Groups Geom. Dyn. 16 (2022), 1461–1514
DOI 10.4171/GGD/692

Simulations and the lamplighter group

Laurent Bartholdi and Ville Salo

Abstract. We introduce a notion of "simulation" for labelled graphs, in which edges of the simulated graph are realized by regular expressions in the simulating graph, and we prove that the tiling problem (a.k.a. the "domino problem") for the simulating graph is at least as difficult as that for the simulated graph. We apply this to the Cayley graph of the "lamplighter group" $L = \mathbb{Z}/2 \wr \mathbb{Z}$, and more generally to "Diestel–Leader graphs". We prove that these graphs simulate the plane, and thus deduce that the seeded tiling problem is unsolvable on the group L. We note that L does not contain any plane in its Cayley graph, so our undecidability criterion by simulation covers cases not addressed by Jeandel's criterion based on translation-like action of a product of finitely generated infinite groups. Our approach to tiling problems is strongly based on categorical constructions in graph theory.

1. Introduction

Let G be a finitely generated group, with finite generating set $S = S^{-1}$. For a given finite set A and a subset $\Pi \subseteq A \times S \times A$ called *tileset*, a *tiling* of G is the choice, for every $g \in G$, of a colour $\phi_g \in A$ such that neighbouring group elements have matching colours: for every $g \in G$, $s \in S$ we have $(\phi_g, s, \phi_{gs}) \in \Pi$.

The *tiling problem* for G asks for an algorithm that, given Π, determines whether there exists such a tiling. A useful variant, the *seeded tiling problem*, asks for an algorithm that, given Π and a seed colour $a_0 \in A$, determines whether there exists a tiling with colour a_0 at the origin.

These problems have attracted much attention since Wang's [21] and Berger's [6] results that the seeded tiling problem, respectively tiling problem are unsolvable for $G = \mathbb{Z}^2$.

More generally, groups that contain sufficiently regular planes also have unsolvable tiling problems; see the next section for more precise statements. It is also easy to see that groups with unsolvable word problem have unsolvable tiling problem. To the best of our knowledge, all groups known up to now to have unsolvable tiling problem fall in one of these two classes.

2020 Mathematics Subject Classification. Primary 03D35; Secondary 05C25, 20E22, 20F65, 20F10, 37B10, 37B50, 68Q10.
Keywords. Lamplighter group, domino problem, tiling problem, decision problems in group theory, wreath products, Diestel–Leader graphs.

On the other hand, straightforward considerations show that both tiling problems are solvable for free groups, and more generally virtually free groups (namely groups with a free subgroup of finite index).

The main contribution of this article is a proof of undecidability for a group not covered by these classes, the *lamplighter group*. This is the wreath product $L = \mathbb{Z}/2 \wr \mathbb{Z}$, and also the group generated by the affine transformations $a(f) = tf$ and $b(f) = tf + 1$ of the ring $\mathbb{F}_2[t, t^{-1}]$; it admits as presentation

$$L = \langle a, b \mid (a^n b^{-n})^2 \text{ for all } n \geq 1 \rangle.$$

This group is of course finitely generated, is not virtually free, does not contain geometrically any plane or hyperbolic plane, and has decidable word problem. We prove:

Theorem A. *The seeded tiling problem for the lamplighter group L is unsolvable.*

We view the tiling problem as a question about marked graphs; in our setting, the Cayley graph of the group L, namely the graph with vertex set L and an edge between g and gs for all $s \in \{a^{\pm 1}, b^{\pm 1}\}$.

We prove Theorem A by introducing a notion of "simulation": even though $\mathbb{Z}^2$ is not contained geometrically in L, it is contained "automatically", in that there is a function $\mathbb{Z}^2 \to L$ which is not Lipschitz but maps edges of the Cayley graph of $\mathbb{Z}^2$ to regular expressions in edges of the Cayley graph of L; these regular expressions are driven by an auxiliary labelling provided to the Cayley graph of L by a subshift of finite type. This is sufficient to reduce the seeded tiling problem of $\mathbb{Z}^2$ to that of L, and therefore conclude with the latter's undecidability.

Our construction is fundamentally graph-theoretical. It applies in particular, with trivial modifications, to all Diestel–Leader graphs $\mathtt{DL}(p, q)$; see Section 7. Up to subtleties regarding edge labellings, $\mathtt{DL}(2, 2)$ is just the Cayley graph of L. We prove in fact:

Theorem B. *On every Diestel–Leader graph $\mathtt{DL}(p, q)$ the seeded tiling problem is unsolvable.*

We leave the following conjecture as an interesting open problem.

Conjecture 1.1. *The tiling problems of the lamplighter group and Diestel–Leader graphs are undecidable.*

Added in print: this will be answered positively in a forthcoming paper of ours.

1.1. Groups with unsolvable (seeded) tiling problem

As we mentioned in the introduction, the tiling problem on G is solvable if G is virtually free [18]. In fact, if a tileset Π tiles at all, then it admits a "rational" tiling, namely there are regular languages $(R_a \subseteq S^*)_{a \in A}$ such that $g \in G$ is coloured a if g is in the image of R_a under the natural evaluation map $S^* \twoheadrightarrow G$. A tiling algorithm therefore lists in parallel tilings of larger and larger balls in G, and tests A-tuples of regular languages in

S^* for valid tilings, and we are guaranteed that one of the processes will stop. See [1, Theorem 9.3.37]. No other groups are known to have solvable tiling problem:

Conjecture 1.2 (Ballier–Stein [4]). *A finitely generated group has solvable domino problem if and only if it is virtually free.*

There has been, since 2013 (when [4] appeared as a preprint), continuous progress towards Conjecture 1.2; here is a brief list of elementary results, in which all groups are assumed to be finitely generated.[1]

- If G has unsolvable word problem (namely there is no algorithm determining, with input a word w in the generators of G, whether $w = 1$ holds in G), then there is *a fortiori* no algorithm determining whether a tileset tiles; see [1, Theorem 9.3.28].
- The fundamental results of Wang and Berger show that $\mathbb{Z}^2$ has unsolvable tiling problem.
- If G, H are commensurable (meaning they have finite-index isomorphic subgroups), then G has solvable tiling problem if and only if H does.
- If H is a finitely generated subgroup of G with unsolvable tiling problem, then the tiling problem of G is also unsolvable; see [1, Proposition 9.3.30].
- If N is a normal finitely generated subgroup of G, and G/N has unsolvable tiling problem, then so does G; see [1, Proposition 9.3.32].

Recall that two finitely generated groups G, H are *quasi-isometric* if there are Lipschitz maps $G \to H$ and $H \to G$ whose compositions are at bounded sup-distance from the identity. Cohen proves that the tiling problem is geometric in the following sense:

Theorem 1.3 ([7]). *If G, H are finitely presented and quasi-isometric, then G has solvable tiling problem if and only if H does.*

Recall next that a finitely generated group H acts *translation-like* on G if the action is free and by self-maps of G at bounded sup-distance from the identity. This notion was introduced by Whyte [22]; Seward [20] proved that a finitely generated group is infinite if and only if it admits a translation-like action of $\mathbb{Z}$.

Theorem 1.4 ([13, Theorem 3]). *If H is finitely presented, has unsolvable tiling problem, and acts translation-like on G, then G has unsolvable tiling problem.*

Thus every group containing a subgroup of the form $H_1 \times H_2$ with H_1, H_2 infinite and finitely generated has unsolvable tiling problem. This applies in particular to "branched groups" such as the first Grigorchuk group. The current "state of the art" also includes the following groups as having unsolvable tiling problem:

[1]The case of infinitely generated G to some extent reduces to the finitely generated case, in the sense that given a tileset Π, a tiling of G exists if and only if a tiling exists on the subgroup generated by elements mentioned in the description of Π.

- Non-virtually-cyclic virtually nilpotent [4] and more generally virtually polycyclic groups [12]; see [1, Theorem 9.3.43].
- Baumslag–Solitar groups [3]; see [1, Theorem 9.3.47].
- Fundamental groups of closed surfaces [2].

Note that all the above examples geometrically contain a plane, either Euclidean or hyperbolic, in their Cayley graph. In fact, Ballier and Stein prove more generally that if G is not virtually cyclic but contains an infinite cyclic central subgroup, then its tiling problem is unsolvable.

We have not yet addressed an important, related question: given a group G, does there exist a finite tileset that admits tilings, but only aperiodic ones? There are several variants to this question, in particular a *weakly aperiodic* tileset is such that all tilings have infinite orbits (equivalently: infinite-index stabilizer) under translation by G, a *strongly aperiodic* tileset is such that all tilings have trivial stabilizer.

These notions are equivalent for $\mathbb{Z}^2$, but differ in general: On $\mathbb{Z}^3$, one can build weakly-but-not-strongly aperiodic tilesets easily from strongly aperiodic tilesets of $\mathbb{Z}^2$, they exist on all Baumslag–Solitar groups which are not virtually abelian [3, 11], and Cohen constructs one for the lamplighter group in [8]. On $\mathbb{Z}^3$ and (at least) amenable Baumslag–Solitar groups, also strongly aperiodic tilesets exist, while on L none are known.

1.2. The general tiling problem of the lamplighter group

It remains open whether the (unseeded) tiling problem of the lamplighter group is solvable. All the standard methods used to prove the undecidability of the tiling problem on the plane and other finitely generated groups seem plausible on the lamplighter group: the lamplighter group is residually finite and one can imagine tiling it by "macro-tiles" that form another copy of the lamplighter (or related) group, and one could try a Robinson-like construction or the fixed-point methods for this. However, the group is too disconnected to make a direct implementation possible. To implement this method, we suspect the obstacle to overcome is to find a strongly aperiodic tileset.

Another approach is the transducer method, which applies to groups equipped with a homomorphism onto $\mathbb{Z}$ ("indicable groups"). On each fibre, we write numbers from a finite set, their average density represents a real number, and a transducer verifies that consecutive fibres have a density related by a piecewise-affine map without periodic points. This strategy succeeds on Baumslag–Solitar groups [3]; on the lamplighter group, it cannot work directly because the fibres are disconnected as graphs (no matter what the generating set is). It would be interesting to understand which $\mathbb{Z}$-subshifts have a sofic pullback to the lamplighter group along the natural homomorphism. We show in Proposition 6.7 that the sunny-side-up on $\mathbb{Z}$ has a sofic lift, while for the transducer method it seems one would need to lift a full shift (though it is not clear this is sufficient).

1.3. Simulations

We consider the (seeded) tiling problem in the context of edge- and vertex-labelled graphs $\mathcal{G}$; when applied to a group $G = \langle S \rangle$, the graph we consider is the Cayley graph of G, which has one vertex per group element and an edge labelled s from g to gs for all $g \in G$ and all $s \in S$. In the case of a seeded tiling problem, the vertex $1 \in G$ furthermore has a special marking. A tileset is then a collection of allowable colourings of vertices and their abutting edges and neighbours.

Our strategy to prove that a graph $\mathcal{G}$ has unsolvable tiling problem can be summarized as follows: define a preorder relation "simulate" on a class of graphs containing $\mathcal{G}$, and then

(1) show that if a graph $\mathcal{A}$ simulates $\mathcal{B}$, and $\mathcal{B}$ has unsolvable tiling problem, then so does $\mathcal{A}$;

(2) choose a graph $\mathcal{H}$ with unsolvable tiling problem;

(3) show that $\mathcal{G}$ simulates $\mathcal{H}$.

This is uninteresting if "simulate" is the equality relation; the coarser "simulate" is, the harder (1) becomes and the easier (3) becomes. Taking "simulate" to mean "admits a translation-like action by" interprets Theorem 1.4 in this context.

We propose a more general notion of "simulation", in which the edges of the simulated graph are not paths of bounded length in the simulating graph (as would be the case in a translation-like action) but rather are given by regular expressions, or more appropriately by graph-walking finite state automata. These automata simulate edges by following edge and vertex labellings on the simulating graph, which typically is drawn by a legal tiling over an auxiliary fixed tile set. An appropriate image, back to the context of groups, is that when G simulates H, there is for every generator t of H an ant that reacts to pheromones deposited on the Cayley graph of G; when started at a vertex representing an element $h \in H$, it will move according to its finite state table and pheromone inputs until it reaches a stop state, which will then represent $ht \in H$.

The "automaton" or "ant" picture is useful to describe actual simulations, but is better expressed in the more formal setting of a "bigraph", namely a graph equipped with two graph morphisms towards finite graphs. These should be thought of as generalizations of graph morphisms (in case one of the morphisms is the identity).

We realize our plan, towards the proof of Theorem A, as follows:

(1) is proven as Theorem 3.4 and Theorem 3.7;

(2) we use the quadrant $\mathbb{N}^2$ or the plane $\mathbb{Z}^2$;

(3) is proven in two different manners: using a "comb", in Proposition 5.2, or a "sea level", in Proposition 6.10.

Both of the simulations require a marking of L by tiles from given "pheromone" tileset (whose size we could bring down to 6 tiles in the case of the "comb"). We have proven, therefore, that the Cayley graph of L, when decorated by this tileset, has unsolvable tiling problem. However, this immediately implies that L itself also has unsolvable seeded tiling problem, since the "pheromone" tileset can be combined to the instance to be solved. We need a seed tile so as to ensure that there is a non-trivial decoration from the "pheromone" tileset; indeed we have not been able to find a tile set on L *all* of whose configurations simulate a quadrant or $\mathbb{Z}^2$.

Summary of the proof of Theorem A. Let Π be an instance of a tiling problem for the quadrant. The "sea level" is a tiling problem $\Pi_{\approx\approx}$ for L, which (after specifying a seed tile) has a unique solution. The grid may be simulated using the labels on L coming from this tiling, leading to a tiling problem Π' for L. If the seeded tiling problem of L were solvable, it would in particular be solvable for the product of $\Pi_{\approx\approx}$ and Π', and therefore lead to a solution of the tiling problem Π on the quadrant, a contradiction.

We could, alternatively, use the "comb" tiling problem instead of the "sea level". Then the induced labelling on L is not unique, and forms a family of tilings, some of which simulate grids. The proof would then proceed identically. ∎

The reason we exhibit two different constructions is that each has its own additional benefits. The comb gives a better proof in terms of conciseness: the proof is shorter, and indeed Section 5 (skipping Proposition 5.2 and Corollary 5.3) gives a self-contained proof of the main theorem using the comb subshift of finite type (SFT).

On the other hand, the sea level SFT admits a unique seeded configuration which simplifies the application of our simulation theorem to it, and is more efficient (an $n \times n$ grid requires roughly $\log(n)n^2$ vertices in the lamplighter graph). In the process, we shall learn more about the "SFT-accessible geometry" of the lamplighter group: recalling that the lamplighter group is a split extension $\mathbb{F}_2[a, a^{-1}] \rtimes \langle a \rangle$, we show that the one-point compactified actions on the coset spaces of the subgroups $\mathbb{F}_2[a, a^{-1}]$ and $\langle a \rangle$ are sofic shifts, and the sea level SFT is obtained by combining their covering SFTs with a simple rule. The sea level construction also does not require the rigidity of a vertex-transitive graph, and with minor modifications applies to the Diestel–Leader graphs $\mathtt{DL}(p, q)$. We feel these ideas are more likely to be useful in further work on the group than the ad hoc tricks used for the comb.

2. Graphs and their simulations

We begin by some general notions on graphs. They rely heavily on categorical ideas, but we have taken care to explain them in direct terms, only hinting at their abstract origin. The reader who is uncomfortable with categorical language is strongly encouraged to ignore every sentence that matches '`.*[Cc]ategor[iy].*`'. These areas are marked by ⚠.

2.1. Graphs

Definition 2.1. A *graph* is a set $\mathcal{G} = V \sqcup E$ partitioned into *vertices* and *edges*, with two maps $e \mapsto e^{\pm}: E \to V$; we call e^+ the *head* of e and call e^- its *tail*. We write $V(\mathcal{G})$ and $E(\mathcal{G})$ for the vertices and edges of a graph $\mathcal{G}$.

Graphs are born oriented. An *unoriented* graph is a graph endowed with an involution $e \mapsto e': E \to E$ on its edges called the *reversal*, satisfying $(e')^+ = e^-$ for all $e \in E$.

The *closure* of an edge $e \in E$ is the subgraph $\overline{e} = \{e, e^+, e^-\}$ in the oriented case and $\overline{e} = \{e, e', e^+, e^-\}$ in the unoriented case, and the closure of a vertex $v \in V$ is $\overline{v} = \{v\}$.

A *morphism of graphs* is a pair of maps between their vertices and edges respectively, which intertwine the head, tail and (in the unoriented case) reversal operations. We use the same symbol for the map on vertices and on edges, and in formulas the conditions for a map $\phi: \mathcal{G} \to \mathcal{H}$ to be a morphism are $\phi(e^{\pm}) = \phi(e)^{\pm}$, $\phi(e') = \phi(e)'$.

A morphism $\phi: \mathcal{G} \to \mathcal{H}$ is *étale* if it is locally injective: if two edges $e, f \in E(\mathcal{G})$ satisfy $\phi(e) = \phi(f)$ and ($e^- = f^-$ or $e^+ = f^+$), then $e = f$. A graph is *weakly étale* if all edges which have the same label and share an endpoint actually share both endpoints.

⚠ Graphs, with graph morphisms, form a category Graph. In the next sections, we shall display important properties of this category. We write $\mathsf{Graph}(\mathcal{G}, \mathcal{H})$ for the set of morphisms $\mathcal{G} \to \mathcal{H}$, or just $\mathrm{Hom}(\mathcal{G}, \mathcal{H})$ if we do not want to precise the subcategory of Graph under consideration.

2.2. Labelled graphs

Consider a graph $\mathcal{G}$, with labels on its vertices and edges. We thus have maps $\alpha: V(\mathcal{G}) \to A$ and $\alpha: E(\mathcal{G}) \to C$ to label sets A and C. Without loss of generality, we may assume that every label $c \in C$ "knows" the label of its extremities, so we have maps $c \mapsto c^{\pm}: C \to A$. Indeed at worst replace C by $A \times C \times A$ and extend the edge labelling to a map $E(\mathcal{G}) \to A \times C \times A$ by $e \mapsto (\alpha(e^-), \alpha(e), \alpha(e^+))$.

Thus a labelling of $\mathcal{G}$ is nothing but a graph morphism $\alpha: \mathcal{G} \to A \sqcup C$. From now on we shall write $\alpha: \mathcal{G} \to \mathcal{A}$ for a labelling on $\mathcal{G}$, with $\mathcal{A} = A \sqcup C$.

Again, in the unoriented case, we assume that there is an involution on C compatible with the edge involution on $\mathcal{G}$.

Example 2.2. Let G be a group with generating set S. An example we shall return to in much more detail is the *Cayley graph* of G: it is the graph $\mathcal{C} = G \sqcup (G \times S)$ with $(g, s)^- = g$ and $(g, s)^+ = gs$. In case $S = S^{-1}$, we also have an edge reversal given by $(g, s)' = (gs, s^{-1})$.

The Cayley graph is naturally an edge labelled graph, in which the edge (g, s) has label s. We express this as a map $\mathcal{C}(G, S) \to 1 \sqcup S$, where $1 \sqcup S$ is the graph with one vertex and S loops at it.

Example 2.3. A bipartite graph is a graph given with a partition of its vertices into two parts $V_0 \sqcup V_1$, and such that all edges cross between V_0 and V_1.

A bipartite graph is then naturally a graph equipped with a labelling to the segment, namely to the graph with vertex set $\{0, 1\}$ and two edges e, e' with $e^+ = (e')^- = 1$ and $e^- = (e')^+ = 0$.

Example 2.4. Let G be a monoid with generating set S. The element $1 \in G$ may be distinguished, leading to a *rooted Cayley graph*, or *sunny-side-up*; this last terminology because one vertex is marked as vitellus (yolk) and all the others as albumen (white). Thus the rooted Cayley graph is expressed by a map

$$\mathcal{C}(G, S) \to \{\mathsf{vitellus}, \mathsf{albumen}\} \sqcup (\{\mathsf{vitellus}, \mathsf{albumen}\} \times S \times \{\mathsf{vitellus}, \mathsf{albumen}\}),$$

with $(x,s,y)^- = x$ and $(x,s,y)^+ = y$; namely the map defined on vertices by $1 \mapsto \mathsf{vitellus}$ and $g \neq 1 \mapsto \mathsf{albumen}$, and likewise on edges.

⚠ In categorical language, the $\mathcal{A}$-labelled graphs, namely the graphs with labelling by a graph $\mathcal{A}$, form the *slice* category $\mathsf{Graph}_{/\mathcal{A}}$; this is the category whose objects are morphisms $\alpha\colon \mathcal{G} \to \mathcal{A}$ and whose morphisms between $\alpha\colon \mathcal{G} \to \mathcal{A}$ and $\beta\colon \mathcal{H} \to \mathcal{A}$ are usual graph morphisms $\gamma\colon \mathcal{G} \to \mathcal{H}$ satisfying $\alpha = \beta \circ \gamma$.

The *simplification* of a labelled graph $\mathcal{G}$ is the graph where all edges e with the same endpoints and label are merged into a single edge, and the *full simplification* is obtained from the simplification by also removing all self-loops. Note that a graph is weakly étale if and only if its simplification is étale.

2.3. Pullbacks and exponentials

Consider two $\mathcal{A}$-labelled graphs $\mathcal{G}_1$ and $\mathcal{G}_2$, namely graphs equipped with labellings $\alpha_i\colon \mathcal{G}_i \to \mathcal{A}$. Their *pullback* is the graph

$$\mathcal{G}_1 \times_{\mathcal{A}} \mathcal{G}_2 = \{(u_1, u_2) \in \mathcal{G}_1 \times \mathcal{G}_2 : \alpha_1(u_1) = \alpha_2(u_2)\},$$

labelled by $\alpha(u_1, u_2) = \alpha_1(u_1)$. The graph structure is given by $(e_1, e_2)^\pm = (e_1^\pm, e_2^\pm)$, and similarly for the reversal.

⚠ Category theory points us to natural constructions, such as the above, that pervade throughout mathematics; and has a way of expressing their important properties. This will be even more so for the pullbacks and exponentials in the next section.

In categorical terms, $\mathcal{G}_1 \times_{\mathcal{A}} \mathcal{G}_2$ is *universal*: it is the "best" way to construct an $\mathcal{A}$-labelled graph equipped with morphisms to $\mathcal{G}_1$ and $\mathcal{G}_2$ (which of course are given by projections to the first and second factor). The universal property says that, for every graph $\mathcal{H}$ also equipped with morphisms $\phi_i\colon \mathcal{H} \to \mathcal{G}_i$, there is a unique map $\eta\colon \mathcal{H} \to \mathcal{G}_1 \times_{\mathcal{A}} \mathcal{G}_2$ with $\phi_i = \pi_i \circ \eta$.

Example 2.5. Consider $\mathcal{G}_2 = \mathcal{A} \sqcup \mathcal{A}$, namely two disjoint copies of the vertices and edges of $\mathcal{A}$. Then $\mathcal{G}_1 \times_{\mathcal{A}} \mathcal{G}_2 \cong \mathcal{G}_1 \sqcup \mathcal{G}_1$.

Consider next $\mathcal{G}_2$ a subgraph of $\mathcal{A}$, for example the subgraph spanned by a subset A' of the vertices of $\mathcal{A}$. Then $\mathcal{G}_1 \times_{\mathcal{A}} \mathcal{G}_2$ is the subgraph of $\mathcal{G}_1$ whose labelling belongs to $\mathcal{G}_2$, for example the subgraph of $\mathcal{G}_1$ spanned by vertices whose label is in A'.

Example 2.6. Let G_1, G_2 be two groups with respective generating sets S_1, S_2. Then the Cayley graph $\mathcal{C}(G_1 \times G_2, S_1 \times S_2)$ is the usual direct product of the graphs $\mathcal{C}(G_1, S_1)$ and $\mathcal{C}(G_2, S_2)$; it is a connected graph precisely when $S_1 \times S_2$ generates $G_1 \times G_2$. This direct product is the pullback $\mathcal{C}(G_1, S_1) \times_{1 \sqcup 1} \mathcal{C}(G_2, S_2)$ over the trivial graph with one vertex and one edge. More generally, let there be maps $f_1 \colon S_1 \to S$ and $f_2 \colon S_2 \to S$; then the Cayley graph of $\mathcal{C}(G_1 \times G_2, S_1 \times_S S_2)$ is $\mathcal{C}(G_1, S_1) \times_{1 \sqcup S} \mathcal{C}(G_2, S_2)$, naturally viewing $\mathcal{C}(G_i, S_i) \in \mathsf{Graph}_{/1 \sqcup S}$ via f_i.

Consider two $\mathcal{B}$-labelled graphs $\mathcal{G}_1, \mathcal{G}_2$, and let $\alpha \colon \mathcal{G}_2 \to \mathcal{A}$ be an $\mathcal{A}$-labelling on $\mathcal{G}_2$. The *exponential*[2] of these graphs is the graph

$$\mathcal{G}_1^{\mathcal{G}_2} = \big\{(f, a) : a \in \mathcal{A},\ f \in \mathsf{Graph}_{/\mathcal{B}}(\alpha^{-1}(\overline{a}), \mathcal{G}_1)\big\}.$$

There is a natural $\mathcal{A}$-labelling on $\mathcal{G}_1^{\mathcal{G}_2}$ by projection to the second coordinate. The graph structure on $\mathcal{G}_1^{\mathcal{G}_2}$ is given by $(f, a)^{\pm} = (f \upharpoonright \alpha^{-1}(a^{\pm}), a^{\pm})$, with $\upharpoonright$ denoting restriction. In the unoriented case, we have $(f, a)' = (f', a')$ with $f'(e) = f(e')$ and $f'(e^{\pm}) = f(e^{\mp})$.

In words, an edge in $\mathcal{G}_1^{\mathcal{G}_2}$ with label $c \in E(\mathcal{A})$ is a function defined on the α-preimages of c and its endpoints, with values in $\mathcal{G}_1$, which preserves the graph structure and the $\mathcal{B}$-labelling.

Example 2.7. Consider the graph $\mathcal{B}$ with vertex set $\{0, 1\}$ and no edge. Thus $\mathsf{Graph}_{/\mathcal{B}}$ is the category of sets S partitioned into two parts $S_0 \sqcup S_1$. Choose $\mathcal{A} = \mathcal{B}$ and let α coincide with the $\mathcal{B}$-labelling. Unwrapping the definition, we see $(S_0 \sqcup S_1)^{T_0 \sqcup T_1} = S_0^{T_0} \sqcup S_1^{T_1}$, where in the last expression $S_i^{T_i}$ is as usual the set of maps $T_i \to S_i$.

⚠ The classical bijection $A^{B \times C} = (A^C)^B$ is a particular case of *adjunction*: there are *functors* on the category of sets, respectively the cartesian product $A \mapsto \mathfrak{F}(A) = A \times C$ and the exponential $A \mapsto \mathfrak{G}(A) = \{f \colon C \to A\}$; and a natural bijection $\mathsf{Set}(\mathfrak{F}(B), A) \leftrightarrow \mathsf{Set}(B, \mathfrak{G}(A))$. One says that $\mathfrak{F}$ and $\mathfrak{G}$ are adjoint. It is a valuable reflex, when thinking categorically, to recognize important functors and enquire whether they have adjoints, and if so what they are. This instinct led us in the right direction in defining $\mathcal{G}_1^{\mathcal{G}_2}$.

Lemma 2.8. *For any labelled graphs $\mathcal{G}_1 \to \mathcal{A}$ and $\mathcal{A} \leftarrow \mathcal{G}_2 \to \mathcal{B}$ and $\mathcal{G}_3 \to \mathcal{B}$ we have a natural bijection*

$$\mathsf{Graph}_{/\mathcal{B}}(\mathcal{G}_1 \times_{\mathcal{A}} \mathcal{G}_2, \mathcal{G}_3) \cong \mathsf{Graph}_{/\mathcal{A}}(\mathcal{G}_1, \mathcal{G}_3^{\mathcal{G}_2}).$$

The proof is essentially "currying": a morphism $\phi \colon \mathcal{G}_1 \times_{\mathcal{A}} \mathcal{G}_2 \to \mathcal{G}_3$ gives rise to a family of morphisms $\phi_x \colon \mathcal{G}_2 \to \mathcal{G}_3$, indexed by $x \in \mathcal{G}_1$, by the formula $\phi_x(y) = \phi(x, y)$; and conversely.

[2]This is not actually an exponential object in a category $\mathsf{Graph}_{/\mathcal{B}}$ because of the different categories $\mathsf{Graph}_{/\mathcal{A}}$, $\mathsf{Graph}_{/\mathcal{B}}$ involved. It does however produce exponential objects, in the sense of adjoints to fibre products, when $\mathcal{A} = \mathcal{B}$ and the $\mathcal{A}$- and $\mathcal{B}$-labellings coincide; see Lemma 2.8.

Proof with the painful details. Write the labellings $\alpha_i : \mathcal{G}_i \to \mathcal{A}$ for $i = 1, 2$ and $\beta_i : \mathcal{G}_i \to \mathcal{B}$ for $i = 2, 3$. We give maps in both directions, and then show that they are inverses of each other.

($\exists \to$) Let $\lambda : \mathcal{G}_1 \times_{\mathcal{A}} \mathcal{G}_2 \to \mathcal{G}_3$ be a graph morphism. Define the map $\rho : \mathcal{G}_1 \to \mathcal{G}_3^{\mathcal{G}_2}$ by $\rho(u_1) = (f_{u_1}, \alpha_1(u_1))$ with $f_{u_1}(u_2) = \lambda(u_1, u_2)$ for all $u_2 \in \alpha_2^{-1}\alpha_1(\overline{u_1})$. Note that $f_{u_1} \in \mathsf{Graph}_{/\mathcal{B}}$ since it is a restriction of $\lambda \in \mathsf{Graph}_{/\mathcal{B}}$. Clearly $\rho(u_1)^{\pm} = \rho(u_1^{\pm})$ and ρ preserves the $\mathcal{A}$-labelling, so $\rho \in \mathsf{Graph}_{/\mathcal{A}}$.

($\exists \leftarrow$) Let $\rho : \mathcal{G}_1 \to \mathcal{G}_3^{\mathcal{G}_2}$ be a graph morphism. Define the map $\lambda : \mathcal{G}_1 \times_{\mathcal{A}} \mathcal{G}_2 \to \mathcal{G}_3$ by $\lambda(u_1, u_2) = f_{u_1}(u_2)$ where $\rho(u_1) = (f_{u_1}, \alpha_1(u_1))$. Note, since $f_{u_1} \in \mathsf{Graph}_{/\mathcal{B}}$, that $\lambda(u_1, u_2)^{\pm} = \lambda(u_1^{\pm}, u_2^{\pm})$ and that λ preserves the $\mathcal{B}$-labelling; so $\lambda \in \mathsf{Graph}_{/\mathcal{B}}$.

($\leftarrow \circ \to$) We next show that the constructions are inverses of each other. Let $\widehat{\rho} \in \mathsf{Graph}_{/\mathcal{A}}(\mathcal{G}_1, \mathcal{G}_3^{\mathcal{G}_2})$ be constructed from $\lambda \in \mathsf{Graph}_{/\mathcal{B}}(\mathcal{G}_1 \times_{\mathcal{A}} \mathcal{G}_2, \mathcal{G}_3)$, which in turn is constructed from $\rho \in \mathsf{Graph}_{/\mathcal{A}}(\mathcal{G}_1, \mathcal{G}_3^{\mathcal{G}_2})$. We show $\widehat{\rho} = \rho$: for $u_1 \in \mathcal{G}_1$ we have

$$\widehat{\rho}(u_1) = (\widehat{f_{u_1}}, \alpha_1(u_1)) \quad \text{and} \quad \rho(u_1) = (f_{u_1}, \alpha_1(u_1)),$$

with $\widehat{f_{u_1}}(u_2) = \lambda(u_1, u_2) = f_{u_1}(u_2)$ directly from the defining formulas, so $\widehat{f_{u_1}} = f_{u_1}$.

($\to \circ \leftarrow$) Conversely, let $\widehat{\lambda} \in \mathsf{Graph}_{/\mathcal{A}}(\mathcal{G}_1 \times_{\mathcal{A}} \mathcal{G}_2, \mathcal{G}_3)$ be the graph constructed from $\rho \in \mathsf{Graph}_{/\mathcal{B}}(\mathcal{G}_1, \mathcal{G}_3^{\mathcal{G}_2})$, which in turn is constructed from $\lambda \in \mathsf{Graph}_{/\mathcal{A}}(\mathcal{G}_1 \times_{\mathcal{A}} \mathcal{G}_2, \mathcal{G}_3)$. Writing $\rho(u_1) = (f_{u_1}, \alpha_1(u_1))$, we have

$$\widehat{\lambda}(u_1, u_2) = f_{u_1}(u_2) = \lambda(u_1, u_2). \qquad \blacksquare$$

Example 2.9. Continuing on Example 2.7, with $\mathcal{G}_1 = S_0 \sqcup S_1$ and $\mathcal{G}_2 = T_0 \sqcup T_1$ and $\mathcal{G}_3 = U_0 \sqcup U_1$, we get

$$\begin{aligned} \operatorname{Hom}(\mathcal{G}_1 \times_{\mathcal{A}} \mathcal{G}_2, \mathcal{G}_3) &= \operatorname{Hom}(S_0 \times T_0 \sqcup S_1 \times T_1, U_0 \sqcup U_1) \\ &= U_0^{S_0 \times T_0} \times U_1^{S_1 \times T_1} \\ &= (U_0^{T_0})^{S_0} \times (U_1^{T_1})^{S_1} \\ &= \operatorname{Hom}(S_0 \sqcup S_1, U_0^{T_0} \sqcup U_1^{T_1}) \\ &= \operatorname{Hom}(\mathcal{G}_1, \mathcal{G}_3^{\mathcal{G}_2}), \end{aligned}$$

in accordance with the previous lemma.

2.4. Path subdivisions

We introduce a geometric operation, that of subdividing edges by replacing them by unbounded paths:

Definition 2.10. Let $\mathcal{B}$ be a graph. Its *path subdivision* is the graph $\mathcal{B}^*$ with vertex set $\mathcal{B} = V(\mathcal{B}) \sqcup E(\mathcal{B})$ and edge set $\{0, 1\} \times E(\mathcal{B}) \times \{0, 1\}$, with extremities

$$\begin{aligned} &(0, e, 0)^{\pm} = e^{\pm}, \qquad &&(0, e, 1)^{-} = e^{-}, \qquad &&(0, e, 1)^{+} = e, \\ &(1, e, 1)^{\pm} = e, \qquad &&(1, e, 0)^{-} = e, \qquad &&(1, e, 0)^{+} = e^{+}. \end{aligned}$$

In the unoriented case the vertex set is $\mathcal{B}/\{e'{=}e \ \forall e \in E(\mathcal{B})\}$, and $(i, e, j)' = (j, e', i)$.

In pictures, we are replacing every edge e, from v to w, by a small graph:

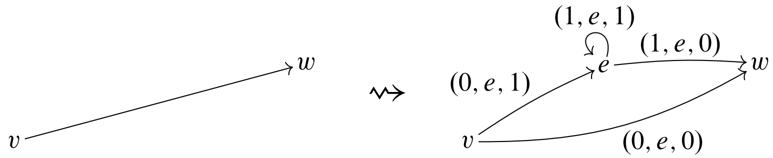

(2.1)

There is a natural map from paths in $\mathcal{B}^*$ to paths in $\mathcal{B}$; namely, replace each path going through an $E(\mathcal{B})$-vertex by the straight edge going directly from v to w in (2.1). There is a natural map between the geometric realization of $\mathcal{B}^*$ to that of $\mathcal{B}$, and a natural embedding of $\mathcal{B}$ in $\mathcal{B}^*$ by mapping every edge e to $(0, e, 0)$. There is no natural graph morphism $\mathcal{B}^* \to \mathcal{B}$. The operation is evidently functorial, i.e. there is a natural way to associate to a graph homomorphism (or labelled graph) $\alpha\colon \mathcal{G} \to \mathcal{B}$ a graph homomorphism (or labelled graph) $\beta\colon \mathcal{G}^* \to \mathcal{B}^*$.

Consider a labelled graph $\alpha\colon \mathcal{G} \to \mathcal{B}^*$. We construct a labelled graph $\beta\colon \mathcal{G}^\flat \to \mathcal{B}$ as follows: its vertex set is $\alpha^{-1}(V(\mathcal{B}))$, with $u \in V(\mathcal{G}^\flat)$ labelled $\alpha(u)$. For each path $(e_1, \dots, e_n)$ in $\mathcal{G}$, say from u to v, such that $\alpha(u), \alpha(v) \in V(\mathcal{B})$ and the labels along $(e_1, \dots, e_n)$ are $(0, c, 1), (1, c, 1), \dots, (1, c, 1), (1, c, 0)$ for some $c \in E(\mathcal{B})$, there is an edge from u to v labelled c. We call paths $(e_1, \dots, e_n)$ as above *coherent paths*. The terminology "$\flat$" is justified by deFLATion of paths into edges.

It may sometimes be necessary to add some extra "sink" vertices to $\mathcal{G}^*$: anticipating Section 2.5, the new vertices in $\mathcal{G}^*$ represent intermediate steps in a calculation; and we want to allow the possibility of calculations that start but do not terminate – no results are produced, but computational resources are consumed. For a labelled graph $\alpha\colon \mathcal{G} \to \mathcal{B}$, define the labelled graph $\beta\colon \mathcal{G}^\sharp \to \mathcal{B}^*$ as follows: its vertex set is

$$V(\mathcal{G}^\sharp) = V(\mathcal{G}^*) \sqcup \{\pm 1\} \times E(\mathcal{B}).$$

Its edge set is the union of $E(\mathcal{G}^*)$ and, for each edge e labelled $c = \alpha(e)$ of $\mathcal{G}$, an edge labelled $(0, c, 1)$ from $e^- \in V(\mathcal{G})$ to $(-, c)$, an edge labelled $(1, c, 0)$ from $(+, c)$ to e^+, and edges labelled $(1, c, 1)$ from $(+, c)$ to $(+, c)$, $(-, c)$ and e, and edges from e and $(-, c)$ to $(-, c)$. In pictures, where all omitted labels are $(1, c, 1)$:

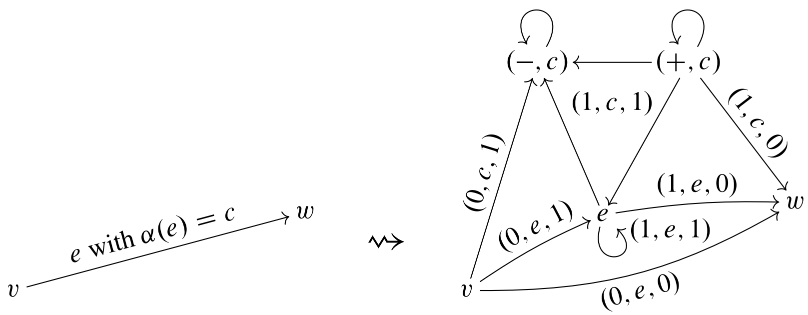

The operations $\mathcal{G} \mapsto \mathcal{G}^\flat$ and $\mathcal{G} \mapsto \mathcal{G}^\sharp$ are close to being adjoint to each other:

Lemma 2.11. *Consider* $\mathcal{G} \in \mathsf{Graph}_{/\mathcal{B}^*}$ *and* $\mathcal{H} \in \mathsf{Graph}_{\mathcal{B}}$. *Then there is a map*

$$\mathrm{Hom}(\mathcal{G}, \mathcal{H}^\sharp) \to \mathrm{Hom}(\mathcal{G}^\flat, \mathcal{H}).$$

If furthermore $\mathcal{G}^\flat$ *is weakly étale, then there is also a map*

$$\mathrm{Hom}(\mathcal{G}, \mathcal{H}^\sharp) \leftarrow \mathrm{Hom}(\mathcal{G}^\flat, \mathcal{H}),$$

so $\mathrm{Hom}(\mathcal{G}^\flat, \mathcal{H})$ *is empty if and only if* $\mathrm{Hom}(\mathcal{G}, \mathcal{H}^\sharp)$ *is empty.*

Proof. Consider first $f: \mathcal{G} \to \mathcal{H}^\sharp$, and define $f^\flat: \mathcal{G}^\flat \to \mathcal{H}$ as follows. Vertices of $\mathcal{G}^\flat$ are in particular vertices of $\mathcal{G}$, so we define $f^\flat$ on $V(\mathcal{G}^\flat)$ as the restriction of f. For an edge $(e_1, \dots, e_n)$ of $\mathcal{G}^\flat$, the path $(f(e_1), \dots, f(e_n))$ is of the form $((\mathbf{1}_{i\neq 1}, e, \mathbf{1}_{i\neq n}))_{1\le i\le n}$, and we set $f^\flat(e_1, \dots, e_n) = e$. Note that in fact $f^\flat$ only depends on the restriction of the image of f to $\mathcal{H}^*$, since there is no coherent path in $\mathcal{H}^\flat$ going through the vertices $\{\pm 1\} \times E(\mathcal{B})$ and starting and ending at vertices of $\mathcal{H}$.

Consider next $g: \mathcal{G}^\flat \to \mathcal{H}$. As a preprocessing step, we may assume without loss of generality that g has the property that for all $p, q \in E(\mathcal{G}^\flat)$ with $p^- = q^-$ and $p^+ = q^+$ and $\alpha(p) = \alpha(q)$ we have $g(p) = g(q)$: simply choose one preferred path $p_{u,c,v}$ between any two $u, v \in V(\mathcal{G}^\flat)$ and for each $c \in E(\mathcal{B})$ with $\alpha(u) = c^-$, $\alpha(v) = c^+$, and redefine $g(q) = g(p_{q^-,\alpha(q),q^+})$ for all $q \in E(\mathcal{G}^\flat)$.

Now define $g^\sharp: \mathcal{G} \to \mathcal{H}^\sharp$ as follows. On vertices of $\mathcal{G}$ with label in $V(\mathcal{B})$, called "vertex vertices", let $g^\sharp$ coincide with g. Consider now a vertex p of $\mathcal{G}$ with label in $V(\mathcal{B}^*) \setminus V(\mathcal{B})$, an "edge vertex". If p lies on a coherent path $(e_1, \dots, e_n)$ in $\mathcal{G}$, set $g^\sharp(p) = g(e_1, \dots, e_n)$ for an arbitrary such path. Since $\mathcal{G}^\flat$ is weakly étale, all coherent paths have the same end points and α-label $\alpha(p)$, and by our preprocessing of g all such paths have the same image, so the definition of $g^\sharp(p)$ is unambiguous.

Consider next an edge vertex p that does not lie on a coherent path. If it admits a path labelled $(1, c, 1), \dots, (1, c, 1), (1, c, 0)$ towards a vertex vertex, map it by $g^\sharp$ to $(+, c)$; if on the other hand there is path from a vertex vertex to p labelled $(0, c, 1), (1, c, 1), \dots, (1, c, 1)$, then map it by $g^\sharp$ to $(-, c)$; if neither path exists, map it (by will) to $(+, \alpha(p))$.

There is a unique way of extending $g^\sharp$ to edges: edges in $\mathcal{G}$ between vertex vertices, namely with label $(0, c, 0)$, are edges $e \in E(\mathcal{G}^\flat)$, so one may set $g^\sharp(e) = (0, g(e), 0)$. Edges along coherent paths are mapped by $g^\sharp$ to the unique possible edges $(0, e, 1), (1, e, 1), \dots, (1, e, 1), (1, e, 0)$ in $\mathcal{H}^* \subset \mathcal{H}^\sharp$. Finally, if an edge of $\mathcal{G}$ starts or ends at a vertex that we already decided to map to $w \in \{+, -\} \times E(\mathcal{B}) \subseteq V(\mathcal{H}^\sharp)$, then map this edge by $g^\sharp$ to a loop at w, except if it respectively ends or starts at a vertex vertex or an edge vertex on a coherent path. ∎

Note that if $\mathcal{G}^\flat$ is actually étale, meaning that in $\mathcal{G}$ there is a unique coherent path with a given label between any two vertices, then the preprocessing step is not needed.

Note also that, in the proof above, the compositions of $f \mapsto f^\flat$ and $g \mapsto g^\sharp$ are not quite the identity: $(g^\sharp)^\flat$ is the identity on vertices, but when $\mathcal{G}^\flat$ is not étale, the composition $\flat \circ \sharp$ identifies images of edges corresponding to coherent paths with the same endpoints and labels. On the other hand, $(f^\flat)^\sharp$ projects components without initial or final vertex to “sinks”. Nevertheless, when restricting to an appropriate subclass of graphs in which all edge vertices lie on coherent paths, one obtains an adjunction between $\mathcal{G} \mapsto \mathcal{G}^\flat$ and $\mathcal{H} \mapsto \mathcal{H}^\sharp$.

2.5. Simulations

We now introduce a more general notion of morphism, in the same way that bimodules generalise algebra morphisms:

Definition 2.12. Let $\mathcal{A}, \mathcal{B}$ be two graphs, which we think of as labellings. A *simulator* is a graph $\mathcal{S}$ equipped with two labellings $\alpha\colon \mathcal{S} \to \mathcal{A}$ and $\beta\colon \mathcal{S} \to \mathcal{B}^*$. We call $\mathcal{S}$, more precisely, an $(\mathcal{A}, \mathcal{B})$-*simulator*.

Let $\mathcal{G}$ be an $\mathcal{A}$-labelled graph. Its *image* under the simulation $\mathcal{S}$ is the graph obtained from $\mathcal{G} \times_{\mathcal{A}} \mathcal{S}$ by replacing each subgraph as on the right of (2.1) by the corresponding edge; namely, the graph $\mathcal{G} \times_{\mathcal{A}} \mathcal{S}$ is $\mathcal{B}^*$-labelled via the labelling β, and we define $\mathcal{G} \rtimes \mathcal{S}$ as the $\mathcal{B}$-labelled graph $(\mathcal{G} \times_{\mathcal{A}} \mathcal{S})^\flat$. We say that $\mathcal{G}$ *simulates* $\mathcal{H}$ if $\mathcal{H} \cong \mathcal{G} \rtimes \mathcal{S}$ for some *finite* simulator $\mathcal{S}$. We say $\mathcal{G}$ *simulates* $\mathcal{H}$ *up to simplification* if $\mathcal{G}$ simulates a graph $\mathcal{K}$ whose simplification is $\mathcal{H}$.

Example 2.13. The *plane* is the Cayley graph of $\mathbb{Z}^2 = \langle \to, \leftarrow, \uparrow, \downarrow \rangle$, and will still be written $\mathbb{Z}^2$. The *quadrant* is the full subgraph of $\mathbb{Z}^2$ with vertex set $\mathbb{N}^2$; its vertices have degree 4, 3 or 2 depending on how many of their coordinates are 0.

The quadrant is naturally vertex-labelled, according to the available directions: $+$, $\vdash$, $\perp$, $\llcorner$; thus the origin is labelled $\llcorner$ in the quadrant, and the line $\{x = 0\}$ is labelled $\vdash$. (For exhaustivity, the plane is also vertex-labelled, with $+$ everywhere). The quadrant's labelling is a bit finer than the sunny-side-up labelling from Example 2.4 in that it remembers the axes too.

The plane and quadrant are $\mathcal{A}$-labelled graphs, for $\mathcal{A}$ the graph with one vertex (written $\cdot$) and four edges $\to, \leftarrow, \uparrow, \downarrow$.

Example 2.14. We claim that the quadrant simulates the plane. For this, it suffices to exhibit a finite simulator, for example the one shown in Figure 1.

In this graph, edges are labelled as '$a|b$' to represent their labels under the left and right map, respectively, towards $\mathcal{A}$. We have only drawn edges corresponding to directions $\uparrow, \to$ for the map to $\mathcal{A}$; the other edges are obtained by flipping the arrows. Since all edges under the right map are of the form $(0, e, 0)$, we have simply written e instead; and since there are no vertex labels in $\mathbb{Z}^2$, we have only indicated the vertex labels from $\mathbb{N}^2$ next to the nodes.

In particular, the root $(0, 0)$ in the quadrant (mapped to the central node of the simulator) corresponds to the root $(0, 0)$ in the plane; the axes in the quadrant are each covered by two half-axes; and the remainder of the quadrant is covered by the four quadrants of the plane. Thus the plane is simulated, within the quadrant, by folding it like a handkerchief.

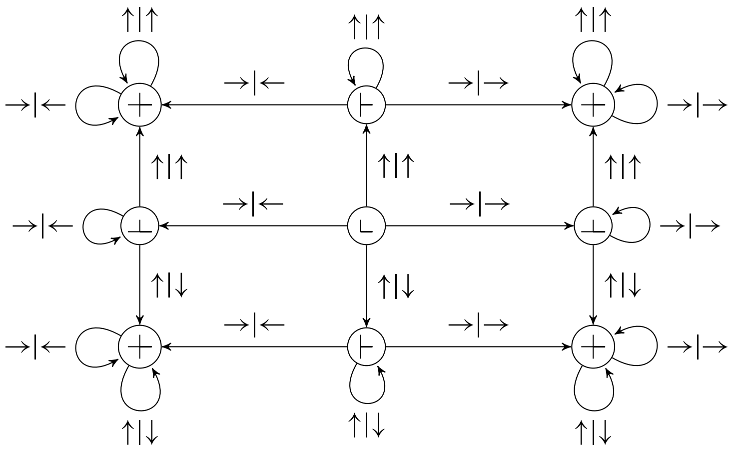

Figure 1. A simulator for the quadrant. The edge label (a, b) indicates that a move by a on the quadrant simulates a move by b on the simulated plane.

Example 2.15. The previous example does not use paths of length > 1 under the right map. It is, in a sense, possible to simulate the plane within the quadrant while never putting two simulated vertices at the same place, at the cost of simulating edges in the plane by paths of length 2 in the quadrant.

For simplicity, we will simulate $\mathbb{Z}$ in $\mathbb{N}$; naturally simulators for $\mathbb{Z}^2$ in $\mathbb{N}^2$ may be obtained by taking products. The simulation from Example 2.14, when restricted to $\mathbb{Z}$, was (now depicted folded, with v. and a. for vitellus and albumen)

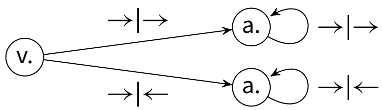

The idea of the following simulator is that $n \in \mathbb{Z}$ is simulated by $2n \in \mathbb{N}$ and $-n \in \mathbb{Z}$ by $2n - 1 \in \mathbb{N}$, for all $n \geq 0$. This is impossible to achieve exactly with our definition, so we produce $\mathbb{Z} \sqcup \mathbb{N} \sqcup -\mathbb{N}$ instead. The vertices of the simulator are now of two kinds: split nodes indicate an actual vertex of the simulated graph (their label in the simulated graph being furthermore indicated by the lower symbol – in this case, $-, \cdot, +$ giving the sign[3] in $\mathbb{Z}$), and unsplit nodes indicate intermediate nodes from $\mathcal{B}^* \setminus \mathcal{B}$:

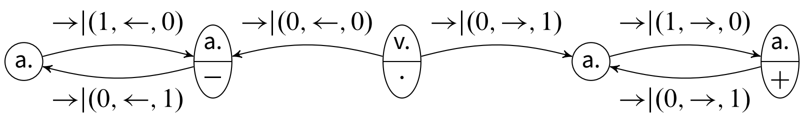

[3]The Cayley graph of $\mathbb{Z}$ does not have these markings, but having them makes the automaton easier to read, and identifying $-$ with $+$ yields a valid simulation of seeded-$\mathbb{Z}$.

Indeed following $(\rightarrow)^{2n}$ in $\mathbb{N}$, starting from the origin, is the same thing as following $((0,\rightarrow,1)(1,\rightarrow,0))^n$ in the pullback, which amounts to following $(\rightarrow)^n$ in $\mathbb{Z}$, and following $(\rightarrow)^{2n+1}$ in $\mathbb{N}$, starting from the origin, is the same thing as following $(0,\leftarrow,0)((0,\leftarrow,1)(1,\leftarrow,0))^n$ in the pullback, which amounts to following $(\leftarrow)^{n+1}$ in $\mathbb{Z}$.

The fibre product $\mathbb{N} \rtimes \mathcal{S}$ also contains vertices

$$\left(2n, \frac{\text{v.}}{-}\right) \quad \text{and} \quad \left(2n+1, \frac{\text{v.}}{+}\right)$$

producing a copy of $-\mathbb{N}$ and a copy of $\mathbb{N}$. If we add vertex labels for odd and even positions in $\mathbb{N}$, then it is possible to eliminate these additional copies of $\mathbb{N}$. Such a labelling can be introduced by an SFT marking (which is always a possibility in our applications, as our simulators run on graphs decorated by SFT configurations). See also Section 6.5, where we give another interpretation.

We adapt the definition of exponential objects for simulators. Let $\mathcal{H}$ be a $\mathcal{B}$-labelled graph, labelled by $\eta\colon \mathcal{H} \to \mathcal{B}$. We define $\mathcal{H}^{\mathcal{S}}$ as the $\mathcal{A}$-labelled graph $(\mathcal{H}^\sharp)^{\mathcal{S}}$, with the morphisms in the exponential belonging to $\mathsf{Graph}_{/\mathcal{B}^*}$.

Lemma 2.16. *Let $\mathcal{G}$ be an $\mathcal{A}$-labelled graph, let $\mathcal{H}$ be a $\mathcal{B}$-labelled graph, and let $\mathcal{S}$ be an $(\mathcal{A},\mathcal{B})$-simulator. Then*

$$\mathsf{Graph}_{/\mathcal{B}}(\mathcal{G} \rtimes \mathcal{S}, \mathcal{H}) \neq \emptyset \quad \textit{whenever} \quad \mathsf{Graph}_{/\mathcal{A}}(\mathcal{G}, \mathcal{H}^{\mathcal{S}}) \neq \emptyset.$$

If additionally $\mathcal{G} \rtimes \mathcal{S}$ is weakly étale, then

$$\mathsf{Graph}_{/\mathcal{B}}(\mathcal{G} \rtimes \mathcal{S}, \mathcal{H}) \neq \emptyset \quad \textit{if and only if} \quad \mathsf{Graph}_{/\mathcal{A}}(\mathcal{G}, \mathcal{H}^{\mathcal{S}}) \neq \emptyset.$$

Proof. This follows immediately from Lemma 2.8 and Lemma 2.11. For the second claim,

$$\mathsf{Graph}_{/\mathcal{B}}(\mathcal{G} \rtimes \mathcal{S}, \mathcal{H}) = \mathsf{Graph}_{/\mathcal{B}}((\mathcal{G} \times_{\mathcal{A}} \mathcal{S})^\flat, \mathcal{H})$$

is empty if and only if

$$\mathsf{Graph}_{/\mathcal{B}^*}(\mathcal{G} \times_{\mathcal{A}} \mathcal{S}, \mathcal{H}^\sharp) \cong \mathsf{Graph}_{/\mathcal{A}}(\mathcal{G}, \mathcal{H}^{\mathcal{S}})$$

is empty. ∎

As could be expected, simulation is transitive:

Lemma 2.17. *If $\mathcal{G}$ simulates a graph $\mathcal{H}$, and $\mathcal{H}$ simulates a graph $\mathcal{K}$, then $\mathcal{G}$ simulates $\mathcal{K}$.*

Proof. The idea of the proof is straightforward: run the simulation of $\mathcal{K}$ as a subroutine within the simulation of $\mathcal{H}$. Here are the details.

Let $\mathcal{S}$ be an $(\mathcal{A},\mathcal{B})$-simulator expressing the simulation $\mathcal{H} \cong \mathcal{G} \rtimes \mathcal{S}$, and let $\mathcal{T}$ be a $(\mathcal{B},\mathcal{C})$-simulator expressing the simulation $\mathcal{H} \rtimes \mathcal{T} \cong \mathcal{K}$. Then $\mathcal{T}^\sharp$ is a $\mathcal{B}^*$- and $\mathcal{C}^{**}$-labelled graph; its labels in $\mathcal{C}^{**}$ take the form $(i,(i',c,j'),j)$. Only for this proof, let $\natural$ be the operation that replaces every such label by $(\max(i,i'), c, \max(j,j'))$, changing the $\mathcal{C}^{**}$-labelling into a $\mathcal{C}^*$-labelling. Set $\mathcal{U} = (\mathcal{S} \times_{\mathcal{B}^*} \mathcal{T}^\sharp)^\natural$, and note that it is an $(\mathcal{A},\mathcal{C})$-simulator. We have, noting that the operations $()^\sharp$ and $()^\natural$ commute with products,

$$\begin{aligned}\mathcal{G} \rtimes \mathcal{U} &= (\mathcal{G} \times_{\mathcal{A}} \mathcal{U})^\flat = (\mathcal{G} \times_{\mathcal{A}} (\mathcal{S} \times_{\mathcal{B}^*} \mathcal{T}^\sharp)^\natural)^\flat \\ &\cong (\mathcal{G} \times_{\mathcal{A}} \mathcal{S} \times_{\mathcal{B}^*} \mathcal{T}^\sharp)^{\natural\flat} \cong ((\mathcal{G} \times_{\mathcal{A}} \mathcal{S})^{\flat\sharp} \times_{\mathcal{B}^*} \mathcal{T}^\sharp)^{\natural\flat} \\ &= (\mathcal{H}^\sharp \times_{\mathcal{B}^*} \mathcal{T}^\sharp)^{\natural\flat} \cong (\mathcal{H} \times_{\mathcal{B}} \mathcal{T})^{\sharp\natural\flat} \\ &\cong (\mathcal{H} \times_{\mathcal{B}} \mathcal{T})^\flat = \mathcal{H} \rtimes \mathcal{T} \cong \mathcal{K}.\end{aligned}$$
■

Note that the $\flat$-operation removes the sinks introduced by $T \mapsto T^\sharp$, so without changing the previous proof, one can shave off a small number of vertices by using $\mathcal{S} \times \mathcal{T}^*$ as the simulator.

Robinson's tileset [19] may be used to prove the undecidability of the tiling problem for $\mathbb{Z}^2$. In Example 3.10 in the next section we show that one can interpret this in terms of the simulation notion, essentially as showing that we can simulate seeded $\mathbb{Z}^2$-tilings by unseeded ones.

There are some subtleties to this, which are explained in the next section, but the main component is the observation that if a set-theoretic rectangle is explicitly marked inside an actual rectangle, then, using simulation, we can interpret the set-theoretic rectangle as an actual rectangle, by deflating the paths consisting of unmarked nodes. We give the details in the following example.

Example 2.18. Let $S = \{(1,0),(0,1),(-1,0),(0,-1)\}$ be the standard generating set for $\mathbb{Z}^2$ and let $\mathcal{B} = 1 \sqcup S$ be the graph used for the Cayley graph labelling. Consider the induced subgraph R with vertex set $\{0,1,\dots,n-1\}^2 \subset \mathbb{Z}^2$. Consider a labelled graph $\alpha\colon R \to \mathcal{A}$, for the product graph

$$\mathcal{A} = (1 \sqcup S) \times \big(\{\text{good},\text{bad}\} \sqcup \{\text{good},\text{bad}\}^2\big).$$

Suppose that the label in $(1 \sqcup S)$ is just the induced Cayley graph labelling from $\mathbb{Z}^2$, and suppose that the subset $S' \subset V(S)$ of good vertices is a set-theoretic rectangle:

$$\exists A, B \subset \{0,1,\dots,n-1\}\colon\ \alpha^{-1}(\{(1,\text{good})\}) = A \times B.$$

Then consider the following $(\mathcal{A},\mathcal{B})$-simulator on S. The vertex labelled "good" has $\mathcal{A}$-label $(1,\text{good})$ and $\mathcal{B}^*$-label $1 \in V(\mathcal{B})$, while the split node labelled "bad/s" has $\mathcal{A}$-label $(1,\text{bad})$ and $\mathcal{B}^*$-label $s \in S = E(\mathcal{B})$; s ranges over S, so the figure is short for a 5-vertex graph.

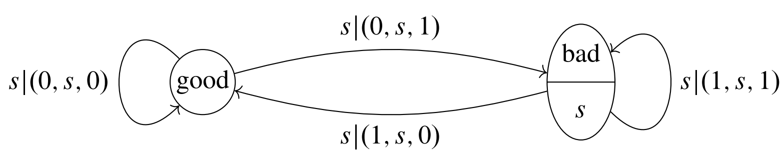

In words, this simulator compresses runs over bad nodes into an edge, and thus compresses the set-theoretic rectangle $A \times B$ into an actual rectangle. The $\mathcal{B}$-labelling is the one induced from $\{0, 1, \dots, |A| - 1\} \times \{0, 1, \dots, |B| - 1\} \subset \mathbb{Z}^2$. In terms of the graph-walking automata introduced in the following section, the automaton "skips over" the bad nodes.

2.6. Graph-walking automata

Simulators can be seen as a compact way of expressing two things at once: vertex duplication (when we want to simulate multiple nodes in one actual node), and coalescence of paths into single edges based on the labels of connecting paths.

In this section, we re-express simulators explicitly in terms of vertex blow-ups, graph-walking automata, and formal languages. We show that these points of view are equivalent to the simulators introduced in the previous section; while simulators tend to be convenient in proofs, graph-walking automata are preferable for describing concrete simulations.

Definition 2.19 (Vertex blow-up). Let $\mathcal{A}$ be a graph, and let $k\colon v(\mathcal{A}) \to \mathbb{N}$ be a function. The *vertex blow-up of* $\mathcal{A}$ by k is the graph with vertices and edges

$$V(\mathcal{A}^k) = \big\{(u, i) \in V(\mathcal{A}) \times \mathbb{N} : i < k(u)\big\},$$
$$E(\mathcal{A}^k) = \big\{(i, e, j) \in \mathbb{N} \times E(\mathcal{A}) \times \mathbb{N} : i < k(e^-), j < k(e^+)\big\}$$

with $(i, e, j)^- = (e^-, i)$ and $(i, e, j)^+ = (e^+, j)$.

We extend this to labelled graphs by blowing-up the graph used for the labels:

Definition 2.20. Let $\alpha\colon \mathcal{G} \to \mathcal{A}$ be an $\mathcal{A}$-labelled graph, and let $k\colon V(\mathcal{A}) \to \mathbb{N}$ be a function. The *vertex blow-up of* $\mathcal{G}$ by k is defined as the graph with vertices and edges

$$V(\mathcal{G}^k) = \big\{(u, i) \in V(\mathcal{G}) \times \mathbb{N} : i < k(\alpha(u))\big\},$$
$$E(\mathcal{G}^k) = \big\{(i, e, j) \in \mathbb{N} \times E(\mathcal{G}) \times \mathbb{N} : i < k(e^-), j < k(e^+)\big\}$$

with $(i, e, j)^- = (e^-, j)$ and $(i, e, j)^+ = (e^+, j)$ and labels in $\mathcal{A}^k$ defined by $\alpha((u, i)) = (\alpha(u), i)$.

It is sometimes convenient to blow up vertices and give them more memorable names than numbers, and use a function $k\colon V(\mathcal{A}) \to \{\text{finite subsets of } U\}$ instead, for some universe U. To interpret this, fix bijections between $k(v)$ and $\{0, \dots, |k(v)| - 1\}$ for all $v \in V(\mathcal{A})$ and conjugate all arguments through this bijection.

Definition 2.21 (Graph-walking automata). Let $\alpha\colon \mathcal{G} \to \mathcal{A}$ be an $\mathcal{A}$-labelled graph. A *graph-walking automaton* (short: GWA) on $\mathcal{G}$ is a tuple $M = (Q, S, F, \Delta)$ with Q a finite set of *states*, $S \subset Q$ a set of *initial states*, $F \subset Q$ a set of *final states*, and $\Delta \subset Q \times E(\mathcal{A}) \times Q$ a *transition relation*. We require $I \cap F = \emptyset$. For $u, v \in V(\mathcal{G})$, an *M-run from u to v* is a sequence $(q_0, e_0, q_1, \dots, e_{k-1}, q_k)$ with, for all $i < k$,

$$q_0 \in I,\ q_k \in F,\ q_i \in Q,\ e_i \in E(\mathcal{G}),\ e_0^- = u,\ e_{k-1}^+ = v,\ (q_i, \alpha(e_i), q_{i+1}) \in \Delta.$$

For $u \in V(\mathcal{G})$, the *M-successors* of u are the nodes $v \in V(\mathcal{G})$ such that there exists an M-run from u to v.

The basic idea is that runs will be replaced by edges, just as paths are replaced by edges in the definition of simulators. Note that $I \cap F = \emptyset$ implies that all runs are of length at least one.

Definition 2.22 (GWA simulator). For an $\mathcal{A}$-labelled graph $\mathcal{G}$ and a finite graph $\mathcal{B}$, a *$\mathcal{B}$-labelled GWA simulator S* consists of the following data:

(1) a function $\lambda\colon V(\mathcal{A}) \to V(\mathcal{B}) \cup \{\perp\}$;

(2) for each $e \in E(\mathcal{B})$, a GWA M_e.

The GWA M_e is required to have the property that if $u \in V(\mathcal{G})$ and $\lambda(\alpha(u)) \neq e^-$, then u has no M_e-successors, and all M_e-successors v of u satisfy $\lambda(\alpha(v)) = e^+$. The GWA simulator S *simulates* the $\mathcal{B}$-labelled graph $\mathcal{H}$ with vertices and edges

$$V(\mathcal{H}) = \{v \in V(\mathcal{G}) : \lambda(v) \neq \perp\},$$
$$E(\mathcal{H}) = \{(u, e, v) \in V(\mathcal{G}) \times E(\mathcal{B}) \times V(\mathcal{G}) : \text{there is an } M_e\text{-run from } u \text{ to } v\}$$

with $(u, e, v)^- = u$, $(u, e, v)^+ = v$ and $\mathcal{B}$-labelling $\beta(v) = \lambda(\alpha(v))$, $\beta((u, e, v)) = e$.

Observe that the extra requirement on the GWA M_e is for convenience only: we can easily modify any GWA so that it satisfies this property.

Definition 2.23. Let $\mathcal{G}$ be an $\mathcal{A}$-labelled graph. We say $\mathcal{G}$ *GWA-simulates* a $\mathcal{B}$-labelled graph $\mathcal{H}$ if there exists a $\mathcal{B}$-labelled GWA simulator that simulates $\mathcal{H}$.

Definition 2.24 (Regular subdivision). Let $\mathcal{G}$ be an $\mathcal{A}$-labelled graph and let $\mathcal{B}$ be a finite graph. Suppose that we are given

(1) a function $\lambda\colon V(\mathcal{A}) \to V(\mathcal{B}) \cup \{\perp\}$, and

(2) for each $e \in E(\mathcal{B})$, a regular language $L_e \subset E(\mathcal{A})^+$ of non-trivial words.

Suppose further that every word $w = w_0 \dots w_{\ell-1} \in L_e$ is a path in $\mathcal{A}$ satisfying $\lambda(w_0^-) = e^-$ and $\lambda(w_{\ell-1}^+) = e^+$. Then the graph $\mathcal{H}$ with vertices and edges

$$V(\mathcal{H}) = \{v \in V(\mathcal{G}) : \lambda(v) \neq \perp\},$$
$$E(\mathcal{H}) = \{(u, e, v) \in V(\mathcal{G}) \times E(\mathcal{B}) \times V(\mathcal{G}) : \text{there is an } L_e\text{-labelled path from } u \text{ to } v\}$$

with $(u, e, v)^- = u$, $(u, e, v)^+ = v$ and $\mathcal{B}$-labelling $\beta(v) = \lambda(\alpha(v))$, $\beta((u, e, v)) = e$, is called a *regular $\mathcal{B}$-subdivision* of $\mathcal{G}$.

With these definitions in place, let us prove the equivalence of these notions:

Lemma 2.25. *Let $\mathcal{G}$ be an $\mathcal{A}$-labelled graph. Then the following are equivalent:*

(1) *$\mathcal{G}$ simulates $\mathcal{H}$ up to full simplification.*

(2) *Some vertex blow-up of $\mathcal{G}$ GWA-simulates $\mathcal{H}$.*

(3) *$\mathcal{H}$ is a regular $\mathcal{B}$-subdivision of a vertex blow-up of $\mathcal{G}$.*

Simplification is only a technicality: it is natural to require it in an automata-theoretic setting, by including only one edge if a run exists, while in the categorical setting of graphs we find it more natural to include an edge for each run. Simplification translates between these conventions. Full simplification (including self-loops) is mainly for notational convenience; self-loops have little effect on the tiling problems we are concerned with, and their presence simplifies some proofs.

Proof. (1) $\Rightarrow$ (2). Let $\mathcal{G}$ simulate $\mathcal{H}$ up to simplification by some simulator $\mathcal{S}$; we denote by $\alpha\colon \mathcal{G} \to \mathcal{A}$ and $\alpha_\mathcal{S}\colon \mathcal{S} \to \mathcal{A}$ and $\beta_\mathcal{S}\colon \mathcal{S} \to \mathcal{B}^*$ the respective labellings. We start by blowing up the vertices of $\mathcal{G}$ using the function $k\colon V(\mathcal{A}) \to \{\text{subsets of } V(\mathcal{S})\}$ defined by

$$k(u) = \{v \in V(\mathcal{S}) : \alpha_\mathcal{S}(v) = u\}.$$

Then the graph $\mathcal{G}^k$ has vertices (w, v) with $w \in V(\mathcal{G})$ and $v \in k(\alpha(w))$. For the function λ we pick $\lambda(v) = \beta_\mathcal{S}(v)$ whenever $\beta_\mathcal{S}(v) \in V(\mathcal{B})$, and $\lambda(v) = \perp$ whenever $v \in V(\mathcal{B}^*) \setminus V(\mathcal{B})$.

The vertices of both $\mathcal{G} \rtimes \mathcal{S}$ and $\mathcal{H}$ are by definition pairs (w, v) where $\alpha(w) = \alpha(v)$ and $\beta(v) \in \mathcal{B}$. No matter how we pick the automata M_e, the vertices of the GWA-simulated graph will be $\{(w, v) : \lambda(v) \neq \perp\}$. These sets are equal, so also in the simulated graph the vertex set will literally be $V(\mathcal{G} \rtimes \mathcal{S})$.

Now, for $e \in E(\mathcal{B})$, we pick the automata M_e. By the definition of simulation up to simplification, in $\mathcal{H}$ there is at most one edge with label e between $w \in V(\mathcal{H})$ and $w' \in V(\mathcal{H})$, and there is such an edge precisely when $w \neq w'$ and we are in one of the following cases:

(1) there is an edge $(e_\mathcal{G}, e_\mathcal{S}) \in \mathcal{G} \rtimes \mathcal{S}$ with $e_\mathcal{G}^- = w$, $e_\mathcal{G}^+ = w'$ and $\beta_\mathcal{S}(e_\mathcal{S}) = (0, e, 0)$, or

(2) there is a path of length at least two in $\mathcal{G} \rtimes \mathcal{S}$ from w to w', such that the following three conditions are satisfied:
   (a) the first edge $(e_\mathcal{G}, e_\mathcal{S})$ on the path satisfies $\beta_\mathcal{S}(e_\mathcal{S}) = (0, e, 1)$,
   (b) the last edge $(e_\mathcal{G}, e_\mathcal{S})$ on the path satisfies $\beta_\mathcal{S}(e_\mathcal{S}) = (1, e, 0)$,
   (c) all other edges $(e_\mathcal{G}, e_\mathcal{S})$ on the path satisfy $\beta_\mathcal{S}(e_\mathcal{S}) = (1, e, 1)$.

We are now ready to construct a GWA $M = M_e$ such that there is an M-run from w to w' if one of these cases occurs. Such a GWA is obtained from the $\mathcal{B}^*$-labelled graph $(\mathcal{S}, \beta_\mathcal{S})$ as follows:

(1) Make three disjoint copies of $V(\mathcal{S})$, say $S_i = V(\mathcal{S}) \times \{i\}$ for $i = 1, 2, 3$.
(2) In S_1 retain only vertices $(v, 1)$ with $\beta_\mathcal{S}(v) = e^-$, remove others.
(3) In S_2 retain only vertices $(v, 2)$ with $\beta_\mathcal{S}(v) = e$, remove others.
(4) In S_3 retain only vertices $(v, 3)$ with $\beta_\mathcal{S}(v) = e^+$, remove others.
(5) Set $Q = S_1 \cup S_2 \cup S_3$ as the set of states.

(6) Construct then $\Delta \subset Q \times E(\mathcal{A}) \times Q$ as follows: for each edge $e_{\mathcal{S}} \in E(\mathcal{S})$, say with $\beta_{\mathcal{S}}(e_{\mathcal{S}}) = (i, e, j)$, include $((e^-, 1+i), \alpha_{\mathcal{S}}(e_{\mathcal{S}}), (e^+, 3-j))$ in Δ.

(7) As initial states choose $I = S_1$, and as final states $F = S_3$.

Setting $M = (Q, I, F, \Delta)$, we see that M-runs are in bijective correspondence with the $\mathcal{G}$-paths and $\mathcal{G}$-edges defining edges of $\mathcal{G} \rtimes \mathcal{S}$ with label e, in particular there is an edge with label e from w to w' if and only if there is an M-run from w to w'. This concludes the proof that if $\mathcal{G}$ simulates $\mathcal{H}$, then a vertex blow-up of $\mathcal{G}$ GWA-simulates $\mathcal{H}$.

(2) ⇒ (1). Recall from Lemma 2.17 that simulation is transitive, so it suffices to show separately that if $\mathcal{H}$ is a vertex blow-up of $\mathcal{G}$, then $\mathcal{G}$ simulates $\mathcal{H}$, and that if $\mathcal{G}$ GWA-simulates $\mathcal{H}$, then it simulates $\mathcal{H}$.

Vertex blow-ups are in fact a special case of simulation: if $k: V(\mathcal{A}) \to \mathbb{N}$ is a function, let $\mathcal{S}$ be the graph $\mathcal{A}^k$; with $\mathcal{B} = \mathcal{A}^k$ and $\alpha: \mathcal{S} \to \mathcal{A}$ the canonical labelling of $\mathcal{A}^k$ and $\beta: \mathcal{S} \to \mathcal{B} \to \mathcal{B}^*$ the identity and natural inclusion, we see that $\mathcal{G} \rtimes \mathcal{S}$ is precisely the vertex blow-up of $\mathcal{G}$ by k.

We next show that if $\mathcal{G}$ GWA-simulates $\mathcal{H}$, then it simulates $\mathcal{H}$. For $e \in \mathcal{B}$, let $Q_e \supseteq I_e, F_e$ be the stateset, initial and final states of the automaton M_e, and assume all Q_e are disjoint. The nodes of the simulator $\mathcal{S}$ will be

$$V(\mathcal{S}) = V(\mathcal{A}) \sqcup \bigsqcup_{e \in \mathcal{B}} Q_e.$$

The $\mathcal{B}^*$-labelling on vertices is given by $\beta(u) = \lambda(u)$ for $u \in V(\mathcal{A})$, and $\beta(q) = e$ if $q \in Q_e$. In $E(\mathcal{S})$ we include an edge with $\mathcal{B}^*$-label $(0, e, 0)$ from e^- to e^+ if there is a run of length one in M_e, with a suitable α-label, namely if there exists $(s, \alpha(e'), t) \in \Delta$ with $s \in I_e$, $t \in F_e$, and $\beta(e') = e$.

For each transition (u, e', u') of M_e (so $u, u' \in Q_e$, $e' \in E(\mathcal{A})$), we include an edge in $\mathcal{S}$ with α-label e' and β-label $(1, e, 1)$ from u to u'. For each $u \in V(\mathcal{A})$, we include an edge with α-label e' and β-label $(0, e, 1)$ from u to $q \in Q_e$ whenever there is $(s, e', q) \in \Delta$ with $s \in I$. Symmetrically, for each $u \in V(\mathcal{A})$, we include an edge with α-label e' and β-label $(1, e, 0)$ from $q \in Q_e$ to u whenever there is $(q, e', s) \in \Delta$ with $s \in F$.

Then the simulated graph $\mathcal{G} \rtimes \mathcal{S}$ can be seen to be isomorphic to the GWA-simulated graph $\mathcal{H}$. We conclude that GWA-simulation and simulation are equivalent concepts.

(2) ⇔ (3). It is enough to show that $\mathcal{H}$ is a regular $\mathcal{B}$-subdivision of $\mathcal{G}$ if and only if $\mathcal{G}$ simulates $\mathcal{H}$. This is clear from the definition of a regular language, because the possible paths corresponding to valid M_e-runs for M_e satisfying $I \cap F = \emptyset$ form precisely a regular language of nonempty words. ∎

3. Subshifts

We are ready to define subshifts using the language of graphs introduced in the previous section.

3.1. Subshifts of finite type

Definition 3.1. Let $\mathcal{G}$ be a graph, possibly labelled. A *directed* Hom-*shift* (short: DHS) with carrier $\mathcal{G}$ is the space $\mathrm{Hom}(\mathcal{G}, \mathcal{F})$, for some finite graph $\mathcal{F}$.

If $\mathcal{G}$ is $\mathcal{A}$-labelled, then $\mathcal{F}$ should also be $\mathcal{A}$-labelled, and the space of homomorphisms should be taken in the appropriate category $\mathsf{Graph}_{/\mathcal{A}}$.

DHSs are nothing more than a formalism for subshifts of finite type, and as explained in Section 3.3 this definition is dynamically entirely equivalent to the more standard definition of a subshift of finite type by finitely many allowed (or forbidden) patterns.

The topology on $\mathrm{Hom}(\mathcal{G}, \mathcal{F})$ is the usual function topology; namely, $\mathrm{Hom}(\mathcal{G}, \mathcal{F})$ is a closed subset of $V(\mathcal{F})^{V(\mathcal{G})} \times E(\mathcal{F})^{E(\mathcal{G})}$, and therefore is compact.

Example 3.2. Consider $\mathcal{G} = \mathbb{N} \sqcup \mathbb{N}$, with $n^+ = n + 1$ and $n^- = n$. Geometrically, it is a one-sided ray. Consider the following finite graphs:

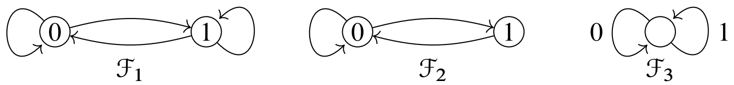

Then, in the standard symbolic dynamics terminology [16],

- $\mathrm{Hom}(\mathcal{G}, \mathcal{F}_1)$ is the full vertex shift $\{0, 1\}^{\mathbb{N}}$;
- $\mathrm{Hom}(\mathcal{G}, \mathcal{F}_2)$ is the "golden mean" shift $\{x \in \{0, 1\}^{\mathbb{N}} : \forall n \in \mathbb{N}\colon\ x_n x_{n+1} = 0\}$;
- $\mathrm{Hom}(\mathcal{G}, \mathcal{F}_3)$ is the full edge shift $\{0, 1\}^{\mathbb{N}}$.

3.2. The tiling problem

Definition 3.3. Let $\mathcal{G}$ be a graph, possibly labelled. The *tiling problem* for $\mathcal{G}$ is the decision problem of, given a finite graph $\mathcal{F}$, deciding whether $\mathrm{Hom}(\mathcal{G}, \mathcal{F})$ is non-empty.

Let now Γ be a family of graphs. The *tiling problem* for Γ is the problem of, given a finite graph $\mathcal{F}$, deciding whether $\mathrm{Hom}(\mathcal{G}, \mathcal{F})$ is non-empty for at least one $\mathcal{G} \in \Gamma$.

We call the tiling problem for $\mathcal{G}$ (respectively Γ) *solvable* if there exists an algorithm that truthfully answers its tiling problem. Next, we show that simulation can be used to transport tiling problems from one graph to another. In the terminology of recursion theory, this is a *many-one* reduction, meaning we algorithmically produce an instance of the tiling problem of one graph given an instance of that of the other (of course while preserving the answer).

Theorem 3.4. *Let $\mathcal{G}, \mathcal{H}$ be labelled graphs, and assume that $\mathcal{G}$ simulates $\mathcal{H}$, and $\mathcal{H}$ is weakly étale. Then the tiling problem of $\mathcal{H}$ many-one reduces to that of $\mathcal{G}$. In particular, if $\mathcal{H}$ has unsolvable tiling problem, then so does $\mathcal{G}$.*

Proof. Let $\mathcal{F}$ be an instance of the tiling problem for $\mathcal{H}$, namely a finite graph with same labelling as $\mathcal{H}$. Since $\mathcal{H}$ is simulated by $\mathcal{G}$, there exists a finite simulator $\mathcal{S}$ with $\mathcal{H} \cong \mathcal{G} \rtimes \mathcal{S}$.

We may algorithmically compute the finite graph $\mathcal{F}^{\mathcal{S}}$, and by hypothesis we may decide whether $\mathrm{Hom}(\mathcal{G}, \mathcal{F}^{\mathcal{S}})$ is non-empty. Now by Lemma 2.16 we have $\mathrm{Hom}(\mathcal{G}, \mathcal{F}^{\mathcal{S}}) \neq \emptyset$ if and only if $\mathrm{Hom}(\mathcal{H}, \mathcal{F}) \neq \emptyset$, so solving the tiling problem for $\mathcal{G}$ on $\mathcal{F}^{\mathcal{S}}$ solves at the same time the tiling problem for $\mathcal{H}$ on $\mathcal{F}$. ∎

This result extends readily to families of graphs; this is the most general result we obtain.

Definition 3.5. For two graphs $\mathcal{G}, \mathcal{H}$, we say that $\mathcal{G}$ *weakly maps* to $\mathcal{H}$ if every finite subgraph of $\mathcal{G}$ maps to $\mathcal{H}$; namely $\mathrm{Hom}(\mathcal{G}', \mathcal{H}) \neq \emptyset$ for all finite subgraphs $\mathcal{G}'$ of $\mathcal{G}$.

In case $\mathcal{H}$ is finite, this is equivalent, by compactness, to $\mathrm{Hom}(\mathcal{G}, \mathcal{H}) \neq \emptyset$, but in general it differs: for example if $\mathcal{G}$ is an infinite ray and $\mathcal{H}$ is a disjoint union of arbitrarily long finite rays.

Definition 3.6. Let Γ, Δ be families of graphs. We say that Γ *weakly simulates* Δ if there exists a finite simulator $\mathcal{S}$ such that

(1) every graph in Δ may be simulated: for every $\mathcal{H} \in \Delta$ there is $\mathcal{G} \in \Gamma$ such that $\mathcal{G} \rtimes \mathcal{S}$ is weakly étale and weakly maps to $\mathcal{H}$;
(2) simulated graphs are images of Δ: for every $\mathcal{G} \in \Gamma$ there is $\mathcal{H} \in \Delta$ that weakly maps to $\mathcal{G} \rtimes \mathcal{S}$.

Theorem 3.7. *Let Γ, Δ be families of graphs, and suppose Γ weakly simulates Δ. Then the tiling problem of Δ many-one reduces to that of Γ. In particular, if Δ has unsolvable tiling problem, then so does Γ.*

Proof. Let $\mathcal{F}$ be an instance of tiling problem for Δ. As in the proof of Theorem 3.4, we may solve the tiling problem for Γ on instance $\mathcal{F}^{\mathcal{S}}$, so it suffices to prove

$$\exists \mathcal{G} \in \Gamma\colon \mathrm{Hom}(\mathcal{G}, \mathcal{F}^{\mathcal{S}}) = \emptyset \iff \exists \mathcal{H} \in \Delta\colon \mathrm{Hom}(\mathcal{H}, \mathcal{F}) = \emptyset.$$

If $\mathrm{Hom}(\mathcal{G}, \mathcal{F}^{\mathcal{S}})$ is non-empty for some $\mathcal{G} \in \Gamma$, then $\mathrm{Hom}(\mathcal{G} \rtimes \mathcal{S}, \mathcal{F}) \neq \emptyset$ by Lemma 2.16, so by the second assumption there is a graph $\mathcal{H} \in \Delta$ such that

$$\mathrm{Hom}(\mathcal{H}', \mathcal{F}) \supseteq \mathrm{Hom}(\mathcal{G} \rtimes \mathcal{S}, \mathcal{F}) \circ \mathrm{Hom}(\mathcal{H}', \mathcal{G} \rtimes \mathcal{S}) \neq \emptyset$$

for all finite $\mathcal{H}' \subseteq \mathcal{H}$. Since $\mathcal{F}$ is finite, $\mathrm{Hom}(\mathcal{H}, \mathcal{F}) = \lim \mathrm{Hom}(\mathcal{H}', \mathcal{F}) \neq \emptyset$ by compactness.

Conversely, if $\mathrm{Hom}(\mathcal{H}, \mathcal{F}) \neq \emptyset$ for some $\mathcal{H} \in \Delta$, then by the first assumption there is a graph $\mathcal{G} \in \Gamma$ such that $\mathcal{G} \rtimes \mathcal{S}$ is étale and $\mathrm{Hom}(\mathcal{G}' \rtimes \mathcal{S}, \mathcal{H}) \neq \emptyset$ for all finite $\mathcal{G}' \subseteq \mathcal{G}$, so $\mathrm{Hom}(\mathcal{G}' \rtimes \mathcal{S}, \mathcal{F}) \neq \emptyset$, so $\mathrm{Hom}(\mathcal{G}', \mathcal{F}^{\mathcal{S}}) \neq \emptyset$ by Lemma 2.16 (because subgraphs of étale graphs are étale). Now $\mathcal{F}^{\mathcal{S}}$ is finite, so $\mathrm{Hom}(\mathcal{G}, \mathcal{F}^{\mathcal{S}}) = \lim \mathrm{Hom}(\mathcal{G}', \mathcal{F}^{\mathcal{S}}) \neq \emptyset$ by compactness. ∎

Remark 3.8. If Γ and Δ are singletons with the same labelling graph $\mathcal{A} = \mathcal{B}$, and $\mathcal{S} = \mathcal{A}$ is the trivial simulator (both labellings are the identity map), the theorem reduces to the fact that if two graphs weakly map to each other, then their tiling problems are equivalent.

The literature often mentions the "seeded tiling problem"; we shall return to it in Section 3.3. It suffices, for now, to define it as a tiling problem for a graph with the sunny-side-up labelling, see Example 2.4.

As a side-note, it is an interesting, and not yet fully understood, problem to determine which groups admit a sunny-side-up labelling defined as a factor (shift-commuting continuous image) of a shift of finite type (equivalently, factor of a DHS), see [9]. Such images are called *sofic*. For us, the sunny-side-up labelling is fixed once and for all on the graph, and exists independently of the tiling problem. The soficity of the sunny-side-up on the lamplighter group is a side-effect of our constructions, see Proposition 6.5.

Our undecidability result rests on the following result of Wang, proven itself by a reduction to the halting problem of Turing machines.

Theorem 3.9 (Kahr–Moore–Wang [14, 21]). *The seeded tiling problem on $\mathbb{Z}^2$ (namely, on the Cayley graph of $\mathbb{Z}^2$ marked by the sunny-side-up) is unsolvable.*

Since the quadrant and half-plane simulate the plane, by Example 2.14, it follows that the tiling problem on the quadrant (with the markings from Example 2.13) is also unsolvable.

One can interpret Robinson's classical proof of undecidability of the tiling problem [19] as simulating seeded-$\mathbb{Z}^2$ on $\mathbb{Z}^2$. We give an informal explanation that concentrates on the link to weak mappings, assuming the reader is familiar with the proof. More specifically, we have in mind Kari's presentation [15, Section 8.2.4]. This continues Example 2.18.

Example 3.10. In Robinson's proof of the undecidability of the tiling problem, one builds a subshift of finite type whose configurations contain drawings of squares (containing squares containing squares ...) around a square grid (possibly with some degenerate squares), so that

- each configuration contains arbitrarily large finite squares;
- the larger a square is, the more "free rows" (resp. free columns) it contains; a free row is one that does not hit (a smaller square inside);
- in some configuration every cell of the grid is contained in a finite square.

One can use additional signals to mark the free rows and free columns inside the rectangles, and the cells that are part of a free row and a free column necessarily form a set-theoretic rectangle. Thus, every square can be seen as being of the type in Example 2.18.

The bottom row and leftmost column can be made visible in each cell, and thus we can modify the construction in Example 2.18 slightly to obtain a simulator $\mathcal{S}$ such that the simulated rectangles $\{0, 1, \dots, |A| - 1\} \times \{0, 1, \dots, |B| - 1\}$ simulate a rectangle on the left corner of $\mathbb{N}^2$ with its natural vertex labelling from Example 2.13.

We claim that the family Γ of all tilings of $\mathbb{Z}^2$ by this tile set then weakly simulates the singleton family $\Delta = \{\mathbb{N}^2\}$ (with its natural labelling). To prove

every graph in Δ may be simulated: for every $\mathcal{H} \in \Delta$ there is $\mathcal{G} \in \Gamma$ such that $\mathcal{G} \rtimes \mathcal{S}$ is weakly étale and weakly maps to $\mathcal{H}$, take $\mathcal{G} \in \Gamma$ from the third item. Then every finite subgraph of $\mathcal{G} \rtimes \mathcal{S}$ is contained in one that is a disjoint union of simulated full finite squares. These are subgraphs of $\mathbb{N}^2$, so to get a graph homomorphism, on each such square separately we can take its graph embedding into $\mathbb{N}^2$. It is clear that all simulated graphs are weakly étale because of the form of the simulator, so in particular $\mathcal{G} \rtimes \mathcal{S}$ is. To prove

> *simulated graphs are images of* Δ*: for every* $\mathcal{G} \in \Gamma$ *there is* $\mathcal{H} \in \Delta$ *that weakly maps to* $\mathcal{G} \rtimes \mathcal{S}$*,*

for any finite subgraph of $\mathcal{H}' \subset \mathcal{H}$, taking a large enough square in $\mathcal{G}$, we see that $\mathcal{G} \rtimes \mathcal{S}$ contains an actual copy of $\mathcal{H}'$, and we can use the graph embedding as the graph homomorphism.

Since $\mathbb{N}^2$ with the labelling from Example 2.13 simulates seeded-$\mathbb{Z}^2$, Γ also weakly simulates {seeded-$\mathbb{Z}^2$}, since weak simulation is easily seen to be transitive.

3.3. General SFTs on Cayley and Schreier graphs

We further develop the link between graphs and groups, sketched in Example 2.2. Consider a group $G = \langle S \rangle$, and a set X on which G acts on the right. We associate with X the *Schreier graph* $\mathcal{X}$ with vertex set X and edge set $X \times S$, with as usual $(x, s)^- = x$ and $(x, s)^+ = xs$. If $S = S^{-1}$, then $\mathcal{X}$ is unoriented with $(x, s)' = (xs, s^{-1})$.

Definition 3.11. A *subshift of finite type* (short: SFT) Ω on X is given by a finite set A called the *alphabet*, an integer n called the *radius*, and a subset Π of $A^{S^{\leq n}}$ called the *allowed patterns*. It is defined as

$$\Omega = \{\alpha \in A^X : \forall x \in X\colon \exists P_x \in \Pi\colon\ P_x(w) = \alpha(xw) \text{ for all } w \in S^{\leq n}\},$$

namely the set of labellings of X by elements of A such that, in every neighbourhood of size n, the labels form an allowed pattern.

The definition above is a generalization of the more classical notion, in which $X = G$ with action by translation. The *tiling problem* for X asks for an algorithm that, given $\Pi \subseteq A^{S^{\leq n}}$, determines whether the corresponding SFT Ω is non-empty.

An important variant is the *seeded tiling problem*, which asks for an algorithm that, given $\Pi \subseteq A^{S^{\leq n}}$, $x_0 \in X$ and $a_0 \in A$, determines whether the corresponding Ω contains a configuration $\alpha \in A^X$ with $\alpha(x_0) = a_0$.

As we shall now see, the tiling problem for X is essentially equivalent to the tiling problem on graphs from Definition 3.3, and the seeded tiling problem is essentially equivalent to the tiling problem on a graph with a marked vertex (as in Example 2.4).

We need a few technicalities. First, SFTs of course lose all edge information, so we need the following definition.

Definition 3.12. A DHS $\mathrm{Hom}(\mathcal{G}, \mathcal{F})$ is *weakly resolving* if $\mathcal{F}$ is weakly étale.

Note that one can make any DHS weakly resolving at the cost of adding a few more vertices to $\mathcal{F}$, without changing the system up to isomorphism (in the sense of the following definition). Even when no information is lost, the constructions between SFTs and DHSs are only inverses of each other up to isomorphism.

We now give a suitable notion of isomorphism:

Definition 3.13. Let Ω_1, Ω_2 be SFTs on a G-set X. A *block map* from Ω_1 to Ω_2 is a map of the form

$$f(\eta)_x = f_{\text{loc}}(\eta(xg_1), \ldots, \eta(xg_k))$$

for some $f_{\text{loc}}\colon A^k \to B$ and some fixed $g_1, \ldots, g_k \in G$. We say two SFTs are *(block map) isomorphic* if there are block maps $f_i\colon \Omega_i \to \Omega_{3-i}$ which are inverses of each other. Similarly one can define block maps and isomorphisms between DHSs $\mathrm{Hom}(\mathcal{X}, \mathcal{F}_1)$ and $\mathrm{Hom}(\mathcal{X}, \mathcal{F}_2)$, as well as between SFTs and DHSs.

In the classical situation $G = X$ with G acting by $g \cdot \eta(h) = \eta(g^{-1}h)$, morphisms between SFTs are just the usual morphisms of topological G-systems, namely shift-commuting continuous functions, and isomorphisms are just the *topological conjugacies*, or shift-commuting homeomorphisms.

In the general situation, there are some subtleties. If X is a G-set, then G also acts on A^X by $g\eta(x) = \eta(xg)$, but the above block maps are not the continuous functions commuting with this action. Indeed, if G acts ∞-transitively on X, then only finitely many continuous functions commute with its natural action on A^X (but there are plenty of morphisms in the above sense); on the other hand if $G = \mathbb{Z}$ and no element in X has infinite orbit but infinitely many elements have non-trivial orbit, then there are uncountably many shift-commuting continuous functions on A^X. We also note that bijectivity of a block map $f\colon A^X \to A^X$ for a transitive G-set X is equivalent to having a block map inverse, but this is no longer true if the action is not transitive.

Proposition 3.14. *Let G be a monoid acting on a set X. SFTs on X are equivalent by block map isomorphisms to weakly resolving DHSs.*

In essence, every SFT can be converted into a DHS, and vice versa; the constructions are defined by local rules, and involve no funny business.

⚠ The precise statement we are referring to is the following: In the proof we construct a mapping F that turns an SFT Ω into a weakly resolving DHS $\mathrm{Hom}(\mathcal{X}, \mathcal{F})$, and give another construction F' for the other direction. These extend to functors between the appropriate categories, when one takes the morphisms to be the block maps, and the functors F and F' give an equivalence of categories.

This equivalence is also "by block map isomorphisms", in that the object mappings of F and F' are themselves given by invertible block maps; in the classical dynamical situation of G-subshifts, the object mappings are topological conjugacies.

Proof. We only give the object mappings and show that they are isomorphisms. The choices of mappings between morphisms are obvious, and the verification that the resulting functors are a categorical equivalence is routine.

In the direction "$\mathcal{X}$ to X": let $\mathcal{F}$ be a finite graph with no vertex labels and edge labels S, and consider the DHS $\operatorname{Hom}(\mathcal{X},\mathcal{F})$ that it defines. We construct an SFT Ω on X as follows: we set $A = V(\mathcal{F})$ and $n = 1$, and define $\Pi \subset A^{S^{\le 1}}$ by taking $P \in \Pi$ if for all $s \in S$, the graph $\mathcal{F}$ has an edge with label s from $P(1)$ to $P(s)$. The next two paragraphs describe maps $\Omega \leftrightarrow \operatorname{Hom}(\mathcal{X},\mathcal{F})$.

Firstly, consider $\eta \in \operatorname{Hom}(\mathcal{X},\mathcal{F})$ and construct $\alpha \in A^X$ by $\alpha(x) := \eta(x)$ for all $x \in X$; namely, $\alpha = \eta \restriction V(\mathcal{X})$. We claim $\alpha \in \Omega$. Indeed consider $x \in X$, and define $P_x \in A^{S^{\le 1}}$ by $P_x(w) := \alpha(xw)$ for all $w \in S^{\le 1}$. Then in $\mathcal{X}$ the edge (x,s) has label s from x to xs, and we have $\eta(x) = \alpha(x) = P_x(1)$ and $\eta(xs) = \alpha(xs) = P_x(s)$, to η map the edge (x,s) to some edge in $\mathcal{F}$ with label s from $P_x(1)$ to $P_x(s)$; in particular such an edge exists, and we have $P_x \in \Pi$.

Secondly, consider $\alpha \in \Omega$, and associate a homomorphism $\eta \in \operatorname{Hom}(\mathcal{X},\mathcal{F})$ with it as follows: We define $\eta(x) := \alpha(x)$ on vertices. Consider an edge (x,s) with label s from x to xs, and define $P_x \in A^{S^{\le 1}}$ by $P_x(w) := \alpha(x)$ for all $w \in S^{\le 1}$. Because $P_x \in \Pi$, there must be an edge from $P_x(1)$ to $P_x(s)$ with label s in $\mathcal{F}$, and since $\mathcal{F}$ is weakly resolving, there is a unique such edge, which we call $\eta(x,s)$. In this manner we defined a graph morphism $\eta\colon \mathcal{X} \to \mathcal{F}$.

The constructions are clearly inverses of each other and are given by block maps.

In the direction "X to $\mathcal{X}$": let Ω be an SFT on X for some alphabet A, some $n \ge 0$ and some $\Pi \subset A^{S^{\le n}}$. We define a weakly resolving DHS via a graph $\mathcal{F}$, which is constructed as follows: $V(\mathcal{F}) = \Pi$, and for each $P, P' \in \Pi$ we include in $\mathcal{F}$ an edge $e = (P, s, P')$ labelled $s \in S$ with $e^+ = P$ and $f^- = P'$ whenever $P'(w) = P(sw)$ for all $w \in S^{\le n-1}$. This graph is obviously weakly resolving, since the labelling map and the head and tail maps are projections. The next two paragraphs describe maps $\operatorname{Hom}(\mathcal{X},\mathcal{F}) \leftrightarrow \Omega$.

Firstly, consider $\alpha \in \Omega$. We construct a homomorphism $\eta \in \operatorname{Hom}(\mathcal{X},\mathcal{F})$ as follows. On vertices $x \in X$, set $\eta(x) := P$ with $P(w) := \alpha(xw)$ for all $w \in A^{S^{\le n}}$. For edges (x,s) of $\mathcal{X}$, on whose extremities we have already defined $\eta(x) = P$ and $\eta(xs) = P'$, note that by definition we have

$$P'(w) = \alpha(xs \cdot w) = \alpha(x \cdot sw) = P(sw)$$

for all $w \in S^{\le n-1}$, so $\mathcal{F}$ has an edge labelled s from P to P'. We let $\eta(x,s)$ be this edge. By definition of $\mathcal{F}$, we have indeed defined a graph morphism $\eta \in \operatorname{Hom}(\mathcal{X},\mathcal{F})$.

Secondly, consider $\eta \in \operatorname{Hom}(\mathcal{X},\mathcal{F})$, and construct a configuration $\alpha \in A^X$ as follows. Set $\alpha(x) := \eta(x)(1)$, namely look at the pattern $\eta(x)$ and extract its symbol at the identity. We claim $\alpha \in \Omega$. To see this, define $P_x \in A^{S^{\le n}}$ by $P_x(w) := \alpha(xw)$ for all $w \in S^{\le n}$. To prove that all P_x belong to Π, it suffices to show $P_x = \eta(x)$, since then $P_x \in V(\mathcal{F}) = \Pi$. We show this simultaneously for all $x \in X$, considering all $w \in S^{\le n}$ in order of increasing length. If $|w| = 0$, this is true by definition, since $P_x(1) = \alpha(x) = \eta(x)(1)$. Supposing the claim is true for w, consider a word sw with $s \in S$. Unwrapping the definitions we have

$$P_x(sw) = \alpha(x \cdot sw) = \alpha(xs \cdot w) = P_{xs}(w) = \eta(xs)(w),$$

so we are led to show $\eta(xs)(w) = \eta(x)(sw)$. The edge $(x,s) \in E(\mathcal{X})$ has label s and extremities $(x,s)^+ = xs$ and $(x,s)^- = x$, so if we write $\eta(x,s) = (P,s,P') \in E(\mathcal{F})$, then $P'(w) = P(sw)$, that is, $\eta(xs)(w) = P'(w) = P(sw) = \eta(x)(sw)$ as required.

The constructions are clearly inverses of each other and are given by block maps. ■

In the seeded case, we fix an "origin" o in the G-set X, and consider the Schreier graph of X with generating set S, with the sunny-side-up labelling where o is mapped to the vitellus, and all others to albumen (edge labellings are uniquely determined). The DHSs on this graph are *seeded-DHSs*. (This is a slight generalization of Example 2.4.) We define similarly a seeded variant of SFTs.

Definition 3.15. A *seeded SFT* Ω on (X,o) is given by a finite set A called the *alphabet*, an integer n called the *radius*, and a subset Π of $(A \times \{0,1\})^{S^{\leq n}}$ called the *allowed patterns*. Writing π_1, π_2 respectively for the pointwise projections to the first and second coordinate of the alphabet, the SFT Ω is defined as

$$\Omega = \big\{\pi_1(\alpha) : \alpha \in (A \times \{0,1\})^X \wedge (\pi_2(\alpha)_x = 1 \iff x = o) \wedge \\ \forall x \in X\colon \exists P_x \in \Pi\colon\ P_x(w) = \alpha(xw) \text{ for all } w \in S^{\leq n}\big\}.$$

In other words, the allowed patterns see the marking at the origin, but this is erased in the actual configurations.

We can define isomorphisms on seeded SFTs and on seeded DHSs by block maps, similarly as in the unseeded case. The only difference is that the block map is allowed to behave differently when near the seed, which can be implemented as in Definition 3.15 (allowing them to see the seed position). The proof of the following proposition is similar to that of Proposition 3.14, and is omitted.

Proposition 3.16. *Let G be a monoid acting on a set X. Seeded SFTs on X are equivalent by block map isomorphisms to weakly resolving seeded DHSs.*

4. The lamplighter group

Our main result applies to a specific example, the "lamplighter group". Write $\mathbb{Z}/2$ for the two-element group. The group L may be defined in various manners: it is the wreath product $L = \mathbb{Z}/2 \wr \mathbb{Z}$, namely the extension of $\{f\colon \mathbb{Z} \to \mathbb{Z}/2 \text{ a finitely supported set map}\}$ by $\mathbb{Z}$ acting by shifts.

Writing a for the generator of $\mathbb{Z}$ and d for the delta-function $f\colon \mathbb{Z} \to \mathbb{Z}/2$ taking value 1 at 0 and 0 elsewhere, we have the presentation

$$L = \langle a, d \mid d^2, [d, d^{a^n}] \text{ for all } n > 0 \rangle.$$

We shall prefer the more symmetric presentation, setting $b = da$,

$$L = \langle a, b \mid (a^n b^{-n})^2 \text{ for all } n > 0 \rangle.$$

The reason L is called the "lamplighter group" is the following. Picture a two-way-infinite street, with a house at every integer, and a lamp between any two neighbouring houses. The "lamplighter" starts at house 0, and has a schedule to follow: turn on some specified lamps, and stop at a given house. This schedule is an element of L. It may be expressed as a word in elementary operations: "move to the next house" (a or a^{-1}, depending on the direction), and "move to the next house, flipping the state of the lamp along the way" (b or b^{-1}).

This description, where the lamps are between houses, avoids the issue of whether the lamplighter flips the lamp before or after moving. We find this convenient, and thus from now on consider the lamps to be on $\mathbb{Z} + 1/2$. We use the notation $\mathbb{E} = \mathbb{Z} + 1/2$ for the half-integers. To avoid confusion between the lamplighter in this mental picture and the group itself, we usually refer to the position of the lamplighter as the *head*.

We shall make use of two representations for elements of L: firstly, words over $\{a^{\pm 1}, b^{\pm 1}\}$ as in the second paragraph of this section; and secondly, as a global description of lamp configurations and final position, as follows: if at the end of its schedule the lamplighter is at position $n \leq 0$ and for all $i \in \mathbb{E}$ the lamp at position i is in state s_i, then the corresponding element of L is written

$$s_{-N} \cdots s_{n-1}{}^{\uparrow}s_n \cdots s_{-1/2}{}_{\curlywedge}s_{1/2} \cdots s_M,$$

with N, M minimal such that s_{-N}, s_M are non-zero. If $n \geq 0$, the same notation is used, but with now the symbol '${}_{\curlywedge}$' to the left of the '${}^{\uparrow}$'. Thus the expressions for the generators are respectively

$$a = {}_{\curlywedge}0^{\uparrow}, \quad b = {}_{\curlywedge}1^{\uparrow}, \quad a^{-1} = {}^{\uparrow}0_{\curlywedge}, \quad b^{-1} = {}^{\uparrow}1_{\curlywedge}.$$

Global descriptions may be multiplied as follows: align the '${}^{\uparrow}$' of the first with the '${}_{\curlywedge}$' of the second, and add bitwise the strings of 0 and 1. The '${}_{\curlywedge}$' and '${}^{\uparrow}$' of the result are respectively the '${}_{\curlywedge}$' of the first and the '${}^{\uparrow}$' of the second operand.

From a description $u^{\uparrow}v_{\curlywedge}w$ or $u_{\curlywedge}v^{\uparrow}w$ one easily reads the final position n of the lamplighter, and the states $(s_i)_{i \in \mathbb{Z}+1/2}$ of the lamps. In that notation, the product of (r, m) and (s, n) is $(t, m + n)$ with $t_i = r_i + s_{i-m}$.

Sometimes, the origin in a global description is unimportant, and is omitted; so we may consider partial descriptions of the form '$u^{\uparrow}v$'. Such partial descriptions may be acted upon by L, by right multiplication. They are naturally identified with the homogeneous space $\langle a \rangle \setminus L$.

4.1. The Cayley graph of L

The Cayley graph of L, in the generating set $\{a, b\}^{\pm 1}$, is a special case of a *horocyclic product*, see [5]. Let first $\mathcal{T}_1, \mathcal{T}_2$ be two 3-regular trees, and choose on each of them a infinite ray $\xi_i : (-\mathbb{N}) \to \mathcal{T}_i$. (These rays define points at infinity ω_i in the respective trees). The corresponding *Busemann functions* are $h_i : \mathcal{T}_i \to \mathbb{Z}$ defined by

$$h_i(v) = \lim_{n \to -\infty} n + d(v, \xi_i(n));$$

the points close to ω_i have very negative h_i, and points with same Busemann function value form horocycles with respect to the boundary points ω_i. Now the *horocyclic product* of these trees, with respect to these Busemann functions, is

$$\mathcal{L} = \{(v_1, v_2) \in \mathcal{T}_1 \times \mathcal{T}_2 : h_1(v_1) + h_2(v_2) = 0\}.$$

Formally speaking, we have defined the vertex set of $\mathcal{L}$ above; there is then an edge between (v_1, v_2) and (w_1, w_2) whenever there are edges in $\mathcal{T}_i$ between v_i and w_i for all $i = 1, 2$. One could also say that h_i is extended linearly to edges, and take the definition of $\mathcal{L}$ above at face value.

A yet equivalent formulation is that the Busemann function h_1 defines a graph morphism $\mathcal{T}_1 \to \mathcal{C}(\mathbb{Z}, \{\pm 1\})$, and $-h_2$ defines likewise a graph morphism $\mathcal{T}_2 \to \mathcal{C}(\mathbb{Z}, \{\pm 1\})$. Then $\mathcal{L}$ is the pullback

$$\mathcal{L} = \mathcal{T}_1 \times_{\mathcal{C}(\mathbb{Z}, \{\pm 1\})} \mathcal{T}_2.$$

Proposition 4.1 ([23, §2]). *The Cayley graph $\mathcal{C}(L, \{a, b\}^{\pm 1})$ is the horocyclic product $\mathcal{L}$ defined above.*

The proof is in fact straightforward: picture $\mathcal{T}_1$ as having its boundary point ω_1 at the bottom and $\mathcal{T}_2$ as having its boundary point ω_2 at the top. With our choice of orientation, a point in $\mathcal{L}$ is then a pair of points (v_1, v_2) of same height.

Label all edges of $\mathcal{T}_i$ by $\{0, 1\}$ with the condition that the edges on the rays ξ_i are all 0. Then a vertex $(v_1, v_2) \in \mathcal{L}$ may be uniquely identified by the following data: a height $n \in \mathbb{Z}$, a finite string $u \in \{0, 1\}^*$ expressing the labels on the geodesic from v_1 to ξ_1, and a finite string $v \in \{0, 1\}^*$ expressing the labels on the geodesic from v_2 to ξ_2. The corresponding element of L is 'reverse(u)↑v', with the extra '⸤' inserted n places to the left of the '↑'.

Consider now a finite subgraph of $\mathcal{L}$ as follows: choose $H \in \mathbb{N}$ and vertices $v_i \in \mathcal{T}_i$ with $h(v_1) + h(v_2) = H$. There are height-H binary trees in $\mathcal{T}_i$ consisting of all vertices w_i with $h(w_i) = d(v_i, w_i) \leq H$, and their product, in $\mathcal{L}$, gives a height-H tetrahedron. Such a tetrahedron is displayed in Figure 2 for $H = 4$.

These tetrahedra in $\mathcal{L}$ are naturally nested: every vertex belongs to increasing sequences of tetrahedra, and every height-H tetrahedron is naturally contained in two height-$(H + 1)$ tetrahedra, one extending above it and one below it.

4.2. The geometry of L

We describe some geometric aspects of the Cayley graph of L; these will not be used elsewhere in the text, and serve as an illustration of the relevance of L to the tiling problem.

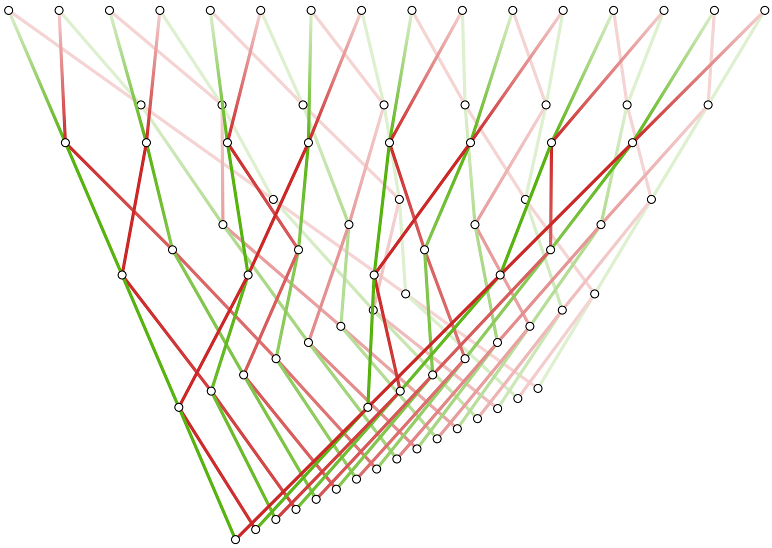

Figure 2. (A portion of) the Cayley graph $\mathcal{L}$ of the lamplighter group, with generators a in green and b in red.

Firstly, $\mathcal{L}$ does not contain any embedded plane, so is a good test case for Conjecture 1.2. Indeed, assume there were an injective Lipschitz map $\mathbb{Z}^2 \to \mathcal{L}$. In particular, the elementary relation $[x, y] = 1$ in $\mathbb{Z}^2$ would map to relations of bounded length in L. Thus we may equivalently ask whether there exists an embedded plane in the finitely presented group

$$L_N = \langle a, d \mid d^2, [d, d^{a^n}] \text{ whenever } 0 < n < N \rangle.$$

(Note that we use, for more convenience, the presentation of L on generators $\{a, d = ab^{-1}\}$). Now free groups do not contain embedded planes, and

Lemma 4.2. *The group L_N is virtually free.*

Proof. Consider the subgroup $K = \langle a, [d, a^N] \rangle$ of L_N. On the one hand, K has index 2^N: every element of L_N may be written in the form $d^{a^{n_1}} \cdots d^{a^{n_k}} a^{\ell}$, and K contains every such expression in which $\#\{i : n_i \equiv m \pmod 2\}$ is even for all $m = 0, \dots, N-1$. On the other hand, one can check using the Reidemeister–Schreier procedure that K is free on its generators [17, §II.4]. ■

The group L is "amenable"; there are numerous equivalent definitions of this property, but the simplest to state is probably the graph-theoretical one: *For every $\varepsilon > 0$, there exists a finite subset $F \subset L$ whose boundary $F \cdot \{a, b\} \setminus F$ has cardinality at most $\varepsilon \# F$.* These subsets may simply be taken to be the tetrahedra mentioned above: a tetrahedron of height H contains $(H + 1) \cdot 2^H$ vertices, and its boundary consists only of the upper and lower strips, so contains $2 \cdot 2^H$ vertices.

4.3. Dominos on L

We now finally begin tiling the lamplighter group. As our application is to the seeded tiling problem, as shown in Proposition 3.16 instead of the sunny-side-up labelling we can simply work with allowed patterns, and specify a different tiling rule near the origin. In practice, we work with tilings of L, and in the end, we pose an additional restriction at the origin, in fact only at the identity element.

We have written the theory of tilings using the DHSs, and in the previous section we explained how to convert between these and the SFTs in the sense of Definition 3.11, which we think of as the most general setting. In the case of the lamplighter group, we introduce two more presentations specific to this group: tetrahedron tilings, and Wang tilings for a particular generating set. Tetrahedron tilings are a special case of SFTs, using the (non-symmetric) generating set $\{e, a, b, ab^{-1}\}$, and they are the most convenient way to present the rules for the sea level construction in Section 6. Wang tiles can be seen as a special case of DHSs, and they are the most convenient way to present the rules for the comb construction in Section 5.

For Wang tiles we always use the generating set $L = \langle a^{\pm 1}, b^{\pm 1}\rangle$. A set of *Wang tiles* consists of a finite set A of edge colours, and a subset $\Pi \subseteq A^{\{a^{\pm 1}, b^{\pm 1}\}}$. Vertices are coloured by elements of Π, and we check that the A-colourings match. We visualize each pattern $\pi \in \Pi$ as a diamond or, equivalently, as a matrix in angle brackets:

$$\pi =$$
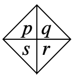
$$= \left\langle \begin{smallmatrix} p & q \\ s & r \end{smallmatrix} \right\rangle \quad \text{means} \quad \pi(a) = p,\ \pi(b) = q,\ \pi(a^{-1}) = r,\ \pi(b^{-1}) = s.$$

The resulting SFT is a subset $\Omega \subseteq \Pi^L$, namely

$$\Omega = \{\eta \in \Pi^L : \forall g \in L\colon\ \eta(g)(a) = \eta(ga)(a^{-1}),\ \eta(g)(b) = \eta(gb)(b^{-1})\}.$$

This can be seen as a special case of the Hom-presentation.

Equivalently, we may specify a colour on every vertex of $\mathcal{L}$, and impose constraints on the edges, or on small subgraphs. We found it most convenient to impose constraints on small, height-1 tetrahedra, as follows. We fix a finite set A of vertex colours, and a subset $\Theta \subseteq A^{\{1, ab^{-1}, a, b\}} \cong A^4$. The resulting *tetrahedron tiling system* is a subset $\Omega \subseteq A^L$, namely

$$\Omega = \{\eta \in A^L : \forall g \in L\colon\ (\eta(g), \eta(gab^{-1}), \eta(ga), \eta(gb)) \in \Theta\}.$$

Note that we may, and do, always assume that Θ is invariant under the permutation $(1,2)(3,4)$ of its coordinates because the condition applied at gab^{-1} is precisely

$$(\eta(gab^{-1}), \eta(g), \eta(gb), \eta(ga)) \in \Theta.$$

For $\theta \in \Theta$ we denote this by $\bar{\theta}$.

We note, even though it is irrelevant to our construction, that the "tetrahedra graph" of $\mathcal{L}$, namely the graph whose vertices are height-1 tetrahedra, and whose edges connect tetrahedra that share a common vertex, is isomorphic to $\mathcal{L}$ (it corresponds to the index-2 subgroup $\langle a, b^2a^{-1}\rangle$ of L, which is isomorphic to L). We may thus equivalently label vertices or tetrahedra by the given tiles.

We can convert between Wang tiles and tetrahedron tilings essentially by the construction of the previous section, and we give the specialized formulas.

It is easy to convert a set of Wang tiles into a set of tetrahedron tiles: assume Ω is given by the Wang tileset $\Pi \subseteq A^{\{a^{\pm1},b^{\pm1}\}}$; then $\Omega \subseteq \Pi^L$ is given by the tetrahedra constraints

$$\Theta = \big\{(\alpha,\beta,\gamma,\delta) \in \Pi^{\{1,ab^{-1},a,b\}} : \alpha(a) = \gamma(a^{-1}),\ \beta(a) = \delta(a^{-1}),\\ \alpha(b) = \delta(b^{-1}),\ \beta(b) = \gamma(b^{-1})\big\}.$$

Conversely, let $\Theta \subseteq A^{\{1,ab^{-1},a,b\}}$ be a collection of tetrahedron tiles. The edge colours will be simply $C = \Theta$. It is easy to see that tilings of the Wang tile set

$$\Pi = \big\{(\theta,\theta,\eta,\bar\eta) : \theta_1 = \eta_3\ [\in A]\big\}.$$

are in one-to-one correspondence (topological conjugacy) to tilings by Θ.

After this section, we will mostly take a more relaxed approach with terminology: "SFT" can refer to DHS, to SFT in the sense of Definition 3.11, or to one of the subclasses from this section. This should always be clear from context, and we have given the formulas for translating between these formalism in Section 3.3 and in the present section.

5. The comb

We construct in $\mathcal{L}$ a geometric structure resembling a "comb": it is an SFT Ω_c marking a bi-infinite line, the *spine* of the comb; rays exiting upwards and downwards from the spine, its *teeth* and *antiteeth*; and some extra synchronizing signals. We then show how, when coloured by a comb, the graph $\mathcal{L}$ simulates the plane $\mathbb{Z}^2$. This comb is defined by the following set Π_c of Wang tiles:

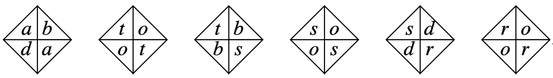

As a graph SFT, it is $\mathrm{Hom}(\mathcal{L},\mathcal{F})$ for the graph $\mathcal{F}$ shown in Figure 3, with generators a in green and b in two-headed red.

Recall our notation for Wang tiles: the edge colour a appears only on the first tile $\left\langle\begin{smallmatrix}a&b\\d&a\end{smallmatrix}\right\rangle$, and implies that each time a vertex $g \in L$ carries the tile $\left\langle\begin{smallmatrix}a&b\\d&a\end{smallmatrix}\right\rangle$, its edges (g, ga) and (g, ga^{-1}) carry the colour a, and thus the whole coset $g\langle a\rangle$ carries the tile $\left\langle\begin{smallmatrix}a&b\\d&a\end{smallmatrix}\right\rangle$; this is the spine of the comb. Likewise, a ray of b's exits in the b direction – the teeth of the comb, and a ray of d's (thought of as reflected b's) exits in the b^{-1} direction – the antiteeth. The tiles $\left\langle\begin{smallmatrix}t&b\\b&s\end{smallmatrix}\right\rangle$ and $\left\langle\begin{smallmatrix}s&d\\d&r\end{smallmatrix}\right\rangle$ share a signal s which propagates via $\left\langle\begin{smallmatrix}s&o\\o&s\end{smallmatrix}\right\rangle$, and are extended on their other ends by respective signals r and t. The remaining a-edges and b-edges may (but are

$\mathcal{F}$:

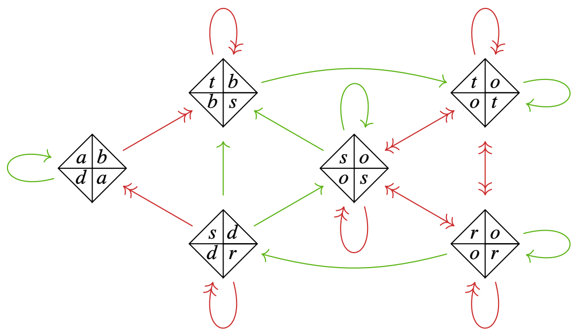

Figure 3. Hom presentation of the comb SFT. (Note that the edges and their colors are fully determined by the edge colors of the Wang tiles.)

not forced to) respectively be labelled t and o. Let us write

$$Z = \{a^n b^m a^k : k, m, n \in \mathbb{Z}\}, \tag{5.1}$$

a union of $\langle a \rangle$-cosets.

Lemma 5.1. *If $\eta \in (\Pi_c)^L$ is a valid Wang tiling and $\eta(1) = \left\langle\begin{smallmatrix} a & b \\ d & a \end{smallmatrix}\right\rangle$, then*

$$\forall n \in \mathbb{Z}\colon \quad \eta(a^n) = \left\langle\begin{smallmatrix} a & b \\ d & a \end{smallmatrix}\right\rangle,$$

$$\forall n \in \mathbb{Z}, m \geq 1\colon \quad \eta(a^n b^m) = \left\langle\begin{smallmatrix} t & b \\ b & s \end{smallmatrix}\right\rangle,$$

$$\forall n \in \mathbb{Z}, m \geq 1\colon \quad \eta(a^n b^{-m}) = \left\langle\begin{smallmatrix} s & d \\ d & r \end{smallmatrix}\right\rangle,$$

$$\forall n \in \mathbb{Z}, 1 \leq k < m\colon \quad \eta(a^n b^m a^{-k}) = \left\langle\begin{smallmatrix} s & o \\ o & s \end{smallmatrix}\right\rangle,$$

$$\forall n \in \mathbb{Z}, 1 \leq k, m\colon \quad \eta(a^n b^{-m} a^{-k}) = \left\langle\begin{smallmatrix} r & o \\ o & r \end{smallmatrix}\right\rangle,$$

$$\forall n \in \mathbb{Z}, 1 \leq k, m\colon \quad \eta(a^n b^m a^k) = \left\langle\begin{smallmatrix} t & o \\ o & t \end{smallmatrix}\right\rangle.$$

Conversely, these formulas, together with $\eta(g) = \left\langle\begin{smallmatrix} t & o \\ o & t \end{smallmatrix}\right\rangle$ whenever $g \notin Z$, define a valid Wang tiling of L.

Proof. Suppose that $\eta \in (\Pi_c)^L$ is a valid Wang tiling with $\eta(1) = \left\langle\begin{smallmatrix} a & b \\ d & a \end{smallmatrix}\right\rangle$. We prove that η satisfies the formulas by a series of Sudoku-style deductions, keeping track of possible values of cells. Since the colour a only appears in the tile $\left\langle\begin{smallmatrix} a & b \\ d & a \end{smallmatrix}\right\rangle$, we have

$$\forall n \in \mathbb{Z}\colon \quad \eta(a^n) = \left\langle\begin{smallmatrix} a & b \\ d & a \end{smallmatrix}\right\rangle.$$

Since the b-colour of $\left\langle\begin{smallmatrix} a & b \\ d & a \end{smallmatrix}\right\rangle$ is b and b only appears as the b^{-1}-colour in the tile $\left\langle\begin{smallmatrix} t & b \\ b & s \end{smallmatrix}\right\rangle$, we then have

$$\forall n \in \mathbb{Z}, m \geq 1: \;\; \eta(a^n b^m) = \left\langle\begin{smallmatrix} t & b \\ b & s \end{smallmatrix}\right\rangle,$$

and by the same argument for d

$$\forall n \in \mathbb{Z}, m \geq 1: \;\; \eta(a^n b^{-m}) = \left\langle\begin{smallmatrix} s & d \\ d & r \end{smallmatrix}\right\rangle.$$

Since the a^{-1}-colour of $\left\langle\begin{smallmatrix} t & b \\ b & s \end{smallmatrix}\right\rangle$ is s, we next have

$$\forall n, m \geq 1: \;\; \eta(a^n b^m a^{-1}) \in \left\{ \left\langle\begin{smallmatrix} s & o \\ o & s \end{smallmatrix}\right\rangle, \left\langle\begin{smallmatrix} s & d \\ d & r \end{smallmatrix}\right\rangle \right\},$$

and since the only tiles with a-colour in $\{r, s\}$ are $\left\{\left\langle\begin{smallmatrix} s & o \\ o & s \end{smallmatrix}\right\rangle, \left\langle\begin{smallmatrix} s & d \\ d & r \end{smallmatrix}\right\rangle, \left\langle\begin{smallmatrix} r & o \\ o & r \end{smallmatrix}\right\rangle\right\}$, and they have their a^{-1}-colours in $\{r, s\}$, we must have

$$\forall n, m \geq 1, k \geq 1: \;\; \eta(a^n b^m a^{-k}) \in \left\{ \left\langle\begin{smallmatrix} s & o \\ o & s \end{smallmatrix}\right\rangle, \left\langle\begin{smallmatrix} s & d \\ d & r \end{smallmatrix}\right\rangle, \left\langle\begin{smallmatrix} r & o \\ o & r \end{smallmatrix}\right\rangle \right\}.$$

Since

$$\eta(a^{n+m} b^{-m}) = \left\langle\begin{smallmatrix} s & d \\ d & r \end{smallmatrix}\right\rangle \quad \text{and} \quad a^n b^m a^{-m} = a^{n+m} b^{-m},$$

the a-colour of $\eta(a^n b^m a^{-m})$ is s, so the a-colour of $\eta(a^n b^m a^{-k})$ is s for all $1 \leq k < m$ and thus

$$\forall 1 \leq k < m: \;\; \eta(a^n b^m a^{-k}) = \left\langle\begin{smallmatrix} s & o \\ o & s \end{smallmatrix}\right\rangle.$$

Finally, the a-colour t of $\eta(a^n b^m)$ forces

$$\forall k \geq 1: \;\; \eta(a^n b^m a^k) = \left\langle\begin{smallmatrix} t & o \\ o & t \end{smallmatrix}\right\rangle,$$

and the a^{-1}-colour r of $\eta(a^n b^{-m})$ forces

$$\forall k \geq 1: \;\; \eta(a^n b^{-m} a^{-k}) = \left\langle\begin{smallmatrix} r & o \\ o & r \end{smallmatrix}\right\rangle.$$

Next, we show that the formulas of the lemma, together with $\eta(g) = \left\langle\begin{smallmatrix} t & o \\ o & t \end{smallmatrix}\right\rangle$ for all $g \notin Z$, indeed define a valid tiling. This means, first, that the point η given by the formulas is well-defined, i.e. exactly one value is given to each coordinate – for this, it suffices to check that the sets

$$Z_1 := \{a^n : n \in \mathbb{Z}\}, \qquad Z_2 := \{a^n b^m : n \in \mathbb{Z}, m \geq 1\},$$
$$Z_3 := \{a^n b^{-m} : n \in \mathbb{Z}, m \geq 1\}, \qquad Z_4 := \{a^n b^m a^{-k} : n \in \mathbb{Z}, 1 \leq k < m\},$$
$$Z_5 := \{a^n b^m a^k : n \in \mathbb{Z}, 1 \leq k, m\}, \qquad Z_6 := \{a^n b^{-m} a^{-k} : n \in \mathbb{Z}, 1 \leq k, m\}$$

are disjoint. This also means that the neighbouring colours are correct whenever two tiles from these sets are adjacent, and that in every other b-direction, their colour is o so that they match with the tile $\left\langle\begin{smallmatrix} t & o \\ o & t \end{smallmatrix}\right\rangle$ used to fill all other cells.

To check these things, we consider the action of the lamplighter group on $\{0,1\}^{\mathbb{E}} \times \mathbb{Z}$, and the orbit of the configuration with the head at the origin and all edges with colour o or t. The set Z defined in (5.1) corresponds to the configurations with a single run of 1s (anywhere). Let us analyse the sets listed above, whose union is Z.

- The set Z_1 is the set of configurations where no 1s have been written.
- The set Z_2 is the set where there is a run of 1s and the head is exactly on the right border of the run.
- The set Z_3 is the set of configurations where there is a run of 1s and the head is exactly on the left border of the run.
- The set Z_4 is the set of configurations where there is a run of 1s and the head is properly inside it.
- The set Z_5 is the set of configurations where there is a run of 1s and the head is strictly to its right.
- The set Z_6 is the set of configurations where there is a run of 1s and the head is strictly to its left.

Obviously these sets are disjoint. (The interpretation of the formula $a^n b^m a^{-m} = a^{n+m} b^{-m}$ used in the proof, in terms of this action, is that the head has either written a run of 1s moving to the right and returned left, or has moved to the right and written a run of 1s while returning.)

Now, let us analyse neighbours of tiles in Z. It is straightforward to verify by a case analysis that whenever a tile in Z has a non-o edge colour in some direction, the corresponding neighbour is also in Z and has the same colour in the opposite direction.

We now study other edges, which must be in direction b since Z is a union of $\langle a \rangle$-cosets. The tile $\left\langle\begin{smallmatrix} s & o \\ o & s \end{smallmatrix}\right\rangle$ appears in Z_4 when the head is properly inside a run of 1s. The b- and b^{-1}-edges have colour o. These neighbours are not in Z: in the action the head either leaves a nonempty run of 1s to the left with a 0 in between, or creates two runs of 1s. The tiles $\left\langle\begin{smallmatrix} t & o \\ o & t \end{smallmatrix}\right\rangle$ and $\left\langle\begin{smallmatrix} r & o \\ o & r \end{smallmatrix}\right\rangle$ in Z_5 and Z_6 are at positions at which the head is out of the run of 1s, so their b- and b^{-1}-neighbours have two runs of 1s and are also not in Z. ■

Note that the Sudoku part of the argument could also be done "bottom up" i.e. starting from $\eta(a^{n+m} b^{-m})$ towards $\eta(a^n b^m)$, in which case one would instead propagate the set

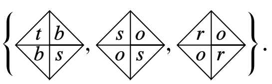

5.1. Simulating $\mathbb{Z}^2$ on $\mathcal{L}$ marked by Ω_c

We consider the Cayley graph $\mathcal{L}$ of the lamplighter group as a labelled graph, the labels given by the "comb" SFT Ω_c described in the previous section. Our aim is to prove the following result.

Proposition 5.2. *The graph $\mathcal{L}$, when labelled by any configuration from the SFT Ω_c with tile $\left\langle\begin{smallmatrix} a & b \\ d & a \end{smallmatrix}\right\rangle$ at the origin, simulates a graph containing the plane $\mathbb{Z}^2$, and the special configuration defined in Lemma 5.1 simulates precisely the plane. The same simulator can be used for each configuration.*

Proof. Recall that the graph $\beta : \mathbb{Z}^2 \to \mathcal{B}$ has edge labels $E(\mathcal{B}) = \{\uparrow, \downarrow, \to, \leftarrow\}$, and vertex labels are trivial, $V(\mathcal{B}) = \{\cdot\}$. The following figure gives the simulator. The split nodes have a vertex label from $\mathcal{L}$ and the trivial vertex label $\cdot$ from $\mathbb{Z}^2$, while the unsplit nodes correspond to vertices of $\mathcal{B}^* \setminus \mathcal{B}$ for the graph $\mathcal{B}$ defining the marking of $\mathbb{Z}^2$ (whose vertex label can be deduced from the edge labels); see Figure 4.

The automaton is undirected: We only write half the edges in the diagram above, those whose label on the $\mathcal{B}^*$-side is $\uparrow$ or $\to$; the missing $\downarrow$ and $\leftarrow$ edges are naturally recovered by applying the involutions of $\mathcal{L}$ and $\mathbb{Z}^2$, namely inverting the generators a, b and reversing the direction of the arrows, switching the source and range tags.

To prove that this simulator indeed produces $\mathbb{Z}^2$, let us explain what it does in terms of graph-walking automata, see Lemma 2.25. First, we recall that $\mathbb{Z}^2$ is simulated, within $\mathcal{L}$, as those vertices on the spine, teeth or antiteeth of the comb; namely, those vertices marked $\left\langle\begin{smallmatrix} a & b \\ d & a \end{smallmatrix}\right\rangle$, $\left\langle\begin{smallmatrix} t & b \\ b & s \end{smallmatrix}\right\rangle$ or $\left\langle\begin{smallmatrix} s & d \\ d & r \end{smallmatrix}\right\rangle$. The bijection associates to $(m, n) \in \mathbb{Z}^2$ the element $a^n b^m$.

The graph walking automaton for the generator '$\to$' of $\mathbb{Z}^2$ is simply "if you're on the spine of the comb, move onto a tooth; if you're on a tooth, move further on the tooth; if you're on an antitooth, move towards the spine". This is realized by the four arrows marked $b/(0, \to, 0)$.

The generator '$\uparrow$' is programmed as follows: "if you're on the spine, move up the spine. If you're on a tooth, follow the 's' signal (using generator a^{-1}) till you reach an antitooth. Then do a step on that antitooth (using generator b), follow the 's' signal back up to a tooth (using generator a) and finally do a step on that tooth. If you're on an antitooth, do the same, except you start by following the 's' signal using generator a till you reach a tooth."

On the one hand, it is easy to see that this is what the above simulator does, following the big hexagon-shaped counterclockwise paths; on the other hand, let us convince ourselves that these operations indeed implement movement on $\mathbb{Z}^2$.

The operation $\uparrow$ is $n \mapsto n + 1$ on $\{(m, n) \in \mathbb{Z}^2\}$. The spine is the subset $Z_1 = \{a^n : n \in \mathbb{Z}\}$ and corresponds to $\{0\} \times \mathbb{Z}$; so $\uparrow$ is simply implemented by following the generator a. Consider now a point $(m, n) \in \mathbb{Z}^2$ with $m > 1$, appearing on a tooth as $a^n b^m \in Z_2$. Then the path along the hexagon leads us successively (using m times a^{-1} along Z_4, following the 's' signal) to $a^n b^m a^{-m} = a^{n+m} b^{-m} \in Z_3$, then to $a^{n+m} b^{-(m-1)} \in Z_3 \cup Z_1$ (we

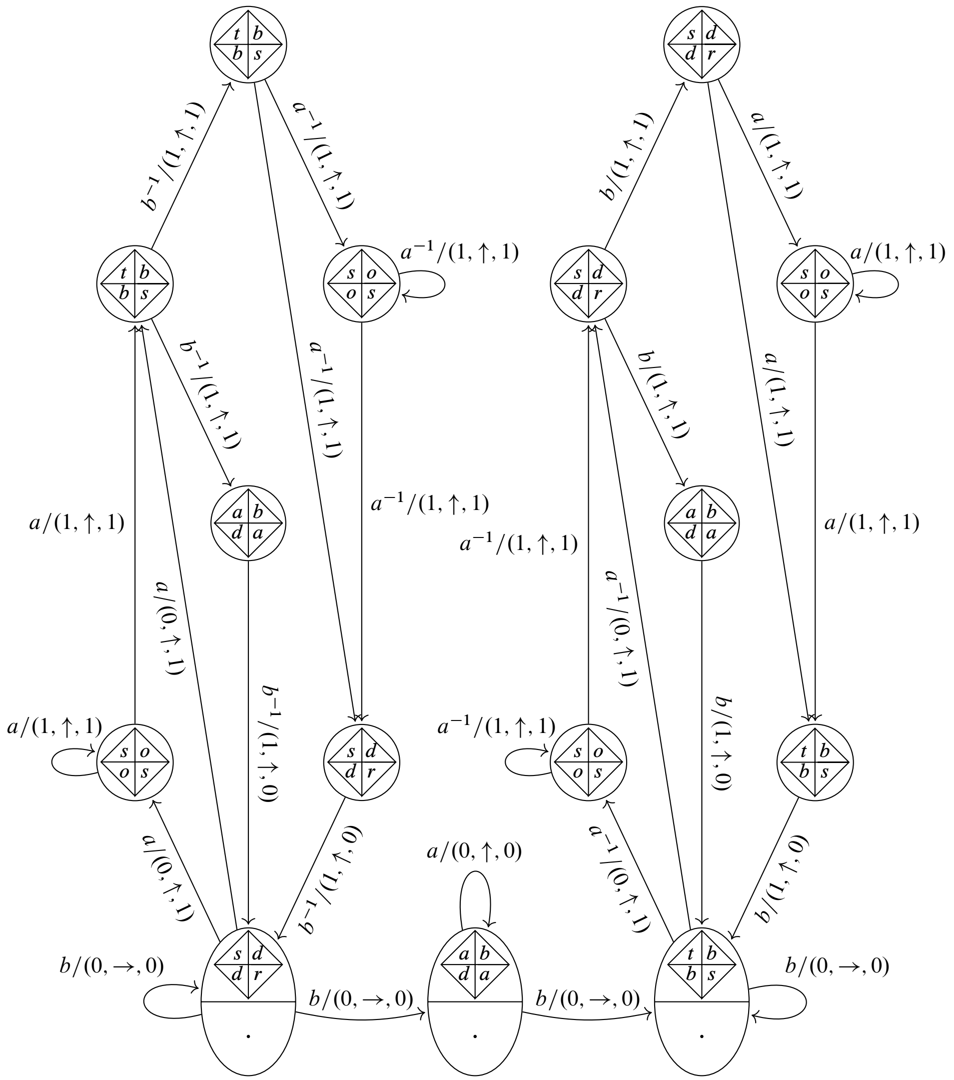

Figure 4. A simulator for $\mathbb{Z}^2$ in $\mathcal{L}$.

reach the spine Z_1 if $m = 1$), then (using $(m-1)$ times a along Z_4 following the 's' signal) to $a^{n+m}b^{-(m-1)}a^{m-1} = a^{n+1}b^{m-1}$, and finally to $a^{n+1}b^m$. This corresponds to the point $(m, n+1) \in \mathbb{Z}^2$ as required. The same argument applies to the antitooth. ■

This gives the first proof of Theorem A:

Corollary 5.3. *The seeded tiling problem on the lamplighter group is undecidable.*

Proof. We apply Theorem 3.7 with Γ the family of all graphs obtained from $\mathcal{L}$ by labelling it by any configuration from the SFT Ω_c with tile $\left\langle\begin{smallmatrix} a & b \\ d & a \end{smallmatrix}\right\rangle$ at the origin, and $\Delta = \{\mathbb{Z}^2\}$. Let $\mathcal{S}$ be the simulator constructed in the previous lemma.

The first condition of Theorem 3.7,

(1) every graph in Δ may be simulated: for every $\mathcal{H} \in \Delta$ there is $\mathcal{G} \in \Gamma$ such that $\mathcal{G} \rtimes \mathcal{S}$ is étale and weakly maps to $\mathcal{H}$,

holds because the unique element $\mathbb{Z}^2 \in \Delta$ is even isomorphic to $\mathcal{G} \rtimes \mathcal{S}$, for $\mathcal{G}$ the special configuration defined in Lemma 5.1. Since $\mathbb{Z}^2$ is étale, so is the simulating graph.

The second condition of Theorem 3.7,

(2) simulated graphs are images of Δ: for every $\mathcal{G} \in \Gamma$ there is $\mathcal{H} \in \Delta$ that weakly maps to $\mathcal{G} \rtimes \mathcal{S}$,

holds because every $\mathcal{G} \in \Gamma$ contains the configuration simulating $\mathbb{Z}^2$ as a subgraph, so we even have $\mathbb{Z}^2 \subset \mathcal{G} \rtimes \mathcal{S}$. ∎

The simulator given above is simple enough that it may be directly translated to a tileset on $\mathcal{L}$, and we do so here, incorporating some ad hoc simplifications.

We construct, from a seeded Wang tileset T on a half-plane, a seeded Wang tileset Π_T on L which tiles if and only if T tiles. The construction could easily be extended to the whole plane, at the cost of extra clutter. Let us consider the half-plane

$$\mathbb{H} = \big\{(m,n) \in \mathbb{Z}^2 : m \geq n\big\}.$$

For a colour set C, a Wang tileset is a subset $T \subseteq C^{\{S,E,N,W\}} = C^4$ interpreted as follows: $(i,j,k,\ell) \in T$ is a square with colours i,j,k,ℓ respectively on the south, east, north, west sides. A valid tiling of the half-plane $\mathbb{H}$ is an assignment $\zeta \colon \mathbb{H} \to T$ with

$$(\zeta_{(m,n-1)})_N = (\zeta_{(m,n)})_S \quad \text{and} \quad (\zeta_{(m,n)})_E = (\zeta_{(m+1,n)})_W$$

for all $(m,n) \in \mathbb{H}$.

It is easy to show that the (un)decidability of the seeded tiling problem on this rotated half-plane is equivalent to that on the standard half-plane. The use of this rotated half-plane $\mathbb{H}$ simplifies somewhat the construction; we use the embedding of $\mathbb{H}$ into $\mathcal{L}$ given by $(m,n) \mapsto a^n b^{m-n}$.

Proposition 5.4. *Let C be a finite set of colours, let $T \subseteq C^{\{S,E,N,W\}} = C^4$ be a Wang tileset for $\mathbb{H}$, and let $t_0 \in T$ be a seed tile. Then a tileset Π_T on L and a seed $\pi_0 \in \Pi_T$ may be algorithmically constructed, such that there exists a valid tiling of $\mathbb{H}$ by T with value t_0 at $(0,0)$ if and only if there exists a valid tiling of L by Π_T with value π_0 at 1.*

Proof. The tileset Π_T will be given by product tiles, with a tile from Π_c in the first layer and a word of length $\in \{0,1,2\}$ over C in the second layer. We denote by ε the empty word (which we of course assume disjoint from C).

For all tiles $(i,j,k,\ell) \in T$ we take the following product tiles in Π_T:

$$\Big(\big\langle\begin{smallmatrix} a & b \\ d & a\end{smallmatrix}\big\rangle, \big\langle\begin{smallmatrix} \varepsilon & j \\ i & \varepsilon\end{smallmatrix}\big\rangle\Big), \quad \Big(\big\langle\begin{smallmatrix} t & o \\ o & t\end{smallmatrix}\big\rangle, \big\langle\begin{smallmatrix} \varepsilon & \varepsilon \\ \varepsilon & \varepsilon\end{smallmatrix}\big\rangle\Big), \quad \Big(\big\langle\begin{smallmatrix} r & o \\ o & r\end{smallmatrix}\big\rangle, \big\langle\begin{smallmatrix} \varepsilon & \varepsilon \\ \varepsilon & \varepsilon\end{smallmatrix}\big\rangle\Big),$$
$$\Big(\big\langle\begin{smallmatrix} t & b \\ b & s\end{smallmatrix}\big\rangle, \big\langle\begin{smallmatrix} \varepsilon & j \\ \ell & ik\end{smallmatrix}\big\rangle\Big), \quad \Big(\big\langle\begin{smallmatrix} s & o \\ o & s\end{smallmatrix}\big\rangle, \big\langle\begin{smallmatrix} ik & \varepsilon \\ \varepsilon & ik\end{smallmatrix}\big\rangle\Big), \quad \Big(\big\langle\begin{smallmatrix} s & d \\ d & r\end{smallmatrix}\big\rangle, \big\langle\begin{smallmatrix} ik & k \\ i & \varepsilon\end{smallmatrix}\big\rangle\Big).$$

The seed tile is $\pi_0 = \big(\langle\begin{smallmatrix} a & b \\ d & a\end{smallmatrix}\rangle, \langle\begin{smallmatrix} \varepsilon & j \\ i & \varepsilon\end{smallmatrix}\rangle\big)$ where $t_0 = (i,j,k,\ell)$.

First, we prove that if there exists a valid tiling $\eta \in (\Pi_T)^L$ with $\eta(1) = \pi_0$, then T admits a valid tiling with seed t_0. To see this, observe that we must have a valid instance of Π_c on the first layer, so we can define the sets $Z = Z_1 \sqcup Z_2 \sqcup Z_3 \sqcup Z_4 \sqcup Z_5 \sqcup Z_6$ as in the proof of Lemma 5.1 and when $g \in Z_i$, we have $\eta(g) = (X_i, t)$ for some t, with the correspondence

$$X_1 = \big\langle\begin{smallmatrix} a & b \\ d & a\end{smallmatrix}\big\rangle, \quad X_2 = \big\langle\begin{smallmatrix} t & b \\ b & s\end{smallmatrix}\big\rangle, \quad X_3 = \big\langle\begin{smallmatrix} s & d \\ d & r\end{smallmatrix}\big\rangle,$$
$$X_4 = \big\langle\begin{smallmatrix} s & o \\ o & s\end{smallmatrix}\big\rangle, \quad X_5 = \big\langle\begin{smallmatrix} t & o \\ o & t\end{smallmatrix}\big\rangle, \quad X_6 = \big\langle\begin{smallmatrix} r & o \\ o & r\end{smallmatrix}\big\rangle.$$

From η we deduce a configuration $\zeta \in T^{\mathbb{H}}$ as follows:

- If $\eta(a^n b^m) = \big(\langle\begin{smallmatrix} t & b \\ b & s\end{smallmatrix}\rangle, \langle\begin{smallmatrix} \varepsilon & j \\ \ell & ik\end{smallmatrix}\rangle\big)$ for $n \in \mathbb{Z}, m \geq 1$, define $\zeta_{(m+n,n)} = (i,j,k,\ell)$.
- If $\eta(a^n) = \big(\langle\begin{smallmatrix} a & b \\ d & a\end{smallmatrix}\rangle, \langle\begin{smallmatrix} \varepsilon & j \\ i & \varepsilon\end{smallmatrix}\rangle\big)$, define $\zeta_{(n,n)} = (i,j,k,\ell)$ for any k,ℓ such that $(i,j,k,\ell) \in T$.
- Set $\zeta_{(0,0)} = t_0$.

By construction we have $\zeta_{(0,0)} = t_0$, and we need to check $(\zeta_{(m,n)})_E = (\zeta_{(m+1,n)})_W$ and $(\zeta_{(m,n-1)})_N = (\zeta_{(m,n)})_S$.

We have $(\zeta_{(m,n)})_E = (\zeta_{(m+1,n)})_W$ directly from the colouring rules of Π_T, since $(\zeta_{(m,n)})_E$ is the b-colour on the second layer of the tile at $\eta(a^n b^{m-n})$ and $(\zeta_{(m+1,n)})_W$ is the b^{-1}-colour on the second layer of the tile at $\eta(a^n b^{m+1-n})$.

We now check the formula $(\zeta_{(m,n-1)})_N = (\zeta_{(m,n)})_S$. If $m > n$ and $\zeta_{(m,n)} = (i,j,k,\ell)$, then the second layer of $\eta(a^n b^{m-n})$ is $\langle\begin{smallmatrix} \varepsilon & j \\ \ell & ik\end{smallmatrix}\rangle$, so its a^{-1}-colour is ik. It follows that in all the tiles $\eta(a^n b^{m-n} a^{-p})$ with $1 \leq p < m-n$ (namely, at positions in Z_4, where the first layer is $\langle\begin{smallmatrix} s & o \\ o & s\end{smallmatrix}\rangle$) the second layer contains the tile $\langle\begin{smallmatrix} ik & \varepsilon \\ \varepsilon & ik\end{smallmatrix}\rangle$. Therefore at $\eta(a^n b^{m-n} a^{n-m}) = \eta(a^m b^{n-m})$ (at a position in Z_3, where the first layer is $\langle\begin{smallmatrix} s & d \\ d & r\end{smallmatrix}\rangle$), by the choice of tiles overlayed on $\langle\begin{smallmatrix} s & d \\ d & r\end{smallmatrix}\rangle$, the tile on the second layer must precisely be $\langle\begin{smallmatrix} ik & k \\ i & \varepsilon\end{smallmatrix}\rangle$. In this manner, we have proven the following property $(\star_{m,n})$ for $m > n$:

- The b-colour of the second layer of $\eta(a^m b^{n-m})$ is $(\zeta_{(m,n)})_N$.
- The b^{-1}-colour of the second layer of $\eta(a^m b^{n-m})$ is $(\zeta_{(m,n)})_S$.

By $(\star_{m,n-1})$ the b-colour at $\eta(a^m b^{n-1-m})$ is $(\zeta_{(m,n-1)})_N$. If $m > n$, then by $(\star_{m,n})$ the b^{-1}-colour at $\eta(a^m b^{n-m})$ is $(\zeta_{(m,n)})_S$; since $a^m b^{n-m}$ is the b-neighbour of $a^m b^{n-1-m}$, we have $(\zeta_{(m,n-1)})_N = (\zeta_{(m,n)})_S$ as required. If $m = n$, then the b^{-1}-colour at $\eta(a^m)$

is $(\zeta_{(n,n)})_S$ and again we have $(\zeta_{(n,n-1)})_N = (\zeta_{(n,n)})_S$. We have proven that a valid π_0-seeded tiling for Π_T yields a valid t_0-seeded tiling of $\mathbb{H}$.

Conversely, if there is a valid tiling $\zeta \in T^{\mathbb{H}}$ with $\zeta_{(0,0)} = t_0$, then the above proof shows rather directly how to construct a valid configuration $\eta \in (\Pi_T)^L$ with $\eta(1) = \pi_0$: set

$$\forall n \in \mathbb{Z}\colon\ \zeta_{(n,n)} = (i,j,k,\ell) \implies \eta(a^n) = \left(\begin{array}{c|c} a & b \\ \hline d & a \end{array}, \begin{array}{c|c} \varepsilon & j \\ \hline i & \varepsilon \end{array}\right),$$

$$\forall m > n \in \mathbb{Z}\colon\ \zeta_{(m,n)} = (i,j,k,\ell) \implies \eta(a^n b^{m-n}) = \left(\begin{array}{c|c} t & b \\ \hline b & s \end{array}, \begin{array}{c|c} \varepsilon & j \\ \hline \ell & ik \end{array}\right),$$

$$\forall m > n \in \mathbb{Z}\colon\ \zeta_{(m,n)} = (i,j,k,\ell) \implies \eta(a^m b^{n-m}) = \left(\begin{array}{c|c} s & d \\ \hline d & r \end{array}, \begin{array}{c|c} ik & k \\ \hline i & \varepsilon \end{array}\right),$$

$$\forall m > p > n \in \mathbb{Z}\colon\ \zeta_{(m,n)} = (i,j,k,\ell) \implies \eta(a^n b^{m-n} a^{p-m}) = \left(\begin{array}{c|c} s & o \\ \hline o & s \end{array}, \begin{array}{c|c} ik & \varepsilon \\ \hline \varepsilon & ik \end{array}\right),$$

$$\forall p > m > n \in \mathbb{Z}\colon\ \eta(a^n b^{m-n} a^{p-m}) = \left(\begin{array}{c|c} t & o \\ \hline o & t \end{array}, \begin{array}{c|c} \varepsilon & \varepsilon \\ \hline \varepsilon & \varepsilon \end{array}\right),$$

$$\forall m > n > p \in \mathbb{Z}\colon\ \eta(a^n b^{m-n} a^{p-m}) = \left(\begin{array}{c|c} r & o \\ \hline o & r \end{array}, \begin{array}{c|c} \varepsilon & \varepsilon \\ \hline \varepsilon & \varepsilon \end{array}\right),$$

$$\forall g \notin Z\colon\ \eta(g) = \left(\begin{array}{c|c} t & o \\ \hline o & t \end{array}, \begin{array}{c|c} \varepsilon & \varepsilon \\ \hline \varepsilon & \varepsilon \end{array}\right). \qquad \blacksquare$$

6. The sea level

In this section, we show, again, that $\mathcal{L}$ can simulate the plane $\mathbb{Z}^2$. We do this in steps:

(1) Using an SFT, we can mark trees in $\mathcal{L}$, or more precisely subgraphs of the form $\mathcal{T} \times_{\mathbb{Z}} \mathbb{Z}$ and of the form $\mathbb{Z} \times_{\mathbb{Z}} \mathcal{T}$ in the pullback description of $\mathcal{L}$. They correspond to limits of faces of tetrahedra.

(2) Using another SFT, we can mark vertices at height 0 in $\mathcal{L}$. A vertex at height 0 is naturally identified with a point in the plane: from $u \!\upharpoonleft\!\!\downharpoonright\! v$ we read u, v as binary expansions of integers x, y respectively, and identify $u \!\upharpoonleft\!\!\downharpoonright\! v$ with (x, y).

(3) Using a simulation, we show how the arithmetic operations $(x, y) \mapsto (x \pm 1, y)$ and $(x, y) \mapsto (x, y \pm 1)$ can be described in $\mathcal{L}$.

In fact, we first explain how the construction lets us simulate the quadrant $\mathbb{N}^2$, and then explain which changes let us simulate a whole plane. This last step is in fact unnecessary, since the quadrant simulates the plane (Example 2.14) and simulation is transitive (Lemma 2.17).

In some sense, this simulation is more efficient than the "comb" from Section 5: the square $[1, 2^n]^2 \subset \mathbb{Z}^2$ is simulated within a tetrahedron of height $2n$, and therefore of size $(2n + 1) \cdot 2^{2n}$.

6.1. Marking a ray $\langle a \rangle$

The SFTs in this section will be given by patterns on tetrahedra. We choose as alphabet the Boolean algebra $\mathbb{B} = \{\bot, \top\}$ meaning false and true, and define $\Pi_{\leftarrow}$ as the following set of patterns:

$$\Pi_{\leftarrow} = \big\{(\alpha, \beta, \gamma, \delta) \in \mathbb{B}^4 = \mathbb{B}^{\{1, ab^{-1}, a, b\}} : \alpha \vee \beta \Longrightarrow \gamma \wedge \delta;\;\; \gamma \vee \delta \Longrightarrow \alpha \neq \beta\big\}.$$

In the geometric model of the Cayley graph, that means that above every $\top$ node both neighbours are $\top$, while below a $\top$ node precisely one of the neighbours is $\top$. We shall see that $\Pi_{\leftarrow}$ forces a ray in $\mathcal{L}$ to be marked $\top$. In passing, we observe another property of the subshift generated by $\Pi_{\leftarrow}$.

Definition 6.1. Let X be a compact metric space. A topological dynamical system $G \curvearrowright X$ is *almost minimal* if there is a unique G-fixed point $0 \in X$, and $\overline{Gx} = X$ for all $x \neq 0$.

We briefly recall and extend our notation for lamplighter group elements that we introduced in Section 4. Recall that the lamps in the lamplighter group are at positions in the half-integers $\mathbb{E} = \mathbb{Z} + \frac{1}{2}$. Every $g \in L$ may be written as $g = (s, n)$ with $s : \mathbb{E} \to \mathbb{Z}/2$ and $n \in \mathbb{Z}$. We write $s_{>n}$ for the right subword $s(n + 1/2)s(n + 3/2)\ldots$ and $s_{<n}$ for the left subword $\ldots s(n - 3/2)s(n - 1/2)$.

There is a natural action of L on $(\mathbb{Z}/2)^{\mathbb{E}}$, if we interpret $(\mathbb{Z}/2)^{\mathbb{E}}$ as bi-infinite strings over $\{0, 1\}$ with a $^{\uparrow}$ marker at position 0. The differences between descriptions of L and $(\mathbb{Z}/2)^{\mathbb{E}}$ is that elements of L also have a $_{\curlywedge}$ marker, but on the other hand are almost everywhere 0.

We shall write, here and throughout this section, $\Omega_{\leftarrow}$ for the subshift of $\mathbb{B}^L$ defined by $\Pi_{\leftarrow}$. (As the astute reader may have guessed, there will soon be a $\Pi_{\rightarrow}$ and $\Omega_{\rightarrow}$.)

Lemma 6.2. *For $t \in (\mathbb{Z}/2)^{\mathbb{E}}$ define $\eta(t) \in \mathbb{B}^L$ as follows:*

$$\textit{for } g = (s, n) \in L\!:\;\; \eta(t)_g \iff s_{>n} = t_{>n}.$$

Then the correspondence $t \mapsto \eta(t)$ is an L-equivariant, surjective map $(\mathbb{Z}/2)^{\mathbb{E}} \twoheadrightarrow \Omega_{\leftarrow}$. It collapses $\{t \in (\mathbb{Z}/2)^{\mathbb{E}} : t_{>0}$ is infinitely supported$\}$ to the point $\bot^L$, and is injective on its complement, so $t \mapsto \eta(t)$ presents $\Omega_{\leftarrow}$ as

$$\begin{aligned}\Omega_{\leftarrow} &= \big\{\eta(t) : t \in (\mathbb{Z}/2)^{\mathbb{E}}\big\}\\ &= (\mathbb{Z}/2)^{\mathbb{E}} / \big(t \sim t' \text{ if both } t_{>0} \text{ and } t'_{>0} \text{ are infinitely supported}\big).\end{aligned}$$

The subshift $\Omega_{\leftarrow}$ is almost minimal.

Note that the mapping $t \mapsto \eta(t)$ is not continuous for the Cantor topology on $(\mathbb{Z}/2)^{\mathbb{E}}$, but it *is* continuous when $(\mathbb{Z}/2)^{\mathbb{E}}$ has the topology with basis

$$\big\{y \in (\mathbb{Z}/2)^{\mathbb{E}} : y_{>n} = x_{>n}\big\} \quad \text{for } n \in \mathbb{Z} \text{ and } x \text{ ranging over } (\mathbb{Z}/2)^{\mathbb{E}}.$$

We may thus compute $\eta(t)$ as $\lim \eta(t_i)$ as long as the t_i agree with t on ever larger right-infinite intervals.

Proof. We first show that indeed $\eta(t)$ belongs to $\Omega_\leftarrow$ for every $t \in (\mathbb{Z}/2)^{\mathbb{E}}$. If $\eta(g)$, then $g = (s, n)$ with $s_{>n} = t_{>n}$. There is then a unique $h \in \{a^{-1}, b^{-1}\}$ such that $\eta(gh)$, namely $h = a^{-1}$ if $s_{n-1/2} = t_{n-1/2}$ and $h = b^{-1}$ otherwise. On the other hand, $\eta(g)$ implies $\eta(ga)$ and $\eta(gb)$ because $ga = (s, n+1)$ and $gb = (s', n+1)$ with $s'_{>n+1} = s_{>n+1} = t_{>n+1}$; so both equations defining $\Omega_\leftarrow$ are satisfied.

We next prove that $t \mapsto \eta(t)$ has image $\Omega_\leftarrow$. Consider $\eta \in \Omega_\leftarrow$; if $\eta = \bot^L$, choose any $t \in (\mathbb{Z}/2)^{\mathbb{E}}$ with infinitely many 1s in its right tail. Otherwise, let $g \in L$ be in the support of η. By the rule '$\gamma \vee \delta \implies \alpha \neq \beta$', precisely one of ga^{-1} and gb^{-1} is in the support of η. Following this path from g, we get a sequence of symbols $w_1, w_2, \ldots \in \{a^{-1}, b^{-1}\}$ such that $gw_1 \cdots w_n$ is in the support of η for all n. The sequence $gw_1w_2\cdots$ converges to a configuration $t \in (\mathbb{Z}/2)^{\mathbb{E}}$. Observe that the head only moves to the left, so $t_i = 0$ for all i large enough.

We show that the t we just constructed is indeed a preimage of η. Consider first some $h = (s, n) \in L$ such that $s_{>n} = t_{>n}$; we prove $\eta(h)$. By definition of t we have $\eta(gw_1 \cdots w_m)$ for some $w_i \in \{a^{-1}, b^{-1}\}^*$ and m arbitrarily large, with $gw_1 \cdots w_m = (t', n')$ and $t'_{>n'} = t_{>n'}$; so $t'_{>n} = t_{>n}$ as soon as $n' \leq n$. By the rule '$\alpha \vee \beta \implies \gamma \wedge \delta$' we have $\eta(gw_1 \cdots w_m u)$ for all $u \in \{a, b\}^*$. If furthermore m is large enough that n' is smaller than all elements in the support of s, then $h \in gw_1 \cdots w_m\{a, b\}^*$ implying $\eta(h)$.

Consider next $h = (s, n) \in L$ such that $s_{>n} \neq t_{>n}$, so $s_i \neq t_i$ for some $i > n$. Define $s' \in (\mathbb{Z}/2)^{\mathbb{E}}$ by $s'_{<n} := s_{<n}$ and $s'_{>n} := t_{>n}$; so $s'_i \neq t_i$ as before. By the rule '$\alpha \vee \beta \implies \gamma \wedge \delta$' we have $\eta(c(s', m))$ for all $m \geq n$. Furthermore, $h \in (s', m)\{a^{-1}, b^{-1}\}^{m-n}$ for all m large enough. Now by induction, using the rule '$\gamma \vee \delta \implies \alpha \neq \beta$', whenever $\eta(f)$ there is for all $m \geq 0$ a unique $w \in \{a^{-1}, b^{-1}\}^m$ such that $\eta(fw)$. Thus we cannot have simultaneously $\eta(s', n)$ and $\eta(s, n)$, so $\neg\eta(h)$.

Consider $t \in (\mathbb{Z}/2)^{\mathbb{E}}$. If $t_{>0}$ is infinitely supported, then $\eta(t) = \bot^L$ since $t_{>n}$ can never agree with $s_{>n}$ for $(s, n) \in L$. Assume then that $t_{>0}$ is finitely supported. Defining $s \in (\mathbb{Z}/2)^{\mathbb{E}}$ by $s_{<0} = 0$ and $s_{>0} = t_{>0}$, we get $\eta(t)_{s,0}$, so $\eta(t) \neq \bot^L$. Consider next $t' \neq t$ such that $t'_{>0}$ is also finitely supported, and let n be such that $t'_{n+1/2} \neq t_{n+1/2}$. Define $(s, n) \in L$ by $s_{<n} = 0$ and $s_{>n} = t_{>n}$; then $\eta(t) \neq \eta(t')$ because

$$\eta(t)_{(s,n)} \implies s_{>n} = t_{>n} \implies s_{>n} \neq t'_{>n} \implies \neg\eta(t')_{(s,n)}.$$

The L-equivariance of η is easily checked; it amounts to checking $a\eta(t) = \eta(t')$ with $t'(i) = t(i-1)$ and $ab^{-1}\eta(t) = \eta(t'')$ with $t''_i = t_i$ for $i \neq 1/2$ and $t''_{1/2} = 1 - t_{1/2}$.

We finally prove that $\Omega_\leftarrow$ is almost minimal. If $\eta(t), \eta(u)$ are different from $\bot^L$, then for every $m \in \mathbb{N}$ we can find $g \in L$ such that $g\eta(t) = \eta(t')$ with $t'_{>-m} = u_{>-m}$. In this manner we can make an arbitrarily large central portion of $\eta(t')$ equal to that of $\eta(u)$, so $L\eta(t)$ approaches $\eta(u)$ arbitrarily closely. The fixed point $0 = \bot^L$ is approached as a limit of $\eta(t)$ with $t_{>0}$ having support of size $\to \infty$. ■

We symmetrically define the SFT $\Omega_\rightarrow$ by switching the roles of left and right; so $\Omega_\rightarrow$ is also almost minimal, and we have

$$\Omega_\rightarrow = \{\eta \in \mathbb{B}^L : \exists t \in (\mathbb{Z}/2)^{\mathbb{E}}\colon\ \forall g = (s, n) \in L\colon\ \eta(g) \iff s_{<n} = t_{<n}\}.$$

We next combine these two SFTs by a product construction:

$$\Omega_{\leftrightarrow} \subset \Omega_{\leftarrow} \times \Omega_{\rightarrow}$$

consists in configurations η such that, writing

$$\big(\eta(g), \eta(gab^{-1}), \eta(ga), \eta(gb)\big) = (\alpha, \beta, \gamma, \delta),$$

we have $\alpha = (\top, \top) \iff \gamma = (\top, \top)$.

Lemma 6.3. *The SFT $\Omega_{\leftrightarrow}$ is the orbit closure of $\eta \in (\mathbb{B} \times \mathbb{B})^L$ defined by*

$$\eta(s, n) = (s_{>n} \equiv 0^{>n}, s_{<n} \equiv 0^{<n}).$$

The configuration η is the only configuration in $\Omega_{\leftrightarrow}$ satisfying $\eta(1) = (\top, \top)$.

Proof. The configuration η belongs to $\Omega_{\leftrightarrow}$: its first projection is in $\Omega_{\leftarrow}$ by Lemma 6.2 with $t = 0^{\mathbb{E}}$, and symmetrically its right projection is in $\Omega_{\rightarrow}$. On the other hand $\eta(g) = (\top, \top)$ happens precisely when $g = (0^{\mathbb{E}}, n)$ for some $N \in \mathbb{Z}$, so $g \in \langle a \rangle$.

Consider now an arbitrary configuration $\zeta \in \Omega_{\leftrightarrow}$. Suppose first $\zeta = (\perp, \perp)^L$; then ζ is in the orbit closure of η, since arbitrarily large $(\perp, \perp)$-balls are seen around elements of the form (s, n) in which s has many 1s in its left and right tails. The set of configurations $\zeta \in \Omega_{\leftrightarrow}$ with second projection $\perp^L$ is precisely $\Omega_{\leftarrow} \times \{\perp^L\}$. To reach these configurations in the orbit closure of η it suffices to find only one of them, by the almost minimality of $\Omega_{\leftarrow}$. Now clearly if $s_{>0} = 0^{>0}$ and $s_{-m-1/2} \neq 0$ for some $m \geq 0$, then $((s, 0) \cdot \eta)_1 = (\top, \perp)$ and the second projection of $(s, 0) \cdot \eta$ tends to $\perp^L$ as $m \to \infty$. Thus indeed $\Omega_{\leftarrow} \times \{\perp^L\}$ is contained in the orbit closure of η. Similarly, $\{\perp^L\} \times \Omega_{\rightarrow}$ is contained in the orbit closure of η.

Consider then $\zeta \in \Omega_{\leftrightarrow}$ whose projections are both $\neq \perp^L$. First, we have $\zeta(g) = (\top, \top)$ for some $g \in L$: indeed, suppose $\zeta(g_1) = (\top, \perp)$ and $\zeta(g_2) = (\perp, \top)$. Then the rules force $\zeta(g_1 v_1) = (\top, *)$ for all $v_1 \in \{a, b\}^*$ and $\zeta(g_2 v_2) = (*, \top)$ for some $v_2 \in \{a, b\}^m$ and any $m \geq 0$; and symmetrically $\zeta(g_1 v_1 w_1) = (\top, *)$ for some $w_1 \in \{a^{-1}, b^{-1}\}^n$ and all $n \geq 0$ while $\zeta(g_2 v_2 w_2) = (*, \top)$ for all $w_2 \in \{a^{-1}, b^{-1}\}^*$. Now every element of L, in particular $g_1^{-1} g_2$, may be written in the form $v_1 w_1 w_2^{-1} v_2^{-1}$ with v_2, w_1 fixed, as soon as they are long enough (depending on the support of g_1, g_2); so we may set $g = g_1 v_1 w_1 = g_2 v_2 w_2$ and note $\zeta(g) = (\top, \top)$. By replacing ζ by a translate, we may assume $\zeta(1) = (\top, \top)$ and it now enough to prove $\zeta = \eta$. By the rule '$\alpha = (\top, \top) \iff \gamma = (\top, \top)$', both η and ζ contain $(\top, \top)$ only on the subgroup $\langle a \rangle$. Now any configuration in $\Omega_{\leftarrow}$ (respectively $\Omega_{\rightarrow}$) is determined by a bi-infinite $\{a, b\}$-labelled path of $\top$s, by Lemma 6.2, so $\zeta = \eta$.

The last claim holds because $a \cdot \eta = \eta$ and η contains $(\top, \top)$ only on the coset $\langle a \rangle$. ■

Generalizing the sunny-side-up shift, it would be interesting to understand *which subgroups can be marked by a sofic shift*. By this we mean the following: consider a countable group G and a subgroup $H \leq G$; then G acts on the space $G/H = \{gH : g \in G\}$ of left cosets of H by left translation, and this action extends to the one-point compactification

$G/H \cup \{\infty\}$ giving it the structure of an expansive zero-dimensional topological dynamical system. Thus it is abstractly a subshift, which we call the *H-coset subshift*. When H is the trivial group, the H-coset subshift is the sunny-side-up; it is still not understood for which groups G it is sofic. Now mapping $(\top, \top)$ to 1 and everything else to 0 produces a subshift of $\{0,1\}^L$, for which we have the following result.

Proposition 6.4. *The $\langle a\rangle$-coset subshift on the lamplighter group is sofic.* ■

It is straightforward to superpose any sofic $\mathbb{Z}$-shift on the $\langle a\rangle$-coset, giving also the following corollary:

Proposition 6.5. *The sunny-side-up subshift on the lamplighter group is sofic.* ■

6.2. Unsynchronized binary trees over the sea surface

We define in this section a "sea level" SFT: it will mark one level of $\mathcal{L}$, namely all $(s,n) \in L$ for some fixed $n \in \mathbb{Z}$, by a "sea level" symbol $\approx$, and mark by "above (respectively below) sea level" symbols all elements (s,n') with $n' > n$ (respectively $n' < n$). The "sea level" appears as a grid in Figure 5.

Additionally, the SFT will mark some binary trees in the "above sea level" portion, that connect columns of the grid together, as well as binary trees in the "below sea level" portion connecting rows of the grid. The SFT will be defined by allowed tetrahedra; we introduce

$$U = \{\nwarrow, \uparrow, \nearrow\}, \quad D = \{\swarrow, \downarrow, \searrow\}, \quad A = U \cup \{\approx\} \cup D,$$

define $\phi \colon A \to \{1,0,-1\}$ by

$$\phi(U) = \{1\}, \quad \phi(\approx) = 0, \quad \phi(D) = \{-1\},$$

and define

$$\Omega_{\approx} \subset A^L$$

as those $\eta \in A^L$ such that, for all $g \in L$, the tetrahedron

$$(\eta(g), \eta(gab^{-1}), \eta(ga), \eta(gb)) = (\alpha, \beta, \gamma, \delta)$$

satisfies

$$\phi(\alpha) = \phi(\beta) \quad \text{and} \quad \phi(\gamma) = \phi(\delta), \tag{$\approx$.1}$$

$$\phi(\gamma) - \phi(\alpha) \in \{0,1\}, \tag{$\approx$.2}$$

$$\alpha = \approx \implies \{\gamma, \delta\} = \{\nwarrow, \nearrow\}, \tag{$\approx$.3}$$

$$\gamma = \approx \implies \{\alpha, \beta\} = \{\swarrow, \searrow\}, \tag{$\approx$.4}$$

$$\alpha \in U \implies \beta = \alpha \wedge \{\gamma, \delta\} = \{\uparrow, \alpha\}, \tag{$\approx$.5}$$

$$\gamma \in D \implies \delta = \gamma \wedge \{\alpha, \beta\} = \{\downarrow, \gamma\}. \tag{$\approx$.6}$$

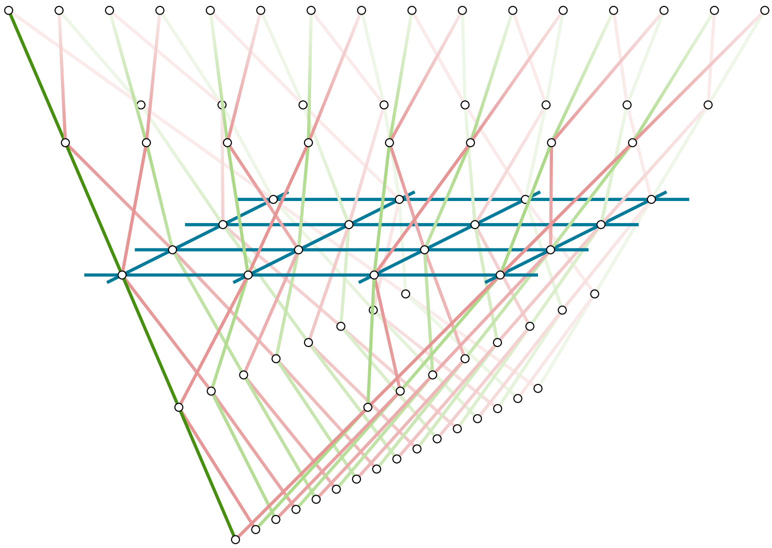

Figure 5. A tetrahedron in $\mathcal{L}$, with in blue the "sea level" grid and in heavy green the marked ray $\langle a\rangle$.

For a finite word v and an infinite (right or left) word u, we write $v \sqsubset u$ if v is at the extremity (prefix, suffix) of u. We recall our notation '$u_\downarrow v^\uparrow w$' for elements of the lamplighter group, introduced in Section 4; in particular the identity is written '$\updownarrow$', and '$u\updownarrow v$' corresponds to elements $(s,0)$. This will be the most convenient notation to describe the structure of $\Omega_{\approx\approx}$.

The following result shows that, in every tiling respecting ($\approx\approx$.1)–($\approx\approx$.6), and containing the symbol $\approx\approx$, there is a "sea level", namely an infinite grid of $\approx\approx$'s; and some binary trees attached above and below it, in directions specified by some rays $s_{u,\nwarrow}$, $s_{u,\nearrow}$, $s_{u,\swarrow}$, $s_{u,\searrow}$:

Lemma 6.6. *Consider $\eta \in A^L$ and suppose $\eta(\updownarrow) = \approx\approx$. Then $\eta \in \Omega_{\approx\approx}$ if and only if the following holds:*

- $\eta(u\updownarrow v) = \approx\approx$ *for all* $u, v \in (\mathbb{Z}/2)^*$*;*
- *for all* $u \in (\mathbb{Z}/2)^*$ *there exist* $s_{u,\nwarrow}$ *and* $s_{u,\nearrow}$ *in* $(\mathbb{Z}/2)^{\mathbb{N}}$ *with* $(s_{u,\nwarrow})_0 \neq (s_{u,\nearrow})_0$*, and* $s_{u,\swarrow}$ *and* $s_{u,\searrow}$ *in* $(\mathbb{Z}/2)^{-\mathbb{N}}$ *with* $(s_{u,\swarrow})_0 \neq (s_{u,\searrow})_0$*, such that*

$$\begin{aligned}
\eta(u_\downarrow v^\uparrow w) &= \nwarrow && \text{if } |v| > 0 \text{ and } v \sqsubset s_{u,\nwarrow},\\
\eta(u_\downarrow v^\uparrow w) &= \nearrow && \text{if } |v| > 0 \text{ and } v \sqsubset s_{u,\nearrow},\\
\eta(u_\downarrow v^\uparrow w) &= \uparrow && \text{if } |v| > 0 \text{ and the previous two cases do not apply},\\
\eta(w^\uparrow v_\downarrow u) &= \swarrow && \text{if } |v| > 0 \text{ and } v \sqsubset s_{u,\swarrow},\\
\eta(w^\uparrow v_\downarrow u) &= \searrow && \text{if } |v| > 0 \text{ and } v \sqsubset s_{u,\searrow},\\
\eta(w^\uparrow v_\downarrow u) &= \downarrow && \text{if } |v| > 0 \text{ and the previous two cases do not apply.}
\end{aligned}$$

A configuration η with $\eta(\updownarrow) = ≋$ is fully determined by the collection, for all $u \in (\mathbb{Z}/2)^*$, of the words $s_{u,\nwarrow}, s_{u,\nearrow} \in (\mathbb{Z}/2)^{\mathbb{N}}$ and $s_{u,\swarrow}, s_{u,\searrow} \in (\mathbb{Z}/2)^{-\mathbb{N}}$.

Proof. It is straightforward to verify that a configuration η satisfying the above for some choices $s_{u,\nwarrow}, s_{u,\nearrow} \in (\mathbb{Z}/2)^{\mathbb{N}}$ and $s_{u,\swarrow}, s_{u,\searrow} \in (\mathbb{Z}/2)^{-\mathbb{N}}$ belongs to $\Omega_{≋}$.

Suppose now $\eta(u\updownarrow) = ≋$ for some $u \in (\mathbb{Z}/2)^*$. Rules (≋.3) and (≋.5) imply, by induction on n, that for some choice $c \in \mathbb{Z}/2$ the pattern $\eta\restriction u_{\curlywedge}c(\mathbb{Z}/2)^{n\uparrow}$ contains exactly one $\nwarrow$ and otherwise only $\uparrow$'s, and $\eta\restriction u_{\curlywedge}(1-c)(\mathbb{Z}/2)^{n\uparrow}$ contains exactly one $\nearrow$ and otherwise only $\uparrow$'s.

We then show by induction on k that $\eta\restriction u_{\curlywedge}c(\mathbb{Z}/2)^{n-k\uparrow}w$ also contains exactly one $\nwarrow$ and otherwise only $\uparrow$ for all $w \in (\mathbb{Z}/2)^k$. First, observe that ≋ cannot appear in this set, as it would imply a $\nearrow$ in $\eta\restriction u_{\curlywedge}c(\mathbb{Z}/2)^{n\uparrow}$. Thus, by (≋.2), no elements of D can appear either. But as long as all symbols in $\eta\restriction u_{\curlywedge}c(\mathbb{Z}/2)^{n-k\uparrow}w$ are in U, it is clear from (≋.5) that in fact the symbol at $\eta(u_{\curlywedge}cv^{\uparrow}w)$ for $v \in (\mathbb{Z}/2)^{n-k}$ is uniquely determined by the counts of symbols $\nwarrow, \uparrow, \nearrow$ in $\eta\restriction u_{\curlywedge}cv^{\uparrow}w\cdot\{a,b\}^k = \eta\restriction u_{\curlywedge}cv(\mathbb{Z}/2)^{k\uparrow}$. We conclude that the symbol $\eta(u_{\curlywedge}cv^{\uparrow}w)$ is independent of w, and a symmetric claim holds for $\eta\restriction u_{\curlywedge}(1-c)v^{\uparrow}w$.

As we observed two paragraphs above, $\eta(u_{\curlywedge}c(\mathbb{Z}/2)^{n\uparrow})$ contains for all $n \in \mathbb{N}$ exactly one $\nwarrow$ and otherwise only $\uparrow$'s. More precisely, (≋.5) implies that there is a unique path $s_{u,\nwarrow}$ such that $\eta(u_{\curlywedge}cv^{\uparrow}) = \nwarrow$ precisely when $v \sqsubset s_{u,\nwarrow}$. The previous paragraph then shows for all $w \in (\mathbb{Z}/2)^*$ that $\eta(u_{\curlywedge}cv^{\uparrow}w) = \nwarrow$ holds precisely when $v \sqsubset s_{u,\nwarrow}$, with $\eta(u_{\curlywedge}cv^{\uparrow}w) = \uparrow$ otherwise. Symmetrically, there is a unique path $s_{u,\nearrow}$ such that $\eta(u_{\curlywedge}(1-c)v^{\uparrow}w) = \nearrow$ if and only if $v \sqsubset s_{u,\nearrow}$, with $\eta(u_{\curlywedge}cv^{\uparrow}w) = \uparrow$ otherwise. Clearly only (≋.3) can apply at $u\updownarrow w$, and we get $\eta(u\updownarrow w) = ≋$ for all $w \in (\mathbb{Z}/2)^*$.

The rules for $\swarrow, \downarrow, \searrow$ are symmetric to those for $\nwarrow, \uparrow, \nearrow$, so if $\eta(\updownarrow u) = ≋$, then there exist $s_{u,\swarrow}, s_{u,\searrow}$ differing at 0 such that $\eta(w^{\uparrow}v_{\curlywedge}u) = \searrow$ if and only if $v \sqsubset s_{u,\searrow}$ and $\eta(w^{\uparrow}v_{\curlywedge}u) = \swarrow$ if and only if $v \sqsubset s_{u,\swarrow}$, with $\eta(w^{\uparrow}v_{\curlywedge}u) = \downarrow$ for all other non-empty v, and thus $\eta(w^{\uparrow}{}_{\curlywedge}u) = ≋$ for all $w \in (\mathbb{Z}/2)^*$.

Now suppose $\eta(\updownarrow) = ≋$. From the above we obtain $\eta(u\updownarrow v) = ≋$ for all $u, v \in (\mathbb{Z}/2)^*$; and the previous analysis applied to $\eta(u\updownarrow) = ≋$ and $\eta(\updownarrow v) = ≋$ proves that η has the claimed form.

Finally, the claim that η is fully determined by $s_{u,\nwarrow}, s_{u,\nearrow} \in (\mathbb{Z}/2)^{\mathbb{N}}$ and $s_{u,\swarrow}, s_{u,\searrow} \in (\mathbb{Z}/2)^{-\mathbb{N}}$ directly follows from the given construction of the values of η. ■

Just as for Lemma 6.3, the previous lemma can be stated in terms of coset subshifts. Indeed, mapping ≋ to 1 and everything else to 0, we obtain the following:

Proposition 6.7. *Let $\phi\colon L \to \mathbb{Z}$ be the homomorphism given by $\phi(a) = \phi(b) = 1$. Then the* $\ker\phi$*-coset subshift on the lamplighter group is sofic.* ■

More generally, let G be a group, let $\phi\colon G \to H$ be a homomorphism, and let $X \subset A^H$ be an H-subshift. The *pullback* of X along ϕ is the subshift

$$\phi^*(X) = \{y \in A^G : \exists x \in X\colon\ \forall g \in G\colon\ y_g = x_{\phi(g)}\}.$$

Rephrasing Proposition 6.7 in these terms, we get the following:

Proposition 6.8. *Let $\phi\colon L \to \mathbb{Z}$ be the homomorphism given by $\phi(a) = \phi(b) = 1$. Then the pullback of the sunny-side-up subshift from $\mathbb{Z}$ along ϕ is sofic.* ∎

We do not know whether the pullback of a full shift on $\mathbb{Z}$ is sofic.

6.3. Synchronizing the trees

We now impose some extra conditions on $\Omega_{≋}$ to synchronize the marked directions, namely to force the binary trees above and below the sea level to lie in specific directions. This is done by combining $\Omega_{\leftrightarrow}$ with $\Omega_{≋}$. Define thus

$$\Omega \subset \Omega_{\leftrightarrow} \times \Omega_{≋}$$

as those $\eta \in \Omega_{\leftrightarrow} \times \Omega_{≋}$ such that, for all $g \in L$, the tetrahedron $(\eta(g), \eta(gab^{-1}), \eta(ga), \eta(gb)) = (\alpha, \beta, \gamma, \delta)$ satisfies

$$\begin{aligned}
\alpha = ((\top, *), ≋) &\implies \gamma = (*, \nwarrow) \wedge \delta = (*, \nearrow),\\
\alpha = ((\top, *), \nwarrow) &\implies \gamma = (*, \nwarrow),\\
\alpha = ((\top, *), \nearrow) &\implies \delta = (*, \nearrow),\\
\gamma = ((*, \top), ≋) &\implies \alpha = (*, \swarrow) \wedge \beta = (*, \searrow),\\
\gamma = ((*, \top), \swarrow) &\implies \alpha = (*, \swarrow),\\
\gamma = ((*, \top), \searrow) &\implies \beta = (*, \searrow).
\end{aligned}$$

This combination of $\Omega_{\leftrightarrow}$ and $\Omega_{≋}$ vastly reduces the size of the SFT Ω; more precisely:

Lemma 6.9. *There is a unique configuration $\eta \in \Omega$ satisfying $\eta(\updownarrow) = ((\top, \top), ≋)$. The projection to $\Omega_{\leftrightarrow}$ of η is the one described by Lemma 6.3, and its projection to $\Omega_{≋}$ is given by the choices*

$$s_{u,\nwarrow} = 0^{\mathbb{N}}, \quad s_{u,\nearrow} = 1^{\mathbb{N}}, \quad s_{u,\swarrow} = 0^{-\mathbb{N}}, \quad s_{u,\searrow} = 1^{-\mathbb{N}}, \tag{6.1}$$

namely the following holds:

$$\eta(u_{\downarrow}v^{\uparrow}w) = \begin{cases} (*, \nwarrow) & \textit{if } v \in 0^+,\\ (*, \nearrow) & \textit{if } v \in 1^+,\\ (*, \uparrow) & \textit{in all remaining cases}, \end{cases}$$

$$\eta(u{\updownarrow}v) = ((\top, \top), ≋),$$

$$\eta(w^{\uparrow}v_{\downarrow}u) = \begin{cases} (*, \swarrow) & \textit{if } v \in 0^+,\\ (*, \searrow) & \textit{if } v \in 1^+,\\ (*, \downarrow) & \textit{in all remaining cases}. \end{cases}$$

Proof. We first show that the given η belongs to Ω. Our choices in (6.1) for all applicable u show that the second projection of η satisfies the characterization of Lemma 6.6. With

$(\eta(g), \eta(gab^{-1}), \eta(ga), \eta(gb)) = (\alpha, \beta, \gamma, \delta)$, we have $\alpha = ((\top, *), *)$ if and only if $g = u_{\uparrow} v^{\upharpoonright}$ or $g = u^{\upharpoonright} 0^n{}_{\uparrow}$ for some u, v, n, since the first projection of η is the configuration described by Lemma 6.3.

We now show that the rules joining $\Omega_{\leftrightarrow}$ and $\Omega_{\approx\!\approx}$ are satisfied by η. If $\alpha = (*, \approx\!\approx)$, then $g = u \updownarrow v$ for some u, v; force $|v| > 0$ by appending 0 if needed. If furthermore $\alpha = ((\top, *), \approx\!\approx)$, then by Lemma 6.3 we may write $v = 0v'$ since $v = 0^{|v|}$, and then indeed we have a '$\nwarrow$' at $(u \updownarrow v) \cdot a = u_{\uparrow} 0^{\upharpoonright} v'$ and a '$\nearrow$' at $(u \updownarrow v) \cdot b = u_{\uparrow} 1^{\upharpoonright} v'$, as required. The verifications when $\alpha = (*, \nwarrow)$ or $(*, \nearrow)$ are similar: by Lemma 6.3 if $\alpha = ((\top, *), *)$, then the bit after the head in g must be 0, and therefore a '$\nearrow$' (respectively a '$\nwarrow$') appears in the correct neighbour. The verifications for $\searrow$ and $\swarrow$ are symmetric because $\alpha = ((*, \top), *)$ implies that the bit to the left of the head is 0.

Next, we show that η is the unique configuration in Ω seeded at $\updownarrow$. Let $\zeta \in \Omega$ be any configuration with $\zeta(\updownarrow) = ((\top, \top), \approx\!\approx)$. By definition of Ω, the first projection of ζ is the unique configuration of $\Omega_{\leftrightarrow}$ described in Lemma 6.3, and its second projection is one of the configurations of $\Omega_{\approx\!\approx}$ described by Lemma 6.6, say by words $s_{u,\nwarrow}, s_{u,\nearrow} \in (\mathbb{Z}/2)^{\mathbb{N}}$ and $s_{u,\swarrow}, s_{u,\searrow} \in (\mathbb{Z}/2)^{-\mathbb{N}}$. We claim that these words must be the ones defining η, which will imply $\zeta = \eta$ because by Lemma 6.6 these words fully determine the configuration.

Suppose for example $s_{u,\nwarrow} \neq 0^{\mathbb{N}}$ for some $u \in (\mathbb{Z}/2)^*$. This means $\zeta(u_{\uparrow} v^{\upharpoonright}) = \nwarrow$ for some $|v| > 0$ and $v \not\sqsubset 0^{\mathbb{N}}$. Let v be of minimal length with this property, so $\zeta(u_{\uparrow} 0^k 1^{\upharpoonright}) = \nwarrow$ for some $k \geq 0$. We then have $\zeta(u_{\uparrow} 0^{k\upharpoonright}) \in \{\approx\!\approx, \nwarrow\}$, so in its first projection $\zeta(u_{\uparrow} 0^{k\upharpoonright}) = ((\top, *), *)$; then the additional rules of Ω require $\zeta(u_{\uparrow} 0^k 0^{\upharpoonright}) = (*, \nwarrow)$. Now this contradicts the defining rules of $\Omega_{\approx\!\approx}$ at $u_{\uparrow} 0^{k\upharpoonright}$ since neither $\approx\!\approx$ nor $\nwarrow$ can have $\nwarrow$ at both its a- and the b-neighbour. Thus $s_{u,\nwarrow} = 0^{\mathbb{N}}$ is the only possibility. The verifications for $s_{u,\nearrow} = 1^{\mathbb{N}}$, $s_{u,\swarrow} = 0^{-\mathbb{N}}$ and $s_{u,\searrow} = 1^{-\mathbb{N}}$ are symmetric. ∎

6.4. Simulating $\mathbb{N} \times \mathbb{N}$ on $\mathcal{L}$ marked by Ω

Consider the graph $\mathcal{L}$ marked by the SFT Ω defined in the previous sections; more precisely, imposing the seed constraint that the origin in $\mathcal{L}$ is labelled $((\top, \top), \approx\!\approx)$, we have by Lemma 6.9 uniquely specified a vertex labelling of $\mathcal{L}$. Its edges retain the Cayley graph labelling by $\{a^{\pm 1}, b^{\pm 1}\}$. We claim:

Proposition 6.10. *The graph $\mathcal{L}$ labelled by the configuration $\eta \in \Omega$ from Lemma 6.9 simulates $\mathbb{N} \times \mathbb{N}$.*

Proof. It suffices to produce a simulator; one is shown in simplified form in Figure 6. The grid $\mathbb{N}^2$ is represented by elements of $\mathcal{L}$ at sea level: given natural numbers m, n, write $m = m_k \dots m_0$ and $n = n_\ell \dots n_0$ in base 2; then the grid point $(m, n) \in \mathbb{N}^2$ is represented by the element $n_\ell \dots n_0 \updownarrow m_0 \dots m_k \in L$.

The transformation $(m, n) \mapsto (m + 1, n)$, corresponding to the generator '$\to$' of $\mathbb{N}^2$, is thus realized by adding 1 with carry to the word on the right of the '$\updownarrow$' mark, and similarly $(m, n) \mapsto (m, n + 1)$ corresponds to the generator '$\uparrow$' and to adding one with carry to the word on the left of the '$\updownarrow$' mark.

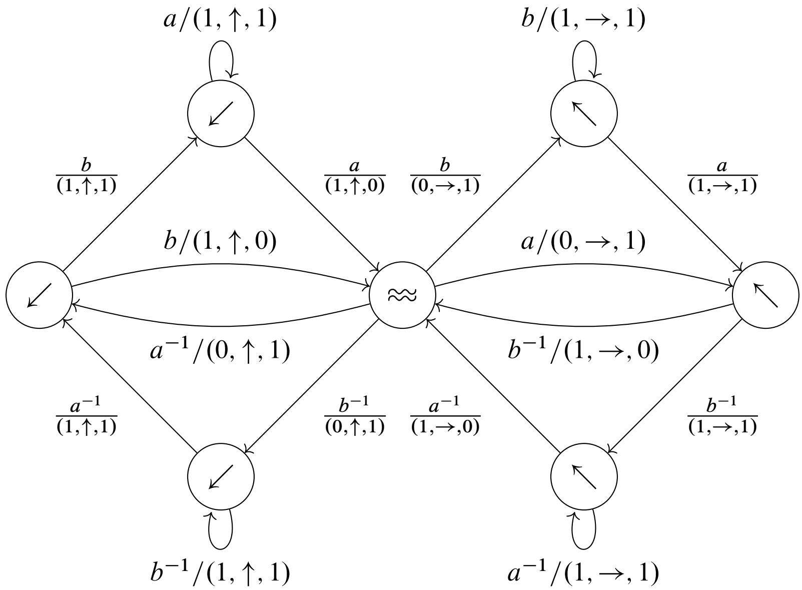

Figure 6. A simulator for $\mathbb{N} \times \mathbb{N}$ in $\mathcal{L}$.

The markings imposed by Ω on $\mathcal{L}$, see Lemma 6.9, force every row of the grid to have a distinguished binary tree marked by the symbol '$\nwarrow$' in the region above it, and force every column of the grid to have a binary tree marked '$\swarrow$' in the region below it. Furthermore, the trees are "synchronized" by $\Omega_{\leftrightarrow}$ in such a way that the path starting from a vertex (m, n) in the grid and going upwards while remaining in the '$\nwarrow$' marked region follows a sequence in $\{a, b\}$ reading the binary expansion of m, and similarly the path going downwards while following '$\swarrow$' follows a sequence in $\{a^{-1}, b^{-1}\}$ reading the binary expansion of n.

The operation of adding one with carry on the word to the right of the mark is therefore realized by following the regular expression $b^*ab^{-1}(a^{-1})^*$, which is precisely the loop followed on the right half of Figure 6; and symmetrically the operation of adding one to the left of the mark is realized by the regular expression $(b^{-1})^*a^{-1}ba^*$, which is the loop followed on the left half of Figure 6.

To obtain a *bona fide* simulator for $\mathbb{N}^2$, three modifications are necessary: Firstly, we have only written the operations $\uparrow$ and $\to$; adding reverse edges gives the operations $\downarrow$ and $\leftarrow$.

Secondly, the simulator in Figure 6 ignores the first two symbols (in $\{\top, \bot\}$) given to $\mathcal{L}$ by Ω. Thus, we should take four copies of the diagram, for all possibilities of elements in $\{\top, \bot\}^2$, and connect them by complete bipartite graphs (namely replace every edge $s \to t$ by sixteen edges for all choices of $\{\top, \bot\}^2$ at source and range).

Thirdly, the simulator in Figure 6 does not recognize the vertex markings of $\mathbb{N}^2$; so again we should take four copies of the simulator, one for each of $\{\llcorner, \vdash, \perp, +\}$, and connect them appropriately: four copies of the '$\uparrow$' circuit on the right of Figure 6 go from $\llcorner$ to $\vdash$, from $\perp$ to $+$, and loop at $\vdash$ and $+$, while four copies of the '$\rightarrow$' circuit on the left of Figure 6 go from $\llcorner$ to $\perp$, from $\vdash$ to $+$, and loop at $\perp$ and $+$. This has the effect that the '$\downarrow$' and/or '$\leftarrow$' operations are unavailable at $\vdash, \perp, \llcorner$; and indeed the regular expression $a^*ba^{-1}(b^{-1})^*$ does not match at $(0, n)$, since it remains stuck on the a^* loop which it reads infinitely. See the next section for another remedy.

It would be tedious to draw the complete simulator, but we have made it available in ancillary computer files, in the language Julia, using which SFTs on the lamplighter group can be explored; see Section 8. ∎

6.5. From the quadrant to the plane

The simulation above implements $\mathbb{N}^2$ inside $\mathcal{L}$; negative coordinates cannot be reached because the operation $\leftarrow$, implemented by the regular expression $a^*ba^{-1}(a^{-1})^*$, reads an infinite string of a's without coming to completion. The cause of this is that we represented integers in binary, with automata implementing addition with carry, and in this notation passing from 0 to -1 causes an infinite sequence of carries.

It is of course possible to simulate $\mathbb{Z}^2$ in $\mathcal{L}$ using the fact that $\mathbb{N}^2$ simulates $\mathbb{Z}^2$ if given suitable markings, see Example 2.14. However, a simple change lets us directly simulate $\mathbb{Z}^2$ in $\mathcal{L}$.

It suffices indeed to represent $\mathbb{Z}$ differently than in usual binary: consider all infinite sequences over $\{0, 1\}$ that are cofinal to $(01)^\infty$, namely all sequences $t_0t_1\cdots$ with $t_n = n \bmod 2$ for all n large enough. Then the operations $+1$ and -1 may be performed on such expressions, with the usual rules for carrying and borrowing bits; and one never encounters an infinite sequence of carries or borrows. (Abstractly, we are working on the coset $\mathbb{Z} + 1/3$ in $\mathbb{Z}_2$.)

This may be realized by making the sequences $s_{u,\nwarrow}$ etc. slightly more complicated: we choose

$$s_{u,\nwarrow} = (01)^{\mathbb{N}}, \quad s_{u,\nearrow} = (10)^{\mathbb{N}}, \quad s_{u,\swarrow} = (01)^{-\mathbb{N}}, \quad s_{u,\searrow} = (10)^{-\mathbb{N}}.$$

We omit the details of the construction of corresponding $\Omega'_{\leftrightarrow}$ and $\Omega' \subset \Omega'_{\leftrightarrow} \times \Omega_{\approx}$, which is only slightly more complicated.

7. Diestel–Leader graphs

The lamplighter group can be seen as a special case of a *Diestel–Leader graph*. As in Section 4.1, consider $p, q \geq 2$ and two trees $\mathcal{T}_p$ and $\mathcal{T}_q$, respectively $(p+1)$-regular and $(q+1)$-regular, and endow each with a Busemann function. Let $\mathtt{DL}(p, q)$ denote their horocyclic product. This is a $(p+q)$-regular graph, endowed with a graph morphism

$h\colon \mathsf{DL}(p,q) \to \mathbb{Z}$; each vertex $v = (v_1, v_2)$ at height n has p neighbours (v_1', v_2') at height $n+1$ with v_1' being one of the p successors of v_1 in $\mathcal{T}_p$ and v_2' being the unique ancestor of v_2 in $\mathcal{T}_q$; and symmetrically q neighbours at height $n-1$.

Vertices of $\mathsf{DL}(p,q)$ may also be described by sequences as in Section 4: these are sequences with a marker at an integer position n, elements of $\{0,\dots,q-1\}$ at all half-integer positions $< n$, almost all 0, and elements of $\{0,\dots,p-1\}$ at all half-integer positions $> n$, also almost all 0.

There is also a notion of tetrahedron for these graphs: for choices of sequences $u \in \{0,\dots,q-1\}^*$ and $v \in \{0,\dots,p-1\}^*$, the associated tetrahedron has $p+q$ vertices '$u{\uparrow}i\,v$' and '$u\,j{\uparrow}v$' for all $i \in \{0,\dots,p-1\}$ and all $j \in \{0,\dots,q-1\}$, and has pq edges between them in a complete bipartite graph.

We consider $\mathsf{DL}(p,q)$ with a natural labelling: Every edge in $\mathsf{DL}(p,q)$ is labelled by $\{0,\dots,p-1\} \times \{0,\dots,q-1\}$: it joins two sequences in which the marker positions differ by 1, say '$\cdots{\uparrow}i\cdots$' and '$\cdots j{\uparrow}\cdots$', and has label (i,j). There are pq different kinds of vertices, depending on the symbols in the sequence immediately left and right of the marker; so there are pq different kinds of immediate neighbourhoods that should be specified in a vertex SFT.

We note that while this labelling is natural, this is not the Cayley graph labelling of the lamplighter group, even when $p = q = 2$. Indeed, Cayley graph labellings are vertex transitive while the labelling above has pq orbits on vertices. Moreover, a remarkable feature of Diestel–Leader graphs [10] is that for $p \neq q$ they are *not* Cayley graphs; the automorphism group of $\mathsf{DL}(p,q)$ acts transitively on vertices, but does not contain a subgroup acting simply transitively (= with trivial stabilizers). It is possible to view $\mathsf{DL}(p,q)$ as the Schreier graph of a $(p+q)$-generated group (say a free group) but, as far as we know, not in any useful way.

All the constructions from this paper work *mutatis mutandis* for subshifts of finite type on these labelled Diestel–Leader graphs.

The comb, ray, and sea level SFTs adapt easily to labelled $\mathsf{DL}(p,q)$ graphs: for instance, since $\mathsf{DL}(p,q)$ contains $\mathsf{DL}(2,2)$ as the subgraph spanned by edges labelled $\{0,1\}^2$: the tiles can be extended from $\mathsf{DL}(2,2)$ to $\mathsf{DL}(p,q)$ by simply ignoring the colours on the extra edges. Furthermore, we obtain tiling systems on the labelled graph $\mathsf{DL}(2,2)$ from tiling systems on the lamplighter group via the map $\mathsf{DL}(2,2) \to \mathcal{L}$ that sends edge labels $(0,0),(1,1)$ to a and $(0,1),(1,0)$ to b.

In the case of the sea level, there is also a natural direct construction on the Diestel–Leader graph $\mathsf{DL}(p,q)$, in which subgroups have simple geometric avatars: The '$\langle a\rangle$' subgroup may be replaced by the ray marked $(0,0)$ through the seed, and corresponds to the all-off lamp configurations with arbitrary marker position. The '$\ker(\phi)$' sea level subgroup corresponds to the lamp configurations with marker at 0. For example, the "ray subshift" $\Pi_\leftarrow$ becomes

$$\Pi_\leftarrow = \Big\{(\alpha_1,\dots,\alpha_p,\beta_1,\dots,\beta_q) \in \mathbb{B}^{p+q} : \bigvee \alpha_i \implies \bigwedge \beta_j;\ \bigvee \beta_j \implies \exists! i\colon\ \alpha_i\Big\}.$$

Theorem 7.1. *For all $p, q \geq 2$, labelled* $\mathsf{DL}(p,q)$ *has undecidable seeded tiling problem.*

One may wonder if the rigidity of the labelling of $\mathsf{DL}(p,q)$ makes its tiling problems easy (to prove undecidable), but this does not seem to be the case: the labelling is highly recurrent, so one cannot use the non-rigidity of the tiling to force a (unique) seed to appear with local rules, and we have not been able to solve the decidability of unseeded tiling problem on these graphs either.

Note that, if $\mathsf{DL}(p,q)$ is considered with the trivial labelling (we write it $\mathcal{D}$ from now on to avoid confusion), then its seeded tiling problem is decidable for uninteresting reasons in the Hom-formalism. To see this, let $\mathcal{B} = \{1,2\} \sqcup \{1,2\}^2$ be the complete graph with self-loops on two vertices $1, 2$, and let $\beta\colon \mathcal{D} \to \mathcal{B}$ be the sunny-side-up labelling marking the origin with 1. We claim that, given a finite graph $\mathcal{F} \in \mathsf{Graph}_{/\mathcal{B}}$, it is decidable whether $\mathsf{Graph}_{/\mathcal{B}}(\mathcal{D}, \mathcal{F}) = \emptyset$. Indeed, let $\mathcal{H}$ be the graph with vertex set $\{1,2,3\}$, edges both ways between i and $i+1$, and labelling induced by $\beta(1) = 1$, $\beta(2) = \beta(3) = 2$. Then $\mathsf{Graph}_{/\mathcal{B}}(\mathcal{D}, \mathcal{F}) = \emptyset$ if and only if $\mathsf{Graph}_{/\mathcal{B}}(\mathcal{H}, \mathcal{F}) = \emptyset$: given $\phi\colon \mathcal{H} \to \mathcal{F}$, lift it to $\mathcal{D}$ by mapping the origin to 1, all other vertices at even height to 3, and vertices at odd height to 2. Conversely, $\mathcal{H}$ is a subgraph of $\mathcal{D}$.

A better question is the following. Let $G = \mathrm{Aut}(\mathcal{D})$ denote the automorphism group of the unlabelled graph $\mathcal{D} = \mathsf{DL}(p,q)$. For a finite alphabet A, consider the set $A^{\mathcal{D}}$ of maps $V(\mathcal{D}) \sqcup E(\mathcal{D}) \to A$, with the natural action of G by precomposition. A subset of $A^{\mathcal{D}}$ is called an *SFT* if it is of the form $\{\eta \in A^{\mathcal{D}} : \forall g \in G\colon\ \eta \circ g \in C\}$ for some clopen subset $C \subset A^{\mathcal{D}}$. One defines seeded SFTs as in Definition 3.15, by conditioning on a sunny-side-up.

Conjecture 7.2. *The unlabelled graph* $\mathsf{DL}(p,q)$ *has undecidable seeded tiling problem.*

8. Electronic resources to manipulate SFTs on the lamplighter group

It is quite entertaining to experiment with SFTs on the lamplighter group. We have written some simple code to help in such experiments; it can be accessed by following the article's DOI.

The Julia module `LL.jl` should be loaded with '`include("LL.jl")`' in a recent Julia distribution, including the packages `Makie` (for 3D visualization) and `CryptoMiniSat` or `PicoSAT` (to compute tilings of tetrahedra using a SAT solver). Elements of L are displayed as sequences over $\{0,1\}$, with an underline or overline between the origin and the marker position: '$u_{\curlywedge}v^{\uparrow}w$' is represented as $\texttt{u}\underline{\texttt{v}}\texttt{w}$ and '$u^{\uparrow}v_{\curlywedge}w$' is represented as $\texttt{u}\overline{\texttt{v}}\texttt{w}$. A sample run could be

```
julia> include("LL.jl")
julia> root = LL.Element(0,0,3)
000
julia> seadict = LL.solve(LL.graph(6),sea,seed=[root=>1]);
```

```
julia> LL.walk(root,seaeast,seadict)
0001
julia> LL.walk(ans,seanorth,seadict)
0011
julia> LL.walk(ans,seawest,seadict)
001
julia> LL.walk(ans,seasouth,seadict)
000
```

Acknowledgements. We are grateful to the Unité de Mathématiques Pures et Appliquée (UMPA) and the Laboratoire de l'Informatique du Parallélisme (LIP) at the École Normale Supérieure de Lyon for their hospitality, and in particular to Nathalie Aubrun for her interest and support.

References

[1] N. Aubrun, S. Barbieri, and E. Jeandel, About the domino problem for subshifts on groups. In *Sequences, groups, and number theory*, pp. 331–389, Trends Math., Birkhäuser/Springer, Cham, 2018 Zbl 1405.20023 MR 3799931

[2] N. Aubrun, S. Barbieri, and E. Moutot, The domino problem is undecidable on surface groups. In *44th International Symposium on Mathematical Foundations of Computer Science*, pp. Art. No. 46, 14, LIPIcs. Leibniz Int. Proc. Inform. 138, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2019 Zbl 07561690 MR 4008435

[3] N. Aubrun and J. Kari, Tiling problems on Baumslag-Solitar groups. In *Proceedings: Machines, Computations and Universality 2013*, pp. 35–46, Electron. Proc. Theor. Comput. Sci. (EPTCS) 128, Open Publishing Association, Waterloo, 2013 Zbl 1469.20028 MR 3594268

[4] A. Ballier and M. Stein, The domino problem on groups of polynomial growth. *Groups Geom. Dyn.* **12** (2018), no. 1, 93–105 Zbl 1386.37017 MR 3781417

[5] L. Bartholdi, M. Neuhauser, and W. Woess, Horocyclic products of trees. *J. Eur. Math. Soc. (JEMS)* **10** (2008), no. 3, 771–816 Zbl 1155.05009 MR 2421161

[6] R. Berger, The undecidability of the domino problem. *Mem. Amer. Math. Soc.* **66** (1966) Zbl 0199.30802 MR 216954

[7] D. B. Cohen, The large scale geometry of strongly aperiodic subshifts of finite type. *Adv. Math.* **308** (2017), 599–626 Zbl 1400.20034 MR 3600067

[8] D. B. Cohen, Lamplighters admit weakly aperiodic SFTs, *Groups Geom. Dyn.* **14** (2020), no. 4, 1241–1252. Zbl 1477.37025 MR 4186473

[9] F. Dahmani and A. Yaman, Symbolic dynamics and relatively hyperbolic groups. *Groups Geom. Dyn.* **2** (2008), no. 2, 165–184 Zbl 1169.20022 MR 2393178

[10] R. Diestel and I. Leader, A conjecture concerning a limit of non-Cayley graphs. *J. Algebraic Combin.* **14** (2001), no. 1, 17–25 Zbl 0985.05020 MR 1856226

[11] J. Esnay and É. Moutot, Weakly and strongly aperiodic subshifts of finite type on Baumslag–Solitar groups. 2020, arXiv:2004.02534

[12] E. Jeandel, Aperiodic subshifts on polycyclic groups. 2015, arXiv:1510.02360

[13] E. Jeandel, Translation-like actions and aperiodic subshifts on groups. 2015, arXiv:1508.06419

[14] A. S. Kahr, E. F. Moore, and H. Wang, Entscheidungsproblem reduced to the ∀∃∀ case. *Proc. Nat. Acad. Sci. U.S.A.* **48** (1962), 365–377 Zbl 0102.00801 MR 169777

[15] J. Kari, Cellular automata, tilings and (un)computability. In *Combinatorics, words and symbolic dynamics*, pp. 241–295, Encyclopedia Math. Appl. 159, Cambridge University Press, Cambridge, 2016 Zbl 1383.68052 MR 3525487

[16] D. Lind and B. Marcus, *An introduction to symbolic dynamics and coding*. Cambridge University Press, Cambridge, 1995 Zbl 1106.37301 MR 1369092

[17] R. C. Lyndon and P. E. Schupp, *Combinatorial group theory*. Reprint of the 1977 edition, Class. Math., Springer, Berlin, 2001 Zbl 0997.20037 MR 1812024

[18] D. E. Muller and P. E. Schupp, Context-free languages, groups, the theory of ends, second-order logic, tiling problems, cellular automata, and vector addition systems. *Bull. Amer. Math. Soc. (N.S.)* **4** (1981), no. 3, 331–334 Zbl 0484.03019 MR 609044

[19] R. M. Robinson, Undecidability and nonperiodicity for tilings of the plane. *Invent. Math.* **12** (1971), 177–209 Zbl 0197.46801 MR 297572

[20] B. Seward, Burnside's Problem, spanning trees and tilings. *Geom. Topol.* **18** (2014), no. 1, 179–210 Zbl 1338.20041 MR 3158775

[21] H. Wang, Proving theorems by pattern recognition, II, *Bell Syst. Tech. J.* **40** (1961), no. 1, 1–41

[22] K. Whyte, Amenability, bi-Lipschitz equivalence, and the von Neumann conjecture. *Duke Math. J.* **99** (1999), no. 1, 93–112 Zbl 1017.54017 MR 1700742

[23] W. Woess, Lamplighters, Diestel-Leader graphs, random walks, and harmonic functions. *Combin. Probab. Comput.* **14** (2005), no. 3, 415–433 Zbl 1066.05075 MR 2138121

Received 1 March 2021.

Laurent Bartholdi
Fachrichtung Mathematik + Informatik, Universität des Saarlandes, 66123 Saarbrücken, Germany; laurent.bartholdi@gmail.com

Ville Salo
Department of Mathematics and Statistics, University of Turku, Turku, Finland; vosalo@utu.fi