On the Undecidability of the Tiling Problem

Jarkko Kari*

Department of Mathematics, FIN-20014 University of Turku, Finland
jkari@utu.fi

Abstract. The tiling problem is the decision problem to determine if a given finite collection of Wang tiles admits a valid tiling of the plane. In this work we give a new proof of this fact based on tiling simulations of certain piecewise affine transformations. Similar proof is also shown to work in the hyperbolic plane, thus answering an open problem posed by R.M.Robinson 1971 [9].

1 Introduction

A *Wang tile* is a unit square tile with colored edges. Tiles are placed on the plane edge-to-edge, under the matching constraint that abutting edges must have the same color. Tiles are used in the given orientation, without rotating. If T is a finite set of Wang tiles, a tiling of the plane is a covering $t : \mathbb{Z}^2 \longrightarrow T$ of the plane by copies of the tiles in such a way that the color constraint is satisfied everywhere.

The *tiling problem* (also known as the domino problem) is the decision problem that asks whether a given finite tile set T admits at least one valid tiling $t : \mathbb{Z}^2 \longrightarrow T$. This problem was proved undecidable by R.Berger in 1966 [1], see also R.M.Robinson [9] for another proof. Both proofs rely on an explicit construction of an *aperiodic* tile set. Set T is called aperiodic if it admits some valid tiling of the plane, but it does not admit a valid periodic tiling, i.e. a tiling that is invariant under some translation. Note that existence of such aperiodic sets is not obvious, and in fact it was conjectured prior to Berger's work that they do not exist. If aperiodic sets did not exist, then the tiling problem would be decidable as one can simply try tilings of larger and larger rectangles until either (1) a rectangle is found that can no longer be tiled, or (2) a tiling of a rectangle is found that can be repeated periodically. Only aperiodic tile sets fail to reach either (1) or (2).

Note that Wang tiles are an abstraction of geometric tiles. Indeed, by using suitable "bumps" and "dents" on the sides to represent different colors, one can effectively replace any set of Wang tiles by a set of geometric tiles (all polygons with rational coordinates) such that the geometric tiles admit a tiling (a non-overlapping covering of the plane) if and only if the Wang tiles admit a tiling. Hence undecidability of the tiling problem by geometric tiles follows from Berger's result.

* Research supported by the Academy of Finland grant 211967.

V. Geffert et al. (Eds.): SOFSEM 2008, LNCS 4910, pp. 74–82, 2008.

In this work we present a new proof for the undecidability of the tiling problem. The proof uses plane tilings to simulate dynamical systems that are based on piecewise affine transformations. The undecidability of the tiling problem will then be concluded from the undecidability of the mortality problem of such dynamical systems.

A particularly nice feature of our proof is the fact that it is purely combinatorial. As a result of this the method generalizes easily to tilings in other lattices as well. In particular, we show that the tiling problem is undecidable in the hyperbolic plane. This resolves an open question asked already by Robinson in 1971 [9], and discussed by him in more details in 1978 [10]. In particular, Robinson proved the undecidability of the origin constrained tiling problem in the hyperbolic plane. This is the easier question where one asks the existence of a valid tiling that contains a copy of a fixed seed tile. We mention that there is a concurrent, independent and unpublished approach by M.Margenstern to the tiling problem in the hyperbolic plane [8]. We have reported our approach previously in [5].

We first discuss piecewise affine transformations and their mortality problem. We then outline the construction of corresponding Wang tiles. We then conclude by providing the analogous construction in the hyperbolic plane.

2 Mortality Problems of Turing Machines and Piecewise Affine Maps

Our proof is based on a reduction from the *mortality problem* of Turing machines. In this question we are given a deterministic Turing machine with a halting state, and the problem is to determine if there exists a non-halting configuration, that is, a configuration of the Turing machine that never evolves into the halting state. Such configuration is called immortal. Note that the Turing machine operates on an infinite tape, and configurations may contain infinitely many non-blank symbols.

Mortality problem of Turing machines. Does a given Turing machine have an immortal configuration ?

The mortality problem was proved undecidable by P.K.Hooper in 1966 [4], the same year that Berger proved his result. The two results have similar flavor, but proofs are independent in the sense that they do not rely on each other in either direction. Note an analogy to aperiodic tile sets: Hooper's result means that there must exist aperiodic Turing machines, that is, Turing machines that have immortal configurations but no immortal configuration repeats itself periodically. Our present proof establishes another connection between Hooper's and Berger's results since we reduce the mortality problem to the tiling problem.

We first consider dynamical systems determined by piecewise affine transformations of the plane. There exists a well known technique to simulate Turing machines by such transformations, see e.g. [2,7]. The idea is to encode Turing machine configurations as two real numbers $(l, r) \in \mathbb{R}^2$, representing the left and the right halves of the infinite tape, respectively. The integer parts of l and r uniquely determine the next rule of the Turing machine to be used. More precisely, suppose the tape and state alphabets of the Turing machine are $A = \{0, 1, \ldots, a\}$ and $Q = \{0, 1, \ldots, b\}$. Let M be an even integer such that $M > a + 1$ and $M > b$. Then we let

$$l = \sum_{i=-1}^{-\infty} M^i t_i, \text{ and}$$

$$r = Mq + \sum_{i=0}^{\infty} M^{-i} t_i,$$

where $q \in Q$ is the current state of the Turing machine and $(t_i)_{i \in \mathbb{Z}}$ is the content of the infinite tape. We use a moving tape model of Turing machines: The Turing machine always reads cell 0 while the tape shifts left or right according to the rules of the Turing machine. Note that $\lfloor r \rfloor = Mq + t_0$ determines the next move of the machine, and that the encoding is one-to-one.

For each transition rule of the Turing machine one can effectively associate a rational affine transformation of $\mathbb{R}^2$ that simulates that transition. The matrix of the transformation is

$$\begin{pmatrix} M & 0 \\ 0 & \frac{1}{M} \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \text{ or } \begin{pmatrix} \frac{1}{M} & 0 \\ 0 & M \end{pmatrix}$$

depending on the direction of the movement associated with the transition. The translation part of the transformation takes care of the changes in the tape symbol in cell 0 as well as the change in the state of the Turing machine.

In this fashion any deterministic Turing machine is converted into a system of finitely many rational affine transformations $f_1, f_2, \ldots, f_n$ of $\mathbb{R}^2$ and corresponding disjoint unit squares $U_1, U_2, \ldots, U_n$ with integer corners. Squares U_i serve as domains for the affine maps: the affine transformation f_i is applied when (l, r) is in the unit square U_i. Together the transformations define a partial function $f : \mathbb{R}^2 \longrightarrow \mathbb{R}^2$ whose domain is $U = U_1 \cup U_2 \cup \ldots \cup U_n$, and whose operation is

$$\vec{x} \mapsto f_i(\vec{x}) \text{ for } \vec{x} \in U_i.$$

Point $\vec{x} \in \mathbb{R}^2$ is called immortal if for every $i = 0, 1, 2, \ldots$ the value $f^i(\vec{x})$ is in the domain U. In other words, we can continuously apply the given affine transformations and the point we obtain always belongs to one of the given unit squares U_i.

The reduction from Turing machines to piecewise affine transformations preserves immortality: the Turing machine has an immortal configuration if and only if the corresponding system of affine maps has an immortal starting point. Hence we conclude from Hooper's result that the following immortality question is undecidable:

Mortality problem of piecewise affine maps: Does a given system of rational affine transformations $f_1, f_2, \ldots, f_n$ of the plane and disjoint unit squares $U_1, U_2, \ldots, U_n$ with integer corners have an immortal starting point ?

3 Reduction into the Euclidean Tiling Problem

Next the mortality question of piecewise affine maps is reduced into the tiling problem of Wang tiles. The idea is very similar to a construction of an aperiodic Wang tile set presented in [6]. In [6] a tile set was given such that every valid tiling is forced to simulate an infinite orbit according to the one-dimensional piecewise linear function $f : [\frac{1}{2}, 2] \longrightarrow [\frac{1}{2}, 2]$ where

$$f(x) = \begin{cases} 2x, & \text{if } x \leq 1, \text{ and} \\ \frac{2}{3}x, & \text{if } x > 1. \end{cases}$$

Function f has no periodic orbits so the corresponding tile set is aperiodic.

The construction needs to be generalized in two ways: (1) instead of linear maps we need to allow more general affine maps, and (2) instead of $\mathbb{R}$ the maps are now over $\mathbb{R}^2$. Fortunately both generalizations are very natural and work without any complications.

The colors of our Wang tiles are elements of $\mathbb{R}^2$. Let $f : \mathbb{R}^2 \longrightarrow \mathbb{R}^2$ be an affine function. We say that tile

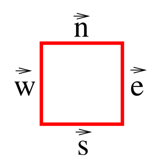

computes function f if

$$f(\vec{n}) + \vec{w} = \vec{s} + \vec{e}.$$

(The "input" $\vec{n}$ comes from north, and $f(\vec{n})$ is computed. A "carry in" $\vec{w}$ from the west is added, and the result is split between the "output" $\vec{s}$ to the south and the "carry out" $\vec{e}$ to the east.)

Suppose we have a correctly tiled horizontal segment of length n where all tiles compute the same f.

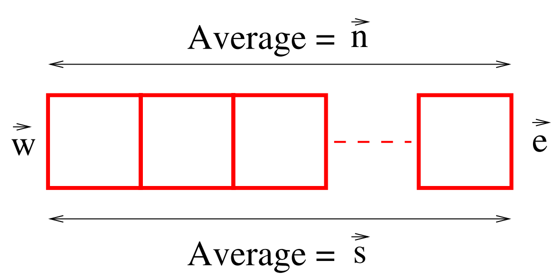

It easily follows that

$$f(\vec{n}) + \frac{1}{n}\vec{w} = \vec{s} + \frac{1}{n}\vec{e},$$

where $\vec{n}$ and $\vec{s}$ are the averages of the top and the bottom labels. As the segment is made longer, the effect of the carry in and out labels $\vec{w}$ and $\vec{e}$ vanish. Loosely speaking then, in the limit if we have an infinite row of tiles, the average of the input labels is mapped by f to the average of the output labels.

Consider now a given system of affine maps f_i and unit squares U_i. For each i we construct a set T_i of Wang tiles that compute function f_i and whose top edge labels $\vec{n}$ are in U_i. An additional label i on the vertical edges makes sure that tiles of different sets T_i and T_j cannot be mixed on any horizontal row of tiles. Let

$$T = \bigcup_i T_i.$$

If T admits a valid tiling then the system of affine maps has an immortal point. Namely, consider any horizontal row in a valid tiling. The top labels belong to a compact and convex set U_i. Hence there is $\vec{x} \in U_i$ that is the limit of the top label averages over a sequence of segments of increasing length. Then $f_i(\vec{x})$ is the limit of the bottom label averages over the same sequence of segments. But the bottom labels of a row are the same as the top labels of the next row below, so $f_i(\vec{x})$ is the limit of top label averages of the next row. The reasoning is repeated for the next row, and for all rows below. We see that $\vec{x}$ starts an infinite orbit of the affine maps, so it is an immortal point.

We still have to detail how to choose the tiles so that any immortal orbit of the affine maps corresponds to a valid tiling. Consider a unit square

$$U = [n, n+1] \times [m, m+1]$$

where $n, m \in \mathbb{Z}$. Elements of

$$\mathrm{Cor}(U) = \{(n,m), (n,m+1), (n+1,m), (n+1,m+1)\}$$

are the *corners* of U. For any $\vec{x} \in \mathbb{R}^2$ and $k \in \mathbb{Z}$ denote

$$A_k(\vec{x}) = \lfloor k\vec{x} \rfloor$$

where the floor is taken for each coordinate separately:

$$\lfloor (x,y) \rfloor = (\lfloor x \rfloor, \lfloor y \rfloor).$$

Denote

$$B_k(\vec{x}) = A_k(\vec{x}) - A_{k-1}(\vec{x}) = \lfloor k\vec{x} \rfloor - \lfloor (k-1)\vec{x} \rfloor.$$

It easily follows that if $\vec{x} \in U$ then

$$B_k(\vec{x}) \in \mathrm{Cor}(U).$$

Vector $\vec{x}$ will be represented as the two-way infinite sequence

$$\ldots B_{-2}(\vec{x}), B_{-1}(\vec{x}), B_0(\vec{x}), B_1(\vec{x}), B_2(\vec{x}), \ldots$$

of corners. It is the *balanced representation* of $\vec{x}$, or the sturmian representation of $\vec{x}$. Note that both coordinate sequences are sturmian.

The tile set corresponding to a rational affine map

$$f_i(\vec{x}) = M\vec{x} + \vec{b}$$

and its domain square U_i consists of all tiles

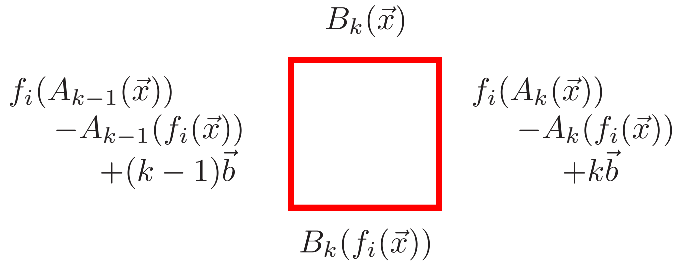

where $k \in \mathbb{Z}$ and $\vec{x} \in U_i$. Observe the following facts:

(1) For fixed $\vec{x} \in U_i$ the tiles for consecutive $k \in \mathbb{Z}$ match in the vertical edges so that a horizontal row can be formed whose top and bottom labels read the balanced representations of $\vec{x}$ and $f_i(\vec{x})$, respectively.
(2) A direct calculation shows that the tile above computes function f_i, that is,

$$f_i(\vec{n}) + \vec{w} = \vec{s} + \vec{e}.$$

(3) Because f_i is rational, there are only finitely many tiles constructed, even though there are infinitely many $k \in \mathbb{Z}$ and $\vec{x} \in U_i$. Moreover, the tiles can be effectively constructed.

Now it is clear that if the given system of affine maps has an immortal point $\vec{x}$ then a valid tiling exists where the labels of consecutive horizontal rows read the balanced representations of the consecutive points of the orbit for $\vec{x}$. We conclude that the tile set we constructed admits a tiling of the plane if and only if the given system of affine maps is immortal. Undecidability of the tiling problem follows from the undecidability of the immortality problem that we established in Section 2

4 Reduction into the Tiling Problem on the Hyperbolic Plane

The method of the previous section works just as well in the hyperbolic plane. Instead of Wang tiles we use hyperbolic pentagons that in the half-plane model of hyperbolic geometry are copies of

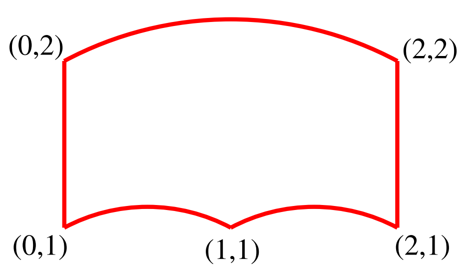

Note that all five edges are straight line segments. These tiles admit valid tilings of the hyperbolic plane in uncountably many different ways

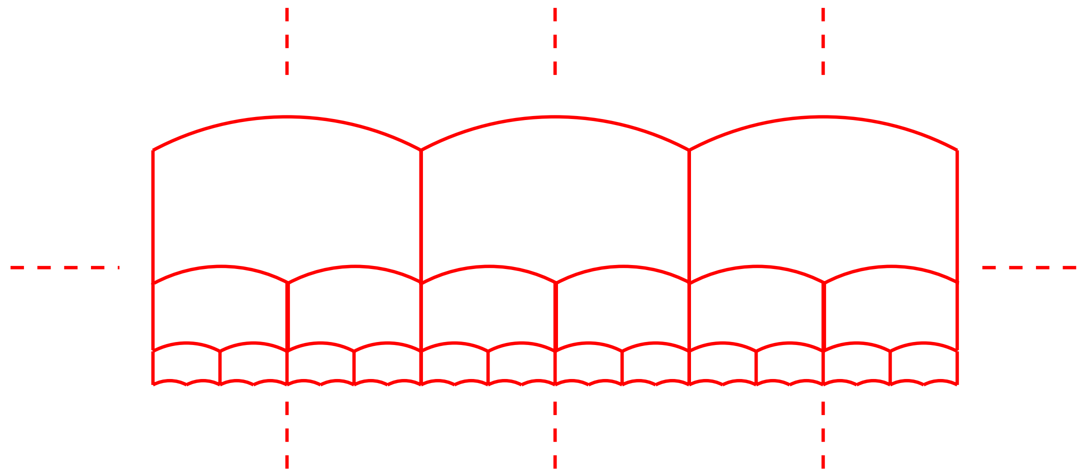

In these tilings the tiles form infinite "horizontal rows" in such a way that each tile has two adjacent tiles in the next row "below".

In the following these hyperbolic pentagons are used instead of the Euclidean square shaped Wang tiles. The five edges will be colored, and in a valid tiling abutting edges of adjacent tiles must match. This is an abstraction – analogous to Wang tiles in the Euclidean plane – that can be transformed into hyperbolic geometric shapes using bumps and dents.

Exactly as in the Euclidean case we color the edges by vectors $\vec{x} \in \mathbb{R}^2$. We say that pentagon

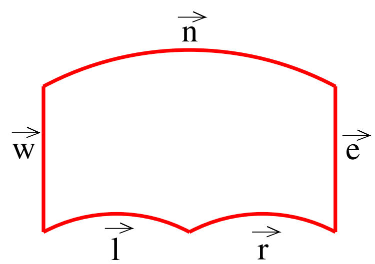

computes the affine transformation $f : \mathbb{R}^2 \longrightarrow \mathbb{R}^2$ if

$$f(\vec{n}) + \vec{w} = \frac{\vec{l} + \vec{r}}{2} + \vec{e}.$$

Note that the difference to Euclidean Wang tiles is the fact that the "output" is now divided between $\vec{l}$ and $\vec{r}$.

Consider a correctly tiled horizontal segment of length n where all tiles compute the same f.

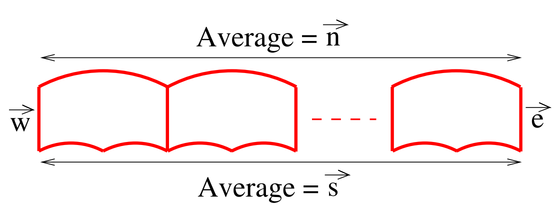

Clearly we have

$$f(\vec{n}) + \frac{1}{n}\vec{w} = \vec{s} + \frac{1}{n}\vec{e},$$

where $\vec{n}$ and $\vec{s}$ are the averages of the top and the bottom labels on the segment.

Analogously to the Euclidean case, given a system of affine maps f_i and unit squares U_i, we construct for each i a set T_i of pentagons that compute function f_i and whose top edge labels $\vec{n}$ are in U_i. It follows, exactly as in the Euclidean case, that if a valid tiling of the hyperbolic plane with such pentagons exists then from the labels of horizontal rows one obtains an infinite orbit in the system of affine maps.

We still have to detail how to choose the tiles so that the converse is also true: if an immortal point exists then its orbit provides a valid tiling. The tile set corresponding to a rational affine map

$$f_i(\vec{x}) = M\vec{x} + \vec{b}$$

and its domain square U_i consists of all tiles

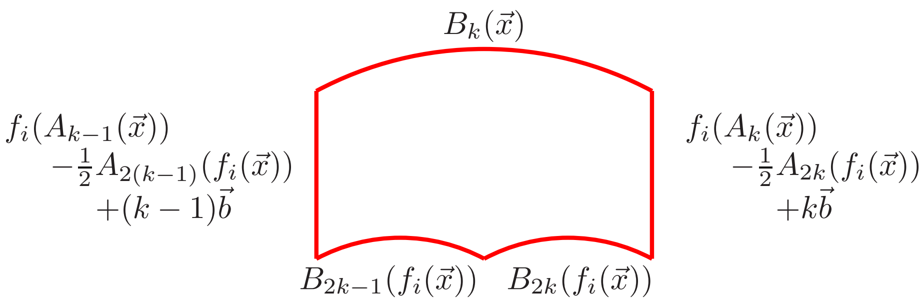

where $k \in \mathbb{Z}$ and $\vec{x} \in U_i$. Now we can reason exactly as in the Euclidean case:

(1) For fixed $\vec{x} \in U_i$ the tiles for consecutive $k \in \mathbb{Z}$ match so that a horizontal row can be formed whose top and bottom labels read the balanced representations of $\vec{x}$ and $f_i(\vec{x})$, respectively.

(2) A direct calculation shows that the tile computes function f_i:

$$f_i(\vec{n}) + \vec{w} = \frac{\vec{l} + \vec{r}}{2} + \vec{e}.$$

(3) There are only finitely many pentagons constructed (because f_i is rational), and they can be formed effectively.

The tiles constructed admit a valid tiling of the hyperbolic plane if and only if the corresponding system of affine maps has an immortal point. So we have proved

Theorem. The tiling problem is undecidable in the hyperbolic plane.

References

1. Berger, R.: Undecidability of the Domino Problem. Memoirs of the American Mathematical Society 66, 72 (1966)
2. Blondel, V., Bournez, O., Koiran, P., Papadimitriou, C., Tsitsiklis, J.: Deciding stability and mortality of piecewise affine dynamical systems. Theoretical Computer Science 255, 687–696 (2001)

3. Goodman-Strauss, C.: A strongly aperiodic set of tiles in the hyperbolic plane. Inventiones Mathematicae 159, 119–132 (2005)
4. Hooper, P.K.: The undecidability of the Turing machine immortality problem. The Journal of Symbolic Logic 31, 219–234 (1966)
5. Kari, J.: The Tiling Problem Revisited (extended abstract). In: Durand-Lose, J., Margenstern, M. (eds.) MCU 2007. LNCS, vol. 4664, pp. 72–79. Springer, Heidelberg (2007)
6. Kari, J.: A small aperiodic set of Wang tiles. Discrete Mathematics 160, 259–264 (1996)
7. Koiran, P., Cosnard, M., Garzon, M.: Computability with low-dimensional dynamical systems. Theoretical Computer Science 132, 113–128 (1994)
8. Margenstern, M.: About the domino problem in the hyperbolic plane, a new solution. Manuscript , 109 (2007), `http://www.lita.univ-metz.fr/~margens/` and also see arXiv:cs/0701096, same title
9. Robinson, R.M.: Undecidability and nonperiodicity for tilings of the plane. Inventiones Mathematicae 12, 177–209 (1971)
10. Robinson, R.M.: Undecidable tiling problems in the hyperbolic plane. Inventiones Mathematicae 44, 259–264 (1978)