Map genus, forbidden maps, and monadic second-order logic

B. Courcelle and V. Dussaux
Laboratoire Bordelais de Recherche en Informatique - UMR 5800
Université Bordeaux I
351, cours de la Libération
F-33405 Talence, France
{Bruno.Courcelle|Valere.Dussaux}@labri.fr

Submitted: January 12, 2001; Accepted: June 8, 2002.
MR Subject Classifications: 05C10, 03B70

Abstract

A map is a graph equipped with a circular order of edges around each vertex. These circular orders represent local planar embeddings. The genus of a map is the minimal genus of an orientable surface in which it can be embedded. The maps of genus at most g are characterized by finitely many forbidden maps, relatively to an appropriate ordering related to the minor ordering of graphs. This yields a "noninformative" characterization of these maps, that is expressible in monadic second-order logic. We give another one, which is more informative in the sense that it specifies the relevant surface embedding, in addition to stating its existence.

Introduction

A graph is a relational structure consisting of a domain which is the set of vertices and a binary "edge-relation". Hence logical formulas written with a binary relation symbol are formal writings of graph properties. For any fixed k, that a graph has degree at most k is easily expressible by a first-order formula. However, first-order logic is weak as a logical language for expressing graph properties. It cannot express a basic property like connectivity.

Second-order logic, its extension with new variables denoting relations and subject to quantifications is much more powerful: most graph properties can be expressed by second-order formulas.

Monadic second-order logic lies between first-order logic and second-order logic. It uses set variables but no variables denoting binary relations or relations of larger arity.

In this language, one can express vertex-colorability properties, path properties, minor inclusion. Hence in particular, by using Kuratowski's theorem one can express that a graph is planar. By using its extension to surfaces by Robertson and Seymour[19, 18] one can express in monadic second-order logic that a graph is embeddable in a given surface.

Monadic second-order logic is especially interesting because every property expressed in this language is evaluable in linear time on graphs, the tree-width of which is bounded by a fixed integer. Thus these properties are fixed-parameter tractable in the sense of Downey and Fellows[11] with tree-width as parameter.

Furthermore, for the study of context-free graph grammars, monadic second-order logic is an essential tool, like finite-state automata for context-free string grammars[5]. There is actually an equivalence between finite-state automata over finite and infinite words and trees established by Rabin[17], which yields many decidability results over infinite structures (surveyed by Gurevich[13]).

Going back to finite graphs of bounded tree-width, the linear evaluability of monadic second-order expressible graph properties extends to monadic second-order expressible optimization functions (like the maximal size of a planar subgraph of a given graph) or counting functions (like the number of paths between two distinguished vertices). See [4, 2, 8, 7].

For all these reasons, it is useful to express graph properties and graph evaluation functions in monadic second-order logic. This requires some efforts in certain cases: we have mentioned above the theorems of Kuratowski and Robertson and Seymour. For another example, the validity of the Strong Perfect Graph conjecture implies that perfectness is monadic second-order expressible. From the definition of perfectness, we only obtain that it is second-order expressible.

We are interested to applying these ideas and tools to other combinatorial structures than graphs. We consider here maps, which represent embeddings of graphs in orientable surfaces. For the purpose of having logical characterizations of embeddability in surfaces we develop a notion of "map-minor" aiming at results similar to those that are known for graphs.

A map is a graph equipped with a circular order of edges around each vertex. These circular orders represent local planar embeddings. The genus of a map is the minimal genus of an orientable surface in which it can be embedded so as to respect the local planar embeddings. A connected map of genus g can be embedded in any surface of genus greater than or equal to g. However the surface of genus g is the unique surface for which the embedding is a two cell embedding[1]: this means that the connected components of the complement of the graph (*i.e.* the faces) are simply connected domains (*i.e.* homeomorphic to a disc). Here an embedding is called *proper* if it is a two cell embedding (any nonproper embedding contains at least a face which is not simply connected. Faces of a nonproper embedding can have a disconnected boundary).

The maps of genus at most g are characterized by finitely many forbidden maps, relatively to an appropriate minor ordering. There exists also a similar statement with the corresponding notion of "topological minor".

[1]This is actually the usual (and more natural) definition of the genus for maps: the genus of a connected map is the genus of the (unique) surface which it tesselates (*i.e.* in which it can be embedded as a two cell embedding).

Robertson and Seymour proved that minor inclusion is a well-quasi order on the set of graphs [20] and the finiteness of the set of minimal forbidden minors for embeddability in a surface follows (it follows actually of a subcase of the Graph minor Theorem, see [19]). We do not know whether the set of maps, even the set of those of fixed genus, is well-quasi ordered for minor inclusion.

In the case of maps of genus at most g the situation is simpler than for graphs because the forbidden minor-maps have only one vertex. They are thus easier to construct than the forbidden minors for graphs of genus at most g. Furthermore, one-vertex maps can be represented by words with two occurrences of each letters (and one letter for each edge), and we will exploit this fact for our logical characterization.

The characterizations of planarity (or embeddability) of graphs and maps by forbidden minors (that are known or that we will obtain in Section 2) are "noninformative" in that, when they hold, they say nothing about embeddings in the considered surfaces. They only guarantee the existence of an embedding, by the nonexistence of a witness of impossibility. We are interested, especially for expressibility in monadic second-order logic by "informative" characterizations that also encode embeddings of the considered graphs. Such an informative monadic second-order expressibility has been established by Courcelle [6] for 3-connected planar graphs and for ordered planar graphs. We consider the same problem for connected maps. By contracting the edges of a spanning tree, we get a one vertex map (*i.e.* a set of loops incident to one vertex) of same genus as the considered graph. In Section 3 we will see that these maps can be represented by certain circular words; a Noetherian and confluent reduction system is given such that each reduction preserves exactly the genus. There are finitely many words in normal form representing one vertex maps of fixed genus. Each of them has an embedding that one can describe in a logical way. Last, we will prove in Section 4 that this description can be transfered to the given graph and yields the desired "informative" characterization by a monadic second-order formula.

1 Preliminaries

1.1 Graphs and maps

All graphs will be undirected and finite. They may have multiple edges and loops. For a graph G, we will denote by V_G its set of vertices and by E_G its set of edges. We define $\underline{\mathrm{vert}}_G(e)$ as being the set of endvertices of (or vertices incident to) the edge $e \in E_G$. This set has cardinality 1 if e is a loop, 2 otherwise.

The book of Mohar and Thomassen [16] will be our reference for definitions concerning surfaces. Let G be a graph properly embedded into a 2-dimensional compact oriented surface Σ simply called a *surface* in the paper. There corresponds a *map* M to this embedding. Let us associate two *darts* (or half-edges) e_1 and e_2 with each edge $e \in E_G$. Formally, if e is a loop on x we let $e_1 = (e, 1)$ and $e_2 = (e, 2)$. These darts are both incident with x. If e is not a loop, if x and y are its incident vertices, we let (e, x) and (e, y) be the two darts respectively incident with x and y (it does not matter which is e_1).

We denote by D_M the set of darts of M. We let α_M maps e_i to e_{3-i} for $i = 1, 2$, $e \in E_G$. We let σ_M associate with a dart incident with a vertex x, the next dart incident with x, where next is relative to a sweep of the surface around x in the direction defined by the orientation (We have $\sigma_M(d) = d$ if x has degree one and d is the unique dart incident to it). As we consider a proper embedding, the genus of M is also the genus of Σ and we call M a Σ-*map*. Hence $M =< D_M, \alpha_M, \sigma_M >$ where D_M is a finite set, α_M, σ_M are permutations, $\alpha_M(d) \neq d$ and $\alpha_M^2 = \underline{Id}$ for all d. Any triple satisfying these conditions is called a *map*. The corresponding graph is G defined as follows:

- $V_G = D_M / \sim_\sigma$,
- $E_G = D_M / \sim_\alpha$,
- $\underline{\text{vert}}_G([d]_{\sim_\alpha}) = \{[d]_{\sim_\sigma}, [\alpha_M(d)]_{\sim_\sigma}\}$

where $d \sim_\sigma d' \Leftrightarrow d' = \sigma_M^n(d)$ for some n (similarly for α).

We denote G by $G(M)$. We say that M is connected if $G(M)$ is. For every connected map M, there exists a surface Σ and a proper embedding of G in Σ such that the corresponding map is M.

Let us consider two maps $M =< D_M, \alpha_M, \sigma_M >$ and $M' =< D_{M'}, \alpha_{M'}, \sigma_{M'} >$. An *isomorphism* of M onto M' is a bijection γ of D_M onto $D_{M'}$ with the following property: If $\gamma(d) = d'$ where $d \in D_M$ and $d' \in D_{M'}$ then also

$$\gamma(\sigma_M(d)) = \sigma_{M'}(d') \text{ and } \gamma(\alpha_M(d)) = \alpha_{M'}(d').$$

If such an isomorphism exists we say that M and M' are *isomorphic* and we write $M \equiv M'$.

Conversely, any two proper embeddings of a connected graph G into a surface Σ having isomorphic maps are homeomorphic. The reader is referred to [16] Theorem 3.2.4 for the proof.

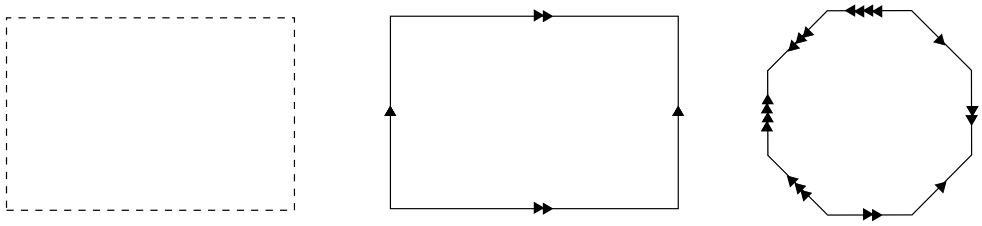

Figure 1: The sphere, the torus, and the torus with 2 holes

We recall (see [15]) that a surface can be represented by a polygon with $4n$ sides, such that the $4n$ sides are organised in $2n$ pairs, the two sides of a pair have equal length and furthermore each side is given a direction, such that two paired sides have opposite directions with respect to a cyclic traversal of the polygon. The surface is obtained by identifying the paired sides while respecting the directions. The polygon is called a *polygonal representation* of the surface. Examples of simple surfaces are shown on Figure 1 (The pairing are represented by identical arrows).

The polygon in the middle of Figure 1 represents the torus. For another example the torus with 2 holes can be defined from the polygon of right of Figure 1 although there are other polygonal representations of this surface.

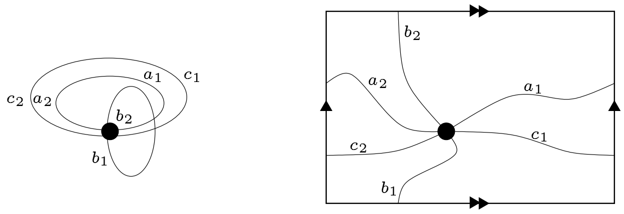

Figure 2: A map represented on the torus

A map on the torus like that of the left side of Figure 2 can be represented on the corresponding polygon as on right of Figure 2.

1.2 Submaps and minors

For two maps M and M', we let $M \subseteq M'$ iff

- $D_M \subseteq D_{M'}$,
- α_M is the restriction of $\alpha_{M'}$ to D_M.
- for $d \in D_M$, $\sigma_M(d) = \sigma^n_{M'}(d)$ where n is the smallest positive integer such that $\sigma^n_{M'}(d) \in D_M$.

If $M \subseteq M'$ then $G(M) \subseteq G(M')$ *i.e.* $G(M)$ is a subgraph of $G(M')$. If M' represents an embedding of $G(M')$ into a surface Σ and $M \subseteq M'$ then M represents the embedding of $G(M)$ into Σ obtained by deleting some curve segments representing edges. Conversely, if M' represents an embedding of $G(M')$ in Σ and $G \subseteq G(M')$ then there is a submap M of M' representing the induced embedding of G in Σ. However the subgraph relation does not preserve the genus and the condition $M \subseteq M'$ does not imply that M represents a proper embedding of $G(M)$ in Σ if M' represents a proper embedding of $G(M')$ in Σ. Indeed, a proper embedding of M may take place into a surface of smaller genus than Σ.

Let $M =< D, \alpha, \sigma >$ and let $X \subseteq D$. We say that X is *α-closed* if $\alpha(X) \subseteq X$. This means that X is the set of all darts associated with a set Y of edges of the graph $G(M)$. The submap M' of M induced by X, is denoted by $M[X]$ and is defined as $< X, \alpha', \sigma' >$ where α' is the restriction of α to X, $\sigma'(x) = \sigma^i(x)$ where i is the smallest $i > 0$ such that $\sigma^i(x) \in X$. Every submap $N =< D', \alpha', \sigma' >$ of a map M is equal to $M[D']$ (Clearly, D' is α-closed in M).

The transformation of M' into $M \subseteq M'$ can be intuitively described as the result of a sequence of *deletions of edges* and of *isolated vertices*. The deleted edges are those of $G(M')$ that are not in $G(M)$ and isolated vertices are systematically removed.

We now define a notion of *edge contraction* for maps. Let M be a map, let $d \in D_M$ and $d' = \alpha_M(d)$. Hence d and d' form an edge e of $G(M)$. There are two cases:

- e is a loop of $G(M)$.
  This loop is said to be *contractible* iff either $d = \sigma_M(d')$ or $d' = \sigma_M(d)$. The result of the contraction of e is the submap M' of M such that $D_{M'} = D_M - \{d, d'\}$. The effect is the same as deleting e. If e is not contractible, then M' is undefined.

- e is not a loop.
  It can be contracted and the result of the contraction is M' such that:

  - $D_{M'} = D_M - \{d, d'\}$,
  - $\alpha_{M'}$ is the restriction of α_M to $D_{M'}$,
  - $\sigma_{M'}$ is defined as follows ($x \notin \{d, d'\}$):
    $\sigma_{M'}(x) = \sigma_M(x)$ if $\sigma_M(x) \notin \{d, d'\}$.
    $\sigma_{M'}(x) = \sigma^2_M(x)$ if $\sigma_M(x) \in \{d, d'\}$ and e is a pending edge.
    $\sigma_{M'}(x) = \sigma_M(\alpha_M(\sigma_M(x)))$ if $\sigma_M(x) \in \{d, d'\}$ and e is not a pending edge.

These cases are illustrated by Figures 3, 4 and 5. Contracting a pending edge (an edge such that one of its endvertices has degree one) is the same as deleting it. In Figure 5 if a and d had already a vertex in common, they yield multiple edges.

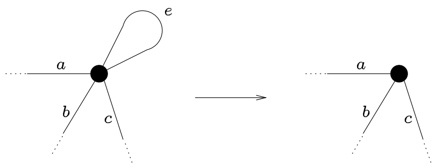

Figure 3: e is contractible

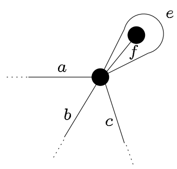

Figure 4: e is not contractible

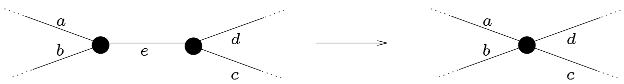

Figure 5: e is not pending

Remark 1.2.1. In all cases where e is contractible, it can be contracted "continuously", *i.e.* by progressive shrinking. In other terms, contraction of contractible edges preserves the genus and sends a proper embedding to a proper embedding. The noncontractible loop e of Figure 4 cannot be shrunk to its endvertex without shrinking also f.

We write $M \underset{e}{\rightarrow} M'$ if M' results from M by the contraction of an edge e. We write $M \rightarrow M'$ if we do not need to specify the contracted edge. We also write $M \xrightarrow{*} M'$ if M' results from M by a sequence of edge contractions.

Let M be a map, let e, e' be two contractible edges. Let $M \underset{e}{\rightarrow} M_1$ and $M \underset{e'}{\longrightarrow} M_2$. Then e' is contractible in M_1 and e is contractible in M_2, and we have a map M' such that $M_1 \underset{e'}{\longrightarrow} M'$ and $M_2 \underset{e}{\rightarrow} M'$. We say that e and e' can be *contracted simultaneously.*

This notion extends to a set of edges. If a set of edges of a map M forms a tree or a forest in $G(M)$, then the edges of this set can be contracted simultaneously. A set of edges is *contractible* if it can be ordered in such a way that it forms a sequence of contractible edges. In Figure 4 e is contractible after f, but not before.

We say that M' is a *minor* of M and we write $M' \trianglelefteq M$ if M' is isomorphic to a map obtained by edge contractions from a submap of M.

Lemma 1.2.1. *Let Σ be a surface of genus g and M, M' be maps.*

1. *If $M' \subseteq M$ and M is a Σ-map then M' is a map with genus at most g.*
2. *If $M \xrightarrow{*} M'$ then M is a Σ-map iff M' is a Σ-map.*

Proof. Assertion 1 is clear. Because 2-cell embeddings in S are preserved, Assertion 2 is clear. □

Remark 1.2.2. Note the "iff" in Assertion 2. It is not true for graphs that if G' results from G by an edge-contraction then G is planar iff G' is planar. We only have the "only if" direction for graphs.

We do not know whether the quasiorder $\trianglelefteq$ on maps is a well-quasi-order (as it is for graphs by the Graph minor Theorem of Robertson and Seymour [20]). However we will prove in Section 2 the following theorem:

Theorem 1.2.1. *For all $g \geq 0$, the set of maps of genus at most g is characterized by a finite set of forbidden minor-maps. In other words, there exists a finite set of maps $\{M_1, \ldots, M_k\}$ such that M has genus at most g iff $M_i \trianglelefteq M$ for no $i = 1, \ldots, k$.*

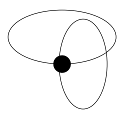

Figure 6: One forbidden minor-map for the sphere

Example 1.2.1. If we consider maps on the sphere (equivalently, on the plane), the corresponding list is reduced to the map $< \{a_1, a_2, b_1, b_2\}, \alpha, \sigma >$ shown on Figure 6 where σ is the circular permutation (a_1, b_1, a_2, b_2) and $\alpha(a_1) = a_2$, $\alpha(b_1) = b_2$.

This fact follows from Theorem 1.2.1 below.

1.3 Topological minors

We write $M \xrightarrow[T]{} M'$ if M' results from the contraction of an edge of M which is not a loop, is not pending and with at least one of its two end-vertices of degree 2. We say that M' is a *topological minor* of M if $M'' \xrightarrow[T]{*} M'_1$ for some submap M'' of M and M'_1 is isomorphic to M'. We write $M' \trianglelefteq_T M$ in such a case. This implies $G(M') \trianglelefteq_T G(M)$ where $\trianglelefteq_T$ is the quasi-order on graphs of topological minor inclusion (see [10] or [16]). For this quasi-order, planar maps are characterized by two forbidden maps, that of last example and that of Figure 7 (See [6] for a complete proof of this result). It is clear that

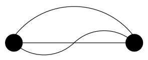

Figure 7: The other forbidden topological minor-map for the sphere

contracting one edge of the map of the Figure 7 yields the first one.

If $M \xrightarrow[T]{*} M'$ we say that M is obtained from M' by edge subdivisions, *i.e.* by substitution of disjoint paths for edges (see [10] for the related notion in graphs).

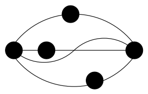

Figure 8: The nonplanar map L of Example 1.3.1

Example 1.3.1. Let L be the nonplanar map of Figure 8. We have $L \supseteq L' \xrightarrow[T]{*} K$ where K is the map of Figure 7. Hence L is nonplanar.

1.4 Parallel edges

Let M be a map. Let e, e' be edges with corresponding unordered pairs of darts $\{e_1, e_2\}$, $\{e'_1, e'_2\}$ such that

$$\sigma(e_1) = e'_1 \text{ and } \sigma(e'_2) = e_2$$

or

$$\sigma(e'_1) = e_1 \text{ and } \sigma(e_2) = e'_2.$$

We say that e and e' form a *pair of parallel edges*.

Lemma 1.4.1. *If e, e' form in M a pair of parallel edges, then M and $M - e'$ have the same genus.*

Proof. It is clear that every embedding of $M - e'$ in a surface Σ can be transformed into an embedding of M in Σ, by the addition of a curve "close to that representing e". And if M is embeddable in Σ then so is the submap $M - e'$. Hence M and $M - e'$ are embeddable in the same surfaces. $\square$

We say that E is a *set of parallel edges* in M if E can be enumerated as $e_1, e_2, \ldots, e_k$ such that each $\{e_i, e_{i+1}\}$, $1 \le i \le k-1$, forms a pair of parallel edges.

It is clear that for each i the maps M and $M - (E - \{e_i\})$ have the same genus. Intuitively $M - (E - \{e_i\})$ is obtained from M by *fusing a set of parallel edges*. See Section 3, Figure 13 for an example.

2 Genus of maps via forbidden submaps and minors

2.1 Forbidden submaps

In this subsection we study the topological minor inclusion on maps. We define for each $g \ge 0$ a finite set $\mathcal{F}_g$ of forbidden maps, and we use it to characterize maps of given genus.

Here *faces* always refer to the simply connected faces of the unique proper embedding.

Definition 2.1.1. Let g be a positive integer. The set of forbidden maps of genus g, denoted by $\mathcal{F}_g$ is the set of maps of genus g with exactly one face and without vertices of degree one or two. Let also $\mathcal{SF}_g$ be the set of subdivisions of maps of $\mathcal{F}_g$, *i.e.* maps of genus g with one face and without vertices of degree one.

Proposition 2.1.1. *For all g, $\mathcal{F}_g$ is finite.*

Proof. Let n and m be the respective numbers of vertices and edges of a map from $\mathcal{F}_g$. Euler's characteristic formula reads $n + 1 = m + 2 - 2g$. As all vertices have degree at least 3, we have $2m \ge 3n$. Therefore $n + 2g - 1 \ge 3n/2$, so that $n \le 4g$ and $m < 6g$. As the number of maps with n vertices and m edges is finite, so is $\mathcal{F}_g$. ☐

Lemma 2.1.1. *Any connected map M of genus g has a submap that belongs to $\mathcal{SF}_g$.*

Proof. Let E be the set of edges of M and $T \subseteq E$ be a minimal (for inclusion) set of edges across which all faces can be connected (*i.e.* T is a spanning tree of the adjacency graph of faces). As deletion in a map of an edge adjacent to 2 faces preserves its genus, then the submap M' of M whose edges are $E - T$ has exactly one face and genus g.

Vertices of degree one in M' can be recursively deleted with their incident edges, until a map of $\mathcal{SF}_g$ is obtained. ☐

Lemma 2.1.2. *Any connected map M of genus $g > 0$ has a submap that belongs to $\mathcal{SF}_{g-1}$.*

Proof. According to Lemma 2.1.1, the map M has a submap M' in $\mathcal{SF}_g$. As $g > 0$, M' contains an edge ℓ that is either a loop or belongs to a simple cycle with at least two edges (otherwise it is a tree and $g = 0$). Deleting ℓ yields a connected map M''. The map M' has only one face so that ℓ is incident twice to the same face. Deleting ℓ thus raises two faces in M''. With one more face, one edge less and the same number of vertices as M', M'' has genus $g - 1$ by Euler's formula. Then Lemma 2.1.1 asserts that M'' has a submap in $\mathcal{SF}_{g-1}$. ☐

From these two lemmas and the fact that the genus of a map is greater or equal to the genus of any of its submap, we deduce the following theorem:

Theorem 2.1.1. *A connected map has a genus at least $g+1$ iff it contains a submap in $\mathcal{SF}_{g+1}$.*

Equivalently a connected map has a genus at most g iff it contains no submap in $\mathcal{SF}_{g+1}$.

2.2 Forbidden minors of maps

In this subsection we consider the minor inclusion for maps.

Definition 2.2.1. Let g be a positive integer. The set of forbidden minors of genus g, denoted by $\mathcal{M}_g$ is the set of maps of genus g with exactly one face and one vertex.

Cori and Marcus have proved in [3], Prop. 6.3. that this set is finite and obtained an exact formula for its cardinality. The first values are 1 (for the sphere, see Figure 6), 4 (for the torus, see Figure 9), 131, 14118, ...

Lemma 2.2.1. *Any connected map M of genus g has a minor that belongs to $\mathcal{M}_g$.*

Proof. According to Lemma 2.1.1, M has a submap M' in $\mathcal{SF}_g$. Consider a spanning tree T of $G(M')$ and contract all edges of T to obtain a minor of M' denoted by M'/T. This contraction is possible as T is a tree, and the resulting map M'' has only one vertex as T is a spanning tree. Like M', M'' has only one face, so that M'' belongs to $\mathcal{M}_g$. $\square$

From this lemma and Lemma 2.1.2, we deduce the following Theorem:

Theorem 2.2.1. *A connected map has a genus at least $g+1$ iff it contains a minor in $\mathcal{M}_{g+1}$. Equivalently, a connected map has a genus at most g iff it contains no minor in $\mathcal{M}_{g+1}$.*

As $\mathcal{M}_{g+1}$ is finite, this achieve to prove Theorem 1.2.1.

Schaeffer has shown in [21] how one could construct the sets $\mathcal{M}_g$ and $\mathcal{F}_g$.

3 Words

A classical tool in combinatorics consists in defining a bijection between a set of combinatorial objects and a set of words. In good cases the corresponding words have a certain structure from which informations on the considered objects can be derived (typically the number of objects of a certain size). Here we encode one vertex maps by words (up to an equivalence relation).

Let A be a countable alphabet. Let $W \subseteq A^*$ be the set of (finite) words such that each letter has 0 or 2 occurrences. We let $\sim$ be the least equivalence relation on W such that:

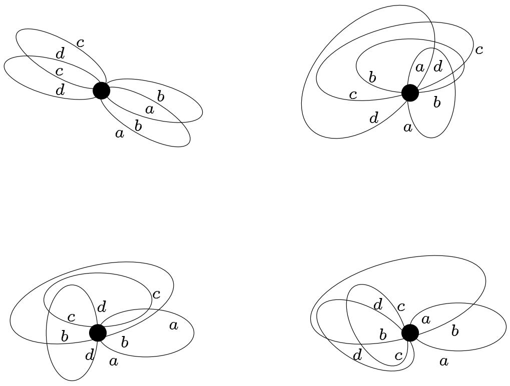

Figure 9: Forbidden maps for the minor relation for the torus

- $uv \sim vu$ for every $u, v \in A^*$, $uv \in W$ (we say that uv and vu are *conjugate* words).
- $u \sim h(u)$ if h is the homomorphism $A^* \rightarrow A^*$ extending an injection $i : A \rightarrow A$.

We show that $W/\sim$ is in bijection with the 1-vertex maps.

Formally, to $w \in W$ corresponds the map $< D, \alpha, \sigma >$ where D is the set of occurrences of letters, say $D = \{1, \ldots, |w|\}$, $\sigma(i) = i + 1$, $\sigma(|w|) = 1$ and $\alpha(i)$ is the unique $j \neq i$ such that i and j are occurrences of the same letter. If $w \sim w'$ then the corresponding maps are isomorphic. Hence we have a mapping of $W/\sim$ to maps (up to isomorphisms).

Conversely, let $< D, \alpha, \sigma >$ be a map. We can enumerate D as $\{1, \ldots, n\}$ in such a way that $\sigma(i) = i + 1$, $\sigma(n) = 1$. Then we define a word $a_1 a_2 \ldots a_n$ where $a_1, \ldots, a_n$ are letters such that $a_i = a_j$ iff $i = j$ or $\alpha(i) = j$ or $\alpha(j) = i$. There are several choices of letters but for any two such choices the corresponding words w and w' are equivalent. There are also several enumerations of D but the corresponding words are also equivalent. Hence $W/\sim$ is in bijection with the set of 1-vertex maps.

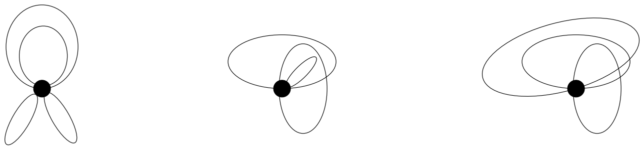

Figure 10: The maps of the Example 3.0.1

Example 3.0.1. The first map on the left of Figure 10 is represented by the equivalent words *aabbcddc*, *ccbbdaad* and *ddcaabbc*. The second, in the middle, is represented by the word *ababcc*. The third is represented by the word *abcabc*.

The genus of a word w in W is the genus of the corresponding map. We denote by W_g the set of words of genus g.

We introduce two sets of rewriting rules over $W/\sim$ that reduce the size of words and preserve the genus. We first define them on W.

- Set L - There are two rules:

$$\begin{array}{llcl} L_1: & w_1 a a w_2 & \xrightarrow[L]{} & w_1 w_2 \\ L_2: & a w_1 a & \xrightarrow[L]{} & w_1 \end{array}$$

for $w_1,\ w_2 \in A^*,\ a \in A$.

- Set P - There are four rules:

$$\begin{array}{llcl} P_1: & w_1 ab w_2 ba w_3 & \xrightarrow[P]{} & w_1 a w_2 a w_3 \\ P_2: & w_1 ab w_2 ba w_3 & \xrightarrow[P]{} & w_1 b w_2 b w_3 \\ P_3: & a w_4 ab w_5 b & \xrightarrow[P]{} & w_4 b w_5 b \\ P_4: & a w_4 ab w_5 b & \xrightarrow[P]{} & a w_4 a w_5 \end{array}$$

for $a,\ b \in A,\ (a \neq b),\ w_1,\ w_2,\ w_3,\ w_4,\ w_5 \in A^*$.

If $w \sim w'$ and $w \xrightarrow[P \cup L]{} w_1$ then $w' \xrightarrow[P \cup L]{} w'_1 \sim w_1$ for some w'_1. Hence these rules preserve $W/\sim$.

It follows also that if $w_1 \to w'_1 \sim w_2 \to w'_2 \sim w_3 \to \ldots \to w'_{n-1} \sim w_n$ then there exist $w''_1, w''_2, \ldots, w''_n$ such that $w_1 \to w''_1 \to w''_2 \to \ldots \to w''_{n-1} \sim w_n$ where $\to = \xrightarrow[P \cup L]{}$.

In terms of maps, Rules L_1 and L_2 consist in removing a contractible loop and Rules P_1 to P_4 consist in fusing two parallel edges. We deduce the following Lemma from this observation (and Subsections 1.2 and 1.4).

Lemma 3.0.2. *Rules L_1, L_2 and P_1 to P_4 preserve the genus of words in W.*

A word $w \in W$ is a *normal form* if neither a rule of L nor a rule of P is applicable. It follows that the set W^{nf} of words in normal form is stable under $\sim$. This set encodes the set of maps without contractible loops and without parallel edges. We say that $w' \in W^{nf}$ is a *normal form of* w if $w \xrightarrow[P \cup L]{*} w'$.

Figure 11: Two normal forms of respective lengths 6 and 8

Example 3.0.2. Up to length 6 the normal forms are ε (in W_0), *abab* and *abcabc* (in W_1, they correspond to maps on the torus. See Figure 11). With four letters, we obtain the different normal forms *abcdabcd*, *ababcdcd* and *abcdadbc*.

We let $W_g^{nf} = W_g \cap W^{nf}$.

Lemma 3.0.3. *Each set W_g^{nf} is finite up to $\sim$.*

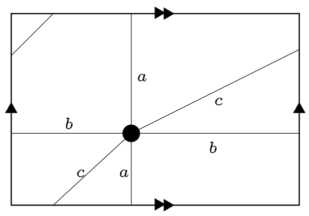

Figure 12: *abcabc* on the torus

Proof. Let w in normal form correspond to a 1-vertex map of genus g. This map has no contractible loop, no two parallel edges, otherwise w is not in normal form. If some face is not a triangle, we can add an edge, such that the corresponding map, still of genus g, has no contractible loop and no parallel edges. By adding edges we get a triangulated map for which $3|F| = 2|E|$ (where F denotes the set of faces and E the set of edges).

By Euler's relation we have $|F| = |E| - 2g + 1$ which gives $|E| = 6g - 3$. Hence w has at most $6g-3$ letters. There are finitely many such words, up to the choice of letters. $\square$

Proposition 3.0.1. *The relation* $\xrightarrow[P\cup L]{}$ *on* $W/\sim$ *is Noetherian and confluent. Hence, every* α *in* $W/\sim$ *has a unique normal form* β *in* $W/\sim$*.*

Proof. The relation $\rightarrow = \xrightarrow[P\cup L]{}$ is *Noetherian* which means that there is no infinite sequence

$$w_1 \rightarrow w_2 \rightarrow w_3 \rightarrow \dots$$

for w_1, w_2, $w_3 \in W/\sim$.

This is clear since each rule $\xrightarrow[L]{}$ or $\xrightarrow[P]{}$ decreases the size of words by 2.

The relation is *confluent* (on $W/\sim$) which means that for every words w_1, w_2, $w_3 \in W$, if $w_1 \xrightarrow{*} w_2$, $w_1 \xrightarrow{*} w_3$ then there exist w_4, $w_5 \in W$ with $w_2 \xrightarrow{*} w_4$, $w_3 \xrightarrow{*} w_5$ and $w_4 \sim w_5$. Since the relation is Noetherian it is enough to prove that it is *locally confluent* (see [9]) *i.e.* that $w_1 \rightarrow w_2$, $w_1 \rightarrow w_3$ implies the existence of w_4, w_5 such that $w_1 \xrightarrow{*} w_4$, $w_1 \xrightarrow{*} w_5$ and $w_4 \sim w_5$.

This follows from lengthy case study of all possibilities for $w_1 \rightarrow w_2$, $w_1 \rightarrow w_3$.

Let us show a few cases.

- First, we consider the case where $w_1 \rightarrow w_2$ (by Rule L_1) and $w_1 \rightarrow w_3$ (by Rule P_1). Hence we have

$$w_1 = w_1' aa w_2' \text{ and } w_2 = w_1' w_2'$$

and

$$w_1 = w_1'' bc w_2'' cb w_3'' \text{ and } w_3 = w_1'' b w_2'' b w_3''.$$

There are 3 possibilities for aa: it can be inside w_1'', w_2'' or w_3''.

Subcase 1 (aa is inside w_2''):

$$w_1 = saatbcw_2''cbw_3''$$

$$w_2 = stbcw_2''cbw_3''$$
$$w_3 = saatbw_2''bw_3''$$
$$w_4 = w_5 = stbw_2''bw_3''$$

The two other subcases are similar.

- We now consider the case where two rules of the set P apply.

  Among the various cases we have:

  – First:
  $$w_1 = w_1'abw_2'cdw_2''baw_3'dcw_3''$$
  $$w_2 = w_1'aw_2'cdw_2''aw_3'dcw_3''$$
  $$w_3 = w_1'abw_2'cw_2''baw_3'cw_3''$$
  $$w_4 = w_5 = w_1'aw_2'cw_2''aw_3'cw_3''$$

  – Here is another case:
  $$w_1 = aw_1'abcw_2'cb$$
  $$w_2 = aw_1'abw_2'b$$
  $$w_3 = w_1'bcw_2'cb$$
  $$w_4 = w_5 = w_1'bw_2'b$$

□

Lemma 3.0.4. *Let $w \in W$ and ε denote the empty word. We have $w \xrightarrow[L]{*} \varepsilon$ iff there do not exist two letters $a,\ b \in A$ such that*

$$w = w_1aw_2bw_3aw_4bw_5 \textit{ for any } w_1,\ w_2\ ,\ w_3,\ w_4,\ w_5 \in A^*. \tag{$*$}$$

Proof. *"If"*. By induction on $|w|$.

The cases $|w| = 0$, $|w| = 2$ are clear.

Let $|w| > 2$. Then $w = aw_1aw_2$. If $w_1 = \varepsilon$ then $w \xrightarrow[L]{} w_2$, w_2 satisfies the condition so $w_2 \xrightarrow[L]{*} \varepsilon$ by induction, thus $w \xrightarrow[L]{*} \varepsilon$. If $w_1 \notin W$ this means that some letter b has one occurrence in w_1 and one occurrence in w_2. So, the condition $(*)$ is violated. Otherwise $w_1 \in W - \{\varepsilon\}$. Then both w_1 and w_2 belong to W, satisfy Condition $(*)$ and we get the result by induction.

"Only if". Easy by induction on the length of rewriting sequence $w \xrightarrow[L]{*} \varepsilon$.

□

We recall that W_0 encodes 1-vertex planar (*i.e.* of genus 0) maps.

Lemma 3.0.5. *Let $w,\ w' \in W$.*

1. *We have $w \in W_0$ iff $w \xrightarrow[L]{*} \varepsilon$.*

2. *We have $w \xrightarrow[L]{*} w'$ iff $w' = a_1a_2\dots a_n$ for some $a_1, a_2 \dots, a_n \in A$ and $w = w_0a_1w_1a_2w_2\dots a_nw_n$ for some words $w_0 \dots w_n$ such that $w_1, w_2, \dots, w_{n-1} \in W_0$, and $w_nw_0 \in W_0$.*

Proof. 1. If $w \xrightarrow[L]{*} \varepsilon$ then w is of genus 0 since loop-contractions preserve genus. Hence $w \in W_0$. Conversely if $w \in W_0$ but $w \xrightarrow[L]{*} \varepsilon$ does not hold then w violates Condition $(*)$ hence the corresponding map contains a nonplanar submap (corresponding to the word $abab$), hence $w \notin W_0$.

2. *"If"*. Clear.

"Only if". By induction on length of a rewriting sequence $w \xrightarrow[L]{*} w'$.

Let us assume $w \xrightarrow[L]{n+1} w'$. Then $w \to w'' \xrightarrow[L]{n} w'$.

By induction hypotheses, $w \to w'' = w_0a_1w_1a_2w_2\dots a_pw_p \xrightarrow{n} w' = a_1a_2\dots a_p$ where $w_1, \dots, w_{p-1} \in W_0$ and $w_pw_0 \in W_0$.

Now $w = saat$ for $ts \in W_0$. We have two cases.

Case 1: $w = saat \xrightarrow[L]{} st = w'' = w_0a_1w_1a_2\dots w_{p-1}a_pw_p$.

We can have

$w_0 = ss'$ for some $s' \in A^*$

or $s = w_0$

or $s = w_0a_1$

or $s = w_0a_1w_1'$ for $w_1 = w_1'w_1''$

or ...

or $s = w_0a_1w_1a_2\dots w_{p-1}a_pw_p'$ for $w_p = w_p'w_p''$.

In each case we have a decomposition of w of the desired form because if

$$xy \in W_0 \cap (A - \{a\})^*$$

then $aaxy, xyaa, xaay \in W_0$.

Case 2: $w = asta \xrightarrow[L]{} st = w'' = w_0a_1w_1a_2\dots w_{p-1}a_pw_p$.

We get a decomposition of w of the desired form because, from

$$w_pw_0 \in W_0 \cap (A - \{a\})^*$$

we get that $w_paaw_0 \in W_0$.

□

Our aim is now to obtain a statement like that of Lemma 3.0.5 (2.) for describing $\xrightarrow[P]{*}$.

Let w, x be nonempty words in W. Let $x = a_1a_2 \dots a_{2n}$. We write $w \geq_P x$ iff for some $u'_1, v'_{2n}, u_1, u_2, \dots, u_{2n}, v_1, v_2, \dots, v_{2n} \in A^*$ we have:

(1) $w = u'_1 a_1 v_1 u_2 a_2 v_2 \dots u_{2n} a_{2n} v'_{2n}$,

(2) $v'_{2n} u'_1 = v_{2n} u_1$,

(3) $u_i a_i v_i = \tilde{v}_j a_j \tilde{u}_j$

for all i, j such that $i < j$ and $a_i = a_j$ ($\tilde{u}$ denotes the mirror image of u). We let also $\varepsilon \geq_P \varepsilon$.

Lemma 3.0.6. *(1) The relation $\geq_P$ is a partial order.*

(2) For $w, x \in W$, we have $w \xrightarrow[P]{*} x$ iff $w \geq_P x$.

Proof. *(1)* The verification that $\geq_P$ is transitive is lengthy. Reflexivity and antisymmetry are clear.

(2) It is clear from the definitions that $w \xrightarrow[P]{} x$ implies $w \geq_P x$. Hence $w \xrightarrow[P]{*} x$ implies $w \geq_P x$ by (1).

From the definition of $\geq_P$ we now verify that $w \geq_P x$ implies $w \xrightarrow[P]{*} x$.

Let $$w = u'_1 a_1 v_1 u_2 a_2 v_2 \dots u_{2n} a_{2n} v'_{2n}$$
where $v'_{2n} u'_1 = v_{2n} u_1$ for words $u'_1, v'_{2n}, u_1, u_2, \dots, u_{2n}, v_1, v_2, \dots, v_{2n}$ such that

$$u_i a_i v_i = \widetilde{v_j} a_j \widetilde{u_j} \text{ if } i < j \text{ and } a_i = a_j.$$

Case 1: $|v'_{2n}| > |v_{2n}|$.
Hence $v'_{2n} = v_{2n} u''_1$, $u_1 = u''_1 u'_1$ and $w = u'_1 a_1 v_1 u_2 a_2 v_2 \dots u_{2n} a_{2n} v_{2n} u''_1$.
Let j be the unique index such that $a_1 = a_j$. Hence

$$w = u'_1 a_1 v_1 \dots \widetilde{v_1} a_j \widetilde{u'_1} \widetilde{u''_1} \dots a_{2n} v_{2n} u''_1.$$

Hence we have (since $a_1 = a_j$)

$$\begin{aligned} w &\xrightarrow[P]{*} a_1 v_1 \dots \widetilde{v_1} a_j \widetilde{u''_1} u_{j+1} a_{j+1} \dots a_{2n} v_{2n} u''_1 && \text{(using Rule } P_1) \\ &\xrightarrow[P]{*} a_1 v_1 \dots \widetilde{v_1} a_j u_{j+1} a_{j+1} \dots a_{2n} v_{2n} && \text{(using Rule } P_4) \\ &\xrightarrow[P]{*} a_1 \dots a_{2n} && \text{(using Rules } P_1 \text{ and } P_2). \end{aligned}$$

The two other cases where $|v'_{2n}| < |v_{2n}|$ and $|v'_{2n}| = |u_{2n}|$ are similar.

□

In terms of the 1-vertex maps M and N represented respectively by w and x, condition (3) of the definition of $\geq_P$ means that each pair $(u_i a_i v_i, \tilde{v}_j a_j \tilde{u}_j)$ with $i < j$ and $a_i = a_j$ represents a set of parallel edges in M that get fused into the single edge (a_i, a_j) of N (See Subsection 1.4).

Example 3.0.3. Here is an example where $a_3 = a_1$, $a_4 = a_2$:

$$w = ca_1dea_2fgda_1cbgfa_2eb,$$

$$x = a_1a_2a_1a_2.$$

To w corresponds the map of Figure 13 and to x corresponds the map of Figure 6 with labelled edges a_1 and a_2.

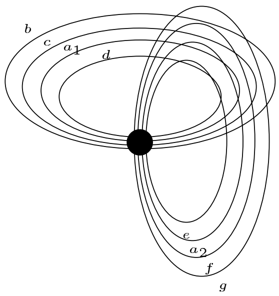

Figure 13: A map with two sets of parallel edges

Proposition 3.0.2. *Let $x \in W^{nf}$ be nonempty and $w \in W$. Then $w \xrightarrow[P\cup L]{*} x$ iff there exists w' such that $w \xrightarrow[L]{*} w' \geq_P x$.*

We will say that these words w are of type x.

Proof.

Fact 3.0.1. *For all words $y_1, y_2, y_3 \in W$ if $y_1 \xrightarrow[P]{} y_2 \xrightarrow[L]{} y_3$ there exists a word $y_4 \in W$ such that $y_1 \xrightarrow[L]{} y_4 \xrightarrow[P\cup L]{} y_3$.*

Proof of Fact 3.0.1. By case analysis as for Proposition 3.0.1. □

"If". Clear.

"Only if". By induction on the length n of a rewriting sequence $w \xrightarrow[L\cup P]{n} x$.

We have $w \xrightarrow[P\cup L]{} w_1 \xrightarrow[L\cup P]{n-1} x$. By induction, we have $w_1 \xrightarrow[L]{*} w_2 \geq_P x$. If $w \xrightarrow[L]{} w_1$ we are done. Otherwise we have $w \xrightarrow[P]{} w_1$ and by Fact 3.0.1 we have $w \xrightarrow[L]{*} w_2$ (and we are done) or $w \xrightarrow[L]{*} w_3 \xrightarrow[P]{} w_2$ and then $w_3 \geq_P x$ (since $w_2 \geq_P x$ and by Lemma 3.0.6) and again we have the desired fact.

□

Proposition 3.0.3. *A word w in W has genus g iff it is of type x for some $x \in {W_g}^{nf}$.*

Proof. Immediate consequence of Lemma 3.0.2 and Proposition 3.0.2. □

4 Monadic second-order logic

Our aim will be to formalize the notions of Sections 2 and 3 in monadic second-order logic. We only recall here that monadic second-order logic is the extension of first-order logic with variables denoting sets of elements of the domain, *i.e.* here sets of vertices or sets of darts for the case of maps. The reader will find precise definitions in [6] (where monadic second-order logic is used to express properties of maps) and in the survey paper [5]. In the sequel, MS will abbreviate "monadic second-order".

4.1 Forbidden submaps

Our first objective is to formalize in MS logic the characterizations of Σ-maps obtained in Section 2.

A map $M =< D, \alpha, \sigma >$ is a logical structure with domain D. One can handle α and σ as binary functional relations (*i.e.* input-output relations of partial functions) on D in order to use relational structures as in [5]. It is convenient to let $M =< D_M, \underline{\mathrm{match}}_M, \underline{\mathrm{suc}}_M >$ where $\underline{\mathrm{suc}}_M = \{(x, \sigma(x)) \mid x \in D_M\}$ and $\underline{\mathrm{match}}_M = \{(x, \alpha_M(x)) \mid x \in D_M\}$. Hence $\underline{\mathrm{match}}_M$ is symmetric.

Let $M =< D, \alpha, \sigma >$ be a map. Let $X \subseteq D$, be α-closed. Then, we get a submap M' of M, denoted by $M[X]$ with set of darts X (see Subsection 1.2).

We recall from [5] that the transitive closure of an MS-expressible binary relation is MS-expressible. Since, if $M' = M[X] \subseteq M$, the relation $\underline{\mathrm{suc}}_{M'}$ is defined in terms of the transitive closure of the relation $\underline{\mathrm{suc}}_M$ (see Subsection 1.2). It follows that one can build an MS formula $\eta(X, x, y)$ such that for every map M, for every α-closed subset X of D_M, for every x, $y \in X$ we have

$$M \models \eta(X, x, y)$$

iff $(x, y) \in \underline{\mathrm{suc}}_{M[X]}$.

We recall that $M \xrightarrow[T]{*} N$ means that the edges of N are replaced by nonempty disjoint paths in M. In particular, a vertex x of N has the same degree as the corresponding vertex $h(x)$ in M, where h is as it is introduced in the following:

Lemma 4.1.1. *Let M and N be maps. We have $M \xrightarrow[T]{*} N'$ where $N' \equiv N$ iff there exists an injective map $h : D_N \to D_M$ satisfying the following conditions:*

1. *For $d \in D_N$, $d' = \sigma_N(d)$ then $h(d') = \sigma_M(h(d))$.*
2. *For $f \in D_M - h(D_N)$ we have $\sigma_M(f) \notin h(D_N)$ and $\sigma_M^2(f) = f$.*
3. *For $d \in D_N$, $d' = \alpha_N(d)$ then $h(d') = (\alpha_M\sigma_M)^n(\alpha_M(h(d)))$ where n is the smallest integer $n \geq 0$ such that $(\alpha_M\sigma_M)^n(\alpha_M(h(d))) \in h(D_N)$.*

Sketch. The first condition says that the circular order of darts in N around a vertex x is mapped by h to the circular order of darts in M around $h(x)$.

The second condition says that the vertices of M not images by h of vertices of N have degree 2 and have two incident darts not in $h(D_N)$.

The third condition says that an edge $e = \{d, d'\}$ corresponds by h to a path in M and that its two end darts are $h(d)$ and $h(d')$ and are the only ones in $h(D_N)$. $\square$

Proposition 4.1.1. *Let N be a map.*

1. *The property of a map M that it is isomorphic to M' such that $M' \xrightarrow[T]{*} N$ is MS-expressible.*

2. *The property of a map M that $M \trianglerighteq_T N$ is MS-expressible.*

Proof. 1. Let us enumerate D_N in a fixed way as $d_1, \ldots, d_{2n}$. One can write an MS formula $\varphi(x_1, x_2, \ldots, x_{2n})$ such that, for every map M, for every $d'_1, d'_2, \ldots, d'_{2n}$ in D_M,

$$M \models \varphi(d'_1, d'_2, \ldots, d'_{2n})$$

iff the mapping $h : d_i \mapsto d'_i$ of $D_N \to D_M$ is injective and satisfies the conditions of Lemma 4.1.1. Hence $M \xrightarrow[T]{*} N$ iff $M \models \exists x_1, x_2, \ldots, x_{2n}.\varphi(x_1, x_2, \ldots, x_{2n})$. We let φ_N be the formula $\exists x_1, \ldots, x_{2n}.\varphi(x_1, \ldots, x_{2n})$.

2. We recall from Section 1 that $M \trianglerighteq_T N$ iff there exists an α-closed subset of D_M such that $M[X] \xrightarrow[T]{*} N$.

   We have observed above that the relation $\mathrm{suc}_{M[X]}(x, y)$ can be expressed in M by an MS formula $\eta(X, x, y)$. Let us consider the property $M[X] \models \varphi_N$. It can be expressed in M by an MS formula $\psi'_N(X)$. This formula is obtained from φ_N by the following transformations:

   - An atomic formula $\underline{\mathrm{suc}}(u, v)$ is replaced by $\eta(X, u, v)$.
   - A quantification $\forall u.\theta$ is replaced by $\forall u(u \in X \Rightarrow \theta')$ where θ' is derived recursively from θ.
   - A quantification $\exists u.\theta$ is replaced by $\exists u(u \in X \wedge \theta')$ where θ' is derived recursively from θ.

   We now let ψ_N be the formula

$$\exists X.(\psi'_N(X) \wedge \text{ “X is } \alpha\text{-closed”}).$$

   It expresses in M the existence of X, α-closed, such that $M[X] \models \psi_N$ hence the fact that $M \trianglerighteq_T N$.

$\square$

Corollary 4.1.1. *For every $g \geq 0$ the property of a map M that it has genus at most g is MS-expressible.*

Proof. We know from Theorem 1.2.1 that there exists a finite set of maps $N_1, \ldots, N_k$ such that an arbitrary map M has genus at most g iff $M \not\trianglerighteq_T N_i$ for each $i = 1, \ldots, k$. Hence M has genus at most g iff it satisfies the MS-formula $\neg\psi_{N_1} \wedge \ldots \wedge \neg\psi_{N_k}$ where the formulas ψ_{N_i} are as in Proposition 4.1.1. $\square$

4.2 Informative characterizations of Σ-embeddability

In Corollary 4.1.1 we have proved that the Σ-embeddability of a given map is MS-expressible by a formula translating conditions of the form: the given map M does not contain as submap a map M' of a certain form.

This characterization is not informative in the sense that when the property is true, its validity gives no hint for constructing an embedding. It only states the existence of an embedding.

By an "informative characterization", we mean one described by a formula of the form:

$$\exists X_1, \ldots, X_n . \mu(X_1, \ldots, X_n)$$

such that every n-tuple satisfying μ in M contains somehow a description of an embedding of M in Σ.

The proof idea is as follows. Let M be a connected map and T be a spanning tree of $G(M)$. Let $M \backslash T$ be the 1-vertex map obtained by the contraction of all edges of T. By Lemma 1.2.1 it has the same genus as M. Simultaneously, from a description of an embedding of $M \backslash T$ in Σ, one can obtain an embedding of M in Σ, as we will explain later. We are thus reduced to handle 1-vertex maps, *i.e.* the circular words with two occurrences of each letter that we considered in Section 3. Hence we give some lemmas that express in MS logic some properties of words developed in Section 3.

Let $w \in A^* - \{\varepsilon\}$, $w = a_1 a_2 \ldots a_m$ with $a_i \in A$ for $1 \leq i \leq m$. We let $\|w\|$ be the relational structure $< \{1, 2, \ldots, m\}, \underline{\text{match}}_w, \underline{\text{suc}}_w >$ such that

$$\underline{\text{suc}}_w = \{(i, i+1) \mid 1 \leq i \leq m\}$$

and

$$\underline{\text{match}}_w = \ \{(i, j) \mid 1 \leq i, j \leq m, i \neq j, a_i = a_j\}.$$

If w, w' are two words in A^* then $\|w\| \equiv \|w'\|$ iff $w' = h(w)$ for some isomorphism $h : A^* \rightarrow A^*$ (Hence $\|w\| \equiv \|w'\|$ implies $w \sim w'$).

If w is as above and $X \subseteq \{1, 2, \ldots, m\}$ we say that X is a set of *positions* in w. A set of *consecutive positions* is one that forms an interval $[i, j]$ for $i < j$. Every set of positions X is the union of a set of sets of consecutive positions denoted by $\text{Consec}(X)$ (where two sets are merged as much as possible).

We let $\text{Consec}(w, X)$ denote the set of factors of w induced by the sets in $\text{Consec}(X)$. Note that the words in $\text{Consec}(w, X)$ need not to be in W, even if $w \in W$.

We let $w - X$ denote the subword of w obtained by deleting the positions in X.

Example 4.2.1. We let $w = abcdcdba$, (note that $w \in W$), $X = \{2, 3, 6, 7\}$, $\text{Consec}(w, X) = \ \{bc, db\}$, $w - X = adca$. Hence $w - X \notin W$.

Lemma 4.2.1. *One can construct MS formulas θ, θ_0, $\theta_1(X)$, $\theta_2(X, x, y)$ such that, for every word $w \in A^* - \{\varepsilon\}$, for every $X \subseteq D_{\|w\|}$ we have:*

1. $\|w\| \models \theta$ *iff* $w \in W$.

2. $\|w\| \models \theta_0$ *iff* $w \in W_0$.

3. $\|w\| \models \theta_1(X)$ *iff* $\mathrm{Consec}(w, X) \subseteq W_0$.

4. *The binary relation on* $D_{\|w\|} - X$ *defined as* $\{(u, v) \mid \|w\| \models \theta_2(X, u, v)\}$ *is the relation* $\underline{\mathrm{suc}}_{\|w-X\|}$.

Proof. 1. From the definition of W it is straightforward to construct θ.

2. The formula θ_0 is then the conjunction of θ and a formula expressing the condition of Lemma 3.0.4 that is equivalent to $w \in W_0$ by Lemma 3.0.5.

3. The construction of $\theta_1(X)$ follows then easily. One writes it in such a way that it is equivalent to: For every maximal subset (for inclusion) Y of X that is a set of consecutive positions, the factor of w induced on Y satisfies the formula θ_0.

4. One writes a formula $\theta_2(X, x, y)$ expressing that x, y are not in X, x is before y, and all positions between x and y are in X.

□

It follows then from Lemma 3.0.5 that if $K \subseteq A^*$ is MS-definable then the set

$$K' = \{w' \in A^* \mid w' \xrightarrow[L]{*} w \text{ for some } w \in K\}$$

is MS-definable.

If ψ describes K then one takes for describing K' a formula ψ' expressing that:

There exists a set of positions X in w'
that satisfies $\theta'_1(X)$ and such that $w' - X$ satisfies ψ

where $\theta'_1(X)$ is like $\theta_1(X)$ except that one checks $w_n w_0 \in W_0$ for w_0, w_n the respectively first and last factors of w' induced on maximal subsets of consecutive positions in X.

In order to express that $w' - X$ satisfies ψ one replaces in ψ the atomic formulas $\underline{\mathrm{suc}}(u, v)$ by $\theta_2(X, u, v)$.

We now want to express the relation $\geq_P$ in a logical way.

We recall that if w and $x = a_1 a_2 \ldots a_{2n}$ are nonempty words in W then

$$w \geq_P x$$

means that

$$w = u'_1 a_1 v_1 u_2 a_2 v_2 \ldots u_{2n} a_{2n} v'_{2n} \text{ where } v'_{2n} u'_1 = v_{2n} u_1 \tag{1}$$

for words $u'_1, u_1, u_2, \ldots, u_{2n}, v_1, v_2, \ldots, v_{2n}, v'_{2n} \in A^*$ such that

$$u_i a_i v_i = \widetilde{v_j} a_j \widetilde{u_j} \text{ if } i < j \text{ and } a_i = a_j. \tag{2}$$

Let us consider a word w. A decomposition of w as in (1) will be formalized by sets of positions $Y''_1, Y'_1, Y_1, Y_2, \ldots, Y_{2n}, Y'_{2n}, Y''_{2n}$ such that:

(i) either

$$Y_1 = Y_1', Y_{2n} = Y_{2n}' \cup Y_{2n}'' \cup Y_1'' \text{ or}$$
$$Y_{2n} = Y_{2n}', Y_1 = Y_{2n}'' \cup Y_1'' \cup Y_1' \text{ or}$$
$$Y_1 = Y_1'' \cup Y_1', Y_{2n} = Y_{2n}' \cup Y_{2n}'',$$

(ii) each of Y_i, $2 \le i \le 2n-1$ is a set of consecutive positions, at least one of Y_1, Y_{2n} is a set of consecutive positions, at most one of Y_1, Y_{2n} is the union of two sets of consecutive positions, $Y_1 \cup Y_{2n}$ is not a set of consecutive positions, and $1, |w| \in Y_1 \cup Y_{2n}$,

(iii) $D_{\|w\|} = \bigcup_{1 \le i \le 2n} Y_i$,

(iv) Z contains one and only one element in each Y_i,

(v) $u_i a_i v_i$ is the factor of w induced by Y_i except for at most one of Y_1, Y_{2n} which induces the word $w''w' = u_1 a_1 v_1$ or $w''w' = u_{2n} a_{2n} v_{2n}$ where w' is a prefix of w and w'' is a suffix of w,

(vi) Z is the set of positions of letters $a_1, \ldots, a_{2n}$ of w.

Lemma 4.2.2. *For every $x \in W$ of length $2n$ $(n \ge 1)$ there exists an MS formula*

$$\omega_x(Y_1'', Y_1', Y_1, \ldots, Y_{2n}, Y_{2n}', Y_{2n}'', Z)$$

such that for every $w \in W$ and $Y_1'', Y_1', Y_1, \ldots, Y_{2n}, Y_{2n}', Y_{2n}'', Z \subseteq D_{\|w\|}$ we have

$$\|w\| \models \omega_x(Y_1'', Y_1', Y_1, \ldots, Y_{2n}, Y_{2n}', Y_{2n}'', Z)$$

iff $Y_1'', Y_1', Y_1, \ldots, Y_{2n}, Y_{2n}', Y_{2n}'', Z$ satisfy Conditions (i)–(vi) for some letters $a_1, a_2, \ldots, a_{2n} \in A$ such that $a_1 a_2 \ldots a_{2n}$ is isomorphic to x and Condition (2) holds.

Proof. We first construct a formula $\omega'(Y, Y', Z)$ such that for $Y, Y', Z \in D_{\|w\|}$, we have

$$\|w\| \models \omega'(Y, Y', Z)$$

iff:

Either

- Y and Y' are disjoint sets of consecutive positions, $Y \cap Z$ and $Y' \cap Z$ are singletons that are occurrences of a same letter, say a, in w,
- the factors $w[Y]$ and $w[Y']$ of w verify $w[Y] = w[\widetilde{Y'}]$.

or

- Y is a disjoint set of consecutive positions, Y' is the union of two sets of consecutive positions Y_1', Y_2', $1 \in Y_1'$, $|w| \in Y_2'$, $Y \cap Z$ and $Y' \cap Z$ are singletons that are occurrences of a same letter, say a, in w,

- the factors $w[Y]$, $w[Y_1']$, $w[Y_2']$ of w verify $w[Y] = w[\widetilde{Y'}]$ for $w[Y'] = w[Y_2']w[Y_1']$.

The construction formalizes the following conditions:

- The first position of $w[Y]$ has the same letter as the last position of $w[Y']$.
- For every two positions y in Y and y' in Y' having the same letter the position z following y (unless y is last in Y) has the same letter as the position z' in Y' preceding y' (unless y' is first in Y').

The construction of ω_x is then straightforward. For expressing (2) we use ω' as follows.

For any two i, j with $1 \leq i, j \leq 2n$ which are positions of a same letter in x, we impose the clause $\omega'(Y_i, Y_j, Z)$. If $1 \leq i < j \leq 2n$ and do not correspond to the same letters in x we write that in w, the positions $Y_i \cap Z$ and $Y_j \cap Z$ do not hold the same letter. □

Hence we have:

Proposition 4.2.1. *For every word $x \in W^{nf}$ there is an MS formula φ_x expressing that M is of type x.*

The formula φ_X is of the form, for $n = |x|/2$:

$$\exists X, Y_1'', Y_1', Y_1, \ldots, Y_{2n}, Y_{2n}', Y_{2n}'', Z.\theta(X, Y_1'', Y_1', Y_1, \ldots, Y_{2n}, Y_{2n}', Y_{2n}'', Z)$$

so that if $\|w\| \vDash \varphi_x$, any $2n + 6$ tuple $(X, Y_1'', Y_1', Y_1, \ldots, Y_{2n}, Y_{2n}', Y_{2n}'', Z)$ satisfying θ reflects the formalization of Proposition 3.0.2. In particular, the set X defines the factors which are deleted by $\xrightarrow[L]{*}$ in the writing $w \xrightarrow[L]{*} w' \geq_P x$.

These factors correspond to sets of contractible loops in the map represented by w. The sets $Y_1, \ldots, Y_{2n}$ represent n sets of parallel edges (possibly reduced to a simple edge) which are fused in the transformation of the map M' (represented by w') into the map N (represented by x). If we know how to embed N on a fixed surface (represented by a polygon, see Section 1), then we know how to embed M': it suffices to replace the edges of N by sets of parallel edges. From this embedding of M' we can obtain an embedding of M by the insertion of loops as they are specified by the factors of w that are deleted from w to form w'. We illustrate this with the following example.

Example 4.2.2. We let

$$w = fecaggdbhhdacbef,$$

$$w' = cadbdacb,$$

$$x = abab.$$

From the embedding of N (represented by x) on the torus of left of Figure 14 we get the embedding of M' (represented by w') of middle of Figure 14 and that of M (represented by w).

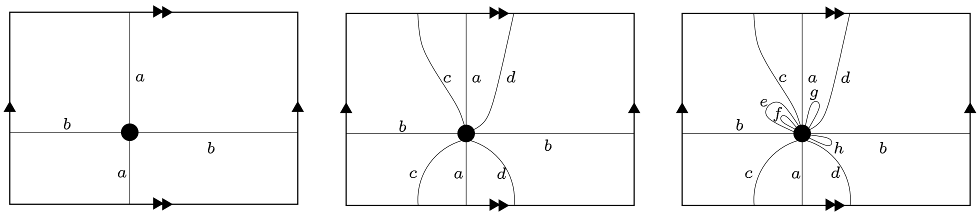

Figure 14: The three torus embeddings of Example 4.2.2

Lemma 4.2.3. *There exists an MS-formula $\gamma(X)$ such that, for every map M, every $X \subseteq D_M$, $M \models \gamma(X)$ iff X is α_M-closed and X defines the set of edges of a spanning tree $T(X)$ of $G(M)$.*

Proof. A set of darts X is α_M-closed iff $\forall x, y\ (x \in X \wedge \underline{\text{match}}(x, y) \Rightarrow y \in X)$ holds. Assuming this, let E be the corresponding set of edges of $G(M)$. A path in E can be defined as an α_M-closed subset Y of X, such that no class $[X]_\sigma$ contains more than two elements of Y and $M[Y]$ is connected (this is also MS expressible). This path is a cycle iff every class $[X]_\sigma$ for $x \in D_M$ contains 0 or 2 elements of Y. From these observations it is easy to express that E is the set of edges of a spanning tree of $G(M)$. $\square$

We will denote by $M\backslash X$ the 1-vertex map resulting from the contraction of the tree $T(X)$ when $\gamma(X)$ holds in M. It is clear that $D_{M\backslash X} = D_M - X$.

Lemma 4.2.4. *There exist MS formulas $\gamma_\sigma(X, x, y)$, $\gamma_\alpha(X, x, y)$ such that for every map M, for every X such that $\gamma(X)$ holds we have*

$$\begin{aligned} M \models \gamma_\alpha(X, x, y) \quad &\Leftrightarrow \quad x, y \notin X,\ y = \alpha_{M\backslash X}(x) \\ \textit{and } M \models \gamma_\sigma(X, x, y) \quad &\Leftrightarrow \quad x, y \notin X,\ y = \sigma_{M\backslash X}(x). \end{aligned}$$

Proof. For $x, y \in D_M - X$ we have $y = \alpha_{M\backslash X}(x)$ iff $y = \alpha_M(x)$. We also have $y = \sigma_{M\backslash X}(x)$ iff $y = \sigma_M(\alpha_M \circ \sigma_M)^n(x)$ where n is the smallest $n \geq 0$ such that $\sigma_M(\alpha_M \circ \sigma_M)^n(x) \notin X$. These conditions are easily MS expressible. $\square$

Proposition 4.2.2. *For every word $w \in W^{nf}$, there exists an MS formula $\theta_w(X)$ such that for every map M, every $X \subseteq D_M$:*

$$M \models \theta_w(X)$$

iff $T(X)$ is a spanning tree of $G(M)$ and $M\backslash X$ has type w.

Proof. Clear from Lemmas 4.2.3, 4.2.4 and Proposition 4.2.1. $\square$

We have thus achieve the goal of this subsection because we have for describing maps of genus at most g a formula of the form

$$\exists X.[\theta_{w_1}(X) \vee \ldots \vee \theta_{w_k}(X)]$$

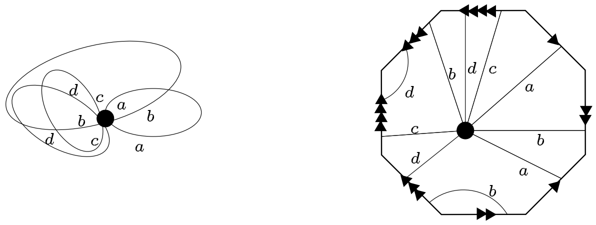

Figure 15: *abacdbcd*, the fourth forbidden configuration for the torus shown on Figure 9

where $w \in W_g$ iff w has type w_i for some $i = 1, \ldots, k$. We have thus the desired informative characterization in MS-logic.

For each w_i forming a type for a given surface, we can fix one representation like on Figure 15. The formula $\theta_{w_i}(X)$, when it holds true, fixes the tree X that can be drawn in the middle of the polygon, the loops around this tree, the sets of parallel edges (there is one set of each letter of w_i) to be drawn across the borders of the polygon. By using Hanusse's algorithms for drawing planar maps (see [14]) we can thus produce a planar drawing inside a polygon which thus specifies by conventions on polygons representing surfaces of Section 1 an embedding into a nonplanar surface. This technique makes it possible to represent nonproper embeddings. It chooses some embeddings but is unable to yield all of them.

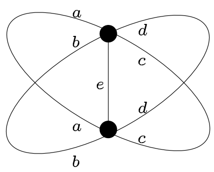

Figure 16: There exist two words associating to this map

Example 4.2.3. *(Two words for a same map)* Let us consider the map of Figure 16.

By contracting e one obtains the word *ababcdcd*. By contracting b one obtains the word *acdeaecd* which is not equivalent to the previous one. Hence the words in W_g^{nf} are not in bijection with the 1-vertex maps of genus g.

Acknowledgments

We thank G. Schaeffer for helpful remarks.

References

[1] D. Archdeacon. A Kuratowski theorem for the projective plane. *J. Graph Theory*, 8:387–398, 1981.

[2] S. Arnborg, J. Lagergren, and D. Seese. Easy problems for tree-decomposable graphs. *Journal of Algorithms*, 12:308–340, 1991.

[3] R. Cori and M. Marcus. Counting non-isomorphic chord diagrams. *Theoret. Comput. Sci.*, 204:55–73, 1998.

[4] B. Courcelle. The monadic second-order logic of graphs I: Recognizable sets of finite graphs. *Information and Computation*, 85:12–75, 1990.

[5] B. Courcelle. *The expression of graph properties and graph transformations in monadic second-order logic in Handbook of graph grammars and computing by graph transformations*, chapter 5, pages 313–400. World Scientific. G. Rozenberg ed., 1997.

[6] B. Courcelle. The monadic second-order logic of graphs XII: Planar graphs and planar maps. *Theoret. Comput. Sci.*, 237:1–32, 2000.

[7] B. Courcelle, J. A. Makowsky, and U. Rotics. On the fixed parameter complexity of graph enumeration problems definable in monadic second order logic. *Discrete Applied Mathematics*, 108:23–52, 2001.

[8] B. Courcelle and M. Mosbah. Monadic second-order evaluations of tree-decomposable graphs. *Theoret. Comput. Sci.*, 109:49–82, 1993.

[9] N. Dershowitz and J. P. Jouannaud. *Handbook of theoretical computer science*, volume B, chapter 6, Rewrite systems. Elsevier Science Publishers, 1990.

[10] R. Diestel. *Graph theory.* Springer, second edition, 1997.

[11] R. G. Downey and M. R. Fellows. *Parameterized complexity.* Springer-Verlag, New-York, 1999.

[12] J. L. Gross and T. W. Tucker. *Topological Graph Theory.* Wiley Interscience Series in Discrete Mathematics and Optimization, 1987.

[13] Y. Gurevich. *Monadic second-order theories in Model-Theoretic Logics*, chapter XIII, pages 479–506. J. Barwise, S. Feferman (Eds.), Springer-Verlag, New-York, 1985.

[14] N. Hanusse. *Cartes, constellations et groupes : questions algorithmiques.* PhD thesis, Université Bordeaux I, December 1997.

[15] D. M. Jackson and T. I. Visentin. *An atlas of the smaller maps in orientable and nonorientable surfaces.* Chapman & Hall/CRC, 2001.

[16] B. Mohar and C. Thomassen. *Graphs on surfaces.* J. Hopkins Univ. Press, 2001.

[17] M. O. Rabin. Decidability of second-order theories and automata on infinite trees. *Trans. Amer. Math. Soc.*, 141:1–35, 1969.

[18] N. Robertson and P. Seymour. Graph minors IX : Disjoint crossed paths. *Journal Comb. Theory*, 49(B):40–77, 1990.

[19] N. Robertson and P. Seymour. Graph minors VIII : A Kuratowski theorem for general surfaces. *Journal Comb. Theory*, 48(B):255–288, 1990.

[20] N. Robertson and P. Seymour. Graph minors XXII: Wagner's conjecture. *Journal Comb. Theory*, (B), to appear.

[21] G. Schaeffer. *Cartes planaires*. PhD thesis, Université Bordeaux I, December 1998.

[22] W.T. Tutte. *Graph theory*. Addison Wesley, 1984.

[23] T. R. S. Walsh and A. B. Lehman. Counting rooted maps by genus I. *Journal Comb. Theory*, 13(3):192–217, December 1972.