Periodic drawings of toroidal triangulations

Let G be a toroidal triangulation. A *non-contractible cycle* of G is a cycle of edges of G that does not delimit a topological disk. Note that a pair Γ_1, Γ_2 of oriented noncontractible
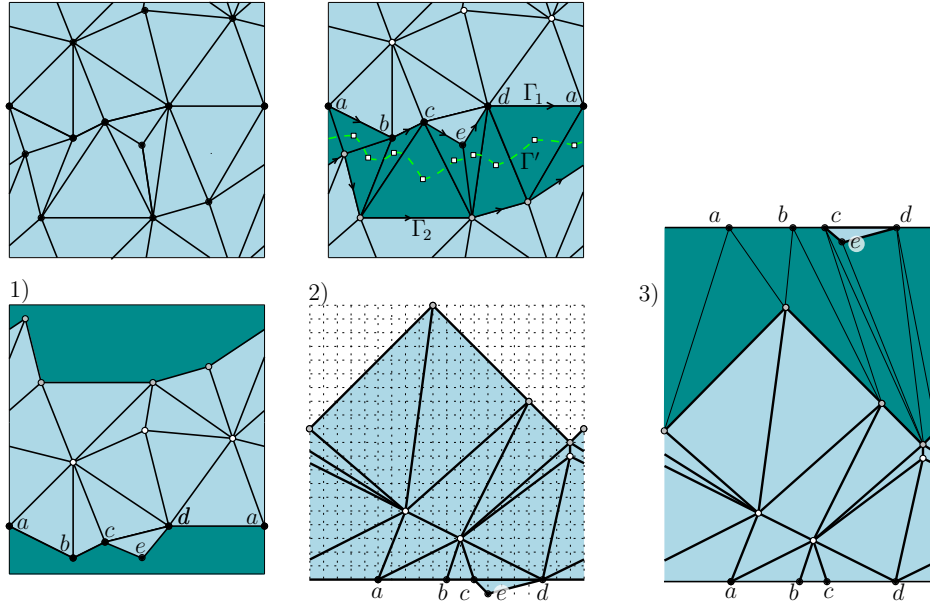
Figure 6: The main steps for drawing of a toroidal triangulation: 1) remove the edges inside a tambourine, (2) draw the obtained cylindric triangulation, 3) insert the edges of the tambourine back into the drawing.

cycles that are parallel (i.e., are homotopic and are oriented in a parallel way) delimits a cylindric triangulation T, which is formed by the faces to the right of Γ_1 and to the left of Γ_2. The pair Γ_1, Γ_2 is called a *tambourine* if the edges dual to those of $T \backslash \{\Gamma_1, \Gamma_2\}$ form a non-contractible cycle Γ' that is homotopic to Γ_1 and Γ_2 (see dark-blue faces and dashed edges in Fig. 6 top-middle). The edges of $T \backslash \{\Gamma_1, \Gamma_2\}$ are said to be *inside* the tambourine. It can be shown (see [2], the master's thesis of the third author Arnaud Labourel, and the next paragraph) that for each non-contractible cycle Γ of G, there exists a tambourine whose two cycles are parallel to Γ. Deleting the edges that are strictly inside the tambourine, one obtains a cylindric triangulation G' with Γ_1 and Γ_2 as the contours of the boundary-faces. Note also that the distance d between Γ_1 and Γ_2 is smaller than the length of a shortest non-contractible cycle not parallel to Γ. We now apply the algorithm of Theorem 1 to G'. If we augment the height h of the drawing to $h' = h + w + 1$, and then wrap the x-periodic grid $\mathbb{Z}/w\mathbb{Z} \times [0..h]$ into a periodic grid $\mathbb{Z}/w\mathbb{Z} \times \mathbb{Z}/h'\mathbb{Z}$, and finally insert the edges inside the tambourine as segments[2], then the slope properties (edges on Γ_1 and Γ_2 have slope at most 1 in absolute value while edges inside the tambourine have slopes greater than 1 in absolute value) ensure that the resulting drawing is crossing-free (see Fig. 6). Observe that we can choose Γ so that the graph-distance between the two boundaries Γ_1 and Γ_2 (in G') is smaller than the length γ of a shortest non-contractible cycle in G; and this choice for Γ can be done without computing a shortest non-contractible cycle. Indeed, let $\{\Gamma_a, \Gamma_b\}$ be a basis of non-contractible cycles (computable in linear time, using for instance a cut-graph). Denoting by $\Gamma_{\min}$ a shortest non-contractible cycle of G, for sure at least one of Γ_a or Γ_b is

---

[2]We insert the edges in the tambourine T in the unique way such that, looking from bottom to top, at least one edge in T goes strictly to the right, and all edges going strictly to the right have x-span at most w; in this way it is easy to check that the x-span of all edges in T is at most w.

9

not parallel to $\Gamma_{\min}$. Hence, for Γ_a or for Γ_b, the distance between the boundary-cycles (after deleting edges of the parallel tambourine) is smaller than $|\Gamma_{\min}|$. In other words if we choose the one cycle (among $\{\Gamma_a, \Gamma_b\}$) that yields the smaller distance between the two boundaries of G', then this distance will be smaller than γ. We obtain:

Theorem 2. *For each toroidal triangulation G, one can compute in linear time a crossing-free straight-line drawing of G on a periodic regular grid $\mathbb{Z}/w\mathbb{Z} \times \mathbb{Z}/h\mathbb{Z}$, where (with n the number of vertices and γ the length of a shortest non-contractible cycle) $w \leq 2n$ and $h \leq 1 + n(2\gamma + 1)$. Since $\gamma \leq \sqrt{2n}$ (as shown in [1]), the grid area is $O(n^{5/2})$.*

Existence of a tambourine. For the sake of completeness we include a proof of existence of a tambourine, which slightly extends the proof given in the master's thesis of Arnaud Labourel. A toroidal map is called *weakly 3-connected* if its periodic representation in the plane is 3-connected. Let G be such a map and let Γ be a non-contractible cycle of G. We are going to show that G has a tambourine parallel to Γ. Let G' be the cylindric map obtained after cutting G along Γ; we take the annular representation of G', calling Γ_1 (resp. Γ_2) the copy of Γ that is the outer (resp. inner) boundary. Let Γ' be the smallest (in terms of the enclosed area) cycle that strictly encloses Γ_2 (i.e., encloses Γ_2 and is vertex-disjoint from Γ_2). Let Γ'' be the largest (in terms of the enclosed area) cycle that is strictly enclosed in Γ' (i.e., is enclosed by Γ' and is vertex-disjoint from Γ'). Note that by minimality Γ' has no chord inside, and by maximality Γ'' has no chord outside. Hence, if we can show that there is no vertex in the area A (strictly) between Γ' and Γ'', then we can conclude that, in G, Γ' and Γ'' form a tambourine parallel to Γ. Assume there is a vertex v in A. Call *vertex of attachment for* Γ' a vertex $w \in \Gamma'$ such that there is a path from v to w visiting only vertices of A before reaching w. Again by minimality of Γ' it is easy to see that there is a unique vertex of attachment v' for Γ'. Similarly (by maximality of Γ'') there is a unique vertex of attachment v'' for Γ''. Then (again by minimality of Γ' and maximality of Γ'') there is a closed curve γ that meets G' only at the vertices v', v'', and such that the interior of γ contains v but does not contain any of the two boundary-faces. Such a curve γ yields a 2-separator in the periodic representation of G, a contradiction.

5   Allowing for non-contractible 1- and 2-cycles

For a cylindric map, a *non-contractible cycle* is a cycle that separates the two boundary-faces (i.e., there is a boundary-face on one side and a boundary-face on the other side of the cycle). Other cycles are called *contractible*. A *weakly simple* cylindric triangulation is a cylindric map whose non-boundary faces are triangles, whose contractible cycles have length at least 3, and such that each vertex has at most one incident non-contractible loop. A *weakly-simple toroidal triangulation* is a toroidal map with triangular faces, with all contractible cycles of length at least 3, and such that any two (non-contractible) loops incident to the same vertex can not be homotopic. These are the necessary and sufficient conditions for the triangulation (whether on the cylinder or on the torus) to have no loops nor multiple edges in the periodic representation; hence these are the conditions under which one can aim at a periodic crossing-free straight-line drawing. For weakly simple cylindric triangulations without loops

Figure 7: Top line: canonical ordering of a weakly simple cylindric triangulation (with some non-contractible 1- and 2-cycles). Middle line: iterative drawing algorithm. Bottom line: dealing with chordal edges.

(nor chords incident to $C(B_1)$), exactly the same shelling procedure and iterative drawing algorithm can be taken as for simple cylindric triangulations. In case there are loops we have to explain how to deal with them (see Fig. 7 for an example). For the shelling procedure, if the current upper boundary C_k is a loop —call v the incident vertex— then one deletes the loop and immediately takes v as the next free vertex (the fact that v is free is due to the fact that there is no other loop at v). In the drawing procedure (how to insert v and its incident loop into the drawing), one first adds v without its loop (by a classical one-step iteration of the drawing algorithm, involving two strip insertions), and then one draws the loop at v as an horizontal segment spreading over the whole width of the current periodic drawing (the loop is added without inserting any vertical strip). Finally one can deal with chords incident to $C(B_1)$ in the same way as for simple cylindric triangulations. About weakly simple toroidal triangulated maps, the procedure is also the same as for simple toroidal triangulations since the above proof of existence of a tambourine holds in that case. And the grid bounds (for the cylindre and for the torus) are the same as for simple triangulations.

References

[1] M. O. Albertson and J. P. Hutchinson. On the independence ratio of a graph. *J. Graph. Theory*, 2:1–8, 1978.

[2] N. Bonichon, C. Gavoille and A. Labourel. Edge partition of toroidal graphs into forests in linear time. In *ICGT*, volume 22, pages 421–425, 2005.

[3] E. Brehm. 3-orientations and Schnyder 3-Trees decompositions. Master's thesis, FUB, 2000.

[4] E. Chambers, D. Eppstein, M. Goodrich, M. Loffler. Drawing Graphs in the Plane with a Prescribed Outer Face and Polynomial Area. *JGAA*, 16(2):243–259, 2012.

[5] L. Castelli-Aleardi, E. Fusy, and T. Lewiner. Schnyder woods for higher genus triangulated surfaces, with applications to encoding. *Discr. & Comp. Geom.*, 42(3):489–516, 2009.

[6] C. Duncan, M. Goodrich, S. Kobourov. Planar drawings of higher-genus graphs. *Journal of Graph Algorithms and Applications*, 15:13–32, 2011.

[7] H. de Fraysseix, J. Pach and R. Pollack. How to draw a planar graph on a grid. *Combinatorica*, 10(1):41–51, 1990.

[8] D. Gonçalves and B. Lévêque. Toroidal maps : Schnyder woods, orthogonal surfaces and straight-line representation. arXiv:1202.0911, 2012.

[9] S. J. Gortler, C. Gotsman and D. Thurston Discrete one-forms on meshes and applications to 3D mesh parameterization. *Computer Aided Geometric Design*, 23 (2): 83–112, 2006.

[10] G. Kant. Drawing planar graphs using the canonical ordering. *Algorithmica*, 16(1):4–32, 1996.

[11] W. Kocay, D. Neilson and R. Szypowski. Drawing graphs on the torus. *Ars Combinatoria*, 59:259–277, 2001.

[12] B. Mohar. Straight-line representations of maps on the torus and other flat surfaces. *Discrete Mathematics*, 15:173–181, 1996.

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

[14] W. Schnyder. Embedding planar graphs on the grid. *SODA*, pp 138-148, 1990.

[15] A. Zitnik. Drawing graphs on surfaces. *SIAM J. Disc. Math*, 7(4):593-597, 1994.