Universal Tilings

Grégory Lafitte[1] and Michael Weiss[2,⋆]

[1] Laboratoire d'Informatique Fondamentale de Marseille (LIF), CNRS – Université de Provence, 39, rue Joliot-Curie, F-13453 Marseille Cedex 13, France
[2] Centre Universitaire d'Informatique, Université de Genève, 24, rue Général Dufour, 1211 Genève 4, Switzerland

Abstract. Wang tiles are unit size squares with colored edges. To know if a given finite set of Wang tiles can tile the plane while respecting colors on edges is undecidable. Berger's proof of this result shows the equivalence between tilings and Turing machines and thus tilings can be seen as a computing model. We thus have tilings that are Turing-universal, but there lacks a proper notion of universality for tilings. In this paper, we introduce natural notions of universality and completeness for tilings. We construct some universal tilings and try to make a first hierarchy of tile sets with a universality criteria.

1 Introduction

Tilings were first introduced by Wang [Wan61]. A tile is a unit size square with colored edges. Two tiles can be assembled if their common edge has the same color. A finite set of tiles is called a tile set. To tile consists of assembling tiles from a tile set on the grid $\mathbb{Z}^2$.

Wang was the first to conjecture that if a tile set tiles the plane, then it tiles it in a periodic way. Another famous problem is to know whether a tile set tiles the entire plane or not. It is called the *domino problem*. Berger proved the undecidability of the domino problem by constructing an aperiodic set of tiles, *i.e.*, a tile set that can generate only non-periodic tilings [Ber66]. Simplified proofs can be found in [Rob71] and later [AD96]. As a corollary, Berger's result shows that Wang's conjecture is false. The main argument of this proof was to simulate the behaviour of a given Turing machine with a tile set, in the sense that the Turing machine M stops on an instance ω if and only if the tile set $\tau_{\langle M,\omega\rangle}$ does not tile the plane. Hanf and later Myers [Mye74, Han74] have strengthened this and constructed a tile set that has only non-recursive tilings.

From this, we have that there exists tile sets that can generate only non-periodic tilings of the plane, and others, that can generate only non-recursive tilings of the plane. But in all those cases, Durand [Dur99] proved that if a tile set tiles the plane, then it can tile the plane in a quasiperiodic way, *i.e.*, a tiling where each pattern appears regularly along the tiling.

⋆ This author has been supported by the FNS grant 200020-105515.

W. Thomas and P. Weil (Eds.): STACS 2007, LNCS 4393, pp. 367–380, 2007.

The core main result from Berger's theorem is that tilings are Turing equivalent and so, there is a notion of computability for tilings. The tilings' computability can be studied from different points of view. One of them is universality. A universal Turing machine is a machine that can simulate all other Turing machines.We have to move away from Turing universality and see more specifically how can universality be defined for tilings. Asking the question of universality for tilings is equivalent to asking the question of how a tiling computes because we expect from a universal tiling to be able to give back the process of computation of each tiling. So, how can a tiling simulate another tiling?

In order to answer our question, we introduce different notions to tackle the search for natural universality. The main thing about tilings' computability consists in the way the tiles are assembled to tile the plane. So, we can see a given tiling T from different scales. The lowest scale is the tiling. By partitioning a given tiling into rectangles of same size, we obtain another tiling, where the tiles are the rectangular patterns. It is the tiling T seen from a higher scale. By this principle, we define strong reductions and weak reductions. By those reductions, we reach our goal : we have a notion of simulation.

We are now able to define universality and completeness for tilings. Since we have two kinds of reductions (*weak and strong*) and that a tiling can be universal for all the tilings (*dense*), or just for one tiling for each tile set (*sparse*), then we have four kinds of universality.

The main goal of this paper is to show the existence or the non-existence of the different universality notions. We will construct a complete tile set, *i.e.*, a tile set such that for each tiling P, our complete tile set can generate a tiling Q with $P \preccurlyeq Q$. We prove that a tile set is complete if and only if it generates a universal weak tiling.

The main result is to construct a strong universal tiling that can simulate, for each tile set τ, a countable infinity of tilings from τ. That is the strongest result we can obtain, because having a strongly dense universal tiling, *i.e.*, a tiling that simulates all tilings, is impossible, because there is a countable set of reductions but an uncountable set of tilings. Those results yield a classification on tile sets that is actually a hierarchy.

In the first part, we introduce the different notions used along the paper (pattern sets, reductions, completeness, universalities) and prove some obvious results from those definitions. The second part concentrates on completeness, by relating it to universality, and by constructing our first complete tile set. The last part explains universality, by constructing weak and strong universal tilings, and by giving some results following from those. Lastly we construct a non-periodic tile set that generates a universal tiling and propose a classification of the different universality and completeness notions seen along the paper.

2 Notions of Tilings' Universality and Completeness

First, we give the definitions of the basic notions of tilings. A tile is an oriented unit size square with colored edges from C, where C is a finite set of colors. A tile set is a finite set of tiles. To tile consists of placing the tiles of a given tile set on the grid $\mathbb{Z}^2$ such that two adjacent tiles share the same color on their common edge. Since a tile set can be described with a finite set of integers, then we can enumerate the tile sets, and τ_i will designate the i^{th} tile set.

Let τ be a tile set. A tiling P generated by τ is called a τ-tiling. It is associated to a tiling function f_P where $f_P(x, y)$ gives the tile at position (x, y) in P. In this paper, when we will say: "Let P be a tiling", we mean that there exists a tile set τ such that P is a τ-tiling of the plane. The set of all tilings is $\mathfrak{T}$. A pattern m is a finite tiling. If it is generated by τ, we call it a τ-pattern. A pattern m is associated to a tiling partial function f_m, defined on a finite subset of $\mathbb{N}^2$ such that $f_m(x, y)$ gives the tile at the position (x, y) in m. A finite set of rectangular τ-patterns is a τ-pattern set. We explain how to tile the plane with a pattern set.

Definition 1. *A pattern set M of $\{a \times b\}$ τ-patterns tiles the plane if there exists $0 \leq c < a$, $0 \leq d < b$ and a function $f_{P_M} : \{\, c + k_1 a \,|\, k_1 \in \mathbb{Z} \,\} \times \{\, d + k_2 b \,|\, k_2 \in \mathbb{Z} \,\} \longrightarrow M$ such that the function $f^{\tau}_{P_M} : \mathbb{Z}^2 \longrightarrow \tau$ defined by: $f^{\tau}_{P_M}(c+k_1a+x, d+k_2b+y) = f_{(f_{P_M}(c+k_1a,d+k_2b))}(x, y)$ for all $0 \leq x < a$ and $0 \leq y < b$, is a τ-tiling function of the plane. Here, $f_{(f_{P_M}(c+k_1a,d+k_2b))}$ is the tiling partial function of the pattern $f_{P_M}(c + k_1a, d + k_2b)$. With the same notation as above, we define the function $s_{P_M} : \mathbb{Z}^2 \rightarrow M$ by $s_{P_M}(k_1, k_2) = f_{P_M}(c + k_1a, d + k_2b)$.*

This definition explains in a formal way what we expect intuitively. To tile the plane with a pattern set consists in putting the patterns side by side in a subgrid of $\mathbb{Z}^2$ in such a way that color matching is respected. By analogy with tilings, we say that P_M is a M-tiling and it is associated to the pattern tiling function f_{P_M}.

The second function in the definition, s_{P_M}, is another way to define P_M: if $s_{P_M}(x, y)$ gives the pattern m_i, then $s_{P_M}(x+1, y)$ gives the pattern that touches m_i on its east side in P_M. The same is true for the south, north and west side of a pattern in P_M.

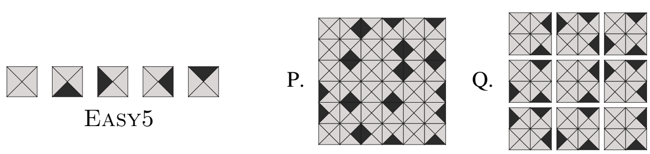

Fig. 1. An EASY5-tiling P and a pattern tiling Q extracted from it

In the previous definition, we used a tiling function $f^{\tau}_{P_M}$. It can be associated to a τ-tiling P because $f^{\tau}_{P_M}$ is defined from $\mathbb{Z}^2$ to τ and respects the color matching. So, we have a strong connection between P and P_M because both of them give rise to the same geometric tiling Q (by geometric tiling, we mean the grid $\mathbb{Z}^2$ filled with the appropriate tiles). That is one of the main notions of this paper. A given tiling can be seen from different heights. Here, P is the smallest height, the unit size level. P_M is a higher level, the $\{a \times b\}$ level. Obviously, a tiling of the plane can be seen with an infinity of different scales. From this infinity of different ways to see a tiling, we expect to obtain a notion of universality for tilings.

More formally, let M be a τ-pattern set, Q be a M-tiling and P be a τ-tiling. We say that Q is extracted[1] from P if $f_Q^{\tau} = f_P$ (figure 1). P describes the geometric tiling with unit size squares while Q describes it with $\{a \times b\}$-patterns.

From the above notions, we are able to define intuitively what we mean by tilings' reductions. From a given tiling P, we can extract an infinity of pattern tilings. Let P' be one of them. By definition, there exists a pattern set M such that P' is a M-tiling. M is a finite set of $\{a \times b\}$ patterns. We can associate to M a tile set τ with a function $R : M \to \tau$ such that two patterns m_1 and m_2 match if and only if the tiles $R(m_1)$ and $R(m_2)$ match in the same way. It can be easily shown that any pattern set can be associated to a tile set with this property. With R we can build a τ-tiling Q defined by: $f_Q(x, y) = R \circ s_{P'}(x, y)$. Since P' is extracted from P, and since Q works as P' works, so Q can be thought to be "easier" than P. That is our idea of reduction, and now we define it more formally:

Definition 2. *Let P be a τ-tiling and Q be a τ'-tiling.*

Q reduces strongly to P, denoted by $Q \trianglelefteqslant P$ if there exists a set of τ-patterns M, a M-tiling P' extracted from P and a function $R : M \longrightarrow \tau'$ such that $\forall (x, y) \in \mathbb{Z}^2$, $R \circ s_{P'}(x, y) = f_Q(x, y)$. R is called the reduction core function and its size is the size of the patterns of M, i.e., size(R)$\in \mathbb{N}^2$.

Q reduces weakly to P,denoted by $Q \preccurlyeq P$, if there exists a set of τ_i-patterns M, a M-tiling P' extracted from P and a function $R : M \longrightarrow \tau'$ such that for any $\{p \times q\}$-pattern m of Q there exists $a, b \in \mathbb{Z}^2$ such that $m(x, y) = R \circ s_{P'}(a + x, b + y)$ for all $0 \leq x < p, 0 \leq y < q$.

When we want to specify R we denote the reduction by $\trianglelefteqslant^R$ or $\preccurlyeq^R$.

As we have seen above, if $Q \trianglelefteqslant P$ then we can extract from P a pattern tiling P' that simulates Q in the sense that the patterns of P' represent the tiles of Q. We will say that Q is the tiling *associated* to the pattern tiling P' or that P' *simulates* Q. The important thing is that different patterns of P' can represent the same tile of Q. But the converse is impossible by definition (a pattern of P' cannot represent different tiles of Q). Concerning the weak reduction, if $Q \preccurlyeq P$ then we can extract from P a pattern tiling P' such that all patterns of Q are simulated somewhere in P'.

We can extend naturally our strong reduction definition to patterns. Let A be a τ-pattern and B be a τ'-pattern. We say that $A \trianglelefteqslant B$ if we can extract from B a M-tiling B' that simulates A, where M is a set of τ'-patterns. For patterns,

[1] We note that in most papers on tilings the word *extract* is already used in the following sense: if a tile set can tile square patterns of ever-increasing sizes, then it can tile the plane; one can *extract* from this set of patterns a tiling of the plane. In this paper, extract will, most of the time, refer to the pattern tiling taken from a given tiling. It is specified when we use *extract* to mean "extraction of a tiling from a set of patterns of ever-increasing sizes".

the weak reduction is equivalent to the strong reduction since weak reduction is locally equivalent to strong reduction.

Those reductions have the following property:

Lemma 1. *$\trianglelefteq$ and $\preccurlyeq$ are preorders on $\mathfrak{T}$, where $\mathfrak{T}$ is the set of all possible tilings.*

We do not have an order on $\mathfrak{T}$, because the antisymmetric property is not respected: we can find two tilings P and Q such that $P \trianglelefteq Q$ and $Q \trianglelefteq P$ but $P \not\equiv Q$ (By $P \not\equiv Q$ we mean that P and Q are not the representation of the same tiling up to a color permutation, or, in an equivalent way, that there is no trivial reduction of size $(1,1)$ between P and Q).

With the definition of reduction, we can now define the notions of completeness and universality.

Definition 3. *Let $\mathcal{A}$ be a set of tilings. A tile set τ is $\mathcal{A}$-*complete *if for any tiling $P \in \mathcal{A}$, there exists a τ-tiling Q such that $P \trianglelefteq Q$. If $\mathcal{A} = \mathfrak{T}$ then τ is called* complete*.*

This completeness notion is natural in the sense that it corresponds to what one would expect: any tiling can be reduced to some instance tiling of our complete tile set in such a way that to answer any question about our tiling it suffices to study the instance tiling of our complete tile set.

We expect from a universal tiling to have in its construction much of the information of all the other tilings. For tilings, we have different ways to define the information contained in a given tiling and can distinguish mainly two kinds of information for a tiling. The first, and the most natural, is the tiling itself (how it is built). The second consists in studying the different patterns that appear in the tiling. Those two ways to consider a tiling's information give rise to two ways to consider universality. Does a tiling contain enough information to explain the behaviour of all other tilings (we call it *strong dense universality*), or only the behaviour of a tiling for each tile set (*strong sparse universality*)? Does a tiling contain enough information to simulate all the patterns of any tiling (*weak dense universality*) or only the patterns of a tiling for each tile set (*weak dense universality*)?

With this motivation, we have the following definitions:

Definition 4. *Let P_u be a τ-tiling. P_u is:*

- strongly dense universal *if for any tiling Q, $Q \trianglelefteq P_u$,*
- strongly sparse universal *if for all τ', there exists a τ'-tiling Q, such that $Q \trianglelefteq P_u$,*
- weakly dense universal *if for any tiling Q, $Q \preccurlyeq P_u$,*
- weakly sparse universal *if for all τ', there exists a τ'-tiling Q, such that $Q \preccurlyeq P_u$.*

We have the following properties:

Lemma 2. *1. P_u strongly (resp. weakly) dense universal $\Rightarrow$ P_u strongly (resp. weakly) sparse universal.*

2. *P_u strongly dense (resp. sparse) universal $\Rightarrow$ P_u weakly dense (resp. sparse) universal.*
3. *$\trianglelefteq$ preserves universalities.*
4. *$\preccurlyeq$ preserves weak universality.*

We will show the existence or the non-existence of these universality notions after the following section.

3 Completeness

The following theorem shows how complete tile sets and universal tilings relate.

Theorem 1. *Assuming the existence of at least one complete tile set and one weakly dense universal tiling, we have:*

Let τ be a tile set. τ is complete if and only if there exists a weakly dense universal τ-tiling.

Proof. [$\Rightarrow$]: Let τ be a complete tile set and P_u be a weakly dense universal tiling. Since τ is complete, there exists a τ-tiling P such that $P_u \trianglelefteq P$. $\trianglelefteq$ preserves weak universality, therefore P is a weakly universal τ-tiling.

[$\Leftarrow$]: Let P_u be a weakly universal τ-tiling. By definition, for any tiling P there exists R such that $P \preccurlyeq^R P_u$. We consider the set of patterns $\{A_i\}_{i>0}$ where A_i is the $\{i \times i\}$-pattern of P centered around $(0,0)$. By definition, there exists a set of patterns $\{B_i\}_{i>0}$ of P_u such that $A_i \trianglelefteq B_i$ for all i. We can extract from the set of patterns $\{B_i\}_{i>0}$ of ever-increasing sizes, a τ-tiling P' such that $P \trianglelefteq P'$; and so, τ is complete.

We now exhibit our first complete tile set. We can easily prove that the tile set Easy5 (figure 1) is complete. In order to do this, we just have to see that we can encode any tile set with square Easy5-patterns such that the Easy5-patterns have a code on their borders that represent the tiles of the tile set. Then, we just have to assemble the patterns in the same way that the tiles, that they represent, are assembled in the tiling.

Theorem 2. Easy5 *is complete.*

We have a stronger result for complete tile sets: if τ is a complete tile set, then for each tiling P, there exists an uncountable set of τ-tilings to which P reduces.

It would be interesting to find a non-trivial complete tile set, *e.g.*, a complete tile set that has only non-periodic tilings. We will construct such a tile set in theorem 6. For now, we construct a more complex complete tile set that we will use later on.

Since Berger's proof of the undecidability of the domino problem, it is known that we can simulate a Turing machine with a tiling. We briefly recall how to do this.

Some tiles are used to transmit a symbol of the alphabet of the Turing machine, some are used to show that the state q_i will act on the symbol a_j at the next step and finally, some are used to represent the transitions of the Turing machine. More details of this construction can be found in [AD96].

We will now build a Turing machine M such that the space×time diagram of the computation of $M(\omega)$ gives a rectangular pattern that simulates a tile from the tile set $\tau_{|\omega|}$ (the $|\omega|^{th}$ tile set).

Our Turing machine works with a semi-infinite tape. A typical input is $\omega = x\$^n$ where $x \in \{0,1\}^*$ and n depends on $|x|$. The length of ω represents the code of the tile set we are working with, and the first part of ω, x, is the code of a color of the set of color of $\tau_{|\omega|}$. We know that we can encode a tile set, *i.e.*, by giving it an unique number or code in a same way that we do for Turing machines. So the first step consists in decoding $|\omega|$ to find the different tiles that compose $\tau_{|\omega|}$. Then we check if x is the code of a color of $\tau_{|\omega|}$. If yes, we choose in a non-deterministic way a tile t of $\tau_{|\omega|}$ such that the color of its south side is x. We can build our Turing machine such that after m steps of computation, the k next steps of the computation are used to write the code of the west/east[2] color of t (n and k depending only on $|\omega|$), *i.e.*, in the space×time diagram of the Turing machine, the first column from time $m+1$ to $m+k$ represents the code of the color of the west side of t. The tiles that are not between the $(m+1)^{th}$ and the $(m+k)^{th}$ lines are all the blank tile $\square$. We do the same for the east/west side (depending on the first choice we made), *i.e.*, in the space×time diagram of the Turing machine, the last column from time $m'+1$ to $m'+k$ represents the code of the color of the east side of t (m' depending only on $|\omega|$). The p last steps of computation are used to write the code of the color of the north side of t completed with \$'s and $\square$'s. For later usage, we precise that our Turing machine M does nothing when the entrance is the empty word. This means that its corresponding space×time diagram will be a rectangular space filled by blank tiles.

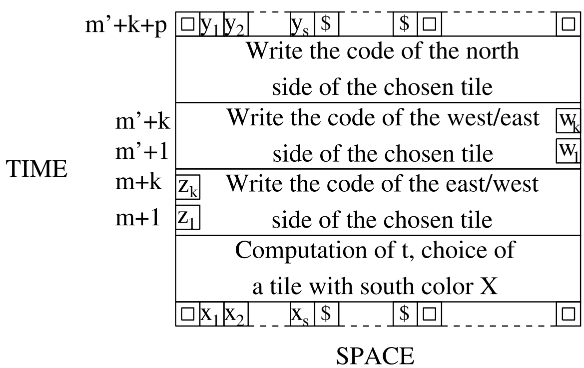

Fig. 2. The space×time diagram of M

The figure 2 represents the space×time diagram of our Turing machine. In addition, we construct the Turing machine such that for two inputs of size p the

[2] Our algorithm chooses in a non-determistic way either the west or the east side. This non-determinism is essential for east/west matching of two diagrams.

machine uses exactly $s(p)$ of space and $t(p)$ of time. With this construction, we are certain that two space×time diagrams have the same size if and only if the inputs have the same lengths. That guarantees that the simulations of two tiles from the same tile set give two diagrams of the same size. So, two diagrams will match on their north/south side if and only if the two tiles that they simulate match on their north/south side. For the east/west border the match rules are different. During the computation, we can choose to write first the east color or the west color. Then, two diagrams will match on their east/west border if and only if the tiles they represent match on their east/west border, and if during the computation, the two Turing machines they represent have done different non-deterministic choices.

The idea is now to associate this Turing machine to its corresponding tiling, called τ_u. By construction, τ_u generates patterns that correspond to the space× time diagram of the simulation of a tile. We have a tiling that gives $\{s \times t\}$-patterns such that two patterns match if and only if they represent two tiles, from the same tile set, that match. Hence, we can simulate with this tile set any behaviour of any tile set. Therefore, τ_u is complete.

4 Universality

We now study the different universality notions defined above. We give some results about universality before constructing our first universal tiling.

Theorem 3. *1. If $P_u \preccurlyeq P$, where P_u is a strongly universal tiling and P is a τ-tiling, then there exists a strongly universal τ-tiling.*
2. Let P_u be a tiling. If P_u is strongly universal, then P_u is non-recursive.

Proof. **1.** If $P_u \preccurlyeq P$, then there exists a reduction R such that for any $\{n_1 \times n_2\}$-pattern A of P_u there exists a $\{m_1 \times m_2\}$-pattern B of P such that $A \trianglelefteq^R B$. We consider the set of patterns $\{A_i\}_{i>0}$ where A_i is the $\{i \times i\}$-pattern of P_u centered around $(0,0)$. So, there exists a set of patterns $\{B_i\}_{i>0}$ of P_u such that $A_i \trianglelefteq B_i \,\forall\, i$. Thus, we can extract from $\{B_i\}_{i>0}$ a τ-tiling Q such that $P_u \trianglelefteq^R Q$. Since $\trianglelefteq$ preserves universality, then Q is a strongly universal tiling.
2. Let P_u be a strongly sparse universal τ-tiling. We will prove that P_u is non-recursive. By Hanf and later Myers [Mye74, Han74], we know that there exists tile sets that produce only non-recursive tilings. Let τ' be such a tile set. Suppose that P_u is recursive, *i.e.*, f_{P_u} is recursive. Let P_{nr} be a non-recursive τ'-tiling such that $P_{nr} \trianglelefteq P_u$. Let $\{R_i\}_{i\geq 1}$ be the family of reduction core functions from a set of τ-patterns to τ'. $\{R_i\}_{i\geq 1}$ is enumerable. If f_{P_u} is recursive, we can compose it with the reductions R_i and obtain the recursive tiling functions of all τ'-tilings that reduce to P_u. By definition P_{nr} reduces strongly to P_u, thus one of those recursive tiling functions defines P_{nr}. This is a contradiction, hence P_u is non-recursive.

The first result we obtain, concerning the different universality notions, is the non-existence of strongly dense universal tilings. This is due to a countability argument. We only have a countable set of possible reductions for a given tiling, but an uncountable set of tilings.

Theorem 4. *Strongly dense universal tilings do not exist.*

We now study the weak version of universality for tilings. The idea of the construction is to build our tiling in the same way that we can construct a Turing machine simulating every step of all the Turing machines by simulating at step i the first i computing steps of the first i Turing machines. Similarly, we construct a weakly dense universal tiling that enumerates all possible patterns of all tilings. They are countable so we can simulate all of them in the same tiling. Thus, we obtain the following result:

Theorem 5. *There exists a weakly dense universal tiling.*

Of course, weak dense universality implies weak sparse universality.

We still have a last universality notion to study: strong sparse universality, *i.e.*, a tiling that can simulate at least one tiling for each tile set. We can still use the Easy5 tile set to show the non-emptiness of this class. In the following theorem, we propose a non-periodic tile set that will generate a strongly sparse universal tiling, that is universal in a more "natural" way than Easy5.

Theorem 6. *There exists a non-periodic tile set that generates a strongly sparse universal tiling.*

Proof. The idea is to simulate a Turing machine in an aperiodic tiling. For this, we use Robinson's tiling. We give some explanations on how to force the computation of a Turing Machine in Robinson's tiling. We again refer the reader to [Rob71] and [AD96] for a detailed construction.

1.

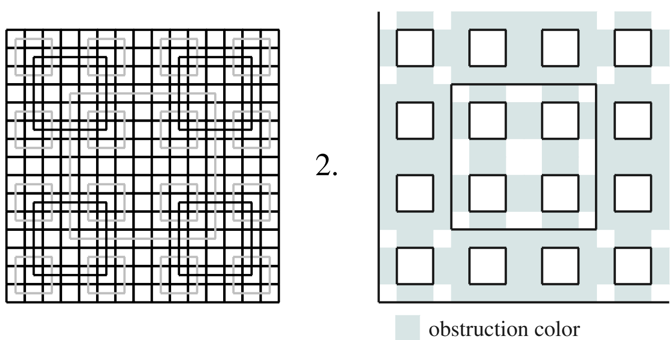

Fig. 3. The hierarchical structure and the obstruction zone in Robinson's tiling

The main idea is to use the hierarchical pattern construction that is generated by Robinson's tiling. It gives rise to $2^n - 1$ square patterns (figure 3.1). We note that two squares of sizes $2^{2n} - 1$ and $2^{2m} - 1$ cannot intersect if $n \neq m$. The idea is to compute a given Turing machine in the squares of size $2^{2n} - 1$.

In order to do this, the squares of size $2^{2n}-1$ send an obstruction color outside their borders. That is an intuitive fact that we can do this with a tiling but it is a quite technical result to prove. Figure 3.2 shows the hierarchical construction with the obstruction tiles.

In this figure there appears white spaces, *i.e.*, areas without obstruction color spaces. Those spaces are called free. Those white spaces are the computation spaces, where the Turing machine M will be simulated. On each free tile, we superimpose the tiles representing our Turing machine as we explained in section 3. The tiles that are obstructed in one direction (horizontally or vertically) will transmit the information of the computation of our Turing machine in the other direction. By construction, the tiles that are obstructed in both direction have no information to transmit. Thus, we can impose to the lowest free spaces of a square to start the simulation of our Turing machine on an empty tape with a given state. Then, by transmission of the information horizontally and vertically, we will simulate in each square of size $2^{2n}-1$, the first steps of our Turing machine.

We now take the Turing machine M described in figure 2. We built it to force that on any input of length n, if M stops, then it stops using exactly $s(n)$ spaces and $t(n)$ steps. We modify it to permit that $t(n) = k \times s(n) = k \times n^2, \forall n$. Since we work in linear space and time, we can modify M to satisfy those conditions.

We also modify Robinson's tiling by simulating any tile of Robinson's tiling with patterns of size $\{1 \times k\}$. The k is the same constant that relates the time function and the space function of M. The tiling that we obtain is Robinson's tiling stretched horizontally with a factor k. Thus, in each square of size $2^{2n}-1$, we will have the equivalent of a square of size $n^2 \times (k\,n^2)$ of free tiles.

Now, we just have to simulate our Turing machine M in these spaces. We can force the south-west free tile of any square pattern to be the tile that simulates M on state q_0. We force that the tiles that touch the south border of the square represent any of the four symbols $\{0, 1, \$, \square\}$. Then, the computation of M on this input will say if it was a correct input, and will halt in exactly $t(n)$ steps and $s(n)$ spaces if it was correct. We can force that the computation tiles match the north board of the square if and only if they are in a final state. Thus, we fill the free tiles of a square if and only if the input was $x\$^n\square^m$, such that x is the code of a color of $\tau_{|x|+n}$ and the computation uses $|x| + n + m$ spaces and halts after $k \times (|x| + n + m)$ steps to give a simulation of a tile of $\tau_{|x|+n}$.

With this construction, we fill any pattern of a given size with the simulations of some tiles from the same tile set. To guarantee that this simulation works, we modify the obstruction color sent by the squares outside their borders. We add four kinds of obstruction colors: $c_0, c_1, c_\$$ and $c_\square$, representing the four symbols $\{0, 1, \$, \square\}$. For example, the obstruction color $c_\$$ will be sent if the first tile inside the square is a computation tile, and is a tile representing the symbol \$. Thus, all the squares of a given size will represent a tiling P of a certain tile set because we have guaranteed that the matching rules were respected. Then, P reduces to our construction.

The last point consists in checking that at least one tiling for each tile set will reduce to our construction. It is the case because in our construction of M, we guarantee that each number represents a tile set, and thus, for each tile set, there exists a unique size of rectangular free spaces where the tile set will be simulated and so, any tile set that tiles the plane has a tiling that reduces to our construction. We have specified that our Turing machine does nothing when the entrance is the empty word. Its special space$\times$time diagram corresponds to the blank tile. Thus, all the free spaces of a given square size will be filled with blank tiles. It is used when a tile set does not tile the plane to guarantee that our final universal tiling will tile the plane.

In this theorem, we constructed a tiling P_1 that simulates a tiling for each tile set. For a given tile set τ_i, we can choose the τ_i-tiling that we will simulate in P_1. At a certain step of our computation, our tiling P_1 will simulate its own tile set. We can imagine that we will simulate a τ_i tiling P_2 which is a strongly sparse universal tiling such that P_2 simulates for any tile set τ_i another τ_i-tiling than the τ_i-tiling simulated in P_1. Thus, by transitivity, with this construction, P_1 simulates at least 2 tilings for each tile set. At a certain point P_2 will also simulate its own tile set, *etc.* By iterating this process, we can build a tiling that simulates for each tile set τ a countable infinity of τ-tilings.

The following theorem gives the conditions needed by a tile set to generate a universal tiling that simulates a countable infinity of tilings for each tile set.

Theorem 7. *Let τ be a tile set. If for any countable set $A = \{P_1, P_2, P_3, \ldots \mid P_j$ is a τ_j-tiling $\forall j\}$ there exists a τ-tiling P_A such that $P_j \trianglelefteq P_A$ for all $P_j \in A$, then there exists a strongly sparse universal τ-tiling P_u such that for all τ_j there exists a countable infinite set $A_j = \{P_{j_1}, P_{j_2} P_{j_3}, \ldots\}$ of τ_j-tilings such that $P_{j_k} \trianglelefteq P_u$ for all j, k. We say that a tiling has the* universal infinity property *(UIP) if it satisfies the conditions of this theorem.*

Proof. Let τ be a tile set that satisfies the hypothesis of the theorem. Since τ is a tile set, there exists i such that $\tau = \tau_i$. We consider the set A_1 composed of P_1^1, P_2^1, $\ldots$ such that, for all j, P_j^1 is a τ_j-tiling and P_i^1 is a τ_i-tiling that simulates all tilings of the set A_2. By induction, we define A_n to be composed of P_1^n, P_2^n, $\ldots$ such that, for all j, P_j^n is a τ_j-tiling and P_i^n is a τ_i-tiling that simulates all tilings of the set A_{n+1}.

If we choose the sets A_n in such a way that $A_n \bigcap A_m = \emptyset$ for all n, m, then by simulating all tilings of the set A_1, P_u will simulate a countable infinity of tilings for each tile set.

Since in theorem 6 we can choose the tiling that we want to simulate for a given tile set, Easy5 and the tile set of theorem 6 have the universal infinity property. We can see that a tile set with *UIP* has the highest class of universality. The following theorem shows that this property is equivalent to other notions mentionned above.

Theorem 8. *The following statements are equivalent:*
1. *τ has the universal infinity property;*
2. *τ is complete;*
3. *τ generates a weakly dense universal tiling.*

Proof. 1 ⇒ 2: Let τ be a tile set with the universal infinity property. Then, for any subset $A = \{P_1, P_2, \ldots\}$, where P_j is a τ_j-tiling, there exists a τ-tiling P such that $P_i \trianglelefteq P$ for all $P_i \in A$. Thus, for any tiling Q, there exists a τ-tiling P such that $Q \trianglelefteq P$. So, τ is complete.

1 ⇐ 2: Let τ be a complete tile set and τ' a tile set with the universal infinity property. Since τ is complete, for any τ'-tiling P_k there exists a τ-tiling Q_k such that $P_k \trianglelefteq Q_k$. Since the theorem 6 shows the existence of at least one tile set with the universal infinity property, then we can reduce all of its tilings to our complete tile set τ and thus, τ has the universal infinity property.

2 ⇔ 3: By theorem 1.

We have shown that completeness, generating weak dense universality and universal infinity property are equivalent. In fact, it is the finest universality class we can get, based on our reduction notion. We call this class $[UIP]$.

The class $[UIP]$ is really interesting in the sense that two tile sets τ and τ' of $[UIP]$ have the following property: for any τ-tiling P, there exists an infinity of τ'-tilings $\{Q_i\}_{i>0}$ such that $P \trianglelefteq Q_i$. In a certain way, the tile sets of $[UIP]$ generate tilings with the same behaviour.

In figure 4, we illustrate the obtained classification of the different universality and completeness notions seen along the paper.

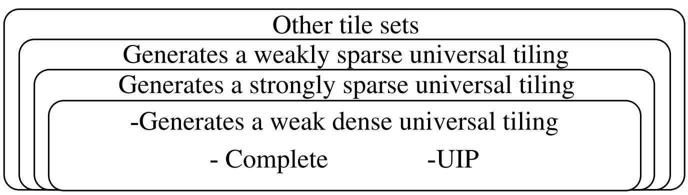

Fig. 4. The universality classification for tile sets

To clarify this classification, we aim to show that some non-periodic tile sets do not belong to $[UIP]$. To prove this, we recall the quasiperiodic function associated to a tiling. For a given tiling P, the quasiperiodic function G_P gives for each n, the smallest m such that any pattern of size n in P appears at least once in any $m \times m$-square pattern of P. Of course, G_P is not a total function for all P. We can have a tiling Q where a given pattern of size s appears only once in Q and thus, $G_Q(s)$ will not be defined. Nevertheless, Durand [Dur99] showed that if a tile set tiles the plane, then it can tile it in a quasiperiodic way, *i.e.*, the quasiperiodic function associated to this tiling is total.

We have the following results:

Theorem 9. *Let P and Q be two tilings of the plane. If there does not exist $c \in \mathbb{N}$ such that $G_Q(cn) > G_P(n)$ $\forall n$, then $P \not\trianglelefteq Q$.*

Proof. Suppose that $P \trianglelefteq Q$ with a reduction of size (a, b). That means that any $\{c \times d\}$-pattern of P is simulated by a $\{ca \times bd\}$-pattern of Q. By the theorem's condition, there exists at least one n such that $G_Q(a \times b \times n) < G_P(n)$. Thus there exists a pattern m of P of size n that appears less frequently in P than any pattern of Q of size $a \times b \times n$ appears in Q. Hence, no pattern of Q of size $a \times b \times n$ can represent the pattern m and thus, $P \not\trianglelefteq Q$.

Since [CD04], we know that there exists tile sets that generate only tilings with a non-recursive quasiperiodic function. Thus, if P is a universal strong tiling, then G_P cannot be recursive. Since Robinson's tile set gives rise only to tilings with recursive quasiperiodic functions, Robinson's tilings are not universal.

5 Concluding Remarks

We have shown that there exists a strongly sparse universal tiling that can simulate a countable infinity of tilings for each tile set. That is the strongest universality notion we can get. In fact, having a tile set with this property is equivalent to completeness and to generating a weakly sparse universal tiling. But we have also shown that there is no strongly universal tiling that simulates all tilings, because of an argument of countability.

The constructions were generated by the trivial tile set EASY5. But we also show that even non-periodic tile sets can generate universality and be complete. There remains the question: are all non-periodic tile sets, that generate only tilings with a non-recursive quasiperiodic function, complete?

Acknowledgements

We warmly thank Jacques Mazoyer for his startling remarks and all his advices, and Bruno Durand for the stimulating discussions.

References

[AD96] Allauzen (C.) et Durand (B.), *The Classical Decision Problem*, appendix A: "Tiling problems", p. 407–420. Springer, 1996.

[Ber66] Berger (R.), « The undecidability of the domino problem », *Memoirs of the American Mathematical Society*, vol. 66, 1966, p. 1–72.

[CD04] Cervelle (J.) et Durand (B.), « Tilings: recursivity and regularity », *Theoretical Computer Science*, vol. 310, n° 1-3, 2004, p. 469–477.

[CK97] Culik II (K.) et Kari (J.), « On aperiodic sets of Wang tiles », in *Foundations of Computer Science: Potential - Theory - Cognition*, p. 153–162, 1997.

[DLS01] Durand (B.), Levin (L. A.) et Shen (A.), « Complex tilings », in *Proceedings of the Symposium on Theory of Computing*, p. 732–739, 2001.

[Dur99] Durand (B.), « Tilings and quasiperiodicity », *Theoretical Computer Science*, vol. 221, n° 1-2, 1999, p. 61–75.

[Dur02] Durand (B.), « De la logique aux pavages », *Theoretical Computer Science*, vol. 281, n° 1-2, 2002, p. 311–324.

[Han74] Hanf (W. P.), « Non-recursive tilings of the plane. I », *Journal of Symbolic Logic*, vol. 39, n° 2, 1974, p. 283–285.

[Mye74] Myers (D.), « Non-recursive tilings of the plane. II », *Journal of Symbolic Logic*, vol. 39, n° 2, 1974, p. 286–294.

[Rob71] Robinson (R.), « Undecidability and nonperiodicity for tilings of the plane », *Inventiones Mathematicae*, vol. 12, 1971, p. 177–209.

[Wan61] Wang (H.), « Proving theorems by pattern recognition II », *Bell System Technical Journal*, vol. 40, 1961, p. 1–41.

[Wan62] Wang (H.), « Dominoes and the $\forall\exists\forall$-case of the decision problem », in *Proceedings of the Symposium on Mathematical Theory of Automata*, p. 23–55, 1962.