International Journal of Foundations of Computer Science

Effective Projections on Group Shifts to Decide Properties of Group Cellular Automata

Pierre Béaur

Laboratoire Interdisciplinaire des Sciences du Numérique
Université Paris-Saclay
pierre.beaur@universite-paris-saclay.fr

Jarkko Kari

Department of Mathematics and Statistics
University of Turku, Finland
jkari@utu.fi

Received (Day Month Year)
Accepted (Day Month Year)
Communicated by (xxxxxxxxxx)

Many decision problems concerning cellular automata are known to be decidable in the case of algebraic cellular automata, that is, when the state set has an algebraic structure and the automaton acts as a morphism. The most studied cases include finite fields, finite commutative rings and finite commutative groups. In this paper, we provide methods to generalize these results to the broader case of group cellular automata, that is, the case where the state set is a finite (possibly non-commutative) finite group. The configuration space is not even necessarily the full shift but a subshift – called a group shift – that is a subgroup of the full shift on $\mathbb{Z}^d$, for any number d of dimensions. We show, in particular, that injectivity, surjectivity, equicontinuity, sensitivity and nilpotency are decidable for group cellular automata, and non-transitivity is semi-decidable. Injectivity always implies surjectivity, and jointly periodic points are dense in the limit set. The Moore direction of the Garden-of-Eden theorem holds for all group cellular automata, while the Myhill direction fails in some cases. The proofs are based on effective projection operations on group shifts that are, in particular, applied on the set of valid space-time diagrams of group cellular automata. This allows one to effectively construct the traces and the limit sets of group cellular automata. A preliminary version of this work was presented at the conference Mathematical Foundations of Computer Science 2020.

Keywords: group cellular automata; group shift; symbolic dynamics; decidability

1. Introduction

Algebraic group shifts and group cellular automata operate on configurations that are colorings of the infinite grid $\mathbb{Z}^d$ by elements of a finite group $\mathbb{G}$, called the state set. The set $\mathbb{G}^{\mathbb{Z}^d}$ of all configurations, called the full shift, inherits the group structure as the infinite cartesian power of $\mathbb{G}$. A subshift (a set of configurations avoiding a fixed set of forbidden finite patterns) is a group shift if it is also a subgroup of $\mathbb{G}^{\mathbb{Z}^d}$. Group shifts are known to be of finite type, meaning that they can be defined by forbidding a finite number of patterns. A cellular automaton is a dynamical system on a subshift, defined by a uniform local update rule of states. A cellular automaton on a group shift is called a group cellular automaton if it is also a group homomorphism.

In this work we demonstrate that group shifts and group cellular automata in arbitrarily high dimensions d are amenable to effective manipulations and algorithmic decision procedures. This is in stark contrast to the general setup of multidimensional subshifts of finite type and cellular automata where most properties are undecidable. Our considerations generalize a long line of past results – see for example [7, 8] and citations therein – on algorithms for linear cellular automata (whose state set is a finite commutative ring) and additive cellular automata (whose state set is a finite abelian group) to non-commutative groups and to arbitrary dimensions, and from the full shift to arbitrary group shifts. Our methods are based on two classical results on group shifts: all group shifts – in any dimension – are of finite type, and they have dense sets of periodic points [19, 28]. By a standard argumentation these provide a decision procedure for the membership in the language of any group shift. We show how to use this procedure to effectively construct any lower dimensional projection of a given group shift (Corollary 10), and to construct the image of a given group shift under any given group cellular automaton (Corollary 11).

To establish decidability results for d-dimensional group cellular automata we then view the set of valid space-time diagrams as a $(d+1)$-dimensional group shift. The local update rule of the cellular automaton provides a representation of this group shift. The one-dimensional projections in the temporal direction are the trace subshifts of the automaton that provide all possible temporal evolutions for a finite domain of cells, and the d-dimensional projection in the spatial dimensions is the limit set of the automaton. These can be effectively constructed. From the trace subshifts – which are one-dimensional group shifts themselves – one can analyze the dynamics of the cellular automaton and to decide, for example, whether it is periodic (Theorem 25), equicontinuous or sensitive to initial conditions (Theorem 27). There is a dichotomy between equicontinuity and sensitivivity (Lemma 26). We can semi-decide negative instances of mixing properties, i.e., non-transitive and non-mixing cellular automata (Theorem 28). The limit set reveals whether the automaton is nilpotent (Theorem 25), surjective or injective (Theorem 21). Note that all these considerations work for group cellular automata over arbitrary group shifts, not only over full shifts, and in all dimensions. We also note that in our setup injectivity implies surjectivity (Corollary 20) and that surjectivity implies pre-injectivity (Theorem 24), with neither implication holding in the inverse direction in general. Moreover, in all surjective cases jointly spatially and temporally periodic points are dense (Corollary 19).

The paper is structured as follows. We start by providing the necessary terminology and classical results about shift spaces and cellular automata; first in the general context of multidimensional symbolic dynamics and then in the algebraic setting in particular. In Section 3 we define projection operations on group shifts and exhibit effective algorithms to implement them. This involves the main technical proof of the paper. In Section 4 we apply the projections on space-time diagrams of cellular automata to effectively construct their traces and limit sets. These are then used to provide decision algorithms for a number of properties concerning group cellular automata.

We presented a preliminary version of this work at the conference Mathematical Foundations of Computer Science (MFCS 2020) [1]. The present article adds the main proof in Section 3 of how the projections can be effectively constructed, and a new part in Section 4 concerning the Garden-of-Eden theorem.

2. Preliminaries

We first give definitions related to general subshifts and cellular automata, and then discuss concepts and properties particular to group shifts and group cellular automata.

Symbolic dynamics

A d-dimensional *configuration* over a finite alphabet A is an assignment of symbols of A on the infinite grid $\mathbb{Z}^d$. We call the elements of A the *states*. For any configuration $c \in A^{\mathbb{Z}^d}$ and any cell $\boldsymbol{u} \in \mathbb{Z}^d$, we denote by $c_{\boldsymbol{u}}$ the state $c(\boldsymbol{u})$ that c has in the cell $\boldsymbol{u}$. For any $a \in A$ we denote by $a^{\mathbb{Z}^d}$ the *uniform* configuration defined by $a^{\mathbb{Z}^d}_{\boldsymbol{u}} = a$ for all $\boldsymbol{u} \in \mathbb{Z}^d$.

For a vector $\boldsymbol{t} \in \mathbb{Z}^d$, the *translation* $\tau^{\boldsymbol{t}}$ shifts a configuration c so that the cell $\boldsymbol{t}$ is pulled to the cell $\mathbf{0}$, that is, $\tau^{\boldsymbol{t}}(c)_{\boldsymbol{u}} = c_{\boldsymbol{u}+\boldsymbol{t}}$ for all $\boldsymbol{u} \in \mathbb{Z}^d$. We say that c is *periodic* if $\tau^{\boldsymbol{t}}(c) = c$ for some non-zero $\boldsymbol{t} \in \mathbb{Z}^d$. In this case $\boldsymbol{t}$ is a *vector of periodicity* and c is also termed $\boldsymbol{t}$*-periodic*. If there are d linearly independent vectors of periodicity then c is called *totally periodic*. We denote by $\boldsymbol{e}_i = (0, \dots, 0, 1, 0 \dots, 0)$ the basic i'th unit coordinate vector, for $i = 1, \dots, d$. A totally periodic $c \in A^{\mathbb{Z}^d}$ has automatically, for some $k > 0$, vectors of periodicity $k\boldsymbol{e}_1, k\boldsymbol{e}_2, \dots, k\boldsymbol{e}_d$ in the d coordinate directions.

Let $D \subseteq \mathbb{Z}^d$ be a finite set of cells, a *shape*. A D*-pattern* is an assignment $p \in A^D$ of symbols in the shape D. A *(finite) pattern* is a D-pattern for some shape D. We call D the *domain* of the pattern. We say that a finite pattern p of shape D *appears* in a configuration c if for some $\boldsymbol{t} \in \mathbb{Z}^d$ we have $\tau^{\boldsymbol{t}}(c)|_D = p$. We also say that c *contains* the pattern p. For a fixed D, the set of D-patterns that appear in a configuration c is denoted by $\mathcal{L}_D(c)$. We denote by $\mathcal{L}(c)$ the set of all finite patterns that appear in c, i.e., the union of $\mathcal{L}_D(c)$ over all finite $D \subseteq \mathbb{Z}^d$.

Let $p \in A^D$ be a finite pattern of a shape D. The set $[p] = \{c \in A^{\mathbb{Z}^d} \mid c|_D = p\}$ of configurations that have p in the domain D is called the *cylinder* determined by p. The collection of cylinders $[p]$ is a base of a compact topology on $A^{\mathbb{Z}^d}$, the *prodiscrete* topology. See, for example, the first few pages of [6] for details. The topology is equivalently defined by a metric on $A^{\mathbb{Z}^d}$ where two configurations are close to each other if they agree with each other on a large region around the cell $\mathbf{0}$. Cylinders are clopen in the topology: they are both open and closed.

A subset X of $A^{\mathbb{Z}^d}$ is called a *subshift* if it is closed in the topology and closed under translations. Note that – somewhat nonstandardly – we allow X to be the empty set. By a compactness argument one has that every configuration c that is not in X contains a finite pattern p that prevents it from being in X: no configuration that contains p is in X. We can then as well define subshifts using forbidden patterns: given a set P of finite patterns we define

$$\mathcal{X}_P = \{c \in A^{\mathbb{Z}^d} \mid \mathcal{L}(c) \cap P = \emptyset\},$$

the set of configurations that do not contain any of the patterns in P. The set $\mathcal{X}_P$ is a subshift, and every subshift is $\mathcal{X}_P$ for some P. If $X = \mathcal{X}_P$ for some finite P then X is a *subshift of finite type* (SFT). For a subshift $X \subseteq A^{\mathbb{Z}^d}$ we denote by $\mathcal{L}_D(X)$ and $\mathcal{L}(X)$ the sets of the D-patterns and all finite patterns that appear in elements of X, respectively. The set $\mathcal{L}(X)$ is called the *language* of the subshift.

A continuous function $F : X \longrightarrow Y$ between d-dimensional subshifts $X \subseteq A^{\mathbb{Z}^d}$ and $Y \subseteq B^{\mathbb{Z}^d}$ is a *shift homomorphism* if it is translation invariant, that is, $\tau_Y^{\boldsymbol{t}} \circ F = F \circ \tau_X^{\boldsymbol{t}}$ for every $\boldsymbol{t} \in \mathbb{Z}^d$, where we have denoted the translations $\tau^{\boldsymbol{t}}$ by a vector $\boldsymbol{t}$ with a subscript that indicates the space. A shift homomorphism from a subshift X to itself (i.e. a shift endomorphism) is called a *cellular automaton* on X. The Curtis-Hedlund-Lyndon-theorem [11] states that shift homomorphisms are precisely the functions $X \longrightarrow Y$ defined by a local rule as follows. Let $N \subseteq \mathbb{Z}^d$ be a finite *neighborhood* and let $f : \mathcal{L}_N(X) \longrightarrow B$ be a *local rule* that assigns a letter of B to every N-pattern that appears in X. Applying f at each cell yields a function $F_f : X \longrightarrow B^{\mathbb{Z}^d}$ that maps every c according to $F_f(c)_{\boldsymbol{u}} = f(\tau^{\boldsymbol{u}}(c)|_N)$ for all $\boldsymbol{u} \in \mathbb{Z}^d$. Shift homomorphisms $X \longrightarrow Y$ are precisely such functions F_f that also satisfy $F_f(X) \subseteq Y$.

The image $F(X)$ of a subshift under a shift homomorphism F is clearly also a subshift. Images of subshifts of finite type are called *sofic*. We refer to [21, 22] for more concepts and results on symbolic dynamics.

Group shifts and group cellular automata

Let $\mathbb{G}$ be a finite (not necessarily commutative) group. There is a natural group structure on the d-dimensional configuration space $\mathbb{G}^{\mathbb{Z}^d}$ where the group operation is applied cell-wise: $(ce)_{\boldsymbol{u}} = c_{\boldsymbol{u}} e_{\boldsymbol{u}}$ for all $c, e \in \mathbb{G}^{\mathbb{Z}^d}$ and $\boldsymbol{u} \in \mathbb{Z}^d$. A *group shift* is a subshift of $\mathbb{G}^{\mathbb{Z}^d}$ that is also a subgroup. In particular, a group shift is not empty. A cellular automaton $F : \mathbb{X} \longrightarrow \mathbb{X}$ on a group shift $\mathbb{X} \subseteq \mathbb{G}^{\mathbb{Z}^d}$ is a *group cellular automaton* if it is a group homomorphism: $F(ce) = F(c)F(e)$ for all $c, e \in \mathbb{X}$. More generally, a shift homomorphism $F : \mathbb{X} \longrightarrow \mathbb{Y}$ that is also a group homomorphism between groups shifts $\mathbb{X}$ and $\mathbb{Y}$ is called a *group shift homomorphism*.

Group shifts have two important properties that are central in algorithmic decidability [18]: every group shift is of finite type, and totally periodic configurations are dense in all group shifts [19, 28].

Theorem 1 ([19]) *Every group shift is a subshift of finite type.*

It follows from this theorem that every group shift $\mathbb{X}$ has a finite representation using a finite collection P of forbidden finite patterns as $\mathbb{X} = \mathcal{X}_P$. This is the representation assumed in all algorithmic questions concerning given group shifts. Also when we say that we effectively construct a group shift $\mathbb{X}$ we mean that we produce a finite set P of finite patterns such that $\mathbb{X} = \mathcal{X}_P$.

Theorem 2 ([19]) *Totally periodic configurations are dense in group shifts, i.e., for every $p \in \mathcal{L}(\mathbb{X})$ there is a totally periodic $c \in \mathbb{X}$ such that $p \in \mathcal{L}(c)$.*

As an immediate corollary of these two fundamental properties we get that the language of a group shift is (uniformly) recursive.

Corollary 3. *There is an algorithm that determines, for any given group shift $\mathbb{X} \subseteq \mathbb{G}^{\mathbb{Z}^d}$ and any given finite pattern $p \in \mathbb{G}^D$ whether p is in the language $\mathcal{L}(\mathbb{X})$ of $\mathbb{X}$.*

Proof. This is a standard argumentation by Hao Wang [29]: There is a (nondeterministic) semi-algorithm for positive membership $p \in \mathcal{L}(\mathbb{X})$ that guesses a totally periodic configuration $c \in \mathbb{G}^{\mathbb{Z}^d}$, verifies that c contains the pattern p, and finally verifies that c does not contain any of the forbidden patterns in the given set P that defines $\mathbb{X} = \mathcal{X}_P$. Such a configuration c exists by Theorem 2 iff $p \in \mathcal{L}(\mathbb{X})$. Conversely, as for any SFT, there is a semi-algorithm for the negative cases $p \notin \mathcal{L}(\mathbb{X})$ that guesses a number n, makes sure that the domain D of $p \in \mathbb{G}^D$ is a subset of $E = \{-n, \ldots, n\}^d$, enumerates all finitely many patterns q with domain E that satisfy $q|_D = p$, and verifies that all such q contain a copy of a forbidden pattern in P that defines $\mathbb{X} = \mathcal{X}_P$. By compactness such a number n exists iff $p \notin \mathcal{L}(\mathbb{X})$. □

The representation of an SFT in terms of forbidden patterns is not unique. However, as soon as the language is recursive, we can effectively test if given representations define the same SFT.

Corollary 4. *There are algorithms to determine*

(a) whether $\mathbb{X}_1 \subseteq \mathbb{X}_2$ holds for given group shifts $\mathbb{X}_1, \mathbb{X}_2 \subseteq \mathbb{G}^{\mathbb{Z}^d}$,
(b) whether $\mathbb{X}_1 = \mathbb{X}_2$ holds for given group shifts $\mathbb{X}_1, \mathbb{X}_2 \subseteq \mathbb{G}^{\mathbb{Z}^d}$,

Proof. To prove (a), let $P = \{p_1, \ldots, p_k\}$ be the given set of forbidden patterns that defines $\mathbb{X}_2 = \mathcal{X}_P$. We have $\mathbb{X}_1 \subseteq \mathbb{X}_2$ if and only if $p_1, \ldots, p_k \notin \mathcal{L}(\mathbb{X}_1)$, so (a) follows from Corollary 3. Now (b) follows trivially from (a) and the fact that $\mathbb{X}_1 = \mathbb{X}_2$ iff $\mathbb{X}_1 \subseteq \mathbb{X}_2$ and $\mathbb{X}_2 \subseteq \mathbb{X}_1$. □

Another important known property is that there are no infinite strictly decreasing chains $\mathbb{X}_1 \supsetneq \mathbb{X}_2 \supsetneq \mathbb{X}_3 \supsetneq \dots$ of group shifts [19]. This is clear as the intersection $\mathbb{X}$ of such a chain is a group shift and hence, by Theorem 1, there is a finite set P such that $\mathbb{X} = \mathcal{X}_P$. If a pattern p is in the languages of all $\mathbb{X}_k$ in the chain then p is also in the language of the intersection $\mathbb{X}$, proving that for large enough k the language of $\mathbb{X}_k$ does not contain any of the forbidden patterns in P. This implies that $\mathbb{X}_k = \mathbb{X}$ and the chain does not decrease any further. (Note, however, that while we presented here the decreasing chain property as a corollary to Theorem 1, in reality the proof is interweaved in the proof of Theorem 1, see [19].)

Theorem 5 ([19]) *There does not exist an infinite chain* $\mathbb{X}_1 \supsetneq \mathbb{X}_2 \supsetneq \mathbb{X}_3 \supsetneq \dots$ *of group shifts* $\mathbb{X}_i \subseteq \mathbb{G}^{\mathbb{Z}^d}$.

We also mention the obvious fact that pre-images of group shifts under group shift homomorphisms $F : \mathbb{X} \longrightarrow \mathbb{H}^{\mathbb{Z}^d}$ are group shifts and they can be effectively constructed. In particular, this applies to the kernel $\ker(F) = F^{-1}(\mathbf{1}_{\mathbb{H}}^{\mathbb{Z}^d})$ of F. (We denote the identity element of any group $\mathbb{G}$ by $\mathbf{1}_{\mathbb{G}}$, or simply by $\mathbf{1}$ if the group is clear from the context.)

Lemma 6. *For any given d-dimensional group shifts* $\mathbb{X} \subseteq \mathbb{G}^{\mathbb{Z}^d}$ *and* $\mathbb{Y} \subseteq \mathbb{H}^{\mathbb{Z}^d}$, *and for a given group shift homomorphism* $F : \mathbb{X} \longrightarrow \mathbb{H}^{\mathbb{Z}^d}$, *the set* $F^{-1}(\mathbb{Y})$ *is a group shift that can be effectively constructed. In particular, the kernel* $\ker(F)$ *is a group shift that can be effectively constructed.*

Proof. The set $F^{-1}(\mathbb{Y})$ is clearly topologically closed, translation invariant, and a group, and therefore it is a group shift. Let P and Q be the given finite sets of forbidden patterns defining $\mathbb{X} = \mathcal{X}_P$ and $\mathbb{Y} = \mathcal{X}_Q$. Let $f : \mathcal{L}_N(\mathbb{X}) \longrightarrow \mathbb{H}$ be the given local rule with neighborhood $N \subseteq \mathbb{Z}^d$ that defines $F = F_f$. For each forbidden $q \in \mathbb{H}^D$ in Q we forbid all patterns $p \in \mathbb{G}^{D+N}$ that the local rule maps to q. We also forbid all patterns $p \in P$. The resulting subshift of finite type is $F^{-1}(\mathbb{Y})$. □

3. Algorithms for group shifts

To effectively manipulate group shifts we need algorithms to perform some basic operations. The main operations we consider are taking projections, either to lower the dimension of the space or to project into a subgroup of the state set but keeping the dimension. As a byproduct we obtain an algorithm to compute the image of a given group shift under a given group cellular automaton. We use derivatives of the symbol π for projections from $\mathbb{Z}^d$ to lower dimensional grids, and derivatives of the symbol ψ for projections that keep the dimension of $\mathbb{Z}^d$ but change the state set.

It is essential that projected group shifts remain as group shifts, thereby ensuring they are of finite type. It should be noted that in the general non-group case, projected subshifts of finite type may not necessaril [12]. Therefore, group shifts behave particularly well with respect to projections.

Notations for projections to lower dimensions

Let us first define the projection operators that cut from d-dimensional configurations $(d-1)$-dimensional slices of finite width in the first dimension. Let $d \geq 1$ be the dimension and $n \geq 1$ the width of the slice. For any d-dimensional configuration $c \in A^{\mathbb{Z}^d}$ over alphabet A the *n-slice* $\pi^{(n)}(c)$ is the $(d-1)$-dimensional configuration over alphabet A^n that has in any cell $\boldsymbol{u} \in \mathbb{Z}^{d-1}$ the n-tuple $(c(1, \boldsymbol{u}), \ldots, c(n, \boldsymbol{u})) \in A^n$. The n-slice of a subshift $X \subseteq A^{\mathbb{Z}^d}$ is then the set $\pi^{(n)}(X)$ of the n-slices of all $c \in X$. Due to translation invariance of X, the fact that we cut slices at first coordinate positions $1, \ldots, n$ is irrelevant: we could use any n consecutive first coordinate positions instead. Clearly $\pi^{(n)}(X)$ is a subshift, and if $\mathbb{X} \subseteq \mathbb{G}^{\mathbb{Z}^d}$ is a group shift then $\pi^{(n)}(\mathbb{X})$ is also a group shift over the group $\mathbb{G}^n = \mathbb{G} \times \cdots \times \mathbb{G}$, the n-fold cartesian power of $\mathbb{G}$.

Patterns in $(d-1)$-dimensional slices of thickness n can be interpreted in a natural way as d-dimensional patterns having the width n in the first dimension. We introduce the notation $\hat{p}$ for such an interpretation of a pattern p. More precisely, for any $D \subseteq \mathbb{Z}^{d-1}$ and a $(d-1)$-dimensional pattern $p \in (G^n)^D$ over the alphabet $\mathbb{G}^n$ we denote by $\hat{p} \in \mathbb{G}^E$ the corresponding d-dimensional pattern over $\mathbb{G}$ whose domain is $E = \{1, \ldots, n\} \times D \subseteq \mathbb{Z}^d$ and $p(\boldsymbol{u}) = (\hat{p}(1, \boldsymbol{u}), \hat{p}(2, \boldsymbol{u}), \ldots, \hat{p}(n, \boldsymbol{u}))$ for every $\boldsymbol{u} \in D$. For a subshift X we then have that $p \in \mathcal{L}(\pi^{(n)}(X))$ if and only if $\hat{p} \in \mathcal{L}(X)$. In particular, using an algorithm for the membership of a pattern in $\mathcal{L}(X)$ we can also decide the membership of any given finite pattern in $\mathcal{L}(\pi^{(n)}(X))$. Based on Corollary 3 we then have immediately the following fact for groups shifts.

Lemma 7. *One can effectively decide for any given d-dimensional group shift $\mathbb{X} \subseteq \mathbb{G}^{\mathbb{Z}^d}$, any given $n \geq 1$ and any given $(d-1)$-dimensional finite pattern $p \in (G^n)^D$ whether $p \in \mathcal{L}(\pi^{(n)}(\mathbb{X}))$.* □

Projections $\pi^{(n)}(X)$ are elementary slicing operations that can be composed together, as well as with permutations of coordinates, to obtain more general projections of subshifts into lower dimensional grids. Very generally, for any subset $E \subseteq \mathbb{Z}^d$ we call the restriction $c|_E$ the *projection* of c on E, and the projection of a subshift X on E is $\pi_E(X) = \{c|_E \mid c \in X\}$. We mostly use operation π_E with sets of type $E = D \times \mathbb{Z}^k$ for some $k < d$ and a finite $D \subseteq \mathbb{Z}^{d-k}$, and we mostly apply π_E to group shifts $\mathbb{X} \subseteq \mathbb{G}^{\mathbb{Z}^d}$. The projection $\pi_E(\mathbb{X})$ is then viewed in the natural manner as the k-dimensional group shift over the finite group $\mathbb{G}^D$. One of the main results of this section is Corollary 10, stating that we can effectively construct $\pi_E(\mathbb{X})$ for given $\mathbb{X}$ and $E = D \times \mathbb{Z}^k$.

Notations for projections that keep the dimension

Let $\mathbb{G} = \mathbb{G}_1 \times \mathbb{G}_2$ be a cartesian product of two finite groups. For any $c \in \mathbb{G}^{\mathbb{Z}^d}$ we let $\psi^{(1)}(c) \in \mathbb{G}_1^{\mathbb{Z}^d}$ and $\psi^{(2)}(c) \in \mathbb{G}_2^{\mathbb{Z}^d}$ be the cell-wise projections to $\mathbb{G}_1$ and $\mathbb{G}_2$, respectively, defined by $c_{\boldsymbol{u}} = (\psi^{(1)}(c)_{\boldsymbol{u}}, \psi^{(2)}(c)_{\boldsymbol{u}})$ for all $\boldsymbol{u} \in \mathbb{Z}^d$. By abuse of notation, for any $c^{(1)} \in \mathbb{G}_1^{\mathbb{Z}^d}$ and $c^{(2)} \in \mathbb{G}_2^{\mathbb{Z}^d}$ we denote by $(c^{(1)}, c^{(2)})$ the configuration

$c \in (\mathbb{G}_1 \times \mathbb{G}_2)^{\mathbb{Z}^d}$ such that $\psi^{(i)}(c) = c^{(i)}$ for $i = 1, 2$. We also use the similar notation on finite patterns and implicitly use the obvious way to identify $\mathbb{G}_1^D \times \mathbb{G}_2^D$ and $(\mathbb{G}_1 \times \mathbb{G}_2)^D$.

Clearly, for any group shift $\mathbb{X} \subseteq \mathbb{G}^{\mathbb{Z}^d}$, the sets $\psi^{(1)}(\mathbb{X})$ and $\psi^{(2)}(\mathbb{X})$ are group shifts over $\mathbb{G}_1$ and $\mathbb{G}_2$, respectively. A pattern $p \in (\mathbb{G}_1)^D$ is in the language of $\psi^{(1)}(\mathbb{X})$ if and only if there is a pattern $q \in (\mathbb{G}_2)^D$ such that $(p, q) \in \mathcal{L}_D(\mathbb{X})$. Therefore we have the following counterpart of Lemma 7.

Lemma 8. *One can effectively decide for any given d-dimensional group shift $\mathbb{X} \subseteq (\mathbb{G}_1 \times \mathbb{G}_2)^{\mathbb{Z}^d}$, and any given d-dimensional finite pattern $p \in (\mathbb{G}_1)^D$ whether $p \in \mathcal{L}(\psi^{(1)}(\mathbb{X}))$.* □

Let D, E be finite sets, $D \subseteq E$, and let $\mathbb{X} \subseteq (\mathbb{G}^E)^{\mathbb{Z}^d}$ be a group shift over the finite cartesian power $\mathbb{G}^E$ of the group $\mathbb{G}$. The group $\mathbb{G}^E$ is isomorphic to $\mathbb{G}^D \times \mathbb{G}^{E \setminus D}$ in a natural manner, and $\psi^{(1)}$ projects then $\mathbb{X}$ into $(\mathbb{G}^D)^{\mathbb{Z}^d}$. We denote this projection by ψ_D. Notice that $\pi_{D \times \mathbb{Z}^k} = \psi_D \circ \pi_{E \times \mathbb{Z}^k}$ so that the projection into $D \times \mathbb{Z}^k$ can be obtained as a composition of projections $\pi^{(n)}$ into slices, permutations of coordinates, and a projection of the type $\psi^{(1)}$.

Effective constructions

Our main technical result is that projections of group shifts can be effectively constructed. We state this as a two-part lemma. Corollaries 10 and 11 that follow the lemma provide clean statements that we use in the rest of the paper.

Lemma 9. *Let $d \geq 1$ be a dimension, and let $\mathbb{G}$ and $\mathbb{G}_1, \mathbb{G}_2$ be finite groups.*

(a) For any given d-dimensional group shift $\mathbb{X} \subseteq \mathbb{G}^{\mathbb{Z}^d}$ and any given $n \geq 1$ one can effectively construct the $d-1$ dimensional group shift $\pi^{(n)}(\mathbb{X}) \subseteq (\mathbb{G}^n)^{\mathbb{Z}^{d-1}}$.

(b) For any given d-dimensional group shift $\mathbb{X} \subseteq (\mathbb{G}_1 \times \mathbb{G}_2)^{\mathbb{Z}^d}$ one can effectively construct the d-dimensional group shift $\psi^{(1)}(\mathbb{X}) \subseteq \mathbb{G}_1^{\mathbb{Z}^d}$.

Proof. The proof is by induction on dimension d. We first prove (a) for dimension d assuming that (b) holds in dimension $d-1$, and then we prove (b) for dimension d assuming (a) holds in dimension d and that (b) holds for dimension $d-1$. To start the induction we observe that (b) trivially holds for dimension $d = 0$: In this case group shifts over $\mathbb{G}$ are precisely subgroups of $\mathbb{G}$.

Proving (a) for dimension d assuming (b) holds for dimension $d-1$: Let a width $n \geq 1$ and a group shift $\mathbb{X} \subseteq \mathbb{G}^{\mathbb{Z}^d}$ be given (in terms of a finite set P of forbidden patterns such that $\mathbb{X} = \mathcal{X}_P$). Let us first assume that n is at least the width of the patterns in P so that we can assume that all patterns in P have the same domain $\{1, \ldots, n\} \times D$ for some finite $D \subseteq \mathbb{Z}^{d-1}$. (Note that we can effectively grow the domain of each forbidden pattern by forbidding instead all patterns with the larger domain that extend the original pattern. Thus a common domain can be taken for all elements in P. We can also shift the domains of the patterns.)

To construct the $(d-1)$-dimensional projection $\mathbb{Y} = \pi^{(n)}(\mathbb{X})$ we effectively enumerate and forbid patterns that are not in the language of $\mathbb{Y}$. We accumulate the forbidden patterns in a set Q that we initialize to be the empty set in the beginning of the process. Let $D_1, D_2, \dots$ be an effective enumeration of all finite subsets of $\mathbb{Z}^{d-1}$ with $D_1 = D$. For each $i = 1, 2, \dots$ in turn we go through all (finitely many) patterns q over $\mathbb{G}^n$ having shape D_i and check, using Lemma 7, whether q is in $\mathcal{L}(\mathbb{Y})$. If not, we add q in the set Q. This way, at any time, Q only contains patterns outside of $\mathcal{L}(\mathbb{Y})$ and hence forbidding patterns in Q gives an upper approximation $\mathcal{X}_Q \supseteq \mathbb{Y}$. Since $\mathbb{Y}$ is a group shift and therefore of finite type, by systematically enumerating the patterns in the complement of $\mathcal{L}(\mathbb{Y})$ we eventually reach a set Q such that $\mathbb{Y} = \mathcal{X}_Q$.

The reason why we process all patterns for each shape D_i before moving to the next shape D_{i+1} is the observation that this way the subshift $\mathcal{X}_Q$ is guaranteed to be a group shift after finishing processing D_i. We have the following general fact:

Claim 1. *Let $\mathbb{X} \subseteq \mathbb{G}^{\mathbb{Z}^d}$ be any group shift in any dimension d, and let $D \subseteq \mathbb{Z}^d$ be finite. For $Q = \mathbb{G}^D \setminus \mathcal{L}_D(\mathbb{X})$ the subshift $\mathcal{X}_Q$ is a group shift and $\mathbb{X} \subseteq \mathcal{X}_Q$.*

Proof of Claim 1. Clearly $\mathcal{X}_Q$ is a subshift and $\mathbb{X} \subseteq \mathcal{X}_Q$. We just have to show that $\mathcal{X}_Q$ is a group. We have $c \in \mathcal{X}_Q$ if and only if $\mathcal{L}_D(c) \subseteq \mathcal{L}_D(\mathbb{X})$. The result now follows from the fact that $\mathcal{L}_D(\mathbb{X})$ is a subgroup of $\mathbb{G}^D$. □

Intersections of group shifts are group shifts so Claim 1 indeed implies that after fully processing any number of domains $D_1, \dots, D_i$ the resulting subshift $\mathcal{X}_Q \supseteq \mathbb{Y}$ is a group shift. Note also that $D_1 = D$ guarantees that already after the first round $i = 1$ we have in Q all the patterns of P.

As mentioned above, we are guaranteed to eventually have enough forbidden patterns in Q to have $\mathbb{Y} = \mathcal{X}_Q$. The problem is to identify when we have enumerated enough patterns and reached such a set Q. Fortunately this can be detected by checking that the left and the right slices of width $n-1$ of the upper approximation $\mathcal{X}_Q$ are identical with each other, as detailed below.

Let us introduce notations ψ_L and ψ_R for the operations of extracting the left and the right slices of width $n-1$. More precisely, for a configuration $c = (c^{(1)}, \dots, c^{(n)}) \in (\mathbb{G}^n)^{\mathbb{Z}^{d-1}}$ of thickness n, where $c^{(i)} \in \mathbb{G}^{\mathbb{Z}^{d-1}}$ are the single cell wide slices of c, we define $\psi_L(c) = (c^{(1)}, \dots, c^{(n-1)})$ and $\psi_R(c) = (c^{(2)}, \dots, c^{(n)})$, respectively. Both are elements of $(\mathbb{G}^{n-1})^{\mathbb{Z}^{d-1}}$.

Claim 2. *$\mathcal{X}_Q = \mathbb{Y}$ if and only if $\psi_L(\mathcal{X}_Q) = \psi_R(\mathcal{X}_Q)$.*

Proof of Claim 2. $\psi_L(\mathbb{Y}) = \psi_R(\mathbb{Y}) = \pi^{(n-1)}(\mathbb{X})$ so the implication from left to right is clear. For the converse direction, let $c \in \mathcal{X}_Q$ be arbitrary. By the assumption $\psi_L(\mathcal{X}_Q) = \psi_R(\mathcal{X}_Q)$ there exists a bi-infinite sequence $\dots, c_{-1}, c_0, c_1, \dots$ of configurations such that $c_0 = c$ and for all $i \in \mathbb{Z}$ we have $c_i \in \mathcal{X}_Q$ and $\psi_R(c_i) = \psi_L(c_{i+1})$. Configurations c_i and c_{i+1} overlap properly so that there is a d-dimensional configuration $c' \in \mathbb{G}^{\mathbb{Z}^d}$ whose consecutive n-slices are $\ldots, c_{-1}, c_0, c_1, \ldots,$ that is, $c_i = \pi^{(n)}(\tau_{i\boldsymbol{e}_1}(c'))$ for all $i \in \mathbb{Z}$. See Figure 1 for an illustration of c'. Since each forbidden pattern in P is also in Q, none of the slices contain such a forbidden pattern and hence $c' \in \mathbb{X}$. Now $c = c_0 = \pi^{(n)}(c')$ so that $c \in \pi^{(n)}(\mathbb{X}) = \mathbb{Y}$. We have shown that $\mathcal{X}_Q \subseteq \mathbb{Y}$. The opposite inclusion holds since $\mathcal{X}_Q$ is an upper approximation of $\mathbb{Y}$. □

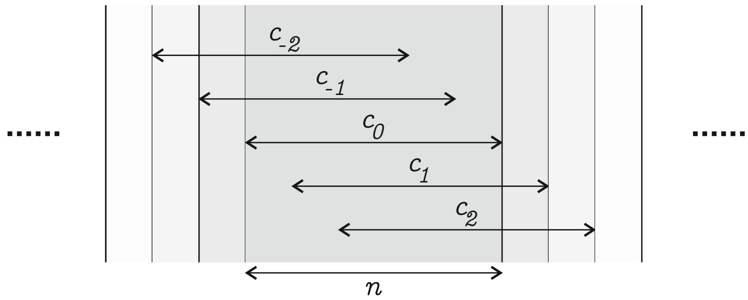

Fig. 1. An illustration of the overlapping n-slices forming the configuration c' in the proof of Claim 2.

Both ψ_L and ψ_R are projection operations of type (b) of the present lemma, so by the inductive hypotheses and the fact that $\mathcal{X}_Q$ is a $(d-1)$-dimensional group shift, the group shifts $\psi_L(\mathcal{X}_Q)$ and $\psi_R(\mathcal{X}_Q)$ can be effectively constructed. Moreover, equality of group shifts is decidable so that the condition $\psi_L(\mathcal{X}_Q) = \psi_R(\mathcal{X}_Q)$ can be effectively checked. In conclusion, each time our algorithm finishes with adding patterns of shape D_i in Q it checks whether $\psi_L(\mathcal{X}_Q) = \psi_R(\mathcal{X}_Q)$ holds for the current group shift $\mathcal{X}_Q$. The algorithm stops and returns set Q once equality is reached. This finishes the description of the algorithm for case (a), provided n is large enough to have all patterns P in a slice of width n. If n is smaller, we first execute the algorithm for large enough width $m > n$ and effectively compute the further projection $\pi^{(n)}(\mathbb{X}) = \psi_L^{m-n}(\pi^{(m)}(\mathbb{X}))$ to slices of width n. The projection can be effectively computed by the inductive hypothesis because it is a $(d-1)$-dimensional operation of type (b) of the present lemma.

Proving (b) for dimension d assuming that (a) holds for dimension d and that (b) holds for dimension $d-1$*:* Let $\mathbb{X} \subseteq (\mathbb{G}_1 \times \mathbb{G}_2)^{\mathbb{Z}^d}$ be given (in terms of a finite set P of forbidden patterns such that $\mathbb{X} = \mathcal{X}_P$). To construct the d-dimensional projection $\mathbb{Y} = \psi^{(1)}(\mathbb{X})$ we – analogously to the proof of case (a) above – use Lemma 8 to effectively enumerate patterns that are not in the language of $\mathbb{Y}$, thus obtaining upper approximations of $\mathbb{Y}$ by subshifts $\mathcal{X}_Q$. We process all patterns of a shape D_i before moving on to the next shape D_{i+1}. This guarantees – as proved in Claim 1 above – that after finishing with each shape D_i the shift $\mathcal{X}_Q$ is a group shift.

We eventually reach a set Q such that $\mathbb{Y} = \mathcal{X}_Q$, but the challenge is again to identify when we have reached such Q. We establish this by proving that we can effectively compute a number n such that $\mathbb{Y} = \mathcal{X}_Q$ if and only if $\pi^{(n)}(\mathcal{X}_Q) = \pi^{(n)}(\mathbb{Y})$.

Once number n is known, the projection $\pi^{(n)}(\mathcal{X}_Q)$ can be effectively constructed by the inductive hypothesis stating that case (a) of the present lemma holds in dimension d. Indeed, $\mathcal{X}_Q$ is a known d-dimensional group shift. The projection $\pi^{(n)}(\mathbb{Y})$ can also be effectively constructed as projections $\psi^{(1)}$ and $\pi^{(n)}$ commute, so that we can first construct $\pi^{(n)}(\mathbb{X})$ (using the inductive hypothesis that case (a) of the present lemma holds in dimension d) and then we apply $\psi^{(1)}$ on the $(d-1)$-dimensional group shift $\pi^{(n)}(\mathbb{X})$ (using the inductive hypothesis that case (b) of the present lemma holds in dimension $d-1$) to obtain $\pi^{(n)}(\mathbb{Y})$.

All that remains is to compute a sufficiently large n for the implication

$$\pi^{(n)}(\mathcal{X}_Q) = \pi^{(n)}(\mathbb{Y}) \Longrightarrow \mathcal{X}_Q = \mathbb{Y}$$

to hold.

First a note on notations: Recall that we denote for any $c \in \mathbb{G}_1^{\mathbb{Z}^d}$ and $e \in \mathbb{G}_2^{\mathbb{Z}^d}$ by (c, e) the configuration in $(\mathbb{G}_1 \times \mathbb{G}_2)^{\mathbb{Z}^d}$ such that $\psi^{(1)}(c, e) = c$ and $\psi^{(2)}(c, e) = e$. We also then denote for any $c \in (\mathbb{G}^n)^{\mathbb{Z}^d}$ and $c' \in (\mathbb{G}^m)^{\mathbb{Z}^d}$ by (c, c') the concatenated configuration in $(\mathbb{G}^{n+m})^{\mathbb{Z}^d}$, by the understanding that $\mathbb{G}^{n+m} = \mathbb{G}^n \times \mathbb{G}^m$. In the following we are going to mix both types of concatenations. For example, for $c \in (\mathbb{G}_1^n)^{\mathbb{Z}^d}$, $c' \in (\mathbb{G}_1^m)^{\mathbb{Z}^d}$, $e \in (\mathbb{G}_2^n)^{\mathbb{Z}^d}$ and $e' \in (\mathbb{G}_2^m)^{\mathbb{Z}^d}$ we may write $((c, e), (c', e'))$ for a concatenated configuration in $((\mathbb{G}_1 \times \mathbb{G}_2)^{n+m})^{\mathbb{Z}^d}$, but also $((c, c'), (e, e'))$ for the same configuration, now expressed in $(\mathbb{G}_1^{n+m} \times \mathbb{G}_2^{n+m})^{\mathbb{Z}^d}$. To help the reader in the task of parsing such expressions, we use the notation $[c|c']$ for the second type of concatenations, with the idea that n-slices can be visualized as strips in the vertical direction and the vertical line $|$ is a "separator" between concatenated vertical strips. So the two examples above will be written as $[(c, e)|(c', e')]$ and $([c|c'], [e|e'])$, respectively.

Also note that for configurations c and e of the same group shift, say $c, e \in (\mathbb{G}^n)^{\mathbb{Z}^d}$, the notation ce is not for the concatenation of the strips but it is for the cell-wise product of the configurations, *i.e.*, for the product in the group $(\mathbb{G}^n)^{\mathbb{Z}^d}$.

For any group shift $\mathbb{U}$ over the alphabet $\mathbb{G}_1 \times \mathbb{G}_2$ we denote by $\text{cut}(\mathbb{U})$ the set of configurations c over $\mathbb{G}_2$ such that $(\mathbf{1}, c) \in \mathbb{U}$. Because $\text{cut}(\mathbb{U}) = \psi^{(2)}(\ker(\psi^{(1)}) \cap \mathbb{U})$ and because projections $\psi^{(i)}$ are group shift homomorphisms the set $\text{cut}(\mathbb{U})$ is a group shift.

Claim 3. *For any given group shift $\mathbb{U} \subseteq (\mathbb{G}_1 \times \mathbb{G}_2)^{\mathbb{Z}^k}$, in any dimension k, one can effectively construct* $\text{cut}(\mathbb{U})$.

Proof of Claim 3. By Lemma 6 we can effectively construct $\ker(\psi^{(1)})$. Intersections of subshifts of finite type can be effectively constructed (simply take the union of the defining sets of forbidden patterns of the two SFTs). This means that $\mathbb{U}' = \ker(\psi^{(1)}) \cap \mathbb{U}$ can be effectively constructed. Let R be the constructed set of finite patterns such that $\mathbb{U}' = \mathcal{X}_R$. All configurations in $\mathbb{U}'$ have $\mathbf{1}$ in their first components so to define $\psi^{(2)}(\mathbb{U}')$ it is enough to forbid for all $(\mathbf{1}, p) \in R$ the pattern p. □

After these notations we can proceed with the proof. Let m be a number such that the forbidden patterns in set P that defines $\mathbb{X}$ fit in a slice of thickness m, that is, the domain of each forbidden pattern in P is a subset of $\{1, \dots, m\} \times \mathbb{Z}^{d-1}$. Let us call a positive integer r a *radius of synchronization* if for all $w \in \mathbb{G}_2^{\{1,\dots,m\}\times\mathbb{Z}^{d-1}}$ holds the implication

$$\begin{aligned}(\exists u, v \in \mathbb{G}_2^{\{1,\dots,r\}\times\mathbb{Z}^{d-1}})\ [u|w|v] \in \mathrm{cut}(\pi^{(m+2r)}(\mathbb{X})) \\ \Longrightarrow \quad w \in \pi^{(m)}(\mathrm{cut}(\mathbb{X})).\end{aligned} \tag{1}$$

(See Figure 2 for an illustration.)

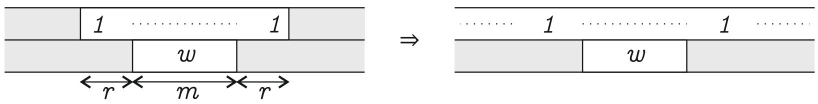

Fig. 2. An illustration of the implication (1). Upper and lower layers are configurations over $\mathbb{G}_1$ and $\mathbb{G}_2$, respectively. Letter r is a radius of synchronization if for every w for which the left situation exists in $\mathbb{X}$ also the right situation exists in $\mathbb{X}$. The picture depicts only the first dimension – each letter represents an entire $(d-1)$-dimensional configuration.

Claim 4. *A radius of synchronization exists, and we can effectively find one.*

Proof of Claim 4. For any r, let us denote by $\mathbb{U}_r$ the set of $w \in \mathbb{G}_2^{\{1,\dots,m\}\times\mathbb{Z}^{d-1}}$ that satisfy the left-hand-side of implication of (1), and by $\mathbb{U}$ the set of those that satisfy the right-hand-side. Now $\mathbb{U} = \pi^{(m)}(\mathrm{cut}(\mathbb{X}))$ and $\mathbb{U}_r = \psi(\mathrm{cut}(\pi^{(m+2r)}(\mathbb{X})))$ where ψ is the projection in the central segment of length m. It follows that $\mathbb{U}_r$ and $\mathbb{U}$ are $d-1$-dimensional group shifts. Group shifts $\mathbb{U}_r$ form a decreasing chain $\mathbb{U}_1 \supseteq \mathbb{U}_2 \supseteq \dots$ so by Theorem 5 there exists k such that $\mathbb{U}_r = \mathbb{U}_k$ for all $r \geq k$. By a simple compactness argument we then also have that $\mathbb{U} = \mathbb{U}_k$: if $w \in \mathbb{U}_k$ then for every $r \geq k$ there exists $c_r \in \mathbb{X}$ as in the left of Figure 2, so that a limit of a converging subsequence of $c_k, c_{k+1}, \dots$ is as in the right of Figure 2, proving that $w \in \mathbb{U}$. This proves that k is a radius of synchronization.

To find a radius of synchronization we enumerate $r = 1, 2, \dots$ and test for each r whether $\mathbb{U}_r = \mathbb{U}$. This can be effectively tested: First, by Claim 3 the set $\mathrm{cut}(\mathbb{X})$ can be constructed and then by the inductive hypothesis that (a) holds in dimension d we can apply $\pi^{(m)}$ to form $\mathbb{U}$. Second, by the inductive hypothesis that (a) holds in dimension d we can construct $\pi^{(m+2r)}(\mathbb{X})$, by Claim 3 we can build $\mathrm{cut}(\pi^{(m+2r)}(\mathbb{X}))$, and finally by the inductive hypothesis that (b) holds in dimension $d-1$ we apply ψ to construct $\mathbb{U}_r$. So both $\mathbb{U}$ and $\mathbb{U}_r$ can be effectively constructed, and by Corollary 4(b) we can test whether they are equal. □

The importance of the radius of synchronization comes from the fact that sufficiently wide slices of identities $\mathbf{1}$ can be extended.

Claim 5. *Let r be a radius of synchronization. Then for any slice $x \in \pi^{(k)}(\mathbb{Y})$ of any width k holds the implication*

$$[x|\mathbf{1}^{m+2r}] \in \pi^{(k+m+2r)}(\mathbb{Y}) \qquad \Longrightarrow \qquad [x|\mathbf{1}^{m+2r+1}] \in \pi^{(k+m+2r+1)}(\mathbb{Y}).$$

Proof of Claim 5. Assume the left-hand-side of the implication. Recalling that $\mathbb{Y} = \psi^{(1)}(\mathbb{X})$ there is a configuration $c \in \mathbb{X}$ such that $\pi^{(k+m+2r)}(c) = [(x,y)|(\mathbf{1}^r,u)|(\mathbf{1}^m,w)|(\mathbf{1}^r,v)]$ for some slices y,u,w,v (of thicknesses k, r, m and r, respectively) over $\mathbb{G}_2$. In particular then $[u|w|v] \in \mathrm{cut}(\pi^{(m+2r)}(\mathbb{X}))$, so that the implication (1) gives that $w \in \pi^{(m)}(\mathrm{cut}(\mathbb{X}))$. By the definition of $\mathrm{cut}(\mathbb{X})$ there is a configuration $e \in \mathbb{X}$ such that $\pi^{(k+m+2r+1)}(e) = [(\mathbf{1}^k,y')|(\mathbf{1}^r,u')|(\mathbf{1}^m,w)|(\mathbf{1}^r,v')|(\mathbf{1},a)]$ for slices y',u',v' and a of thicknesses k, r, m, r and 1, respectively. The forbidden patterns in the set P that defines $\mathbb{X}$ have thickness at most m, so we can cut and exchange tails at the common slice $(\mathbf{1}^m,w)$ of c and e without introducing any forbidden patterns. See Figure 3 for an illustration of the cut and exchange between c and e along their common slice. This implies that the slice $[(x,y)|(\mathbf{1}^r,u)|(\mathbf{1}^m,w)|(\mathbf{1}^r,v')|(\mathbf{1},a)]$ of thickness $k+m+2r+1$ is in $\pi^{(k+m+2r+1)}(\mathbb{X})$, providing the result that $[x|\mathbf{1}^{m+2r+1}] \in \pi^{(k+m+2r+1)}(\mathbb{Y})$. □

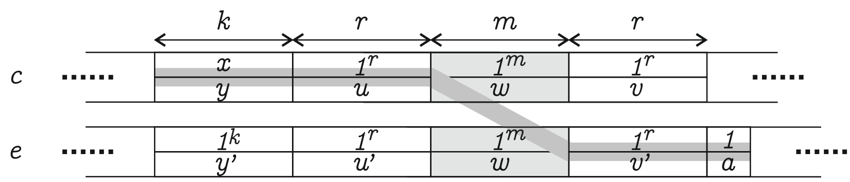

Fig. 3. An illustration of cutting and reconnecting halves of configurations c and e along a common slice of width m in the proof of Claim 5.

Let $n = m + 2r + 1$ where r is the radius of synchronization that we computed for $\mathbb{X}$. This turns out to be a sufficient thickness for our purpose of halting the algorithm.

Claim 6. *If $\pi^{(n)}(\mathcal{X}_Q) = \pi^{(n)}(\mathbb{Y})$ then $\pi^{(k)}(\mathcal{X}_Q) = \pi^{(k)}(\mathbb{Y})$ for all $k \geq n$.*

Proof of Claim 6. We prove this by induction on k. Case $k = n$ is clear. For the inductive step, suppose that $\pi^{(k)}(\mathcal{X}_Q) = \pi^{(k)}(\mathbb{Y})$ is known for some $k \geq n$ and consider slices of width $k+1$. Containment $\pi^{(k+1)}(\mathbb{Y}) \subseteq \pi^{(k+1)}(\mathcal{X}_Q)$ is clear since $\mathbb{Y} \subseteq \mathcal{X}_Q$. We just need to prove that $\pi^{(k+1)}(\mathcal{X}_Q) \subseteq \pi^{(k+1)}(\mathbb{Y})$.

Let $c \in \pi^{(k+1)}(\mathcal{X}_Q)$ so that $c = [a|x|b]$ for $a, b \in \pi^{(1)}(\mathcal{X}_Q)$ and $x \in \pi^{(k-1)}(\mathcal{X}_Q)$. We have $[a|x], [x|b] \in \pi^{(k)}(\mathcal{X}_Q)$ so that by the inductive hypothesis $[a|x], [x|b] \in \pi^{(k)}(\mathbb{Y})$. Because $[x|b]$ is a slice in a configuration of $\mathbb{Y}$, there exists $a' \in \pi^{(1)}(\mathbb{Y})$ such that $[a'|x|b] \in \pi^{(k+1)}(\mathbb{Y})$. Because $\mathbb{Y}$ is a group shift the product $[a|x]\,[a'|x]^{-1} = [aa'^{-1}|\mathbf{1}^{k-1}]$ is in $\pi^{(k)}(\mathbb{Y})$. Because $k-1 \geq m+2r$ we get from Claim 5 that $[aa'^{-1}|\mathbf{1}^{k}]$ is in $\pi^{(k+1)}(\mathbb{Y})$. But this is all we need: we get $[aa'^{-1}|\mathbf{1}^{k}]\,[a'|x|b] = [a|x|b] = c$ is in $\pi^{(k+1)}(\mathbb{Y})$ as claimed. □

It is now a simple compactness argument to show that if $\pi^{(k)}(\mathcal{X}_Q) = \pi^{(k)}(\mathbb{Y})$ for all $k \geq n$ then $\mathcal{X}_Q = \mathbb{Y}$. So our algorithm constructs sets Q until condition $\pi^{(n)}(\mathcal{X}_Q) = \pi^{(n)}(\mathbb{Y})$ is satisfied for $n = m + 2r + 1$. At that time we can stop because we know that we have reached the situation $\mathcal{X}_Q = \mathbb{Y}$. This completes the proof of Lemma 9. □

Lemma 9 is used in the rest of the paper via the following two corollaries. The first corollary states that arbitrary projections can be effectively implemented on group shifts.

Corollary 10. *Given a d-dimensional group shift $\mathbb{X} \subseteq \mathbb{G}^{\mathbb{Z}^d}$ and given $k < d$ and a finite $D \subseteq \mathbb{Z}^{d-k}$ we can effectively construct the k-dimensional group shift $\pi_{D\times\mathbb{Z}^k}(\mathbb{X}) \subseteq (\mathbb{G}^D)^{\mathbb{Z}^k}$.*

Proof. By shift invariance of $\mathbb{X}$ we arbitrarily translate D, so we may assume without loss of generality that D is a subset of $E = \{1, \dots, n\}^{d-k}$ for some n. By applying $d-k$ times Lemma 9(a), permuting the coordinates as needed, we can effectively construct $\mathbb{X}' = \pi_{E\times\mathbb{Z}^k}(\mathbb{X})$. Now $\pi_{D\times\mathbb{Z}^k}(\mathbb{X}) = \psi_D(\mathbb{X}')$, and by Lemma 9(b) the projection ψ_D from $\mathbb{G}^E$ to $\mathbb{G}^D$ can be effectively implemented. □

The second corollary tells that images of group shifts under group cellular automata can be also effectively constructed.

Corollary 11. *Given a d-dimensional group shift $\mathbb{X} \subseteq \mathbb{G}^{\mathbb{Z}^d}$ and given a group shift homomorphism $F : \mathbb{X} \longrightarrow \mathbb{H}^{\mathbb{Z}^d}$ one can effectively construct the group shift $F(\mathbb{X}) \subseteq \mathbb{H}^{\mathbb{Z}^d}$.*

Proof. Let $\mathbb{X} = \mathcal{X}_P$ where P is the given finite set of forbidden patterns that defines $\mathbb{X}$, and let $F = F_f$ where $f : \mathcal{L}_N(\mathbb{X}) \longrightarrow \mathbb{H}$ is the given local rule of F with a neighborhood N. We can pad symbols to patterns to grow their domains, so we can assume without loss of generality that all patterns in P have the same domain D, that the neighborhood is the same set $N = D$, and that $\mathbf{0} \in D$.

We first effectively construct $\mathbb{X}' = \{(c, F(c)) \mid c \in \mathbb{X}\} \subseteq (\mathbb{G} \times \mathbb{H})^{\mathbb{Z}^d}$. This is a group shift over group $\mathbb{G} \times \mathbb{H}$ because F is a homomorphism. It is defined by forbidding all patterns $(p, q) \in (\mathbb{G} \times \mathbb{H})^D$ where $p \notin \mathcal{L}_D(\mathbb{X})$, or $p \in \mathcal{L}_D(\mathbb{X})$ but $q(\mathbf{0}) \neq f(p)$. So $\mathbb{X}'$ can indeed be effectively constructed. By Lemma 9(b) we can then effectively compute the second projection $F(\mathbb{X}) = \psi^{(2)}(\mathbb{X}')$. □

4. Algorithms for group cellular automata

In this part we apply the algorithms developed for group shifts to analyze group cellular automata. The basic idea is to view the set of space-time diagrams as a higher dimensional group shift and to effectively compute one-dimensional projections in the temporal direction. This way, trace subshifts are obtained. As these are one-dimensional group shifts, and hence of finite type, the long term dynamics can be analyzed. A projection in the spatial dimensions provides the limit set of the cellular automaton.

We first define the central concepts of space-time diagrams, traces and limit sets, and show that they can be effectively constructed. Then we use this to prove properties and algorithms concerning several dynamical properties of group cellular automata. We refer to [15, 21] for more details and known results on the dynamical properties we consider.

Space-time diagrams

Let $\mathbb{X} \subseteq \mathbb{G}^{\mathbb{Z}^d}$ be a d-dimensional group shift and let $F : \mathbb{X} \longrightarrow \mathbb{X}$ be a group cellular automaton on $\mathbb{X}$. A bi-infinite *orbit* of F is a sequence $\dots c^{(-1)}, c^{(0)}, c^{(1)}, \dots$ of configurations $c^{(i)} \in \mathbb{X}$ such that $c^{(i+1)} = F(c^{(i)})$ for all $i \in \mathbb{Z}$. Such an orbit can be viewed as the $(d+1)$-dimensional configuration $c \in \mathbb{G}^{\mathbb{Z}^{d+1}}$ by concatenating the configurations c_i one after the other along the additional dimension, that is, $c_{\boldsymbol{u},i} = c^{(i)}_{\boldsymbol{u}}$ for all $i \in \mathbb{Z}$ and $\boldsymbol{u} \in \mathbb{Z}^d$. The first d dimensions are spatial dimensions while the $(d+1)$st dimension is the temporal dimension. The configuration c is a *space-time diagram* of the cellular automaton F. Note that the orbits and space-time-diagrams are temporally bi-infinite. The set of all space-time diagrams of F is denoted by $\mathrm{ST}(F)$. Because F is a group homomorphism we have the following.

Lemma 12. $\mathrm{ST}(F) \subseteq \mathbb{G}^{\mathbb{Z}^{d+1}}$ *is a group shift.* □

Given $\mathbb{X}$ and F we can effectively construct $\mathrm{ST}(F)$. Indeed, we just need to forbid in spatial slices all the forbidden patterns that define $\mathbb{X}$, and in temporally consecutive pairs of slices patterns where the local update rule of F is violated. More precisely, let P be the given finite set of forbidden patterns that defines $\mathbb{X} = \mathcal{X}_P$, and let $f : \mathcal{L}_N(\mathbb{X}) \longrightarrow \mathbb{G}$ be the given local update rule that defines F with the finite neighborhood $N \subseteq \mathbb{Z}^d$. For any $p \in P$ we forbid the $(d+1)$-dimensional pattern $\hat{p}$ over the domain $D \times \{0\}$ with $\hat{p}(\boldsymbol{u}, 0) = p(\boldsymbol{u})$ for all $\boldsymbol{u} \in D$, i.e., the spatial slices are forced to belong to $\mathbb{X}$, and for any neighborhood pattern $q \in \mathcal{L}_N(\mathbb{X})$

and for any $a \in \mathbb{G}$ such that $a \neq f(q)$ we forbid the pattern q'_a with the domain $N \times \{0\} \cup \{(\mathbf{0}, 1)\}$ where $q'_a(\boldsymbol{u}, 0) = q(\boldsymbol{u})$ for all $\boldsymbol{u} \in N$ and $q'_a(\mathbf{0}, 1) = a$, i.e. consecutive slices are prevented from having an update error according to the local rule f. Let P' be the set of all $\hat{p}$ and q'_a. Then clearly $\mathrm{ST}(F) = \mathcal{X}_{P'}$.

Lemma 13. *Given* $\mathbb{X}$ *and* F *one can effectively construct* $\mathrm{ST}(F)$. □

Traces

Let $D \subseteq \mathbb{Z}^d$ be finite. For any orbit $\ldots, c^{(-1)}, c^{(0)}, c^{(1)}, \ldots$ the sequence $\ldots, c^{(-1)}|_D, c^{(0)}|_D, c^{(1)}|_D, \ldots$ of consecutive views in the domain D is a *D-trace*. Each $c^{(i)}|_D$ is an element of the finite group $\mathbb{G}^D$, and hence the trace is a one-dimensional configuration over the group $\mathbb{G}^D$. Let us denote by $\mathrm{Tr}_D(F) \subseteq (\mathbb{G}^D)^{\mathbb{Z}}$ the set of all D-traces of F.

Lemma 14. $\mathrm{Tr}_D(F)$ *is a one-dimensional group shift over* $\mathbb{G}^D$. *It is the projection of* $\mathrm{ST}(F)$ *on* $D \times \mathbb{Z}$. □

We call the set $\mathrm{Tr}_D(F)$ the *D-trace subshift* of F, or simply a trace subshift of F. It can be effectively constructed: Given $\mathbb{X}$ and F we can use Lemma 13 to effectively construct the group shift $\mathrm{ST}(F)$ of space-time diagrams, and then by Corollary 10 we can effectively construct the projection $\mathrm{Tr}_D(F)$ of $\mathrm{ST}(F)$ on $D \times \mathbb{Z}$.

Lemma 15. *Given* $\mathbb{X}$ *and* F *and any finite* $D \subseteq \mathbb{Z}^d$, *one can effectively construct* $\mathrm{Tr}_D(F)$. □

Limit sets

The *limit set* Ω_F of a cellular automaton F consists of all configurations $c^{(0)} \in \mathbb{X}$ that are present in some bi-infinite orbit $\ldots c^{(-1)}, c^{(0)}, c^{(1)}, \ldots$ In other words, Ω_F is the set of the d-dimensional slices of thickness one of $\mathrm{ST}(F)$ in the d spatial dimensions. As a projection of the group shift $\mathrm{ST}(F)$, the set Ω_F is a group shift.

Lemma 16. Ω_F *is a d-dimensional group shift over* $\mathbb{G}$. *It is the projection of* $\mathrm{ST}(F)$ *on* $\mathbb{Z}^d \times \{0\}$. □

Using Corollary 10 we immediately get an algorithm to construct the limit set.

Lemma 17. *Given* $\mathbb{X}$ *and* F, *one can effectively construct* Ω_F. □

By definition it is clear that $F(\Omega_F) = \Omega_F$ so that F is surjective on its limit set. By a simple compactness argument we have that $\Omega_F = \bigcap_{n \in \mathbb{N}} F^n(\mathbb{X})$, stating that any configuration that has arbitrarily long sequences of pre-images has an infinite sequence of pre-images. Note that $\mathbb{X} \supseteq F(\mathbb{X}) \supseteq F^2(\mathbb{X}) \supseteq$ is a decreasing chain of group shifts. By Theorem 5 there are no infinite strictly decreasing chains of group shifts, so we have that $F^{k+1}(\mathbb{X}) = F^k(\mathbb{X})$ holds for some k. Then $F^j(\mathbb{X}) = F^k(\mathbb{X})$ for all $j > k$ so that $\Omega_F = F^k(\mathbb{X})$. So all group cellular automata reach their limit set after a finite time:

Group cellular automata $F : \mathbb{X} \longrightarrow \mathbb{X}$ are stable in the sense that there exists $k \in \mathbb{N}$ such that $F^k(\mathbb{X}) = \Omega_F$. □

Periodic points

A well-known open problem asks whether every surjective cellular automaton on a full shift (even in one-dimension) has a dense set of temporally periodic points. This has been proved to be the case for one-dimensional cellular automata that are closing [3] or have equicontinuity points [2]. It is also the case for one-dimensional group CA [3] and, in any dimension d, for group CA over the cyclic group $\mathbb{G} = \mathbb{Z}_m$ [5]. It is an immediate corollary of Theorem 2 that the latter results can be generalized to all group cellular automata, in any dimension and on any group shift, not just the full shift. Even jointly periodic configurations are dense: a configuration is called *jointly periodic* for a cellular automaton if it is temporally periodic and also totally periodic in space.

Corollary 19. *Let $F : \mathbb{X} \longrightarrow \mathbb{X}$ be a group cellular automaton on a d-dimensional group shift $\mathbb{X}$. Jointly periodic configurations are dense in Ω_F. In particular, if F is surjective then they are dense in $\mathbb{X}$.*

Proof. By Lemma 12 the set $\mathrm{ST}(F)$ of space-time diagrams is a $(d+1)$-dimensional group shift, and by Theorem 2 totally periodic elements are dense in $\mathrm{ST}(F)$. The projection $\pi(c)$ of a totally periodic space-time diagram c on the domain $\mathbb{Z}^d \times \{0\}$ is a totally periodic element of Ω_F that is also temporally periodic. The density of totally periodic space-time diagrams c in $\mathrm{ST}(F)$ implies the density of their projections $\pi(c)$ in $\Omega_F = \pi(\mathrm{ST}(F))$. If F is surjective then $\Omega_F = \mathbb{X}$. □

Injectivity and surjectivity

Another immediate implication of Theorem 2 is a *surjunctivity* property: every injective group cellular automaton $F : \mathbb{X} \longrightarrow \mathbb{X}$ is surjective.

Corollary 20. *Let $F : \mathbb{X} \longrightarrow \mathbb{X}$ be a group cellular automaton on a d-dimensional group shift $\mathbb{X}$. If F is injective then it is surjective.*

Proof. If F is injective then it is injective among totally periodic configurations of $\mathbb{X}$. For any fixed $k > 0$ there are finitely many configurations in $\mathbb{X}$ that are $k\boldsymbol{e}_i$-periodic for all $i \in \{1, \dots, d\}$. These are mapped by F injectively to each other. Any injective map on a finite set is also surjective, so we see that F is surjective among totally periodic configurations of $\mathbb{X}$. By Theorem 2 the totally periodic configurations are dense in $\mathbb{X}$ so that $F(\mathbb{X})$ is a dense subset of $\mathbb{X}$. By the continuity of F it is also closed which means that $F(\mathbb{X}) = \mathbb{X}$. □

We have that every injective group cellular automaton is bijective. Recall that a bijective cellular automaton F is automatically reversible, meaning that the inverse F^{-1} is also a cellular automaton. If F is a reversible group cellular automaton then clearly so is F^{-1}. Reversible cellular automata are of particular interest due to their relevance in modeling microscopic physics and in other application domains [16]. While it is decidable if a given one-dimensional cellular automaton is injective (=reversible) or surjective, the same questions are undecidable for general two-dimensional cellular automata [14]. As expected, the situation is different for group cellular automata.

Theorem 21. *It is decidable if a given group cellular automaton $F : \mathbb{X} \longrightarrow \mathbb{X}$ over a given d-dimensional group shift $\mathbb{X}$ is injective (surjective).*

Proof. By Lemma 17 one can effectively construct the limit set Ω_F. The CA F is surjective if and only if $\Omega_F = \mathbb{X}$. As equality of given group shifts is decidable (Corollary 4(b)), it follows that surjectivity is decidable.

For injectivity, recall that a group homomorphism F is injective if and only if $\ker(F) = \{\mathbf{1}_{\mathbb{X}}\}$. Since $\ker(F)$ is a group shift that can be effectively constructed (Lemma 6), we can check injectivity by checking the equality of the two group shifts $\ker(F)$ and $\{\mathbf{1}_{\mathbb{X}}\}$. □

The Garden-of-Eden-theorem

The Garden-of-Eden-theorem is among the oldest results in the theory of cellular automata. It links injectivity and surjectivity. Let us call two configurations $c, e \in A^{\mathbb{Z}^d}$ *asymptotic* if their *difference set* $\text{diff}(c, e) = \{\boldsymbol{u} \in \mathbb{Z}^d \mid c_{\boldsymbol{u}} \neq e_{\boldsymbol{u}}\}$ is finite. A cellular automaton $F : X \longrightarrow X$ on a subshift X is called *pre-injective* if for any asymptotic c, e the following holds: $c \neq e \implies F(c) \neq F(e)$. So injectivity is only required among mutually asymptotic configurations. Trivially every injective cellular automaton is pre-injective but the converse implication is not true. In fact, the classical Garden-of-Eden-theorem states that on full shifts in any dimension pre-injectivity is equivalent to surjectivity.

Theorem 22 (the Garden-of-Eden-theorem [25, 26]) *A cellular automaton $F : A^{\mathbb{Z}^d} \longrightarrow A^{\mathbb{Z}^d}$ is pre-injective if and only if it is surjective.*

That surjectivity implies pre-injectivity was first proved by E.F.Moore [25], and the converse implication a year later by J.Myhill [26]. Later the theorem has been extended to many other settings. For example, it is known that the Garden-of-Eden-theorem holds for cellular automata over so-called *strongly irreducible* subshifts of finite type [10].

Note that the Myhill direction implies surjunctivity: if a cellular automaton is injective then it is pre-injective and by Myhill's theorem surjective. For group shifts we proved surjunctivity differently in Corollary 20, using the density of periodic points. There is a good reason for this: the Myhill direction of the Garden-of-Eden-theorem is namely not true for all group cellular automata over group shifts, as shown by the following trivial example.

Example 23. *Let* $\mathbb{X} = \{0^{\mathbb{Z}}, 1^{\mathbb{Z}}\}$ *be the two-element group shift over the two-element cyclic group* $\mathbb{Z}_2$*, and let* $F : \mathbb{X} \longrightarrow \mathbb{X}$ *be the group cellular automaton* $F(0^{\mathbb{Z}}) = F(1^{\mathbb{Z}}) = 0^{\mathbb{Z}}$*. Then* F *is pre-injective but not surjective.*

Recall that it is decidable whether a given group cellular automaton is surjective (Theorem 21). Since surjectivity and pre-injectivity are not equivalent for all group cellular automata, a natural follow up question is to determine if a given group cellular automaton is pre-injective. The decidability status of this question remains open.

Question 1. *Is it decidable if a given group cellular automaton is pre-injective ?*

Next we show that the Moore direction of the Garden-of-Eden-theorem holds for all group cellular automata. The proof is based on the fact that all surjective cellular automata preserve entropy while group cellular automata that are not pre-injective do not preserve it. The topological *entropy* of a d-dimensional subshift X is defined as

$$h(X) = \lim_{n\to\infty} \frac{\log|\mathcal{L}_{B_n}(X)|}{|B_n|}$$

where $B_n = \{1, \dots, n\}^d$ is the d-dimensional box of size $n \times \cdots \times n$. It is well known that the limit exists: this can be proved, for example, using a multidimensional version of Fekete's subadditive lemma [4], or using the Ornstein–Weiss lemma [20] in the special case of a grid $\mathbb{Z}^d$.

Theorem 24. *Let* $F : \mathbb{X} \longrightarrow \mathbb{X}$ *be a group cellular automaton over a group shift* $\mathbb{X} \subseteq \mathbb{G}^{\mathbb{Z}^d}$*. If* F *is surjective then* F *is pre-injective.*

Proof. Suppose F is not pre-injective so $F(x) = F(y)$ for an asymptotic pair $x, y \in \mathbb{X}$, $x \neq y$. Then $c = xy^{-1} \in \mathbb{X}$ is asymptotic with $\mathbf{1}_{\mathbb{X}}$ while $c \neq \mathbf{1}_{\mathbb{X}}$ and $F(c) = F(\mathbf{1}_{\mathbb{X}}) = \mathbf{1}_{\mathbb{X}}$. It follows from this fact that the entropy of the kernel of F is strictly positive, $h(\ker(F)) > 0$. However, by Theorem 14.1 in [28] the entropies of the group shifts $\mathbb{X}$, $F(\mathbb{X})$ and $\ker(F)$ satisfy the following addition formula: $h(\mathbb{X}) = h(F(\mathbb{X})) + h(\ker(F))$. We then have that $h(\mathbb{X}) > h(F(\mathbb{X}))$, implying that $\mathbb{X} \neq F(\mathbb{X})$, *i.e.*, that F is not surjective. □

Nilpotency, equicontinuity and sensitivity

A cellular automaton is called *nilpotent* if there is only one configuration in the limit set Ω_F. (Clearly the limit set is never empty.) Nilpotency is undecidable even for cellular automata over one-dimensional full shifts [13, 27]. In the case of group cellular automata the identity configuration is a fixed point and hence automatically in the limit set. Nilpotency of group cellular automata can be easily tested by effectively constructing the limit set (Lemma 17) and testing equivalence with the singleton group shift $\{\mathbf{1}_{\mathbb{X}}\}$.

More generally, a cellular automaton F is *eventually periodic* if $F^{n+p} = F^n$ for some n and $p \geq 1$, and it is *periodic* if F^p is the identity map for some $p \geq 1$. Nilpotent cellular automata are clearly eventually periodic with $p = 1$. Note that eventually periodic cellular automata are periodic on the limit set and, conversely, if F is periodic on its limit set then it is eventually periodic on $\mathbb{X}$ because $\Omega_F = F^n(\mathbb{X})$ for some n by Lemma 18.

Theorem 25. *It is decidable for a given group cellular automaton $G : \mathbb{X} \longrightarrow \mathbb{X}$ on a given d-dimensional group shift $\mathbb{X}$ whether F is nilpotent, periodic or eventually periodic.*

Proof. We have that F is

- nilpotent if and only if $\Omega_F = \{\mathbf{1}_{\mathbb{X}}\}$,
- eventually periodic if and only if $\mathrm{Tr}_{\{\mathbf{0}\}}(F)$ is finite,
- periodic if and only if it is injective and eventually periodic.

Group shifts Ω_F and $\mathrm{Tr}_{\{\mathbf{0}\}}(F)$ can be effectively constructed (Lemma 15 and Lemma 17). Equivalence of Ω_F and $\{\mathbf{1}_{\mathbb{X}}\}$ can be tested (Corollary 4(b)) and finiteness of a given one-dimensional subshift of finite type is easily checked, so nilpotency and eventual periodicity are decidable. By Theorem 21 injectivity of F is decidable so also periodicity can be decided. □

A configuration $c \in \mathbb{X}$ is an *equicontinuity point* of $F : \mathbb{X} \longrightarrow \mathbb{X}$ if for every finite $D \subseteq \mathbb{Z}^d$ there exists a finite $E \subseteq \mathbb{Z}^d$ such that $e|_E = c|_E$ implies $F^n(e)|_D = F^n(c)|_D$ for all $n \geq 0$. Orbits of equicontinuity points can hence be reliably simulated even if the initial configuration is not precisely known. Let $\mathrm{Eq}(F) \subseteq \mathbb{X}$ be the set of equicontinuity points of F. We call F *equicontinuous* if $\mathrm{Eq}(F) = \mathbb{X}$.

Cellular automaton $F : \mathbb{X} \longrightarrow \mathbb{X}$ is *sensitive to initial conditions*, or just *sensitive*, if there exists a finite observation window $D \subseteq \mathbb{Z}^d$ such that for every configuration $c \in \mathbb{X}$ and every finite $E \subseteq \mathbb{Z}^d$ there is $e \in \mathbb{X}$ with $e|_E = c|_E$ but $F^n(e)|_D \neq F^n(c)|_D$ for some $n \geq 0$. Clearly c cannot be an equicontinuity point so for all sensitive F we have $\mathrm{Eq}(F) = \emptyset$. For group cellular automata also the converse holds.

Lemma 26. *Let $F : \mathbb{X} \longrightarrow \mathbb{X}$ be a group cellular automaton over a d-dimensional group shift $\mathbb{X}$. Then exactly one of the following two possibilities holds:*

- $\mathrm{Eq}(F) = \mathbb{X}$ *and F is equicontinuous, or*
- $\mathrm{Eq}(F) = \emptyset$ *and F is sensitive.*

Proof. Assume that some $c \notin \mathrm{Eq}(F)$ exists, which means that there exists a finite $D \subseteq \mathbb{Z}^d$ such that for all finite $E \subseteq \mathbb{Z}^d$ there is $e \in \mathbb{X}$ and $n \geq 0$ with $e|_E = c|_E$ but $F^n(e)|_D \neq F^n(c)|_D$. Consider an arbitrary $c' \in \mathbb{X}$. For $c'' = c'ec^{-1} \in \mathbb{X}$ we then have that $c''|_E = c'|_E$ but $F^n(c'')|_D \neq F^n(c')|_D$. This proves that $c' \notin \mathrm{Eq}(F)$.

We can conclude that for group cellular automata either $\mathrm{Eq}(F) = \mathbb{X}$ or $\mathrm{Eq}(F) = \emptyset$. By definition, $\mathrm{Eq}(F) = \mathbb{X}$ is equivalent to equicontinuity of F.

If F is sensitive then $\mathrm{Eq}(F) = \emptyset$ holds. Conversely, if F is not sensitive then, by definition, for all finite $D \subseteq \mathbb{Z}^d$ there exists $c \in \mathbb{X}$ and a finite $E \subseteq \mathbb{Z}^d$ such that $e|_E = c|_E$ implies that $F^n(e)|_D = F^n(c)|_D$ for all $n \geq 0$. As above, we can replace c by any other configuration c', which implies that all configurations are equicontinuity points, i.e., $\mathrm{Eq}(F) \neq \emptyset$. □

We can decide equicontinuity and sensitivity.

Theorem 27. *It is decidable for a given group cellular automaton $G : \mathbb{X} \longrightarrow \mathbb{X}$ on a given d-dimensional group shift $\mathbb{X}$ whether F is equicontinuous or sensitive to initial conditions.*

Proof. By the dichotomy in Lemma 26 it is enough to decide equicontinuity. Let us show that F is equicontinuous if and only if it is eventually periodic, after which the decidability follows from Theorem 25.

If F is eventually periodic then it is trivially equicontinuous since there are only finitely many different functions F^k, $k \geq 0$, and all these functions are continuous. Conversely, if F is equicontinuous then one easily sees that there are only finitely many different traces in $\mathrm{Tr}_{\{\mathbf{0}\}}(F)$. Indeed, equicontinuity at configuration c implies that there is a finite set $E \subseteq \mathbb{Z}^d$ such that $e|_E = c|_E$ implies that $F^n(e)_{\mathbf{0}} = F^n(c)_{\mathbf{0}}$ for all $n \geq 0$. As in the proof of Lemma 26 we see that the same set E works for all configurations c. But then $|\mathcal{L}_E(\mathbb{X})|$ is an upper bound on the number of different traces in $\mathrm{Tr}_{\{\mathbf{0}\}}(F)$ because $c|_E$ uniquely identifies the positive trace of c (and by the translation invariance of the trace subshift any k different traces can be shifted to provide k different positive traces.)

Finiteness of $\mathrm{Tr}_{\{\mathbf{0}\}}(F)$ implies that all traces are periodic with a common period, so that cellular automaton F is periodic on its limit set. Hence F is eventually periodic. □

Mixing properties

A cellular automaton $F : \mathbb{X} \longrightarrow \mathbb{X}$ is *transitive* if there is an orbit from every non-empty open set to every non-empty open set, that is, if for any finite $D \subseteq \mathbb{Z}^d$ and all $p, q \in \mathcal{L}_D(\mathbb{X})$ there exists $c \in \mathbb{X}$ and $n \geq 0$ such that $c|_D = p$ and $G^n(c)|_D = q$. It is *mixing* if there exists such c for every sufficiently large n, that is, if for all D, p and q as above there is m such that for every $n \geq m$ there exists $c \in \mathbb{X}$ such that $c|_D = p$ and $G^n(c)|_D = q$.

For these properties we obtain only semi-algorithms for the negative instances. Decidability remains open.

Theorem 28. *It is semi-decidable for a given group cellular automaton $G : \mathbb{X} \longrightarrow \mathbb{X}$ on a given d-dimensional group shift $\mathbb{X}$ whether F is non-transitive or non-mixing.*

Proof. A non-deterministic semi-algorithm guesses a finite $D \subseteq \mathbb{Z}^d$, forms the trace subshift $\mathrm{Tr}_D(F)$, and verifies that the trace subshift is not transitive (not mixing, respectively). Clearly F is not transitive (not mixing, respectively) if and only if such a choice of D exists. For one-dimensional subshifts of finite type, such as $\mathrm{Tr}_D(F)$, it is easy to decide transitivity and the mixing property [22]. □

Question 2. *Is it decidable if a given group cellular automaton is transitive (or mixing) ?*

5. Conclusions

We have demonstrated how the "swamp of undecidability" [23] of multidimensional SFTs and cellular automata is mostly absent in the group setting. For general cellular automata nilpotency [13, 27], as well as eventual periodicity, equicontinuity and sensitivity [9] are undecidable on one-dimensional full shifts, and periodicity [17], as well as sensitivity, mixingness and transitivity [24] are undecidable even among reversible one-dimensional cellular automata on the full shift; injectivity and surjectivity are undecidable for two-dimensional cellular automata on the full shift [14]. Algorithms and characterizations have been known for linear and additive cellular automata (on full shifts, sometimes depending on the dimension [7, 8]). Our results improve these to the greater generality of non-commutative groups and cellular automata on higher dimensional subshifts. However, it should be noted that the existing characterizations in the literature typically provide easy to check conditions on the local rule of the cellular automaton for the considered properties, while algorithms extracted from our proofs are impractical and only serve the purpose of proving decidability.

It remains open whether it is decidable if a given group cellular automaton is pre-injective, transitive or mixing.

References

[1] P. Béaur and J. Kari, Decidability in group shifts and group cellular automata, *45th International Symposium on Mathematical Foundations of Computer Science, MFCS 2020, August 24-28, 2020, Prague, Czech Republic*, eds. J. Esparza and D. Král' *LIPIcs* **170**, (Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020), pp. 12:1–12:13.

[2] F. Blanchard and P. Tisseur, Some properties of cellular automata with equicontinuity points, *Annales de l'Institut Henri Poincare (B) Probability and Statistics* **36**(5) (2000) 569–582.

[3] M. Boyle and B. Kitchens, Periodic points for onto cellular automata, *Indagationes Mathematicae* **10**(4) (1999) 483–493.

[4] S. Capobianco, Multidimensional cellular automata and generalization of Fekete's lemma, *Discrete Mathematics & Theoretical Computer Science* **10**(3) (2008).

[5] G. Cattaneo, A. Dennunzio and L. Margara, Solution of some conjectures about topological properties of linear cellular automata, *Theoretical Computer Science* **325**(2) (2004) 249–271.

[6] T. Ceccherini-Silberstein and M. Coornaert, *Cellular Automata and Groups*Springer Monographs in Mathematics, Springer Monographs in Mathematics (Springer Berlin Heidelberg, 2010).

[7] A. Dennunzio, E. Formenti, L. Manzoni, L. Margara and A. E. Porreca, On the dynamical behaviour of linear higher-order cellular automata and its decidability, *Information Sciences* **486** (2019) 73–87.

[8] A. Dennunzio, E. Formenti, D. Grinberg and L. Margara, Dynamical behavior of additive cellular automata over finite abelian groups, *Theoretical Computer Science* **843** (2020) 45–56.

[9] B. Durand, E. Formenti and G. Varouchas, On undecidability of equicontinuity classification for cellular automata, *Discrete Models for Complex Systems, DMCS'03, Lyon, France, June 16-19, 2003*, eds. M. Morvan and E. Rémila *DMTCS Proceedings* **AB**, (DMTCS, 2003), pp. 117–128.

[10] F. Fiorenzi, Cellular automata and strongly irreducible shifts of finite type, *Theoretical Computer Science* **299**(1) (2003) 477 – 493.

[11] G. A. Hedlund, Endomorphisms and automorphisms of the shift dynamical systems, *Mathematical Systems Theory* **3**(4) (1969) 320–375.

[12] M. Hochman, On the dynamics and recursive properties of multidimensional symbolic systems, *Inventiones Mathematicae* **176**(1) (2009) 131–167.

[13] J. Kari, The nilpotency problem of one-dimensional cellular automata, *SIAM Journal on Computing* **21**(3) (1992) 571–586.

[14] J. Kari, Reversibility and surjectivity problems of cellular automata, *Journal of Computer and System Sciences* **48**(1) (1994) 149 – 182.

[15] J. Kari, Theory of cellular automata: A survey, *Theoretical Computer Science* **334**(1-3) (2005) 3–33.

[16] J. Kari, Reversible cellular automata: From fundamental classical results to recent developments, *New Generation Computing* **36**(3) (2018) 145–172.

[17] J. Kari and N. Ollinger, Periodicity and immortality in reversible computing, *Mathematical Foundations of Computer Science 2008, 33rd International Symposium, MFCS 2008, Torun, Poland, August 25-29, 2008, Proceedings*, eds. E. Ochmanski and J. Tyszkiewicz *Lecture Notes in Computer Science* **5162**, (Springer, 2008), pp. 419–430.

[18] B. Kitchens and K. Schmidt, Periodic points, decidability and Markov subgroups, *Dynamical Systems*, ed. J. C. Alexander (Springer Berlin Heidelberg, Berlin, Heidelberg, 1988), pp. 440–454.

[19] B. Kitchens and K. Schmidt, Automorphisms of compact groups, *Ergodic Theory and Dynamical Systems* **9**(4) (1989) p. 691–735.

[20] F. Krieger, *Le lemme d'Ornstein–Weiss d'après Gromov*, *Dynamics, Ergodic Theory and Geometry* ed. B. HasselblattMathematical Sciences Research Institute Publications Mathematical Sciences Research Institute Publications, (Cambridge University Press, 2007), pp. 99–112.

[21] P. Kůrka, *Topological and Symbolic Dynamics*Collection SMF, Collection SMF (Société mathématique de France, 2003).

[22] D. Lind and B. Marcus, *An Introduction to Symbolic Dynamics and Coding* (Cambridge, 1995).

[23] D. A. Lind, Multi-dimensional symbolic dynamics, *Symbolic Dynamics and its Applications*, ed. S. Williams *AMS Short Course Lecture Notes*, (American Mathematical Society, 2004), pp. 61–79.

[24] V. Lukkarila, Sensitivity and topological mixing are undecidable for reversible one-dimensional cellular automata, *Journal of Cellular Automata* **5**(3) (2010) 241–272.

[25] E. F. Moore, Machine models of self-reproduction, *Proceedings Symposium Applied Mathematics* **14** (1962) 17–33.
[26] J. Myhill, The converse to Moore's garden-of-eden theorem, *Proceedings of the American Mathematical Society* **14** (1963) 685–686.
[27] H. L. S. Aanderaa, Linear sampling and the ∀∃∀ case of the decision problem, *The Journal of Symbolic Logic* **39** (1974) 519–548.
[28] K. Schmidt, *Dynamical systems of algebraic origin*Progress in mathematics, Progress in mathematics (Birkhäuser Verlag, 1995).
[29] H. Wang, Proving theorems by pattern recognition – II, *The Bell System Technical Journal* **40**(1) (1961) 1–41.