

Knight tiles: particles and collisions in the realm of 4-way deterministic tilings

Bastien Le Gloannec and Nicolas Ollinger

Univ. Orléans, INSA Centre Val de Loire, LIFO EA 4022, FR-45067 Orléans, France
{Bastien.Le-Gloannec,Nicolas.Ollinger}@univ-orleans.fr

Abstract. Particles and collisions are convenient construction tools to compute inside tilings and enforce complex sets of tilings with simple tile-sets. Locally enforceable particles being incompatible with expansivity in the orthogonal direction, a compromise has to be found to combine both notions in a same tileset. This paper introduces knight tiles: a framework to construct 4-way deterministic tilings, that is tilings completely determined by any infinite diagonal of tiles, for which local particles and collisions with many slopes can still be constructed while being expansive in infinitely many directions. The framework is then illustrated by an elegant yet simple construction to mark a diagonal with a 4-way deterministic knight tileset.

Key words: Deterministic tiles; Domino Problem; Tiling Problem; Expansive Subdynamics; Particles; Wang tiles

A tiling is a coloring of the discrete plane \mathbb{Z}^2 assigning a tile from a finite tileset to each position so that the local tiling rules associated to the tileset are satisfied for every tile in its neighborhood. Wang tiles provide a convenient and universal syntactic description of tiling rules: a Wang tile is a unit square with colored edges and the tiling rule requires adjacent squares to share a same color along their common edge – tiling constraints are readable directly on the tiling. Starting with the study of the Domino Problem [1,2], computations have been successfully embedded into tilings to enforce computational phenomena and prove undecidability and complexity results. Following what is done in the case of cellular automata, particles and collisions have been successfully used as a tool to transmit information quanta through tilings and mark positions.

A tiling is 4-way deterministic if any infinite diagonal strip of tiles uniquely determines the whole tiling. A syntactic way to enforce such a property on a tileset is to require any consecutive pair of colors along its edges to uniquely determine a tile. Kari and Papasoglu [3] were able to prove that it is still possible to construct aperiodic tilesets under such constraint as are every family of tilings generated by substitution systems [4]. Ten years later, Lukkarila [5] showed that Turing machines can still be simulated in this setting and proved that the Domino Problem is still undecidable for 4-way deterministic tilesets. The construction is subtle and quite involved, building on previous construction

by Kari and Papasoglu [3]. The construction has to be quite involved in particular because such a syntactic constraint enforces the expansivity [6] of the tilings in every directions but two, prohibiting the use of locally defined particles that are neither horizontal nor vertical. In this paper, inspired by the partitioned cellular automata [7], we propose to replace Wang tiles with knight tiles, keeping a syntactic condition to enforce 4-way determinism and gaining back particles and collisions.

Similar to a Wang tile, a knight tile is a colored unit square sharing colors not only with its 4 direct neighbors but also with the 8 neighbors located at a chess knight move. Determinism along one diagonal is then simply expressed as a syntactic condition on tuples of colors. While keeping expansive in infinitely many directions, knight tiles can be combined to construct particles along slopes that are not too steep. The interest of this family of tilesets is illustrated by considering two problems handled with complicated constructions in Lukkarila [5].

The first construction is the key ingredient to the proof of the undecidability of the Domino Problem: solving the seeded Domino Problem by simulating a Turing machine.

The second construction is a key ingredient for the previous construction in the classical case: marking a diagonal — a slope that cannot be obtained directly by knight tiles using a particle. Indeed, Lukkarila [5] asks “Could there be a significantly simpler tile set for drawing a single diagonal line 4-way deterministically?” The construction provided here is significantly simpler, combining time-symmetry and the infinite Firing Squad technique by Kari [8] with properties of the Thue-Morse substitution to enforce 4-way determinism.

1 Preliminary definitions

Wang tiles and deterministic tilesets Given a finite alphabet of *colors* \mathcal{C} , a *Wang tile* is an oriented unit square tile with one color on each side. Formally, it is a quadruple $t \in \mathcal{C}^4$ whose four components are identified with the four directions $\{W, S, E, N\}$ and, for convenience, denoted as t_W, t_S, t_E and t_N . A *Wang tileset* τ is a finite set of Wang tiles. A *tiling* by τ is a map $T : \mathbb{Z}^2 \rightarrow \tau$ associating a tile of τ to each cell of the discrete plane \mathbb{Z}^2 such that two adjacent tiles (for the 4-connectedness) share the same color on their common edge. More formally, a tiling T satisfies the following constraints, for all $(x, y) \in \mathbb{Z}^2$: $T(x, y)_W = T(x - 1, y)_E$, $T(x, y)_S = T(x, y - 1)_N$, $T(x, y)_E = T(x + 1, y)_W$, $T(x, y)_N = T(x, y + 1)_S$. The set of all tilings by a tileset τ is denoted as \mathcal{X}_τ .

The tileset τ is *NE-deterministic* if for all couples of tiles $(w, s) \in \tau^2$, there exists at most one tile $t \in \tau$ simultaneously compatible to the west with w and to the south with s : $w_E = t_W$ and $s_N = t_S$. In the case of Wang tiles, one can equivalently say that any tile $t \in \tau$ is uniquely identified by its couple of colors (t_W, t_S) . *{SW, SE, NW}-determinism* is defined symmetrically. A tileset is *4-way deterministic* if it is simultaneously SW, SE, NE and NW-deterministic.

Subshifts and expansiveness Given a finite alphabet Σ , a Σ -coloring (or simply *coloring* in the absence of ambiguity) of the discrete plane is a map $c : \mathbb{Z}^2 \rightarrow \Sigma$. For all $u \in \mathbb{Z}^2$, we define the translation σ_u over colorings by $\sigma_u(c)(x) = c(x - u)$ for all $c \in \Sigma^{\mathbb{Z}^2}$ and $x \in \mathbb{Z}^2$. A coloring $c \in \Sigma^{\mathbb{Z}^2}$ is *periodic* of period $p \in \mathbb{Z}^2$ if $\sigma_p(c) = c$. The set $\Sigma^{\mathbb{Z}^2}$ of all Σ -colorings is endowed with the product (over \mathbb{Z}^2) topology of the discrete topology (over Σ). A *subshift* \mathcal{X} is a (topologically) closed and translation-invariant ($\forall u \in \mathbb{Z}^2, \sigma_u(\mathcal{X}) = \mathcal{X}$) subset of $\Sigma^{\mathbb{Z}^2}$. In particular, the set of tilings by a tileset τ is a subshift of $\tau^{\mathbb{Z}^2}$.

For all *slope* $\alpha \in \mathbb{R} \cup \{\infty\}$, let us denote as l_α the real line of slope α going through the origin. For all *radius* $\rho > 0$, let us define $L_\alpha(\rho) = (l_\alpha + [-\rho, \rho]\mathbf{u}) \cap \mathbb{Z}^2$ for $\mathbf{u} \in \mathbb{R}^2$ a unit vector orthogonal to l_α . $L_\alpha(\rho)$ is the discrete thick line of slope α and width 2ρ centered on the origin. A slope α is a *direction of expansiveness* of a subshift $\mathcal{X} \subseteq \Sigma^{\mathbb{Z}^2}$ if there exists a radius $\rho > 0$ such that for all $x, y \in \mathcal{X}$, $x|_{L_\alpha(\rho)} = y|_{L_\alpha(\rho)} \implies x = y$. In particular, if τ is a 4-way deterministic Wang tileset, \mathcal{X}_τ is expansive in (at least) all directions of $\mathbb{R} \setminus \{0\}$.

2 From Wang tiles to knight tiles

The notion of 4-way determinism introduces a very strong constraint on the tileset. Indeed, it strongly limits its capacity to construct particles locally. To mix particles and determinism, one might loosen the constraints by generalizing the notion of local determinism to a broader *radius* $r \geq 1$ of determinism.

Radius of determinism A tileset τ is *NE deterministic with radius* r if for all *valid* (i.e. containing no tiling error along its inner edges) $(2r + 1) \times (2r + 1)$ square pattern by τ , the center tile is perfectly determined by the $2r$ tiles at positions $(1, 2r), (2, 2r - 1), \dots, (2r - 1, 1)$. Determinism with radius r in the three other diagonal directions (SW, SE and NW) is defined symmetrically. The tileset is *4-way deterministic with radius* r if it is simultaneously deterministic with radius r in the four diagonal directions.

The different radiuses can be compared through the expansiveness of their tilings as pointed out by the following proposition.

Proposition 1. *If τ is a 4-way deterministic tileset with radius r , then \mathcal{X}_τ is (at least) expansive in directions $\left] -\frac{r}{r-1}, -\frac{r-1}{r} \right[\cup \left] \frac{r-1}{r}, \frac{r}{r-1} \right[$.*

The expansiveness directions are actually tightly related to the *particles* that can be locally realized. Expansiveness is, up to a change in the radius ρ , independent of the profile chosen for the slope. At a fixed radius of determinism, the ability to mark a line however depends on its profile. Let us denote the horizontal and vertical unit vectors of \mathbb{Z}^2 as $\mathbf{e}_1 = (1, 0)$ and $\mathbf{e}_2 = (0, 1)$. For any slope $\alpha \in \mathbb{Q} \cup \{\infty\}$, a (*periodic*) *profile* for α is a finite sequence $(u_0, \dots, u_{k-1}) \in \{\mathbf{e}_1, \mathbf{e}_2\}^k$ such that $\alpha = \frac{\sum_i u_i \cdot \mathbf{e}_2}{\sum_i u_i \cdot \mathbf{e}_1}$. Given a profile $P = (u_0, \dots, u_{k-1})$ for α , we define the associated subshift $\mathcal{L}_P \subseteq \{0, 1\}^{\mathbb{Z}^2}$ as the subshift generated by (i.e. the smallest

subshift containing) the configuration c_P containing a discrete 4-connected line of slope α drawn following the profile P , i.e. formally defined by $c_P(x) = 1$ if $x = \sum_{0 \leq i \leq n} u_{i \bmod k}$ or $x = -\sum_{-n \leq i < 0} u_{i \bmod k}$ for some $n \geq 0$, and $c_P(x) = 0$ for all other $x \in \mathbb{Z}^2$. Note that \mathcal{L}_P exactly contains all translated versions of c_P , plus the blank coloring (all 0s) for compactness reasons. A coloring $c \in \Sigma^{\mathbb{Z}^2}$, where Σ contains an identified blank color $c_0 \in \Sigma$, is a *particle* of direction $\alpha \in \mathbb{Q} \cup \{\infty\}$ if it is periodic in direction α (i.e. c admits a periodicity vector $(p_x, p_y) \in \mathbb{Z}^2$ such that $\alpha = \frac{p_y}{p_x}$) and ultimately constant¹ equal to c_0 in any other direction. This equivalently means that all non-blank cells of c are contained in $L_\alpha(\rho) + \Delta$ for some $\rho > 0$ and $\Delta \in \mathbb{Z}^2$. We say that a slope $\alpha \in \mathbb{Q} \cup \{\infty\}$ is *locally realized* by a tileset τ if there exists a projection $\pi : \tau \rightarrow \{0, 1\}$ (naturally extended to colorings $\pi : \tau^{\mathbb{Z}^2} \rightarrow \{0, 1\}^{\mathbb{Z}^2}$) such that there exists a profile P for the slope α such that $\pi(\mathcal{X}_\tau) = \mathcal{L}_P$ and every tiling of \mathcal{X}_τ is a particle. The previous definition introduces a reinforced notion of *soficity* for \mathcal{L}_P .

The following results conclude our remarks on radiuses of determinism.

Proposition 2. *If a slope α is locally realizable by a tileset τ then α is not a direction of expansiveness for \mathcal{X}_τ .*

Proposition 3. *The slopes that are locally realizable by 4-way deterministic tilesets at radius r are exactly $\mathbb{Q} \cap \left(\left[-\infty, -\frac{r}{r-1} \right] \cup \left[-\frac{r-1}{r}, \frac{r-1}{r} \right] \cup \left[\frac{r}{r-1}, +\infty \right] \right)$.*

Knight tiles Whereas classical 4-way determinism is a purely syntactic property of the tileset, to check determinism at radius r one has to consider all tilings of $(2r + 1) \times (2r + 1)$ squares. We introduce deterministic knight tiles as a convenient and purely syntactic notion of determinism at radius 2 that can be checked directly on the tiles.

Given a finite alphabet \mathcal{C} of *colors*, a *knight tile* is formally a 12-tuple of \mathcal{C}^{12} and a *knight tileset* is a finite set of knight tiles. For convenience, each of the twelve components of a knight tile will be identified by a direction among $\{W, S, E, N, WS, SW, SE, ES, EN, NE, NW, WN\}$ and for a knight tile $T \in \mathcal{C}^{12}$, we will for instance denote by $T_W \in \mathcal{C}$ the corresponding w color. A *tiling* of the discrete plane by a knight tileset τ is a map $t : \mathbb{Z}^2 \rightarrow \tau$ satisfying the following constraints, for all $(x, y) \in \mathbb{Z}^2$: $t(x, y)_W = t(x - 1, y)_E$, $t(x, y)_S = t(x, y - 1)_N$, $t(x, y)_E = t(x + 1, y)_W$, $t(x, y)_N = t(x, y + 1)_S$, $t(x, y)_{WS} = t(x - 2, y - 1)_{EN}$, $t(x, y)_{SW} = t(x - 1, y - 2)_{NE}$, $t(x, y)_{SE} = t(x + 1, y - 2)_{NW}$, $t(x, y)_{ES} = t(x + 2, y - 1)_{WN}$, $t(x, y)_{EN} = t(x + 2, y + 1)_{WS}$, $t(x, y)_{NE} = t(x + 1, y + 2)_{SW}$, $t(x, y)_{NW} = t(x - 1, y + 2)_{SE}$, and $t(x, y)_{WN} = t(x - 2, y + 1)_{ES}$. This means that each tile of a tiling shares each of its different color components with one of the twelve neighboring tiles represented on the figure 1(a).

A knight tileset τ is *NE-deterministic* if any tile $t \in \tau$ is uniquely identified by its quadruple of colors $(t_{WN}, t_W, t_S, t_{SE})$, i.e. there is at most one tile in the set that is compatible with the four colors pointed out on the figure 1(b). Determinism

¹ Note that all our results would remain true replacing “ultimately constant” by “ultimately periodic” in the definition of a particle.




Fig. 1. Knight tiles and determinism

in the three other diagonal directions (SW, SE and NW) is defined symmetrically. A knight tileset is *4-way deterministic* if it is simultaneously deterministic in the four diagonal directions SW, SE, NE and NW.

The rest of the paper is dedicated to advocate for the use of this notion of radius-2 determinism by showing that constructions that are painful to handle in the classical 4-way deterministic setting can be treated with particles and collisions in the 4-way deterministic knight setting.

3 The seeded Knight Domino Problem is undecidable

Simulating a Turing machine with classical 4-way tiles requires a complicated machinery [5] involving particular aperiodic tilesets. Knight colors provide enough flexibility to handle it directly and prove the undecidability of the Domino Problem with a seed tile in a classical way.

Problem 1 (Domino Problem with a seed tile). Given a tileset τ and a specified seed tile $t_0 \in \tau$, does there exist a tiling of \mathbb{Z}^2 by τ using (at least once) the tile t_0 ?

The reader is assumed familiar with Turing machines. As every Turing computation can be made reversible [9,10], we directly work with reversible Turing machines. A *reversible Turing machine* (RTM) is a 5-tuple $(\Sigma, Q, \overleftarrow{q}_i, F, \delta)$ where Σ is the tape (finite) alphabet, Q the finite set of states. Before defining the remaining elements of the tuple, let us state that the head always moves at each transition and we denote its two possible moves by $\{\leftarrow, \rightarrow\}$. Let us also define $\overleftrightarrow{Q} = Q \times \{\leftarrow, \rightarrow\}$ whose elements will be written \overrightarrow{q} (resp. \overleftarrow{q}) to denote (q, \rightarrow) (resp. (q, \leftarrow)). Then we define the **partial injective** transition map $\delta : \overleftrightarrow{Q} \times \Sigma \rightarrow \overleftrightarrow{Q} \times \Sigma$. Finally $\overleftarrow{q}_i \in \overleftrightarrow{Q}$ is the initial (oriented) state and $F \subset \overleftrightarrow{Q}$ is the set of final (oriented) states (and we will assume that δ is not defined on states of F). This is an acceptable model for reversible Turing machines that will be furthermore well-fitted to our purpose.

The tileset described on figure 2, where colors are conveniently represented by arrows that must be go uninterrupted across tile edges, is a 4-way deterministic knight tileset that simulates a given RTM $(\Sigma, Q, \overleftarrow{q}_i, F, \delta)$. Firstly note that

the only knight color components used in this simulation are $\{\text{SW}, \text{SE}, \text{NE}, \text{NW}\}$ hence there is no ambiguity on the knight colors represented on the figure 2. To be able to use efficiently these knight colors for determinism, one must slow down the simulation. To that purpose, each transition of the Turing machine is decomposed into three steps in this fixed order: 1. a state transition step where the transition is done but the head does not yet move, though its move is already contained in the state of \vec{Q} (the corresponding tiles are represented on the third line of the figure); 2. a waiting step where nothing happens (second line of the figure); 3. a move step where the head finally moves (first line of the figure). The Wang constraints to enforce this order for the steps of the computation process are not explicitly represented on the tiles of figure 2, but it is a simple counting modulo 3 in the vertical colors. Assuming that a Turing computation is correctly initialized in a tiling, then each line of the space-time diagram of the Turing machine can be read once every three lines of the tiling.




Fig. 2. Deterministic knight tiles simulating a RTM

One can enforce the initialization of the Turing computation using a seed tile and the tiles of the figure 3(a) that force a blank tape containing a unique Turing machine head to appear in the tilings. The labels “ ≤ -2 ”, “ -1 ”, “ 0 ”, “ 1 ” and “ ≥ 2 ” appearing on the tiles denote some specifications for the horizontal colors: the tile with index 0 is only compatible to the left (resp. right) with the tile with index -1 (resp. 1) and, to the left (resp. right) of the -1 (resp. 1) tile, only the “ ≤ -2 ” (resp. “ ≥ 2 ”) tile can appear. However, these tiles cannot directly be added to the previous simulation tile set without losing the 4-way determinism. Indeed, it not difficult to see that, as is, predicting the tiles with index -1 or 1 is not possible in every direction. To solve that problem, one can add a very simple layer of information represented by the Wang tiles of the figure 3(b): a blank tile g_0 and a tile g_1 carrying a vertical “ghost” signal (that must be vertically preserved along columns of the tilings). Denoting as τ_T the Turing simulation

tiles of the figure 2, τ_i^1 the three initialization tiles with indices $\{-1, 0, 1\}$ of the figure 3(a), τ_i^2 the two initialization tiles with labels “ ≤ -2 ” and “ ≥ 2 ” of the figure 3(a) and b the blank tile of the figure 3(a), we define the two-layered knight tileset $\tau = ((\tau_T \cup \{b\}) \times \{g_0, g_1\}) \cup (\tau_i^1 \times \{g_1\}) \cup (\tau_i^2 \times \{g_0\})$, which can be interpreted as a standard knight tileset on couples of colors of each layer. In a tiling, the three consecutive columns containing the initialization tiles with indices $\{-1, 0, 1\}$ are the only one carrying a ghost signal (tile g_1) on the second layer. This allows to predict these tiles in every direction, hence the knight tileset τ is 4-way deterministic. As the transition tiles are not defined on final states, τ tiles the plane if and only if the simulated Turing machine does not halt on the blank tape. It is then straightforward to derive the following result.

Theorem 1. *The Domino Problem with a seed tile is undecidable for 4-way deterministic knight tilesets.*

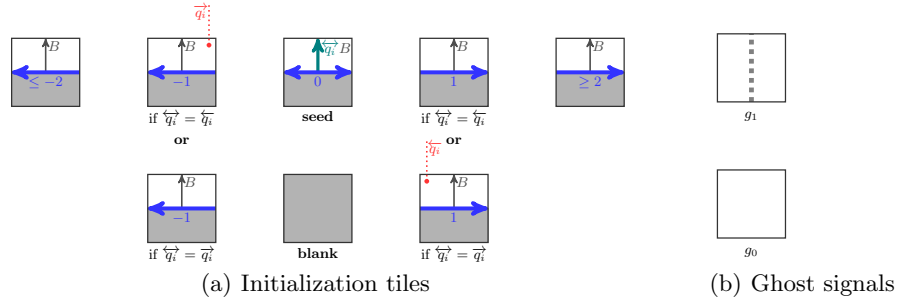


Fig. 3. Additional knight tiles to initialize the Turing computation

4 Marking a diagonal

Considering deterministic knight tiles allows to easily build particules that are not realizable in the usual deterministic Wang framework (where the only realizable slopes are 0 and ∞ by proposition 3) as illustrated by the following result.

Proposition 4. *The slopes that are locally realizable using 4-way deterministic knight tiles are exactly $\mathbb{Q} \cap ([-\infty, -2] \cup [-\frac{1}{2}, \frac{1}{2}] \cup [2, +\infty])$.*

Our main construction will rely on these realizable slopes to mark a diagonal in a lighter way than the construction of [5] in the Wang case. The general idea is to use a fundamentally one-dimensional hierarchical structure of signals similar to an infinite version of Minsky’s classical solution to the *Firing Squad* synchronization problem on cellular automata, as what is done for instance in [8].

This structure has moreover to be reversible so that the tileset can be made deterministic in two opposite directions, say NE and SW. The structure also has to be compatible with determinism for knight tiles in the two orthogonal directions NW and SE. The reader might keep in mind that only knight colors $\{\text{WS, SW, EN, NE}\}$ will be used. Thus the tileset will be deterministic at radius 1 (classical Wang tiles determinism) in directions NE and SW, while it will be deterministic at radius 2 (using knights) in directions NW and SE.

The general structure is described on figure 4 where the red dots represent the diagonal to be marked, referred to as the *fire line*. For convenience, the four kinds of signals are named H, H', V and V' according to the figure. Their respective slopes are $0, \frac{1}{3}, \frac{2}{3}$ and ∞ . A binary hierarchical structure is used to mark some points along successive *front lines* until the granularity of space allows all the marked points of a front to be sufficiently close to decide by a local rule to *fire* at that point (and consequently mark the diagonal). Let us enumerate the front lines by their *rank* starting at 0 for the fire line. The points of the front of rank k are regularly arranged and two consecutive points are separated by $2^k - 1$ non-marked positions. The marked points of the front of rank k will be referred to as *pillars*, which correspond to the points that are marked on the front of rank $k + 1$ (i.e. points that were already marked seen things as a firing squad with time running towards NE), and *middles*, which exactly correspond to the middle positions between pillars. Referring to the figure 4 for the orientation, building a local rule for the construction to be deterministic in the NE direction is not particularly difficult and goes straightforward using the represented signals. By contrast, making this construction deterministic in the SW direction is more challenging as one has to be able to locally and deterministically distinguish between pillars and middles. For this, one need an infinite source of well-structured alternating bits sequences. A way to get such information properly is to resort to a sequence that is a fixed-point of a well-chosen substitution.

Let us define the *Thue-Morse substitution* (see [11]) $s : \{0, 1\} \rightarrow \{0, 1\}^*$ by $s(0) = 01$ and $(1) = 10$. We naturally extend the definition of s to finite words of $\{0, 1\}^*$ (resp. infinite words of $\{0, 1\}^{\mathbb{N}}$) defining, for all $u \in \{0, 1\}^*$ (resp. $u \in \{0, 1\}^{\mathbb{N}}$), $s(u)_{2i+k} = s(u_i)_k$ for all $i \in \{0, \dots, |u| - 1\}$ (resp. $i \in \mathbb{N}$) and $k \in \{0, 1\}$. The *Thue-Morse word* $\mathcal{T} \in \{0, 1\}^{\mathbb{N}}$ is the unique fixed-point of s starting with the letter 0: $\mathcal{T} = \lim_{n \rightarrow +\infty} s^n(0)$ that exists as $s(0)_1 = 0$ and is infinite as $|s(0)| > 1$. As the Thue-Morse word is a fixed-point for s , it can be unsubstituted into itself by s placing bars every two letters starting with a bar at position 0. A given factor v of the Thue-Morse word is said *even* if there exists a position $p \in \mathbb{N}$ even such that $\mathcal{T}_p \cdots \mathcal{T}_{|v|} = v$, which means that there is a bar just before its first letter v_0 (i.e. between \mathcal{T}_{p-1} and \mathcal{T}_p in the Thue-Morse word) in one of the possible decompositions of v . Complementarily, a factor is said *odd* if it appears at an odd position p in the Thus-Morse word, which means that the bar is between v_1 and v_2 in a decomposition of v . A factor can simultaneously be even and odd if it appears in the Thue-Morse word at both even and odd positions: for instance 011 appears at positions 1 and 7. However,




Fig. 4. General scheme of the structure

the following lemma points out the fact that a sufficiently large factor cannot be simultaneously even and odd.

Lemma 1. *Any factor of \mathcal{T} of length at least four is either even or odd, which means that it admits a unique decomposition.*

Hence it is sufficient to look at factors of size four to be able to determine unambiguously their alignment. Observe that moving letter by letter a window of size four over the Thue-Morse word, one sees an alternating sequence of even/odd factors. This will constitute the alternating bits sequence we require.

Now that all the required objects are defined, each signal of the construction will henceforth carry a factor of size four of the Thue-Morse sequence. The signals H and V only carry the even words while H' and V' only carry odd words. When two signals meet on a marked point, we require them to carry the same factor u ($|u| = 4$) and we substitute it by s ($|s(u)| = 8$) to derive the four factors carried by the four “outgoing” (seeing time going NE) signals: the produced type V signal must carry $s(u)_1 \cdots s(u)_4$, the type V' must carry $s(u)_2 \cdots s(u)_5$, the type H' must carry $s(u)_4 \cdots s(u)_7$ and the type H must carry $s(u)_5 \cdots s(u)_8$. Firstly remark that knowing $s(u)$, one uniquely deduces u . Although there is no “middle signal” to carry $s(u)_3 \cdots s(u)_6$, it can easily be deduced from the four other derived words we dispose of.

This defines a coherent structure. Indeed, let us consider any bi-infinite word of $\{0, 1\}^{\mathbb{Z}}$ such that every factor appears in the Thue-Morse sequence², “written” along a front line of rank $k > 1$ in the following sense: two consecutive

² For topological reasons, such a word exists. The set of all these words is actually a non-empty subshift of $\{0, 1\}^{\mathbb{Z}}$. All of them are fixed-point for the Thue-Morse substitution.




Fig. 5. Marking a diagonal

marked points hold two consecutive factors of size 4 with a one-letter shift. The previously described local mechanisms of substitution/unsubstitution on marked points enforce the same word (as it is a fixed-point) to appear on all the other front lines and to be used as a resource to properly alternate between middles and pillars on a front line, coherently with other lines so that a proper global binary hierarchical structure is enforced.

A realization of this structure using knight tiles is given on the figure 5 where the colors are represented using some convenient signal notations. The knight colors are not represented on this figure but are all illustrated on the figure 6. H and V signals, of respective slopes 0 and ∞ , carry even words and meet on pillars (solid black squares on the figures). The pilar tiles force their two incoming H and V signals to carry the same factor u . As the four outgoing factors are derived from this common word, a pilar tile is perfectly determined by u , which is explicitly written next to each pillar on the figure 5. Similarly, H' and V' signals, of respective slopes $\frac{1}{3}$ and $\frac{2}{3}$ realized using particular profiles for reasons that will be exposed later and two different types of knight colors to distinguish the two different types of angles that appear in that profile as illustrated by figures 6(a) and 6(b), carry odd factors. These signals meet on middles, that are represented by gray squares on the figures. The middle tiles force their two incoming H' and V' signals to carry the same factor u . As the four outgoing factors are derived from this common word, a middle tile is perfectly determined by u , which is explicitly written next to each middle on the figure 5.

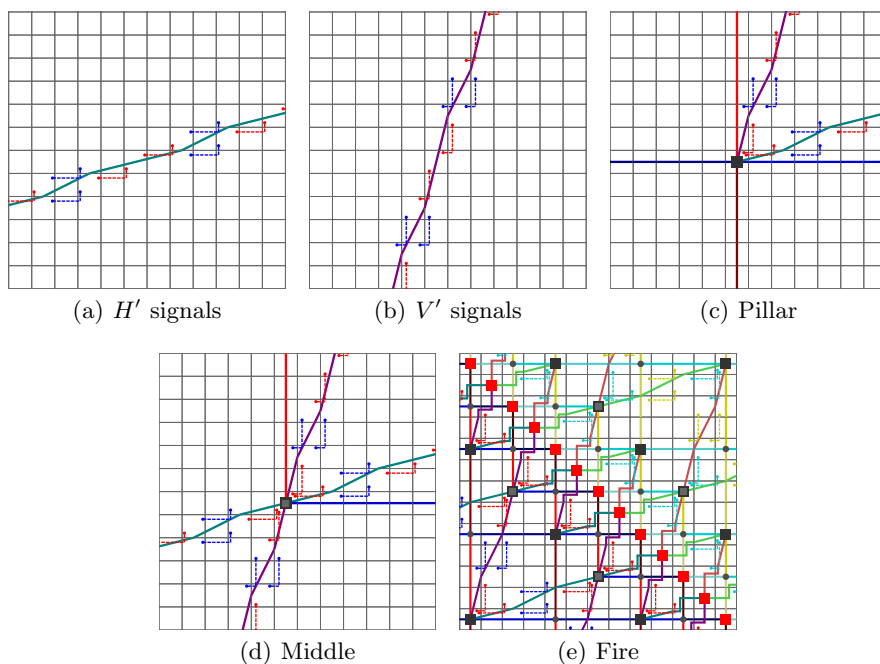


Fig. 6. Knight colors of the tileset

We do not dispose of enough space to formally describe the whole tileset. Although most of it is a rather straightforward translation of the previously exposed scheme (figure 4 and its symmetric along the fire line) into a tileset, several points certainly require explanations. This is what the remaining paragraphs are dedicated to.

First, when a H signal meets a V , they do not simply cross each other, there is actually a *collision* occurring (grey dot on the figures) and there are actually two kinds of H/V signals. When two signals of the first kind meet, they cross each other and by the way change into the second kind (darker color on the figures). When two signals of the second kind meet, they produce a pillar that will in turn produce two signals of the first kind. This collision has to be treated particularly when it is the last collision before the fire line, which can be detected as the collision tile also has to contain some H'/V' signals when it is near the fire line (see figure 6(e)). In that case, the first kind of H/V signals is changed into a third kind that produces a fire tile (instead of a pillar) when meeting.

The case of H' and V' signals requires some careful attention as we need to use knight colors for them to be deterministic in the NW and SE directions. The knight colors used basically contain the type of signal and the word it carries. However a knight must absolutely not cross a front line as in this case the determinism in one of the orthogonal NE or SW directions would fail. Without entering into the details, our solution consists in using the $(e_1, e_1, e_2, e_1, e_1, e_1, e_1, e_2)$

profile (alternating between 2 and 4 horizontal steps) such that one can make sure that the tiles that appear at both end of a signal are tiles with blank knight colors (hence no knight crosses the front line). We use the symmetrical profile for V' signals. All situations are illustrated by figure 6.

The fire line must appear when marked points are sufficiently close (and as for the front line, no knight should cross the fire line). We resort for that to several *ad-hoc* tiles that appear only near the fire line, as this was already illustrated by the third kind of H/V signals. This is illustrated by the figure 6(e).

On the other side of the fire line, say the NE side, the hierarchical structure is symmetrically dismantled. This is done by duplicating all the colors used in the SW part and defining a NE tile from each SW tile by applying a symmetry along the NW–SW diagonal on the tile and replacing each SW color by its NE duplicate. Only the fire tiles, that make the junction between both sides, simultaneously hold SW and NE colors. That way, the resulting tileset is *time-symmetric*: it is its own “inverse” up to a swap in colors.

To conclude, let us denote as κ the previously described knight tileset. The following theorem underlines the properties κ was builded for.

Theorem 2. *κ is a 4-way deterministic knight tileset and its set of tilings projects onto the diagonal subshift, i.e. $\{0, 1\}$ -colorings containing at most one diagonal of 1 and 0 everywhere else.*

References

1. Berger, R.: The Undecidability of the Domino Problem. PhD thesis, Harvard University (1964)
2. Robinson, R.: Undecidability and nonperiodicity for tilings of the plane. *Inventiones mathematicae* **12**(3) (1971) 177–209
3. Kari, J., Papasoglu, P.: Deterministic aperiodic tile sets. *Geometric And Functional Analysis* **9** (1999) 353–369
4. Le Gloannec, B., Ollinger, N.: Substitutions and strongly deterministic tilesets. In Cooper, S.B., Dawar, A., Löwe, B., eds.: *CiE. Volume 7318 of Lecture Notes in Computer Science.*, Springer (2012) 462–471
5. Lukkarila, V.: The 4-way deterministic tiling problem is undecidable. *Theor. Comput. Sci.* **410**(16) (2009) 1516–1533
6. Boyle, M., Lind, D.: Expansive subdynamics. *Transactions of the American Mathematical Society* **349**(1) (1997) 55–102
7. Morita, K.: Reversible simulation of one-dimensional irreversible cellular automata. *Theoretical Computer Science* **148**(1) (1995) 157 – 163
8. Kari, J.: Rice’s theorem for the limit sets of cellular automata. *Theor. Comput. Sci.* **127**(2) (1994) 229–254
9. Lecerf, Y.: Machines de Turing réversibles. *Comptes rendus de l’Académie française des Sciences* **257** (October 1963) 2597–2600
10. Bennett, C.H.: Logical reversibility of computation. *IBM J. Res. Dev.* **17**(6) (November 1973) 525–532
11. Pytheas Fogg, N.: *Substitutions in Dynamics, Arithmetics and Combinatorics.* Lecture Notes in Mathematics. Springer (2002)

Appendix: Proofs and complementary material

The complementary material for the reviewers follows the order of the article.

Proof of proposition 1 A perfectly formalized proof of this result would require to define some formal tools to properly manipulate the objects we are dealing with, like the *coding relation* of [6] (definition 3.1). For brevity and simplicity, we will only give a slightly informal yet convincing presentation of our arguments.

Proof sketch. First observe that, for $\alpha \in \mathbb{R} \cup \{\infty\}$, if there exist $\rho, \varepsilon > 0$ such that $L_\alpha(\rho) \subsetneq L_\alpha(\rho+\varepsilon)$ and for all $x \in \mathcal{X}_\tau$, $x|_{L_\alpha(\rho)}$ perfectly determines $x|_{L_\alpha(\rho+\varepsilon)}$, then \mathcal{X}_τ is expansive in direction α (as, “iterating” the argument, $x|_{L_\alpha(\rho)}$ perfectly determines x). One can refer to lemma 3.3 of [6] for more details.

For a tilingset SE and NE-deterministic at radius r , the two red cells of the subfigures of the figure 7 are uniquely determined by the $2r$ white ones. For any slope $\alpha \in \left] \frac{r-1}{r}, \frac{r}{r-1} \right[$, i.e. strictly between the two limit slopes in blue on the figures, there exists an appropriate $\rho_0 > 0$ such that, up to a translation, the red cells (i.e. their centers) can be made outside the “expansiveness window” $L_\alpha(\rho_0)$ while the (centers of the) white ones all are inside. This argument applied along any expansiveness window of radius $\rho \geq \rho_0$ ensures that it can be locally extended on both sides using the local rule of determinism. Said otherwise, as illustrated by the figure 8 where the rule of determinism is represented by the red lozenge lying on the line of slope 1, there exists $\varepsilon > 0$ such that $L_\alpha(\rho) \subsetneq L_\alpha(\rho+\varepsilon)$ and for all $x \in \mathcal{X}_\tau$, $x|_{L_\alpha(\rho)}$ perfectly determines $x|_{L_\alpha(\rho+\varepsilon)}$ which implies the expansiveness in direction α .

The other interval of directions of expansiveness is obtained symmetrically considering the two other directions of determinism. ■

Proof of proposition 2

Proof. For any $\rho > 0$, Consider the tiling that is projected to the blank coloring (all 0s) and a tiling that projects onto a line of slope α sufficiently far from the origin so that only blank tiles are visible in the “expansiveness window” $L_\alpha(\rho)$. Both tilings coincide on $L_\alpha(\rho)$ but are distinct, hence \mathcal{X}_τ cannot be expansive in direction α . ■

Proof of proposition 3

Proof. The realizability part is proved in details for radius 2 in the more constraining case of knight tiles (keep the same profile, simply remove the additional knight colors that are not required anymore as the local rule of determinism can already “see” the cells “emitting” these colors) in the proof of proposition 4. This generalizes straightforwardly to the case of deterministic Wang tiles of arbitrary radius.

The “exactly” part is a consequence of propositions 1 and 2. ■




Fig. 7. Limit expansive directions for determinism radius r




Fig. 8. Expansiveness in a direction $1 < \alpha < \frac{r}{r-1}$




Fig. 9. Example of tiling simulating a Turing computation

Proof of proposition 4

Proof. Firstly recall that the “exactly” part is due to propositions 1 and 2. Also remark that the cases of slopes 0 and ∞ are trivial, the case of the interval $[-\frac{1}{2}, 0]$ is equivalent to the case of the interval $]0, \frac{1}{2}]$ up to a horizontal symmetry on tiles and the case of the interval $[2, +\infty[$ (resp. $] -\infty, 2]$) is equivalent to the case of the interval $]0, \frac{1}{2}]$ (resp. $[-\frac{1}{2}, 0[$) up to a symmetry along the SW–NE (resp. NW–SE) diagonal on tiles. Hence it is enough to consider the case of a slope $\alpha \in \mathbb{Q} \cap]0, \frac{1}{2}]$. For such a slope, the determinism in directions NE and SW can be obtained without effort (as one directly sees the tiles of the slope in these directions), hence only knight colors of types $\{\text{WS}, \text{SW}, \text{EN}, \text{NE}\}$ could be useful here, and we will actually only use $\{\text{WS}, \text{EN}\}$. The difficulty is to be able to predict the appearance of the angles of the slope in the NW and SE directions. All we need is to provide a profile (considered cyclically) for α whose angles (i.e. transitions $\mathbf{e}_1/\mathbf{e}_2$) can be predicted at radius 2 using knight colors. It is not difficult to see that as long as no sub-sequence of the form $(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_1, \mathbf{e}_2)$ (resp. $(\mathbf{e}_2, \mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_1)$) appears in the profile, it is always possible to “see” the line with knight colors of types $\{\text{SW}, \text{EN}\}$ (resp. $\{\text{WS}, \text{NE}\}$) when going towards NW (resp. SE). If we limit ourselves to the single knight color EN (resp. WS), we have the more restrictive $(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_1, \mathbf{e}_2)$ (resp. $(\mathbf{e}_2, \mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_1)$) forbidden pattern. Let us write $\alpha = \frac{p}{q}$ with $q \geq 2p > 0$ and consider the profile $P = (\underbrace{\mathbf{e}_1, \mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_1, \mathbf{e}_1, \mathbf{e}_2}_{p \text{ times}}, \underbrace{\mathbf{e}_1, \dots, \mathbf{e}_1}_{q-2p \text{ times}})$.

It satisfies the required property and using a special knight color to mark each of the angles of the line, one straightforwardly builds 4-way deterministic knight tiles to locally realize the line. \blacksquare

Parabola We give here without entering into the details an example of a non-trivial system of particles and collisions: the parabola of the figure 10. Such a system would be technically very complex to realize using 4-way deterministic Wang tiles. It is however rather simple to achieve in the 4-way deterministic knight setting. Only the knight colors $\{\text{SW}, \text{SE}, \text{NE}, \text{NW}\}$ are used. The tiles can directly be read on the figure.

Proof of lemma 1

Proof. Let us first state one classical fact about the Thue-Morse word: \mathcal{T} does not contain any triple factor www where $w \in \{0, 1\}^*$. A proof (and more about the Thue-Morse substitution) can for instance be found in [11] (chapter 5). Also observe that if 00 or 11 appears, then the bar is necessarily in the middle (0|0 or 1|1) as they do not belong to $s(\{0, 1\}) = \{01, 10\}$ (and this observation is not ambiguous as the triple factors 000 and 111 do not appear). Among all 16 words of $\{0, 1\}^4$, only 10 do not contain a triple factor (and they all actually appear in the Thue-Morse word). All of them contain 00 or 11 (thus are unambiguous) except 0101 and 1010. As they are symmetric, we only consider 0101. If it were odd, i.e. 0|10|1, then we would necessarily have 10|10|10 considering the previous and the next letter of the factor in \mathcal{T} , which is contradictory as this gives




Fig. 10. Deterministic marking of a parabola

birth to a 10 triple factor. Hence it is even. For the sake of completeness, the following 4 factors are even $\{0101, 1001, 1010, 0110\}$, while the other 6 are odd $\{1011, 0010, 1100, 0011, 0100, 1101\}$. The result holds for factors of size greater than 4 as they contain a factor of size 4. ■

More detailed explanations for H' and V' signals The case of H' and V' signals requires some careful attention. Indeed we need to use knight colors for them to be deterministic in the NW and SE directions. The knight colors used basically contain the type of signal and the word it carries. However there is a constraint on the possible positions for a knight color: a knight must absolutely not cross a front line as in this case the determinism in one of the orthogonal NE or SW directions would fail. Indeed, seeing time going NE, this could happen in the two following situations:

1. The knight links a tile from an incoming H' or V' signal to a tile that is on the NE side of the front line. Then the factor carried by this knight would not be predictable in the SW direction as the four derived factor would not all be “readable” by the deterministic rule.
2. The knight links an outgoing H' or V' tile to a tile that is on the SW side of the front line. Then this tile would not be predictable in the NE direction.

Using the quite natural (e_1, e_1, e_1, e_2) profile (3 horizontal steps for each vertical one) for the slope $\frac{1}{3}$ of H' signals would make us fall into one of these situations. The adopted solution consists in using the $(e_1, e_1, e_2, e_1, e_1, e_1, e_1, e_2)$ profile (alternating between 2 and 4 horizontal steps). That way, there are exactly three tiles in the signal that do not hold a knight color: the tile that is in the middle of the 3 tiles corresponding to the 2 horizontal steps and the two tiles surrounding the middle of the 5 tiles corresponding to the 4 horizontal steps. We use the symmetrical profile for V' signals. Both cases are illustrated on the figures 6(a) and 6(b). Moreover, by choosing the right starting tile for the emission of a signal on a pillar/middle, one can make sure that when two H' and V' signals meet to produce a middle, they had always just used that knight-less tile, i.e. the western and southern neighbors of a middle are always the two knight-less tiles of the H' and V' signal respectively. Said otherwise, for H' signals (the V' case is symmetric), the marked points always appear at the position of the middle tile in the 5 tiles corresponding to the 4 horizontal steps, hence they are directly surrounded in the signal by tiles that hold no knight color. We prevent that way the first failure as illustrated by the figure 6(d): the knights from the incoming signals do not cross the front line. The choice done also prevents the second failure as illustrated by the figures 6(c) and 6(d): the knight from the outgoing signals do not cross the front line.

Complete description of the tiling The tiling κ contains 1151 tiles that can be classified into the 41 categories of figure 11. The colors are denoted by lines drawn on the tiles that must continue uninterrupted across tile edges, except for the H' and V' signals where, due to their particular profiles, some discontinuities

appear along edges marked by the special colors $\#_1$ and $\#_2$. The H and V signals can carry any of the 4 even subwords $\{0101, 1001, 1010, 0110\}$. The H' and V' signals can carry any of the 6 odd subwords $\{1011, 0010, 1100, 0011, 0100, 1101\}$. The knight colors only appear on tiles containing a H' or V' signal and carry the same odd subword as the associated H' or V' signal. On tiles containing several H' and V' signals, the direction of each knight color easily permits to identify unambiguously the associated H' or V' signal. Pillar and middle tiles were already comprehensively described in the article (the two incoming signals carry the same subword, the four outgoing words are derived from the substitution of this subword).

Only the tiles of the SW half plane delimited by the fire line are represented on the figure 11. The NE tiles appearing on the other side of the fire line are obtained from the SW tiles by applying the transformation represented on the figure 12 (color swap and symmetry along the NW–SW diagonal of the tiles), where, for each SW color c , we denote by c' its duplicate for the NE tileset. The same notation is used on figure 11 for the fire line tiles, which must themselves respect this symmetry as they make the junction between both symmetric tilesets. The transformation is applied to all the tiles to obtain their symmetrical counterpart, except for the fire line tiles (that are fixed points for this transformation).

A large version of the figure 5 is provided on figure 13.

Proof of theorem 2

Proof sketch. The first assertion is a straightforward syntactic verification on tiles. Note that κ is even deterministic at radius 1 (standard Wang tiles determinism) in directions NE and SW (as the knight colors for these directions are not used here).

On the other hand clear, κ tiles the plane (blank tiling) and there exists $\pi : \kappa \rightarrow \{0, 1\}$ (naturally extended to colorings $\pi : \kappa^{\mathbb{Z}^2} \rightarrow \{0, 1\}^{\mathbb{Z}^2}$) such that $\pi(\mathcal{X}_\kappa) = \mathcal{L}_{-1}$, where \mathcal{L}_{-1} denote the subshift generated by the $\{0, 1\}$ -coloring c such that, for all $(x, y) \in \mathbb{Z}^2$, $c(x, y) = 1$ if $x + y = 0$ and $c(x, y) = 0$ otherwise. Indeed, picking an arbitrarily long factor of the Thue-Morse sequence and writing it with fire tiles on the diagonal of an empty square, as on the figure 5, one easily sees that the local NE and SW determinisms can be applied to turn the whole square into a valid square pattern. Hence there exist arbitrarily large valid square patterns marking their finite diagonal. For compactness reasons, there must exist a tiling of the whole plane containing a marked bi-infinite diagonal. This implies in turn the left-right inclusion by translation invariance. To prove the converse, observe that a tiling containing a fire tile must contain a full fire line and the whole associated structure, hence all “degenerated” tilings are projected to the blank coloring. ■

Finally, observe that the set of tilings by κ is at least expansive in directions $]-\infty, 0[\cup]\frac{1}{2}, 2[$.

Blank tile:




H and V signals and collision:




H' signal:




V' signal:




V/H' and H/V' crossings:



Pillar, signal emission tiles and middle:



Particular collisions and crossings near the fire line:



Fire:




Fig. 11. 41 types of tiles in κ




Fig. 12. Symmetric transformation on tiles

