The completion problem

Let P be a fixed finite set (of prototiles).

A **finite pattern** is an assignment $p : D \longrightarrow P$ of tiles into cells of a finite domain $D \subseteq \mathbb{Z}^2$.

We say that p is a **sub-pattern** of a configuration $c : \mathbb{Z}^2 \longrightarrow P$ if

$$c|_D = p.$$

The completion problem of tile set P
Instance: A finite domain $D :\subseteq \mathbb{Z}^2$ and a finite pattern $p : D \longrightarrow P$
Positive instance: D and p such that there exists a valid tiling $c : \mathbb{Z}^2 \longrightarrow P$ that has p as a subpattern.

The completion problem is different for different tile sets. The tile set is fixed, and the input is a finite pattern. For some tile sets the completion problem is decidable. But we can prove that there are also tile sets whose completion problem is undecidable.

The situation is analogous to the decision problem:

Halting problem of TM $M = (S, \Gamma, \delta, s_0, s_h, b)$
Instance: A finite tape content $f \in \Gamma^{\mathbb{Z}}$
Positive instance: f such that $(s_0, 0, f) \vdash \cdots \vdash (s_h, i, g)$ for some and $i \in \mathbb{Z}$ and $g \in \Gamma^{\mathbb{Z}}$

This is also decidable for some Turing machines M, but:

Theorem [Turing 1936]. There exists a TM M such that **Halting problem of TM** M is undecidable.

We call a Turing machine M **universal** if **Halting problem of** M is undecidable.

Example. The following universal Turing machine (due to Y.Rogozhin) with $4+1$ states and 6 tape symbols has an undecidable halting problem:

$$M = (\{q_1, q_2, q_3, q_4, h\}, \{1, b, >, <, 0, c\}, \delta, q_1, h, b),$$

and δ is given by the following table

	q_1	q_2	q_3	q_4
1	$(q_1, <, L)$	$(q_2, 0, R)$	$(q_3, 1, R)$	$(q_4, 0, R)$
b	$(q_1, >, R)$	$(q_3, >, L)$	$(q_4, <, R)$	(q_2, c, L)
$>$	(q_1, b, L)	$(q_2, <, R)$	(q_3, b, R)	$(q_4, <, R)$
$<$	$(q_1, 0, R)$	$(q_2, >, L)$	h	h
0	$(q_1, <, L)$	$(q_2, , 1, L)$	(q_1, c, R)	(q_2, c, L)
c	$(q_4, 0, R)$	(q_2, b, R)	$(q_1, 1, R)$	(q_4, b, R)

where the item on column q, row x is $\delta(q, x)$.

Theorem. There exists a finite set P of Wang prototiles such that **The completion problem of tile set P** is undecidable.

Proof. In the homework assignments.

Beyond aperiodicity: arecursive tile sets

Robinson's tile set is aperiodic: no periodic tilings exist. Still, valid tilings that are **"simple"** exist. Using the special $2^n - 1$ -squares one can algorithmically construct bigger and bigger portions of a fixed valid tiling.

But there exist tile sets that only admit very complicated tilings: only tilings that cannot be algorithmically constructed.

We call a configuration $c : \mathbb{Z}^2 \longrightarrow T$ **recursive** if there exists an algorithm that outputs $c(i, j)$, when given arbitrary integers i, j as input. If no such algorithm exists then c is called **non-recursive**.

Analogously, a tape content

$$f : \mathbb{Z} \longrightarrow \Gamma$$

of a Turing machine is **recursive** if there exists an algorithm that returns $f(i)$ for any given input $i \in \mathbb{Z}$. Otherwise f is **non-recursive**.

Any **strongly periodic** configuration c is recursive: The pattern of a fixed $n \times n$ square period can be hard coded inside the algorithm as a lookup table. The symbol in any position i, j is then

$$c(i,j) = c(i \pmod n, j \pmod n).$$

Analogously every periodic tape content $f : \mathbb{Z} \longrightarrow \Gamma$ is recursive.

But there are of course many non-periodic recursive configurations also (e.g. some valid tilings by the Robinson tile set).

Analogously to aperiodicity, we define the concept of arecursivity as follows: Wang tile set P is **arecursive** if and only if

(i) it admits valid tilings, and

(ii) it does not admit any recursive valid tiling.

Clearly any arecursive tile set is aperiodic, but the converse is not true since Robinson's tile set is not arecursive. Arecursivity is a stronger property than aperiodicity.

A simpler concept is a tile set with a **seed tile** such that any valid tiling that contains a seed tile is non-recursive.

Theorem. There exists a finite set P of Wang prototiles and a seed tile $t \in P$ such that every valid tiling that contains t is non-recursive, and there are valid non-recursive tilings that contain t.

The proof is based on an analogous concept on Turing machines:

Lemma. There exists a Turing machine M that halts when started on any recursive initial tape, but for some non-recursive initial tape M does not halt.

Proof of the lemma is sketched in the notes. (We skip the proof here as it is necessarily incomplete unless we dig deep into the topics of the "Automata and Formal Languages" -course.)

Using the arecursive TM M of the lemma we construct a tile set as follows:

We take the machine tiles P_M as before, except that the start tiles

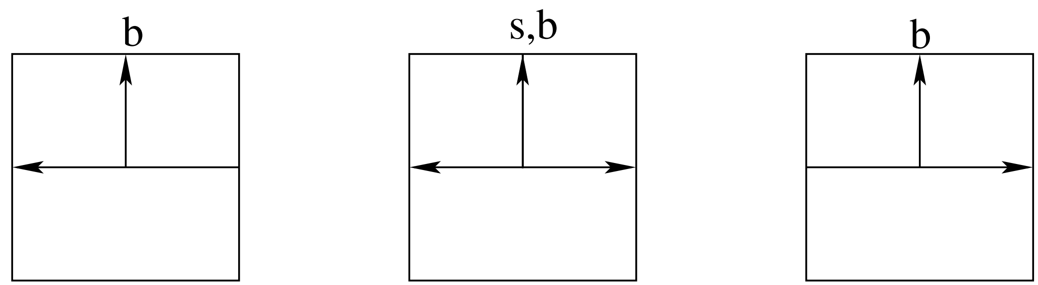

that represent the initial empty tape are replaced by the tiles

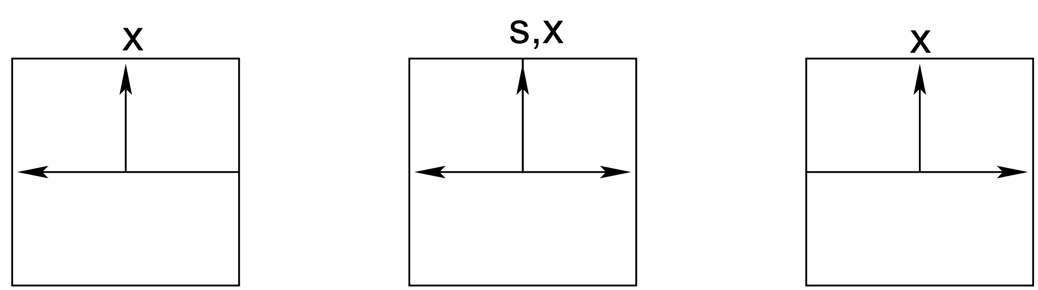

for all tape letters $x \in \Gamma$ and the initial state $s \in S$.

The new start tiles allow non-blank initial tapes. As the seed tile t we select the start tile

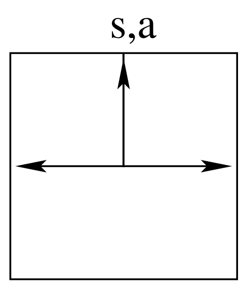

where letter a is chosen so that there is an initial tape content $f : \mathbb{Z} \longrightarrow \Gamma$ with $f(0) = a$ such that M does not halt when started on tape f.

The new start tiles allow the horizontal row with t to contain any initial tape content f with $f(0) = a$. The machine tiles then force the rows above to simulate machine M. If M halts then the tiling becomes impossible.

We have: The tiles admit a tiling containing a horizontal row r of initialization tiles $\iff$ the initial configuration of M represented by row r does not lead to halting of M.

The tiles admit a valid tiling containing t: We choose an initial row that represents a tape content f from which M does not halt.

The tiles do not admit a recursive tiling containing t: If the tiling is recursive then the horizontal row containing t is recursive, so we would have a recursive initial tape from which M does not halt, a contradiction with the property of M

It is possible to modify Robinson's tile set so that the seed tile restriction can be removed. Robinson's tiles form nested boards, and a simulation of the tile set of the previous theorem is done on all boards, with the seed tile at the center of each board.

Theorem. There exist arecursive sets of Wang tiles.

Proof. Skipped.

A main new problem in the proof: the tilings on all boards should be central patterns of the **same** valid tiling around a seed tile. Otherwise, if different boards are allowed to contain pieces of different tilings then a recursive tiling can be easily built.

Solution: new signals that carry the information about the initial tape content of the simulated TM between boards of different sizes, so that different boards are forced to be consistent with each other.