A TILING PROBLEM

NISHANT CHANDGOTIA

Abstract. In this essay, we will discuss some results in R.M. Robinson's paper "Undecidability and Nonperiodicity for Tilings of the Plane" .

1. Introduction

Suppose we are given a finite set of square tiles with coloured edges and unit length sides. Any tile mentioned in the essay will be of this type. A tiling of the plane is defined as a covering of the plane with infinite copies of the given tiles. The plane here refers to $\mathbb{R}^2$. The rules of tiling are that the edges of the tiles must be either horizontal or vertical and abutting edges must have the same colour. We are not allowed to rotate or reflect the tiles, only translation is allowed. The tiling problem is, whether or not we can tile the plane with a given set of tiles. The question was initially raised by Hao Wang in [Wang]. Let us understand his motivation.

Hao Wang was interested in analysing mathematical statements using predicate calculus. More specifically he wanted to understand provability and satisfiability of a given formula in predicate calculus. By the undecidability of the halting problem it is clear that there cannot exist a finite procedure for this decision problem. So he wished to study certain classes of formulas to say whether or not there exists a finite procedure for those formulas. One such class considered are the formulas with simple quantifier prefix AEA. In [Wang] he proved that the problem of deciding this class is the same as deciding the tiling problem mentioned above. He conjectured :

"A finite set of plates is solvable(has at least one solution) if and only if there exists a cyclic rectangle of plates, or, in other words a finite set of plates is solvable if and only if it has atl east one periodic solution "

Plates are the tiles, solvability refers to the existence of a tiling of the plane by a given set of tiles and periodic solutions refer to tilings which are invariant under a certain non-zero horizontal and vertical shift. Suppose that the conjecture is true. Then we can prove that the tiling problem is decidable.

Suppose we are given a finite set of tiles $a_1, a_2, a_3 \dots a_n$. The set can tile the plane if and only if it can tile every finite rectangle with integral dimensions. That is, it cannot tile the plane if and only if there exists a rectangle of size $M \times N$ for some $M, N > 1$ such that it cannot be tiled by the set. By the conjecture, the set can tile the plane if and only if there exists a periodic tiling or equivalently there exists a tiling of rectangle of size $M1 \times N1$ for some $M1, N1 > 1$ such that the first row is the same as the last row and the first column is the same as the last column. The machine taking the decision starts listing out all possible tilings of finite rectangles one by one. It has to eventually reach a rectangle of size greater

than 2×2 which cannot be tiled or the first row equals the last row and the first column equals the last column. By checking each tiling of the rectangles for the given two conditions that it lists out, the machine can check whether or not the plane can be tiled. Hence the tiling problem will be decidable.
This conjecture was eventually disproved by his PhD student Robert Berger in [Berger] and he also proved that the tiling problem is undecidable. Eventually the solution was simplified and was published by Raphael M. Robinson in [Robinson]. Recently, there has been further simplification by Jarkko Kari in [Kari]. He further proves that the decision problem regarding tiling the hyperbolic plane is also undecidable. This question was raised in [Robinson]. However the author is not very familiar with this work.

The essay is organised in the following way. We will begin by stating some things related to the result. Then we show an example of a set of tiles forcing non-periodicity. The essay will end with the proof of the undecidability of the tiling problem.

2. Some Things Related or Vaguely Related

We will first discuss how this helps us understand a certain physical phenomena and then try to see how it is related to Symbolic Dynamics.

C.Radin mentions in [Radin] curious relations between aperiodic tiling and crystalline structures. There the tiles considered are not necessarily rectangular or oriented along the axis. I will not discuss them since it will take us away from the topic however the reader is encouraged to have a look at this paper.

Now we will look at it from the symbolic dynamics context. Suppose we name the tiles $1, 2 \dots n$. The colouring gives us adjacency rules on what symbol can lie next to what. Notice that restricting that the centres lie on lattice points does not change the tiling problem in anyway. So with this restriction, we can look at every tiling as a map from $\mathbb{Z}^2$ to $\{1, 2, \dots n\}$ with the restrictions on adjacency of symbols decided by the colours of the corresponding tiles. The set of these maps form nearest neighbours shifts of finite type in 2- dimensions and the correspondence thus described is bijective. Nearest neighbour shifts of finite types are central objects in Symbolic Dynamics. Thus asking the question whether or not a certain set of tiles can tile the plane is same as asking whether or not the corresponding shift of finite type is empty. By Berger's Theorem(i.e. the tiling problem is undecidable), non-emptiness of nearest neighbour shift of finite type is undecidable.

For contrast, we will look into 1 dimensional nearest neighbour shifts of finite type.This corresponds to tilings of a unit strip given by $\{(x, y) | x \in \mathbb{R} \text{ and } -1 \leq y \leq 1\}$ in the same way as above. Formally, let $\mathcal{A} = \{1, 2 \dots n\}$ be a finite alphabet and $\mathcal{F} \subset \mathcal{A}^2$.

$$\text{Define a shift } X_{\mathcal{F}} = \{x \in \mathcal{A}^{\mathbb{Z}} : x_i x_{i+1} \notin \mathcal{F} \text{ for all } i \in \mathbb{Z}\}$$

Suppose $x \in X_{\mathcal{F}}$. Since the alphabet is finite, there exists $i < j$ such that $x_i = x_j$ where x_r is the r^{th} coordinate of x. Since the rules forbid only which symbols can be next to the other, this means that the sequence $x_i \dots x_j$ can be repeated indefinitely to the right and to the left. This shows a strip can be tiled if and only if there exists a periodic tiling therefore proving Wang's Conjecture for strips. Hence there is a decision method for the tiling problem on strips. This argument can be naturally extended to thicker strips. The proof actually proves far more than Wang's Conjecture. It tells us that every allowable block in the the shift belongs to a periodic point. This is quite like the pumping lemma for testing whether a language is regular.

This argument gives us a possible intuition behind Wang's Conjecture. So what goes wrong while tiling a plane? The reason is the boundary effect. One can argue in the following way, although the rows can attain periodicity, the entries force vertical restrictions which in turn come round and force aperiodicity.

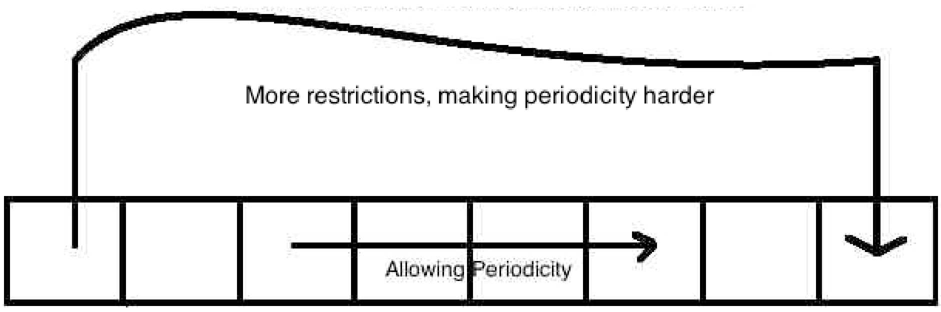

Figure 1. Boundary effects

Initially, the author had got the impression that any set of tiles would allow periodic tilings. However after trying to prove it for a long time, he was relieved to find that there exists a set of tiles which force aperiodicity.

We end this section with the following quote by C. Radin in [Radin]. *"As in the movie of the title of this article there can be a significant advantage in considering various points of view of a complicated phenomenon, and it is not surprising that the further separated the worlds from which the views originate, the more useful is the contrast."*

3. An Aperiodic Tiling

The following construction is mentioned in [Robinson]. We will start with certain parity markings which will force a certain periodic tiling of the plane. After this we will superimpose some direction tiles on to these which force a certain tile to occur only as vertices of a square. This shall force non-periodicity. We will not mention any colours till maybe the end of the this section.

We start with Parity markings. i.e. tiles marked with elements of $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$. The rules of tiling are, any tile (a, b) to the left of tile (c, d) satisfies $a = c - 1$ and $b = d$ and any tile (h, k) above (r, s) satisfies $k = s + 1$ and $h = r$. All computations are being made in $\mathbb{Z}/2\mathbb{Z}$.

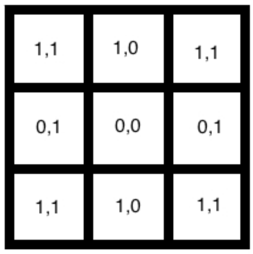

FIGURE 2. Tiling of a 3×3 area

Now consider direction tiles. They are like network signals travelling in various directions. The 5 basic direction tiles are:

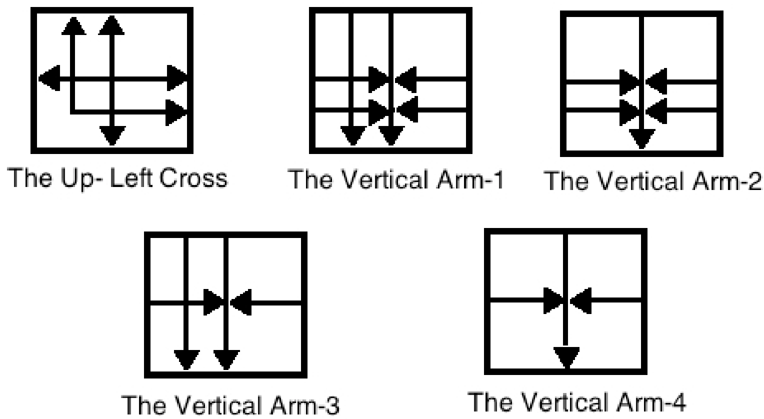

FIGURE 3. The 5 Direction Tiles

These are rotated and reflected to give us 28 tiles. There will be 4 types of crosses

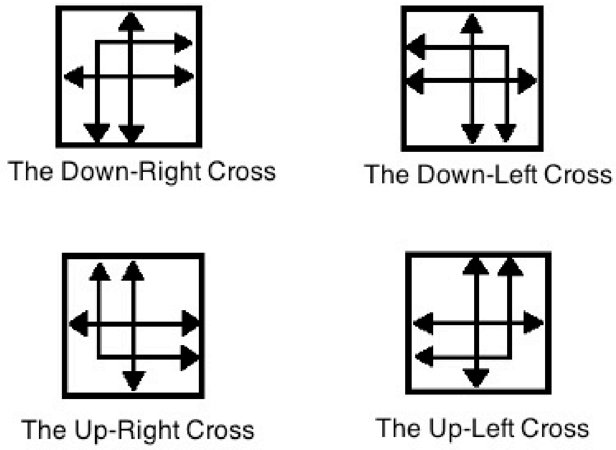

FIGURE 4. The Crosses

and a lot of horizontal and vertical arms. The arms are those tiles which have an arrow running through it with arrow tail on one edge and the arrow head on the other. The vertical arms are those with the concerned arrows oriented vertically and the horizontal arms are those which have the concerned arrows oriented horizontally.

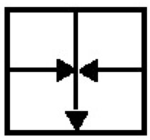

FIGURE 5. A Vertical Arm

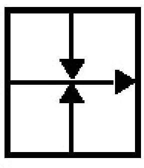

FIGURE 6. A Horizontal Arm

The rules governing the tiling by these tiles are that the directions should match.

In the sense, an arrow should meet the appropriate tail. For example,

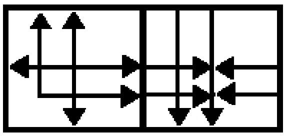

FIGURE 7

is correct, however

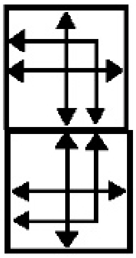

FIGURE 8

is incorrect.

Notice, that if we have two facing crosses e.g. down-right and down left in the same row, and no other crosses in between then the configuration has to look like

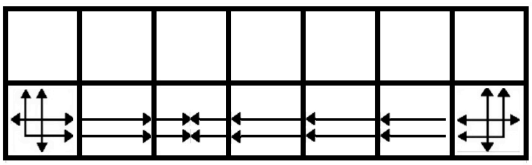

FIGURE 9

The (1,2) box would be forced to have two incoming arrows from the bottom. Browsing through all the choices available, this would mean that the (1,2) box would have a arrow-tail on its right edge.This would force the (2,2) to be a cross since it would have to give atleast 1 arrow to the (1,2) box and atleast 1 arrow to the (2,1) box. This process would continue forcing every alternate box in the upper row to be a cross with some orientation. Hence if there was some way of forcing the cross to occur at every alternate position we would get the crosses to occur as vertices of a square. This property is quite important.

All direction tiles will be superimposed with certain parity tiles and the rules of adjacency will be the superimposition of the rules of the tiles being superimposed. The crosses will be superimposed with (0,0) tile. The vertical arms will be combined with the (0,1) tile and the horizontal arm will be superimposed with the (1,0) tile. All the direction tiles will be superimposed with the (1,1) tile. This will force crosses to occur at alternate positions in alternate rows. The total number of tiles is 56.

Lemma 3.1. *The tiles can tile the whole plane. However they can only tile it aperiodically.*

Proof. Firstly let us show that it can tile the whole plane. We will hide the parity markings since in most cases it will be obvious.Suppose we are given the up-left tile with the (0,0) marking. This would force a configuration of the type

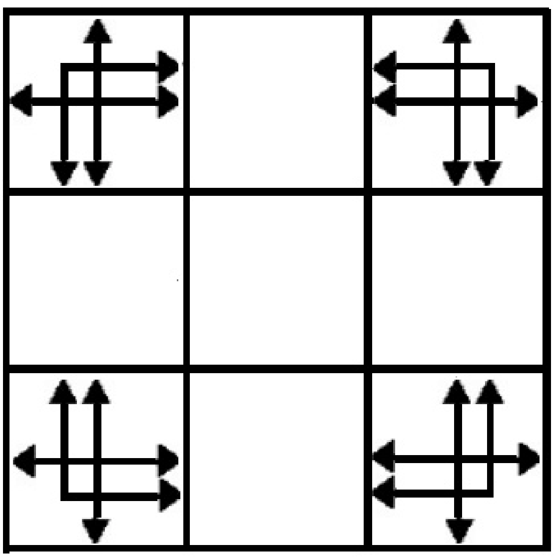

Figure 10

Putting a left up cross in the centre would again force a configuration of the type

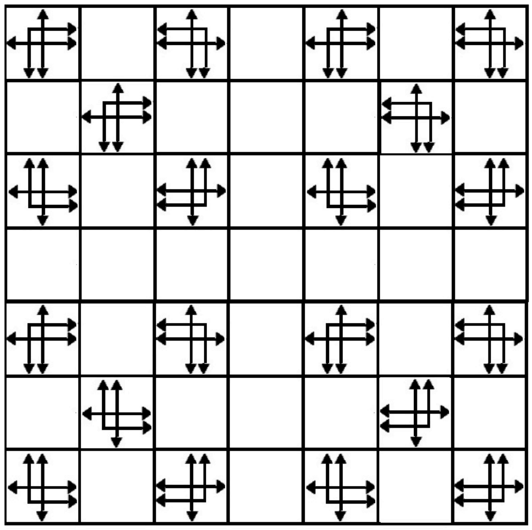

FIGURE 11

In figures 10 and 11, the centres have to be crosses and once we decide upon the orientation of the cross in the centre all the other boxes can be uniquely filled in by arms. We can hence recursively construct a $2^n - 1$ square for every n. Since we can tile arbitrary regions of the plane, therefore we can tile the whole plane. By the forcing itself, depending on the choices of the orientation of the central crosses one can tile the quarter plane, the half plane or the full plane. Thus any tiling will be either two half planes with a column/row of arrows in between, four quarter planes, two quarter planes and a half plane or a whole plane tiling by the choices we make. Now, we are left to prove that any tiling is non-periodic. In figure 11, the only cross in the 4th row has to be the one in the centre. Similarly in our construction the only cross in the middle row of the square of size $2^n - 1$ has to be the one in the centre.Any tiling has to contain a cross. Any cross will eventually give us a square of the above type. Hence any tiling of the plane by these squares is non-periodic.

We are yet to understand how this is related to coloured tiles that we mentioned in the introduction. Let $X = \{x_1, \dots x_{56}\}$ be a set of 56 tiles. Apply the same adjacency rules on it as given for the 56 tiles above. Suppose x_i can be placed left to the tile x_j. Then colour the right edge of x_i and left edge of x_j by the same colour. If two tiles cannot be placed in this way, then the corresponding colours must be different. To be precise let

$$\mathbb{H} = \{(i, j)|x_i \text{ can be placed to the left of } x_j\}$$

Then

$$\mathbb{H} = \bigcup (A_r \times B_r)$$

for some partitions $\{A_r\}, \{B_r\}$ of $\{1, 2 \dots 56\}$ Therefore we can colour the right edges of x_i and left edges of x_j by the colour r for all $(i, j) \in A_r \times B_r$ for all r.

Colour the upper and the lower edges similarly. This colouring gives the required adjacency rules. □

4. Proof of the Berger's Theorem

The way this proof will proceed is the following. We will first find a way of embedding any given turing machine in the tilings given in the previous section. This embedding will give rise to a different set of tiles for every turing machine and the question whether the tiles can tile the whole plane will be the same as the question whether the corresponding turing machine does not halt beginning on an empty tape. Hence if there is a turing machine which decides whether or not a given set of tiles can tile a plane, it also decides the halting problem. Since the halting problem is undecidable, this would prove that the tiling problem is undecidable and hence prove Berger's Theorem.

Theorem. *The Tiling problem is undecidable.*

Proof. We will begin with 56 tiles constructed in the beginning of the earlier section. We will begin by colouring certain double arrows. There are three types of tiles in consideration:
1)Crosses
2)Tiles which are not crosses but have double arrows in both directions and
3)Tiles which which have double arrows in only one direction.
Make another copy of the crosses superimposed on (1,1) and colour the double arrows red.
Make another copy of tiles of type 2. Colour the horizontal double arrows on the first copy and the vertical double arrows on the second copy red.
Make another copy of tiles of type 3. Colour the double arrows red.
This will force a tiling like:

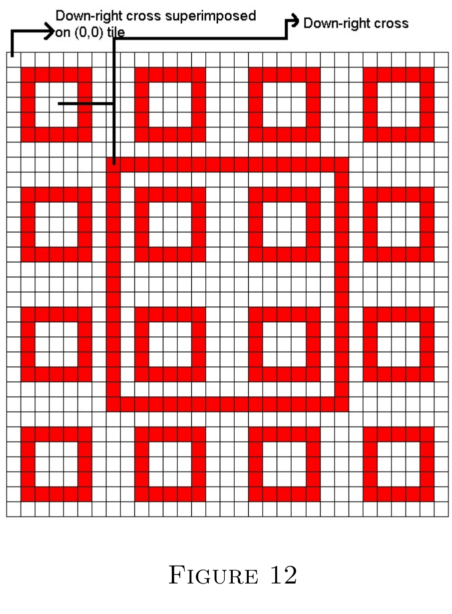

Figure 12

The crosses superimposed on (0,0) are forced to be black. This along with the colouring in the type 2 crosses forces this kind of a pattern to emerge. In our construction described in the earlier section, every alternate sized box will be coloured red. There will be larger and larger 'red-edged squares' and two red-edged squares will either contain one another or not intersect at all. The entire plane can be tiled with these tiles.

The size of a red edged box will be $2^n + 1$ where n is even. A typical red edged square will look like:

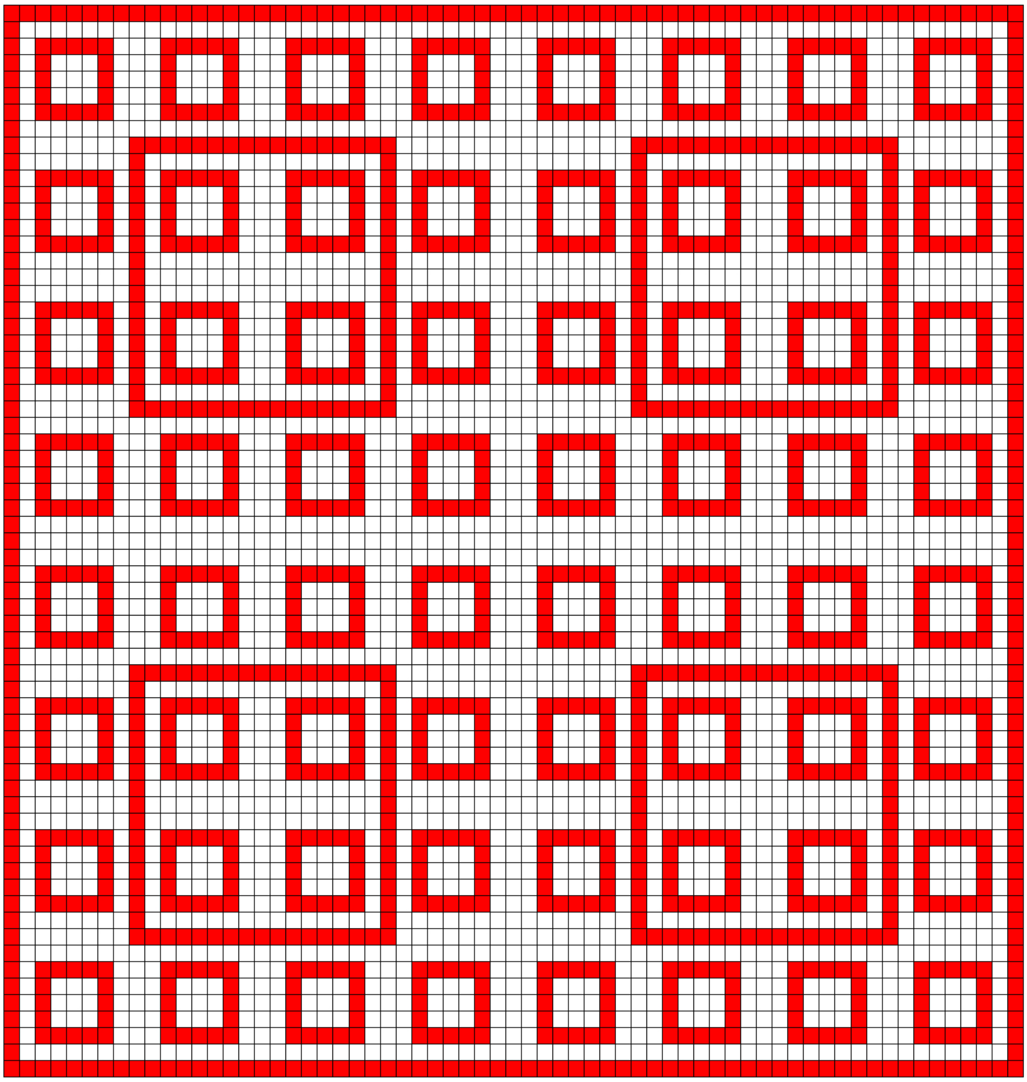

Figure 13

One way of thinking of this, is once inside a red-edged square the tiles forget that there is anything outside and arrange among themselves in the most natural way. Once we are given a big red-edged square, if we forget about the colouring there is a unique way to fill the inside part modulo the middle row and column. We start with a Down-Right Cross on the upper right corner. This forces a 3×3 box. Since we cannot turn back this forces a Down-Right Cross of red colour in the centre of this box. This in-turn forms a 7-box of red colour as in Figure-11 and so on. Ultimately we are left with a choice for the orientation of the central cross which is black. This also implies that the a $2^{2r} + 1$ red-edged square occurs with a period of 2^{2r+1} in a bigger red-edged square

The turing machine will work along unobstructed rows and columns running from the inner edge to the opposite inner edge of a red-edged square. Call these rows and columns, working rows and working columns. The other rows and columns will be called carrier row and carrier columns. For example, the number of working rows in Figure 10 are 9. So as to make sure that the turing machine can work as long as it wants, we need to make sure that the number of such rows and columns keeps on increasing with the size of the box. Let the number of working columns and rows in a $2^{2k} + 1$ red-edged square be F_k.

Lemma 4.1. $F_k = 2^k + 1$ *for all* $k \in \mathbb{N}$.

Proof. $F_1 = 3$ can be checked by mere inspection. Let us assume that this is true for some $k \in \mathbb{N}$. We will prove it for k+1.Let us look at a red-edged square of size $2^{2k+2} + 1$. By arguments above,there will be 4 red-edged squares of size $2^{2k} + 1$ inside such a square. The 2^{2k-1} upper left block of the $2^{2k+2} + 1$ square is the same as the 2^{2k-1} upper left block of the $2^{2k} + 1$ square.This is because we are forced to make the same choices for the orientation of crosses in these portions. By vertical periodicity of squares, we get that the number of working columns in these portions is $\frac{F_k - 1}{2}$. By repeating this on the left and right sides of the upper two squares of size $2^{2k} + 1$ and the middle working column we get that the total number of working columns is $4\frac{(F_k-1)}{2} + 1$. By assumption that $F_k = 2^k + 1$, we can deduce that $F_{k+1} = 2^{k+1} + 1$. Therefore, we are done. $\square$

Till now there is no natural way to decide whether or not the rows and columns are carriers or not. That is, if we were somewhere in a middle of a tiling we would not know whether or not we are in a carrier row or not. We will only be concerned with the portion in a red-edged square which is outside any other red-edged square inside it.

For every tile with a red arrow marking attach a direction symbol on the edges not abutting another red tile. These will be given by D for being on the lower edge, U for being on the upper edge, R for being on the right edge and L for being on the left edge of a red edged square. For example,

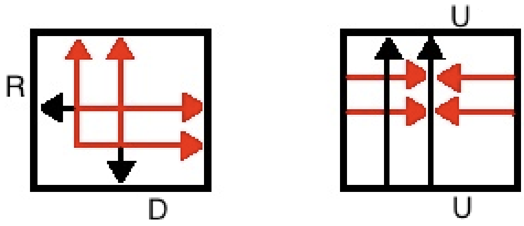

FIGURE 14

Make more copies of the tiles with the black arrows and give them markings D, U, R and L on appropriate positions. Opposite edges must have the same markings, and unlike the tiles with red-arrows it can be marked DU on the upper and lower sides and LR on both left and right sides. The adjacency rules are that adjacent tiles with black arrows must have the same direction symbols on abutting edges. When meeting an outer edge of a tile with red arrows the direction symbol on the red tile should also be present on the black tile . When meeting an inner edge of a red tile then the symbol should exactly match or must have no symbol on them at all.Thus with this convention it is clear that these direction symbols will force a situation like:

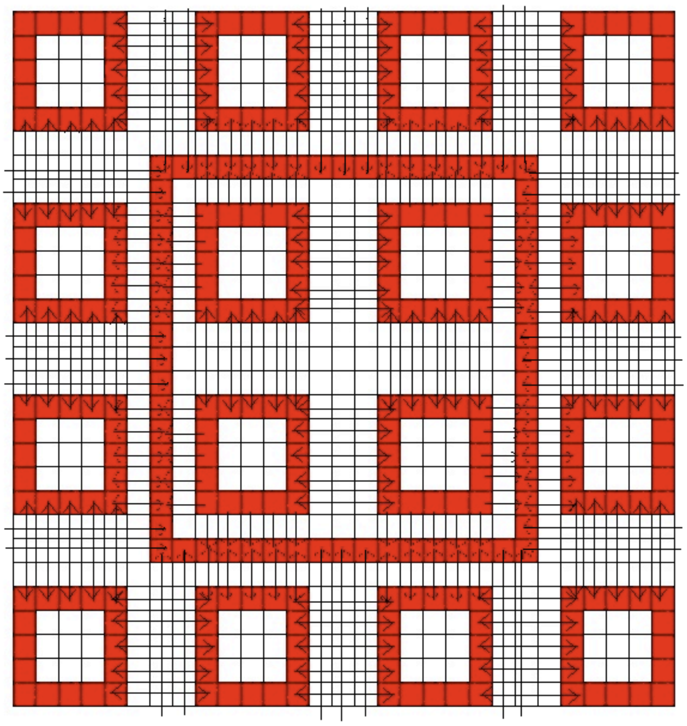

We have thus managed to distinguish between carrier tiles and working tiles. Now we are ready to introduce the turing machine into our system. Take any turing machine as defined in Chapter 3 of [Sipser]. Basically it is a turing machine with a one-sided tape who's head has to move at every step to the one tile to the left or to the right.
Take an arbitrary turing machine $T = (Q, \Sigma, \Gamma, \delta, q_o, q_{accept}, q_{reject})$ where Q, Σ and Γ are all finite sets and

(1) Q is the set of states,
(2) Σ is the input alphabet not containing the special blank symbol $\sqcup$,
(3) $\Gamma$is the tape alphabet, where $\{\sqcup\} \in \Gamma$ and $\Sigma \subset \Gamma$,
(4) $\delta : Q \times \Gamma \longrightarrow Q \times \Gamma \times \{L, R\}$ is the transition function,
(5) $q_o \in Q$ is the start state,
(6) $q_{accept} \in Q$ is the accept state, and
(7) $q_{reject} \in Q$ is the reject state where $q_{reject} \neq q_{accept}$

The form of the halting problem that we shall use is the following.
The language consisting of turing machine encodings which do not halt on an empty input is undecidable.
The turing machine tiles will be the following:

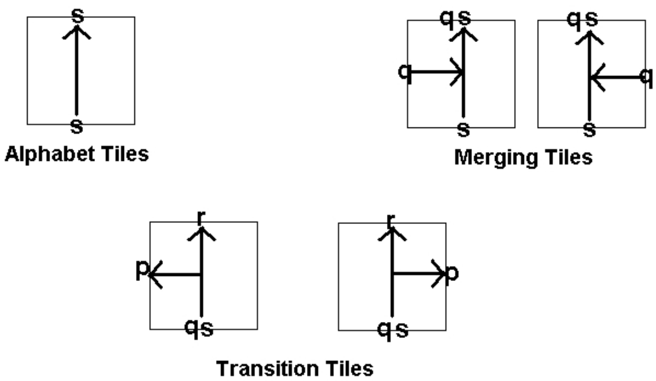

FIGURE 15

where $q \in Q$ is a state, $s \in \Gamma$ is an alphabet and $\delta(q, s) = (p, r, L)$ then we have the first transition tile, and $\delta(q, s) = (p, r, R)$ if we have the second one. The beginning configuration is given by the tiles:

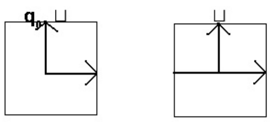

FIGURE 16. Starting Tiles

The adjacency rules among the tiles is that the arrows and the signs on the arrow heads should match. A portion of the tiling corresponding to a turing machine which keeps on writing 0s and 1s alternately while alternating between states q_0 and q_1 is given by:

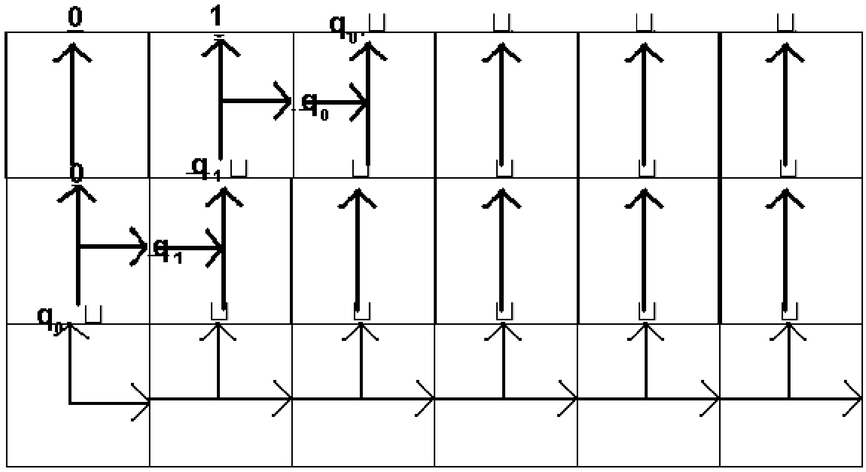

Figure 17

Now superimpose these tiles with the tiles we already had. The alphabet tiles, the merging tiles and the action tiles need to be superimposed on the tiles with black arrows without any direction symbols that carrier tiles have. The lower row 2nd left most red tile in a red-edged square receives a signal from the up-right red cross and will be superimposed the first starting tiles. All other lower edge tiles on a red-edged square except the left-up cross and those adjacent to tiles with D markings are superimposed with the 2nd starting tile. The carrier tiles carry the turing machine signals unchanged horizontally and vertically depending on the direction symbol that they have. Tiles with coloured edges as required can be obtained as in the last section.

Now for every turing machine we have a set of finite tiles. These tiles can tile the plane if and only if the corresponding turing machine does not halt on the blank tape.

Now suppose we are given a finite set of tiles. It is easy to decide whether or not the tiles are an encoding of a turing machine. Therefore if one can decide whether or not the tiles can tile the whole plane, one will be able to conclude whether or not a turing machine halts starting on a blank tape. Since the latter is undecidable we have that the former is too. Hence Berger's Theorem is proved. □

A careful study of the proof gives us that the completion problem is also undecidable. That is given a beginning tiling whether the whole plane can be tiled or not.

5. Acknowledgments

Many thanks to Prof. Brian Marcus for teaching me Symbolic Dynamics among a lot of other things, encouraging me and helping me grow, Tom Meyerovitch for introducing Robinson's paper to me and clarifying many doubts that I had along the way and Prof. Joel Friedman for giving me the right opportunity and incentive to understand the material better.

References

[Wang] Hao Wang, *Proving Theorems by Pattern Recognition-II*, Bell System Technical Journal,1961

[Berger] Robert Berger, *The Undecidability of the Domino Tiling*, Memoirs of American Mathematical Society, 1966

[Robinson] Raphael M. Robinson, *Undecidability and Nonperiodicity of Tilings of the Plane*, Inventiones mathematicae, 1971

[Kari] Jarkko Kari, *The Tiling Problem Revisited*, Machines, Computations, and Universality, 2007

[Radin] C. Radin, *Rashomon, Pavages et rotations* Gazette des mathmaticiens: bulletin de liaison de la . . . ,1999

[Sipser] Michael Sipser, *Introduction to the Theory of Computation* Second Edition,2005