Finite systems of forbidden patterns

(aka **subshifts of finite type**)

To make constructing tile sets easier we may replace the edge coloring based matching condition of Wang tiles by specifying a finite collection of **forbidden patterns**. A configuration is then a valid tiling if and only if it does not contain a forbidden pattern.

Take

- a finite set A of symbols,
- a finite $N \subseteq \mathbb{Z}^2$, a **neighborhood**,
- a set $R \subseteq A^N$ of **allowed patterns**

The complement $F = A^N \setminus R$ is the corresponding set of **forbidden patterns**.

A configuration $c \in A^{\mathbb{Z}^2}$ is in the **SFT (subshift of finite type)** defined by allowed pattern $R \subseteq A^N$ (or forbidden patterns $F \subseteq A^N$) if for every $\vec{a} \in \mathbb{Z}^2$

$$\tau_{\vec{a}}(c)|_N \in R,$$

where $\tau_{\vec{a}}$ is the translation of the configuration by vector $\vec{a}$:

$$\forall \vec{b} \in \mathbb{Z}^2 \quad \tau_{\vec{a}}(c)(\vec{b}) = c(\vec{b} - \vec{a}).$$

In other words: c does not contain any of the forbidden patterns in any position $\vec{a}$.

Tilings by any Wang tile set can be expressed as an SFT with the neighborhood

$$N = \{(0,0),(0,-1),(0,1),(-1,0),(1,0)\},$$

and the set R that contains all those cross shaped patterns where the colors of the neighboring tiles match.

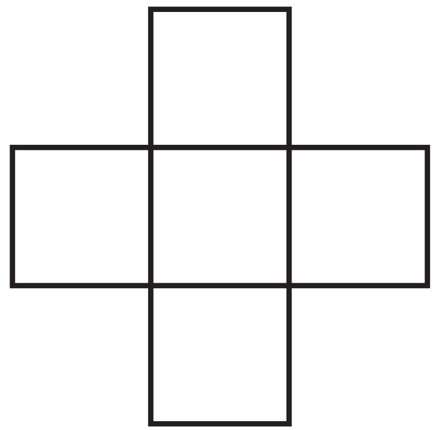

There is also a correspondence to the **other direction**: for any given SFT we can effectively (=algorithmically) construct a Wang tile set T such that there is a natural correspondence between configurations in the SFT and valid tilings by T.

Let $R \subseteq A^N$ be the given **allowed** patterns. We construct Wang tiles as follows:

(i) First note that if $N \subseteq N'$ for another neighborhood N' then we can construct allowed patterns $R' \subseteq A^{N'}$ so that R and R' define the same SFT.

Indeed: we take in R' all patterns $p \in A^{N'}$ such that $p|_N \in A$.

Because N is a subset of some square N', we thus convert R into an equivalent set R' of square shapes patterns.

Let $R \subseteq A^N$ be the given **allowed** patterns where N is an $m \times m$ **square**.

(ii) Construct a set P of Wang tiles as follows: The tiles are the allowed $m \times m$ square patterns over A, that is, $P = R$.

Colors of $t \in R$ are obtained by erasing one boundary column or row from it:

For example, the following figure illustrates the tile corresponding to a 3×3 allowed pattern:

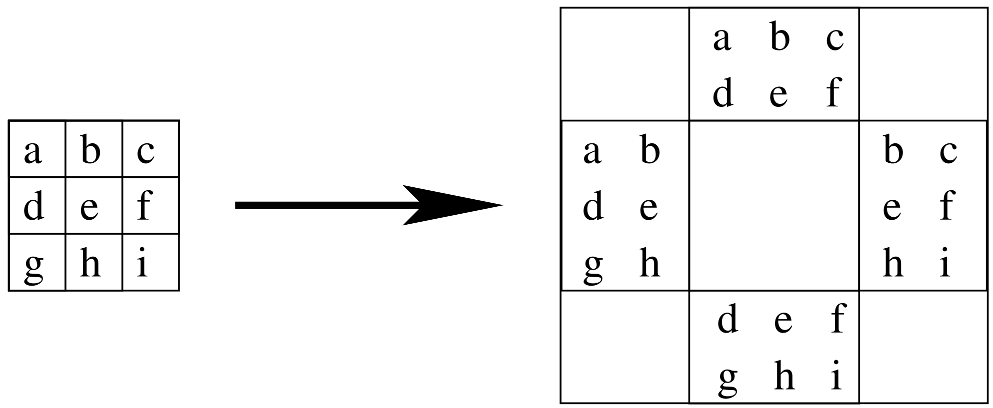

Allowed 3x3 pattern　　　　Wang tile

Two adjacent Wang tiles match if and only if the $m \times m$ patterns they represent have the correct $(m-1) \times m$ or $m \times (m-1)$ overlap when the tiles are placed next to each other.

Let $c \in A^{\mathbb{Z}^2}$ be any configuration. Associate to it the element

$$\Psi(c) \in (A^{m \times m})^{Z^2}$$

where for each $(i, j) \in \mathbb{Z}^2$ we have in position (i, j) the $m \times m$ block that appears in c in position (i, j) (lower left corner at position (i, j).) This is known as a **higher block presentation** of c.

- The map $c \mapsto \Psi(c)$ is one-to-one:

- If c only contains allowed patterns then $\Psi(c) \in P^{\mathbb{Z}^2}$ and neighboring tiles match in color. So $\Psi(c)$ is a valid tiling by Wang prototiles P.

- Conversely, any valid tiling t is $\Psi(c)$ for a suitable c: Such c is obtained by pasting, for each $(i, j) \in \mathbb{Z}^2$, tile $t(i, j)$ as the $m \times m$ block in c with lower left corner in position (i, j). Clearly then c only contains allowed patterns since all pasted patterns are allowed patterns.

Remark: c is well-defined as the pasted $m \times m$ blocks are compatible; if two pasted blocks overlap they assign the same symbols in the overlapping region.

In summary, we have the following directions:

(i) For every Wang protoset P one can effectively construct an SFT over P such that $c \in P^{\mathbb{Z}^2}$ is in the SFT if and only if c is a valid Wang tiling.

(ii) Conversely, for every SFT one can effectively construct a Wang protoset P such that P admits a (periodic) tiling if and only if the SFT contains a (periodic) tiling.

In the following decision problems we can describe tiles in any terms that locally determine which tiles are allowed to be next to each other. Such tiles can anyway be effectively converted into an equivalent set of Wang tiles. This substantially simplifies the constructions.

The periodic tiling problem

Recall the following decision problem.

Periodic tiling problem

Instance: A finite set $\mathcal{T}$ of Wang tiles

Positive instance: $\mathcal{T}$ admits a periodic tiling of the plane

We have seen the that problem is semi-decidable. Next we show that the problem is undecidable (which then implies that the negative instances are not semi-decidable.)

The proof is a reduction from **Halting from blank tape**: For any given TM M we construct a tile set that admits a periodic tiling if and only if M halts from the blank initial tape.

As expected, an aperiodic prototile set will be needed in the construction.

Theorem. The problem **Periodic tiling problem** is undecidable.

Proof.

For any given Turing machine $M = (S, \Gamma, \delta, s, h, b)$, we construct three layers of tiles.

(i) First layer: we start with the same Wang set $\mathcal{P}_M$ that simulates M as before.

We add to these tiles the following halting tiles for every tape letter $x \in \Gamma$:

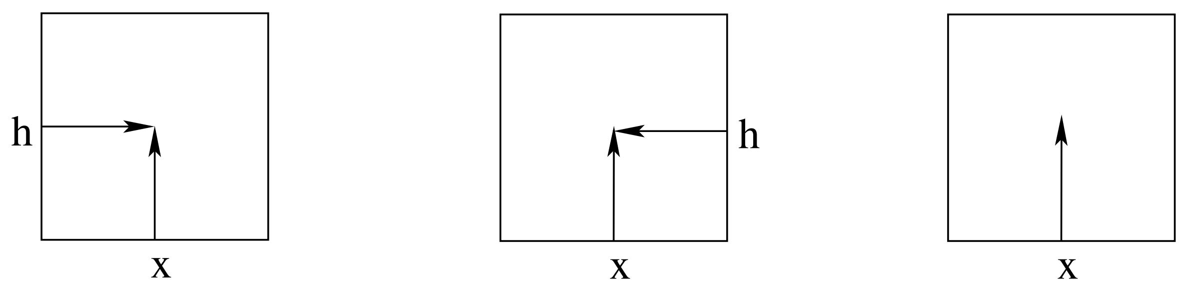

Here h is the halting state. Now a tiling becomes possible even if the Turing machine halts:the state component simply disappears from the configuration. Using the third tile, the entire configuration can then disappear.

We also add the following tile

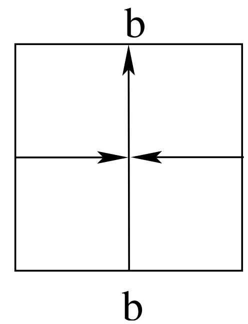

to the start tiles of the Turing machine. This tile allows the same horizontal row to contain several copies of the start configuration of the Turing machine.

With these tiles a periodic tiling becomes possible if the Turing machine halts from the blank initial tape:

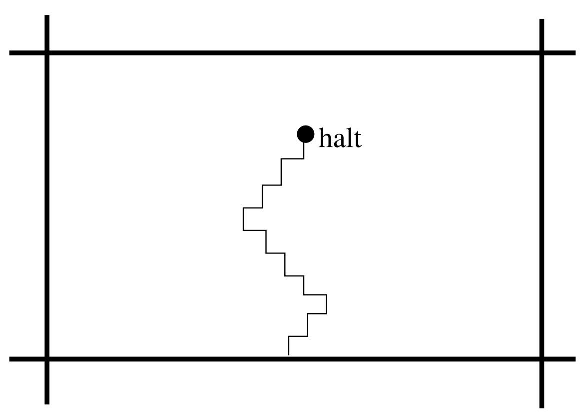

(ii) Second layer: We take one fixed aperiodic tile set $\mathcal{P}$. This can be, for example, Robinson's aperiodic tile set.

(iii) Third layer: A fixed set $\mathcal{Q}$ of four tiles:

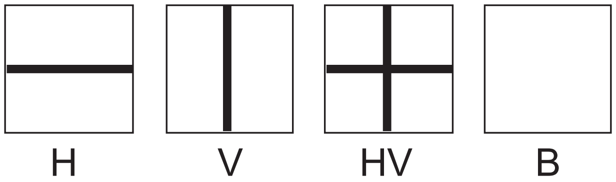

The black lines are called **fault lines** and these indicate places where a tiling error is allowed on the aperiodic tiles of layer two. The fault lines continue across tile boundaries so the lines cut through the plane horizontally or vertically:

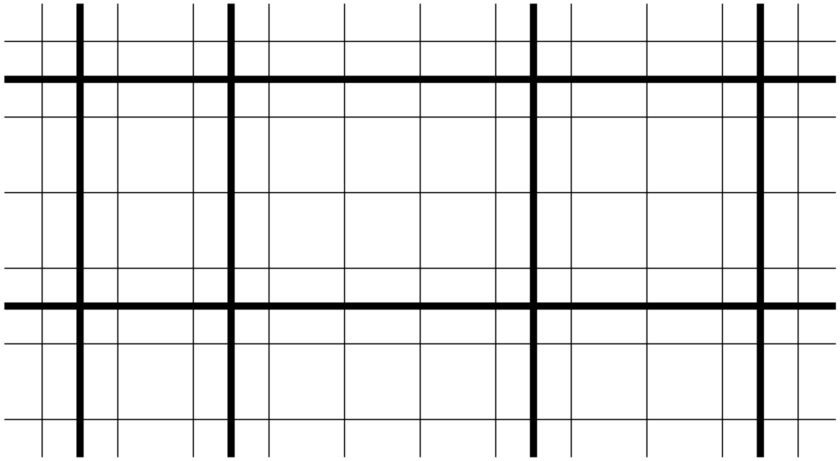

Using the SFT formalism we add further constraints on the third layer to force the following property:

(*h*) If a tiling contains at least two horizontal fault lines then it also contains at least two vertical fault lines.

To establish property (*h*) we make two versions of the empty tile B without fault lines: one called **red** and the other one called **blue**, and add the following local constraints on validity of tilings:

• The north neighbor of a horizontal fault line (tile H) must be a blue version of B,

• The south neighbor of a horizontal fault line (tile H) must be a red version of B,

• The north neighbor of a horizontal row of three blue B's is a blue B,

• The south neighbor of a horizontal row of three red B's is a red B.

These local constraints can be satisfied in $2n \times 2n$ square patterns of fault lines, for any n:

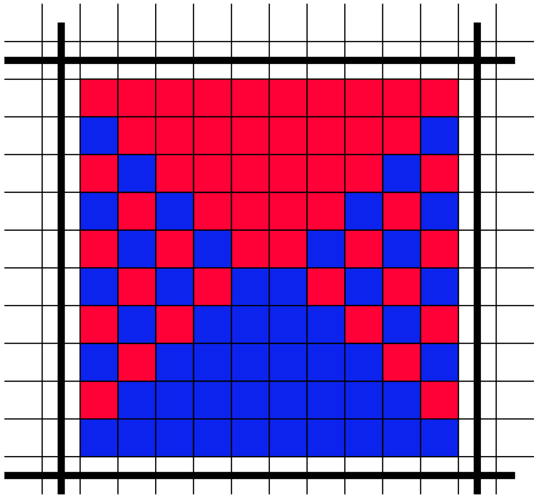

(*h*) If a tiling contains at least two horizontal fault lines then it also contains at least two vertical fault lines.

Proof. If two horizontal fault lines at vertical distance n, and a gap of horizontal length $2n$ without vertical fault lines

$\Longrightarrow$ a contradiction

We make an analogous red/blue coloring also for the perpendicular direction (so the no-fault-line tile comes in four color combinations.) Then we also have

(*v*) If a tiling contains at least two vertical fault lines then it also contains at least two horizontal fault lines.

Combining colorings: We have four versions of tile B (red/red, red/blue, blue/red, blue/blue), yielding a set $\mathcal{Q}$ that satisfies:

(*) If a tiling contains at least two parallel fault lines then it also contains at least two fault lines in the perpendicular direction.

Combining the three layers gives the final tile set

$$\mathcal{P}_M \times \mathcal{P} \times \mathcal{Q}.$$

Tiles are thus triplets (a, b, c).

Each layer has its own matching conditions as discussed above.

Next we add "inter-layer" local matching conditions that tie the layers together.

In tile (a, b, c):

(1) if c contains a fault line then the tiling rule is not enforced on the second layer b. The idea is to allow periodic tilings (even though $\mathcal{P}$ is aperiodic) in the presence of fault lines.

Combining the three layers gives the final tile set

$$\mathcal{P}_M \times \mathcal{P} \times \mathcal{Q}.$$

Tiles are thus triplets (a, b, c).

Each layer has its own matching conditions as discussed above.

Next we add "inter-layer" local matching conditions that tie the layers together.

In tile (a, b, c):

(2) if c contains only the horizontal fault line then the first component a must be one of the start tiles

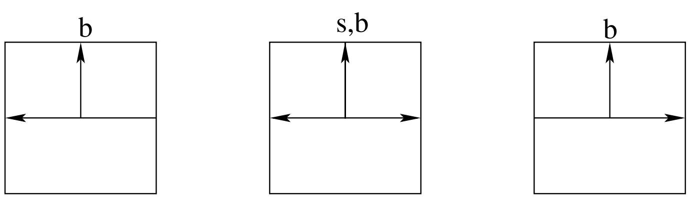

Combining the three layers gives the final tile set

$$\mathcal{P}_M \times \mathcal{P} \times \mathcal{Q}.$$

Tiles are thus triplets (a, b, c).

Each layer has its own matching conditions as discussed above.

Next we add "inter-layer" local matching conditions that tie the layers together.

In tile (a, b, c):

(3) if c contains only the vertical fault line then the first component a must be

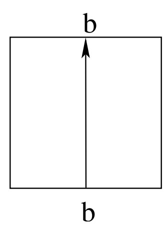

Combining the three layers gives the final tile set

$$\mathcal{P}_M \times \mathcal{P} \times \mathcal{Q}.$$

Tiles are thus triplets (a, b, c).

Each layer has its own matching conditions as discussed above.

Next we add "inter-layer" local matching conditions that tie the layers together.

In tile (a, b, c):

(4) if c contains both horizontal and vertical fault lines then a must be

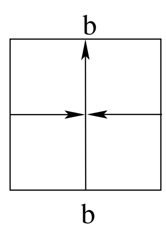

Consider a rectangle bordered by fault lines. Constraints **(2)**, **(3)** and **(4)** force the lower border to contain the blank tape and a single Turing machine in its initial state s. The vertical fault lines are forced to contain the black symbol only (no TM state component), and thus Turing machine is never allowed to reach a vertical fault line.

The construction of the tile set is now complete. The tile set can be effectively constructed for any given Turing machine. It is enough to prove that:

TM M halts from the blank initial tape
$\Longleftrightarrow$
the tile set admits a periodic tiling.

Proof. ($\Longrightarrow$)

Suppose that M halts in n steps. Then the tiles admit a valid periodic tiling with the horizontal and vertical period $2n$:

• On the **third layer** the fault lines partition the space into squares of size $2n \times 2n$.

The construction of the tile set is now complete. The tile set can be effectively constructed for any given Turing machine. It is enough to prove that:

TM M halts from the blank initial tape
$\Longleftrightarrow$
the tile set admits a periodic tiling.

Proof. ($\Longrightarrow$)

Suppose that M halts in n steps. Then the tiles admit a valid periodic tiling with the horizontal and vertical period $2n$:

• The **second layer** contains a correctly tiled $2n \times 2n$ square, repeated inside the squares between the fault lines. The tiling of the second layer fails on some tiles along the fault lines, but that is allowed.

The construction of the tile set is now complete. The tile set can be effectively constructed for any given Turing machine. It is enough to prove that:

TM M halts from the blank initial tape
$\Longleftrightarrow$
the tile set admits a periodic tiling.

Proof. $(\Longrightarrow)$

Suppose that M halts in n steps. Then the tiles admit a valid periodic tiling with the horizontal and vertical period $2n$:

- The **first layer** consists of the halting simulation of the Turing machine M. The start of the simulation begins at the bottom of each $2n \times 2n$ square. The entire simulation fits inside the $2n \times 2n$ square, because the machine halts after n steps. The halting tiles allow the disappearance of the Turing machine configuration before the next vertical line is reached. Hence a periodic tiling is admitted.

TM M halts from the blank initial tape
$\Longleftrightarrow$
the tile set admits a periodic tiling.

Proof. $(\Longleftarrow)$

Suppose a periodic tiling exists.

• Because $\mathcal{P}$ is an aperiodic set, there must be a place on the tiling where the tiling in the second layer is incorrect. This is possible only if there is a fault line in that location. Because the tiling is periodic, this implies the existence of infinitely many parallel fault lines. By property (*) this further implies the presence of a rectangle bordered by fault lines.

• Consider the first layer of one such rectangle. The bottom is forced to contain a (finite segment) of the start configuration of the Turing machine. The tiles in $\mathcal{P}_M$ force the following rows to simulate the Turing machine moves one-by-one. The Turing machine must halt before the simulation reaches the upper border of the rectangle.