Subshifts with sparse projective subdynamics

Ville Salo

University of Turku
vosalo@utu.fi

August 1, 2016

Abstract

We study two-dimensional subshifts whose projective subdynamics contains only points of finite support. Our main result is a classification result for such subshifts satisfying a minimality property. As corollaries, we obtain new proofs for various known results on projective subdynamics of SFTs, nilpotency and decidability of cellular automata, topological full groups and the subshift of prime numbers. We also construct various (sofic) examples illustrating the concepts.

1 Introduction

Multidimensional subshifts are $\mathbb{Z}^d$-actions by translations on closed subspaces of $\Sigma^{\mathbb{Z}^d}$. The best-known examples of such are subshifts of finite type (SFTs), which are isomorphic to sets of tilings by square tiles with adjacency contraints. Much of the theory of multidimensional subshifts revolves around the phenomenon that the set of tilings, while defined by finitely many local constraints, can have complex dynamical and computational properties [7, 33, 25, 10, 22, 14]. While SFTs can be highly complex, they have severe limitations that follow trivially from their definition by local rules – for example, an SFT containing a nontrivial finite configuration (that is, a configuration of finite support) has positive entropy.

Another popular class are the sofic shifts, images of SFTs through factor maps. These subshifts are already much more general. For example, substitutive subshifts (satisfying some technical conditions) are sofic [29], and the projective subdynamics (sets of rows appearing in configurations) of sofic shifts exactly coincide with the class of computable subshifts [20, 13, 2]. The latter in particular shows that the theory of multidimensional SFTs, and especially sofics, leads quite naturally to general computable subshifts. While not all computable subshifts are sofic, understanding which of them are is an active research area [17, 30], and many simple-looking questions are open.

Cellular automata are another point of view to SFTs, and cellular automata are, essentially, SFTs with a rational deterministic direction. Similarly as with SFTs, there is much freedom in constructing cellular automata; many properties turn out to be undecidable, and many kinds of behavior are possible for cellular automata [24, 26].

However, also limitations are known. In this paper, we concentrate on limitations arising from having very simple projective subdynamics,[1] and from various types of nilpotency. A well-known result of this type is that the projective subdynamics of an infinite SFT cannot be sparse, that is, it cannot consist of only configurations with finitely many nonzero symbols. This is a corollary of the full characterization of possible sofic projective subdynamics of SFTs given in [32].

Other similar results in the nilpotency framework appear in [18] and [36], where the assumption of SFTness is in some sense replaced with determinism, and the multidimensional subshift is not discussed explicitly. In [18] it is in particular shown that on certain one-dimensional SFTs, asymptotically nilpotent cellular automata are nilpotent, and in [36] a result of similar flavor is proved for CA on countable sofic shifts, in particular showing that nilpotency is decidable on these subshifts.

This paper is an attempt to clarify and unify such results; our main result gives new proofs for the results of [32, 18, 36] mentioned in the previous paragraphs.[2] The result applies to all subshifts with sparse projective subdynamics, not only SFTs or sofic subshifts. We also contribute two techniques that we find quite universally helpful in structuring such proofs, namely restricting to an almost minimal subsystem and studying paths drawn on configurations.

The corollaries listed above are discussed in more detail in Section 1.2. In addition to them, we obtain for example results that apply to subshifts where the projective subdynamics is sparse in an irrational direction and some observations about expansive directions in sparse subshifts. We also extract results for topological full groups, and show that the subshift of prime numbers contains either a finite point or a blob fractal for purely topological reasons.

We also make tiny contributions also to the construction side of multidimensional subshifts, and show that general subshifts (and already sofic shifts) with sparse projective subdynamics can be quite nontrivial. We summarize them in Section 1.3.

1.1 The main theorem

Every minimal subshift can be endowed with a substitutive structure by finding its words of length k and building its longer words by gluing together words of smaller length. There are many ways to formalize this, for example Bratteli-Vershik diagrams [19].

An *almost minimal* subshift is one containing a fixed-point for the dynamics, with the property that every point except the fixed one generates the whole subshift. It is easy to give a substitutive structure to such subshifts, in the following sense (proved as Theorem 11), which is a one-dimensional version of our main theorem:

Theorem 1. *Let X be any almost minimal one-dimensional subshift. Then X is the orbit-closure of a finite point or a blob fractal.*

Our main result is Theorem 2 below (proved in Theorem 12). It classifies two-dimensional subshifts that are almost minimal, and whose projective

[1]Of course there are many other known restrictions that force subshifts and cellular automata to behave, for example countability [3, 4], expansivity of CA [8] and algebraicity [39].

[2]All three papers also contain many other results, which we do not reproduce.

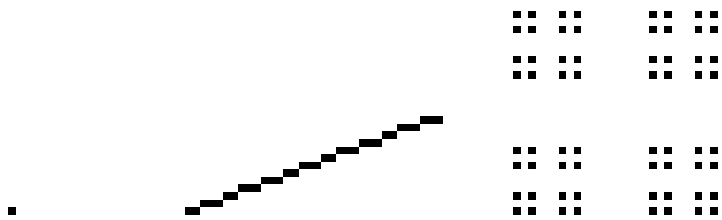

(a) A finite configuration: the configuration with a single nonzero symbol.

(b) A highway configuration: the discretization of an irrational line.

(c) A blob fractal: a subshift based on Cantor's dust.

Figure 1: Folklore examples of the three cases that can appear in Theorem 2. The projective subdynamics of the blob fractal given here is not actually sparse, as we are not aware of such examples in the literature – see Section 7 for a blob fractal with sparse projective subdynamics.

subdynamics is *sparse*, that is, contains only configurations with finitely many nonzero symbols.

Theorem 2. *Let X be any almost minimal two-dimensional subshift with sparse projective subdynamics. Then X is the orbit-closure of a finite point, a highway, or a blob fractal.*

As a corollary, we find these kinds of configurations in all subshifts where there is suitable directional convergence to a fixed point. The following is our main extraction result of this type, proved in Theorem 13:

Theorem 3. *Let X be any two-dimensional subshift whose projective subdynamics Y satisfies one of the following:*

- *every point in Y is eventually zero to the right, or*
- *Y is countable, and the only periodic point in Y is $0^{\mathbb{Z}}$.*

Then X contains a finite point, a highway, or a blob fractal whose projective subdynamics is sparse.

In these theorems, a *finite point* is a point whose support is finite (that is, all but finitely many cells contain the symbol 0). A *highway* is a *uniformly recurrent ascending path configuration*, that is, a configuration whose support consists of a single ascending path of bounded width whose movement is guided by a uniformly recurrent sequence. *Blob fractals* are recursively defined configurations where finite patterns padded with zeroes (blobs) are collected into larger finite patterns with larger and larger separation.[3] For the definition of blob fractals, see Section 2, and for paths and highways see Section 4. See Figure 1 for illustrations.

The way we prove Theorem 2 is by taking an arbitrary nonzero point in the subshift X (as every point generates the subshift), and attempting to build a

[3]In this sense blob fractals resemble the observable universe where planets orbit stars, which are collected into galaxies, which form galaxy groups, and then superclusters, though we have opted for less intimidating terminology.

blob fractal structure for the point. This process ends if the point is finite, or if its support contains an infinite path. We then show that every sparse almost minimal subshift where the support of some configuration contains an infinite path in fact contains only uniformly recurrent ascending paths of uniformly bounded width, i.e. highways, which is the most technical part of the proof.

Intuitively, once we have a configuration containing a path, we first show that it can be made ascending, and then a bounded-width ascending path configuration is found as follows: Since the subshift has sparse projective subdynamics, there is a global limit on the number of parallel paths one can find in a configuration. So we take a maximal number of paths that can be laid parallel to each other (with arbitrarily large separation) in a configuration. Then we observe that if the width of one of these paths is not bounded, we can extract one more path, which is a contradiction. It follows that all the paths must be of bounded width, and we have found our path. By finding a uniformly recurrent point in the orbit-closure of the path, we obtain a highway.

Our formalization of the notion of 'path' is quite explicit. Namely, we study dynamical systems of paths as objects in their own right in Section 4, and then study the path covers of subshifts, where paths are overlaid on top of nonzero symbols of the configuration. This allows a relatively direct translation of the above idea to a proof.

1.2 Corollaries

In this section, we briefly discuss four theorems from the literature that are given new proofs in this paper using Theorem 2. The first one is the following theorem of Pavlov and Schraudner [32, Theorem 6.4]:

Theorem 4. *If a $\mathbb{Z}$-subshift Y has universal period and is not a finite union of periodic points, then it is not the $\mathbb{Z}$-projective subdynamics of any $\mathbb{Z}^2$-SFT X.*

Universal period means that every point is periodic, except possibly at a bounded number of cells. Reducing this to Theorem 2 is straightforward, but somewhat technical, and we do this in Theorem 15. However, the following weaker version of Theorem 4 is a direct corollary of Theorem 2:

Theorem 5. *If a nontrivial $\mathbb{Z}$-subshift Y is sparse, then it is not the $\mathbb{Z}$-projective subdynamics of any $\mathbb{Z}^2$-SFT X.*

This follows from Theorem 2 by restricting to an almost minimal subshift of X and observing that horizontal translates of configurations of any of the three types in Theorem 2 (finite configurations, paths or blob fractals) can be freely glued within the SFT to obtain non-sparse projective subdynamics (even positive entropy). This is proved in Theorem 14.

Nilpotency is an important notion in the theory of cellular automata, as it is perhaps the best-known undecidable property for cellular automata on the full shift [23, 1]. The next theorem is due to Guillon and Richard [18, Theorem 4]. A CA is *nilpotent* if every point is mapped in finite time to the all-zero configuration, and *asymptotically nilpotent* if every configuration tends to the all-zero configuration.

Theorem 6. *Let X be a one-dimensional transitive SFT and $f : X \to X$ a cellular automaton. Then f is nilpotent if and only if it is asymptotically nilpotent.*

This is a direct corollary of the following, proved in Proposition 9, which in turn is a direct corollary of Theorem 2. By a *glider* we mean a finite configuration that is shifted by some power of the cellular automaton.

Proposition 1. *If $f : X \to X$ be a cellular automaton on a one-dimensional subshift X. If either f is asymptotically nilpotent, the limit set of f is sparse, or the closure of the asymptotic set of f is sparse, then f has a glider.*

An asymptotically nilpotent CA on an SFT cannot have a glider, as gluing infinitely many gliders together would clearly give a non-nilpotent configuration. Thus Theorem 6 follows from Proposition 1.

We also state two decidability results obtained by the author and coauthors. The following is one of the main decidability results in [36].

Theorem 7. *If X is a one-dimensional countable sofic shift, then nilpotency of cellular automata on X is decidable.*

After characterizing countable sofic shifts as ones consisting of periodic patterns and finitely many transitions between them, this theorem follows from Proposition 1 by observing that the existence of gliders for a CA on a sofic shift is semidecidable, and nilpotency is semidecidable.

Finally, we mention an observation made in [5] about the topological full group of a full shift, which follows from Theorem 2, after interpreting the topological full group of a one-dimensional subshift as a two-dimensional subshift in a suitable way.

Theorem 8. *If X is a full shift, then the torsion problem of the topological full group of X is decidable.*

The one-dimensional version of our main theorem also has some corollaries of interest. We use it to show that one can say something about the possible patterns in prime numbers using almost none of their specific properties. Let $X_{\mathcal{P}}$ be the subshift whose language contains $w \in \{0,1\}^*$ if and only if for infinitely many n, we have $\forall a \in [0, |w|-1] : n + a \text{ is prime} \iff w_a = 1$.

Proposition 2. *The subshift $X_{\mathcal{P}}$ contains either a finite point or a blob fractal.*

We also give a number-theoretic proof of this fact.

1.3 Constructions

In Section 7, we give some constructions of subshifts with sparse projective subdynamics and path spaces.

To prove the theorems listed in the previous section, we only use the first two cases of the classification result Theorem 2, as in all those cases blob fractals turn out to be impossible. However, in general blob fractals do exist, and in some sense they are the more common case, as the other two cases are in some sense degenerate blob fractals (finite points are blob fractals where the separation between blobs is infinite, and infinite paths are blobs of infinite size). We show in particular that subshifts with sparse subdynamics generated by blob fractals need not contain any nonzero configurations where every row has bounded width.

We also discuss the computational side of subshifts with sparse projective subdynamics. We show that a subshift where the support of every point is an ascending path is Π^0_1 if and only if it is sofic when the movement speed of the path is bounded, but construct a Π^0_1 non-sofic subshift with sparse projective subdynamics.

Our main technical tool besides almost minimality are uniformly recurrent paths. Paths are a standard object in the dynamical systems literature, and usually go by the name *cocycle* [15, Definitions 2.1]. We prove a classification of minimal path spaces, and show that there are four types of such spaces: every path is ascending, every path is descending, every path is bounded, or some path enters the origin infinitely many times. In Section 7, we show how to build representative examples of paths in each of these classes, and show that the last class splits further into subclasses.

1.4 Unpublished related work

We also mention some unpublished but related results. In [12], a result orthogonal to Theorem 2 is proved: if $f : X \to X$ is a CA on a one-dimensional subshift X and there is a point $x \in X$ that never maps to the all-zero point, and there is a positive lower density of times $N \subset \mathbb{N}$ such that $f^n(x)$ has at most k nonzero symbols, then f has a glider. The proof is of similar flavor as that of [36, 32] and the one in this paper.

An unpublished result of Pierre Guillon and Pierre-Étienne Meunier states the following: Let f be a cellular automaton on a one-dimensional full shift whose neighborhood does not grow in cardinality when the CA is iterated.[4] Then the neighborhoods of f^n are contained in finitely many arithmetic progressions.

2 Definitions

Points are elements of $\Sigma^{\mathbb{Z}^d}$. A *unary point* is a point $x \in \Sigma^{\mathbb{Z}^d}$ with $\forall i \in \mathbb{Z}^d : x_i = a$. We write this as $a^{\mathbb{Z}^d}$. The *shift*[5] or *translation* on $\Sigma^{\mathbb{Z}^d}$ is $\sigma^v(x)_u = x_{u+v}$. We also call points *configurations*.

See [27] for a reference on symbolic dynamics. A topologically closed shift-invariant subset of $\Sigma^{\mathbb{Z}^d}$ is called a *subshift*, and more generally we call expansive actions on Cantor spaces subshifts. Two subshifts are *conjugate* if there is a shift-commuting homeomorphism between them. Conjugate subshifts have the same dynamical properties. Most subshifts considered in this article contain the unary point $0^{\mathbb{Z}^d}$. In particular, the alphabets of our subshifts almost always contain 0, and in almost all notions where a 'zero symbol' is needed, this zero symbol is 0.[6] We write $\Sigma = \Sigma_+ \cup \{0\}$ where $\Sigma_+ \not\ni 0$ is the set of *nonzero symbols*. In this article, a *nontrivial* subshift is one containing at least two points. A *finite point* is one where only finitely many nonzero symbols occur. Such points are also known as *homoclinic*.

[4] $\exists m : \forall n : f^n(x)_0$ depends on at most m not necessarily adjacent cells of x.

[5] Our shift convention is a standard one in the symbolic dynamics of abelian groups, where the contents of points are actually shifted in direction $-v$ by the action of the vector v. This action is contravariant in general, but it is of course covariant in the abelian case.

[6] One may find it useful to think of our subshifts as pointed, but note that our special point is always a fixed-point.

Perhaps the most important subshifts are the *all-zero* or *trivial* subshift $\{0^{\mathbb{Z}}\}$, and the subshift $\overline{\mathcal{O}(\cdots 00100 \cdots)}$, called the *sunny-side-up subshift*, *one-one subshift* or *one-or-less subshift*, consisting of binary points with at most one nonzero symbol.

A *pattern* is a function $P : D \to \Sigma$ where $D = D(P) \subset \mathbb{Z}^d$ is a (possibly infinite) subset called the *domain* of P. Patterns are closely related to *cylinders*. If P is a pattern, we define the cylinder $[P] = \{x \mid x|_{D(P)} = P\}$. However, we more commonly use the terminology that a pattern *occurs* or *appears* in a point, or in another pattern, and by this we mean that a translate of the point is in the cylinder defined by the pattern. We denote this by $P \sqsubset x$. We say P *occurs at* v *in* x if $\sigma_v(x)|_{D(P)} = P$. This allows the use of descriptive terminology such as 'P occurs at a bounded distance from every occurrence of Q', which is an important concept for us. In the one-dimensional case, patterns are usually called *words*, and the set of words occurring in a one-dimensional subshift (or a point) is called its *language*.

Patterns with domain $\{i\} \times \mathbb{Z}$ are called *columns* and ones with domain $\mathbb{Z} \times \{i\}$ are called *rows*. We identify rows and columns with one-dimensional configurations in an obvious way. The *projective subdynamics* of a subshift $X \subset \Sigma^{\mathbb{Z}^2}$ is the set of configurations $y \in \Sigma^{\mathbb{Z}}$ such that y is a row in a point of X. We mostly stay in the two-dimensional case, but for $\mathbb{Z}^d$-subshifts, by projective subdynamics we mean the one-dimensional subshift on the first axis; in particular one-dimensional subshifts are their own projective subdynamics.

The projective subdynamics is not the same as the dynamical system obtained by considering a subaction by a subgroup of $\mathbb{Z}^2$ – the projective subdynamics is always expansive, but subactions need not be. Note that the projective subdynamics is not a conjugacy-invariant, but most properties of the projective subdynamics that are of interest to us – for example, sparseness, boundedness, countability and points being eventually zero – are.

On $\mathbb{Z}^d$, we use the metric $d(u, v) = \sum_{i=1}^{d} |u_i - v_i|$, the word metric when $\mathbb{Z}^d$ is seen as a Cayley graph with the standard generators. Write $B_r(v)$ for the ball of radius r around $v \in \mathbb{Z}^d$, $B_r(v) = \{u \mid d(u, v) \leq r\}$. Write $B_r(S) = \bigcup_{s \in S} B_r(s)$. If P is a pattern, write $S(P) \subset D(P)$ for the *support* of P, that is, the set of cells in $D(P)$ where P has a nonzero symbol.

A subset S of $\mathbb{Z}^d$ is *r-connected* if the undirected graph with nodes S and edges $\{(u, v) \mid d(u, v) \leq r\}$ is connected. The *r-components* of a subset of $\mathbb{Z}^d$ are defined as its maximal r-connected subsets.

A pattern P is *r-padded* if $D(P) \supset B_r(S(P))$, that is, P contains all cells that are at most r away from the support of P. An *r-blob* is an r-padded pattern

- whose the domain's nonzero cells are r-connected and
- whose domain's zero cells are all at most r away from the support.

Note that a blob is uniquely determined by its support. In every configuration, every nonzero cell i is in precisely one r-blob, for every r: the blob is the r-connected component of the nonzero cells that i belongs to, r-padded with 0.

The *zero-gluing* of two patterns P, Q is a pattern obtained by taking their 'disjoint' union configuration, in the sense that only zero symbols may come from both patterns. More precisely, R is a zero-gluing of P and Q if we have $D(R) = D(P) \cup D(Q)$, $P|_{D(P) \cap D(Q)} = Q|_{D(P) \cap D(Q)} = 0^{D(P) \cap D(Q)}$, and

$$R_v = a \iff P_v = a \vee Q_v = a.$$

We sometimes write this as $R = P + Q$, and write $x + y$ for the gluing of two points when the full patterns they define on $\mathbb{Z}^2$ can be glued.

A subshift *allows zero-gluing* if there exists r such that the zero-gluing of any two r-padded patterns that occur in the subshift occurs in the subshift. Note that every SFT allows zero-gluing. Not every strongly irreducible subshift allows zero-gluing: On $\mathbb{Z}$, the *even shift* [27] is such an example, and in higher dimensions we can take the subshift where every row independently contains points from the even shift.

The *essential width* of a two-dimensional configuration x is the minimal m such that for some r the support $S \subset \mathbb{Z}$ of every row of x is contained in the union of m r-balls, that is, $\exists n_1, \dots, n_m : S \subset \bigcup_i n_i + B_r$.

Definition 1. *A configuration $x \in (\Sigma \cup \{0\})^{\mathbb{Z}^d}$ is a* blob fractal *if there is an increasing sequence $r_1, r_2, r_3, \dots$ and for all r_i a finite set B_i of r_i-blobs such that for all i,*

- *each blob in B_{i+1} is obtained by zero-gluing from translates of B_i-blobs,*
- *each blob in B_{i+1} contains a translate of every B_i-blob,*
- *each blob in B_{i+1} contains at least two B_i-blobs with disjoint supports, and*
- *the support of x is contained in the limit of B_j-blobs as $j \to \infty$.*

A *cellular automaton* is a shift-commuting continuous function $f : X \to X$ on a subshift X. We say f is *nilpotent* if $\exists n : \forall x \in X : f^n(x) = 0^{\mathbb{Z}}$ and *asymptotically nilpotent* if $f^n(x) \to 0^{\mathbb{Z}}$ for all $x \in X$. The *limit set* of f is $\bigcap_{n=0}^{\infty} f^n(X)$, and the *asymptotic set* is $\bigcup_{x \in X} \bigcap_{n=0}^{\infty} \overline{\bigcup_{m \geq n} f^m(x)}$.

The *spacetime subshift* of a cellular automaton $f : X \to X$ is the subshift Y whose projective subdynamics is the limit set of f, and in every configuration, the $(i+1)$th row is the f-image of the ith row for every $i \in \mathbb{Z}$. A subshift X is *deterministic* in direction $d \in \mathbb{R}^2$ if, for every half-plane $H \subset \mathbb{R}^2$ whose boundary is perpendicular to d, the map $x \mapsto x_{H \cap \mathbb{Z}^2}$ is injective on X. A *half-plane in direction* d is a half-plane whose boundary is orthogonal to d and whose interior contains all vectors $-rd$ for large enough $r \in \mathbb{N}$.

We sometimes use English words for directions: $(1, 0)$ is right, $(0, 1)$ is up. With cellular automata, we also use the terminology that application of the CA is 'temporal movement', and applying the shift means 'spatial' movement. In particular, *spatial periodicity* for $x \in X$ means $\exists n : \sigma^n(x) = x$, and *temporal periodicity* means $f^n(x) = x$ for some $x \in X$. When we say x is *eventually periodic* we typically refer to spatial periodicity, and mean that $x = \dots uuuvwww \dots$ for some words u, v, w.

3 Almost minimality and projective subdynamics

A subshift is minimal if every point generates it. Because rowsparse subshifts always contain the all-zero point, they clearly cannot be minimal. The next best thing is almost minimality – every point except the all-zero point generates the subshift. It turns out that we can generally extract almost minimal subshifts from our subshifts of interest, after which obtaining the main results of this paper is just a matter of characterizing almost minimal rowsparse subshifts.

Definition 2. *Let X be a compact zero-dimensional metric space and let G be a finitely-generated discrete group acting on X by a continuous map $\sigma : G \times X \to X$. Then (X, G, σ) is called a* G-system. *The* subsystem poset *$\mathcal{S}(X)$ of X is the poset whose elements are subsystems of X with order $Y \leq Z \iff Y \subset Z$. We say X is* minimal *if it is an atom of this poset, that is, $\mathcal{S}(X) = \{\emptyset, X\}$, and* essentially almost minimal *if $\mathcal{S}(X) = \{\emptyset, Y, X\}$ for some one-point system $Y = \{y\}$.*

Almost minimality is discussed in [11]. We called systems with finite $\mathcal{S}(X)$ *quasiminimal* in [34]. Systems with a unique minimal subsystem are called *essentially minimal* [19]. Clearly almost minimal subshifts are essentially minimal and quasiminimal. We mention that in [11], it is shown that minimal locally compact but non-compact dynamical systems on Cantor spaces are in one-to-one correspondence with almost minimal compact Cantor systems, through the one-point compactification (although we make no direct use of this fact).

Definition 3. *A $\mathbb{Z}^d$-subshift is* uniformly rowsparse *if its projective $\mathbb{Z}$-subdynamics is k-sparse for some k, where a $\mathbb{Z}$-subshift is k*-sparse *if every configuration has at most k nonzero symbols. If every row has finitely many nonzero symbols, we say it is* (nonuniformly) rowsparse. *We say that a subshift is* rowcountable *if its projective subdynamics is a countable subshift.*

A rowcountable subshift obviously need not be rowsparse. A rowsparse subshift need not be uniformly rowsparse:

Example 1: There is a sparse subshift that is not uniformly sparse: Let $x^n \in \{0,1\}^{\mathbb{Z}}$ be the point where $x^n_i = 1 \iff i \in n\mathbb{Z} \cap [0, n^2]$. Then $\overline{\mathcal{O}(\{x^n \mid n \in \mathbb{N}\})}$ is sparse, but not uniformly sparse. △

In the next sections, we give some technical tools that are helpful when applying the main theorem outside the almost minimal sparse setting – namely by giving sufficient conditions that allow us to extract almost minimal sparse subshifts from a multidimensional subshift. In Section 3.1, we give sufficient conditions for the extractability of an almost minimal subsystem. In Section 3.2, we show that among almost minimal subshifts, sparsity is implied by various weaker properties. These conditions are somewhat different, and their intersection gives the following lemma.

Lemma 1. *Let a subshift $X \subset \Sigma^{\mathbb{Z}^2}$ with projective subdynamics Y satisfy one of the following:*

- *Y is sparse,*
- *Y is countable and the only periodic point in Y is $0^{\mathbb{Z}}$, or*
- *every point of Y is eventually zero to the right.*

Then X contains an almost minimal uniformly rowsparse subsystem.

3.1 Extracting an almost minimal subsystem

Lemma 2. *Let X be an essentially minimal $\mathbb{Z}^d$-subshift containing the point $\{0^{\mathbb{Z}^d}\}$. Then X contains an almost minimal subsystem.*

Proof. Since $\{0^{\mathbb{Z}^d}\}$ is minimal, it is the unique minimal subsystem. If $(X_i)_{i\in I}$ is a decreasing net of subshifts containing a nonzero symbol, then by the pigeonhole principle some symbol $a \in \Sigma$ occurs in these subshifts for arbitrarily large $i \in I$, and thus for all $i \in I$. The intersection $C_i = X_i \cap [a]$ is a nonempty compact set for all i, and $(C_i)_{i\in I}$ is then a decreasing net of nonempty closed sets. By compactness there exists $x \in \bigcap_i C_i \subset X_i$. The result follows from Zorn's lemma. □

Lemma 3. *Let a subshift $X \subset \Sigma^{\mathbb{Z}^d}$ with projective subdynamics Y satisfy one of the following:*

- *Y is countable and the only periodic point in Y is $0^{\mathbb{Z}}$,*
- *every point of Y is eventually zero to the right*
- *the upper Banach density of nonzero symbols is 0 in every configuration of Y, or*
- *every point of X contains arbitrarily large balls of zeroes.*

Then X contains an almost minimal subsystem.

Proof. Let us show that the first condition implies the fourth (so that each of the first three conditions implies the last one). Every subsystem of Y is countable, and thus contains a minimal subsystem that must be countable. Applying this to orbit closures of points, it follows that every point has $0^{\mathbb{Z}}$ in its orbit closure, and thus contains arbitrarily long words over 0. Since Y^k is countable for all k, it also has only the periodic point $(0^{\mathbb{Z}})^k$, and we obtain the fourth condition.

If $Z \subset X$ is minimal, and contains a configuration with a nonzero symbol, then that symbol appears with bounded gaps (in a syndetic subset of coordinates) in every configuration of Z. This clearly contradicts the last condition. It follows that if any of the four conditions holds, the only minimal subsystem of X is $\{0^{\mathbb{Z}^d}\}$, so we can apply the previous lemma. □

One can strengthen this lemma to saying that every nontrivial subsystem of X with one of these properties has an almost minimal subsystem (though it may not be the same almost minimal subsystem – consider any union of two almost minimal subsystems sharing the fixed point).

One can ask, then, to what extend this stronger version has a corollary, that is, if every subsystem of X has an almost minimal subsystem, which of the conditions does it satisfy? The fourth condition is clearly also necessary: if there is a configuration not containing a ball of zeroes of radius r, then its orbit closure cannot be almost minimal (with fixed point $0^{\mathbb{Z}^2}$). The first two conditions are obviously not necessary:

Example 2: There exists a subshift $X \subset \Sigma^{\mathbb{Z}^2}$ which is almost minimal and whose projective subdynamics is $\{0^{\mathbb{Z}}, 1^{\mathbb{Z}}\}$. Namely, take any one-dimensional almost minimal subshift Y (for example, a subshift generated by a finite point) and consider the two-dimensional subshift with constant rows, and columns taken from Y. △

Slightly more interestingly, one can also build almost minimal subshifts where the upper Banach density of nonzero symbols is positive in every nonzero point of the projective subdynamics (see Section 7).

The first two conditions of Lemma 3 are not comparable for general subshifts: The subshift generated by the characteristic function of $\{\pm 2^i \mid i \in \mathbb{N}\}$ is countable, has only the periodic point $0^{\mathbb{Z}}$, and contains points that are not eventually zero in either direction. Conversely, for every $y \in \{0,1\}^{\mathbb{N}}$ construct the point $x^y \in \{0,1\}^{\mathbb{Z}}$ with support $\{-2^{n+1} + y_n \mid n \in \mathbb{N}\}$. The set of these points generates an uncountable subshift where every point is eventually zero to the right.

3.2 Rowsparsity of almost minimal subshifts

Lemma 4. *Let X be an almost minimal subshift with countable projective subdynamics Y such that Y has no isolated periodic points. Then X is uniformly rowsparse.*

Proof. Let Y be the projective subdynamics of X. Because Y is countable, it is not perfect, and thus contains an isolated point y with isolating pattern w which appears in a point of X. By the assumption, w is not all zero, since $0^{\mathbb{Z}}$ is a non-isolated periodic point in Y. Then by almost minimality, whenever a point $x \in X$ has a nonzero symbol on a row, w must appear a bounded distance away, forcing that row to be a translate of y. If $z \in Y$ is not a finite configuration, and z is a row in x, then at a bounded distance of each of the infinitely many nonzero symbols of z in x the isolating pattern w appears. In particular by the pigeonhole principle y must be periodic. But then y is an isolated periodic point, which contradicts the assumption. Thus Y must be uniformly sparse. □

It is trivial to find subshifts with countable projective subdynamics that do not contain almost minimal subshifts. An example is the two-element group $\mathbb{Z}_2$ with the $\mathbb{Z}^2$-action $\sigma_{(m,n)}(a) = a + m$.

Lemma 5. *Let X be an almost minimal subshift where every row is eventually zero to the right. Then X is uniformly rowsparse.*

Proof. Let Y be the projective subdynamics of X. If every word $w \sqsubset Y$ can be extended to the right by a nonzero symbol, then clearly Y contains a configuration that is not eventually zero to the right. Thus, there is a word w to the right of which only 0 can occur. Such a word occurs at a bounded distance from every nonzero symbol of every configuration of X, and we see as in the proof of the lemma above that X is rowsparse. □

4 Paths and path covers

4.1 Paths

Let `Path` be the space of all functions $p : \mathbb{Z} \to \mathbb{Z}$ such that $p(0) = 0$ and $|p(i+1) - p(i)|$ is bounded. We give it the product topology and the shift action $\tau : \texttt{Path} \to \texttt{Path}$ defined by $\tau(r)(i) = r(i+1) - r(1)$. Let `DPath` be the space of all functions $p : \mathbb{Z} \to \mathbb{Z}$ with $p(\mathbb{Z})$ finite. We give it the product topology and the usual shift action. The two spaces are conjugate by the discrete derivative $\phi : \texttt{Path} \to \texttt{DPath}$ defined by $\phi(p)(i) = p(i+1) - p(i)$. All definitions apply to both kinds of paths, in the sense that if we have defined a property P for elements of Path a path $p \in$ DPath has property P if $\phi^{-1}(p)$ has property P, and vice versa.

As a dynamical system, DPath is just a direct union of full shifts, and Path is then also conjugate to such a union. As the name implies, we think of an element of Path as a two-way infinite path in $\mathbb{Z}$, and DPath is the sequence of moves that it follows.

An *ascending path* is a path $p \in$ DPath such that for some $m \in \mathbb{N}$, $\sum_{i=j}^{j+m} p(i) > 0$ for all j. A path $p \in$ Path is *bounded-to-one* if for some $m \in \mathbb{N}$, $|p^{-1}(n)| \leq m$ for all $n \in \mathbb{Z}$. A *finite-to-one* path is a path $p \in$ Path with $p^{-1}(n)$ finite for all $n \in \mathbb{Z}$. A *bounded path* is a path $p \in$ Path with $p(\mathbb{Z})$ finite. Uniform recurrence and minimality, for paths, mean the usual dynamical notions, and in DPath have the same interpretations as the usual ones for subshifts. Note that an ascending path can be recurrent, since recurrence means recurrence of the sequence of moves, not the sequence of cells the path visits.

In the multidimensional case, we define analogous systems $\mathtt{Path}^d$ and $\mathtt{DPath}^d$ of paths $p : \mathbb{Z} \to \mathbb{Z}^d$. For such paths, we define ascending and descending paths differently: to a path $p : \mathbb{Z} \to \mathbb{Z}^d$, we associate its *height path* $p' : \mathbb{Z} \to \mathbb{Z}$ by forgetting the first $d-1$ coordinates of $p(n)$ for all n. An *ascending path* $p \in \mathbb{Z} \to \mathbb{Z}^d$ is one where the height path $p' : \mathbb{Z} \to \mathbb{Z}$ is ascending, and *descending paths* are defined symmetrically.

We give a classification theorem for minimal path spaces.

Theorem 9. *Let $X \subset \mathtt{Path}_r$ be a minimal path space. Then exactly one of the following holds:*

- *every path in X is uniformly ascending,*
- *every path in X is uniformly descending,*
- *every path in X is bounded, or*
- *every path in X is unbounded, and some path in X is infinite-to-one.*

Proof. A word $w \sqsubset X$ is a *cut path* if whenever $p \in X$, $p_{[0,|w|-1]} = w$, we have $p_j \notin [0, r-1]$ for all $j \notin [0, |w|-1]$. It is easy to verify that if X has a cut path, then we have one of three cases: One possibility is that in every point, to the right of every cut path, the path stays above the origin, and stays below it on the left side, in which case every point is ascending. Symmetrically, if every path is above the origin on the left and below it on the right, every path is descending. In the remaining case, both can happen. But then some paths in X must stay between two cut paths, and we obtain a contradiction, since the cut path occurs syndetically.

Suppose then that X does not contain a cut path. Then it is easy to construct a path that is infinite-to-one: start with an arbitrary finite path w_1. Since it is not a cut path, we can continue it to a longer path w_2 that re-enters $[0, r-1]$. Continuing inductively, in the limit we obtain a path that enters one of the coordinates in $[0, r-1]$ infinitely many times. Thus, we are in the third or the fourth case. What remains is to show that if not every path in X is bounded, then every path in X is unbounded. This holds since not being bounded by m is an open condition and X is minimal. □

Path spaces of the last type are not encountered in the proofs of our main results, but it is perhaps the most interesting one, and it splits into further classes. The present classification is enough for our purposes, but for the interested reader, we give some examples of the different possible types of class four minimal path spaces in Section 7.

We need the following generalization in the proof of Proposition 8. The proof is exactly the same.[7]

Theorem 10. *Let $R \subset [-r, r]$ be a finite set of real numbers and let $X \subset R^{\mathbb{Z}}$ be a minimal subshift. Define*

$$Y = \{f : \mathbb{Z} \to \mathbb{R} \mid f(0) = 0 \wedge \exists x \in X : \forall i : f(i+1) = f(i) + x_n\}.$$

Then with the obvious definitions, exactly one of the following holds:

- *every path in Y is uniformly ascending,*
- *every path in Y is uniformly descending,*
- *every path in Y is bounded, or*
- *every path in Y is unbounded, and some path $p \in Y$ satisfies $p_i \in [0, r]$ for infinitely many i.*

4.2 Path covers

An *r-road (configuration)* is one whose support, up to translation, is the range of a two-way infinite ascending r-path. If the path can be chosen to be uniformly τ-recurrent, then we call the configuration an *r-highway.* The essential width of such a configuration is 1. An *r-preroad (resp. prehighway) configuration* is one whose nonzero support contains a translated image of an infinite ascending (resp. ascending and recurrent) r-path, but may contain other nonzero symbols too. The essential width of such a configuration is at least 1, but need not be finite. Note that r-prehighways do not, in the general case, have anything to do with paths – the all-1 point is an r-prehighway for every r.

The *(acyclic) r-path cover* of a subshift $X \subset \Sigma^{\mathbb{Z}^d}$ is the subshift $\mathrm{PC}_r(X) \subset X \times (\{\#\} \cup [-r, r]^d)^{\mathbb{Z}^d}$ of configurations (x, y) such that $y_i \neq \# \implies x_i \neq 0$ for all i, such that in the directed graph $G(Y)$ with nodes $\{i \mid y_i \neq \#\}$ and edges $\{(i, i+v) \mid y_i = v\}$, every node has in- and out-degree at most one, and there are no cycles. In $G(Y)$, every node is part of a unique maximal (infinite, one-way infinite or two-way infinite) path. If $(x, y) \in PC_r(X)$ is such that $0 \in G(Y)$, the path through 0 is two-way infinite and all nodes of $G(Y)$ are on this path, then (x, y) is a *marked path configuration*, and we write $\mathrm{MPC}_r(X)$ for the set of configurations, which is a compact (possibly empty) space.

Lemma 6. *Let X be a subshift. If r is such that for every n, X has a configuration whose support contains an r-component of size at least n, then* $\mathrm{MPC}_r(X)$ *is nonempty.*

Proof. Suppose that X has, for all n, some configuration x whose support has an r-component of size at least n. Then there are arbitrarily long distances between cells in the arbitrary large components, so consider an r-path of minimal

[7]Note that possible finite sequences of moves are still a discrete set, even though paths themselves may explore a dense set of positions.

length between two cells at distance at least $2n+1$ along the support of x, and construct a pair $(x, y) \in \mathrm{PC}_r(X)$ where a path between them is drawn on the y-component. Translating such pairs so that the middle of the path is at the origin, we obtain configurations in $\mathrm{PC}_r(X)$ where the maximal path through 0 is of length at least n in both directions. Any limit point of such pairs as $n \to \infty$ gives a point of $\mathrm{MPC}_r(X)$. □

The importance of the space $\mathrm{MPC}_r(X)$ is that it inherits the dynamics of $\mathtt{Path}^d$ in an obvious way, the action of $n \in \mathbb{Z}$ following the path written on the right component for n steps. More precisely, if $(x, y) \in \mathrm{MPC}_r(X)$ and $y_0 = v$, define $\tau(x, y) = \sigma^v(x, y)$, so that $(\mathrm{MPC}_r(X), \tau)$ becomes a $\mathbb{Z}$-system.

Lemma 7. *Let X be a rowsparse subshift. If $\mathrm{MPC}_r(X)$ is nonempty, then X contains a preroad.*

Proof. Take any minimal subsystem of $\mathrm{MPC}_r(X)$ in the $\mathbb{Z}$-action τ and a point (x, y) in it. Then in the projection to the right component, 0 is contained in a uniformly recurrent path p. It follows from Theorem 9 that p is either ascending, descending or not finite-to-one. The last case is impossible because X is rowsparse. If the path is descending, its reversal is ascending, so in both of the other two cases we obtain that the support of x contains the range of an infinite ascending path. □

We also need the following simple observations.

Lemma 8. *Let x be a preroad and $p : \mathbb{Z} \to \mathbb{Z}^2$ an ascending path whose range is contained in the support of x. If the distance from the range of p to every cell in the support of x is bounded, then x is a road.*

Proof. Suppose that m is such that $x_v \neq 0 \implies \exists n : |p(n) - v| \leq m$. Then let $a = |[-m, m]^2|$ and construct a q as follows: Let $n \in \mathbb{Z}$ and enumerate all cells in the support of x at distance at most m from $p(n)$ as $v_0, v_1, \ldots, v_{\ell-1}$, where $\ell \leq a$. Define $q(na + j) = v_j$ for all $j \in [0, \ell - 1]$ and $q(na + j) = v_1$ for $j \in [\ell, a - 1]$. Clearly the support of x is precisely the range of q, and the path q is clearly ascending since p is (though possibly with different ascension speed). □

Lemma 9. *Let X be almost minimal and $x \in X$ be a road. Then x is a highway.*

Proof. Let $p : \mathbb{Z} \to \mathbb{Z}^2$ be an ascending path whose range is contained in the support of x. Shift x so that p goes through the origin, and is thus an element of $\mathtt{Path}^2$. Take a path q in the orbit closure of p in $\mathtt{Path}^2$, and let $y \in X$ be the corresponding configuration. Then y contains the same patterns as x. It follows that x is a highway: every finite pattern in x is traced by a subpath of q (up to translation), and thus the support of x is the range of a path in the orbit closure of q. □

Note that the path p need not be uniformly recurrent – it is always possible to parametrize the range of a uniformly recurrent path in a non-recurrent way. The lemma states that there is always automatically also a uniformly recurrent parametrization.

5 The main theorem

In this section, we prove our main theorem, the characterization of almost minimal rowsparse subshifts. Before this, as a warm-up, we give a characterization of one-dimensional almost minimal subshifts.

Theorem 11. *Let a subshift $X \subset \Sigma^{\mathbb{Z}}$ be almost minimal and nontrivial. Then exactly one of the following holds:*

- *X is the orbit closure of a finite point, or*
- *X is the orbit closure of a blob fractal.*

Proof. Fix a nonzero point $x \in X$ with nonzero support $S \subset \mathbb{Z}$. If the language of x does not contain arbitrarily long words of the form 0^n, then the orbit closure does not contain $0^{\mathbb{Z}}$, contradicting almost minimality. It follows that for every $r \in \mathbb{N}$, there exists $n \in \mathbb{N}$ such that for all $i \in \mathbb{Z}$ the r-connected component of $i \in \mathbb{Z}$ in S is of size at most n. Collecting the subwords of x corresponding to these connected components, we obtain a set of r-blobs, which must be finite, since the blobs are words of length at most $n + 2r$.

Now, we start building a blob fractal structure for x. Let $r_1 \in \mathbb{N}$ be arbitrary, and let B_1 be the finite collection of blobs obtained across all $i \in S$ as in the previous paragraph. If x is not a finite point, then $|i - j|$ can be arbitrarily large for $i, j \in S$. From this, it follows that for large enough r, there is an r-blob b whose nonzero support is strictly larger than that of any of the blobs in B_1. By picking r large enough, this is true for all r-blobs, because b appears at a bounded distance from every nonzero symbol by almost minimality. By increasing r yet more, every blob in B_1 appears as a subpattern of every r-blob. Pick $r_2 \gg r_1$ with these properties, and define B_2 as the set of r_2-blobs.

Continuing inductively, we obtain that x is a blob fractal. □

Lemma 10. *Let $X \subset \Sigma^{\mathbb{Z}^2}$ be the orbit closure of a finite configuration, a highway or a blob fractal. Then X is almost minimal.*

Proof. The case of finite configurations is trivial.

For highways x, let P be a nonzero pattern in x with a connected domain. Pick an arbitrary nonzero cell v in P and follow the path p giving the support of x forward and backward from the cell v for t steps, to obtain some finite subpath q of p of length $2t + 1$ in the support of x. Because the support of x is the range of p, there is a function f of the size of the domain of P such that if $t > f(|D(P)|)$ then p visits all cells of P and can never again visit the domain $D(P)$. Since p is uniformly recurrent, the subpath q is traced regularly, so in particular P occurs at a bounded distance from every cell in the support of x. Since p uniformly recurrent, the same holds for all nonzero points of X

For blob fractals, let B_i be the sets of r_i-blobs as in the definition. If $x \in X$ and P is a nonzero pattern in x, then let i be such that r_i is larger than the diameter of the domain of P. Then P is contained in a single r_i-blob $B \in B_i$ in x. Now, let y be any point in X and w any nonzero cell. Since y is a limit of blobs, w must be contained in some B_{i+1}-blob. This blob contains B, and thus the pattern P. □

Finite configurations and path configurations of course always have essential width one. Blob fractals can have infinite essential width, and their essential width can be any finite number even in the rowsparse case, see Section 7.

The following lemma is easy to prove.

Lemma 11. *Let X be the orbit closure of a finite configuration (highway or blob fractal, respectively). Then every point in X is a finite configuration (highway or blob fractal, respectively).*

Theorem 12. *Let X be almost minimal and rowsparse. Then exactly one of the following holds:*

- *X is the orbit closure of a finite configuration,*
- *X is the orbit closure of a highway, or*
- *X is the orbit closure of a blob fractal.*

Proof. By the previous lemma, $X \subset \Sigma^{\mathbb{Z}^2}$ cannot be the orbit closure of two of these types of configurations. Now, by the assumption of almost minimality, we only have to find a point of one of these types in X. We begin with exactly the same argumentation as in Theorem 11. Take a point $x \in X$ and start building a blob fractal structure for it by increasing r, collecting the r-blobs that occur in x, and iterating.

If we can continue this for an infinite number of steps, and at each step we obtain finitely many distinct blobs that each contain at least two blobs of the previous size, we get in the limit that X is the orbit-closure of a blob fractal (the fact that all B_{i+1}-blobs contain all B_i-blobs is automatic if we increase r fast enough, again due to almost minimality). This process can stop for two reasons. First, it can stop because B_i is a singleton, and all $r_{i+1} > r_i$ only give padded versions of the blob in B_i. In this case, x is a finite configuration. In the other case, we have infinitely many r-blobs in x, for some r. In particular, configurations of X contain arbitrarily large r-components for some r. By Lemma 6, X then contains a preroad configuration. This is the interesting case, and in the rest of the proof we show that then X is the orbit closure of a highway.

Define the set $Y_m \subset (\Sigma^{\mathbb{Z}^2})^m$ as the set of m-tuples of configurations such that for any n, the $n \times n$ central patterns of those configurations can be found in the same n rows of a configuration of X, separated by at least distance n (but allowing said configuration to contain more nonzero symbols on those rows, and anything outside these m $(n \times n)$-rectangles). More precisely,

$$(x_1, \dots, x_m) \in Y_m \iff \forall n : \exists x \in X : \exists n_1, \dots, n_m \in \mathbb{Z} : \forall i \neq j : \\ |n_i - n_j| \geq 2n \wedge \sigma^{(n_i,0)}(x)_{[0,n-1]^2} = (x_i)_{[0,n-1]^2}$$

Note that every Y_m is closed under simultaneous translation in all coordinates $\sigma_v(x_1, \dots, x_m) = (\sigma_v(x_1), \dots, \sigma_v(x_m))$. It is also closed under horizontal translation of any individual component: $(x_1, \dots, x_i, \dots x_m) \in Y_m \implies (x_1, \dots, \sigma_{(j,0)}(x_i), \dots x_m) \in Y_m$.

What we have shown is that in $Y_1 = X$, we find an r-preroad configuration for some r. Let m be maximal such that in Y_m we find an m-tuple $(x_1, \dots, x_m)$ of r-preroad configurations.[8] Naturally, there is an upper bound on m as a

[8] The r must be the same, but the ascension speeds of the paths can a priori be different.

function of r and k, since all paths go through one of the rows $0, 1, \ldots, r-1$. Now, for each $i \in [1, m]$ fix an ascending path $p_i : \mathbb{Z} \to \mathbb{Z}^2$ whose image is contained in the support of x_i. If every nonzero symbol in x_i is at most a bounded distance away from the range of p, then x_i is a path configuration by Lemma 8, and we are done.

Suppose then that x^i is not a path configuration for some i, and let $x = x^i$ and $p = p_i$. We claim that m is not maximal. Namely, let $\ell \in \mathbb{N}$, take a finite subpath q of length ℓ in p, and let Q be the corresponding pattern. Take h such that a translate of Q appears in the h-ball around every cell in the support of every configuration in X (by almost minimality). Now, take a cell v in the support of x that is at distance at least $\ell + r\ell + h$ from the central path p. Then a translate of Q occurs in x in the h-ball around v, and thus every coordinate of the translate of Q is at distance at least ℓ from the path p.

Now, we are in the following situation: The pattern Q appears in x in some position (a_{m+1}, b), and on the same row in x, there is a nonzero coordinate $(a_i, b) + r_i$ that lies on the central path, where $r_i \in [-r, r]^2$ (because the path is ascending, and thus visits a syndetic set of rows), and $|a_i - a_{m+1}| \geq \ell$. Similarly, the paths p_j in the other points x_j visit some coordinates $(a_j, b) + r_j$ where $r_j \in [-r, r]^2$.

By translating $(x_1, \ldots, x_m)$ by $\sigma_{(a_i,b)+r_i}$, and taking ℓ large enough, for any $\ell' \in \mathbb{N}$ we may assume that a translate of Q occurs in the support of x at distance at least ℓ' from the origin, visiting some cell $v_{\ell'}$ on one of the rows $0, 1, \ldots, r-1$ and extending ℓ' steps both below and above row zero, and the central path p goes through the origin of x. By applying the horizontal translations $\sigma^{(a_j,0)}$ to the jth component for all $j \neq i$, we may assume the path p_j in x_j, for all $j \in [1, m]$, visits one of the coordinates in $[-r, r]^2$.

It follows that for every ℓ' we can find a tuple $(y_1, \ldots, y_m) \in Y_m$ where each y_j contains a path that visits one of the cells in $[-r, r]^2$, and in y_i, a finite ascending path extends ℓ' rows upward and downward from some cell $v_{\ell'}$ arbitrarily far from the origin, but at vertical distance at most r from row zero. As $\ell' \to \infty$, we then find in the limit a tuple $(z_1, \ldots, z_{m+1}) \in Y_{m+1}$ where all of the z_i are preroads. This is a contradiction, and thus x must have been a road. By Lemma 9, x is a highway. □

6 Corollaries

In this section, we give a list of corollaries of Theorem 12. The most general extraction theorem obtained from Theorem 12 and Lemma 1 is the following, which we mentioned in the introduction.

Theorem 13. *Let X be any two-dimensional subshift whose projective subdynamics Y satisfies satisfies one of the following:*

- *Y is sparse,*
- *every point in Y is eventually zero to the right, or*
- *Y is countable and the only periodic point in Y is $0^{\mathbb{Z}^2}$.*

Then X contains a finite point, a highway, or a rowsparse blob fractal.

Proof. In any of the three cases, Lemma 1 applies and we find an almost minimal rowsparse subsystem. The result follows from Theorem 2. □

6.1 SFTs, sofics and zero-gluing

Proposition 3. *Let X be a nontrivial rowsparse subshift such that for some r, in all configurations of X all r-components are infinite. Then X contains a highway.*

Proof. If X is rowsparse, then it contains an almost minimal rowsparse nontrivial subsystem by Lemma 3. Let Y be an almost minimal subshift of X. Then Y is the orbit closure of a finite configuration, a highway or a blob fractal by Theorem 2. Only highways have an infinite r-component for some r. □

Proposition 4. *No rowsparse nontrivial subshift X allows zero-gluing.*

Proof. As above, in all three cases of the classification theorem, applied to an almost minimal subsystem Y of X, we can use zero-gluing to show that the projective subdynamics is uncountable: For finite and highway configurations, simply glue them to their horizontal translates. For blob fractals, take an r_i-blob for r_i larger than the gluing radius, and the blob must extend to a finite point because we can legally glue it to the all-zero point. □

Theorem 14. *No nontrivial two-dimensional SFT is rowsparse.*

Proof. An SFT is clearly zero-gluing. □

Proposition 5. *Let X be a two-dimensional sofic shift with SFT cover Y such that the only preimage of $0^{\mathbb{Z}^2}$ is $0^{\mathbb{Z}^2}$. Then X does not have sparse projective subdynamics.*

Proof. Since the only preimage of a large ball of zeroes in the covering map is necessarily a large ball of zeroes, such a sofic shift is zero-gluing. □

General sofic shifts can have sparse projective subdynamics. See Section 7.

A $\mathbb{Z}$-subshift has *universal period n* if there exists a bound M such that for every point $x \in X$ there is a finite set $F_x \subset \mathbb{Z}$ of coordinates with $|F_x| \leq M$ and a periodic point $y \in X$ with $\sigma^n(y) = y$ such that $x|_{\mathbb{Z}\setminus F} = y|_{\mathbb{Z}\setminus F}$. (see Definition 4.3 in [31])

The following is Theorem 6.4 in [32], and it is the non-implementability result for zero-entropy sofic shifts in their classification of sofic projective subdynamics of two-dimensional subshifts of finite type.

Theorem 15. *If a $\mathbb{Z}$-subshift Y has universal period and is not a finite union of periodic points, then it is not the $\mathbb{Z}$-projective subdynamics of any $\mathbb{Z}^2$-SFT X.*

Proof. Suppose Y is the projective subdynamics of a $\mathbb{Z}^2$-SFT X. Consider the *blocking* of the subshift X, defined as the subshift with space X and $\mathbb{Z}^2$-dynamics $\gamma_{(a,b)}(x) = \sigma_{(na,b)}(x)$. Picking n to be the universal period of Y, this turns Y into a finite union of sparse subshifts. Thus, we may assume that X is an SFT whose projective subdynamics is a finite union of sparse subshifts with possibly different zero-symbols. Let Z be the factor of (X, γ) obtained by mapping all the distinct zero symbols on every row of $x \in X$ to the same unary point $0^{\mathbb{Z}}$, so that no non-unary point is mapped to $0^{\mathbb{Z}}$. Note that this is possible because there is a uniform bound on the number of coordinates where a point $y \in Y$ can differ from a unary point. Now, Z is clearly rowsparse, so by Lemma 3, it contains an almost minimal subsystem W.

By Theorem 12, W contains either a finite configuration, a highway configuration or a blob fractal. If $z \in Z$ is finite or a highway, then clearly $\sum_{i \in \mathbb{Z}} \sigma_{(ik,0)}(z) \in X$ for some k, since we can simply glue the preimages of z together as they all have the same zero symbol on the left and right side of the support of z (either the finite blob or the infinite ascending path). If z is a blob fractal, then it is easy to see by analyzing the covering point $x \in X$ that Z also contains a finite configuration, which is a contradiction. □

More generally, we obtain that there are singly periodic points in all SFTs whose projective subdynamics are countable sofic shifts. We say a $\mathbb{Z}$-subshift Y is *bounded* if its language is bounded, where a bounded language is any sublanguage of $w_1^* w_2^* \cdots w_\ell^*$ for words w_i. In [35] it is shown that all countable sofic shifts are bounded. A *singly periodic* configuration is a configuration with a nontrivial period whose orbit in the shift action is infinite. The following proposition should be constrasted with [3, Theorem 3.11], where it is shown that all countable two-dimensional SFTs contain a singly periodic point.

Proposition 6. *If X is an SFT whose projective subdynamics is a bounded infinite subshift, then it contains a singly periodic point.*

Proof. Let n be the least common multiple of the lengths of the words w_i in the definition of boundedness. As in the previous proof, consider the subshift Z obtained from X by blocking and factoring so that the projective subdynamics becomes sparse. By the assumption on the projective subdynamics of X, the projective subdynamics of Z is not $\{0^{\mathbb{Z}}\}$. As above, take an almost minimal subsystem of Z.

If Z contains a highway z, then from a covering configuration in X of z, we obtain a singly periodic point in X by the pigeonhole principle, since the preimage of a finite point in the projective subdynamics of Z is clearly an eventually periodic point in the projective subdynamics of X, and there are only finitely many points with given left and right periods and central pattern length.

It is easy to see that from a blob fractal, we obtain a finite configuration (again by the pigeonhole principle). For a finite configuration we can find a covering point of the following form: for some n, m, $x_{(a,b)} = x_{(a,b+n)} = x_{a+n,b)}$ for all $|b| \geq m$, and $x_{\mathbb{Z} \times [-m,m]}$ is eventually periodic as a one-dimensional configuration. From such a configuration, we obtain either a finite configuration or a singly periodic point by compactness. Of course, an SFT with bounded projective subdynamics cannot contain a finite point. □

6.2 Subshifts with deterministic or expansive directions

SFTness can often be replaced by determinism. The main thing we need is that if d is a deterministic direction, then a half-plane in direction d[9] containing only

[9]Recall that by this we mean that the boundary is perpendicular to d – we think of the half-plane as moving in direction d and eating up the configuration.

0s, or that is otherwise doubly periodic, can only be continued periodically. If X contains a configuration contradicting this, we say *something appears from nothing.*

Proposition 7. *If X is a subshift with a vertical deterministic direction whose projective subdynamics is a bounded infinite subshift, then it contains a singly periodic point.*

Proof. Suppose X is deterministic upward. In the argument of the previous proposition, one can then replace the pigeonhole principle by determinism. Namely, let $x \in X$ be a covering configuration of a path configuration in Z. Determinism means that a lower half-plane containing half of the path will be uniquely filled to a configuration of X. Since the determinism is given by a local rule (by compactness), and every row of X is eventually periodic, it is easy to see that the direction of the path is rational.

Finite configurations are impossible, because something cannot appear from nothing. Similar reasoning applies to blob fractals: in the orbit closure of a blob fractal, for any d one can find a half-plane in direction d of all zeroes whose boundary contains a nonzero symbol. □

When the deterministic direction is not vertical, X need not contain a singly periodic point. See Example 7 in Section 7. However, in the rowsparse case, any direction of determinism implies the existence of a singly periodic point. A well-studied concept in multidimensional dynamics are the directions of expansivity [9], a concept similar to determinism. We say $d \in \mathbb{R}^2 \setminus \{(0)\}$ is an *expansive direction*[10] if, defining $L_{d,r} = \mathbb{R}d + B_r \subset \mathbb{R}^2$, we have

$$\exists \epsilon > 0 : \exists r > 0 : (\forall v \in L_{d,r} : d(\sigma_v(x), \sigma_v(y)) < \epsilon) \implies x = y.$$

Determinism and expansivity are connected by the following well-known lemma.

Lemma 12. *Let X be a two-dimensional subshift. Then d is an expansive direction for X if and only if both of the directions orthogonal to d are deterministic.*

Proof. Let D be the set of deterministic directions, and E the set of expansive ones (which is of course symmetric). We will show $d \notin D \implies (d^T \notin E \wedge -d^T \notin E)$ and $d \notin E \implies (d^T \notin D \vee -d^T \notin D)$, from which the lemma follows easily.

If d is a direction of nondeterminism, then we have two points x, y that agree on a d-directional half-plane, but not everywhere, and then thick strips around the boundary of this half-plane imply the directions orthogonal to d are non-expansive. Conversely, if d is a direction of nonexpansivity, then we have arbitrarily thick strips $L_{d,r}$ that extend in two ways in at least one of the directions. Translating this difference to the origin, in the limit as $r \to \infty$ we obtain two points that agree on a half-plane in one of the two directions orthogonal to d, but not everywhere. □

In [9], $\mathbb{Z}^2$-subshifts with almost any[11] set of expansive directions is constructed, the only exception being a single irrational expansive direction, which was left open. Since a rational line in direction d will have d as the unique

[10] Of course, a more natural way to define directional expansivity is to talk about expansive subspaces of $\mathbb{R}^2$, but for the purpose of our discussion, we find directions notationally easier.

[11] Up to the obvious restrictions.

nonexpansive direction, the obvious way to try to solve this is a single irrational line. The next result, Proposition 8, shows that such an idea cannot work.[12] See [21] for a correct implementation of a unique nonexpansive direction.

We need the following well-known fact, which can be shown analogously as the closedness of expansive directions in [9].

Lemma 13. *Let $X \subset \Sigma^{\mathbb{Z}^2}$ be a subshift. Then the set of nondeterministic directions is closed in the unit circle.*

Lemma 14. *If X is nontrivial, almost minimal and rowsparse, and has a direction of determinism, then it is the orbit closure of a singly periodic point, and thus its set of nondeterministic directions is $\{d, -d\}$ where $d \in \mathbb{Q}^2$.*

Proof. As above, we see that if there is a deterministic direction, then X cannot contain a finite configuration or a blob fractal, and thus must consist of r-paths for some fixed r. By the previous lemma, there is a rational direction d, which we may (without loss of generality) assume to lie in the top right quadrant and to have integer coordinates, $d \in \mathbb{N}^2$. By blocking X into blocks of size (d_1, d_2) and restricting to an almost minimal subshift, we can assume X has direction $(1, 1)$ of determinism, and by skewing, we may assume it is deterministic upward. There result follows from Proposition 7. □

Proposition 8. *If X is nontrivial and rowsparse and has a direction of determinism, then it contains a singly periodic point, and thus has a rational nondeterministic direction.*

6.3 Cellular automata

Surjective cellular automata are nothing but subshifts with a rational deterministic direction, and thus it is clear that we obtain corollaries for cellular automata.

The main technical tool used in [37] is the Starfleet Lemma for cellular automata on countable sofic shifts, which states that, starting from any configuration, a 'fleet of spaceships' appears infinitely often, in the same configuration. This result shows that all cellular automata on countable sofic shifts at least occasionally behave like counter machines. We do not state this result precisely, but we reproduce one of the main corollaries obtained from it in [37], namely the decidability of the nilpotency problem.

More precisely, let $X \subset \Sigma^{\mathbb{Z}}$ be a countable sofic shift. A cellular automaton $f : X \to X$ is *nilpotent* if for some $n \in \mathbb{N}$, $f^n(x) = 0^{\mathbb{Z}}$ for all $x \in X$. A *spaceship* for f is a non-periodic configuration $x = {}^{\infty}uvw^{\infty}$ such that $f^n(x) = \sigma^m(x)$ for some $m \in \mathbb{Z}, n \geq 1$. If $u, w \in 0^*$, then x is called a *glider*.

Theorem 16. *Let X be a one-dimensional subshift with bounded projective subdynamics and $f : X \to X$ a cellular automaton. Then either there exists k such that $f^k(X)$ is spatially periodic, or there exists a spaceship for f. In particular, nilpotency is decidable for cellular automata on countable sofic shifts.*

Proof. If $f^k(X)$ is not finite for any k, then the projective subdynamics of the spacetime subshift of f is infinite. By the Proposition 8, X contains a singly periodic point. Such a configuration is precisely a spaceship for f. □

[12]The fact that this idea does not work is folklore, but we do not have a reference.

Proposition 9. *Let $X \subset \Sigma^{\mathbb{Z}}$ be any one-dimensional subshift, and let $f : X \to X$ be a CA such that either f is asymptotically nilpotent or the limit set of f contains only configurations that are eventually zero to the right. Then f is either nilpotent or has a glider.*

Proof. If f is asymptotically nilpotent, then consider its spacetime subshift Y where f is run to the right, so that rows are eventually zero to the right. By Lemma 1 we can extract an almost minimal rowsparse subshift Z from Y. Then Z is a rowsparse almost minimal subshift with a deterministic direction, and thus contains a singly periodic point, which corresponds to a glider. The proof in the case of a limit set with only one-sided configurations is the same, though the direction of determinism is different. □

We can also obtain the following proposition.

Proposition 10. *Let $X \subset \Sigma^{\mathbb{Z}}$ be any subshift, and let $f : X \to X$ be a CA such that the closure of the asymptotic set of f contains only configurations that are eventually zero to the right. Then f is either nilpotent or has a glider.*

Proof. If the asymptotic set of f contains only the all-zero point, then f is asymptotically nilpotent, and the claim follows from the previous theorem. Otherwise, the closure Z of the asymptotic set of f is a subshift on which f is surjective and that contains at least one nonzero point. Then the spacetime subshift of f where the projective subdynamics are restricted to be in Z is a two-dimensional nonzero subshift where every row is eventually zero to the right, and the claim follows as in the previous theorem. □

In the case where X is an SFT, we of course cannot have any such gliders in the nilpotent case, and we obtain the following theorem (which is also proved in [18]).

Theorem 17. *Let X be a one-dimensional SFT and $f : X \to X$ a cellular automaton. Then f is nilpotent if one (equivalently all) of the following holds:*

- *f has a sparse limit set,*
- *f is asymptotically nilpotent,*
- *the (one- or two-sided) trace subshift of f is sparse, or*
- *the asymptotic set of f is sparse.*

6.4 Irrational sparseness

Since Theorem 2 applies to any subshift, it in particular applies to binary subshifts that are coding some property of points in another system. We show one application of this observation, namely that in Theorem 2 it is enough to assume rowsparsity in any (possibly irrational) direction:

Proposition 11. *Let X be a $\mathbb{Z}^2$-subshift and let L be a line with irrational slope. Suppose that for some $r > 0$, the strip $v + L + B_r$ in x contains finitely many nonzero symbols for every $v \in \mathbb{R}^2$ and every $x \in X$. Then X satisfies the conclusion of Theorem 2.*

Proof. Let d be a unit vector giving the slope of L, and let d^T be the unit vector orthogonal to it. Let $R \subset \mathbb{R}^2$ be the square having one corner at $(0,0)$ and with sides d and d^T. Now, to a point $x \in X$ whose nonzero support is $S \subset \mathbb{Z}^2$, associate the binary point $\psi(x) \in \{0,1\}^{\mathbb{Z}^2}$ defined by

$$\psi(x)_v = 0 \text{ iff } S \cap (v_1 d + v_2 d^T + R) \neq \emptyset.$$

Then $\psi(x)$ is a coding of a rotated version of x. Let Y be the subshift generated by the points $\psi(x)$. Now, it follows from the assumption that the subshift $\psi(X)$ is rowsparse, and contains a point of one of the three types in Theorem 2. It is easy to check that finite points must come from finite points through a rotation, and paths must come from paths, which we can make uniformly recurrent by again passing to an almost minimal subshift. Blob fractals need not a priori come from blob fractals (since our definition of blob fractal implies almost minimality). However, it is again clear that any almost minimal subsystem will be supported by a blob fractal. □

6.5 Topological full groups

The *topological full group* of a subshift X is the group of homeomorphisms $h : X \to X$ such that for some *cocycle*[13] $c : X \times G \to \mathbb{Z}$, $c(h,x) = n \implies h \cdot x = \sigma^n(x)$ and c is continuous in its left argument [16].

Lemma 15. *Let g be an element of the topological full group G of a subshift $X \subset \Sigma^{\mathbb{Z}}$ with cocycle c. If g has infinite order, then there is a point $x \in X$ and $k \geq 1$ such that $n \mapsto c(g^{kn}, x)$ is increasing or decreasing.*

Proof. Take the $\mathbb{Z}^2$-subshift Y where rows are unary and columns are configurations of X. Add another layer on top of Y, where in cell $v \in \mathbb{Z}^2$ of $x \in Y$, we write the vector $(1, c(g, \sigma^v(x)))$ (note that by continuity in the left argument, there are finitely many such vectors), and call the resulting subshift Z. Now, the function $n \mapsto c(g^{kn}, x)$ tells us the vertical movement of a path in a configuration of Z. If the path visits the same row twice, then it is periodic. Because g is of infinite order, we must find paths that do not revisit a cell in arbitrarily many steps.

Thus, forbidding revisits, we still obtain infinite paths. Now, consider a factor of this subshift where the Y-component is removed. This subshift is rowsparse, and thus contains an almost minimal subshift, which must consist of a single uniformly ascending or descending path. By looking at the corresponding Y-configuration, the claim is proved. □

Theorem 18. *If X is a full shift, then the torsion problem of the topological full group of X is decidable.*

Proof. Clearly the torsion problem is semidecidable. On the other hand, if g has infinite order, then let c be as above. By the previous lemma there is a point $x \in X$ and $k \geq 1$ such that $n \mapsto c(x, g^{kn})$ is increasing or decreasing. By the pigeonhole principle, we can find such periodic x. By changing a single coordinate in the tail of x that g does not see, we find a spatially eventually periodic but non-periodic configuration that g shifts periodically in one direction. This is clearly semidecidable. □

[13] Such c is a cocycle in the sense of cohomology, that is, $c(g \circ h, x) = c(h,x) + c(g, h(x))$.

6.6 Patterns in primes

Unlike our main theorem, the one-dimensional version Theorem 11 applies very generally (though of course also its conclusion is rather trivial). We show one application of it – finding 'infinite patterns' in primes. Let $N \subset \mathbb{N}$ be arbitrary, and $x_N \in \{0,1\}^{\mathbb{Z}}$ its characteristic function. We call

$$X_N = \bigcap_m \overline{\{\sigma_n(x) \mid n \geq m\}}$$

the *N-subshift*; it is always a subshift, and it is the set of two-way infinite configurations whose subwords appear arbitrarily late to the right in the point x_N. Setting N to be set $\mathcal{P}$ of prime numbers we obtain what we call the *prime subshift.*

Lemma 16. *Let p_i be the ith prime number. Then*

$$\liminf_{n\to\infty} p_{n+1} - p_n < \infty \tag{1}$$

if and only if the prime subshift is not equal to the sunny-side-up subshift $\mathcal{O}(\ldots 00100\ldots)$.

Proof. Since there are infinitely many primes, $X_{\mathcal{P}}$ is not the all-zero subshift by compactness. If (1) holds, then there is some $k \in \mathbb{N}$ such that $p_{n_i+1} - p_{n_i} = k$ for some sequence $n_i \to \infty$. Let y be a limit point of the sequence $\sigma_{n_i}(x)$ as $i \to \infty$. Then $y_{[0,k]} = 10^{k-1}1$ and $y \in X_{\mathcal{P}}$.

Suppose then that $X_{\mathcal{P}}$ is not the sunny-side-up. Since $X_{\mathcal{P}}$ is not the all-zero subshift, some point in it contains a word with two ones, and thus a word of the form $10^{k-1}1$. But then $\liminf_{n\to\infty} p_{n+1} - p_n \leq k$. □

The fact that indeed $\liminf_{n\to\infty} p_{n+1} - p_n < \infty$ was shown in 2014 in a revolutionary paper by Zhang [40]. Since then, much more has been learned about the words appearing in $X_{\mathcal{P}}$, and in particular [28] shows that $1w1 \sqsubset X_{\mathcal{P}}$ holds for some $|w| < 600$. For a discussion of the subshift point of view to the study of divisibility, see [6] and references thereof.

The following is a corollary of Theorem A in [6], where it is shown that every *$\mathcal{B}$-free subshift* is essentially minimal. Our definition of the prime subshift is slightly different from theirs, so we give the proof.

For a set $N \subset \mathbb{N}$, write $\pi_N(n) = |N \cap [0,n]|$, and define $\pi = \pi_{\mathcal{P}}$.

Lemma 17. *The prime subshift is essentially minimal, and its only minimal subshift is $\{0^{\mathbb{Z}}\}$.*

Proof. Let n be arbitrary. Let $I = \{0,1,\ldots,n\}$ and let $\phi : I \to P$ be an injection from I into the prime numbers. Let $N = \prod_{i\in I} \phi(i)$. Let $k = \sum_{i\in I} a_i N/\phi(i)$ where the a_i are chosen so that $k \equiv -i \bmod \phi(i)$ for all $i \in I$. Then $k+i \equiv i-i \equiv 0 \bmod \phi(i)$ for all $i \in I$. It follows that $x_{\mathcal{P}}$ contains the word 0^n for arbitrarily large n. If $(x_{\mathcal{P}})_{[k,k+n-1]} = 0^n$, then $k+i$ is divisible by p_i for $i \in [0, n-1]$ for some primes p_i. It follows that 0^n occurs in every subword of $x_{\mathcal{P}}$ of length $\prod_{i=0}^{n-1} p_i + 2n$. □

We note that the prime number theorem $\pi(n) \sim n/\log n$ or any other density-related statement is *not* enough on its own, and it is rather the periodic structure of the zeros that is important – by perturbing the primes by inserting intervals $[\ell, \ell+n]$ very sparsely, but for arbitrarily large n, one obtains a set $N \subset \mathbb{N}$ with $\pi_N \sim \pi$ such that X_N is not essentially minimal. By alternating long intervals of zeroes and ones, it is also possible to build sets $N \subset \mathbb{N}$ such that $\pi_N \sim \pi$ and $X_N = \{0^{\mathbb{Z}}, 1^{\mathbb{Z}}\}$.

Proposition 12. *The prime subshift contains either a finite point or a blob fractal.*

Proof. By the previous lemma, $X_{\mathcal{P}}$ is essentially minimal with $\{0^{\mathbb{Z}}\}$ the only minimal subsystem. By Lemma 3, $X_{\mathcal{P}}$ contains an almost minimal subsystem, and the claim follows from Theorem 11. □

What is interesting about the proposition is mainly that we obtained it for essentially abstract reasons from only the Chinese remainder theorem and the infinitude of the primes, while proving that $X_{\mathcal{P}}$ contains a particular finite point or blob fractal is quite difficult from first principles. We do not know whether the prime subshift contains a blob fractal. If it did, $X_{\mathcal{P}}$ would be uncountable. We note that at least the conclusion of Theorem 1.2 of [28] allows the prime subshift to be countable, as it is easy to show that the subshift

$$X = \{x \in \{0,1\}^{\mathbb{Z}} \mid (|x| \geq n \wedge x_i = x_j = 1) \implies (i = j \vee |i-j| \geq n)\},$$

where $|x|$ is the number of nonzero symbols in x, satisfies its conclusion (and it still does if one further restricts X to contain only admissible patterns).

We remark that one *can* show that $X_{\mathcal{P}}$ contains the sunny-side-up with a more sophisticated but standard number-theoretic tool, namely Dirichlet's theorem on arithmetic progressions.

Proposition 13. *The prime subshift contains the sunny-side-up subshift.*

Proof. Let n be arbitrary. Let $I = \{-n, \dots, -1\} \cup \{1, \dots, n\}$ and let $\phi : I \to P$ be an injection from I to the prime numbers all larger than $2n$. Let $N = \prod_{i \in I} \phi(i)$. Let $k = \sum_{i \in I} a_i N/\phi(i)$ where the a_i are chosen so that $k \equiv i \bmod \phi(i)$ for all $i \in I$. Then k and N are coprime, so by Diriclet's theorem on arithmetic progressions, there is a prime number $p = k + \ell N$ for some $\ell \in \mathbb{N}$. Then $p - i \equiv i - i \equiv 0 \bmod \phi(i)$ for $i \in I$. It follows that the characteristic sequence of the primes contains the word $0^n 1 0^n$ at the prime p. Taking a suitable limit, we obtain that the sunny-side-up is contained in X_P. □

7 Examples

In this section, we give some examples of almost minimal and sparse subshifts, to illustrate to what extent the various assumptions are needed in Theorem 2, and in particular show what kind of behaviors can happen in sofic shifts. We construct these examples using existing constructions of sofic shifts as black-boxes. In particular we use the result of Mozes [29] that fixed-point subshifts of many substitutions (in particular rectangular deterministic ones) are sofic, and the result of various authors that projective subdynamics of sofics can be arbitrary computable subshifts [20, 13, 2]. We also give some examples of path spaces.

We note that many examples of subshifts (in particular sofic shifts and SFTs) with bounded projective subdynamics are constructed in [36, 38], and bounded subshifts can be turned into sparse ones by blocking and factoring (see the proof of Theorem 15). The proof of Lemma 15 shows that every element of a topological full group gives an example of a rowsparse subshift.

7.1 Subshifts

The following example is our first example of a rowsparse blob fractal, and thus shows that in Theorem 2, the case of blob fractals must indeed be included. Less trivially, it also shows that while essential width is always one in the cases of finite configurations and path configurations, the essential width of an almost minimal rowsparse subshift can in general be arbitrarily large and can match the sparsity constant.

Example 3: For every k, there exists an almost minimal two-dimensional sofic subshift X with k-sparse projective subdynamics such that every configuration has essential width k.

See Figure 3 for an illustration of a typical pattern in X.

Proof. We inductively construct k patterns $P_{i,1}, \dots, P_{i,k}$ of shape $[0, m_i - 1]^2$ such that the bottom row of each contains exactly one 1, which is at the bottom left corner: $(P_{i,j})_{0,0} = 1$. The patterns also satisfy that if $j \neq j'$ then the only row that is nonzero in both $P_{i,j}$ and $P_{i,j'}$ is row 0 and that no row of $P_{i,j}$ contains more than k nonzero symbols. Furthermore, each of the patterns $P_{i,j}$ has at least one row with k nonzero symbols that are pairwise separated by at least distance m_{i-1}.

Pick the matrices $P_{1,j}$ and m_1 arbitrarily so that the assumptions are satisfied (for example, use the matrix below with $P_{0,j} = 1$ for all j). Now, for $j \in \{1, \dots, k\}$, build $P_{i+1,j}$ from the patterns $P_{i,1}, \dots, P_{i,k}$ as follows. First, set $m_{i+1} = (k+1)m_i$, so that $[0, m_{i+1} - 1]^2$ partitions into a $(k+1) \times (k+1)$ grid of translates of the squares $[0, m_i - 1]^2$. If $i \geq 1$, we define

$$P_{i+1,j} = \begin{bmatrix} 0 & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 0 \\ 0 & P_{i,1} & P_{i,2} & \cdots & P_{i,k} \\ 0 & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 0 \\ P_{i,j} & 0 & 0 & \cdots & 0 \end{bmatrix}$$

as the block matrix where the slice $0 \; P_{i,1} \; P_{i,2} \; \cdots \; P_{i,k}$ appears on the jth row from the bottom. It is easy to see that the induction hypothesis is satisfied.

We define X as the limit of these patterns. Since every nonzero pattern occurring in the subshift always occurs in one of the patterns $P_{i+1,j}$, and each such pattern contains $P_{i,j'}$ for every j', we can conclude that the subshift is almost minimal and a row containing k ones with pairwise distances at least m_{i-1} occurs at a bounded distance from every nonzero symbol for every i. Since $m_{i-1} \to \infty$, the subshift has essential width k.

Finally, we observe that X is sofic by Mozes' theorem [29], since it is defined by a deterministic primitive substititution. □

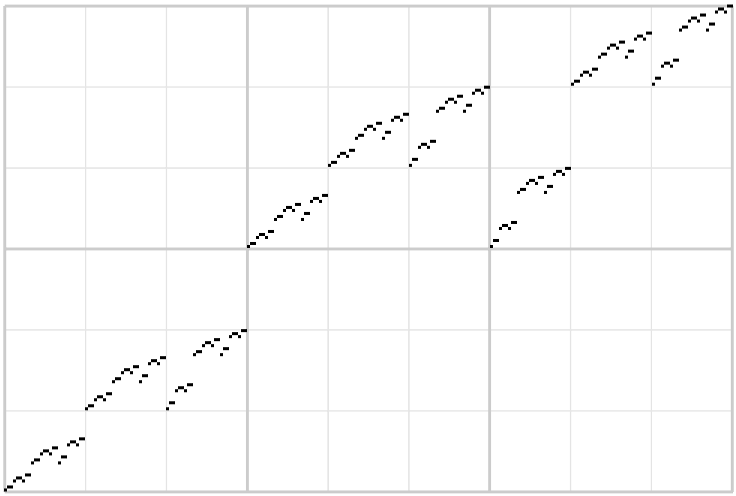

Figure 2: The pattern $P_{5,1}$ for $k = 2$ constructed in Example 3.

There are certainly sparse subshifts that are not almost minimal. Thus the assumption of almost minimality is necessary in Theorem 2.

We show that in Theorem 2, the assumption of rowsparsity is necessary.

Example 4: There is an almost minimal binary sofic subshift where the connected component of every 1 is infinite.

Proof. Use the substitution $\blacksquare \mapsto$  . See Figure 3. $\square$

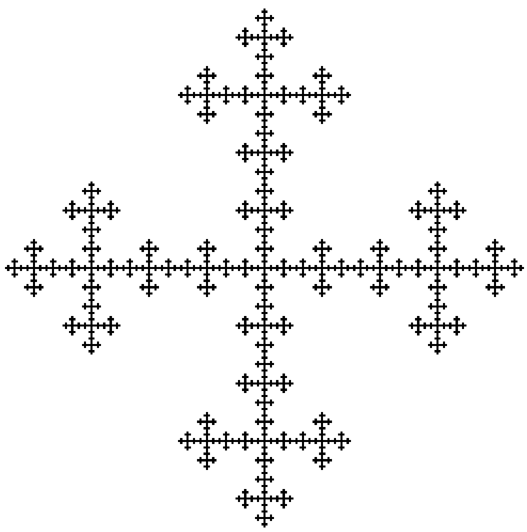

Figure 3: The almost minimal subshift of Example 4.

By varying the substitution, it is even possible to make nonzero symbols appear with positive density on every nonzero line (showing that almost minimal subshifts can support nontrivial shift-invariant probability measures).

Example 5: For $n \in \mathbb{N}$ define the substitution $\tau_n(0) = 0^{2^n}$ and $\tau_n(1) = 1^{2^n-1}0$. Consider $w_k = \tau_2(\cdots \tau_k(1) \cdots)$. This word is of length $\prod_{i=2}^k 2^i$ and the number of 1-symbols in it is $\prod_{i=2}^k (2^i - 1)$, so the density of nonzero symbols in it is

$$\prod_{i=2}^{k} (1 - 1/2^i) \geq 1 - \sum_{i=2}^{k} 1/2^i > 1/2.$$

Moreover, if $\ell \geq k$, every nonzero symbol occurring in w_ℓ is contained in a copy of w_k, and thus the ball of radius $|w_k|$ around it has nonzero symbols with upper Banach density at least 1/4. It follows that in the limit as $k \to \infty$ we obtain a subshift where every point containing a nonzero symbol contains nonzero symbols with upper Banach density 1/4. △

SFTs containing a path must contain a periodic path by the pigeonhole principle. We now show that sofic shifts can consist of arbitrarily complex ascending paths.

Proposition 14. *Let X be a Π^0_1-subshift where the nonzero support of every nonzero configuration is a uniformly ascending r-path. Then X is sofic.*

Proof. Let $\{0\} \cup \Sigma$ be the alphabet of X. Let r be such that the nonzero support of every nonzero configuration is the range of an r-path, and let m be the ascension constant of the path (so starting from v, any path stays above v after m steps). Let $n = rm$.

Let Y be a one-dimensional Π^0_1-subshift over the alphabet $D = \# \cup ([-n, n] \times [1, n] \times (\{0\} \cup \Sigma)^{[-n,n]^2})$. From such a subshift, we construct a subshift where the points of Y give directions for an ascending path, as follows. Let Z be the two-dimensional subshift over D where all rows are constant and columns are points of Y. On every non-# row i, if the number in the $[1, n]$-component is a, then call $i + a$ the *successor row* of row i. Require that the successor of every row is non-#, the rows between a row an its successor contain #, and that every row has a unique successor. Thus, the $[1, n]$-components trace a vertical ascending path on the configuration.

Now, add a binary layer where on each row not containing #, at most one cell contains 1, and rows containing # are all-0. In other words, we overlay the sunny-side-up on the non-#-rows. Call cells containing 1 on this layer *marked cells*. If on the Z-layer the $([-n, n] \times [1, n])$-component contains u at v, then $v + u$ is called the *successor* of the cell v. We require that v is marked if and only if its successor is marked.

At this point, we have a subshift with at most one single ascending n-path on the binary layer, which must travel according to the directions given on the Z-layer. Now, consider the factor map ϕ defined as follows: If v contains 1 on the binary layer, then write the nonzero support of the pattern in the $(\{0\} \cup \Sigma)^{[-n,n]^2}$-component around v. Write 0 in every remaining cell. We add to our subshift the final SFT constraint that this factor map is well-defined. We obtain a subshift W that has as a factor a $(\{0\} \cup \Sigma)$-subshift where the nonzero support of every configuration is the range of a $3n$-path, whose movement is guided by Y.

Now, observe that for every Y, the subshift W constructed above is a sofic shift. For this, simply observe that Z is sofic, that checking that every row contains at most one binary symbol is doable since the sunny-side-up is a sofic shift, and that all other constraints we added were local. Of course, then also $\phi(W)$ is sofic.

Finally, we need to show that some Π_1^0-subshift Y yields $\phi(W) = X$. For this, simply define Y as the Π_1^0-subshift where for every forbidden pattern of X we forbid every set of directions that would yield a forbidden pattern of X in the ϕ-image. Then $\phi(W) \subset X$ because any forbidden pattern would have been traced by a path, guided by a word w of Y and then w would have been forbidden in Y. On the other hand, for any configuration x of X, one can easily find an n-path through the configuration and construct the corresponding $(\{0\} \cup \Sigma)^{[-n,n]^2}$-patterns, to obtain a legal guiding sequence in Y. □

It is easy to construct minimal Π_1^0 path spaces without a well-defined linear direction, and thus we obtain the following corollary.

Example 6: There exists a sofic shift X whose configurations are two-way infinite uniformly ascending paths, but are not periodic and their support does not fit any strip of the form $L + B_r$ where L is a straight line. △

Similarly by using Sturmian subshifts with computable angles as the guiding sequences, we obtain sofic shifts where the nonzero support of every configuration does fit a straight strip $L + B_r$, but no such rational line. With a similar idea, we also obtain an example of a sofic shift with bounded projective subdynamics with only four nondeterministic directions, two of which are irrational.

Example 7: Let X be a Sturmian subshift (with any computable irrational angle). Let Y be the sofic shift where each row is constant, and columns are points of X. Construct $Z \subset Y \times \{0,1\}^{\mathbb{Z}^2}$ similarly as in Proposition 14: allow at most one 1 on each row, and if $(y, z) \in Z$ and $x_z = 1$, require $x_{z+(y_v,1)} = 1$ and that either $x_{v+(0,-1)} = 1$ or $x_{v+(-1,0)} = 1$.

Now, we claim that the nondeterministic directions are precisely the vertical ones and the ones parallel to the irrational slope. First, both of them are easily seen to be nondeterministic, since in the vertical directions it is impossible the know the continuation of the Y-component, and in a direction orthogonal to the path, it is impossible to know whether the continuation of the half-plane eventually hits the path.

In half-planes of any other slope, we see the color of every column, and thus the contents of the half-plane determines the Y-component. If the path visits the half-plane, we can uniquely fill its trajectory based on the Y-component, and every path that fits in a strip $L + B_r$ where L is a straight line visits every half-plane not parallel to L. △

If we remove the guiding rows Y, it is tempting to think that the only nondeterministic directions are the ones parallel to the irrational line, but this is not the case. Subshifts with a single irrational nonexpansive direction are not easy to construct even in theory [21], and making such sofic shifts and SFTs is another challenge still [41]. And of course, if the guiding rows are removed, Z becomes sparse, and we know it then has a singly periodic point by Proposition 8.

We do not believe it is essential that the paths be ascending in Proposition 14, as every row carries little enough information (the finitely many bounded offsets of finitely many paths) that it can be carried on each row. However, adapting the proof of the theorem above to this general case seems to require additional tricks, if it is possible at all, and perhaps one would need a radically different idea.

Conjecture 1. *If X is a uniformly rowsparse Π^0_1-subshift where the nonzero support of every configuration is the range of an r-path, then X is sofic.*

We also show that sparseness and computability alone are not enough to guarantee soficness, by a slight adaptation of the well-known mirror subshift.

Proposition 15. *There exists a 1-sparse non-sofic Π^0_1-subshift.*

Proof. We construct such a binary subshift. First, forbid every configuration where the set of rows containing a nonzero symbol intersects every residue class modulo 3, and every configuration where some row contains two ones. If two rows whose distance is not divisible by 3 contain 1 at coordinates (n_1, m_1) and (n_2, m_2), respectively, and $m_2 - m_1 \equiv 1 \bmod 3$, then require that $n_1 < n_2$ and that if $|n_1 - n_2| = m$ then there exist two infinite vertical (necessarily nonoverlapping) strips of width at most $m/4$ containing the support of the configuration, and that both of these strips contain a nonzero symbol on exactly every third row. Finally, require that for every pair of adjacent nonzero rows, the horizontal distance between the nonzero symbols on them is the same. This is clearly a 1-sparse Π^0_1 subshift, and it is easily seen to be non-sofic by the pigeonhole principle. □

With a slightly more elaborate construction, one can find even a 1-sparse non-sofic Π^0_1-subshift where every configuration is either finite or an infinite ascending path (though of course they are not all r-paths for any fixed r).

7.2 Paths

In Theorem 9, we showed that there are four kinds of minimal path spaces. It is easy to find examples of the first three cases. Path spaces X where all paths are bounded can be found by taking a minimal subshift Y with alphabet contained in $\mathbb{Z}$, and taking the paths of X to be the discrete derivatives (differences between consecutive cells) of points of Y.[14] Ascending path spaces can be constructed by summing points of minimal subshifts with alphabet contained in $\mathbb{Z}$ where the sum of every long enough word is positive, and descending path spaces can be found by inverting ascending paths.

The fourth case, that every path is unbounded-to-one and some path is infinite-to-one was not really needed in our characterizations, as we have concentrated on rowsparse spaces. However, this case is quite interesting, and we show by examples that it splits into three subcases.[15] In the rest of this section, we give representatives of each of these classes. All of the examples are substitutive, and thus also give examples of rowsparse sofic shifts when the paths are drawn on $\mathbb{Z}^2$.

The three subcases of minimal path spaces we exhibit are ones where some path visits some cell infinitely many times, and

[14] Such paths are known as *coboundaries*.

[15] Similar examples can be found in the ascending and descending case when the paths are graphed as two-dimensional configurations and we consider non-horizontal rows.

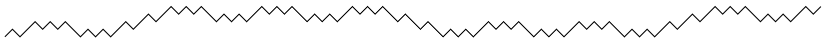

Figure 4: Part of a typical path in Example 8.

Figure 5: Part of a typical path in Example 9.

- every path visits every cell (that it visits at all) infinitely many times (Example 8, Figure 4),
- some path visits every cell finitely many times (Example 9, Figure 5), or
- some path visits some cell finitely many times, but every path visits some cell infinitely many times (Example 10, Figure 6).

Write $\nearrow = 1$ and $\searrow = -1$.

Example 8: There is a minimal subsystem of `Path` where every path visits every cell that it visits either zero or infinitely many times, and always visits infinitely many cells. Consider the substitution τ_1 defined by

$$\nearrow \mapsto \nearrow\nearrow\searrow\searrow\nearrow\nearrow; \quad \searrow \mapsto \searrow\searrow\nearrow\nearrow\searrow\searrow.$$

Let Y be the orbit closure of the fixed-point of this substitution, considered as a minimal subshift of `DPath`. Then $X = \phi^{-1}(Y) \subset$ `Path` has the property that every path visits every cell (that it visits at all) infinitely many times.

Namely, to each path $w \in \{\nearrow, \searrow\}^\ell$, associate a point $z_w \in \mathbb{N}^\mathbb{Z}$ where $(z_w)_h$ records the number of times the path w visits $h \in \mathbb{Z}$, relative to the starting point of w. More precisely, $(z_w)_h = k$ if there are exactly k distinct $j \in [0, \ell]$ such that $\sum_{i=0}^{j} w_i = h$. By induction, we see that for $w = \tau_1^n(\nearrow)$ we have $(z_w)_i \geq 2^n$ or $(z_w)_i = 0$ for all $i \in \mathbb{Z}$, and similarly the support of z_w doubles in size after each substitution.

From this, it easily follows that in every path in X, every cell (that is visited at all) is visited infinitely many times, and that no path in X is bounded. △

Example 9: Consider the substitution τ_2 defined by

$$\nearrow \mapsto \nearrow\nearrow\searrow\nearrow\nearrow; \quad \searrow \mapsto \searrow\searrow\nearrow\searrow\searrow.$$

Define X as in the previous example. It is clear that paths in X are not uniformly ascending or bounded. Nevertheless, there is a path in X where every cell is visited a finite number of times. Namely, the limit of $\tau_2^m(\nearrow).\tau_2^m(\nearrow)$ as $m \to \infty$ is such a path, where the decimal point denotes the center of the path. △

Example 10: Consider the substitution τ_3 defined by

$$\nearrow \mapsto \nearrow\nearrow\searrow; \quad \searrow \mapsto \nearrow\searrow\searrow.$$

Figure 6: Part of the path in Example 10 where one height is visited only once, and all other heights are visited infinitely many times. The gray line marks the height visited only once.

Define X as in the previous examples. It is clear that paths in X are not uniformly ascending or bounded. It is clear that no path of X is bounded or ascending. Nevertheless, there is a path in X where every cell is visited a bounded number of times: as in the previous example, consider the configurations $z \in \mathbb{N}^{\mathbb{Z}}$ corresponding to $\tau^n(\nearrow)$ or $\tau^n(\searrow)$. Then by induction one can show that the support of z is the interval $[0, n+1]$, $z_0 = 1$ and $z_i \geq 2^{n-1}$ for $i \in [1, n+1]$.

It easily follows that all paths in X have at most one cell that is not visited infinitely many times. A path where 0 is visited once, negative integers are never visited, and positive integers are visited infinitely many times, is obtained as the limit of $\tau^n(\searrow).\tau^n(\nearrow)$. △

8 Acknowledgements

The author would like to thank Pierre Guillon, Mariusz Lemańczyk, Kaisa Matomäki, Ronnie Pavlov, Siamak Taati, Ilkka Törmä and Charalampos Zinoviadis for many interesting discussions.

References

[1] Stål O. Aanderaa and Harry R. Lewis. Linear sampling and the ∀∃∀ case of the decision problem. *The Journal of Symbolic Logic*, 39:519–548, 9 1974.

[2] Nathalie Aubrun and Mathieu Sablik. Simulation of effective subshifts by two-dimensional subshifts of finite type. *Acta Appl. Math.*, 126(1):35–63, August 2013.

[3] Alexis Ballier, Bruno Durand, and Emmanuel Jeandel. Structural aspects of tilings. In Pascal Weil Susanne Albers, editor, *Proceedings of the 25th Annual Symposium on the Theoretical Aspects of Computer Science*, pages 61–72, Bordeaux, France, February 2008. IBFI Schloss Dagstuhl. 11 pages.

[4] Alexis Ballier and Emmanuel Jeandel. Structuring multi-dimensional subshifts. *ArXiv e-prints*, September 2013.

[5] Sebastián Barbieri, Jarkko Kari, and Ville Salo. *The Group of Reversible Turing Machines*, pages 49–62. Springer International Publishing, Cham, 2016.

[6] A. Bartnicka, S. Kasjan, J. Kułaga-Przymus, and M. Lemańczyk. $\mathscr{B}$-free sets and dynamics. *ArXiv e-prints*, September 2015.

[7] Robert Berger. The undecidability of the domino problem. *Mem. Amer. Math. Soc. No.*, 66, 1966. 72 pages.

[8] Mike Boyle and Bruce Kitchens. Periodic points for onto cellular automata. *Indagationes Mathematicae*, 10(4):483 – 493, 1999.

[9] Mike Boyle and Douglas Lind. Expansive subdynamics, 1997.

[10] Mike Boyle, Ronnie Pavlov, and Michael Schraudner. Multidimensional sofic shifts without separation and their factors. *Transactions of the American Mathematical Society*, 362(9):4617–4653, 2010.

[11] Alexandre I. Danilenko. Strong orbit equivalence of locally compact Cantor minimal systems. *International Journal of Mathematics*, 12(01):113–123, 2001.

[12] Benjamin Hellouin de Menibus, Ville Salo, and Guillaume Theyssier. Work in progress.

[13] Bruno Durand, Andrei Romashchenko, and Alexander Shen. Effective closed subshifts in 1D can be implemented in 2D. In *Fields of logic and computation*, volume 6300 of *Lecture Notes in Comput. Sci.*, pages 208–226. Springer, Berlin, 2010.

[14] Bruno Durand, Andrei Romashchenko, and Alexander Shen. Fixed-point tile sets and their applications. *J. Comput. System Sci.*, 78(3):731–764, 2012.

[15] Alan Forrest. Symmetric cocycles and classical exponential sums. *Colloquium Mathematicae*, 84/85(1):125–145, 2000.

[16] Thierry Giordano, Ian F. Putnam, and Christian F. Skau. Full groups of cantor minimal systems. *Israel Journal of Mathematics*, 111(1):285–320, 1999.

[17] Pierre Guillon and Emmanuel Jeandel. Infinite Communication Complexity. *ArXiv e-prints*, January 2015.

[18] Pierre Guillon and Gaétan Richard. Asymptotic behavior of dynamical systems and cellular automata. *ArXiv e-prints*, April 2010.

[19] Richard H. Herman, Ian F. Putnam, and Christian F. Skau. Ordered bratteli diagrams, dimension groups and topological dynamics. *International Journal of Mathematics*, 03(06):827–864, 1992.

[20] Michael Hochman. On the dynamics and recursive properties of multidimensional symbolic systems. *Invent. Math.*, 176(1):131–167, 2009.

[21] Michael Hochman. Non-expansive directions for $\mathbb{Z}^2$ actions. *Ergodic Theory and Dynamical Systems*, 31:91–112, 1 2011.

[22] Michael Hochman and Tom Meyerovitch. A characterization of the entropies of multidimensional shifts of finite type. *Ann. of Math. (2)*, 171(3):2011–2038, 2010.

[23] Jarkko Kari. The nilpotency problem of one-dimensional cellular automata. *SIAM J. Comput.*, 21(3):571–586, 1992.

[24] Jarkko Kari. Theory of cellular automata: a survey. *Theoret. Comput. Sci.*, 334(1-3):3–33, 2005.

[25] Jarkko Kari. *SOFSEM 2008: Theory and Practice of Computer Science: 34th Conference on Current Trends in Theory and Practice of Computer Science, Nový Smokovec, Slovakia, January 19-25, 2008. Proceedings*, chapter On the Undecidability of the Tiling Problem, pages 74–82. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.

[26] Jarkko Kari. Decidability and undecidability in cellular automata. *Int. J. General Systems*, 41(6):539–554, 2012.

[27] Douglas Lind and Brian Marcus. *An introduction to symbolic dynamics and coding*. Cambridge University Press, Cambridge, 1995.

[28] James Maynard. Small gaps between primes. *Annals of Mathematics*, 181(1):383–413, 2015.

[29] Shahar Mozes. Tilings, substitution systems and dynamical systems generated by them. *J. Analyse Math.*, 53:139–186, 1989.

[30] Nic Ormes and Ronnie Pavlov. Extender sets and multidimensional subshifts. *Ergodic Theory and Dynamical Systems*, 36:908–923, 5 2016.

[31] Ronnie Pavlov and Michael Schraudner. Classification of sofic projective subdynamics of multidimensional shifts of finite type. submitted.

[32] Ronnie Pavlov and Michael Schraudner. Classification of sofic projective subdynamics of multidimensional shifts of finite type. *Transactions of the American Mathematical Society*, 367(5):3371–3421, 2015.

[33] Raphael M. Robinson. Undecidability and nonperiodicity for tilings of the plane. *Invent. Math.*, 12:177–209, 1971.

[34] Ville Salo. Decidability and Universality of Quasiminimal Subshifts. *ArXiv e-prints*, November 2014.

[35] Ville Salo. *Subshifts with Simple Cellular Automata*. PhD thesis, University of Turku, 2014.

[36] Ville Salo and Ilkka Törmä. Computational aspects of cellular automata on countable sofic shifts. *Mathematical Foundations of Computer Science 2012*, pages 777–788, 2012.

[37] Ville Salo and Ilkka Törmä. Computational aspects of cellular automata on countable sofic shifts. *Mathematical Foundations of Computer Science 2012*, pages 777–788, 2012.

[38] Ville Salo and Ilkka Törmä. Constructions with countable subshifts of finite type. *Fundam. Inf.*, 126(2-3):263–300, April 2013.

[39] Klaus Schmidt. *Dynamical systems of algebraic origin*, volume 128 of *Progress in Mathematics*. Birkhäuser Verlag, Basel, 1995.

[40] Yitang Zhang. Bounded gaps between primes. *Annals of Mathematics*, 179(3):1121–1174, 2014.

[41] C. Zinoviadis. Hierarchy and Expansiveness in Two-Dimensional Subshifts of Finite Type. *ArXiv e-prints*, March 2016.