

Fixed Parameter Undecidability for Wang Tilesets

Emmanuel Jeandel*

LIRMM, CNRS UMR 5506 - CC 477
161 rue Ada, 34 095 Montpellier Cedex 5, France
jeandel@lirmm.fr

Nicolas Rolin

LIP6
4 place Jussieu, 75252 Paris cedex 05, France
nicolas.rolin@ens-cachan.fr

Deciding if a given set of Wang tiles admits a tiling of the plane is decidable if the number of Wang tiles (or the number of colors) is bounded, for a trivial reason, as there are only finitely many such tilesets. We prove however that the tiling problem remains undecidable if the difference between the number of tiles and the number of colors is bounded by 43.

One of the main new tool is the concept of Wang bars, which are equivalently inflated Wang tiles or thin polyominoes.

Introduction

Wang tiles are a model of computation introduced by Wang [6] to study decision procedures for some fragments of first-order logic. The model is quite simple: We are given a finite set of tiles, i.e. a squares with colored edges, and we look at a way to tile the plane with the tiles so that contiguous edges have the same color. Berger, a student of Wang, proved [1] that the problem is algorithmically undecidable: there is no way to decide, given a set of tiles, whether it can tile the plane. One of the reasons for the difficulty of the problem is the existence of *aperiodic* sets of tiles, for which it is possible to find a tiling, but no *periodic* one.

If we happen to bound the number of different tiles, say by 100, then the problem becomes trivially decidable, as there are only a finite number of such sets of tiles. The same is true if we bound the number of colors. We will try here to find a good parameter, for which this trivial situation does not happen: We prove in this article that the problem remains undecidable if the difference between the number of tiles and the number of colors is bounded by a constant, here 43. Furthermore, we prove that if the difference is small enough, then the tiling problem is actually decidable, because there are no aperiodic tilesets with such a difference. In another way, if there are too many colors (compared to the number of tiles), the set of tiles cannot be aperiodic.

This last result has pragmatic implications that motivated this research. The first author is at present conducting experiments to find an aperiodic set of tiles with as few tiles as possible (actually set of tiles that are *candidates* for being aperiodic, as the problem is not decidable. . .). The main bottleneck for this kind of approach is indeed the number of such sets of tiles, thus proving that we need a small number of colors to be aperiodic may critically reduce the size of the search space.

The article is organised as follows. We introduce the relevant definitions in the first section. We then prove that the problem under consideration is equivalent to the question of tileability with a fixed number of *Wang bars*, which are an intermediate between Wang tiles and polyominoes. We then proceed to prove the upper and lower bound on the parameter that is needed to have an undecidability result.

*The first author is partially supported by ANR-09-BLAN-0164. Part of this work was done during a traineeship of the second author.



Figure 1: A Wang tile

1 Definitions

1.1 Wang tiles

A Wang tile t is a square tile with colored edges, as represented in Fig.1. Formally, it is given by a quadruplet (t_e, t_w, t_n, t_s) of symbols, called colors. A *tileset* τ is a finite set of Wang tiles. A *tiling* of the plane by τ is a map f from the discrete plane \mathbb{Z}^2 to τ so that two tiles that share a common edge agree on the color: For all integers i, j , we have $f(i, j)_e = f(i + 1, j)_w$ and $f(i, j)_n = f(i, j + 1)_s$.

The main question we study in this article is the Domino Problem: To decide, given a tileset τ , if there is a tiling by τ . This problem was proven undecidable by Berger [1].

A reason for the complexity of the problem is the existence of *aperiodic* tilesets. A tiling t is *periodic* if there exists p so that $t(i, j) = t(i, j + p) = t(i + p, j)$ for all i, j . If our tileset τ admits a periodic tiling, this tiling is easy to find, by just testing all possible finite maps from $[0, n]^2$ to τ .

A tileset is said to be periodic if it admits a periodic tiling. It can be proven equivalent to the existence of a tiling t which is only horizontally periodic, that is $t(i, j) = t(i + p, j)$ for all i, j and for some p . However there exist *aperiodic* tilesets, that is tilesets that tile the plane, but that admit no periodic tiling. Such a tileset, due to Culik [3] based on work by Kari [4] is depicted in Fig. 2 However, every such tileset




Figure 2: An aperiodic tileset with 13 tiles from Kari and Culik

admits tilings with some specific regularity properties. We will use the following lemma in some of the proofs:

Lemma 1.1. *Let τ be a tileset that tiles the plane. Then there exists a tiling f by τ with the following property: For every tile t that appears in f , there exist a row where t appears at least twice*

We can prove much more (for example the existence of a quasiperiodic configuration [2]), but this will be sufficient for this article.

Proof. Take a tiling f with the minimal number of different tiles. If some tile t appears at most once on each row of f , then we can find in f big squares where t does not appear. By a compactness argument, we can extract from f a tiling f' without t , hence with strictly less different tiles, a contradiction. ■




Figure 3: The tileset of Kari and Culik represented as a labeled graph.

1.2 The graph approach

If we look at tilings of only one row by τ , a good way to understand the tilings is using a (labeled) graph: Represent each color by a vertex, and for each Wang tile t with west color c and east color c' , add an edge from c to c' labeled with its north and south color. An example corresponding to Fig.2 is provided in Fig.3. Note that for this particular example the graph has two connected components, that are strongly connected (the relevant vocabulary on graph theory will be defined below).

It is quite clear that a tiling of a row is equivalent to a biinfinite path in this graph. We will now describe more precisely the connection.

First, some definitions. A labeled (multi)graph where each edge is labeled by a pair of symbols will be called a 1-labeled graph.

Given a 1-labeled graph G , a pair of biinfinite words (u^1, u^2) is *compatible* if (u^1, u^2) correspond to a biinfinite path on the graph. That is, there exists vertices $(v_i)_{i \in \mathbb{Z}}$ so that for all $i \in \mathbb{Z}$ there is an edge from v_i to v_{i+1} labeled (u_i^1, u_i^2) . A *tiling* by G is a biinfinite sequences of biinfinite words $(u^j)_{j \in \mathbb{Z}}$ so that (u^j, u^{j+1}) is compatible for all j .

Now it is clear that, upto notations, a 1-labeled graph G with n edges and k vertices is exactly the same as a set τ of n Wang tiles with k colors on its east/west wide, and that a tiling by G exists if and only if a tiling by τ exists.

1.3 The parameter

For a given set of Wang tiles τ , we denote by $n(\tau)$ (or simply n) the number of tiles of τ and by $c(\tau)$ (or simply c) the maximum over all four sides of the number of colors. By definition, it is clear that $c(\tau) \leq n(\tau)$.

By rotating the Wang tiles, we may suppose w.l.o.g., and we do from now, that the maximum number of colors is reached on the west side.

If we view the set of Wang tiles as a 1-labeled graph G , $n(\tau)$ represents the number of edges of G , and $c(\tau)$ the number of vertices of G (more precisely the number of vertices with non-zero outdegree).

It is obvious that $n(\tau)$ and $c(\tau)$ are not good parameters, in the sense that the Domino problem is trivially decidable for a fixed n (resp. c), as there are only finitely many tilesets with n tiles (resp. using at most c colors on each side).

However we now observe that $n(\tau) - c(\tau)$ might be a good parameter. Indeed, if the number of colors on the west side is roughly the same as the number of tiles, then this means that for almost every tile, there usually will be at most one tile that can be put on its right. The tiling problem becomes highly constrained and we might expect the problem to be easier. This is indeed the case:

Theorem 1.1. *The domino problem for tilesets of parameter $n - c \leq 1$ is decidable. More precisely, a tileset with parameter $n - c \leq 1$ admits a tiling if and only if it admits a periodic tiling.*

On the other hand, if $n - c$ is big, we have more choices for the next tile to put. We will prove:

Theorem 1.2. *The domino problem for tilesets of parameter $n - c = 43$ is undecidable. It remains undecidable for higher parameters.*

What happens between 1 and 43 is unknown. What is clear is that the situation becomes more complex very quickly:

Proposition 1.1. *There exists a tileset τ of parameter $n - c = 8$ that admits a tiling but no periodic tiling (i.e. that is aperiodic).*

This is exactly the tileset of Fig.2.

The rest of the paper will be organized as follows. In the next section, we will introduce a new object, called Wang bars, and prove that tileability with a tileset of parameter $n - c = k$ is (somewhat) equivalent to tileability with $O(k)$ Wang bars. Then we will prove in the following sections that: (a) tileability with 44 Wang bars is undecidable (thus proving Theorem1.2) and that (b) tileability with 2 Wang bars is decidable (thus proving Theorem1.1).

2 Wang bars

A Wang bar is a Wang tile which is bigger horizontally. Formally, a Wang bar b is a quadruplet (b_e, b_w, b_n, b_s) where b_n and b_s are words of the same length over some alphabet C . We denote by $|b|$ the length of the word b_n . Letters of b_n are numeroted from 1 to $|b|$. It is clear from the definition that a Wang tile is a Wang bar of length 1.

A *barset* is a set of Wang bars. See Fig. 4 for an example.

A tiling of the plane by Wang bars is a partition of the plane by Wang bars so that consecutive Wang bars have the same colors on their shared edge. Here is a formal yet nonintuitive definition.

Definition 2.1. *Let B be a barset. A tiling of the plane by B is a pair of a map f from \mathbb{Z}^2 to B ($f(i, j)$ is called the bar at (i, j)), and a map p from \mathbb{Z}^2 to \mathbb{N} ($p(i, j)$ is the position inside the bar (i, j)) so that:*

- $0 < p(i, j) \leq |f(i, j)|$ (the position inside the bar is less than the length of the bar)
- If $p(i, j) < |f(i, j)|$ then $p(i + 1, j) = p(i, j) + 1$ and $f(i + 1, j) = f(i, j)$ (if we are inside the bar, this is still the same bar)
- If $p(i, j) = |f(i, j)|$ (we are at the end of the bar) then $p(i + 1, j) = 1$ (a new bar starts) and $f(i, j)_e = f(i + 1, j)_w$ (horizontal colors match)
- $(f(i, j)_n)_{p(i, j)} = (f(i, j + 1)_s)_{p(i, j + 1)}$ (vertical colors match)




Figure 4: A set of three Wang bars. The first one is formally defined as $(0, 0, 102, 010)$.




Figure 5: A fragment of a tiling by the three bars of Figure 4. The acute reader may try to convince himself that there exists a periodic tiling.

An example is provided in Fig.5.

Wang bars have of course a graph counterpart. We define a ω -labeled graph to be a directed (multi)graph where every edge is labeled by a pair of words of the same length. We do not define explicitly how a ω -labeled graph tiles the plane, but it should be clear. An example of such a graph is provided in Fig. 6

It is quite clear that tileability with Wang bars can be reduced to tileability with Wang tiles: Just consider a Wang bar of length p as p Wang tiles.

Theorem 2.1. *Tileability with k bars is many-one reducible to tileability of Wang tiles with parameter at most $k - 1$. That is, every barset B with k bars can be transformed into a tileset W with parameter at most $k - 1$ so that B tiles the plan iff W tiles the plane.*




Figure 6: The example of Fig. 4 in the graph formalism

Proof. Let c be the number of colors that appear in the east side of a Wang bar of B and k be the number of bars.

We now split each Wang bar into Wang tiles, introducing new colors each time. If we see a Wang tile as a Wang bar, it is easy to see that each time we split a Wang bar into p Wang tiles(bars), we increase the number of bars by $p - 1$ and the number of colors by $p - 1$, hence the difference between the number of bars and the number of colors on the east side stays constant.

As a consequence, the tileset τ we obtain at the end has $n(\tau) = p + k$ tiles and $p + c$ colors in the east side (so that $c(\tau) \geq p + c$), hence is of parameter $n(\tau) - c(\tau) \leq k - c \leq k - 1$. ■

The converse is less clear. We will use the graph formalism to proceed. We start from a tileset with parameter $n - c = k$. We suppose w.l.o.g. that the maximum number of colors is reached on the west side. In the graph formalism, this means that G has n edges, and c vertices with at least one outgoing edge.

First we recall some definitions from digraph theory. The *outdegree* (resp. *indegree*) of a vertex v is the number of labeled edges that start from v (resp. end in v)

A (directed) *path* in G from u to v is a sequence $u = u_1 \dots u_p = v$ (possibly with $p = 1$) of vertices so that for all $i < p$ there exists an edge from u_i to u_{i+1} .

Proposition 2.1. *Let τ be a tileset represented as a graph G . Suppose that there exists two vertices u, v so that:*

- *There exists an edge from u to v , corresponding to the tile t*
- *There exists no path from v to u .*

Then τ tiles the plane if and only if $\tau - \{t\}$ tiles the plane.

Proof. One direction is clear. Suppose that τ tiles the plane. We know from Lemma 1.1 that there exists a tiling f , where each tile that appears in f appears twice in some row. t cannot appear in such a tiling, as it is not possible to go back from its east side (the color v) to the west side (the color u). Hence the tiling f does not use t and $\tau - \{t\}$ tiles the plane. ■

We can now suppose w.l.o.g that there is no such tile t in τ . In particular, in the graph, every vertex is of outdegree (resp. indegree) at least one. Moreover, by a straightforward induction, this implies that if there is a path from u to v , then there is a path from v to u (In digraph terms, all connected components are strongly connected).

For two vertices u and v , we say that $u \sim_G v$ if there exists a path from u to v . \sim_G is in our case an equivalence relation (it is reflexive because every vertex is of outdegree at least one). Note that on each row of a tiling by τ , the horizontal colors belong to the same equivalence class.

These equivalence classes are usually called strongly connected components in the theory of digraphs.

Proposition 2.2. *Let τ be a tileset represented as a graph G for which \sim_G is an equivalence relation.*

Let H be an equivalence class so that every vertex in H is of outdegree exactly one, and T the tiles corresponding to H .

Then:

- *Either $\tau - T$ tiles the plane*
- *Or there is a periodic tiling by τ*
- *Or τ does not tile the plane*

Proof. Suppose that τ tiles the plane. By Lemma 1.1, there exists a tiling where each tile t that appears in f appears in at least two different rows (as it appears twice on some column. . .). If no tile of T appear in f we are done. Otherwise, notice that each row where a tile from T appears must be tiled periodically, of period $|T|$. Now, there exist two rows with a tile from T , say 0 and $p > 0$. We can now obtain a new tiling f' which is periodic vertically by $f'(i, j) = f(i, j \bmod p)$. But this implies that there exists a tiling which is periodic. ■

We now are ready to prove the theorem.

Theorem 2.2. *Tileability of tilesets with parameter at most k is reducible to tileability with at most $2k$ bars. More precisely, using tileability with at most $2k$ bars as on oracle, we can design an algorithm that, given a tileset W with parameter at most k , decides whether W tiles the plane, asking at most one question to the oracle.*

Note: The reduction we have is **not** a many-one reduction. It is usually called a weak-truth-table reduction.

Proof. We start from a tileset W with parameter at most k . If we rotate the tileset, we can suppose w.l.o.g that the maximum number of colors is reached on the west side. An example of the whole algorithm is provided on Fig.7 and Fig.8.

Using Proposition 2.1 and Proposition 2.2, we obtain a tileset $V \subseteq W$ so that:

- Either W tiles the plane periodically
- Or V tiles the plane
- Or W does not tile the plane

By definition of the parameter, it is clear that the parameter of V is less or equal to the parameter of W . We now look at V as a graph G . Let $e(G)$ be the number of edges of G and $v(G)$ its number of vertices. By definition of G , $e(G) - v(G) \leq k$. We now look at each equivalence class (strongly connected component) for \sim_G . We now execute the following algorithm on G , changing it into a ω -labeled graph:

- If the number of edges in each equivalence class is at least twice the number of vertices, then we are done.
- Otherwise let C be such an equivalence class, and $e(C)$ (resp. $v(C)$) its number of edges (resp. vertices). We now have $e(C) < 2v(C)$.
- This implies that $v(C) > 1$, otherwise all vertices of C would have outdegree 1, and such equivalences classes have been eliminated.
- By a simple counting argument, there exists a vertex of outdegree less than 2 (the sum of all outdegrees is equal to $e(C)$), that is of outdegree 1.
- Let u be this vertex. The edge leaving u cannot be from u to u : There exists some other vertex v in C , and there is a path from u to v , which would be impossible if the only vertex leaving u came back to u .




Figure 7: The two first steps in the algorithm of Theorem 2.2. (1) corresponds to the representation of the tileset (I) . (2) deletes edges (u, v) from (1) if no path from v to u exists. (3) deletes equivalence classes (connected components) from (2) which are cycles. The transform from (2) to (3) does **not** preserve tileability.

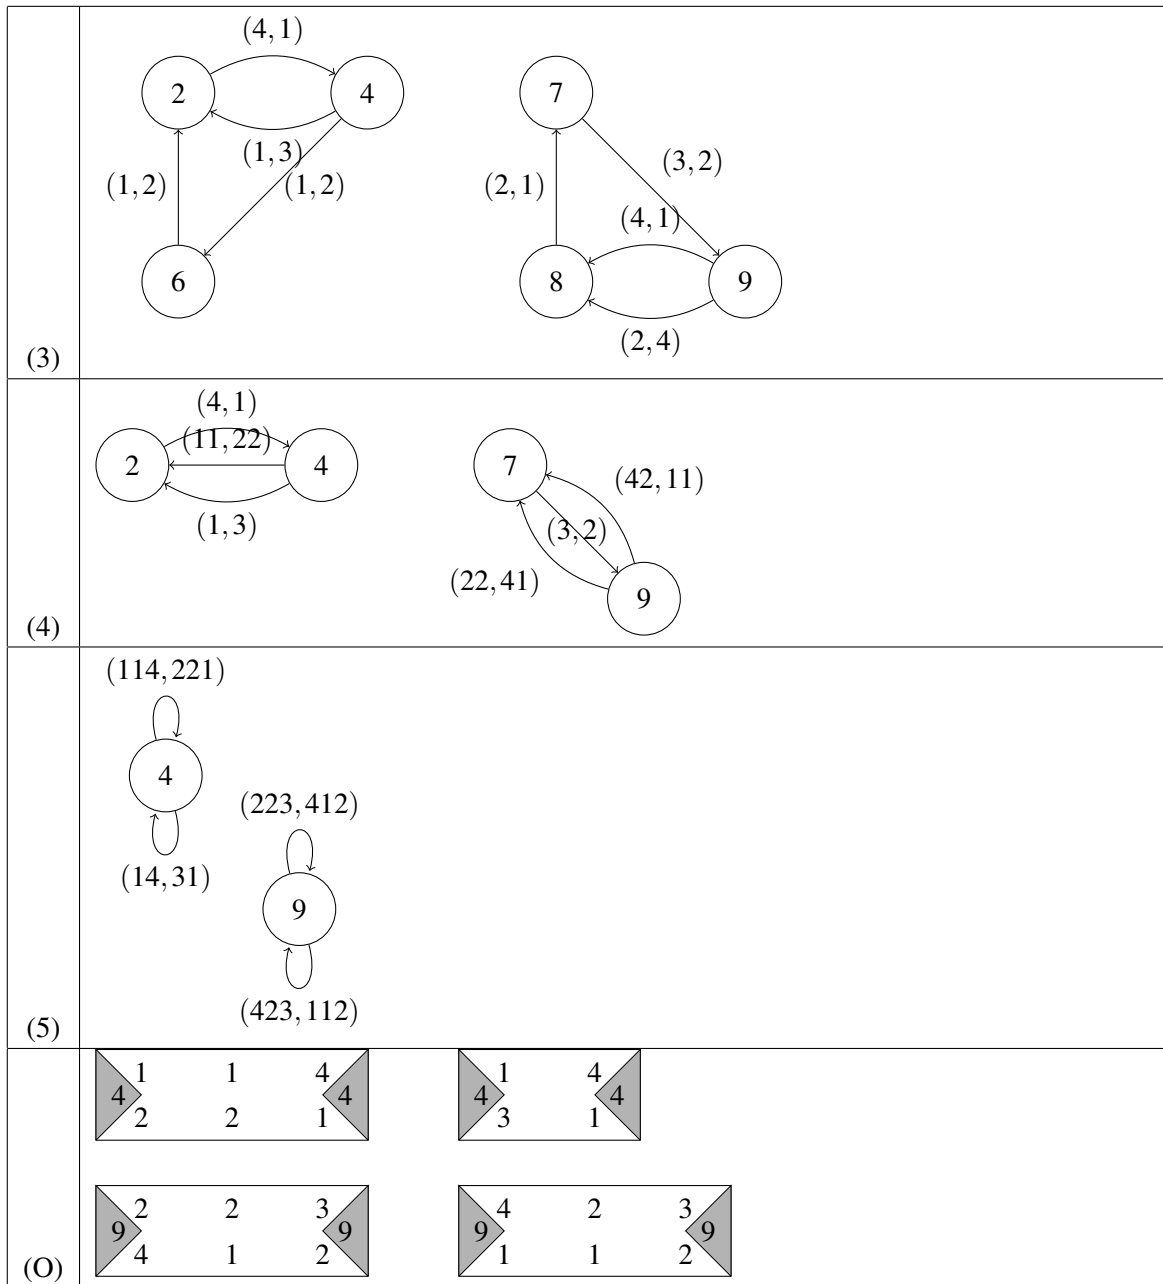


Figure 8: The last steps in the algorithm of Theorem 2.2. We basically contract vertices of outdegree one until there are no more such vertices. We contract vertices 6 and 8 from (3) to (4) and (2) and (7) from (4) to (5). We then obtain the 4 Wang bars in (O), from a initial parameter of $12 - 9 = 3$.

- The edge from u is then going to some $v \neq u$. Let (x,y) be its label. We now change the graph in the following way:
 - We delete the edge leaving from u .
 - For each edge from w to u labeled by (u_1, u_2) , we delete this edge and add a new edge from w to v , labeled with (u_1x, u_2y) . It is clear that this does not change tileability by the ω -labeled graph.
 - The new graph has exactly one less vertex and one less edge than the previous one.
 - Doing this does not change the property of the graph: It is still true that a path from u to v implies a path from v to u , and that there exist no equivalence class where every vertex is of outdegree one.
- We repeat until the number of edges in each equivalence class is at least twice the number of vertices.

Now the ω -labeled graph G' we obtain has n edges and p vertices, with $n \geq 2p$. Furthermore, as we always delete a vertex and a edge at the same time, $n - p \leq k$. We deduce that $n = 2n - 2p + 2p - n \leq 2k + 0 \leq 2k$. So G' represents a set of at most $2k$ bars. Let B be this set.

Now the barset B has the following property:

- Either W tiles periodically
- Or B tiles the plane (in which case W tiles the plane)
- Or W does not tile the plane

So we first ask the oracle whether B tiles the plane. If it does, then W tiles the plane. Otherwise we know that either W tiles the plane periodically, or it does not tile the plane. We can test for the two concurrently until one of them halts. ■

To recap, if we are able to decide if a tileset of parameter k tiles the plane, we can decide whether a barset of $k + 1$ Wang bars tiles the plane. Conversely, if we are able to decide if a barset of k Wang bars tiles the plane, we can decide if a tileset of parameter $k/2$ tiles the plane.

3 Tileability for 44 Wang Bars is undecidable

We will now prove that there is no algorithm that decides whether a set of 44 Wang Bars tiles the plane. This implies that there exists no algorithm that decide whether a tileset of parameter at most 43 tiles the plane.

We will do a reduction from the Domino Problem. That is we will explain how to encode any tileset τ into a barset B_τ of 44 Wang bars (and thus into a tileset of parameter at most 43) so that τ tiles the plane iff B tiles the plane.

The reduction we use here is heavily inspired by the transformation by Ollinger [5] from any tileset to a set of 11 polyominoes. The main difference is that we need here to make the polyominoes “flat”.

Let τ be a set of n Wang tiles. To define B_τ easier, all bars of B_τ will have the property that the color on the east and west side will be the same for all bars, that is the only constraint to put two bars together are vertical constraints. This means that we can see a bar b as a pair of words of the same length.

Now we give the transformation. Let n be the number of tiles in τ and C an upper bound on the number of colors, and suppose the colors are numbered from 1 to C .

The barset will first contain 16 bars, that do not depend on the tileset τ . These bars are depicted in Fig.9. The colors a, b, c, d, e, f, g that appear on these bars appear only in these bars, so the only way to arrange the bars in a tiling is as depicted on the figure. We will omit the labels a, b, c, d, e, f, g on subsequent pictures.




Figure 9: The first 16 bars

The next tiles are almost the same for any tileset, and they just depend on the number of tiles n , and the number of colors C . They are depicted in Fig. 10. As before, lowercase letters and greek letters appear in only two tiles so that the only way to arrange the bars is as depicted. We will not see them as 27 bars, but as 5 big blocks. The first one will be called the *box*, the two next ones *fillers*, and the last ones *handles*.

And now, to finish, the last bar, which depends on the exact tiles we use. For each tile t of τ with north,west, east, south color respectively (a, b, c, d) , we will first consider the following pair of words $(n_t, s_t) = (0^{a-1}Y0^{C-a}B0^{b-1}Y0^{C-b}, 2^{c-1}Y2^{C-c}B2^{d-1}Y2^{C-d})$ over the alphabet $\{0, Y, B\}$ of length $2C + 1$. Basically, we encode the colors in *unary*.

For a given tileset τ we now consider the bar obtained by concatenating all pair of words (n_t, s_t) for tiles $t \in \tau$, adding to them two blank symbols to each extremity. Fig. 11 gives an example. This last bar will be called the *content*.

Now that we have defined all the bars, we can explain how the simulation works. The idea is all these bars combine to form n polyominoes, represented in Fig. 12, that acts like the n original Wang tiles.

The north and east side are encoded in unary above the polyomino, and the west and south side below, so that from an original tiling of the plane by the Wang tiles, we can obtain a tiling by the polyominoes, and thus by the bars. An example is depicted in Fig. 13. It should now be clear that if there is a tiling by the original tileset, then there is a tiling by the set of Wang bars.

We will now explain why a tiling by the bars has to be of this form. We will prove only that it must look as depicted in Fig.12, the rest of the proof being easy.




Figure 10: The 27 next bars.




Figure 11: How the last bar is built

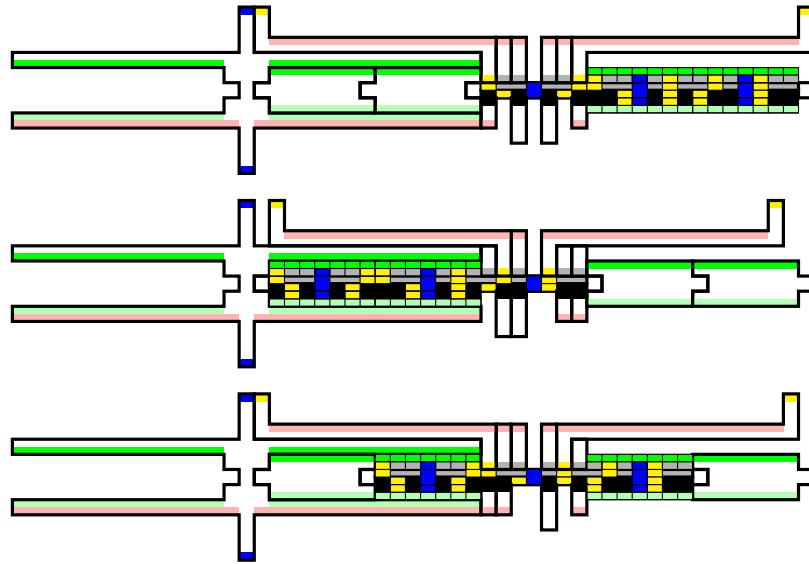


Figure 12: How the bars can represent the original Wang tiles




Figure 13: How a tiling by the Wang tiles becomes a tiling by the bars

The different steps in the proof are depicted in Fig.14. We will first explain why the bar labeled *A* must be there, then the bar labeled *B*, etc.

First we look at all the bars, et we prove that we need to use the box. It is quite clear:

- The fillers must be attached to the box
- The first four polyominoes from the first set must be attached to the content, and the last ones to the box
- The handles, due to the yellow color, must be attached either to the content, or to a (green,yellow) bar that can only be attached to the box.
- Finally, the content, due to the blue color, cannot be attached to itself, so it must either be attached to the small (blue, light green) bar (that in turn must be attached to the box), or to the box directly.

In all cases, we have seen that any bar is linked, sooner or later, to the box.

So we can now suppose that a box appear somewhere (*A* on Fig.14)

We know look at how to fill its upper right half. If we do not use a handle, then the only possibility is to use the fourth tile from Fig.9. But we will obtain a row of at least $(2C + 1)(n - 1)$ black colors, and the only bar with a black color on its south side (the content) has at most $2C$ consecutive black colors. So there must be a handle (*B* on Fig.14).

Now we look at how to fill the right part of the box. It is clear it can be only filled with fillers or with the content (attached to small tiles). Now we can use at most $(n - 1)$ fillers: If we use one more, this filler will be outside of the box, which is not possible as it will collide with the handle.

So there must be somewhere a content (*C* and *D* on the figure). Due to the handle, the content cannot be entirely inside boxes. So some of it is outside, and in particular there is a blue color on the north side outside. The only thing we can attach to this blue color is a box (*E* on the figure). In particular, there can be only **one** blue color appearing outside a box, which means the only possible situation is the one depicted in Fig. 14.




Figure 14: How to prove the bars must behave correctly.

It is then clear that this situation must force the bars (and therefore the tiles they simulate) to be aligned correctly, so that the proof is done. We can now state the theorem and its corollary.

Theorem 3.1. *The tileability problem is undecidable for 44 bars.*

Corollary 3.1. *The domino problem is undecidable for tilesets of parameter 43.*

It is quite clear that the parameter 43 is not optimal. There are indeed a few way to gain bars in the construction. For instance, the two bars with upper side v and z can be identified without changing anything. However we see no need in this article to insert any reference to The Hitchhiker's Guide to the Galaxy, and will stay with a parameter of 43.

4 Tileability with 2 Wang bars is decidable

In this last section, we will prove that we can decide whether a set of two Wang bars tiles the plane. More precisely

Theorem 4.1. *A set of two Wang bars B tiles the plane iff there is a periodic tiling by B .*


The rest of the section is devoted to the proof.

We start from two Wang bars α and β , et we will assume that there exists a tiling with α and β , but no periodic tiling, and reach a contradiction.

First, we see that the colors on the east and west side of the bars have to be the same: Otherwise, either we cannot put α and β on the same line (so that every tiling must be periodic horizontally of period $lcm(|\alpha|, |\beta|)$, hence there is a periodic tiling), or we must always put α after β and β after α on the same line (in which case every tiling is periodic horizontally of period $|\alpha| + |\beta|$, hence there is a periodic tiling). So, as in the previous section, the only constraints to put the two bars together are vertical constraints, and we can see the bars as two words $\alpha = (\alpha_N, \alpha_S)$, $\beta = (\beta_N, \beta_S)$.

W.l.o.g., we will assume $|\alpha| \geq |\beta|$. We can also suppose that we cannot tile using only α (resp. β), as this implies that there exists a periodic tiling.

Lemma 4.1. *If the following pattern appears in some tiling, then there exists a periodic tiling*




where γ is one of the two bars α or β .

Proof. It is quite clear from the picture. Let n be the length of the bars, and p the horizontal distance between the lower left bar and the upper bar.


Looking at the upper bar, we write $\gamma_S = uv$ where u is of size $n - p$ and v of size p . Looking at the two other bar, we see that the last p symbols of γ_N are u , and the first $n - p$ symbols of γ_N are v , so that $\gamma = (vu, uv)$. So a row containing only γ will have the same upper and lower side, up to a shift of p , and hence we can obtain a periodic tiling by repeating this row with a suitable shift every time. ■

Lemma 4.2. *If there is a tiling where the following pattern (and its horizontal symmetry) does **not** appear, then there exists a periodic tiling.*




Proof. We consider this tiling, and $p \geq 1$ the minimal number of consecutive bars of type α in such a tiling. (it is clear that p cannot be infinite. Otherwise this imply that a bar of type α can be followed (or preceded) at most once by a bar of type β on each row, hence by the same argument as Lemma 1.1, there exists a tiling where this doesn't appear, i.e. where every row consists either entirely of bars of type α or entirely of bars of type β , which would imply the existence of a periodic tiling).


Now we look at a position where this minimum number appears. We take $p = 2$ for the picture.




By definition of p , these tiles are surrounded horizontally by bars of type β .




Now by the hypothesis of the lemma, all bars above the bars of type α must be β tiles:



Now we look at the bar we have to put above the lower rightmost bar of type β . If it is of type β , we obtain the pattern from the previous lemma and we are done. Otherwise it is of type α , and by the minimality of p , there are at least p such bars:




To finish, we have again by the hypothesis of the lemma, that the bars below it are of type β :




Now we look at the picture without the bottom leftmost bar of type β . We obtain a pattern of k bars of type β followed by p bars of type α on the first row, and of p bars of type α followed by k bars of type β on the second row (it is easy to see it must be the exact same k). Now this pattern can be repeated periodically: The last few symbols of the north part of the bar β are indeed the same as the first symbols of the south part of the bar β , as is witnessed by the bottom leftmost bar of type β we just deleted. Hence this pattern of k bars of type β followed by p bars of type α can be repeated periodically horizontally and vertically up to shift, and we obtain a periodic tiling. ■

So we are done to our last case: The following pattern appears somewhere:




We look at some occurrence of this pattern. The lower left bar of type α is perhaps preceded by some other bars of type α so that there are p_1 of them in total, and the same goes for the upper right bar, which might be constituted of p_2 bars of type α . Upto a rotation of the bars, we suppose that $p = \min(p_1, p_2) = p_1$. We will now not care at all about the remaining $p_2 - p$ bars of type α that might appear on the upper right. The figure is now as follows (we take $p = 2$):



(The leftmost bar of type α must of course be preceded by a bar of type β by the definition of p)


Now to complete each bar of type α , we need to put at least one bar of type β , otherwise we would obtain the figure from Lemma 4.1. We thus put q_1 bars of type β at the top left of the picture, and q_2 bars at the bottom right. There are two cases:

- There are too many bars of type β on both side, so that they completely cover the p bars of type α . Let $q < \min(q_1, q_2)$ be such that the first $q - 1$ bars of type β will not cover all bars of type α but q such bars will. We now have the following picture (with $q = 3$)




But by the same reasoning as the previous lemma, this implies that we can tile using q bars of type β and p bars of type α periodically

- There is at least one side on which there are not enough bars of type β to cover the bars of type α . We take again $q = \min(q_1, q_2)$ and suppose that $q = q_1$ (the other case is similar, as our argument will not use the bottom left bar of type β), and we completely forget about the $q_2 - q_1$ remaining bars of type β . We take $p' \leq p$ so that the $p' - 1$ th bar is entirely covered by the bars of type β but not the p' th, to obtain the following picture ($p' = 1$ on the picture)



Now by definition of q , there must be a α -bar on the top left:



By the same argument as the previous lemma, this implies that we can tile using q bars β followed by p' bars α periodically in each row.

The proof is now finished, and in all cases, we proved that there exists a periodic tiling.

5 Conclusion

There is no algorithm to decide whether a set of 44 Wang Bars tiles the plane. On the other hand, a set of 2 Wang Bars tiles the plane if and only if it tiles it periodically. It would be interesting to reduce the gaps between these two values. It is clear that 44 is not optimal, and that we can reduce it. While we did not work out the details, it is safe to assume we can obtain something around 35 using the same idea. However, it is much more interesting to try to prove that the problem is decidable for 3 or 4 Wang bars. While it is our opinion that the result holds, this situation became dramatically more complex with even 3 Wang Bars, and we do not know of any approach to solve this problem.

Acknowledgements The first author thanks Daniel Gonçalves and Pascal Vanier for some interesting discussions that lead to the proof of the case of 2 Wang Bars.

References

[1] Robert Berger (1964): *The Undecidability of the Domino Problem*. Ph.D. thesis, Harvard University.

[2] Bruno Durand (1999): *Tilings and Quasiperiodicity*. *Theoretical Computer Science* 221(1-2), pp. 61–75, doi:10.1016/S0304-3975(99)00027-4.

[3] Karel Culik II (1996): *An aperiodic set of 13 Wang tiles*. *Discrete Mathematics* 160, pp. 245–251, doi:10.1016/S0012-365X(96)00118-5.

[4] Jarkko Kari (1996): *A small aperiodic set of Wang tiles*. *Discrete Mathematics* 160, pp. 259–264, doi:10.1016/0012-365X(95)00120-L.

[5] Nicolas Ollinger (2009): *Tiling the Plane with a Fixed Number of Polyominoes*. In: *Proceedings of the Third International Conference on Language and Automata Theory and Applications (LATA)*, LNCS 5457, pp. 638–647, doi:10.1007/978-3-642-00982-2_54.

[6] Hao Wang (1961): *Proving theorems by Pattern Recognition II*. *Bell Systems technical journal* 40, pp. 1–41.