

Tessellation and visibility representations of maps on the torus

Bojan Mohar¹

Department of Mathematics
University of Ljubljana, Slovenia

and

Pierre Rosenstiehl²

Centre d'Analyse et de Mathématique Sociales
EHESS, Paris, France

Abstract

The model of the torus as a parallelogram in the plane with opposite sides identified enables us to define two families of parallel lines and to tessellate the torus, then to associate to each tessellation a toroidal map with an upward drawing. It is proved that a toroidal map has a tessellation representation if and only if its universal cover is 2-connected. Those graphs that admit such an embedding in the torus are characterized.

1 Introduction

Given a graph G , let $V(G)$ be the set of vertices of G , $E(G)$ the set of edges of G . For $A \subseteq V(G)$ we denote by $E(A)$ the set of edges of G with both ends in A .

A map M on a surface Σ is a connected graph G together with a 2-cell embedding of G in Σ . Two maps are *equivalent* if there is a homeomorphism

¹Supported in part by the Ministry of Science and Technology of Slovenia, Research Project P1-0210-101-94.

²This work was partially supported by the ESPRIT Basic Research Action No. 7141 (ALCOM II).

of Σ mapping the graph of the first map onto the graph of the second. It is well-known that the equivalence classes of maps on orientable surfaces correspond naturally to rotation systems on the underlying graphs [4]. Let us recall that a *rotation system* on a graph G is a set of cyclic permutations π_v ($v \in V(G)$) where π_v cyclically permutes edges emanating from v ; π_v corresponds to the cyclic order of these edges around v on the surface. A similar combinatorial representation can be given for maps on nonorientable surfaces (cf., e.g., [4] or [10]). The rotation system around the faces of M defines the *dual map* M^* . A map M and its dual M^* can be simultaneously drawn in Σ such that each vertex of M^* corresponds to an interior point of the corresponding face of M (and *vice versa*), and such that precisely dual pairs of edges cross each other.

If M is a map in Σ whose underlying graph has n vertices, e edges and f faces, then $n - e + f = \chi(\Sigma)$ where $\chi(\Sigma)$ is the *Euler characteristic* of the surface Σ .

In this paper we only consider maps on *closed surfaces* (i.e. compact surfaces without boundary) and maps in the plane. Most attention is given to the torus which can be represented by a parallelogram Q in the plane whose opposite sides are pairwise identified. This representation of the torus carries the local geometry of the plane and hence we call it the *flat torus*.

From now on we shall assume that Σ is a closed surface, M a map on Σ , and G the underlying graph of M . An *angle* of M (respectively, M^*) is a pair of consecutive arcs at a vertex v of M (respectively, a face f of M). The *angle map* of M is a map \mathcal{A} on Σ whose vertices are the vertices of M plus the vertices of M^* (i.e., the faces of M), and whose edges are the angles of M , each angle being incident with the corresponding vertex and face of M [12]. The set of angles incident with a vertex v of M has a local rotation determined by M , and the set of angles incident with a face of M inherits the local rotation from M^* . The angle map \mathcal{A} is bipartite and each face of \mathcal{A} is a quadrangle whose diagonals are a pair of dual edges of M and M^* . An example is shown in Figure 1(b) where the dotted lines represent the edges of M and they are not part of \mathcal{A} . The dual map of \mathcal{A} is known as the *medial map* of M .

An important concept related to drawings of graphs in the plane is the notion of upward drawings where each edge of a graph is oriented and drawn in the plane \mathbf{R}^2 so that the ordinate monotonically increases when we traverse the edge according to its orientation. In order to define an upward drawing of a graph G on the torus, we first need a definition of a monotone arc. Consider, without loss of generality, the flat torus obtained from




Figure 1: A toroidal map and its angle map

a parallelogram Q in the plane. The lines parallel to the sides of Q oriented according to the usual axes determine *horizontal* and *vertical* circuits on the torus. At each point of the torus two circuits cross, the vertical one being crossed always by the horizontal one, for instance, from left to right. A (polygonal) arc on the torus is *monotone* if it can be oriented such that by traversing the arc in the chosen direction, one crosses horizontal circles only from bottom to top. In particular, the vertical circuits are examples of monotone arcs. Let us observe that a monotone arc may cross a horizontal circuit more than once by “winding” around two or more times. The definition of horizontal circles is easily extended to the case when the horizontal direction is not necessarily parallel to a side of Q . Monotone arcs are defined accordingly. An *upward drawing* of a toroidal map M is a map equivalent to M drawn on the torus with monotone (polygonal) arcs such that at each vertex v of M , at least one edge incident to v enters v from below (with respect to the chosen horizontal direction) and at least one edge enters v from above (see Figure 1(a)). The authors have shown in [9] by using network flow techniques that every toroidal map with 2-connected universal cover admits a drawing in the flat torus that is an upward drawing and, moreover, the corresponding drawing of the dual map is upward at the same time (with respect to the vertical circuits instead of the horizontal ones).

A stronger concept than upward drawings is the tessellation representation of a map. This concept was investigated in case of planar graphs by Tamassia and Tollis [14]. Given horizontal and vertical directions Δ and




Figure 2: A tessellation representation of a toroidal map

Δ' of the flat torus (possibly distinct from the directions of the sides of the fundamental parallelogram Q), we consider a family of horizontal line segments called the vertices, and a family of vertical line segments called the faces, all disjoint except that each segment of one family touches at each end a segment of the other family at one of its internal points. A map is associated with a tessellation in the following way. Horizontal segments are the vertices, vertical segments the faces. Each of the obtained quadrangular regions determines an edge of the map; it is, by definition, incident with two vertices by its two horizontal sides, and with two faces by its two vertical sides. The local rotation π_v around the vertex v is determined by the order of edges obtained by traversing the boundary of a small neighborhood of the vertical segment corresponding to v in the clockwise direction. This rule determines the map up to equivalence. A *tessellation* of the torus is then a representation of a toroidal map defined by its vertices, faces and edges as introduced above. An example of a tessellation and the corresponding map is shown in Figure 2 and a tessellation representation of the map from Figure 1(a) is presented in Figure 3. Our main result is that a toroidal map admits a tessellation representation if and only if it is essentially 2-connected. This, in particular, implies the above mentioned result of [9] about the simultaneous upward drawings of the map M and its dual M^* .

In Section 2 we present several characterizations of essentially 2-connected maps on the torus and show that they can be obtained from two minimal maps by vertex splitting and creating digons. Section 3 contains the proof of the main theorem about existence of tessellation representations. Part 4 is devoted to the visibility representation and grid contact graphs as corollaries

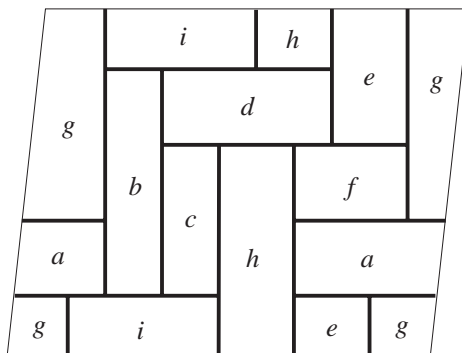


Figure 3: A tessellation representation of the map from Figure 1

of the main theorem. In Part 5, a characterization of those graphs that can be embedded as an essentially 2-connected map on the torus is presented.

Similar questions as addressed in this paper for the flat torus have been considered in the case of the plane by Fraysseix, O. de Mendez, and Pach [3], and for the torus by Kratochvil and Przytycka [6].

2 Essentially 2-connected maps

The *universal cover* of the surface Σ is the simply connected surface $\tilde{\Sigma}$ together with a covering projection $p: \tilde{\Sigma} \rightarrow \Sigma$. We refer to [7] for the general theory of covering spaces and for basic properties of universal covering spaces. If G is the graph of a map on Σ , then $p^{-1}(G)$ defines a map in $\tilde{\Sigma}$. By the homotopy lifting property of covering maps [7], a lifting of a closed walk W of G into the cover $\tilde{\Sigma}$ is a closed walk in $p^{-1}(G)$ if and only if W is contractible in Σ . In particular, every *facial walk* of G (i.e., the walk in G corresponding to the traversal of the boundary of a face) is lifted to facial walks in the covering map.

The universal cover of the flat torus represented by the parallelogram Q is the plane paved with replicates of Q , and the lifting of a map in the torus is an infinite plane graph. A part of such a pavement is shown in Figure 4.

A map is *essentially 2-connected* if the graph of its universal cover map is 2-connected. Figure 4 shows an example (of a part) of the universal cover of an essentially 2-connected toroidal map whose graph is not 2-connected.

Essentially 2-connected maps can be characterized in several other ways.




Figure 4: A part of the universal cover of a toroidal map

Lemma 2.1 *Given a map M of a graph G on a surface Σ whose Euler characteristic $\chi(\Sigma)$ is not positive, the following conditions are equivalent:*

- (i) *M is essentially 2-connected.*
- (ii) *No facial walk f of M contains a proper closed subwalk which is contractible on Σ .*
- (iii) *There are no planar separations in the graph G of M , i.e., if G_1 and G_2 are graphs each having at least one edge, and such that $G = G_1 \cup G_2$ and $G_1 \cap G_2 = \{v\}$, where v is a vertex, then G_1 and G_2 each contain a circuit that is non-contractible on Σ .*
- (iv) *The angle map \mathcal{A} of M has no homotopic pair of parallel edges, i.e., a pair of edges bounding a disk.*
- (v) *For every subset S of the vertices of the angle map \mathcal{A} of M we have:*

$$2|S| - |E(S)| \geq 2\chi(\Sigma). \quad (1)$$

Proof. (i) \Leftrightarrow (ii): Since $\chi(\Sigma) \leq 0$, the underlying surface of the universal covering map \tilde{M} of M is the plane. It is well known that a (possibly infinite) graph in the plane is 2-connected if and only if every facial walk is a simple cycle (not meeting the same vertex twice). Since every facial walk (and every contractible subwalk of a facial walk) in M lifts into \tilde{M} to a facial walk of the same length (a subwalk of a facial walk, respectively), (i) and (ii) are easily seen to be equivalent.

(ii) \Rightarrow (iii): Suppose that G admits a planar separation $G = G_1 \cup G_2$, $G_1 \cap G_2 = \{v\}$, and suppose that G_1 contains only circuits that are contractible on Σ . Consider a facial walk f of M which is not contained entirely in $E(G_1)$ or in $E(G_2)$. Then f is the union of nonempty closed segments from $E(G_1)$ and from $E(G_2)$. By our choice, G_1 contains only contractible circuits. Therefore, any closed segment of f in G_1 is contractible. This contradicts (ii).

(iii) \Rightarrow (iv): Suppose that the angle map \mathcal{A} of M has a pair of parallel edges bounding an open disk D . Let the endpoints of these edges of \mathcal{A} be v , a vertex of M , and f , a face of M . Note that $\partial\overline{D} \cap G = \{v\}$ where $\partial\overline{D}$ is the boundary of the closure of D in Σ . Since \mathcal{A} contains only quadrangular faces, D is not a face of \mathcal{A} . Therefore the subgraph $G_1 = G \cap \overline{D}$ of G contains at least one edge. If $G_2 = G \cap (\Sigma \setminus D)$, then also G_2 is nontrivial by the same reason. Since G_1 is contained in a disk, it only contains contractible circuits, so the decomposition $G = G_1 \cup G_2$ is a planar separation which contradicts (iii).

(iv) \Rightarrow (v): Pick $S \subseteq V(\mathcal{A})$. The angle graph \mathcal{A} being bipartite, so is the subgraph $\mathcal{A}(S)$ of \mathcal{A} induced on S . To prove the inequality $2|S| - |E(S)| \geq 2\chi(\Sigma)$, it suffices to consider the case when $\mathcal{A}(S)$ is connected and when $|S| \geq 3$ since $\chi(\Sigma) \leq 0$. (Let us note that condition $\chi(\Sigma) \leq 0$ is really needed only in cases when $\mathcal{A}(S)$ is just a vertex or an edge.) Consider the induced embedding of $\mathcal{A}(S)$ in Σ . By (iv), and our assumptions on connectivity and $|S| \geq 3$, $\mathcal{A}(S)$ contains no digonal faces. Then it is easy to see by using the Euler's formula and the bipartiteness of $\mathcal{A}(S)$, that we get the required inequality.

(v) \Rightarrow (ii): Suppose that a proper closed subwalk Q of a facial walk f of M is contractible on Σ . Let x be the endvertex of $Q \subset f$, and let α_1, α_2 be the angles in f corresponding to the appearances of x as an end of Q . Let e_1, e_2 be the edges in \mathcal{A} corresponding to α_1 and α_2 , respectively. Note that e_1 and e_2 are parallel edges, joining x and f in \mathcal{A} . They determine a closed curve C on Σ which is homotopic to Q since it can be "pushed" by a homotopy within f onto Q . Since Q is contractible, so is C , and thus C

bounds an open disk D in Σ . Let S be the set of vertices of \mathcal{A} that do not lie in D , and consider the corresponding submap $\mathcal{A}(S)$ of \mathcal{A} . Since D is a disk, the underlying surface of $\mathcal{A}(S)$ is also Σ . The faces of $\mathcal{A}(S)$ coincide with faces of \mathcal{A} except for the digonal face D . Thus $2|E(S)|$ equals four times the number of faces of $\mathcal{A}(S)$ minus 2 (due to the digon). Euler's formula then shows that $2|S| - |E(S)| = 2\chi(\Sigma) - 1$. This contradicts (v). \square

Note that, since a map M and its dual map M^* have the same angle map, if one of M, M^* is essentially 2-connected, so is the other. The essential 2-connectivity of toroidal maps appears to be an appropriate generalization of the 2-connectivity planar maps.

Let M be an essentially 2-connected map on the torus. An edge e of M is *contractible* (resp. *removable*), if M/e (resp. $M - e$) is an essentially 2-connected map. Note that if e is contractible (resp. removable), then its dual edge e^* is removable (resp. contractible) in the dual map M^* .

Lemma 2.2 *Let M be an essentially 2-connected map on the torus and let $e \in E(M)$. Then:*

- (a) *If e is incident with a vertex of degree 2, then e is contractible.*
- (b) *If e is on a digonal face, then e is removable.*
- (c) *If M has no digonal faces and has at least two vertices, then M has a contractible edge.*
- (d) *If M has no vertices of degree 2 and has at least two faces, then M has a removable edge.*

Proof. Statements (a) and (b) are obvious, and (c) and (d) are dual to each other. So it suffices to prove (c). By (a) we may assume that M has no vertices of degree 2.

Since M has two or more vertices, there are non-loop edges in M . Consider e being one of them. If e is not contractible, its contraction results in a map that is not essentially 2-connected. According to Lemma 2.1(iv), the corresponding angle map contains a digon bounding a disk D . Suppose that the contracted edge e was chosen in such a way that the number of vertices in D is as small as possible. By our assumptions (no degree 2 vertices, no digons), D contains an edge of the map M . This edge is then contractible with respect to M since its contraction yielding a digon in the corresponding

angle map would contradict the minimality of the originally chosen edge e and the corresponding disk D . \square




Figure 5: Minimal essentially 2-connected maps M_0 and M_1

Lemma 2.2 shows that the only minimal essentially 2-connected map (with neither contractible nor removable edges) on the torus is the map M_0 of Figure 5. Lemma 2.2 also implies that every essentially 2-connected toroidal map M can be reduced to one of the two maps M_0 and M_1 from Figure 5 after a sequence of edge contractions and deletions of digon edges such that all intermediate maps are essentially 2-connected. Keeping track of that sequence, we may reconstruct M from the map M_0 or M_1 obtained at the end by using the sequence backwards: a deletion being replaced by the addition of a homotopic parallel edge, a contraction being replaced by a vertex splitting. In other words:

Corollary 2.3 *Every essentially 2-connected toroidal map can be obtained from one of the maps M_0 and M_1 by a sequence of homotopic parallel edge additions and vertex splittings such that all intermediate maps are essentially 2-connected.*

3 Tessellation representations

From now on we shall only consider maps on the torus. In this section we shall prove that a toroidal map admits a tessellation representation on the torus if and only if it is essentially 2-connected.

Suppose that we have a tessellation representation of a map M . This defines an orientation of edges of M : each edge e is oriented from the vertex

corresponding to the base of the rectangle R representing e towards the end of e corresponding to the upper side of R . For an arbitrary given orientation of the edges of a map M , an angle at a vertex v is *lateral* if one of its arcs is incoming and the other is outgoing at v . Otherwise, if both arcs of an angle are incoming or both are outgoing at v , then the angle is *extremal*. Having an orientation of edges of a map M , we get an orientation of the angle map \mathcal{A} as follows. Given an angle α incident with a vertex v and a face f of M , the angle α as an edge of \mathcal{A} is oriented from f to v if α is lateral, and from v to f if α is extremal. We say that an orientation of edges of M has the *upward property* at v (resp. at f) if there are exactly two lateral angles at v (resp. two extremal angles at f). Equivalently, in the oriented angle map \mathcal{A} , the indegree of every vertex (either v or f) is equal to two (see Figure 1(b)). Cf. also [11]. An orientation of M , which has the upward property at every vertex and every face, is called an *upward orientation*. By the above remark, an upward orientation induces an indegree-two orientation of the angle map \mathcal{A} . (It is an easy corollary of Euler's formula that only maps with Euler characteristic zero admit upward orientations.) Having indegree-two orientation of \mathcal{A} , we can use the following lemma:

Lemma 3.1 *Let k be an integer. A graph G admits an orientation of its edges such that each vertex has indegree at most k if and only if the average degree of the vertices of any subgraph of G is at most $2k$.*

Proof. Let H be an arbitrary subgraph of G and let $d^-(v)$ denote the indegree of $v \in V(G)$ with respect to an orientation Ω of edges of G . If $d^-(v) \leq k$ for every vertex v , then

$$|E(H)| \leq \sum_{v \in V(H)} d^-(v) \leq k|V(H)|,$$

so the average degree in H is at most $2k$.

Conversely, suppose that $|E(H)| \leq k|V(H)|$ for every subgraph H of G . Let Ω be an orientation of edges of G such that the number

$$S(\Omega) = \sum_{d^-(v) > k} (d^-(v) - k)$$

is minimal. If $S(\Omega) = 0$, we are done. So, assume that v_0 is a vertex with $d^-(v_0) > k$, and let H be the subgraph induced by the vertices v such that there exists a directed path from v to v_0 . Then

$$|E(H)| = \sum_{v \in V(H)} d^-(v) \leq k|V(H)|.$$

Since $d^-(v_0) > k$, there is a vertex $v_1 \in V(H)$ such that $d^-(v_1) < k$. By reversing the arcs on a directed path from v_1 to v_0 we get an orientation Ω' satisfying $S(\Omega') = S(\Omega) - 1$, a contradiction. \square




Figure 6: Tessellation representations of maps M_0 and M_1

Now we state the main result of this section.

Theorem 3.2 *A toroidal map admits a tessellation representation if and only if it is essentially 2-connected.*

Proof. Suppose that we have a tessellation representation of a map M . Then we easily get an upward drawing of M which induces an upward orientation of edges of M . This in turn defines an orientation of edges of the angle map \mathcal{A} such that every vertex of \mathcal{A} has indegree exactly 2 (cf. the introduction). By Lemma 3.1, every subgraph of \mathcal{A} has average degree at most 4. In our case this means that for every subset S of the vertices of \mathcal{A} we have

$$2|S| - |E(S)| \geq 0. \quad (2)$$

By Lemma 2.1(v), M is essentially 2-connected.

To prove the converse, we have to construct a tessellation representing an arbitrary given essentially 2-connected map M . This is done step by step as M can be obtained (according to Lemma 2.2) from M_0 and M_1 by a sequence of parallel edge additions and vertex splittings. We start with a tessellation representation of M_0 or M_1 shown in Figure 6.




Figure 7: Extending a tessellation after a vertex splitting

Let us now consider the two types of generating rules. Adding a parallel edge is simple: the rectangle corresponding to the edge e splits into two by adding a vertical line in the middle. Of course, the new vertical segment corresponds to the digonal face.

The vertex splitting operation is slightly more involved. Up to symmetries, we have to consider four cases as displayed in Figure 7. Note that in the first vertex-splitting rule, the orientation of the edge e depends on the relative position of the two faces involved in the splitting. These rules give the procedure how to obtain a tessellation representation of M starting from M_0 or M_1 . The proof is complete. \square

The proof of Theorem 3.2 also yields a polynomial time algorithm for constructing tessellation representations of essentially 2-connected toroidal maps. Since a tessellation representation determines an upward orientation, we also get a polynomial time algorithm for upward orientations.

There are other possibilities for tessellation representations of M_0 and M_1 . However, in none of them is any direction Δ or Δ' of the tessellating rectangles parallel to the sides of the fundamental parallelogram Q of the flat torus. It would be interesting to know which toroidal maps have a tessellation representation in the square model of the torus with the directions Δ and Δ' parallel to the sides of the fundamental square.

4 Some corollaries on graph drawing

4.1 Upward drawing

A tessellation representation of a toroidal map M determines an upward drawing and an upward orientation of M . In proving that the maps which admit tessellation representations are essentially 2-connected, we used Lemma 3.1 where only their upward orientation was needed. Therefore we also have:

Corollary 4.1 *A toroidal map admits an upward orientation if and only if it is essentially 2-connected.*

4.2 Visibility representations

Let M be an essentially 2-connected toroidal map. By Theorem 3.2 it has a tessellation representation. Let us consider the horizontal segments H_v , $v \in V(M)$. A vertex v is adjacent to a vertex u in M if and only if H_v contains a segment S of positive length such that by shifting S in the vertical direction Δ' we bump into H_u (or vice versa). We say that the segment H_u is ε -visible from H_v . The intervals H_v , $v \in V(M)$, thus uniquely determine M . Such a representation of M is called an ε -visibility representation of M . For the ε -visibility it is not important whether H_v are open or closed segments. If we take as H_v half-open segments (without their left endpoints, say), then the tessellation of M yields a visibility representation for M where a point of H_v can see a point of H_u in the vertical direction if and only if $uv \in E(M)$. Such notion of visibility (in the plane) was considered by Melnikov [8]. Tamassia and Tollis [13] and independently Wismath [15] characterized which graphs have such a visibility representation in the plane. Corollary 4.3 below gives a toroidal analogue of that result.

Tessellation representation is a stronger concept than ε -visibility representation since it gives simultaneous ε -visibility representations for M and

its dual map M^* . However, our next lemma shows that an ε -visibility representation of a map M can be transformed into a tessellation representation of M .

Lemma 4.2 *Every ε -visibility representation of a toroidal map M without digonal faces can be transformed into a tessellation representation of M .*

Proof. Let H_v , $v \in V(M)$, be the horizontal segments of an ε -visibility representation of M . For each edge $uv \in E(M)$, let Q_{uv} be a quadrangle determined by ε -visibility of H_u and H_v . The base of Q_{uv} is a segment $S(u, v) \subseteq H_u$ (say), and the upper side of Q_{uv} is a segment $S(v, u) \subseteq H_v$. Suppose that the quadrangles Q_{uv} are chosen such that the segments $S(u, v)$ and $S(v, u)$ are maximal. Now, the vertical sides of the quadrangles Q_{uv} , $uv \in E(M)$, give rise to vertical segments in the torus. None of these segments can become a closed circle since what we started with was an ε -visibility representation of a 2-cell embedded graph. Now it is easy to see that we got a tessellation representation of M . \square

Theorem 3.2 and Lemma 4.2 yield:

Corollary 4.3 *A toroidal map admits an ε -visibility representation in the flat torus (or has a visibility representation with half-open horizontal segments) if and only if it is essentially 2-connected.*

4.3 Toroidal grid contact graphs

Two families of horizontal and vertical segments of the grid of the flat torus, each one being disjoint from the others except for some contact points between two segments of different families, define a bipartite toroidal graph, called a *grid contact graph*. Relying on our tessellation theorem we get a characterization of the graphs which can be represented as toroidal grid contact graphs on the flat torus.

Theorem 4.4 *A graph H can be represented as the contact graph of straight line segments on the flat torus where all segments corresponding to the same bipartition class of H are mutually parallel if and only if H is a bipartite graph that can be embedded in the torus in such a way that no pair of parallel edges bounds a disk.*

Proof. A toroidal grid contact graph representation of H on the torus determines an embedding in the torus without parallel edges bounding discs. Conversely, let us embed H in the torus without homotopic parallel edges. By repeatedly adding paths of length 2 or 3 between vertices of the graph it is possible to get an embedded graph \tilde{H} such that:

- (a) H is an induced subgraph of \tilde{H} .
- (b) \tilde{H} is bipartite.
- (c) \tilde{H} is 2-cell embedded (and thus also connected).
- (d) \tilde{H} has no homotopic parallel edges.
- (e) No face of \tilde{H} has repeated vertices.

Now, beside the white bipartition class W and the black class B of vertices of \tilde{H} , consider a third class R of red vertices: add a new vertex $r \in R$ in every face f of \tilde{H} , and join r by red edges to every white vertex of f . The resulting map \overline{H} has bipartition $(W, B \cup R)$ and also satisfies (a)–(e). By Lemma 2.1 (equivalence of conditions (i) and (iv)), \overline{H} is the angle graph of an essentially 2-connected map G , and by Theorem 3.2 we can get a tessellation representation of G . It is clear that the tessellation representation yields the required grid contact representation of \overline{H} . By (a), the restriction to the segments corresponding to vertices of H yields the required representation for H . \square

It is worth pointing out that Theorem 4.4 is closely related to results of Bellantoni et al. [2] who considered the grid dimension of graphs. By taking a parallel to results of [2], Theorem 4.4 characterizes graphs whose “toroidal grid dimension” is at most two.

5 Graphs of essentially 2-connected maps

It is of some interest to know which abstract graphs can be obtained as graphs of essentially 2-connected toroidal maps. It is easy to find some sufficient conditions for a graph G to have a representation as an essentially 2-connected toroidal map. For example, if G is 2-connected, then every 2-cell embedding of G in the torus yields an essentially 2-connected map. We claim that every 2-connected graph (except graphs $C_n, n \geq 3$) with genus at most 1 has a 2-cell embedding in the torus. This is clear if the graph G has genus 1. On the other hand, if G is planar, consider one of its plane embeddings. Let $e = uv$ be an arbitrary edge of G with one of its endvertices, say u , having degree at least 3. Such an edge exists if and

only if G is not a cycle. Let F_1, F_2 be the two faces in the plane containing e on the boundary. By replacing the unbounded face by a disk and adding a “handle” joining F_1 and F_2 we get an embedding of G in the torus. If $f = uw \in \partial F_1$ is an edge at u distinct from e , we can re-embed it by using the handle so that it attaches u from “inside” of F_2 . It is clear that this gives rise to a 2-cell embedding of G in the torus.

Recall that a graph is *non-separable* if it is either a vertex, a loop, a bond ($p \geq 1$ parallel edges), or a 2-connected graph.

Let G_1, \dots, G_k ($k \geq 2$) be non-separable graphs, each containing at least one edge and such that at least one of G_i ($1 \leq i \leq k$) is not isomorphic to K_2 . A graph G is a *cyclic amalgamation* of G_1, \dots, G_k if there are vertices $u_i, v_i \in V(G_i)$ (possibly $u_i = v_i$ if $G_i \neq K_2$), $i = 1, \dots, k$, and G is isomorphic to the graph obtained from the (disjoint) union of G_1, \dots, G_k after identifying u_i with v_{i+1} (index modulo k), $i = 1, \dots, k$.

Theorem 5.1 *A graph G can be represented as an essentially 2-connected toroidal map if and only if it satisfies one of the following conditions:*

- (a) *G is 2-connected, it has an embedding in the torus, and it is not a cycle C_n for some $n \geq 3$.*
- (b) *G is a cyclic amalgamation of non-separable planar graphs G_1, \dots, G_k where $k \geq 2$, $G_1 \neq K_2$, and each of G_1, \dots, G_k contains at least one edge.*

Proof. We have already demonstrated sufficiency of (a). Let us now prove that graphs satisfying (b) have an essentially 2-connected toroidal representation. For $i = 1, \dots, k$ where $G_i \neq K_2$, embed G_i in the (closed) cylinder Q_i in such a way that ∂Q_i is in two distinct faces, say F_1, F_2 , $G_i \cap \partial Q_i = \emptyset$, and $u_i \in \partial F_1, v_i \in \partial F_2$. Such embeddings always exist. For that purpose take distinct faces F'_1, F'_2 of a planar embedding of G_i such that $u_i \in \partial F'_1$ and $v_i \in \partial F'_2$. Since G_i is non-separable and $G_i \neq K_1, K_2$, such faces exist. After removing an open disk from each of F'_1 and F'_2 , we get a required cylinder embedding. By cyclically identifying boundaries of these cylinders we get a torus embedding of $G_1 \cup \dots \cup G_k$ with vertices that need to be identified in common faces. If at least one pair u_i, v_i are distinct vertices, then it is obvious that we can get a 2-cell embedding of G in the torus (after adding the remaining graphs G_j that are isomorphic to K_2 , and making the appropriate identifications of vertices). If $u_1 = v_1, u_2 = v_2, \dots, u_k = v_k$,




Figure 8: The case of a single cutvertex

then we get a 2-cell embedding as indicated in Figure 8. It is easy to see by using Lemma 2.1(ii) that the resulting maps are essentially 2-connected.

We will now show that conditions (a), (b) are also necessary. Let G be the graph of a given essentially 2-connected map. If G is 2-connected, we have (a). If G is not 2-connected, let B_0, \dots, B_a ($a > 1$) be its blocks. Since the map is essentially 2-connected, each endblock B_i contains a cycle C_i that is non-contractible on the surface.

Suppose first that G has two distinct cutvertices. Then distinct cutvertices v, v' of G can be chosen such that an endblock, say B_1 , contains v and another endblock, say B_2 , contains v' . Then C_1 and C_2 are disjoint non-contractible cycles. Therefore they cut the torus into two cylinders, say Q, Q' . Any other block of G intersects $C_1 \cup C_2$ only at v or v' . Since the map is 2-cell, there is a path $P \subseteq G$ in Q joining v and v' . There is a similar path P' in Q' . Let B_0 be the block of G containing the cycle $C_0 = P \cup P'$. Now, each remaining block of G is embedded entirely in Q or in Q' . Since each endblock B_i ($i \geq 3$) contains a noncontractible cycle, it must intersect C_0 in a cutvertex of G . Now it is easy to see that G fits case (b) with each of the graphs G_1, \dots, G_k being either one of B_1, \dots, B_a , or a part of B_0 “between” consecutive vertices of G on $P \cup P'$ that are either cutvertices of G or cutvertices of $B_0 \cap Q$ or cutvertices of $B_0 \cap Q'$.

The remaining possibility is when G has exactly one cutvertex v . Then all blocks are endblocks. If one of them contains a non-contractible cycle which does not pass through v , then the proof is similar to the above. Thus, we may assume that every non-contractible cycle of G passes through v . Then $G - v$ consists of several components corresponding to the blocks of G . Each of such components D is plane embedded (contains only contractible circuits) and all edges from D to v attach D at the boundary of the “outer”

face. Therefore all blocks of G are planar graphs. Their identification at v can also be interpreted as a special case of a cyclic amalgamation ($u_i = v_i$ for every i). Hence we have (b). \square

It is worth mentioning that recognizing graphs satisfying (a) or (b) can be done in linear time. For (a), testing 2-connectivity is easy by a depth-first search [1], while checking if G has genus at most 1 can be performed in linear time by a recent algorithm of Juvan, Marinček, and Mohar [5]. To test if G satisfies (b) we first determine all blocks of G and test their planarity. There must be a block B_0 containing all cutvertices (otherwise (b) is not satisfied) and such that all blocks distinct from B_0 are planar. Blocks distinct from B_0 will appear in the cyclic amalgamation with $u_i = v_i$, while B_0 itself is a cyclic amalgamation of (at least 1) planar graphs. If there are 2 or more cutvertices, it is easy to see how to get the corresponding decomposition. Having just one cutvertex, we simply apply the algorithm of [5].

Acknowledgement. We are greatly indebted to Jan Kratochvíl and Hubert de Fraysseix for pointing out a possibility in applying Theorem 3.2 towards Theorem 4.4.

References

- [1] A. V. Aho, J. E. Hopcroft, J. D. Ullman, The design and analysis of computer algorithms, Addison-Wesley, 1976.
- [2] S. Bellantoni, I. B.-A. Hartman, T. Przytycka, S. Whitesides, Grid intersection graphs and boxicity, *Discrete Math.* 114 (1993) 41–49.
- [3] H. de Fraysseix, P. O. de Mendez, J. Pach, Representation of planar graphs by segments, in “Intuitive Geometry” (K. Boroczky, G. Fejes Toth, Eds.), North-Holland, Amsterdam, 1994, pp. 109–117.
- [4] J. L. Gross, T. W. Tucker, *Topological Graph Theory*, John Wiley and Sons, 1987.
- [5] M. Juvan, J. Marinček, B. Mohar, Embedding a graph into the torus in linear time, submitted.
- [6] J. Kratochvíl, T. Przytycka, Grid intersection and box intersection graphs on surfaces (Extended Abstract), in “Graph Drawing” (F. J. Brandenburg Ed.), Springer-Verlag, Berlin, 1996, pp. 365–372.

- [7] W. S. Massey, Algebraic Topology: An Introduction, Harcourt, Brace and World, 1967.
- [8] L. A. Melnikov, Problem at the 6th Hungarian Colloquium on Combinatorics (1981).
- [9] B. Mohar, P. Rosenstiehl, A flow approach to upward drawings of toroidal maps, in “Graph Drawing”, R. Tamassia, I. G. Tollis (Eds.), Lect. Notes in Computer Science, Vol. 894, Springer, 1995, pp. 33–39.
- [10] B. Mohar, C. Thomassen, Graphs on Surfaces, Johns Hopkins Univ. Press, to appear.
- [11] P. Ossona de Mendez, Orientations bipolaires, thèse de doctorat, E.H.E.S.S., Paris, 1994.
- [12] P. Rosenstiehl, Embedding in the plane with orientation constraints: The angle graph, Annals New York Academy of Sciences, 1989.
- [13] R. Tamassia, I. G. Tollis, A unified approach to visibility representations of planar graphs, Discr. Comput. Geom. 1 (1986) 321–341.
- [14] R. Tamassia, I. G. Tollis, Tessellation representation of planar graphs, in Proc. Twenty-Seventh Annual Allerton Conference on Communication, Control and Computing (1989) 48–57.
- [15] S. K. Wismath, Characterizing bar line-of-sight graphs, Proc. ACM Symp. on Comput. Geometry, Baltimore, Maryland (1985) 147–152.