



HAL
open science

Canonical Ordering for Graphs on the Cylinder with Applications to Periodic Straight-line Drawings on the Flat Cylinder and Torus

Luca Castelli Aleardi, Olivier Devillers, Eric Fusy

► **To cite this version:**

Luca Castelli Aleardi, Olivier Devillers, Eric Fusy. Canonical Ordering for Graphs on the Cylinder with Applications to Periodic Straight-line Drawings on the Flat Cylinder and Torus. *Journal of Computational Geometry*, 2018, 9 (1), pp.391 - 429. 10.20382/jocg.v9i1a14 . hal-01959590

HAL Id: hal-01959590

<https://inria.hal.science/hal-01959590>

Submitted on 18 Dec 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

CANONICAL ORDERING FOR GRAPHS ON THE CYLINDER, WITH APPLICATIONS TO PERIODIC STRAIGHT-LINE DRAWINGS ON THE FLAT CYLINDER AND TORUS *

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

ABSTRACT. We extend the notion of canonical ordering, initially developed for planar triangulations and 3-connected planar maps, to cylindric triangulations and more generally to cylindric 3-connected maps. This allows us to extend the incremental straight-line drawing algorithm of de Fraysseix, Pach and Pollack and of Kant from the planar triangulated case and the 3-connected case to this setting. Precisely, for any cylindric essentially 3-connected map G with n vertices, we can obtain in linear time a straight-line drawing of G that is periodic in x -direction, crossing-free, and internally (weakly) convex. The vertices of this drawing lie 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 face-distance between the two boundaries. This also yields an efficient periodic drawing algorithm for graphs on the torus. Precisely, for any essentially 3-connected map G on the torus (i.e., 3-connected in the periodic representation) with n vertices, we can compute in linear time a periodic straight-line drawing of G that is crossing-free and (weakly) convex, on a periodic regular grid $\mathbb{Z}/w\mathbb{Z} \times \mathbb{Z}/h\mathbb{Z}$, with $w \leq 2n$ and $h \leq 1 + 2n(c + 1)$, where c is the face-width of G . Since $c \leq \sqrt{2n}$, the grid area is $O(n^{5/2})$.

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 the canonical ordering [5, 10, 15, 18, 20] are typically incremental, adding vertices one by one while keeping the drawing planar. Algorithms based on Schnyder woods [3, 14, 23] 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 $O(n) \times O(n)$. They can also both be extended [14, 18] to obtain (weakly) convex drawings of

*This work is supported by the ANR grant “EGOS” 12-JS02-002-01 and the ANR grant “GATO” ANR-16-CE40-0009-01. A preliminary version appeared at the 20th International Symposium on Graph Drawing and Network Visualization (GD’12).

[†]LIX, École Polytechnique, France. <http://www.lix.polytechnique.fr/~amturing/>

[‡]INRIA Nancy - Grand est, France. www.loria.fr/~odevil

[§]LIX, École Polytechnique, France. <http://www.lix.polytechnique.fr/~fusy/>

3-connected maps on a grid of size $O(n) \times O(n)$. The problem of obtaining planar drawings of higher genus graphs has been addressed less frequently [9, 13, 17, 19, 21, 22, 24], 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 $O(n^3)$ in the worst case have been described [9, 13]. Such methods unfold the graph in the plane along a *cut-graph*. However, these methods do not yield easily periodic representations: for example, in the case of a torus, the boundary vertices might not be aligned, so that the drawing does not give rise to a periodic drawing. For the torus, a recent article [8] achieves an adaptation of these methods to get alignment of opposite vertices while keeping the size of the periodic grid polynomial, but with the drawback of having a quite large exponent, the guaranteed grid size being $O(n^4) \times O(n^4)$. Another method for drawing toroidal graphs with polynomial grid size is the algorithm of Gonçalves and Lévêque [16], which is an adaptation of Schnyder's drawing principles to the torus¹. It achieves both the periodicity requirement and polynomial grid-size; precisely the size of the periodic regular grid is $O(n^2) \times O(n^2)$ for simple toroidal triangulations (no loops or multiple edges) and is $O(n^4) \times O(n^4)$ for essentially simple (simple in the periodic representation) toroidal triangulations. Gonçalves and Lévêque also extend these ideas to 3-connected toroidal maps (similarly, Schnyder woods for plane triangulations have been extended to plane 3-connected maps [14]), but as opposed to the planar case [14], the periodic drawings of 3-connected toroidal maps they obtain do not necessarily have the desired convexity property.

The main contributions of this article are algorithms to obtain crossing-free (weakly) convex straight-line drawings using a small grid-size of essentially 3-connected maps on the cylinder and then on the torus. The key idea here is to adapt the principles of the iterative algorithms based on canonical orderings [15, 18] to the cylinder², first in the case of triangulations, then to the case of 3-connected maps. As in the planar case, the 3-connected case is technically more involved. Precisely, we first adapt the shelling procedure yielding a canonical ordering and the incremental straight-line drawing algorithm of de Fraysseix, Pach, and Pollack [15] (shortly called FPP algorithm thereafter) to triangulations on the cylinder (Section 3). Then, more generally we can also extend the notion of canonical ordering and convex straight-line drawing of Kant [18] to 3-connected maps on the cylinder (Section 4). Precisely, for any essentially internally 3-connected map G on the cylinder, our algorithm yields in linear time a crossing-free internally convex 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 face-distance³ between the two boundaries of G .

Then (in Section 5), we explain how to obtain periodic drawings on the torus by a reduction to the cylindric case using the notion of a *tambourine* [4]. For any essentially

¹Their method relies on the existence of certain orientations where every vertex has outdegree 3. They have also recently applied these methods to design a bijective encoding scheme for toroidal triangulations [12]. Such technique may induce generalisation to higher genus since it has been proved [1] that a genus g triangulation always admits an orientation where every vertex has for outdegree a non-zero multiple of 3.

²Another notion of canonical ordering for toroidal triangulations has been introduced [7] (this actually works in any genus and yields an efficient encoding procedure) but we will not use it here.

³The face distance is the smallest possible number of faces traversed by any curve connecting the two boundaries and meeting G only at vertices.




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

3-connected toroidal map G with n vertices and *face-width* c (the smallest possible number of points of G met by a non-contractible curve), we can compute in linear time a weakly convex periodic straight-line drawing of G on a regular grid on the flat torus of size $w \times h$, with $w \leq 2n$ and $h \leq 1 + 2n(c + 1)$. Since $c \leq (2n)^{1/2}$ [2], we have $h \leq (2n)^{3/2}$, so that the grid area is $O(n^{5/2})$. This improves upon the previously best known grid size for the torus, of $O(n^2) \times O(n^2)$, by Gonçalves and Lévêque [16], and also always gives a weakly convex drawing, which was not guaranteed in [16].

2 Preliminaries

2.1 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 $S \setminus G$ 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 G with two marked faces B_{inn} and B_{ext} whose boundaries C_{inn} and C_{ext} are simple cycles (C_{inn} and C_{ext} might share vertices and edges). The faces B_{inn} and B_{ext} are respectively called the *inner boundary-face* and the *outer boundary-face* (we will often consider cylindric maps in the annular representation where B_{ext} is the outer face, as shown in Fig. 1 left-part). The other faces are called *internal faces*. Boundary vertices and edges are those belonging to C_{inn} (gray circles in Fig. 1) or C_{ext} (black circles in Fig. 1); the other ones are called *internal vertices* (white circles in Fig. 1) and edges. The notations G , B_{inn} , B_{ext} , C_{inn} , C_{ext} will be used throughout the article.

2.2 Periodic drawings.

Here we consider the problem of drawing a cylindric map on the flat cylinder and drawing a toroidal map on the flat torus. For $w > 0$ and $h > 0$, the *flat cylinder* of width w and




Figure 2: Left: an essentially 3-connected toroidal map G . Right: a weakly convex straight-line drawing of G on a periodic regular grid of size 8×7 .

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 map 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 $\mathbb{Z}/w\mathbb{Z} \times [0, h]$. Similarly, for a toroidal map 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 $\mathbb{Z}/w\mathbb{Z} \times \mathbb{Z}/h\mathbb{Z}$. A periodic straight-line drawing on the flat torus is said to be *weakly convex* if all corners have angle at most π , see Fig. 2 for an example. Note that a drawing of a toroidal triangulation is automatically weakly convex, so that convexity becomes a constraint only when there are faces of degree larger than 3.

3 Periodic drawings of cylindric triangulations

In this section we describe an algorithm to obtain periodic (in x) drawings of cylindric triangulations. We extend these results to 3-connected maps on the cylinder in Section 4. We start with the case of triangulated maps for pedagogical reasons: the different steps are the same as the ones to be used in the more general 3-connected case; but the arguments at each step are simpler in the triangulated case.

3.1 Definitions and statement of the result

A *cylindric triangulation* is a cylindric map T such that all internal faces are triangles. A cylindric map is called *simple* if it has no loops nor multiple edges, and is called *essentially simple* if it has no loops nor multiple edges in the periodic representation. Note that an essentially simple cylindric map might have 2-cycles and 1-cycles (loops), which have to be

non-contractible (they have B_{inn} on one side and B_{ext} on the other side), and two loops cannot be incident to the same vertex. We also define a *chordal edge*, or *chord*, at C_{inn} as an edge not on C_{inn} but with its two ends on C_{inn} . Similarly a *chord* at C_{ext} is an edge not on C_{ext} but with its two ends on C_{ext} . For a cylindric map, the *edge-distance* d between the two boundaries is the length of a shortest possible path starting from a vertex of C_{inn} and ending at a vertex of C_{ext} (possibly $d = 0$). The main result obtained in this section is the following:

Theorem 1. *For each essentially simple 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 edge-distance between the two boundaries— $w \leq 2n$ and $h \leq 2n(d + 1)$. In the drawing, the upper (resp. lower) boundary is a broken line monotone in x , formed by segments of slope at most 1 in absolute value.*

As a first step we will restrict to the case with no chordal edge at C_{inn} :

Proposition 2. *For each essentially simple cylindric triangulation G with no chordal edge at C_{inn} , 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 edge-distance between the two boundaries— $w \leq 2n$ and $h \leq n(2d + 1)$, such that the upper boundary is a broken line monotone in x formed by segments of slope in $\{+1, -1, 0\}$ and the lower boundary is an horizontal line.*

To prove Proposition 2 we will start with the subcase of cylindric simple triangulations. In that case we will introduce a notion of canonical ordering (where it is necessary that C_{inn} has no chordal edge). This makes it possible to design an incremental periodic drawing algorithm, which can be seen as the cylindric counterpart of the FPP algorithm. Then we will extend the canonical ordering and periodic drawing algorithm to cylindric essentially simple triangulations with no loop. We will then explain how to deal with non-contractible loops. This will establish Proposition 2. Finally Proposition 2 and handling chordal edges at C_{inn} will yield Theorem 1 as explained in Section 3.6.

3.2 Canonical ordering for cylindric simple triangulations with no chord at C_{inn} .

We first introduce a notion of canonical ordering (classically studied on plane graphs) for cylindric simple triangulations:

Definition 3. Let G be a cylindric simple triangulation with no chordal edge at C_{inn} . An ordering $\pi = \{v_1, v_2, \dots, v_n\}$ of the vertices of $G \setminus C_{\text{inn}}$ is called a *canonical ordering* if it satisfies:

- For each $k \in [0..n]$ the map G_k induced by C_{inn} and by the vertices $\{v_1, \dots, v_k\}$ is a cylindric triangulation. The outer boundary-face of G_k is denoted C_k .
- For each $k \in [1..n]$, the vertex v_k is on C_k , and its neighbours in G_{k-1} are consecutive on C_{k-1} (see Fig. 3).




Figure 3: From G_{k-1} to G_k in a canonical ordering for a cylindric simple triangulation (annular representation, only the boundaries and next added vertex and added edges are shown).

The notion of canonical ordering makes it possible to construct a cylindric triangulation G incrementally, starting from $G_0 = C_{\text{inn}}$ and adding one vertex at each step. This is similar to canonical orderings for planar triangulations, as introduced by de Fraysseix, Pach and Pollack [15]; the main difference is that, for a planar triangulation, one starts with G_0 being an edge, whereas here one starts with G_0 being a cycle, seen as a cylindric map with no internal face.

The computation of such an ordering is done by a shelling procedure similar to the one considered in the planar case [6, 15]. At each step the graph formed by the remaining vertices is a cylindric triangulation, the inner boundary remains C_{inn} all the way, while the outer boundary (initially C_{ext}) has its contour, denoted by C_k , getting closer to C_{inn} . A vertex $v \in C_k$ is *free* if v is incident to no chord of C_k and if $v \notin C_{\text{inn}}$. The shelling procedure goes as follows, with n the number of vertices in $G \setminus C_{\text{inn}}$:

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 [6]. First, since there is no chord at C_{inn} , then as long as $C_k \neq C_{\text{inn}}$ there is at least one vertex on $C_k \setminus C_{\text{inn}}$. If there is no chord for C_k , then any vertex $v \in C_k \setminus C_{\text{inn}}$ is free. If there is at least one chord $e = \{u, v\}$ for C_k , let P_e be the path connecting u and v on C_k such that the cycle $P_e + e$ does not enclose the inner boundary-face in the annular representation, and let d_e be the length of P_e (note that $d_e \geq 2$). Let $e = \{u, v\}$ be a chord such that d_e is smallest possible. Then any vertex in $P_e \setminus \{u, v\}$ is free. Since there exists a free vertex at each step, the procedure terminates. Let us now justify that the shelling procedure has linear time complexity. Note that an outer vertex $v \in C_k \setminus C_{\text{inn}}$ is free if and only if the number $N(v)$ of neighbours of v belonging to C_k is equal to two. So we just have to mark the vertices when they become outer vertices of the current triangulation and maintain $N(v)$ for these outer vertices. Notice that when v become an outer vertex, it remains an outer vertex and its value $N(v)$ can only decrease. The free vertices (those




Figure 4: Shelling procedure to compute a canonical ordering of a cylindric simple triangulation (at each step the next shelled vertex is surrounded). 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.

for which $N(v) = 2$) are put in a buffer, picking any element of the buffer at each step. All this can be done in time $O(|E|)$, with $|E| \leq 3n$ the number of edges of the cylindric triangulation. To sum up:

Proposition 4. *Any cylindric simple triangulation G with no chordal edge at C_{inn} admits a canonical ordering that can be computed in linear time by a shelling procedure.*

Underlying forest and dual forest. Given a cylindric simple triangulation G with no chordal edge at C_{inn} endowed with a canonical ordering π , we define the *underlying forest* F for π as the oriented subgraph of G where each vertex $v \in C_{\text{ext}}$ has outdegree 0, and where each $v \notin C_{\text{ext}}$ has exactly one outgoing edge, which is connected to the adjacent vertex u of v of largest label in π . The forest F can be computed on the fly during the shelling procedure: when processing an admissible 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 and has its roots on C_{ext} . The *augmented map* \widehat{G} (\widehat{G} has to be seen as a map on the sphere) is obtained from G by adding a vertex w_1 inside B_{inn} , a vertex w_2 inside B_{ext} , and connecting all vertices around B_{inn} to w_1 and all vertices around B_{ext} to w_2 , thus triangulating the interiors of B_{inn} and B_{ext} . We denote by \widehat{F} the forest F plus all edges incident to w_1 and all edges incident to w_2 . The *dual forest* F^* for π is defined 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 C_{inn} , and the edges of the tree can be oriented toward this root-vertex, see Fig. 4 bottom right. Each edge e^* of F^* is in a certain tree-component T^* rooted at a vertex v_0 in front of a certain edge of C_{inn} . Let P be the path from e^* to v_0 in T^* ; P is shortly called *the root-path of e^** .

3.3 Periodic drawing algorithm for cylindric simple triangulations, no chord at C_{inn} .

Given a cylindric simple triangulation G with no chord at C_{inn} , we first compute a canonical ordering of G , and then draw G in an incremental way. We start with a cylinder of width $2|C_{\text{inn}}|$ and height 0 (i.e., a circle of length $2|C_{\text{inn}}|$) and draw the vertices of C_{inn} equally spaced on the circle: space 2 between two consecutive vertices⁴.

Then the strategy for each $k \geq 1$ is to compute the drawing of G_k out of the drawing of G_{k-1} . Note that the set of vertices of C_{k-1} that are neighbours of v_k forms a path γ on C_{k-1} . Traversing γ with the outer face of G_{k-1} to the left, let e_ℓ be the first edge of γ and e_r be the last edge of γ (note that $e_\ell = e_r$ if v_k has only two neighbours on C_{k-1}). Let also a_k be the starting vertex and let b_k be the ending vertex of γ . Two cases can occur.

(1) If, in the drawing of G_{k-1} obtained so far, $\text{slope}(e_\ell) < 1$ and $\text{slope}(e_r) > -1$, then we can directly insert v_k in the drawing. We place 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.

(2) If $\text{slope}(e_\ell) = 1$ or $\text{slope}(e_r) = -1$, then we cannot directly insert v_k as done in Case (1), because the edges e_ℓ and $\{a_k, v_k\}$ would overlap if $\text{slope}(e_\ell) = 1$, or the edges e_r and $\{b_k, v_k\}$ would overlap if $\text{slope}(e_r) = -1$. We first have to perform stretching operations (thereby increasing the cylinder width by 2) to make the slopes of e_ℓ and e_r smaller than 1 in absolute value. 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} . Let P_ℓ (resp. P_r) be the root-path of e_ℓ^* (resp. e_r^*) in F^* . We stretch the cylinder by inserting a vertical strip of length 1 along P_ℓ and another along P_r , see Fig. 5. This comes down to

⁴It is also possible to start with any configuration of points on a circle such that any two consecutive vertices are at even distance




Figure 5: 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.

increasing by 1 the x -span of each edge of G_{k-1} 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,⁵ whose effect is to make the slopes of e_ℓ and e_r strictly smaller than 1 in absolute value, we insert, as in Case (1), 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.

Note that in the two cases (1) and (2), the two rays from a_k and b_k actually intersect at a grid point since the Manhattan distance between any two vertices on C_{k-1} is even. Fig. 6 shows the execution of the algorithm on the example of Fig. 4.

The fact that the drawing remains crossing-free relies on the fact that all edges of the upper boundary have slope at most 1 in absolute value, and on the following inductive property (similar to the one used in [15]), which is easily shown to be maintained at each step k from 1 to n :

PI: for each edge e on C_k (the upper boundary of G_k), let P_e be the root-path of e^* in F^* , 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 crossing-free after 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, where we call w the width and h the

⁵In 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. In addition it also gives rise to a very easy linear-time algorithm (another linear-time version of the FPP algorithm is given in [11]).

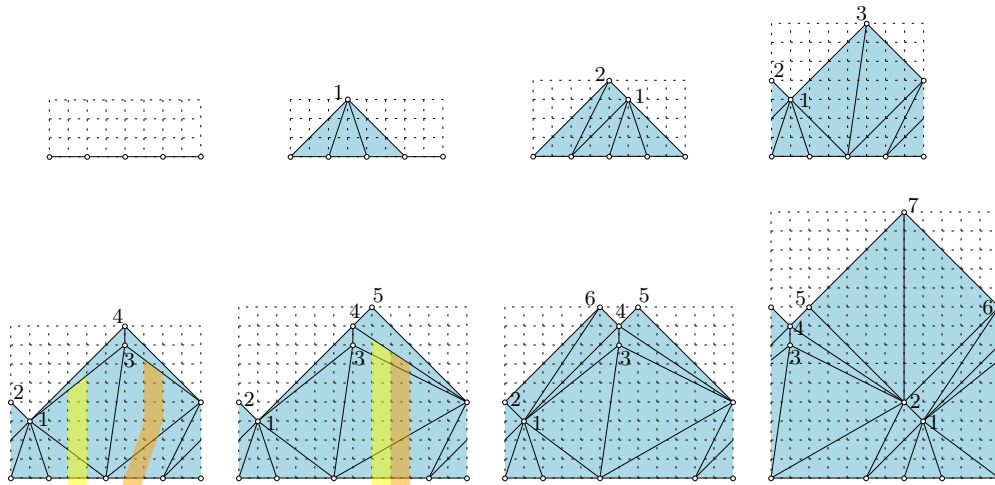


Figure 6: Complete execution of the algorithm computing an x -periodic drawing of a cylindrical simple triangulation with no chordal edge at B_{inn} . The vertices are processed in increasing label (the canonical ordering is the one computed in Fig. 4).

height of the cylinder on which G is drawn. If $|C_{\text{inn}}| = t$ then the initial cylinder is $2t \times 0$; and at each vertex insertion, the grid-width grows by 0 or 2. Hence $w \leq 2n$. In addition, due to the slope conditions (slopes of boundary-edges are at most 1 in absolute value), the 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_{ext} that is closest (at distance d) from C_{inn} , then the ordinate of v is at most $d \cdot (2n)$. And due to the slope conditions, the vertical span of C_{ext} in the drawing is at most $w/2 \leq n$. Hence the grid-height is at most $n(2d + 1)$. The linear-time complexity is shown next.

Linear-time complexity. 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 of the edge 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 -coordinates 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; let s_e be the total further increase undergone by e . Note that for each edge e not in F , if $e \notin C_{\text{ext}}$ there is a certain step k such that $e \in C_{k-1}$ and $e \notin C_k$. Let $w_e \in \{0, 1, 2\}$ be defined as the stretch (increase of x -span) that e undergoes just before adding v_k to the drawing; in case $e \in C_{\text{ext}}$ no such step k exists and we assign $w_e = 0$. We call w_e the *weight* of e (the quantities w_e can be computed in a first pass, together with the quantities r_e). Let P be the root-path of e^* . When stretching e just before adding v_k , all edges dual to edges of P undergo the same stretch, by w_e . In other words, if we denote by T_e^* the subtree of F^* hanging from e^* (including e^*), and denote by W_e the total weight of the dual of the edges in T_e^* , then $s_e = W_e$. Hence the total x -span of each edge $e \in G$ is




Figure 7: The FPP algorithm for simple planar quasi-triangulations is recovered from our algorithm by adding a vertex of degree 2 to complete the inner boundary.

given by $r_e + s_e$, where $s_e = 0$ if $e \in F$ or $e \in C_{\text{ext}}$, and $s_e = W_e$ if $e \notin F$ and $e \notin C_{\text{ext}}$. Since all quantities s_e can easily be computed in linear time from the quantities w_e , starting from the leaves and going up to the roots of F^* , this gives a linear-time algorithm.

To sum up, we have proved Proposition 2 for cylindric simple triangulations with no chord at C_{inn} .

Remark 5. For each edge e of C_{inn} , let r_e be the initial horizontal stretch of e in the drawing procedure (an even number, classically $r_e = 2$ to have a compact drawing). And let t_e be the final horizontal stretch in the drawing procedure. The vectors $R = (r_e)_{e \in C_{\text{inn}}}$ and $T = (t_e)_{e \in C_{\text{inn}}}$ are called *initial-stretch* and *final-stretch* vectors relative to the drawing of G (endowed with a given canonical ordering). Then the vector $S := T - R$ is an invariant, it does not depend on R since $s_e = t_e - r_e$ just depends on the underlying forest and dual forest given by the canonical ordering. Hence, if with R as initial stretch-vector we obtain a drawing with final-stretch vector T , then for any vector T' of the form $T + 2V$ —with V a vector of non-negative integers—we can obtain a drawing with final-stretch vector T' , by taking $R' = R + 2V$ as initial-stretch vector instead of R .

Remark 6. Note that our algorithm can be seen as an extension of the FPP algorithm, which works for simple planar quasi-triangulations, i.e., simple graphs embedded in the plane with triangular inner faces and a polygonal outer face. If we are given a simple planar quasi-triangulation Q , we can turn it into a cylindric simple triangulation G by adding a vertex of degree 2 connected to the two ends of the root-edge. Then the FPP drawing of Q is recovered from the periodic drawing of Q upon deleting the added vertex, see Fig. 7.

3.4 Allowing for non-contractible 2-cycles (no chord at C_{inn}).

Our method based on a canonical ordering and incremental drawing algorithm is easily




Figure 8: When 2-cycles are allowed, the additional case shown here can occur for the transition from G_{k-1} to G_k .

extended to essentially simple cylindric triangulations with no loop but possibly with non-contractible 2-cycles. Let G be such a cylindric map with no chord at C_{inn} . The definition of canonical ordering for G is exactly the same as for simple cylindric triangulations, adding the possibility that G_k is obtained from G_{k-1} as shown in Fig. 8. Such a canonical ordering can be computed by a shelling procedure that extends the one of Section 3.2. A 2-cycle is called *internal* if its two incident vertices are not both on the outer boundary. This time, a vertex on the outer boundary and not on the inner boundary is called *free* if it is not incident to a chord nor incident to an internal 2-cycle.

The shelling procedure consists in choosing a free vertex at each step, and deleting it together with its incident edges, until there just remains the inner boundary. The existence of a free vertex, when the cylindric map is not reduced to a cycle, is proved as follows. First, since there is no chord at C_{inn} , as long as $C_{\text{ext}} \neq C_{\text{inn}}$ there is at least one vertex on $C_{\text{ext}} \setminus C_{\text{inn}}$. If there is no chord nor internal 2-cycle incident to a vertex on C_{ext} , then any vertex on $C_{\text{ext}} \setminus C_{\text{inn}}$ is free. If there is at least one chord at C_{ext} , for each chord e at C_{ext} let P_e be the path on C_{ext} such that $P + e$ does not enclose B_{inn} , and let d_e be the length of P_e . Let $e = \{u, v\}$ be a chord at C_{ext} such that d_e is smallest possible. Then any vertex $v \in P_e \setminus \{u, v\}$ is admissible. If there is no chord but there is at least one internal 2-cycle, consider the largest internal 2-cycle (in terms of containment, noticing that the 2-cycles are nested in the annular representation). Since there is no chord at C_{inn} , at least one outer vertex v is strictly exterior to this 2-cycle, hence is free.

A linear time algorithm is also readily obtained by maintaining, for each outer vertex, how many neighbours on C_{ext} it has and how many internal 2-chords it is incident to. Note that such a canonical ordering also induces an underlying forest F and the dual forest F^* . These can be computed on the fly during the shelling procedure, in the same way as for simple cylindric triangulations. Finally, the incremental drawing algorithm (and linear complexity using the dual forest) works exactly in the same way as for cylindric simple triangulations. An example is shown in Fig. 9. The grid bounds are also the same as for simple triangulations (the arguments to obtain the bounds in the simple case did not use the fact that there are no 2-cycles). So this gives Proposition 2 for essentially simple triangulations with no loops. Finally, note that Remark 5 still holds here (the arguments are the same).




Figure 9: Left-side: the shelling procedure for an essentially simple loopless cylindrical triangulation G with no chord at C_{inn} ; the last drawing shows the underlying forest and dual forest. Right-side: the incremental drawing algorithm.




Figure 10: Drawing algorithm when the inner boundary is the unique loop.

3.5 Allowing for non-contractible loops (no chord at C_{inn}).

We finally explain how to deal with non-contractible loops. Our strategy is not to extend the notion of canonical ordering but simply to decompose (at the loops, which are nested) such a cylindrical map into a “tower” of components, where the only loops in each component are at the boundary-faces. Let G be an essentially simple cylindrical triangulation with n vertices and at least one loop. There are a few cases to consider:

(a) C_{inn} is the unique loop of G . In that case, the algorithm of Section 3.2 (canonical ordering, shelling procedure, and incremental drawing procedure) works in the same way, see Fig. 10 for an example. Let $2m$ be the width of the drawing. By the arguments of Remark 5, for any $m' \geq m$, G has a periodic drawing of width $2m'$ and height at most $m'(2d + 1)$, with d the edge-distance between the two boundaries.

(b) C_{ext} is a loop, C_{inn} is possibly a loop, and there are no other loops. Assume C_{ext} is a loop and G is not reduced to that loop (i.e., $C_{\text{ext}} \neq C_{\text{inn}}$). Let u be the vertex incident to the loop (note that u is not on C_{inn} since there is no chord at C_{inn}), and let c be the innermost 2-cycle incident to u . Cutting along c (see Fig. 11), we obtain two components: a

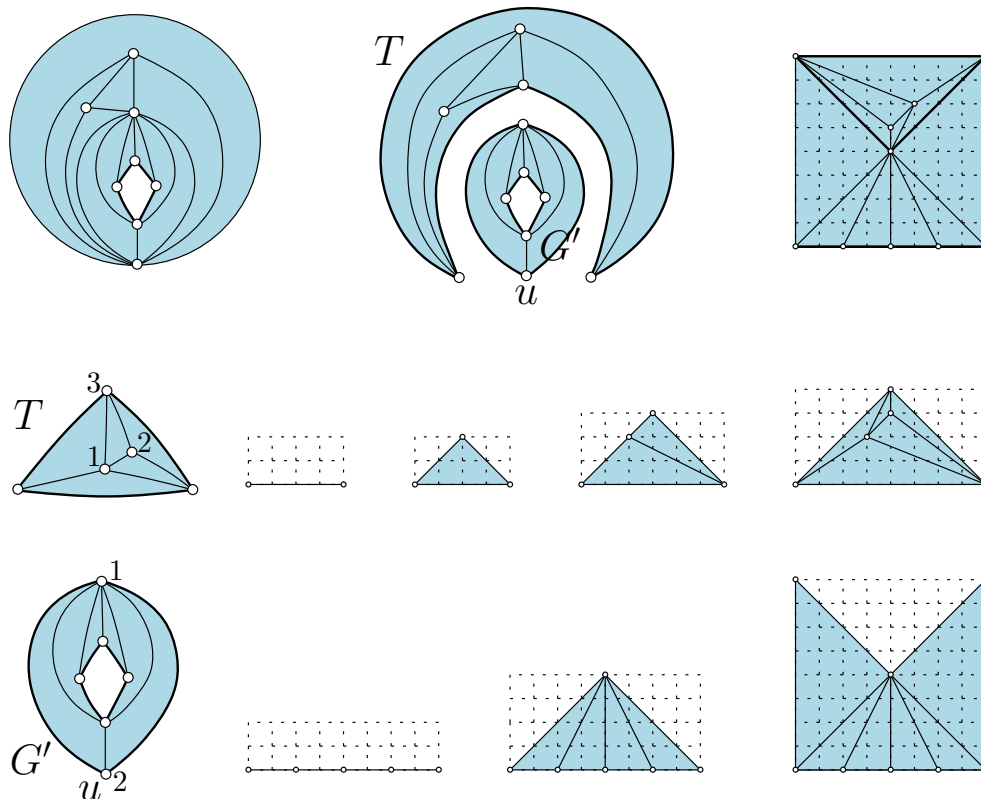


Figure 11: Drawing algorithm when the outer face contour is a loop (and there is no other loop except possibly at C_{inn}): the map is split along the innermost 2-cycle (the initial x -span of T is taken to be 4 instead of 2 so that the drawings of Q and G' fit together).

planar triangulation T and a cylindric essentially simple triangulation G' such that: G' has no loop except possibly at C_{inn} , G' has outer degree 2 and u is a free vertex for G' . Hence there is a canonical ordering for G' such that u is the first shelled vertex. Take a periodic drawing of G' for this canonical ordering and take an FPP drawing of T . The widths of the respective drawings are even numbers, denoted $2n_1$ and $2n_2$. Let $m = \max(n_1, n_2)$. By the arguments of Remark 5 it is possible to redraw the graph that has the smaller grid-width so that it gets width $2m$, after which both drawings are of width $2m$. Then the drawing of T (taken upside down) fits into the upper boundary of G' yielding a periodic drawing of G of width $2m$, see Fig. 11. The height of the drawing is at most $2dm \leq 2dn$, where d is the edge-distance between the two boundaries (indeed, in the usual bound $h \leq (2d + 1)m$, the $+1$ in the parenthesis is due to the vertical extension of the upper boundary, which is 0 here).

(c) *General case (with no chord at C_{inn}).* We can assume there is at least one loop (the case with no loop has been covered in the last section). Let ℓ_1, \dots, ℓ_r be the sequence of nested loops of G , with ℓ_1 the innermost loop and ℓ_r the outermost loop; and let $G^{(0)}, \dots, G^{(r)}$ be the $r + 1$ components that result from cutting successively along all




Figure 12: Drawing an essentially simple triangulation G (no chord at C_{inn}); G is first decomposed at its loops; each component is drawn so that the component-drawings have the same width, and can be stacked up to obtain a periodic drawing of G .

these loops. For $i \in [0..r]$ let d_i be the edge-distance between the two boundaries in $G^{(i)}$; and let d be the edge-distance between the two boundaries of G . Note that $d = \sum_i d_i$. Each component G_i has loops only at the boundary-face contours, hence has a periodic drawing (according to cases (a) and (b)) such that boundaries that are loops are drawn as horizontal lines. Let $2n_1, \dots, 2n_r$ be the widths of the drawings of $G^{(1)}, \dots, G^{(r)}$ thus obtained. Let $m = \max(n_1, \dots, n_r)$. By the arguments of Remark 5, each of the graphs $G^{(i)}$ can be redrawn so as to have width $2m$. Stacking up all these drawings we obtain a periodic drawing of G of width $2m$, see Fig. 12. Regarding the grid size, the width is $2m$, with clearly $m \leq n$, and the height of the drawing of G_i is at most $2md_i$ for $i \in [0..r-1]$ and at most $m(2d_r + 1)$ for $i = r$. Hence the total height is at most $m(2d + 1) \leq n(2d + 1)$. This establishes Proposition 2.

3.6 Allowing for chords at C_{inn} .

We finally explain how to draw a cylindric essentially simple triangulation when allowing for chords incident to C_{inn} . It is good to view B_{ext} as the top boundary-face and B_{inn} as the bottom boundary face, and imagine a standing cylinder. For each chord e at the cycle C_{inn} , the *component under e* , denoted Q_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_{inn} is *maximal* if the component Q_e under e is not strictly included in the component under another chord at C_{inn} . The *FPP-size* $|e|$ of such an edge e is defined as the width of the FPP drawing of Q_e . If we delete the component under each maximal chordal edge (i.e., delete everything from the component except for the chordal

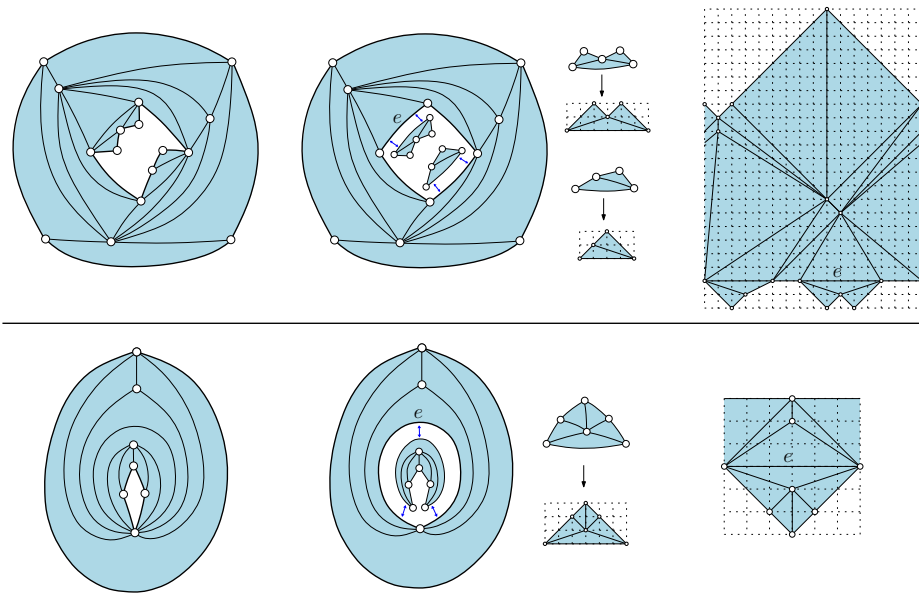


Figure 13: Drawing a cylindric triangulation with chords at C_{inn} (the top-line example is simple, the bottom-line example is essentially simple with loops and 2-cycles in the annular representation). In the top-line example, to make enough space to place the component under e , one takes 6 (instead of 2) as the initial x -span of e . In the bottom-line example, to make enough space to place the component under e , one takes 4 (instead of 2) as the initial x -span of e .

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 2. Let w_e be the width of each edge e of C'_0 in this drawing. According to Remark 5, we can redraw G' such that each edge $e \in C'_0$ that is chordal in G has width $\ell(e)$, with $\ell(e)$ defined as the smallest integer that is at least $\max(w_e, |e|)$ and such that $\ell(e) - w_e$ is even (note that $\ell(e) \leq \max(w_e, |e| + 1)$).

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 in ± 1 . We shift the left-extremity of e to the left 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. 13. The overall drawing of G is clearly planar.

We now give bounds on the grid-size of the overall drawing. Let S be the sum of the FPP-sizes over all maximal chords e at C_{inn} , and let n' be the number of vertices of G' . Clearly the width w of the drawing of G satisfies $w \leq 2n' + \sum_{e \in C'_0} \ell(e) - w_e \leq 2n' + S$. For each maximal chord e at C_{inn} let $n_e + 2$ be the number of vertices of the component Q_e under e . Let N be the sum of the quantities n_e over all maximal chords at C_{inn} , so that $n = n' + N$. Since the FPP drawing of a quasi-triangulation with $p \geq 3$ vertices has width at most $2p - 4$, we have $|e| \leq 2n_e$ for each maximal chord at C_{inn} . Hence $S \leq 2N$, so that $w \leq 2n$. Regarding the height of the drawing, by the same arguments as in Section 3.3, the height of the drawing of G' is at most $n(2d + 1)$, with d the edge-distance between the two




Figure 14: Left: a cylindric map G with a (contractible) 1-separating curve and a (non-contractible and nondegenerate) 2-separating curve (in bold line). Right: a cylindric map with a contractible degenerate 2-separating curve (the enclosed area appears in yellow).

boundaries. After adding the components under the chords at C_{inn} , the lower boundary is not horizontal anymore, but since it is made of segments of slope at most 1 in absolute value, it has vertical extension at most $w/2 \leq n$. Hence the overall height of the drawing is at most $n(2d + 1) + n = 2n(d + 1)$. This finally yields the result pursued in this section, Theorem 1.

4 Periodic drawing of 3-connected maps on the cylinder

We now extend the results obtained in Section 3 to the more general 3-connected case. The approach is completely parallel to the one used in Section 3, but the arguments at each step are more technical.

4.1 Definitions and statement of the result

Let G be a cylindric map, embedded in the plane using the annular representation (with B_{ext} as the outer face and B_{inn} as the marked inner face). We define a *1-separating curve* as a closed curve γ not meeting any edge and intersecting G in exactly one vertex, the unique visited face being an internal face, and such that the area enclosed by γ contains at least one edge. Such a curve γ is called *non-contractible* if the area enclosed by γ entirely contains B_{inn} . We now define a *2-separating curve* as a closed curve γ intersecting G in two vertices (not meeting any edge), visiting exactly two faces, and such that the area enclosed by γ strictly contains at least one vertex; again γ is said to be *non-contractible* if it entirely encloses B_{inn} . We also allow for the degenerate situation of a contractible 2-separating curve where the two incident vertices are equal (it is still required that the enclosed area strictly contains at least one vertex), see Fig. 14 for an example. Two 1-separating curves (resp. two 2-separating curves) are considered as equivalent if they are isotopic. It is convenient—and will always be assumed from now on—to discard non-contractible curves passing by the outer face.

We consider in this section certain cylindric maps G with some marked vertices on C_{ext} and some marked vertices on C_{inn} which are called *active* (if a vertex is on $C_{\text{inn}} \cap C_{\text{ext}}$

it might be active for C_{inn} or C_{ext} , or both, or none). As we will see, when all active vertices are for C_{ext} , these vertices are to be the ones that are allowed to be selected during the shelling procedures to compute a canonical ordering in the more general 3-connected case presented here. The terminology of active vertices will also be useful when presenting the drawing algorithm for toroidal 3-connected maps via a reduction to the cylindrical case (we will delete certain edges of the toroidal map to make it a cylindrical map, and declare as active the vertices of the cylindrical map incident to at least one deleted edge).

A cylindrical map G with some active vertices is called *internally 3-connected* if it has no 1-separating curve, and any 2-separating curve is contractible (and nondegenerate) and strictly encloses at least one active vertex (see Fig. 17(a) for an example). And G is called *essentially internally 3-connected* if there is no contractible 1-separating curve, and any contractible 2-separating curve (possibly degenerate) strictly encloses at least one active vertex. Being essentially internally 3-connected can also be conveniently characterized in the x -periodic representation \hat{G} of G : a 1-separating (resp. 2-separating) curve in \hat{G} is a simple closed curve not meeting any edge of \hat{G} , meeting \hat{G} at exactly 1 vertex (resp. 2 vertices), and whose interior strictly contains at least one vertex. Then G is essentially internally 3-connected if and only if, in \hat{G} , there is no 1-separating curve and any 2-separating curve strictly encloses at least one active vertex.

For a cylindrical map G with active vertices, a periodic straight-line drawing of G is called *weakly convex* if all corners have angle at most π , except possibly for corners of B_{inn} (resp. B_{ext}) at an active vertex for C_{inn} (resp. for C_{ext}), whose angle in the drawing can be larger than π . For a cylindrical map G , the *face-distance* between the two boundaries is the minimal possible integer q such that there is a curve in G starting from a vertex of C_{inn} , ending at a vertex of C_{ext} , not meeting any edge, and passing by q (internal) faces of G . The main result obtained in this section is the following:

Theorem 7. *For each essentially internally 3-connected cylindrical map G with at least one active vertex, one can compute in linear time a periodic weakly convex 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 face-distance between the two boundaries— $w \leq 2n$ and $h \leq 2n(d + 1)$. In the drawing, the upper (resp. lower) boundary is a broken line monotone in x formed by segments of slope at most 1 in absolute value.*

Having a convexity condition (angle at most π) at all corners except possibly at the active vertices is important in view of our drawing algorithm for toroidal 3-connected maps. Indeed, we will obtain a convex periodic toroidal drawing by adding edges to a periodic cylindrical drawing, and the corners not at active vertices will be the ones kept unchanged (not receiving any additional edge), it is thus necessary that these have angle at most π already in the cylindrical drawing.

As a first step to show Theorem 7, we will prove the result when there is no active vertex on C_{inn} :

Proposition 8. *For each essentially internally 3-connected cylindrical map G where C_{ext} has at least one active vertex and C_{inn} has no active vertex, one can compute in linear time a periodic weakly convex 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 face-distance between the two boundaries— $w \leq 2n$ and $h \leq n(2d + 1)$. In the drawing, the upper boundary is a broken line, monotone in x , formed by segments of slope at most 1 in absolute value, and the lower boundary is an horizontal line.

Note that Theorem 7 and Proposition 8 are respectively extensions of Theorem 1 and Proposition 2. Indeed, for an essentially simple cylindric triangulation G , making all boundary-vertices of G active yields an essentially internally 3-connected cylindric map, and if G has no chord at C_{inn} , then making all vertices of C_{ext} active yields an essentially internally 3-connected cylindric map with no active vertex at C_{inn} . In addition, for an essentially simple cylindric triangulation, the face-distance between the two boundaries coincides with the edge-distance between the two boundaries. To prove Proposition 8 (the strategy is parallel to the one we have followed to prove Proposition 2 for cylindric triangulations) we will start with the subcase where G is internally 3-connected. In that case we will introduce a notion of canonical ordering, which extends both the canonical ordering for cylindric triangulations introduced in Section 3.2, and the canonical ordering for internally 3-connected plane graphs introduced by Kant [18]. This makes it possible to design an incremental periodic drawing algorithm, which is the cylindric counterpart of the algorithm introduced by Kant [18] in the planar case (which itself extends the FPP algorithm to 3-connected planar maps). Then we will extend the canonical ordering and drawing algorithm to the subcase where there is no 1-separating curve; after which we will explain how to deal with 1-separating curves. This will establish Proposition 8. We will then explain how to deal with active vertices at C_{inn} . This will yield Theorem 7.

4.2 Restatement of the definitions in terms of the corner-map

We provide here a classical reformulation of the 3-connectedness conditions in terms of the so-called corner-map, which provides a more combinatorial way of viewing 2-separating curves. Given a cylindric map G (whose vertices are considered as white), in its annular representation, the *corner-map* S of G is obtained by inserting a black vertex v_f in each internal face f of G and connecting v_f to all the corners around f ; S is the graph made of black and white vertices and of the (newly added) edges between black and white vertices. The *completed map* \widehat{G} of G is defined as G superimposed with S (see Fig. 15 for an example). A *separating 4-cycle* in S is a 4-cycle containing at least one vertex in its interior; it is called *non-contractible* if it completely encloses B_{inn} and contractible otherwise. Note that a separating 4-cycle exactly corresponds to a separating 2-curve of G passing by two internal faces. We define (in \widehat{G}) a *2-chord* γ for C_{ext} as a path e_1, e_2 of length two in S starting from a vertex u of C_{ext} and ending at a vertex $v \neq u$ of C_{ext} . We denote by P_γ the path from u to v on C_{ext} such that the cycle $C_\gamma := P_\gamma \cup \gamma$ does not contain B_{inn} ; C_γ is called the *cycle enclosed by the 2-chord*. The 2-chord γ is called *separating* if P_γ is of length larger than 1, i.e., has at least one non-extremal vertex. The non-extremal vertices of P_γ are said to be *enclosed* by γ . A 2-chord or separating 2-chord at C_{inn} is defined analogously. Note that a separating 2-chord at C_{ext} (resp. at C_{inn}) exactly corresponds to a 2-separating curve of G passing by B_{ext} and by an internal face (resp. passing by B_{inn} and by an internal face). If C_{ext} meets C_{inn} , we define an intersection-vertex as a vertex




Figure 15: Left: a cylindric map G .

Right: the associated corner-map S , with, in bolder line, a non-contractible separating 4-cycle, which corresponds to a 2-separating curve of G passing by two internal faces, and a separating 2-chord.

Middle: the completion map \widehat{G} , which is obtained by superimposing G and S .

of $C_{\text{ext}} \cap C_{\text{inn}}$; G can be seen as a cyclic sequence of elementary blocks which are attached along the intersection-vertices. Such elementary blocks are called the *portions* of G , and the two intersection-vertices delimiting the portion are called the *extremal vertices* of the portion. A portion is said to be *non-trivial* if it is not reduced to an edge. Note that a non-trivial portion delimited by two distinct intersection-vertices exactly corresponds to a 2-separating curve whose two incident faces are B_{inn} and B_{ext} . For instance, in Figure 17, the cylindric map in (a) and in (b) has only one portion; in (c) it has two portions (one being trivial), and in (d), (e) and (f) it has 3 portions (one being trivial).

Given this discussion it is clear that G is internally 3-connected iff S has no 2-cycle nor separating 4-cycle, any separating 2-chord γ at C_{ext} (resp. at C_{inn}) encloses at least one active vertex, and any non-trivial portion delimited by two distinct intersection-vertices has at least one non-extremal vertex that is active.

We now reformulate the condition of being essentially internally 3-connected in terms of the corner-map⁶. A cylindric map G with active vertices, with S its corner-map, is essentially internally 3-connected iff S has no contractible 2-cycle nor contractible 4-cycle (including the degenerate situation of a contractible 4-cycle with two vertices repeated), every separating 2-chord (u, e_1, e_2, v) at C_{ext} (including the degenerate situation where $u = v$, in which case P_γ is taken to be the whole contour C_{ext}) encloses at least one active vertex, every separating 2-chord (u, e_1, e_2, v) at C_{inn} (including the degenerate situation where $u = v$, in which case P_γ is taken to be the whole contour C_{inn}) encloses at least one active vertex, and every non-trivial portion has at least one non-extremal vertex that is active.

⁶This reformulation relies on arguments similar to the ones used in [22] (see Lemma 2.1) for dealing with the *essentially 2-connectedness* of maps on surfaces having non positive Euler characteristic.




Figure 16: The two possible transitions from G_{k-1} to G_k in a canonical ordering of an internally 3-connected cylindric map.

4.3 Canonical ordering

We first introduce a notion of canonical ordering for internally 3-connected cylindric maps. Before that, we need a bit of terminology. Given a pair (u, w) of outer active vertices, the *outer path* for (u, w) is the path, denoted $\gamma(u, w)$, from u to w on the outer face contour, having the outer face on its left. The pair (u, w) is called *consecutive* if there is no active vertex in $\gamma \setminus \{u, w\}$.

Definition 9. Let G be an internally 3-connected cylindric map with no active vertex for C_{inn} and with at least one active vertex for C_{ext} . A *canonical ordering* for G is a growing sequence G_0, G_1, \dots, G_p of internally 3-connected cylindric maps such that:

- Initially, $G_0 = C_{\text{inn}}$ with all vertices active, at the end $G_p = G$. All along, the inner boundary-face of G_k is B_{inn} . The outer boundary of G_k is denoted C_k .
- For $k \in [0..p]$, the set of active vertices on C_k is the union of the set of those with at least one neighbour in $G \setminus G_k$ and the set of those that are on C_{ext} and are active for G .⁷
- For $k \in [1..p]$, G_k is obtained from G_{k-1} either by:
 - choosing a non-consecutive pair (u, w) of distinct active vertices on C_{k-1} and connecting all active vertices on $\gamma(u, w)$ to a newly added vertex in the outer face (see Fig. 16(a));
 - or choosing a consecutive pair (u, w) of distinct active vertices on C_{k-1} and adding a path of at least two edges in the outer face connecting these two vertices (see Fig. 16(b)).
- Additionally, for $k \in [1..p]$, the vertices of $G_k \setminus G_{k-1}$ must be active (for G_k).

Given a canonical ordering of G , the *rank* of a vertex $v \notin C_{\text{inn}}$ is the smallest k such that $v \in G_k$. We prove the existence of such a canonical ordering by induction on the

⁷With these conditions it is easy to check that, since G has at least one active vertex for C_{ext} , then G_k must have at least one active vertex for C_k .




Figure 17: Shelling procedure to compute a canonical ordering of an internally 3-connected cylindric map (at each step the next deleted vertices are surrounded, the ranks of successively shelled vertices are also indicated). The last drawing shows the cylindric map with the underlying forest and dual forest.

number of internal faces, which also yields a shelling procedure similar to the one described in the triangulated case. A first remark is that if G is not reduced to a cycle, then G must have an active vertex for C_{ext} that is not on $C_{\text{ext}} \setminus C_{\text{inn}}$; indeed, if all the active vertices of G were on $C_{\text{ext}} \cap C_{\text{inn}}$, it would easily yield a 2-separating curve not enclosing any active vertex, a contradiction. To describe the shelling procedure we need a bit of terminology. An internal face f is called *separating* if there is a separating 2-chord (always for C_{ext} here) passing by f . A vertex on the outer face is called *admissible* if it is active, not on C_{inn} , and not incident to any separating face. Note that, by deleting an admissible vertex and declaring all its neighbours as active, one gets a cylindric map G' such that G is obtained from G' by applying the operation of Fig. 16(a). We now prove that either there is an admissible vertex or it is possible to get G from a cylindric map G' with one less internal face by applying the operation of Fig. 16(b). For a given separating face f , the *maximal separating 2-chord* for f is the separating 2-chord c passing by f and such that the cycle enclosed by c is maximal (for the containment relation on the enclosed area). Consider the set E of maximal separating 2-chords associated to all separating internal faces of G . If E is not empty (i.e., there is at least one separating internal face), let c be a separating 2-chord

in E that is minimal (for the containment relation) and let f be the internal face to which c belongs. Two cases can occur. If the cycle enclosed by c contains no other internal face than f , then the subgraph G_c of G in the cycle enclosed by c (including the boundary of the cycle) is a path P that has at least two edges. Then by deleting the edges and internal vertices of P , and declaring as active the two extremities of P , we obtain a cylindric map G' such that G is obtained from G' by applying the operation of Fig. 16(b). Otherwise G_c contains at least one internal face f' . Among the separating 2-chords passing by f and whose enclosed cycle contains f' , let c' be the minimal one (for the containment relation). Let P be the path on C_{ext} inside c' , let u, w be the extremities of P . By minimality of c' , all vertices of $P \setminus \{u, w\}$ are not incident to f . In addition, at least one of these vertices has to be active (since G is internally 3-connected). Let v be such a vertex. By minimality of c , all internal faces (except for f) in the cycle enclosed by c are non-separating, hence v is admissible since it is not incident to f . If E is empty, then there is no separating face. As we have seen, there is at least one active vertex on $C_{\text{ext}} \setminus C_{\text{inn}}$. Since there is no separating face, this vertex is admissible. To sum up, in all cases, it is possible to obtain G from a smaller cylindric map G' by applying the operation of Fig. 16(a) or Fig. 16(b). It is also readily checked that the outer face of G' is a simple cycle and that G' satisfies the conditions of Definition 9. So we can continue inductively (starting from G') until there just remains the cycle C_{inn} . This yields a shelling procedure to compute a canonical ordering satisfying Definition 9. An example is shown in Fig. 17.

Let us now justify that the shelling procedure has linear time complexity. At each step, for each internal face f whose contour meets the current outer boundary C_k , we denote by $V(f)$ the number of outer vertices incident to f and by $E(f)$ the number of outer edges on the contour of f . Similarly as in Kant [18] for the planar case, we note that, at each step, an internal face is non-separating if and only if $E(f) \leq 1$ and $V(f) = E(f) + 1$; otherwise in the case where $E(f) \geq 2$ and $V(f) = E(f) + 1$, then the face f can be shelled, corresponding to the reverse of the transition in Fig. 16(b). At each step, for a current active outer vertex v , let $N(v)$ be the number of separating faces incident to v . Note that v is admissible if and only if $N(v) = 0$ and $v \notin C_{\text{inn}}$, in which case v can be shelled, corresponding to the reverse of the transition from Fig. 16(a). By maintaining the quantities $E(f)$ and $V(f)$ for all internal faces touching the outer face, one can also maintain the quantities $N(v)$ for all outer vertices (as well as their status: active or non-active). The shelling is done using two buffers S, S' : in S are stored the current admissible vertices, in S' are stored the current faces f for which $E(f) \geq 2$ and $V(f) = E(f) + 1$. At each step, at least one of the two buffers is non-empty (as we have proved above); so it suffices to shell the top-vertex from S or the top-face from S' , and then update the quantities $E(f), V(f), N(v)$. Maintaining all these informations over the shelling procedure takes time $O(|E|)$ (with $|E|$ the number of edges of the cylindric map), which is also $O(n)$ since $|E| \leq 3n$. More details on implementing such a procedure in linear time are given by Kant [18] for the shelling procedure in the planar case.

Underlying forest and dual forest. Given an internally 3-connected cylindric map G with no chord at C_{inn} , and endowed with a canonical ordering π , we define the *underlying forest* F for π as the oriented subgraph of G where each active vertex $v \in C_{\text{ext}}$ has outdegree

0, and where each other vertex has exactly one outgoing edge, which is connected to the adjacent vertex u of v of largest rank in π . Since the edges are oriented in increasing labels, F is a spanning (oriented) forest; each component of the forest is rooted at each of the active vertices on C_{ext} . The *augmented map* \widehat{G} (seen as a map on the sphere) is defined as the map obtained from G by adding a vertex w_1 inside B_{inn} , a vertex w_2 inside B_{ext} , and connecting all vertices around B_{inn} to w_1 and all active vertices around B_{ext} to w_2 . We define \widehat{F} as F plus all edges (of \widehat{G}) incident to w_1 and all edges incident to w_2 . And we 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} . Each of the trees (connected components) of F^* is rooted at a vertex “in front of” each edge of B_{inn} , and the edges of the tree can be oriented toward this root-vertex (see Fig. 17 bottom right). Similarly as in the triangulated case, for $e^* \in F^*$, we call *root-path* of e^* the path from e^* to the root of the tree-component of F^* to which e^* belongs.

4.4 Periodic drawing algorithm for internally 3-connected cylindric maps with no active vertex for C_{inn} .

Given an internally 3-connected cylindric map G with no active vertex for C_{inn} , we first compute a canonical ordering of G , and then draw G in an incremental way, similarly as in the triangulated case. A first useful remark is that, at any step k , if we look at a path P of edges on C_k connecting two consecutive active vertices for C_k (i.e., P starts at an active vertex for C_k , ends at an active vertex for C_k , and all internal vertices of P are non-active), then P contains exactly one edge not in the underlying forest; this edge is called the *bottom-edge* of P . We start with a cylinder of width $2|C_{\text{inn}}|$ and height 0 (i.e., a circle of length $2|C_{\text{inn}}|$) and draw the vertices of C_{inn} equally spaced on the circle: space 2 between two consecutive vertices (as in the triangulated case, it is possible to start with any configuration of points on a circle such that any two consecutive vertices are at even distance). Then the strategy for each $k \geq 1$ is to compute the drawing of G_k out of the drawing of G_{k-1} . The difference with the triangulated case is that there are now two cases (those shown in Fig. 16).

(a) *Addition of one vertex and several internal faces.* Consider the case of Fig. 16(a); let v be the new added vertex, i.e., the unique vertex of G_k not in G_{k-1} . Let u_1, \dots, u_s (with $s \geq 2$) be the neighbours of v on C_{k-1} , such that the path γ of C_{k-1} from u_1 to u_s has the outer face on its left. Let e_1 be the first edge on γ and let e_2 be the last edge on γ . There are two subcases. (1) If in the drawing of G_{k-1} obtained so far, $\text{slope}(e_1) < 1$ and $\text{slope}(e_2) > -1$, then, as in the triangulated case, we place v at the intersection of the ray of slope 1 starting from u_1 and the ray of slope -1 starting from u_s ; and we draw all the edges from v to u_1, \dots, u_s as segments. (2) If $\text{slope}(e_1) = 1$ or $\text{slope}(e_2) = -1$, let e_ℓ be the bottom-edge of the part of γ between u_1 and u_2 and let e_r be the bottom-edge of the part of γ between u_{s-1} and u_s . Let P_ℓ (resp. P_r) be the root-path of e_ℓ^* (resp. e_r^*) in F^* . Similarly as in the triangulated case (see Fig. 5), we stretch the cylinder by inserting a vertical strip of length 1 along P_ℓ and another one along P_r . After this, we insert (as in subcase (1)) the vertex v at the intersection of the ray of slope 1 starting from u_1 and the

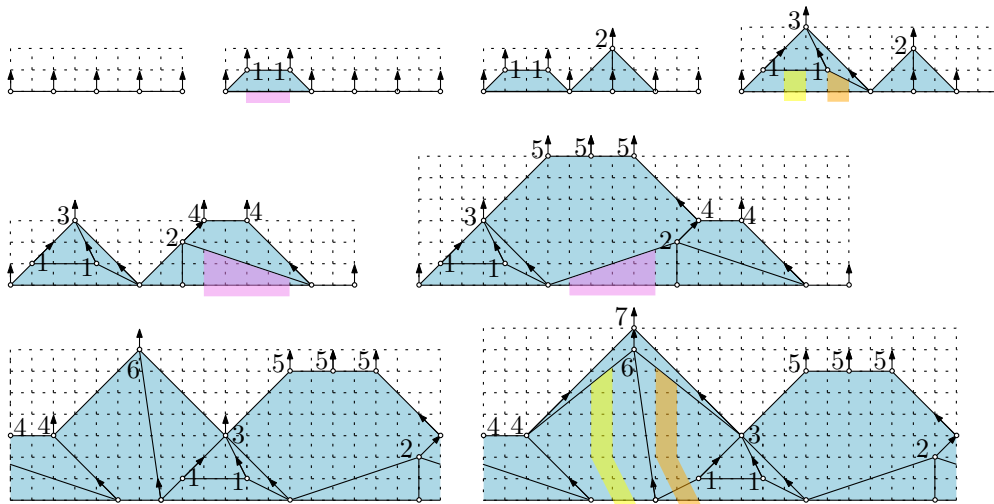


Figure 18: Complete execution of the algorithm computing an x -periodic drawing of an internally 3-connected cylindric map (no active vertex for C_{inn}). The steps follow the canonical ordering (in increasing order) computed in Fig. 17.

ray of slope -1 starting from u_s , and we connect v to all vertices u_1, \dots, u_s by segments. The two rays from u_1 and u_s actually intersect at a grid point since the Manhattan distance between any two vertices on C_{k-1} is even.

(b) *Addition of one internal face.* Consider the case of Fig. 16(b), where we denote u, v_1, \dots, v_s, w ($s \geq 1$) the vertices of the new added chain, such that the path γ on C_{k-1} from u to w has the outer face on its left. Let e_1 be the first edge on γ and e_2 the last edge on γ (note that e_1 might be equal to e_2). Let e be the bottom-edge of γ , and let P be the root-path of e^* in F^* . There are two subcases. (1) If, in the drawing of G_{k-1} obtained so far, $\text{slope}(e_1) < 1$ and $\text{slope}(e_2) > -1$, we insert a vertical strip of width $2s - 2$ along P , increasing by $2s - 2$ the x -span of each edge of G_{k-1} dual to an edge in P . (2) If $\text{slope}(e_1) = 1$ or $\text{slope}(e_2) = -1$, we insert a vertical strip of width $2s$ along P , increasing by $2s$ the x -span of each edge of G_{k-1} dual to an edge in P . Then we insert the vertices v_1, \dots, v_s into the drawing as follows. Let R_u be the ray of slope $+1$ from u and let R_w be the ray of slope -1 from w . Let q be the intersecting point of the two rays, denote by $y(q)$ its y -coordinate; let S be the horizontal segment connecting R_u to R_w at ordinate $y(q) - s + 1$. Note that S has length $2s - 2$. Then we insert v_1, \dots, v_s equally spaced (space 2 between two consecutive vertices) on S , with v_1 at the left extremity and v_s at the right extremity of S . And we draw the edges of the chain u, v_1, \dots, v_s, w as segments.

Fig. 18 shows the execution of the algorithm on the example of Fig. 17. The fact that the whole drawing of G_k remains crossing-free and weakly convex relies on the following inductive property, which is easily shown to be maintained at each step k from 1 to n :

PI: In the upper boundary part γ between two consecutive active vertices —written (from left to right) as $\gamma = P_1, e, P_2$ with e the bottom-edge of γ — the edges of P_1 have slope -1 , e has slope in $\{-1, 0, 1\}$, and the edges of P_2 have slope $+1$.

For each bottom-edge e on C_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 bottom-edges $e \in C_k$.

Remark 10. Clearly the resulting drawing has also the property that for each edge $e \in G$, the absolute value of the slope of e is at most 1 if $e \notin F$ and is at least 1 if $e \in F$.

We now prove the bounds on the grid-size, where w is the width and h the height of the cylinder on which G is drawn. If $|C_{\text{inn}}| = t$ then the initial cylinder is $2t \times 0$; and at each vertex insertion, the grid-width grows by 0 or 2. Hence $w \leq 2n$. In addition, due to the slope conditions, the vertical span of every internal face f is not larger than the current width at the time when f is inserted in the drawing. Hence, if we denote by v the vertex of C_{ext} that is closest (at face-distance d) to C_{inn} , then the y -coordinate of v is at most $d \cdot (2n)$. And due to the slope conditions, the vertical span of C_{ext} is at most $w/2 \leq n$. Hence the grid-height is at most $n(2d + 1)$. The linear-time complexity is shown next.

Linear-time complexity. The algorithm is completely similar to the one in the triangulated case. In a first pass one computes the y -coordinates of vertices and the x -span r_e of each edge $e \in G$ at the time $t = k$ when e appears in G_k (as well one gets to know which extremity of e is the left-end vertex). Afterwards if $e \notin F$ and $e \notin C_{\text{ext}}$, the x -span of e might further increase due to the insertion of new vertices. Let s_e be the total further increase of the x -span undergone by e after its appearance. Let $w_e \geq 0$ be the increase of the x -span undergone by e at the step k such that $e \in C_{k-1}$ and $e \notin C_k$ if such a step exists (otherwise $w_e = 0$). The quantity w_e is called the *weight* of e ; the w_e 's can be computed in a first pass together with the quantities r_e . Let T_e^* be the subtree hanging from e^* (including e^*) in the dual forest F^* , and let W_e be the sum of the weights of the dual of all edges in T_e^* . Then, as in the triangulated case, $s_e = W_e$. Since the quantities s_e are easily computed in linear time from the quantities w_e , starting from the leaves and going up to the roots of F^* , this gives a linear time algorithm.

To sum up, we have proved Proposition 8 for internally 3-connected cylindric maps with no active vertices for C_{inn} .

Remark 11. As in the triangulated case (Remark 5), if we consider the *initial-stretch vector* $R = (r_e)_{e \in C_{\text{inn}}}$ and the *final-stretch vector* $T = (t_e)_{e \in C_{\text{inn}}}$, then the vector $S := T - R$ is an invariant (it does not depend on R), because $s_e = t_e - r_e$ only depends on the canonical ordering. Hence if a certain drawing yields final-stretch vector T , starting from initial-stretch vector R , then for any vector T' of the form $T + 2V$ —with V a vector of non-negative integers— one can redraw the graph to have final-stretch vector T' , by taking $R' = R + 2V$ as initial-stretch vector instead of R .




Figure 19: Kant's algorithm for internally 3-connected plane graphs is recovered from our algorithm by adding a vertex of degree 2 to complete the inner boundary.

Remark 12. Our algorithm extends Kant's algorithm [18], which works for internally 3-connected planar maps, i.e., planar maps with a simple outer face contour and where all 2-separating curves have to pass by the outer face. Indeed, if we are given an internally 3-connected plane graph G (with an outer root-edge) we can turn it into a cylindric internally 3-connected map \tilde{G} by adding a vertex of degree 2 connected to the two extremities of the root-edge; we also declare as active all outer vertices of \tilde{G} that are not one of the three vertices of the inner boundary. Then, Kant's drawing of G is recovered from the periodic drawing of \tilde{G} upon deleting the added vertex of degree 2, see Fig. 19.

4.5 Allowing for non-contractible 2-separating curves (no active vertex on C_{inn}).

The method based on a canonical ordering and incremental drawing algorithm is easily extended to essentially internally 3-connected maps with no 1-separating curve. Additionally we require here that at least two outer vertices are active (the case of one outer vertex active will be handled in the next section). Let G be such a cylindric map. The definition of canonical ordering for G is exactly the same as for internally 3-connected cylindric maps, adding the possibilities that G_k is obtained from G_{k-1} as shown in Fig. 20.

Such a canonical ordering can be computed by a shelling procedure that extends the one of Section 4.3. Call a 2-separating curve *internal* if it passes by two internal faces and does not pass by two outer vertices. In terms of the associated corner-map S such a 2-separating curve corresponds to a separating 4-cycle where the two incident white vertices (vertices of the cylindric map) are not both on C_{ext} . This time, a vertex on the outer boundary and not on the inner boundary is called *admissible* if it is not incident to a separating face nor incident to an internal 2-separating curve. By the same arguments as in Section 4.3, one can show that if there is a separating face, then either there is an admissible vertex (which can be shelled) or there is an internal face sharing a path of $r \geq 2$ edges with




Figure 20: When non-contractible 2-separating curves are allowed (and there are at least two active vertices for C_{ext}), the additional cases shown here might occur for the transition from G_{k-1} to G_k .

the outer boundary (which can be shelled). If there is no separating face but there is at least one internal 2-separating curve, one can check that, if two vertices of C_{ext} are incident to an internal 2-separating curve, then one must be in the situation in the right-part of Fig. 20. Otherwise at most one vertex of C_{ext} is incident to a 2-separating curve, and since there are at least two outer vertices active, at least one of them is admissible. Finally, in case there is no separating face nor 2-separating curve, then any active outer vertex is admissible and can be shelled. One can continue shelling iteratively until there just remains C_{inn} (since there is no 1-separating curve, it is easy to see that the property of having at least two active outer vertices is automatically maintained). A linear time algorithm is also readily obtained by maintaining, for each outer vertex, how many separating faces and how many internal 2-separating curves it is incident to. Note that such a canonical ordering also induces an underlying forest F and a dual forest F^* . Finally, the incremental drawing algorithm (and linear complexity using the dual forest) works in the same way as for internally 3-connected cylindric maps. An example is shown in Fig. 21, where the situation corresponding to the right-part of Figure 20 is shown in the last step of the drawing.

The grid bounds are also the same as for internally 3-connected cylindric maps (the arguments to obtain the bounds in the simple case did not use the fact that there are no 2-separating curves). So this gives Proposition 8 for essentially internally 3-connected cylindric maps with no 1-separating curve and with at least two outer vertices active. Finally, Remark 11 still holds here (the arguments are the same).

4.6 Allowing for 1-separating curves (no active vertex for C_{inn}).

We finally explain how to deal with 1-separating curves. Similarly as in the triangulated case, we do not extend the notion of canonical ordering but simply decompose (at 1-separating curves turned into loops) such a cylindric map into a “tower” of components separated by loops. Let G be an essentially internally 3-connected cylindric map with n vertices. Note that if G has a loop e , then there is a 1-separating curve “just inside” e (if e is not C_{inn}) and there is a 1-separating curve “just outside” e (if e is not C_{ext}).

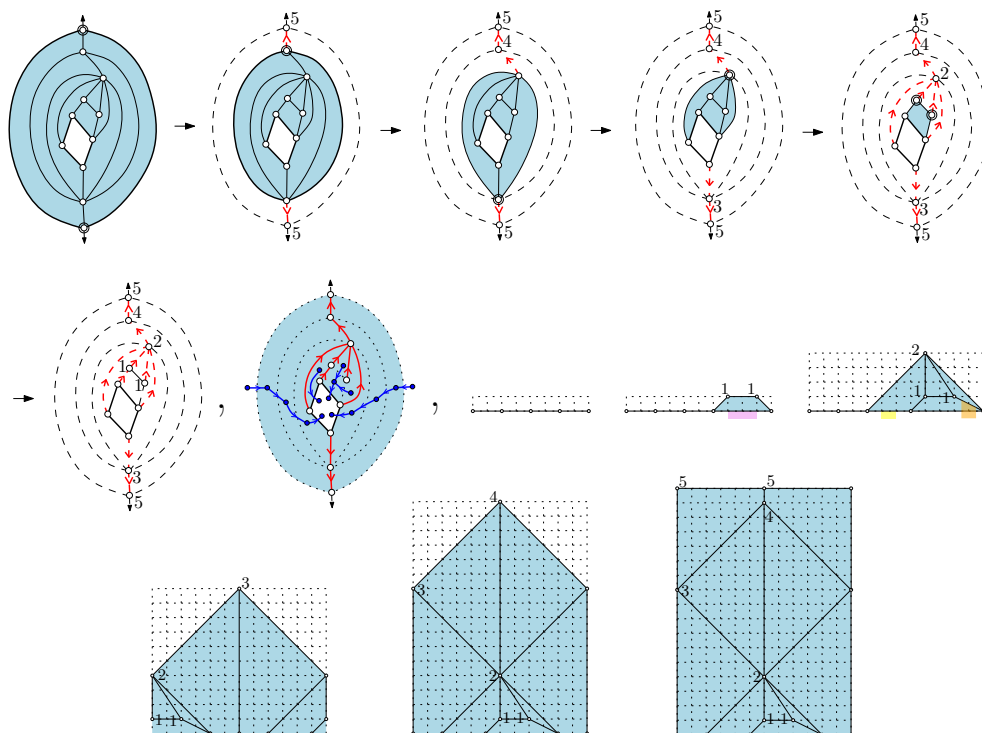


Figure 21: (Before first comma): The shelling procedure for an essentially internally 3-connected cylindric map G with at least two active vertices for C_{ext} , with no active vertex for C_{inn} nor 1-separating curve; (after first comma) the underlying forest and dual forest; (after second comma) the incremental drawing algorithm.




Figure 22: Drawing algorithm when there is a unique 1-separating curve that is just outside the inner boundary (which is a loop).

Similarly as for cylindric triangulations, there are a few cases to consider:

(a) C_{inn} is a loop, the unique 1-separating curve is the one just outside C_{inn} , and there are at least two active vertices. In that case, the algorithm of Section 4.5 (canonical ordering, shelling procedure, and incremental drawing procedure) works in the same way, see Fig. 22 for an example. In addition, by the arguments of Remark 11, for any $n' \geq n$, G has a periodic drawing of width $2n'$ and height at most $n'(2d + 1)$, with d the face-distance between the two boundaries.

(b) C_{ext} is a loop, C_{inn} is possibly a loop, and the only 1-separating curves are just inside C_{ext} and just outside C_{inn} . In that case, there is the possibility that C_{inn} and C_{ext} are both loops and are connected by one edge e , in which case the drawing is very easy (draw C_{inn} and C_{ext} as horizontal segments, and e as a vertical segment). Otherwise, let u be the unique outer vertex; it is incident to a 1-separating curve; equivalently it is incident to a non-contractible 2-cycle c' in the corner-map S . Then, since we are not in the first case above, it is easy to see that, in S , there is a non-contractible 4-cycle “directly inside” c' . Hence the set of non-contractible 4-cycles is non-empty. Let c be the “innermost” non-contractible 4-cycle, which we now consider in G in the form of a 2-separating curve. Let f_1, f_2 the two internal faces visited by c and let v the vertex different from u and visited by c . If there is an edge of f_1 connecting u and v , call e_1 this edge, otherwise draw an edge, again called e_1 , inside f_1 that connects u and v . Similarly if there is an edge of f_2 connecting u and v , call e_2 this edge, otherwise draw an edge e_2 inside f_2 that connects u and v . Call \tilde{c} the 2-cycle formed by e_1 and e_2 . Cutting along \tilde{c} we obtain two components: a plane 3-connected map T of outer degree 3, and a cylindric essentially internally 3-connected map G' of outer degree 2 and such that G' has no internal 2-separating curve incident to u , and no 1-separating curve, except possibly just outside C_{inn} if C_{inn} is a loop. Hence, if we declare both outer vertices of G' as active, then u is admissible and we can apply the results of Section 4.5: there is a canonical ordering of G' such that u is the first shelled vertex; and this canonical ordering yields a periodic drawing of G' where the upper boundary is made of e_1 , of slope -1 , and e_2 , of slope $+1$. In addition, denoting by $2m$ the maximum of the widths of the drawings of T and of G' , one can redraw the drawing of smaller width so that both drawings have width $2m$. Then the drawing of T (taken upside down) fits into the upper boundary of G' , yielding a periodic drawing of G of width $2m$, see Fig. 23. In that




Figure 23: Drawing algorithm when C_{ext} is a loop and the unique 1-separating curve is just inside C_{ext} : the map is split along the innermost 2-separating curve.

case, the height is at most $2dm$ where d is the face-distance between the two boundaries (indeed, in the usual bound $h \leq (2d + 1)m$ the $+1$ in the parenthesis is due to the vertical extension of the upper boundary, which is 0 here). More generally, by the arguments of Remark 11, for any $n' \geq m$ there is a periodic drawing of G with horizontal lower and upper boundaries, of width $2n'$ and height at most $2dn'$. It remains to show that, if e_1 (similarly e_2) was not present in G , then the deletion of e_1 keeps the drawing weakly convex. To get convinced of it, one just has to notice that, if e_1 is not in G and thus is to be deleted, then the properties of the drawings (Remark 10) guarantee that the $(\pi/2)$ sector between e_1 and e_2 in the drawing of T contains at least one edge (different from e_1, e_2); and in the drawing of G' the $\pi/2$ sector around v starting (in ccw direction) from e_1 also contains at least one edge (different from e_1). Hence the angle at v left by the deletion of e_1 is at most π . Similarly the angle at u left by the deletion of e_1 is also at most π ; and the angles at u, v left by the deletion of e_2 (if e_2 is to be deleted) are at most π .

(c) *General case (no active vertex at C_{inn})*. For each 1-separating curve c , if c is not just inside nor just outside a loop, then we draw a loop-edge along c . Also, if there is just one active vertex and C_{ext} is not a loop, then we add a loop at the unique active vertex, so that C_{ext} becomes a loop. After doing this, let r be the number of loops (note that the case $r = 0$ has already been addressed in the previous sections, so we can assume $r \geq 1$), and let ℓ_1, \dots, ℓ_r be the sequence of nested loops of G (from innermost to outermost) and let $G^{(0)}, \dots, G^{(r)}$ be the $r + 1$ components that result from cutting successively along all these loops, see Fig. 24. For $i \in [0..r]$ let d_i be the face-distance between the two boundaries in $G^{(i)}$; and let d be the face-distance between the two boundaries of G . Note that $d = \sum_i d_i$. By construction, each component G_i has its 1-separating curves either just inside the outer boundary (if the outer boundary is a loop, which occurs if $i \in [0..r - 1]$) or just outside the inner boundary (if the inner boundary is a loop, which occurs if $i \in [1..r]$). Hence each component G_i has a periodic drawing (according to cases (a) and (b)) whose boundaries are drawn as horizontal lines, except possibly for the outer boundary of G_r . Let $2n_1, \dots, 2n_r$ be the widths of the drawings of $G^{(1)}, \dots, G^{(r)}$ thus obtained. Let $m = \max(n_1, \dots, n_r)$. By the arguments of Remark 11, each of the graphs $G^{(i)}$ can be redrawn so as to have width $2m$. Stacking up all these drawings, we obtain a periodic drawing of G of width $2m$ and height at most $m(2d + 1)$.




Figure 24: Drawing an essentially internally 3-connected cylindric map G (no active vertex for C_{inn}); loops (dashed in the left drawing) are drawn along 1-separating curves, then G is decomposed at its loops; each component is redrawn so that the component-drawings have the same width, and can be stacked up to obtain a periodic drawing of G .

It remains to check that the deletion of the added loops keeps the drawing weakly convex. In the periodic drawing the loop is drawn as an horizontal segment $S = [a, b]$, where we denote by a the left end and b the right end. The properties of the drawings (Remark 10) easily imply that there is at least one (non-loop) edge in the $\pi/4$ sector around a starting (in counterclockwise direction) at S , and there is at least one edge in the $3\pi/4$ sector around a starting (in clockwise direction) at S . Hence the angle at a left by the deletion of S is at most π . Similarly the angle at b left by the deletion of S is at most π .

This establishes Proposition 2.

4.7 Allowing for active vertices at C_{inn} .

We finally explain how to draw an essentially internally 3-connected cylindric map G in the general case, i.e., when allowing for active vertices at C_{inn} . The only requirement here is that there is at least one active vertex (for C_{ext} or C_{inn}). We can assume that there is also at least one active vertex for C_{ext} , otherwise the situation would be symmetric to the case where C_{inn} has no active vertex, considered in the previous sections.

4.7.1 The boundaries C_{ext} and C_{inn} are disjoint

Let G be an essentially internally 3-connected cylindric map, with at least one active vertex for C_{ext} , and such that C_{ext} and C_{inn} are disjoint (as in the triangulated case, it is good here to imagine a standing cylinder, so that C_{inn} is seen as the lower boundary and C_{ext} as the




Figure 25: Drawing an internally 3-connected cylindrical map with active vertices (and 2-chords) at C_{inn} .

upper boundary). As in the previous section, we can consider the corner-map S of M , and define a *lower 2-chord* as a separating 2-chord for C_{inn} . A lower 2-chord is called *maximal* if it is maximal for the containment relation of enclosed cycles. For each maximal lower 2-chord, we can add an edge e “just above” such a 2-chord (in the internal face it passes by) if e is not already present. For each such edge e , let Q_e be the so-called *component under e* , which is the subgraph of G induced by e and by the vertices and edges in the component below e . Detaching all such components we obtain a reduced cylindrical map G' such that, even when declaring all vertices of C_{inn} as non-active, G' is essentially internally 3-connected. We can thus draw G' using the procedure of Proposition 8. And we draw all detached components Q_e using Kant’s algorithm. According to Remark 11 we can then redraw G' so that each edge e has width large enough to plug the drawing of Q_e (recall that, if $|e|$ is the width of Kant’s drawing of Q_e and w_e is the width of e in the first obtained drawing of G' , then G' can be redrawn so that the width $\ell(e)$ of e is the smallest integer at least $\max(|e|, w_e)$ and congruent to w_e modulo 2). Then for each such edge e , one can shift the left-extremity of the drawing of Q_e so that the width of the bottom-edge of Q_e is $\ell(e)$, and then (after rotation of angle π) plug the drawing of Q_e into the lower segment for e in the drawing of G' . The obtained drawing of G is clearly crossing-free, see Fig. 25 for an example.

We also have to check that, after deleting the edges that have been added (one such edge just above each maximal lower 2-chord, when needed), all angles corresponding to corners in internal faces are at most π . Let e be such a deleted edge; note that (before being deleted), e is horizontal in the drawing; let u be its right extremity and v its left extremity. The properties of the drawings easily imply that the sector of angle $3\pi/4$ around u starting (in counterclockwise direction) at e contains an edge, indeed u is either the left extremal vertex of A or otherwise it has one incident edge in the underlying forest for A' . Moreover, since u is the left extremal vertex of Q_e , the sector of angle $\pi/4$ around u starting (in clockwise direction) at e contains an edge. We conclude that the corner at u left by the deletion of e has angle at most π in the drawing. Similarly the angle at v left by the deletion of e has angle at most π . And by similar arguments as for cylindrical triangulations, the width of the grid is at most $2n$, with n the number of vertices of G , and the height is at most $2n(d+1)$, with d the face-distance between the two boundaries.




Figure 26: An internally essentially 3-connected cylindric map G where C_{ext} and C_{inn} intersect at 3 vertices. Accordingly G decomposes as a cyclic sequence of 4 portions (one of which is trivial). Each portion can be drawn (after adding an edge to complement it into a cylindric map) by a decomposition at maximal lower 2-chords. The drawings of the portions are then attached sequentially along the horizontal axis.

4.7.2 The boundaries C_{ext} and C_{inn} meet

We now handle the case where C_{ext} and C_{inn} intersect. Recall that in that case G decomposes along such intersections as a cyclic sequence of elementary blocks called *portions*, where each portion is delimited by two intersection-vertices that are said to be *extremal* (for the considered portion). A portion A is *complemented* by adding an outer edge e_{out} connecting the two extremal vertices, so as to obtain a cylindric map \tilde{A} . Note that A , if non-trivial, has at least one non-extremal vertex that is active, either at C_{inn} or at C_{ext} ; without loss of generality (one can exchange the roles of the outer and inner boundary for \tilde{A}) we can assume that A has at least one active vertex for C_{ext} . We can thus draw \tilde{A} by a procedure similar to the one described in the previous section (by detaching components under maximal lower 2-chords).

The overall drawing of A is periodic planar, with e_{out} drawn as an horizontal segment of width 2. In addition, by similar arguments as in the previous section it can be checked that the drawing is still weakly convex after deleting the added edges (edges added just above maximal lower 2-chords).

We can finally attach together (along the horizontal axis) the drawings of all portions thus obtained (where the complementation edge is removed) to obtain an x -periodic drawing of G . By similar arguments as in the triangulated case it can be checked that the width and height of the drawing are at most $2n$, with n the number of vertices in G . This concludes the proof of Theorem 7.

5 Periodic drawings on the torus

5.1 Definitions and statement of the results

A toroidal triangulation is a map on the torus with only triangular faces. A toroidal map is called *simple* if it has no loop nor multiple edges, and is called *essentially simple* if its

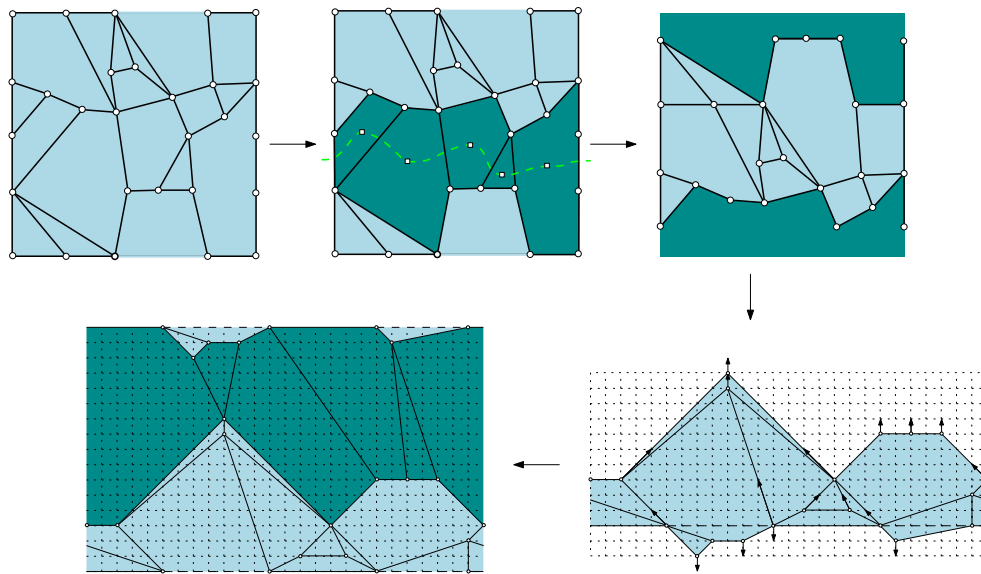


Figure 27: The successive steps to draw an (essentially) 3-connected toroidal map: 1) remove the edges inside a tambourine, (2) draw the obtained (essentially) internally 3-connected cylindric map, 3) insert the edges of the tambourine back into the drawing.

periodic representation in the plane is simple. A toroidal map is called *3-connected* if it is 3-connected as a graph, and is called *essentially 3-connected* if its periodic representation⁸ in the plane is 3-connected. Let G be a toroidal map. On the torus a non-contractible curve is a closed curve that cannot be continuously deformed to a point. Such a curve is called *proper* if it meets G only at vertices (not at edges); the *length* of a proper non-contractible curve is the number of vertices it meets. The *face-width* of G is the minimum of the lengths of the proper non-contractible curves for G . A *non-contractible cycle* of G is a (simple) cycle of edges of G that forms a non-contractible curve (the length of such a cycle is its number of edges). The *edge-width* of G is the minimum of the lengths of the non-contractible cycles of G . In this section we obtain the main result of the article:

Theorem 13. *For each essentially 3-connected toroidal map G , one can compute in linear time a weakly convex 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 c the face-width— $w \leq 2n$ and $h \leq 1 + 2n(c + 1)$. Since $c \leq \sqrt{2n}$, the grid area is $O(n^{5/2})$.*

The fact that $c \leq \sqrt{2n}$ follows from [2]. Indeed, an essentially simple toroidal graph G (with n the number of vertices) can be triangulated, by adding edges but not adding vertices, so as to yield an essentially simple toroidal triangulation \tilde{G} . As shown in [2], \tilde{G}

⁸ The periodic representation (also called universal cover) is obtained by gluing an infinite number of translated copies of our drawing, assembling the right (resp. upper) side of one copy with the left (resp. lower) side of another. Then a non contractible loop become an edge linking two different translated copies of the same point.

has a non-contractible cycle γ of length at most $\sqrt{2n}$, and it is easy to see that this cycle can be locally deformed into a proper non-contractible curve meeting G at the vertices on γ .

5.2 Tambourine: definition, existence, and computation

Let G be a toroidal map, and let Γ_1, Γ_2 be a pair of homotopic non-contractible cycles of G that are oriented in the same direction. The pair Γ_1, Γ_2 is called a *tambourine* if the area A between Γ_1 and Γ_2 —on the right of Γ_1 and on the left of Γ_2 —is a “ribbon of faces”, i.e., A is face-connected and the set of edges dual to edges in $A \setminus \{\Gamma_1, \Gamma_2\}$ forms a non-contractible cycle homotopic to Γ_1 and Γ_2 . The second drawing of Fig. 27 give an example of such a configuration. Note that Γ_1 and Γ_2 might share vertices and edges. It is shown in the master’s thesis of Arnaud Labourel (see also [4]) that for each non-contractible cycle Γ of a simple toroidal triangulation G , there exists a tambourine whose two cycles are homotopic to Γ . We show here that this holds more generally for essentially 3-connected toroidal maps.

Lemma 14. *Let G be an essentially 3-connected toroidal map and let Γ be a non-contractible cycle of G . Then there is a tambourine Γ_1, Γ_2 of G whose two cycles are homotopic to Γ . Moreover one can compute such a tambourine in linear time.*

Proof. Cutting along Γ we obtain a cylindric map \widehat{G} , which we consider in its annular representation. Let C be the inner boundary of \widehat{G} . Let Γ_1 be the smallest (in terms of the enclosed area) cycle that strictly encloses C (i.e., encloses C and is vertex-disjoint from C). Then, let Γ_2 be the largest (in terms of the enclosed area) cycle that is strictly enclosed in Γ_1 (i.e., is enclosed by Γ_1 and is vertex-disjoint from Γ_1). Orient Γ_1 and Γ_2 clockwise (in the annular representation). Let A be the annular area between Γ_1 and Γ_2 . Note that, by minimality of the enclosed area, Γ_1 has no chord inside. Similarly, by maximality of the enclosed area, Γ_2 has no chord outside; hence both cycles delimiting A have no chordal edge in A . It just remains to show that there is no vertex in the strict interior of A . Assume the set E of such vertices is not empty, and let H be a connected component of the subgraph induced by E . Call *vertex of attachment for Γ_1* (resp. for Γ_2) a vertex $w \in \Gamma_1$ (resp. $w \in \Gamma_2$) such that there is a path from a vertex of H to w that visits only vertices of H before reaching w . By minimality of Γ_1 , it is easy to see that there is a unique vertex of attachment v_1 for Γ_1 . Similarly, by maximality of Γ_2 , there is a unique vertex of attachment v_2 for Γ_2 . Let H_2 be the graph made by adding to H the vertex v_2 as well as all edges connecting v_2 to vertices from H . The annular representation is actually an embedding in the plane, so the faces of H_2 partition the plane. Hence the area corresponding to the inner boundary-face of G is either inside an inner face f of H_2 or is inside the outer face of H_2 . In the first case, the contour of f yields a cycle that is vertex-disjoint from Γ_1 and whose interior strictly contains the interior of Γ_2 . This is impossible by maximality of Γ_2 . In the second case (B_{inn} in the outer face of H_2), let H_{12} be the map obtained from H_2 by adding the edges from H to v_1 . Two subcases can arise. If the area of B_{inn} is in the outer face of H_{12} , then there is a 2-separating curve (for \widehat{G}) passing by v_1 and v_2 and enclosing H but not enclosing the inner boundary-face, which contradicts the fact that G is essentially 3-connected; if the area of B_{inn} is in a bounded face of H_{12} then there must exist a cycle C passing by v_1 and by vertices of H such that C encloses B_{inn} , which contradicts the

minimality of Γ_1 . In all cases we reach a contradiction, so Γ_1, Γ_2 form a tambourine.

Let us now justify that Γ_1, Γ_2 can be computed in linear time. Let \widehat{G}_1 be \widehat{G} where vertices and edges of C have been deleted. Let f be the inner face of \widehat{G}_1 whose interior contains the interior of C (such a face is unique, even when \widehat{G}_1 has several connected components), and let c be the unique simple cycle extracted from the contour of f and such that the interior of c contains the interior of C . Then Γ_1 equals c (indeed c is vertex-disjoint from C , the interior of c contains the interior of C , and no other cycle enclosed by c can have these two properties). Now it is clear that computing c (i.e., finding f and then extracting c from f) can be done in linear time. The computation of Γ_2 is very similar. Let \widehat{G}_2 be the subgraph of \widehat{G} where all vertices and edges on Γ_1 and exterior to Γ_1 have been deleted. Let c' be the unique cycle extracted from the outer face contour of \widehat{G}_2 and such that the interior of c' contains the interior of C . Then Γ_2 equals c' , so Γ_2 is also easily computable in linear time. \square

5.3 The drawing algorithm

Let G be an essentially 3-connected toroidal map with n vertices, and let Γ_1, Γ_2 be a tambourine of G . By deleting the edges that are strictly inside the tambourine, one obtains a cylindrical map G' with Γ_1 the outer boundary and Γ_2 the inner boundary. Declare as active for Γ_1 each vertex $v \in \Gamma_1$ incident to at least an edge e inside the tambourine (with an incidence on the right-side of Γ_1). And declare as active for Γ_2 each vertex $v \in \Gamma_2$ incident to at least an edge e inside the tambourine (with an incidence on the left-side of Γ_2). In the periodic representation of G' , if there is a 2-separating curve with no active vertex strictly inside, then it yields a 2-separating curve in the periodic representation of G , a contradiction. Hence G' is essentially internally 3-connected. We can apply the drawing algorithm of Theorem 7 to obtain a weakly convex periodic drawing of G on an x -periodic grid $\mathbb{Z}/w\mathbb{Z} \times [0..h]$. 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 (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). 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 and weakly convex. Note that it is actually enough to augment the height by the least value such that the edges of T have slope greater than 1 in absolute value. See Fig. 27 for an example, where the height is augmented by 4 (whereas $w = 26$).

We now argue that we can find a tambourine Γ_1, Γ_2 so that the face-distance between the two boundaries Γ_1 and Γ_2 (in G') is smaller than the face-width of G ; and we can find it *without having to compute* a curve Γ_{\min} realizing the face-width (this is crucial to obtain a linear-time complexity, since it is not known how to find such a curve Γ_{\min} in linear time). Indeed, let $\{\Gamma_a, \Gamma_b\}$ be a basis of non-contractible cycles of G (computable in linear time, using for instance a cut-graph). Then at least one of Γ_a or Γ_b is not homotopic (parallel) to Γ_{\min} . Let Γ be one among $\{\Gamma_a, \Gamma_b\}$ that is not homotopic to Γ_{\min} , and let Γ_1, Γ_2 be a

(computable) tambourine parallel to Γ . Since we are on the torus, Γ_{\min} has to cross the tambourine Γ_1, Γ_2 . Hence, the distance between the boundary-cycles (after deleting edges in the tambourine Γ_1, Γ_2) is smaller than the face-width. In other words, if we choose the one cycle among $\{\Gamma_a, \Gamma_b\}$ that yields the smaller face-distance between the two boundaries of G' , then this distance d is smaller than the face-width c of G . The grid-size of the drawing of G' satisfies $w \leq 2n$ and $h \leq 2n(d+1)$, and the grid-size of the drawing of G is w, h' where $h' \leq h + w + 1 \leq 2n(d+1) + 2n + 1 = 1 + 2n(d+2)$. Since $d < c$ we conclude that the grid-height of the drawing of G is bounded by $1 + 2n(c+1)$. This concludes the proof of Theorem 13.

Acknowledgments. The authors thank Nicolas Bonichon, Daniel Gonçalves, Benjamin Lévêque, and Bojan Mohar for interesting discussions.

References

- [1] B. Albar, D. Gonçalves and K. Knauer. Orienting Triangulations. *J. Graph Theory*, 83(4): 392-405, 2016.
- [2] M. O. Albertson and J. P. Hutchinson. On the independence ratio of a graph. *J. Graph. Theory*, 2:1-8, 1978.
- [3] N. Bonichon, S. Felsner and M. Mosbah. Convex Drawings of 3-Connected Plane Graphs. *Algorithmica*, 47(4): 399-420, 2007.
- [4] N. Bonichon, C. Gavaille and A. Labourel. Edge partition of toroidal graphs into forests in linear time. In *ICGT*, volume 22, pages 421-425, 2005.
- [5] F.-J. Brandenburg. Drawing Planar Graphs on $(8/9)*n^2$ Area. *Electronic Notes in Discrete Mathematics*, 31:37-40, 2008.
- [6] E. Brehm. 3-orientations and Schnyder 3-Trees decompositions. Master's thesis, FUB, 2000.
- [7] L. Castelli Aleardi, É. Fusy, and T. Lewiner. Schnyder woods for higher genus triangulated surfaces, with applications to encoding. *Discr. & Comp. Geom.*, 42(3):489-516, 2009.
- [8] L. Castelli Aleardi, É. Fusy and A. Kostygin. Periodic planar straight-frame drawings with polynomial resolution. Proceedings of *Latin'14*, Lecture Notes in Computer Science Vol. 8392, pages 168-179, 2014.
- [9] E. Chambers, D. Eppstein, M. Goodrich and M. Löffler. Drawing Graphs in the Plane with a Prescribed Outer Face and Polynomial Area. *JGAA*, 16(2):243-259, 2012.
- [10] M. Chrobak and G. Kant. Convex Grid Drawings of 3-Connected Planar Graphs. *Int. J. Comput. Geometry Appl.*, 7(3):211-223, 1997.
- [11] M. Chrobak and T. H. Payne. A Linear-Time Algorithm for Drawing a Planar Graph on a Grid. *Inf. Process. Lett.*, 54(4): 241-246, 1995.
- [12] V. Despré, D. Gonçalves and B. Lévêque. Encoding toroidal triangulations. *Discrete and Computational Geometry*, 57: 507-544, 2017.
- [13] C. Duncan, M. Goodrich and S. Kobourov. Planar drawings of higher-genus graphs. *Journal of Graph Algorithms and Applications*, 15:13-32, 2011.

- [14] S. Felsner. Convex Drawings of Planar Graphs and the Order Dimension of 3-Polytopes. *Order*, 18(1): 19-37, 2001.
- [15] H. de Fraysseix, J. Pach and R. Pollack. How to draw a planar graph on a grid. *Combinatorica*, 10(1):41-51, 1990.
- [16] D. Gonçalves and B. Lévêque. Toroidal maps : Schnyder woods, orthogonal surfaces and straight-line representations. *Discrete & Computational Geometry*, 51: 67-131, 2014.
- [17] 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.
- [18] G. Kant. Drawing planar graphs using the canonical ordering. *Algorithmica*, 16(1):4-32, 1996.
- [19] W. Kocay, D. Neilson and R. Szypowski. Drawing graphs on the torus. *Ars Combinatoria*, 59:259-277, 2001.
- [20] K. Miura, S.-I. Nakano and T. Nishizeki. Grid Drawings of 4-Connected Plane Graphs. *Discrete & Computational Geometry*, 26(1):73-87, 2001.
- [21] B. Mohar. Straight-line representations of maps on the torus and other flat surfaces. *Discrete Mathematics*, 15:173-181, 1996.
- [22] B. Mohar and P. Rosenstiehl. Tessellation and visibility representations of maps on the torus. *Discrete & Comput. Geom.*, 19:249-263, 1998.
- [23] W. Schnyder. Embedding planar graphs on the grid. *SODA*, pp 138-148, 1990.
- [24] A. Zitnik. Drawing graphs on surfaces. *SIAM J. Disc. Math*, 7(4):593-597, 1994.