

Subshifts with sparse traces

Ville Salo

University of Turku
vosalo@utu.fi

February 5, 2019

Abstract

We study two-dimensional subshifts whose horizontal trace (a.k.a. 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 traces 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 constraints. 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, 38, 27, 10, 24, 14]. While two-dimensional 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 in two dimensions [35], and the (horizontal) traces, i.e. sets of rows appearing in configurations, of sofic shifts exactly coincide with the class of computable subshifts [22, 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 [18, 36], and many simple-looking questions are open.

Cellular automata are another point of view to SFTs, in the sense that in many situations they correspond to 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 [26, 28].

However, also limitations are known. In this paper, we concentrate on limitations arising from having very simple horizontal traces,¹ and from various types of nilpotency. A well-known result of this type is that the horizontal trace 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 traces of SFTs given in [37].

Other similar results in the nilpotency framework appear in [19] and [41], where the assumption of SFTness is in some sense replaced with determinism, and the multidimensional subshift is not discussed explicitly. In [19] it is in particular shown that on certain one-dimensional SFTs, asymptotically nilpotent cellular automata are nilpotent, and in [41] 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 [37, 19, 41] mentioned in the previous paragraphs.² However, our result is much more general, in that it applies to all subshifts with sparse traces, 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 trace 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 cluster fractal for purely topological reasons.

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

1.1 The main theorem


In this section, we state the main structure theorems, and need some terminology for points of various forms. A *finite point* is a point whose support is finite (that is, all but finitely many cells contain the symbol 0). A *highway* (Definition 4.4) is a configuration whose support consists of a single ascending path of bounded width whose movement is guided by a uniformly recurrent sequence. *Cluster fractals* (Definition 2.2) are recursively defined configurations where finite patterns padded with zeroes (clusters) are collected into larger finite clusters with larger and larger separation. See Figure 1 for illustrations.

When studying a minimal subshift, it is useful to express it in a substitutive way, by listing its words of some length, organizing them into longer words that occur in the subshift, each long enough to contain all previous ones, and continuing inductively. This is often formalized by Bratteli-Vershik diagrams [21].

An *almost minimal* subshift (Definition 3.1) is one containing a fixed-point for the dynamics, with the property that every point except the fixed one gener-

¹Of 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 [44].

²All 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 cluster fractal: a subshift based on Cantor's dust.

Figure 1: Folklore examples of the three cases that can appear in Theorem 1.2. The trace of the cluster fractal given here is not actually sparse, as we are not aware of such examples in the literature – see Section 7 for a cluster fractal with a sparse trace.

ates the whole subshift. We can give a substitutive structure to such subshift in the following sense (proved as Theorem 5.1), which is a one-dimensional version of our main theorem:

Theorem 1.1. *Let X be any almost minimal one-dimensional subshift. Then X is the orbit-closure of a finite point or a cluster fractal.*

Our main result is Theorem 1.2 below (proved as Theorem 5.4). It classifies two-dimensional subshifts that are almost minimal, and whose trace is *sparse* (Definition 3.3), that is, contains only configurations with finitely many nonzero symbols.

Theorem 1.2. *Let X be any almost minimal two-dimensional subshift with a sparse trace. Then X is the orbit-closure of a finite point, a highway, or a cluster 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 as Theorem 6.1:

Theorem 1.3. *Let X be any two-dimensional subshift whose trace 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 cluster fractal whose trace is sparse.

The way we prove Theorem 1.2 is by taking an arbitrary nonzero point in the subshift X (as the orbit of any such point is dense in the subshift), and attempting to build a cluster 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 trace, 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 1.2. The first one is the following theorem of Pavlov and Schraudner [37, Theorem 6.4]:

Theorem 1.4. *If a \mathbb{Z} -subshift Y has universal period and is not a finite union of periodic points, then it is not the horizontal trace 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 1.2 is straightforward, but somewhat technical, and we do this in Theorem 6.6. However, the following weaker version of Theorem 1.4 is a direct corollary of Theorem 1.2:

Theorem 1.5. *If a nontrivial \mathbb{Z} -subshift Y is sparse, then it is not the horizontal trace of any \mathbb{Z}^2 -SFT X .*

This follows from Theorem 1.2 by restricting to an almost minimal subshift of X and observing that horizontal translates of configurations of any of the three types in Theorem 1.2 (finite configurations, highways or cluster fractals) can be freely glued within the SFT to obtain a non-sparse trace (even positive entropy). This is proved as Theorem 6.4.

Nilpotency is an important notion in the theory of cellular automata, and it is perhaps the best-known undecidable property for cellular automata on the full shift [25, 1]. The next theorem is due to Guillon and Richard [19, 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 1.6. *Let X be a one-dimensional transitive SFT and $f : X \rightarrow X$ a cellular automaton. Then f is nilpotent if and only if it is asymptotically nilpotent.*

This is a direct corollary of the following, proved as Proposition 6.15, which in turn is a direct corollary of Theorem 1.2. By a *glider* we mean a non-trivial finite configuration that is shifted by some power of the cellular automaton.

Proposition 1.7. *Let $f : X \rightarrow X$ 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 with bounded gaps would clearly give a configuration that does not tend to the all-zero configuration. Thus Theorem 1.6 follows from Proposition 1.7.

We also state two decidability results obtained by the author and coauthors, for which we obtain new proofs. The following is one of the main decidability results in [41].

Theorem 1.8. *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 1.7 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 1.2, after interpreting the topological full group of a one-dimensional subshift as a two-dimensional subshift in a suitable way. See Theorem 6.20 for the proof.

Theorem 1.9. *If X is a one-dimensional 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 the equivalence $\forall a \in [0, |w| - 1] : n + a \text{ is prime} \iff w_a = 1$. See Proposition 6.23 for the proof.

Proposition 1.10. *The subshift $X_{\mathcal{P}}$ contains either a finite point or a cluster 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 traces and path spaces.

To prove the theorems listed in the previous section, we only use the first two cases of the classification result Theorem 1.2, as in all those cases cluster fractals turn out to be impossible. However, in general cluster fractals do exist, and in some sense they are the more common case, as the other two cases are in some sense degenerate cluster fractals (finite points are cluster fractals where the separation between clusters is infinite, and infinite paths are clusters of infinite

size). We show in particular that subshifts with sparse traces generated by cluster fractals need not contain any nonzero configurations where the diameter of the support of every row is bounded. In fact, this fails as badly as possible, see Example 7.1.

We also discuss the computational side of subshifts with sparse traces. We show that a subshift where the support of every point is an ascending path is Π_1^0 if and only if it is sofic when the movement speed of the path is bounded, but construct a Π_1^0 non-sofic subshift with sparse traces. (See Section 2 or [11] for definitions.)

Our main technical tool besides almost minimality are uniformly recurrent paths. Paths are (up to a small change in viewpoint) 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 no path is bounded but 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.

2 Definitions

See [31] for a reference on symbolic dynamics. A topologically closed shift-invariant subset of $\Sigma^{\mathbb{Z}^d}$ is called a $(\mathbb{Z}^d\text{-})$ subshift, and more generally we call expansive actions on closed subsets of Cantor space subshifts. Two subshifts are *conjugate* if there is a shift-commuting homeomorphism between them. Conjugate subshifts have the same dynamical properties. Cantor space is metrizable, and on occasion we use a metric d on this space (the fact d has multiple meanings should not cause confusion).

Points are elements of $\Sigma^{\mathbb{Z}^d}$; we also call them *configurations*. A *unary point* is a point $x \in \Sigma^{\mathbb{Z}^d}$ with $\forall \mathbf{v} \in \mathbb{Z}^d : x_{\mathbf{v}} = a$, for some fixed $a \in \Sigma$. We write this as $a^{\mathbb{Z}^d}$. The *shift*³ or *translation* on $\Sigma^{\mathbb{Z}^d}$ is $\sigma^{\mathbf{v}}(x)_{\mathbf{u}} = x_{\mathbf{u}+\mathbf{v}}$. We also call points *configurations*.

Most subshifts considered in this article contain the unary point $0^{\mathbb{Z}^d}$, and are *pointed*, in the sense that this *zero point* $0^{\mathbb{Z}^d}$ is thought of as being part of the structure (in particular morphisms, by default, map the zero point to the zero point). 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. 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*.

A subshift is *SFT* (subshift of finite type) if it is defined by a finite set of forbidden patterns (in the sense that it is the largest subshift of some full shift that misses a defining clopen set), and *sofic* if it is the image of an SFT under a shift-commuting continuous function.

³Our shift convention is a standard one in the symbolic dynamics on abelian groups, where the contents of points are actually shifted in direction $-\mathbf{v}$ by the action of the vector \mathbf{v} . This action is contravariant in general, but it is of course covariant in the abelian case.

A Π_1^0 -*subshift* is a subshift $X \subset \Sigma^{\mathbb{Z}}$ that is on the Π_1^0 level of the arithmetical hierarchy (equivalently the lightface Borel hierarchy). In other words, there exists a Turing machine that outputs a sequence $w_1, w_2, \dots, w_i \in \Sigma^*$, such that X is defined by forbidding the words w_i . See [11].

In the study of nilpotency and pointed subshifts, two particularly important subshifts are the *all-zero* or *trivial* subshift $\{0^{\mathbb{Z}}\}$, and the subshift $X_{\leq 1} = \overline{\mathcal{O}(\dots 00100\dots)}$, called the *sunny-side-up subshift*, *one-one subshift* or *one-or-less subshift*, consisting of two-sided sequences over a binary alphabet with at most one nonzero symbol. The *two-dimensional sunny-side-up* is defined similarly as the binary subshift whose configurations sum to at most one.

A (d -dimensional) *pattern* is a function $P : D \rightarrow \Sigma$ 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 (or appears) at \mathbf{v} in x* if $\sigma^{\mathbf{v}}(x)|_{D(P)} = P$. This allows the use of descriptive terminology such as ‘ P occurs at a bounded distance from every occurrence of Q ’. For a subshift X we define $P \sqsubset X \iff \exists x \in X : P \sqsubset x$.

In the one-dimensional case, patterns with a contiguous domain are usually called *words*, and the set of words occurring in a one-dimensional subshift (or a point) is called its *language*.

Abusing notation, if $P : D \rightarrow \Sigma$ is a pattern and $E \subset \mathbb{Z}^d$, we write $P|_E : E \rightarrow \Sigma$ for the pattern defined by $P|_E(\mathbf{v}) = P|_D(\mathbf{v})$ for $v \in D \cap E$, $P|_E(\mathbf{v}) = 0$ otherwise. In particular for a pattern $P : D \rightarrow \Sigma$, $P|_{\mathbb{Z}^d}$ is the corresponding configuration in $\Sigma^{\mathbb{Z}^d}$ padded with zeroes.

Two-dimensional 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 (*horizontal*) *trace* 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 , and the row y defined by $y_j = x_{j,i}$ is denoted by x_i . The trace is a \mathbb{Z} -subshift. We mostly stay in the two-dimensional case, but for \mathbb{Z}^d -subshifts, by the trace we mean the one-dimensional subshift whose configurations are seen on the first axis; in particular one-dimensional subshifts are their own traces. On occasion, we refer to traces in other rational directions, and the meaning should be clear (configurations seen along stripes).

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

A *direction* is an element $\mathbf{v} \in \mathbb{R}^2 \setminus \{(0, 0)\}$, under the equivalence relation that \mathbf{v} and $r\mathbf{v}$ are considered equal for all $r \in \mathbb{R}^+ = \{r \in \mathbb{R} \mid r > 0\}$. A direction $\mathbf{v} \in \mathbb{R}^2$ is *rational* if it is equal (as a direction) to a point in \mathbb{Q}^2 , equivalently one of \mathbb{Z}^2 . Note that any direction is defined by a point of the unit circle up to translation, but a rational direction may not be equal to a point of $\mathbb{Q}^2 \cap S^1$.

A *cellular automaton* is a shift-commuting continuous function $f : X \rightarrow X$ on a \mathbb{Z} -subshift X . We say f is *nilpotent* if $\exists n \in \mathbb{N} : \forall x \in X : f^n(x) = 0^{\mathbb{Z}}$

and *asymptotically nilpotent* if $f^n(x) \rightarrow 0^{\mathbb{Z}}$ 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 $f : X \rightarrow X$ is the \mathbb{Z}^2 -subshift Y whose trace is the limit set of f , and in every configuration,⁴ the $(i+1)$ th row is the f -image of the i th row for every $i \in \mathbb{Z}$. A *half-plane in direction* $\mathbf{v} \in \mathbb{R}^2$ is a half-plane whose boundary is orthogonal to \mathbf{v} and whose interior contains all vectors $-\mathbf{r}\mathbf{v}$ for large enough $r \in \mathbb{N}$. A subshift X is *deterministic* in direction \mathbf{v} if, for every half-plane $H \subset \mathbb{R}^2$ in direction \mathbf{v} , the map $x \mapsto x|_{H \cap \mathbb{Z}^2}$ is injective on X .

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 $\exists n \in \mathbb{Z} : \sigma^n(x) = x$, and *temporal periodicity* means $f^n(x) = x$ for some $x \in X, n > 0$. When we say x is *eventually periodic* we typically refer to spatial periodicity, and mean that $x = \dots uuuvww \dots$ for some words u, v, w . For such x , the minimal choice of $|v|$ is the *preperiod* of x .

On \mathbb{Z}^d , we use the metric $d(\mathbf{u}, \mathbf{v}) = \max |\mathbf{u}_i - \mathbf{v}_i|$ and sometimes the ℓ^2 -norm $\|\mathbf{u}\|_2 = \sqrt{\sum \mathbf{u}_i^2}$. The zero-vector in every dimension is called $\mathbf{0}$. Write $B_r(\mathbf{v})$ for the ball of radius r around $\mathbf{v} \in \mathbb{Z}^d$, $B_r(\mathbf{v}) = \{\mathbf{u} \mid d(\mathbf{u}, \mathbf{v}) \leq r\}$, and $B_r = B_r(\mathbf{0})$. 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 $\{(\mathbf{u}, \mathbf{v}) \mid d(\mathbf{u}, \mathbf{v}) \leq r\}$ 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-cluster*⁵ is an r -padded pattern for which

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

Note that a cluster is uniquely determined by its support and need not be finite (though we mostly deal with finite clusters). In every configuration, every nonzero cell i is in precisely one r -cluster, for every r : the r -cluster is the r -connected component of the nonzero cells that i belongs to, r -padded with 0-symbols.

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 $R|_{D(P)} = P$ and $R|_{D(Q)} = Q$. We sometimes write this as $R = P + Q$, and write $x + y$ for the zero-gluing of two points when the full patterns they define on \mathbb{Z}^2 can be zero-glued. For an infinite sequence of configurations x^i , we also use the notation $\sum_i x^i$ for the limit of the finite gluings $\sum_{i=1}^k x^i$ when it exists.

A subshift *allows zero-gluing* if there exists r such that when P and Q are r -padded patterns that occur in the subshift, and their zero-gluing exists, then

⁴Note that the *trace of a cellular automaton* refers to the vertical trace of the spacetime subshift in most references, but this should not cause confusion.

⁵These were called ‘blobs’ in a previous version.

the zero-gluing 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* [31] is such an example, and in higher dimensions we can take the subshift where every row independently contains points from the even shift.

Remark 2.1. *The definition of zero-gluing given here differs from that in earlier versions of the paper, where we instead defined zero-gluing as the weaker notion of being able to glue two configurations together as long as their supports are sufficiently separated. Zero-gluing in its weaker form is not equivalent, and the proof of Proposition 6.3 uses the strong notion given above.*

The *essential rowsparseness* of a two-dimensional configuration x is the minimal m such that for some fixed r the support $S \subset \mathbb{Z}$ of every row of x is contained in the union of m many r -balls, that is, $\exists n_1, \dots, n_m : S \subset \bigcup_i n_i + B_r(\mathbf{0})$.

Definition 2.2. *A configuration $x \in (\Sigma_+ \cup \{0\})^{\mathbb{Z}^d}$ is a cluster fractal if it is nontrivial and there is an increasing sequence r_1, r_2, r_3, \dots of natural numbers and for all r_i a finite set B_i of finite r_i -clusters such that for all i ,*

- *each cluster in B_{i+1} is obtained by zero-gluing translates of clusters in B_i and adding the r_{i+1} -padding,*
- *each cluster in B_{i+1} contains a translate of every cluster in B_i ,*
- *each cluster in B_{i+1} contains at least two translates of clusters in B_i with disjoint supports, and*
- *$x = \lim_{j \rightarrow \infty} \sigma^{\mathbf{v}_j}(b_j|_{\mathbb{Z}^d})$ for some sequences $\mathbf{v}_j \in \mathbb{Z}^d$, $b_j \in B_j$.*

3 Almost minimality and traces

A subshift is minimal if it is the orbit closure of every point that it contains. Because trace sparse subshifts (see Definition 3.3) always contain the all-zero point by compactness, they clearly cannot be minimal (unless trivial). The next best thing is almost minimality – every point except the all-zero point generates the subshift in its orbit closure. It turns out that we can generally extract almost minimal subshifts from our subshifts of interest.

Definition 3.1. *Let X be a zero-dimensional compact metric space and let G be a finitely-generated discrete group acting on X by a continuous map $\sigma : G \times X \rightarrow X$. 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 almost minimal if $\mathcal{S}(X) = \{\emptyset, Y, X\}$ for some one-point system $Y = \{y\} \neq X$.*

Remark 3.2. *Note that in our definition a trivial system is not almost minimal. Almost minimality is discussed in [12]. We called systems with finite $\mathcal{S}(X)$ quasiminimal in [39]. Systems with a unique minimal subsystem are called essentially minimal [21]. Clearly almost minimal subshifts are essentially minimal and quasiminimal, but neither essential minimality nor quasiminimality implies almost minimality by easy examples, and essential minimality and quasiminimality are orthogonal concepts.*

Definition 3.3. A subshift $X \subset \Sigma^{\mathbb{Z}^d}$ is k -sparse if every configuration $x \in X$ satisfies

$$|\{\mathbf{v} \in \mathbb{Z}^d \mid x_{\mathbf{v}} \neq 0\}| \leq k.$$

A subshift is sparse if it is k -sparse for some $k \in \mathbb{N}$. A \mathbb{Z}^d -subshift is (nonuniformly) trace sparse if its trace is sparse, and uniformly trace sparse if its trace is k -sparse for some k . We say that a \mathbb{Z}^d -subshift is trace countable if its trace is a countable subshift.

A (trace) countable subshift obviously need not be (trace) sparse. A (trace) sparse subshift need not be uniformly (trace) sparse:

Example 3.4: There is a sparse \mathbb{Z} -subshift that is not uniformly sparse: Let $x^n \in \{0, 1\}^{\mathbb{Z}}$ be the point where $x_i^n = 1 \iff i \in n\mathbb{Z} \cap [0, n^2]$. Then $\mathcal{O}(\{x^n \mid n \in \mathbb{N}\})$ is sparse, but not uniformly sparse. \circ

In the next sections, we give some technical tools that are helpful when applying the main theorem outside the almost minimal trace sparse setting—namely we give sufficient conditions that allow us to extract almost minimal trace 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, trace sparseness is implied by various weaker properties. These conditions are somewhat different, and their intersection gives the following lemma. A configuration $x \in \Sigma^{\mathbb{Z}}$ is *eventually zero to the right* if $x_i = 0$ for all large enough $i \in \mathbb{N}$.

Lemma 3.5. Let a subshift $X \subset \Sigma^{\mathbb{Z}^2}$ with trace 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 trace sparse subsystem.

Proof. If Y is sparse, apply any of the four items of Lemma 3.7 to obtain an almost minimal subsystem $Z \subset X$, and then apply Lemma 3.9 or Lemma 3.10 to conclude Z is uniformly trace sparse.

If Y is countable and the only periodic point is $0^{\mathbb{Z}}$, apply item 1 of Lemma 3.7 to obtain an almost minimal subsystem $Z \subset X$, and then apply Lemma 3.9 to conclude that Z is uniformly trace sparse. For this, observe that the only periodic point $0^{\mathbb{Z}}$ cannot be isolated in the trace of Z as $0^{\mathbb{Z}^2}$ is in the orbit closure of every point of Z .

If every point of Y is eventually zero to the right, apply item 2 or item 4 of Lemma 3.7 and then Lemma 3.10. \square

3.1 Extracting an almost minimal subsystem

Lemma 3.6. 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 nonzero 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 $x \in \bigcap_i C_i \subset X_i$, so also the intersection contains a nonzero symbol. The result follows from Zorn's lemma. \square

Lemma 3.7. *Let a subshift $X \subset \Sigma^{\mathbb{Z}^d}$ with trace 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. \square

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 extent this stronger version has a converse, 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 not necessary:

Example 3.8: There exists a subshift $X \subset \Sigma^{\mathbb{Z}^2}$ which is almost minimal and whose trace 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 . \circ

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 trace (see Section 7).

The first two conditions of Lemma 3.7 are not comparable for general subshifts: The subshift that is the orbit closure of 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 Trace sparseness of almost minimal subshifts

Lemma 3.9. *Let X be an almost minimal \mathbb{Z}^2 -subshift with countable trace Y such that Y has no isolated periodic points. Then X is uniformly trace sparse.*

Proof. Let Y be the trace 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 . Identify w with the corresponding two-dimensional pattern of shape $|w|$ -by-1. By almost minimality, there is a bound $C > 0$ such that whenever a point $x \in X$ has a nonzero symbol $x_{(i,j)} \neq 0$, w appears at $x_{(i',j')} \in \mathbb{Z}$ for some $d((i',j'),(i,j)) < C$. this forces the row $x_{j'}$ to be a translate of y .

If $z \in Y$ has a support of at least size $(2C+1)^2$, and z is a row in x , then at a bounded distance of each of the many nonzero symbols of z in x the isolating pattern w appears. In particular by the pigeonhole principle, two occurrences of w are on the same row, and thus y must be periodic, since we find two occurrences of the isolating pattern w in it, at a nonzero offset. But then y is an isolated periodic point, which contradicts the assumption. Thus Y must be uniformly sparse. \square

Lemma 3.10. *Let X be an almost minimal subshift with trace Y such that every point of Y is eventually zero to the right. Then X is uniformly trace sparse.*

Proof. If every word $w \sqsubset Y$ can be extended to the right by a word $0^k a$ with $a \neq 0$, 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 an infinite word of all zeroes 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 uniformly trace sparse. \square

4 Paths and path covers

4.1 Paths

Let \mathbf{Path}_r be the space of all functions $p : \mathbb{Z} \rightarrow \mathbb{Z}$ such that $p(0) = 0$ and $|p(i+1) - p(i)| \leq r$ for all i , and $\mathbf{Path} = \bigcup_{r \in \mathbb{N}} \mathbf{Path}_r$. We give \mathbf{Path} the product topology and the shift action $\tau : \mathbf{Path} \rightarrow \mathbf{Path}$ defined by $\tau(p)(i) = p(i+1) - p(1)$ (which is a homeomorphism). Let \mathbf{DPath}_r be the space of all functions $p : \mathbb{Z} \rightarrow \mathbb{Z}$ with $p(\mathbb{Z}) \subset [-r, r]$ and $\mathbf{DPath} = \bigcup_{r \in \mathbb{N}} \mathbf{DPath}_r$. We give it the product topology and the usual shift action, also denoted σ . The two spaces are conjugate by the discrete derivative $\phi : \mathbf{Path} \rightarrow \mathbf{DPath}$ defined by $\phi(p)(i) = p(i+1) - p(i)$.

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

All of the following definitions will apply to both kinds of paths, in the sense that if we have defined a property P for elements of \mathbf{Path} we will say a path $p \in \mathbf{DPath}$ has property P if $\phi^{-1}(p)$ has property P , and vice versa.

An *ascending path* is a path $p \in \mathbf{DPath}$ such that for some $m \in \mathbb{N}$, $\sum_{i=j}^{j+m} p(i) > 0$ for all j , and symmetrically, we define *descending paths*. A path $p \in \mathbf{Path}$ 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 $p \in \mathbf{Path}$ with $p^{-1}(n)$ finite for all $n \in \mathbb{Z}$, and otherwise p is *infinite-to-one*. A *bounded path* is a path $p \in \mathbf{Path}$ with $p(\mathbb{Z})$ contained in some $[-m, m]$. A set of paths $X \subset \mathbf{DPath}$ is *uniformly ascending* (resp. *descending*, *bounded-to-one*, *bounded*) if every path in X has the property, and the constant m can be chosen independently of the path. In the case of ascending paths, we call this the *ascension constant*.

Uniform recurrence and minimality, for paths and systems of paths, mean the usual dynamical notions, and in \mathbf{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 two-dimensional case, we define analogous systems \mathbf{Path}^2 and \mathbf{DPath}^2 of paths $p : \mathbb{Z} \rightarrow \mathbb{Z}^2$. For such paths, we define ascending and descending paths differently: to a path $p : \mathbb{Z} \rightarrow \mathbb{Z}^2$, we associate its *height path* $p' : \mathbb{Z} \rightarrow \mathbb{Z}$ by forgetting the first coordinate of $p(n)$ for all n . An *ascending path* $p : \mathbb{Z} \rightarrow \mathbb{Z}^2$ is one where the height path $p' : \mathbb{Z} \rightarrow \mathbb{Z}$ is ascending, and *descending paths* are defined symmetrically.

A (Σ) -*colored path* is $p : \mathbb{Z} \rightarrow \mathbb{Z} \times \Sigma$ or $p : \mathbb{Z} \rightarrow \mathbb{Z}^2 \times \Sigma$ for a finite set of colors Σ . We write \mathbf{Path}_Σ and \mathbf{Path}_Σ^2 for the spaces of Σ -colored path. All notions defined above generalize in an obvious way to colored paths (for geometric definitions we forget the colors, and for dynamical ones, we include them).

Especially in informal geometric explanations, we will also discuss finite subpaths of infinite paths, which need not go through the origin. These are simply functions $p : \{a, \dots, b\} \rightarrow \mathbb{Z}$ or $p : \{a, \dots, b\} \rightarrow \mathbb{Z}^2$ for some $a, b \in \mathbb{Z}$.

We give a classification theorem for minimal path spaces.

Theorem 4.1. *Let $X \subset \mathbf{Path}_r$ 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 uniformly 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 by the syndeticity of w (by minimality) 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, there is

a cut path that forces the rest of the path to stay below or above it. This is easily seen to contradict minimality.

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. \square

Path spaces of the last type are not encountered in the proofs of our main results, but it is perhaps the most interesting case, and it splits into further subtypes. 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 6.13. The proof is exactly the same.⁶

Theorem 4.2. *Let $R \subset [-r, r]$ be a finite set of real numbers and let $X \subset R^{\mathbb{Z}}$ be a minimal subshift. Define*

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

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 uniformly 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

The (acyclic) r -path cover of a subshift $X \subset \Sigma^{\mathbb{Z}^d}$ is the subshift $\text{PC}_r(X) \subset X \times (\{\#\} \cup [-r, r]^d)^{\mathbb{Z}^d}$ of configurations (x, y) such that

$$y_{\mathbf{u}} \neq \# \implies x_{\mathbf{u}} \neq 0 \wedge y_{\mathbf{u}+\mathbf{v}} \neq \#$$

for all $\mathbf{u} \in \mathbb{Z}^d$, such that in the directed graph $G(y)$ with nodes $\{\mathbf{u} \mid y_{\mathbf{u}} \neq \#\}$ and edges $\{(\mathbf{u}, \mathbf{u} + \mathbf{v}) \mid \mathbf{v} \neq \mathbf{0}, y_{\mathbf{u}} = \mathbf{v}, y_{\mathbf{u}+\mathbf{v}} \neq \#\}$, every node has in-degree at most one (the out-degree is automatically also at most one), and there are no cycles. In $G(y)$, every node is part of a unique maximal (finite, one-way infinite or two-way infinite) path. If $(x, y) \in \text{PC}_r(X)$ is such that $\mathbf{0} \in G(y)$, the path through the origin $\mathbf{0}$ 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 $\text{MPC}_r(X)$ for the set of such configurations, which is a compact (possibly empty) space.

⁶Note that the possible finite sequences of moves still form a discrete set, even though paths themselves may explore a dense set of positions.

Lemma 4.3. *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 $\text{MPC}_r(X)$ 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 arbitrarily large components, so consider an r -path of minimal length between two cells at distance at least $2n + 1$ along the support of x , and construct a pair $(x, y) \in \text{PC}_r(X)$ 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 $\text{PC}_r(X)$ where the maximal path through $\mathbf{0}$ is of length at least n in both directions. Any limit point of such pairs as $n \rightarrow \infty$ gives a point of $\text{MPC}_r(X)$. \square

See Section 7.2 for an example of such a path extraction.

Let us now concentrate on the case $d = 2$. The importance of the space $\text{MPC}_r(X)$ is that it inherits the dynamics of Path^2 in an obvious way, the action of $n \in \mathbb{Z}$ following the path written on the second component for n steps. More precisely, if $(x, y) \in \text{MPC}_r(X)$ and $y_0 = \mathbf{v}$, define $\tau(x, y) = \sigma^{\mathbf{v}}(x, y)$, so that $(\text{MPC}_r(X), \tau)$ becomes a \mathbb{Z} -system.

Definition 4.4. *An r -road (configuration) x is one whose support, up to translation, is the range of a two-way infinite ascending r -path $p : \mathbb{Z} \rightarrow \mathbb{Z}^2$. If we can choose p so that the corresponding colored path q defined by $q(i) = (p(i), x_{p(i)})$ is uniformly recurrent, then we call the configuration x an r -highway. An r -preroad (resp. r -prehighway) configuration is one whose nonzero support contains a translated image of an infinite ascending (resp. ascending and uniformly recurrent) r -path.*

The essential trace sparseness of a road configuration is 1, and the essential trace sparseness of a preroad 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 .

Lemma 4.5. *Let X be a trace sparse subshift. If $\text{MPC}_r(X)$ is nonempty, then X contains an r -preroad.*

Proof. Take any minimal subsystem of $\text{MPC}_r(X)$ in the \mathbb{Z} -action τ and a point (x, y) in it. Then the support of y is the range of a uniformly recurrent path $p \in \text{Path}^2$. It follows from Theorem 4.1 that p is either ascending, descending or not finite-to-one. The last case is impossible because X is trace sparse. 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. \square

We also need the following simple observations.

Lemma 4.6. *Let x be a preroad and $p : \mathbb{Z} \rightarrow \mathbb{Z}^2$ 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 $x_v \neq 0 \implies \exists n : |p(n) - \mathbf{v}| \leq m$. Then let $a = \lceil [-m, m]^2 \rceil$ and construct a path 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 $\mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_{\ell-1}$, where $\ell \leq a$. Define $q(na + j) = \mathbf{v}_j$ for all $j \in [0, \ell - 1]$ and $q(na + j) = \mathbf{v}_1$ 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 constant). \square

Lemma 4.7. *Let X be almost minimal and $x \in X$ be a road. Then x is a highway.*

Proof. Shift x suitably to obtain an ascending path $p \in \text{Path}^2$ with $p(0) = \mathbf{0}$ whose range is equal to the support of x . Let p' be the corresponding colored path $p'(i) = (p(i), x_{p(i)})$. Take a uniformly recurrent path q in the orbit closure of p' in Path_{Σ}^2 . Then q is obtained by taking a limit of some sequence of translates of p' . Define $y \in \Sigma^{\mathbb{Z}^2}$ by $y_{\mathbf{v}} = a$ if $q(n) = (\mathbf{v}, a)$ for some $n \in \mathbb{Z}$, $a \in \Sigma$, and $y_{\mathbf{v}} = 0$ otherwise. One can show that y is well-defined and is in the orbit closure of x by following the translates of p' that tend to q (we need the fact that p is ascending).

Since q is uniformly recurrent, y is a highway. By almost minimality y contains the same patterns as x . It follows that also x must be 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 colored path in the orbit closure of q . \square

Note that the paths p, p' need not be uniformly recurrent. It is always possible to parametrize the range of a uniformly recurrent path in a non-recurrent way, for example for a uniformly recurrent injective $p \in \text{Path}^2$ consider $q(n) = p(n - 1)$ for $n \geq 1$, $q(n) = p(n)$ otherwise. The lemma states that there is always automatically also a uniformly recurrent parametrization (for q , one such parametrization is p).

5 The main theorem

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

Theorem 5.1. *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 cluster fractal.

Proof. Fix a nonzero point $x \in X$ with 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 S is of diameter at most n . Collecting the subwords of x corresponding to these connected components, we obtain a set of finite r -clusters, which must be finite, since the clusters are words of length at most $n + 2r$.

Now, we start building a cluster fractal structure for x . Let $r_1 \in \mathbb{N}$ be arbitrary, and let B_1 be the finite collection of clusters 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 -cluster b whose nonzero support is strictly larger than that of any of the clusters in B_1 . By picking r large enough, this is true for all r -clusters, because b appears at a bounded distance from every nonzero symbol by almost minimality. By increasing r yet more, every cluster in B_1 appears as a subpattern of every r -cluster. Pick $r_2 \gg r_1$ with these properties, and define B_2 as the set of r_2 -clusters.

Continuing inductively, we obtain that x is a cluster fractal. \square

Lemma 5.2. *Let $X \subset \Sigma^{\mathbb{Z}^2}$ be the orbit closure of a nontrivial finite configuration, a highway or a cluster 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 \mathbf{v} in P and follow the colored path p giving the support of x forward and backward from the cell \mathbf{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 and p is ascending, there is a function f of the size of the domain $D(P)$ of P such that if $t > f(|D(P)|)$ then q visits all cells of P and p can never again visit the domain $D(P)$. Since p is uniformly recurrent, the subpath q is traced every m steps for some m , so in particular P occurs at a bounded distance from every cell in the support of x , and thus the same is true for its orbit closure X .

For cluster fractals, let B_i be the sets of r_i -clusters 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 -cluster $B \in B_i$ in x . Now, let y be any point in X and \mathbf{v} any nonzero cell in y , i.e. $y_{\mathbf{v}} \neq 0$. Since y is in the orbit closure of x , it is also a limit of clusters, and thus \mathbf{v} must be contained in the support of a translate of a B_{i+1} -cluster. This cluster contains a copy of B , and thus the pattern P . \square

Finite configurations and highways (even roads) of course always have essential trace sparseness one. Cluster fractals can have infinite essential trace sparseness, and their essential trace sparseness can be any finite number even in the trace sparse case, see Section 7.

The following lemma is easy to prove.

Lemma 5.3. *Let X be the orbit closure of a finite configuration (highway or cluster fractal, respectively). Then every nonzero point in X is a finite configuration (highway or cluster fractal, respectively).*

Theorem 5.4. *Let X be an almost minimal and trace sparse \mathbb{Z}^2 -subshift. 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 cluster fractal.

Proof. By the previous lemma, $X \subset \Sigma^{\mathbb{Z}^2}$ cannot be the orbit closure of two of these types of configurations. An almost minimal trace sparse subshift is uniformly trace sparse by Lemma 3.10. Suppose the trace is k -sparse. 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 5.1. Take a point $x \in X$ and start building a cluster fractal structure for it by increasing r , collecting the r -clusters 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 clusters that each contain at least two clusters of the previous size, we get in the limit that X is the orbit-closure of a cluster fractal (the fact that all B_{i+1} -clusters contain all B_i -clusters 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 cluster in B_i . In this case, x is a finite configuration. In the other case, we have either infinitely many r -clusters or a single infinite r -cluster in x , for some r . In particular, supports of configurations of X contain arbitrarily large r -components for some r . By Lemma 4.3, X then contains a preroad configuration. 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 many $(n \times n)$ -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 $\sigma^{\mathbf{v}}(x^1, \dots, x^m) = (\sigma^{\mathbf{v}}(x^1), \dots, \sigma^{\mathbf{v}}(x^m))$. 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.⁷ Naturally, there is an upper bound on m as a function of r and k , since all paths go through one of the rows $0, 1, \dots, r-1$. Now, for each $i \in [1, m]$ fix an ascending path $p_i : \mathbb{Z} \rightarrow \mathbb{Z}^2$ 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_i , then x^i is a path configuration by Lemma 4.6, 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 \mathbf{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 \mathbf{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 a nearby row in x , there is a nonzero coordinate $(a_i, b) + \mathbf{r}^i$ that lies on the central path, where $\mathbf{r}^i \in [-r, r]^2$ (because the path

⁷The r must be the same, but the ascension constants of the paths can a priori be different.

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 $(a_j, b) + \mathbf{r}^j$ where $\mathbf{r}^j \in [-r, r]^2$.

By translating (x^1, \dots, x^m) by $\sigma^{(a_i, b) + \mathbf{r}^i}$, 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 $\mathbf{v}^{\ell'}$ on one of the rows $0, 1, \dots, r-1$ 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 j th 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 $(y^1, \dots, y^m) \in Y_m$ 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 $\mathbf{v}^{\ell'}$ arbitrarily far from the origin, but at vertical distance at most r from row zero. As $\ell' \rightarrow \infty$, we then find in the limit a tuple $(z^1, \dots, z^{m+1}) \in Y_{m+1}$ where all of the z^i are preroads. This is a contradiction with the maximality of m , and thus x must have been a road. By Lemma 4.7, x is a highway. \square

6 Corollaries

In this section, we give a list of corollaries of Theorem 5.4. The most general extraction theorem obtained from Theorem 5.4 and Lemma 3.5 is the following, which we mentioned in the introduction.

Theorem 6.1. *Let X be any two-dimensional subshift whose trace Y satisfies one of the following:*

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

Then X contains a finite point, a highway, or a (trace sparse) cluster fractal.

Proof. In any of the three cases, Lemma 3.5 applies and we find an almost minimal trace sparse subsystem. The result follows from Theorem 5.4. \square

6.1 SFTs, sofics and zero-gluing

Proposition 6.2. *Let X be a nontrivial trace sparse \mathbb{Z}^2 -subshift such that for some r , in the supports of all configurations of X all r -components are infinite. Then X contains a highway.*

Proof. If X is trace sparse, then it contains an almost minimal trace sparse nontrivial subsystem by Lemma 3.7. Let Y be an almost minimal subshift of X . Then Y is the orbit closure of a finite configuration, a highway or a cluster fractal by Theorem 5.4. Only highways have an infinite r -component for some r . \square

Proposition 6.3. *No trace sparse nontrivial \mathbb{Z}^2 -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 trace is not sparse: For finite and highway configurations, simply glue them to their horizontal translates. For cluster fractals, take an r_i -cluster for r_i larger than the gluing radius, and the cluster must extend to a finite point because we can legally zero-glue it to the all-zero point. \square

Theorem 6.4. *No nontrivial two-dimensional SFT is trace sparse.*

Proof. An SFT clearly allows zero-gluing. \square

Proposition 6.5. *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 a sparse trace.*

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 allows zero-gluing. \square

General sofic shifts can have sparse traces. See Section 7.

A \mathbb{Z} -subshift has *universal period* n ($n \in \mathbb{N}, n \geq 1$) if there exists a bound $M \in \mathbb{N}$ 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 $x|_{\mathbb{Z} \setminus F_x} = y|_{\mathbb{Z} \setminus F_x}$. We say it has *universal period* if it has universal period n for some n . (see Definition 4.3 in [37])

The following is Theorem 6.4 in [37], and it is the non-implementability part of their characterization of zero-entropy sofic traces of two-dimensional subshifts of finite type.

Theorem 6.6. *If a \mathbb{Z} -subshift Y has a universal period and is not a finite union of periodic points, then it is not the trace of any \mathbb{Z}^2 -SFT X .*

Proof. Suppose Y is the trace 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 Y_1, \dots, Y_m . Since there is a uniform bound on the number of coordinates where a point $y \in Y$ can differ from a unary point, we may assume that the zero symbol of Y_i cannot appear in Y_j for $j \neq i$.

Thus, we may assume that X is an SFT whose trace is a finite union of sparse subshifts Y_1, \dots, Y_m with the above disjointness property on zero-symbols. Let Z be a factor of X obtained by mapping all the distinct zero symbols on every row of $x \in X$ to the same symbol 0. Note that no non-unary row is mapped to $0^{\mathbb{Z}}$. Now, Z is clearly trace sparse and nontrivial, so by Lemma 3.7, it contains an almost minimal subsystem W .

By Theorem 5.4, W contains either a finite configuration, a highway configuration or a cluster fractal. If $z \in Z$ is finite or a highway, then clearly the zero-gluing $\sum_{i \in \mathbb{Z}} \sigma^{(ik,0)}(z)$ is in X for some k , since we can simply glue the preimages of z together as on each row they use the same preimage of the zero-symbol on the left and right side of the support of z . This contradicts the assumption that the original subshift had a universal period.

If z is a cluster fractal, then pick any covering point $x \in X$ and some r -cluster b in x with r larger than the maximal size of a forbidden pattern of the SFT, and replace every cell outside b that is not one of the zero symbols with the zero-symbol used on that row of x . This configuration is legal, and all but

finitely many of its rows are periodic. As in the previous paragraph, we can glue together a legal configuration contradicting the universal period. \square

More generally, we obtain that there are singly periodic points in all SFTs whose traces 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 [40] 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 contrasted with [3, Theorem 3.11], where it is shown that all countably infinite two-dimensional SFTs contain a singly periodic point.

A configuration $x \in \Sigma^{\mathbb{Z}^2}$ is *(totally) periodic* if $\{\mathbf{u} \in \mathbb{Z}^2 \mid \sigma^{\mathbf{u}}(x) = x\}$ spans \mathbb{R}^2 as an \mathbb{R} -vector space. We say a point $x \in \Sigma^{\mathbb{Z}^2}$ is *asymptotic to a periodic point* if there exists a totally periodic $y \in \Sigma^{\mathbb{Z}^2}$ such that $x_{\mathbf{v}} = y_{\mathbf{v}}$ for all but finitely many \mathbf{v} . We say x is *strictly asymptotic to a periodic point* if it is asymptotic to a periodic point but is not periodic. The following lemma is easy to prove.

Lemma 6.7. *Suppose $X \subset \Sigma^{\mathbb{Z}^2}$ is an SFT containing a point that is strictly asymptotic to a periodic point. Then the trace of X is uncountable.*

Proposition 6.8. *If $X \subset \Sigma^{\mathbb{Z}^2}$ is an SFT whose trace is a bounded infinite subshift, then it contains a singly periodic point.*

Proof. Suppose X is an SFT whose trace is a bounded infinite subshift, and suppose it contains no singly periodic point. First, we show that X contains only finitely many points which have a horizontal period: The trace of X contains only finitely many periodic points. Considering the periodic rows as a finite alphabet, the set of points with all rows periodic is an SFT under the vertical action. If this one-dimensional SFT were infinite, it would contain an eventually periodic aperiodic point, giving a singly periodic point in X . Thus this one-dimensional SFT is finite, implying that indeed there are only finitely many points with a horizontal period.

Let now n be such that every point with some horizontal period has horizontal and vertical period dividing n , and suppose that X is defined by forbidden patterns of size at most n -by- n . Suppose $0 \notin \Sigma$ and consider the subshift Z obtained from X as the image of the factor map $f : X \rightarrow (\Sigma \cup \{0\})^{\mathbb{Z}^2}$ defined by $f(x)_{\mathbf{v}} = 0$ if $x_{\mathbf{v}+(a,b)} = x_{\mathbf{v}+(a+n,b)} = x_{\mathbf{v}+(a,b+n)} = x_{\mathbf{v}+(a+n,b+n)}$ for all $0 \leq a, b < n$, and $f(x)_{\mathbf{v}} = x_{\mathbf{v}}$ otherwise. The trace of Z is then sparse. By the assumption on the trace of X , the trace of Z is not $\{0^{\mathbb{Z}}\}$. We get that Z contains a nontrivial finite point, a highway or a cluster fractal.

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 trace of Z is clearly an eventually periodic point in the trace of X , and there are only finitely many points with a given period n for the left and right tails, and with a given central pattern length.

For a nontrivial finite configuration we can find a covering point of the following form: for some m , $x_{\mathbf{v}} = x_{\mathbf{v}+(n,0)} = x_{\mathbf{v}+(0,n)} = x_{\mathbf{v}+(n,n)}$ for all $|\mathbf{v}| \geq m$, but x is not periodic. But then x is strictly asymptotic to a periodic point, a contradiction with Lemma 6.7.

Now, let z be a cluster fractal in Z and let $x \in X$ be a covering configuration. We claim that it is enough to show that around any 1-component of the support of z (note that for the definition of 1-components we use the *king grid*, that is, the graph whose path metric is the ℓ^∞ -metric d), we can draw a path along zeroes around it with movements in cardinal directions only. More precisely, that we can find a simple cycle p in the *standard grid* ($\mathbb{Z}^2, \{(u, v) \mid |u_1 - v_1| + |u_2 - v_2| = 1\}$) whose interior (as a Jordan curve) contains the component, and $z_{p(i)} = 0$ for all i .

To see that this suffices, once we have such a path, by the assumption on the factor map and on n , throughout this cycle, we see a horizontal and vertical period of n in x , meaning that the unique point y in X with horizontal and vertical period n that agrees with x in some $2n$ -by- $2n$ square with bottom left corner on the path p agrees with x in all the $2n$ -by- $2n$ blocks on the path p . If b is the cluster corresponding to the component, then by the assumption that its support is contained in the interior of p , we can ignore the configuration outside b , and continue the period to obtain a point strictly asymptotic to the periodic point y , which is a contradiction by Lemma 6.7.

For the existence of the path, consider the family of counterclockwise cycles in the standard grid, whose interior contains the support of b . One such Jordan curve exists because b is finite, and thus there exists such a curve p of minimal area. A simple case analysis shows that every cell $p(i)$ is in the king grid neighborhood of the support of b , or we could decrease the area of the curve by cycling around more internal nodes.⁸ \square

Of course this implies also that two-dimensional SFTs with bounded trace contain doubly periodic points, but this has an easier direct proof by passing to a minimal subsystem.

6.2 Subshifts with deterministic or expansive directions

SFTness can often be replaced by determinism. The main thing we need is that if \mathbf{v} is a deterministic direction, then a half-plane in direction⁹ \mathbf{v} containing only 0s (or other doubly periodic content), can only be continued periodically. If X contains a configuration contradicting this, we say *something appears from nothing*. The *vertical deterministic directions* are $\{r(0, 1), r(0, -1) \mid r \in \mathbb{R}_+\}$.

Proposition 6.9. *If $X \subset \Sigma^{\mathbb{Z}^2}$ is a subshift with a vertical deterministic direction whose trace is a bounded infinite subshift, then it contains a singly periodic point.*

Proof. Suppose X is deterministic upward.

The first two paragraphs of the proof of the previous proposition go through with minor modifications, in the first observing that a one-dimensional subshift with a deterministic direction is an SFT, and in the second picking n to be the radius of the local rule of upward determinism instead of the SFT window. Let Z be as in that proof.

⁸For a more rigorous proof, one needs that the two sides of a curve are really on different sides, so that always we have an internal side to decrease in area (and that we have a notion of area at all). The polygon version of the Jordan curve theorem suffices for this [45, Lemma 1].

⁹Recall that by this we mean that the boundary is perpendicular to \mathbf{v} – we think of the half-plane as moving in direction \mathbf{v} and eating up the configuration.

Let now $x \in X$ be a covering configuration of a highway configuration $z \in Z$. Determinism means that a lower half-plane containing half of the path will have a unique extension to a configuration of X . Since the determinism is given by a local rule, in fact a finite number ℓ of rows already determines the next row upward. Since every row of x is eventually periodic with uniformly bounded preperiod (by the definition of a highway, and the definition of the factor Z), there are, up to shifting, only finitely many ℓ -tuples of consecutive rows. It follows from the pigeonhole principle that the direction of the path is rational, and thus x is singly periodic.

There cannot be any finite configurations in Z because otherwise something appears from nothing in X . Similar reasoning applies to cluster fractals: in the orbit closure of a cluster fractal, for any direction $\mathbf{v} \in \mathbb{R}^2 \setminus \{\mathbf{0}\}$ (in particular for $\mathbf{v} = (0, 1)$) one can find a half-plane in direction \mathbf{v} of all zeroes whose boundary contains a nonzero symbol, and again something appears from nothing in X . \square

When the deterministic direction is not vertical, X need not contain a singly periodic point. See Example 7.6 in Section 7. However, we will show next that in the trace sparse case, any direction of determinism implies the existence of a singly periodic point. A much studied concept in multidimensional dynamics are the directions of expansivity [9], a concept similar to determinism. We say $\mathbf{v} \in \mathbb{R}^2 \setminus \{\mathbf{0}\}$ is an *expansive direction*¹⁰ if, defining $L_{\mathbf{v},r} = \mathbb{R}\mathbf{v} + B_r(\mathbf{0}) \subset \mathbb{R}^2$, we have

$$\exists \epsilon > 0 : \exists r > 0 : (\forall \mathbf{u} \in L_{\mathbf{v},r} \cap \mathbb{Z}^2 : d(\sigma^{\mathbf{u}}(x), \sigma^{\mathbf{u}}(y)) < \epsilon) \implies x = y.$$

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

Lemma 6.10. *Let X be a two-dimensional subshift. Then \mathbf{v} is an expansive direction for X if and only if both of the directions orthogonal to \mathbf{v} 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 $\mathbf{v} \notin D \implies (\mathbf{v}^T \notin E \wedge -\mathbf{v}^T \notin E)$ and $\mathbf{v} \notin E \implies (\mathbf{v}^T \notin D \vee -\mathbf{v}^T \notin D)$, from which the lemma follows easily.

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

In [9], \mathbb{Z}^2 -subshifts with almost any¹¹ set of expansive directions are constructed, the only exception being a single irrational expansive direction, which was left open. Since a rational line in direction \mathbf{v} will have \mathbf{v} as the unique nonexpansive direction, the obvious way to try to solve this is a single irrational line. The next result, Proposition 6.13, shows that such an idea cannot

¹⁰Of 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.

¹¹Up to natural topological restrictions.

work. See [23] for a correct implementation of a unique nonexpansive irrational direction.

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

Lemma 6.11. *Let $X \subset \Sigma^{\mathbb{Z}^2}$ be a subshift. Then the set of nondeterministic directions (scaled onto the unit circle) is closed in the unit circle.*

Lemma 6.12. *If a subshift $X \subset \Sigma^{\mathbb{Z}^2}$ is nontrivial, almost minimal and trace sparse, and has a direction of determinism, then it is the orbit closure of a singly periodic point, and its set of nondeterministic directions is $\{r\mathbf{v}', -r\mathbf{v}' \mid r \in \mathbb{R}_+\}$ for some $\mathbf{v}' \in \mathbb{Z}^2$.*

Proof. Again since X has a deterministic direction, it has a rational deterministic direction \mathbf{v} by the previous lemma. It cannot contain a finite configuration or a cluster fractal, and thus must consist of r -highways for some fixed r . Let H be a half-plane in direction \mathbf{v} whose boundary is at the origin. Now, consider a minimal subsystem Y of $\text{MPC}_r(X)$. To each path $p : \mathbb{Z} \rightarrow \mathbb{Z}^2$ corresponding to a point in Y , associate the sequence of dot products $n \mapsto \langle p(n), \mathbf{v} \rangle$. This sequence measures how far the path is from the boundary of H .

Since the paths are uniformly recurrent, it follows from Theorem 4.2 that either some path enters some strip around the boundary of H infinitely many times, or every path deviates from the boundary with some uniformly lower bounded velocity (in the sense that the sequence of dot products is an ascending path).

First, suppose some path visits some strip around the boundary of H infinitely many times. Then it is easy to find a configuration where something appears from nothing, by looking at the times when the path is maximally far from the boundary in the direction $-\mathbf{v}$. Thus, this case is impossible.

Suppose then that $n \mapsto \langle p(n), \mathbf{v} \rangle$ is a uniformly ascending for every path p corresponding to a point in $\text{MPC}_r(X)$, that is, there exist t, m such that if p is any such path and $p(t)$ is in the inner m -border of H , i.e. $p(t) \in A = (\partial H + B_m(\mathbf{0})) \cap H$ then for all $t'' \notin [t' - t, t' + t]$, $p(t) \notin A$.

Now, observe that the border of any translate of H is periodic, in the sense that for any $\mathbf{u} \in \mathbb{R}^2$ the characteristic function of $(\mathbf{u} + \mathbb{R}\mathbf{v}^T + B_{m/2}(\mathbf{0})) \cap \mathbb{Z}^2$ is in the orbit of one of finitely many singly periodic configurations in $\{0, 1\}^{\mathbb{Z}^2}$, and uniform ascension of paths implies that there are only finitely many distinct configurations that can be seen on the border of H . From determinism in direction \mathbf{v} it follows that every path is periodic, and almost minimality then implies that the subshift is the orbit closure of a singly periodic point. Pick $\mathbf{v}' = \mathbf{v}^T$. \square

By extracting an almost minimal subsystem and applying the previous lemma, we obtain the following.

Proposition 6.13. *If a subshift $X \subset \Sigma^{\mathbb{Z}^2}$ is nontrivial and trace sparse 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 essentially just subshifts with a rational deterministic direction, and thus it is clear that we obtain corollaries for cellular automata.

The main technical tool used in [42] 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 [42], namely the decidability of the nilpotency problem.

More precisely, let $X \subset \Sigma^{\mathbb{Z}}$ be a countable sofic shift. A cellular automaton $f : X \rightarrow X$ 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 eventually 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 6.14. *Let X be a bounded one-dimensional subshift and $f : X \rightarrow X$ 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 horizontal trace of the spacetime subshift of f is infinite. By Proposition 6.9, X contains a singly periodic point. Such a configuration is precisely a spaceship for f . The decidability of nilpotency follows because both “ $\exists k : |f^k(X)| < \infty$ ” and “ x is a spaceship for f ” are semidecidable on sofic shifts, and in the former case we can easily check nilpotency, while the latter implies non-nilpotency. \square

Proposition 6.15. *Let $X \subset \Sigma^{\mathbb{Z}}$ be any one-dimensional subshift, and let $f : X \rightarrow X$ 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 3.5 we can extract an almost minimal trace sparse subshift Z from Y . Then Z is a trace sparse 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 all configurations eventually zero to the right is the same, but using a vertical deterministic direction. \square

We can also obtain the following proposition. We note that this is stronger than the previous proposition in the sense that even on full shifts, there are cellular automata such that the closure of their asymptotic set is a proper subshift of their limit set.

Proposition 6.16. *Let $X \subset \Sigma^{\mathbb{Z}}$ be any subshift, and let $f : X \rightarrow X$ 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 horizontal traces 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. \square

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 [19]).

Theorem 6.17. *Let X be a one-dimensional SFT and $f : X \rightarrow X$ a cellular automaton. Then f is nilpotent if one (equivalently all) of the following holds:*

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

Here, the two-sided vertical trace subshift of f is the vertical trace of the spacetime subshift of f , and the one-sided vertical trace is $\{y \in \Sigma^{\mathbb{N}} \mid \exists x \in X : \forall i : y_i = f^i(x)_0\}$.

6.4 Irrational sparseness

Since Theorem 1.2 applies to any almost minimal 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 1.2 it is enough to assume sparseness in any (possibly irrational) direction.

If $\mathbf{v} \in \mathbb{R}^2 \setminus \{(0, 0)\}$ is a direction and X a subshift, we say $x \in X$ is a *road in direction \mathbf{v}* if the support of x is the range of a path $p : \mathbb{Z} \rightarrow \mathbb{Z}^2$ such that the sequence of dot products $n \mapsto \langle p(n), \mathbf{v} \rangle : \mathbb{Z} \rightarrow \mathbb{R}$ is a uniformly ascending path. It is a *highway in direction \mathbf{v}* if the colored path q defined by $q(i) = (p(i), x_{p(i)})$ is uniformly recurrent.

Proposition 6.18. *Let X be a \mathbb{Z}^2 -subshift and let $L = \mathbb{R}\mathbf{v}$ be a line with arbitrary slope. Suppose that for some $r > 0$, the strip $\mathbf{u} + L + B_r(0)$ in x contains finitely many nonzero symbols for every $\mathbf{u} \in \mathbb{R}^2$ and every $x \in X$. Then X contains a finite point, a cluster fractal or a highway in direction \mathbf{v}^T .*

Proof. Let \mathbf{v} be the unit vector giving the slope of L . Let $R \subset \mathbb{R}^2$ be the closed square having one corner at $(0, 0)$ and with sides \mathbf{v} and \mathbf{v}^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

$$\psi(x)_{\mathbf{u}} = 0 \text{ iff } S \cap (\mathbf{u}_1\mathbf{v} + \mathbf{u}_2\mathbf{v}^T + R) \neq \emptyset.$$

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 trace sparse, and contains a point of one of the three types

in Theorem 5.4. It is easy to check that finite points must come from finite points through a rotation, and highways must come from roads in direction d^T , which we can turn into highways in direction d^T by again passing to an almost minimal subshift. Cluster fractals need not a priori come from cluster fractals (since our definition of cluster fractal implies almost minimality). However, it is again clear that any almost minimal subsystem will be supported by a cluster fractal. \square

6.5 Topological full groups

The *topological full group* of a subshift X is the group G of homeomorphisms $h : X \rightarrow X$ such that for some cocycle $c : G \times X \rightarrow \mathbb{Z}$, $c(h, x) = n \implies h \cdot x = \sigma^n(x)$ and c is continuous in its right argument [17]. Such a c is a cocycle in the sense of cohomology, that is, $c(g \circ h, x) = c(h, x) + c(g, h(x))$.

Lemma 6.19. *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 . For $x \in X$, write \bar{x} for the unique configuration of Y with central column x . Add another layer on top of Y , where in cell $\mathbf{v} \in \mathbb{Z}^2$ of $\bar{x} \in Y$, we write the vector $(1, c(g, \sigma^{\mathbf{v}^2}(x)))$. To see that this gives a subshift, observe that by continuity $c(g, \sigma^{\mathbf{v}^2}(x))$ is determined by $\sigma^{\mathbf{v}^2}(x)|_F$ for a finite $F \subset \mathbb{Z}$, thus by $x|_{\mathbf{v}^2+F}$ in a shift-invariant way, thus $(1, c(g, \sigma^{\mathbf{v}^2}(\bar{x})))$ is determined by the values in $\bar{x}|_{\mathbf{v}+\{0\} \times F}$ in a shift-invariant way (since rows are constant). Call the resulting subshift Z .

(The relation between Z and the path cover $\text{PC}_r(Y)$ (with $r = \max |c(f, x)|$) is that Z is defined similarly as $\text{PC}_r(Y)$ except cycles are not forbidden, but additionally the symbol $\#$ is forbidden and furthermore the path is entirely determined by the local rule of the cocycle c .)

Now, the function $n \mapsto c(g^{kn}, x)$ tells us the vertical movement of a path in a configuration of Z , when we interpret it as a graph $G(Y)$ by following the vectors written on the second layer as in Section 4.2. This correspondence follows from the cocycle formula: If $\bar{x}_{\mathbf{v}} = (1, c(g, \sigma^{\mathbf{v}^2}(x))) = (1, n)$ then $g \cdot \sigma^{\mathbf{v}^2}(x) = \sigma^{\mathbf{v}^2+n}(x)$. The arrow at $\bar{x}_{\mathbf{v}}$ points to the cell $\mathbf{v} + (1, n)$, which corresponds to the central coordinate of $\sigma^{\mathbf{v}^2+n}(x)$.

If a path visits the same row twice, then it is periodic. Because g is of infinite order, we must find paths that do not revisit any cell in arbitrarily many steps, meaning that the corresponding drawn path does not visit the same row twice.

Now, consider a subshift $W \subset Z \times \{0, 1\}^{\mathbb{Z}^2}$ where we allow coloring exactly one of the paths in Z with 1s. Forbid, in W , revisits of the colored path to the same row. By the assumption of the previous paragraph W still contains configurations with a colored path. Now, consider a factor of W where the Z -component is removed. This subshift is trace sparse, 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. \square

Of course the proof adapts easily also to topological full monoids, but we omit the definition. We obtain another proof of a result of [5]. The *torsion*

problem of a recursively presented group G is the problem of, given $g \in G$, determining whether the order of g is finite. The topological full group of a full shift (more generally any Π_1^0 subshift) has an obvious recursive presentation by local rules, see [5].

Theorem 6.20. *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 $n \mapsto c(g^{kn}, x)$ 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. \square

6.6 Patterns in primes

Unlike our main theorem, the one-dimensional version Theorem 5.1 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

$$X_N = \bigcap_m \overline{\{\sigma^n(x_N) \mid n \geq m\}}$$

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 the set \mathcal{P} of prime numbers we obtain what we call the *prime subshift* $X_{\mathcal{P}}$.

Lemma 6.21. *Let p_i be the i th prime number. Then*

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

if and only if the prime subshift is not equal to the sunny-side-up subshift.

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 $n_i \rightarrow \infty$. Let y be a limit point of the sequence $\sigma^{p_{n_i}}(x)$ as $i \rightarrow \infty$. Then $y_{[0,k]} = 10^{k-1}1$ and $y \in X_{\mathcal{P}}$.

Suppose then that $X_{\mathcal{P}} \neq X_{\leq 1}$. 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 $\liminf_{n \rightarrow \infty} p_{n+1} - p_n \leq k$. \square

The fact that indeed $\liminf_{n \rightarrow \infty} p_{n+1} - p_n < \infty$ was shown in 2014 in a revolutionary paper by Zhang [46]. Since then, much more has been learned about the words appearing in $X_{\mathcal{P}}$, and in particular [33] shows that $10^n 1 \sqsubset X_{\mathcal{P}}$ holds for some $n < 600$. 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 (see their paper for the definition) is essentially minimal. Our definition of the prime subshift is slightly different from theirs, so we give the proof.

Lemma 6.22. *The prime subshift is essentially minimal, and its only minimal subshift is $\{0^{\mathbb{Z}}\}$.*

Proof. Let n be arbitrary. Let $I = \{0, 1, \dots, n\}$ and let $\phi : I \rightarrow \mathcal{P}$ 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 $k \equiv -i \pmod{\phi(i)}$ for all $i \in I$. Then $k+i \equiv i-i \equiv 0 \pmod{\phi(i)}$ 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 prime p_i . It follows that 0^n occurs in every subword of $x_{\mathcal{P}}$ of length $\prod_{i=0}^{n-1} p_i + 2n$. \square

For a set $N \subset \mathbb{N}$, write $\pi_N(n) = |N \cap [0, n]|$, and define $\pi = \pi_{\mathcal{P}}$. 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 $X_N = \overline{\mathcal{O}(\dots 000111\dots)} \cup \overline{\mathcal{O}(\dots 111000\dots)}$.

Proposition 6.23. *The prime subshift contains either a finite point or a cluster fractal.*

Proof. By the previous lemma, $X_{\mathcal{P}}$ is essentially minimal with $\{0^{\mathbb{Z}}\}$ the only minimal subsystem. By Lemma 3.7, $X_{\mathcal{P}}$ contains an almost minimal subsystem, and the claim follows from Theorem 5.1. \square

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 cluster fractal is quite difficult from first principles.

We do not know whether the prime subshift contains a cluster fractal. If it did, $X_{\mathcal{P}}$ would be uncountable. We note that at least the conclusion of Theorem 1.2 of [33] 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, in the sense that if you pick for a large enough set of integers, then a randomly chosen m -tuple taken from those integers will with high probability be the support of a legal configuration of X . (The same is true even if one further restricts X to contain only admissible patterns.)

We remark that one *can* show that $X_{\mathcal{P}}$ contains $X_{\leq 1}$ using a more sophisticated but standard number-theoretic tool, namely Dirichlet's theorem on arithmetic progressions.

Proposition 6.24. *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 $\phi : I \rightarrow \mathcal{P}$ 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 $k \equiv i \pmod{\phi(i)}$ for all $i \in I$. Then k and N are coprime, so by Dirichlet's theorem on

arithmetic progressions, there is a prime number $p = k + \ell N$ for some $\ell \in \mathbb{N}$. Then $p - i \equiv i - i \equiv 0 \pmod{\phi(i)}$ for $i \in I$. It follows that the characteristic sequence of the primes contains the word $0^n 1 0^n$ around the prime p . Taking a suitable limit, we obtain that $X_{\leq 1} \subset X_{\mathcal{P}}$. \square

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 1.2, 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 [35] that fixed-point subshifts of many substitutions (in particular rectangular deterministic ones) are sofic, and the result of various authors that traces of vertically constant sofics can be arbitrary computable subshifts [22, 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 traces are constructed in [41, 43], and bounded subshifts can be turned into sparse ones by blocking and factoring (see the proof of Theorem 6.6). The proof of Lemma 6.19 shows that every element of a topological full group of a one-dimensional subshift gives an example of a trace sparse two-dimensional subshift.

7.1 Subshifts

The following is our first example of a trace sparse cluster fractal, and thus shows that in Theorem 1.2, the case of cluster fractals must indeed be included. Less trivially, it also shows that while essential trace sparseness is always 1 in the cases of finite configurations and path configurations, the essential trace sparseness of an almost minimal trace sparse subshift can in general be arbitrarily large and can match the sparseness constant.

Example 7.1: For every k , there exists an almost minimal two-dimensional sofic subshift $X \subset \{0, 1\}^{\mathbb{Z}^2}$ with k -sparse trace such that every nonzero configuration has essential trace sparseness k .

See Figure 2 for an illustration of a typical pattern in the subshift X we construct.

Proof. For $k = 1$, the two-dimensional sunny-side-up subshift has this property.

Let $k \geq 2$. 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 as the patterns $P_{1,j}$ the following matrices (using $P_{0,j'} = 1$ for all j'), and observe that the assumptions are then satisfied for $i = 1$. 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 j th row from the bottom. It is easy to see that the induction hypothesis is satisfied. Observe that this is simply a $(k+1)$ -by- $(k+1)$ substitution.

Define X as the limit of the patterns $P_{i,1}$. More precisely, take the subshift containing those configurations x such that for all finite $D \subset \mathbb{Z}^2$, $x|_D$ is a subpattern of $P_{i,1}$ for some i . Since every nonzero pattern occurring in X occurs in one of the patterns $P_{i,1}$ by definition, we can conclude that its trace is k -sparse (by the inductive assumption).

Now we prove¹² almost minimality. Obviously the all-zero point is in the subshift. We next note that every m_i -by- m_i block (not necessarily $(0 \bmod m_i)$ -aligned) that appears in some $P_{i',1}$ is actually contained already in $P_{i+2,1}$. This follows from the fact that all $(0 \bmod m_i)$ -aligned $2m_i$ -by- $2m_i$ blocks in $P_{i+3,1}$ already appear in $P_{i+2,1}$, which is clear from drawing the support of $P_{i+3,1}$. If a $(0 \bmod m_i)$ -aligned $2m_i$ -by- $2m_i$ subpattern Q in X has a nonzero cell, then since every configuration consists of a $(0 \bmod m_i)$ -aligned grid of empty blocks and $P_{i,j}$ -blocks for varying j , Q contains some pattern $P_{i,j'}$ entirely. It follows that it contains every pattern $P_{i-1,j''}$ for $j'' \in \{1, 2, \dots, k\}$, thus every m_{i-3} -by- m_{i-3} block that appears in X by the previous argument.

The previous argument also shows that, for every i , every configuration contains a row with k ones with pairwise distances at least m_{i-1} , at a bounded distance from every nonzero symbol. Since $m_{i-1} \rightarrow \infty$, the subshift has essential trace sparseness k .

Finally, we observe that X is sofic by Mozes' theorem [35], since it is defined by a deterministic primitive substitution. \square

There are certainly sparse subshifts that are not almost minimal. Thus the assumption of almost minimality is necessary in Theorem 5.4. The assumption of trace sparseness is also necessary:

Example 7.2: There is an almost minimal binary sofic \mathbb{Z}^2 -subshift where the connected component of every 1 is infinite, but which does not contain a highway in any direction.

Proof. Use the substitution $\blacksquare \mapsto \blacksquare \oplus$. See Figure 3. Almost minimality is shown as in the previous example. To see the other properties, observe there

¹²This mimics the standard argument for the minimality of a primitive (uniform) one-dimensional substitution [30, Proposition 4.15] and only needs the ‘‘almost primitivity’’ of the substitution. The number 3 in ‘‘ m_{i-3} -by- m_{i-3} ’’ in the proof plays the role of the usual ‘‘delay until small patterns that eventually appear start appearing’’.




Figure 2: The pattern $P_{5,1}$ for $k = 2$ in Example 7.1 (top third cut off).

is an exponential, thus superlinear, lower bound on the number of black cells connected to it in any ball centered at a black cell. \square




Figure 3: A pattern from the almost minimal subshift of Example 7.2.

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 7.3: 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 $k \rightarrow \infty$ we obtain a subshift where every point containing a nonzero symbol contains nonzero symbols with upper Banach density at least $1/4$. \circ

An SFT containing a road must contain a periodic road by the pigeonhole principle. We now show that sofic shifts can consist of computationally complex ascending paths.

Proposition 7.4. *Let X be a two-dimensional Π_1^0 -subshift where every nonzero configuration is an r -road with uniform ascension constant. 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 up to translation the range of an r -path with ascension constant m (so starting from \mathbf{v} , the path stays above \mathbf{v} after m steps). Let $n = rm$.

Let Y be a one-dimensional Π_1^0 -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 \mathbb{Z}^2 -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 and its successor contain $\#$, and that every row has a predecessor. Thus, the $[1, n]$ -components describe a vertical (strictly) 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 $X_{\leq 1}$ 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 \mathbf{u} at \mathbf{v} , then $\mathbf{v} + \mathbf{u}$ is called the *successor* of the cell \mathbf{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 \mathbf{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 \mathbf{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 subshift over the alphabet $(\{0\} \cup \Sigma_+)$ 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 by [22, 13, 2], that checking that every row contains at most one binary symbol is doable since $X_{\leq 1}$ 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 . \square

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 7.5: There exists a two-dimensional sofic shift X whose configurations are r -roads with uniform ascension rate, but are not periodic and their support does not fit any strip of the form $L + B_r(0)$ where L is a straight line. \circ

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 $L + B_r(0)$, but no such rational line. With a similar idea, we also obtain an example of a sofic shift with a bounded trace, and with only four nondeterministic directions, two of which are irrational.

Example 7.6: Let X be a Sturmian subshift (with any computable irrational angle). Let Y be the two-dimensional 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 7.4: allow at most one 1 on each row, and if $(y,z) \in Z$ and $z_{\mathbf{v}} = 1$, require $z_{\mathbf{v}+(y_{\mathbf{v}},1)} = 1$ and that either $z_{\mathbf{v}+(0,-1)} = 1$ or $z_{\mathbf{v}+(-1,0)} = 1$. See Figure 4.

Now, we claim that the nondeterministic directions are precisely the vertical ones and the ones parallel to the irrational slope.¹³ First, both of them are easily seen to be nondeterministic, since in the vertical directions it is impossible to know the continuation of the Y -component, and in a direction orthogonal to the path (as it appears in configurations), 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 $L + B_r(0)$ where L is a straight line visits every half-plane not parallel to L . \circ

If we project away 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. If the guiding rows are removed, Z becomes sparse, so if it had a deterministic direction then it would have a singly periodic point by Proposition 6.13, so in fact *every* direction becomes nondeterministic and nonexpansive when the guiding rows are removed. Subshifts with a single irrational nonexpansive direction are not easy to construct [23], and making such sofic shifts and SFTs is another challenge still [47].

¹³We refer to the slope as it appears graphically. This is also the slope of the unique path drawn in Z if angle is interpreted suitably: if X corresponds to a *mechanical word* (see [32]) of slope $\alpha \in (0,1)$, then the path will always stay at a bounded distance from some translate of the line $\mathbb{R}(\alpha,1)$.

ascending path (though of course they are not all r -paths for any fixed r): pick a slowly-growing recursive function $g : \mathbb{N} \rightarrow \mathbb{N}$ with $g(n) \rightarrow \infty$ and additionally require that if there are two 1s at distance at most $g(n)$ from each other, then the two-columns property of the subshift holds with respect to some $m \leq n$.

7.2 Path covers

We give an example of extraction of highways through path covers.

Example 7.9: Fix some digital approximation scheme that produces subsets of \mathbb{Z}^2 as approximations to real Euclidean circles with integer centers and integer radii. We choose for each circle the approximation where in the octant between angles $\pi/2$ and $\pi/4$ (in the standard orientation of the complex plane \mathbb{C} or the unit circle), for each horizontal coordinate m we include the point (m, n) where n is chosen so that the distance to the corresponding octant of the real circle is minimal along the vertical line through $(m, n) \in \mathbb{Z}^2$. Since

$$\sqrt{r^2 - m^2} = n + \frac{1}{2} \implies r^2 - m^2 = n^2 + n + \frac{1}{4},$$

the real circle cannot go through $(m, n + 1/2)$, and thus the approximation is unique. The rest of the points are produced by dihedral symmetry. This is the scheme produced by the Bresenham circle-drawing algorithm and the midpoint circle algorithm [20].

Say $x \in \{0, 1\}^{\mathbb{Z}^2}$ is a *circles configuration* if there exists a finite set of disjoint circles with centers in \mathbb{Z}^2 and integral radii, such that $x_{\mathbf{v}} = 1$ if and only if the digital approximation for one of the circles contains \mathbf{v} , and all the digital approximations are 1-clusters in x . Let X be the smallest subshift containing every circles configuration, call this the *circles subshift*. Since discretizations of circles can have arbitrarily large support in configurations of X , Lemma 4.3 shows that $\text{MPC}_1(C)$ is nonempty.

The proof of Lemma 4.3 in the case of X is carried out in Figure 5. Say the *parameters* of a circle are the vector $(a, b, c) \in \mathbb{Z}^2 \times \mathbb{N}$, where (a, b) is the center and c the radius. The subfigures show a series of pairs of circles, one getting larger and surrounding the origin. The circle getting larger always contains the coordinate $(3, 3)$ and its parameters are given in the captions. The smaller circle has parameters $(-7, -6, 3)$ in each configuration. We trace finite circular paths partially around the big circle to obtain configurations of $\text{PC}_1(X)$. These paths get arbitrarily long for the growing circle, thus in the suitable limit (which looks like the last picture) one can trace a bi-infinite path. Shifting it by $(3, 3)$, we obtain a configuration $x \in \text{MPC}_1(X)$.

Observe that x is not a highway, as we obtained it as a limit of configurations containing a “bump”, which appears only once, see [34, 29] and OEIS sequence A055979 for details. Note that x is, however, a road by Lemma 4.6, as we can trace the path through the finite circle as well. The path proving 7-roadness is shown in Figure 6. (Since X is zero-gluing, we could also simply remove the circle to get a 1-road.) By passing to a minimal subsystem of $\text{MPC}_7(X)$, we obtain a highway, a diagonal path through the origin. By positioning the circle differently, one can obtain also irrational lines, as constructed in Example 7.6.

○

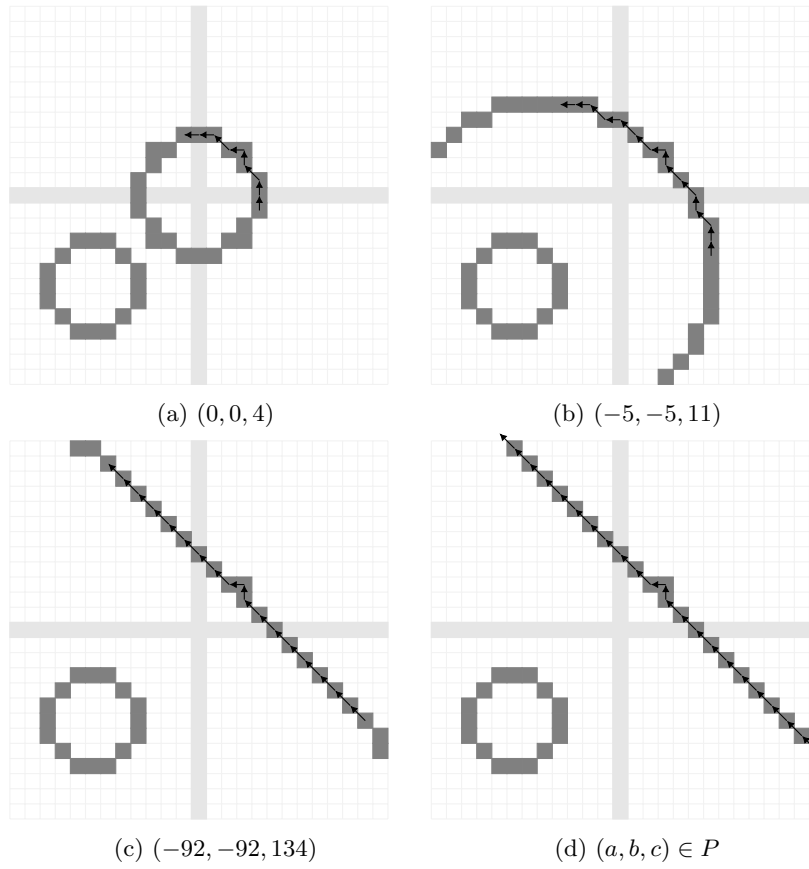


Figure 5: Illustration of how to obtain nonemptiness of $\text{MPC}_1(X)$ for the circles subshift. In the last subfigure $P = \{(a, a, r) \mid a = -\lfloor r/\sqrt{2} \rfloor + 2, \exists n \in \mathbb{N} : r^2 = 2n^2 - n + 1, r \geq 373\}$ [29].




Figure 6: The path illustrating 7-roadness of x , and a point of $\text{MPC}_7(X)$ (even $\text{MPC}_1(X)$) in its orbit closure under τ .

7.3 Paths

In Theorem 4.1, 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, where 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 trace sparse spaces. However, this case is quite interesting, and we show by examples that it splits into three subtypes.¹⁵ In the rest of this section, we give representatives of each of these subtypes. All of the examples are substitutive, and thus also give examples of trace sparse 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

- every path visits every cell (that it visits at all) infinitely many times (Example 7.10, Figure 7),
- some path visits all cells finitely many times (Example 7.11, Figure 8), or
- some path visits some cell finitely many times, but every path visits some cell infinitely many times (Example 7.12, Figure 9).

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

Example 7.10: There is a minimal subsystem of \mathbf{Path} where every path visits every cell 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 a two-sided fixed-point of this substitution, considered as a minimal subshift of \mathbf{DPath} . Then $X = \phi^{-1}(Y) \subset \mathbf{Path}$ 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. \circ

¹⁴Such paths are known as *coboundaries*.

¹⁵Similar 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 7: Part of a typical path in Example 7.10.



Figure 8: Part of a typical path in Example 7.11.

Example 7.11: 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. From the fact that the total offset $f(n) = 3^n$ of the path $\tau_2^n(\nearrow)$ is unbounded but $f(n)/5^n \rightarrow 0$ and the primitiveness of the substitution, 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 $\tau_2^m(\nearrow) \cdot \tau_2^m(\searrow)$ as $m \rightarrow \infty$ is such a path, where the decimal point denotes the center of the path. \circ

Example 7.12: Consider the substitution τ_3 defined by

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

Define X as in the previous examples. As in the above example, it is clear that paths in X are not uniformly ascending or bounded. As in Example 7.10, consider the configurations $z \in \mathbb{N}^{\mathbb{Z}}$ corresponding to $\tau^n(\nearrow)$. Then by induction one can show that the support of z is the interval $[0, n+1]$, $z_0 = 1$ and $z_i \geq n$ for $i \in [1, n+1]$. (Seeing z as an element of $\mathbb{Z}^{\mathbb{N}}$, the recurrence is $z \mapsto z + 2 \cdot 0z0^\omega - 0110^\omega$, with cellwise addition.)

It easily follows that all paths in X have at most one cell that is not visited infinitely many times, as if m is visited at most k times, then for $n > k$, the property of $z \in \mathbb{N}^{\mathbb{Z}}$ corresponding to $\tau^n(\nearrow)$ or $\tau^n(\searrow)$ established in the previous paragraph shows that m must be the beginning of a $\tau^n(\nearrow)$ -path and the end of a $\tau^n(\searrow)$ -path (as every path can be written as compositions of such subpaths). Conversely, indeed 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 $\tau^n(\searrow) \cdot \tau^n(\nearrow)$. \circ




Figure 9: Part of the path in Example 7.12 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.

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. We thank an anonymous referee for suggesting many improvements.

References

- [1] Stål O. Aanderaa and Harry R. Lewis. Linear sampling and the $\forall\exists\forall$ 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. \mathcal{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] D. Cenzer and J. B. Remmel. Π_1^0 classes in mathematics. In *Handbook of recursive mathematics, Vol. 2*, volume 139 of *Stud. Logic Found. Math.*, pages 623–821. North-Holland, Amsterdam, 1998.
- [12] Alexandre I. Danilenko. Strong orbit equivalence of locally compact Cantor minimal systems. *International Journal of Mathematics*, 12(01):113–123, 2001.
- [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] Dora Giammarresi and Antonio Restivo. Handbook of formal languages, vol. 3. chapter Two-dimensional languages, pages 215–267. Springer-Verlag New York, Inc., New York, NY, USA, 1997.
- [17] Thierry Giordano, Ian F. Putnam, and Christian F. Skau. Full groups of cantor minimal systems. *Israel Journal of Mathematics*, 111(1):285–320, 1999.
- [18] Pierre Guillon and Emmanuel Jeandel. Infinite Communication Complexity. *ArXiv e-prints*, January 2015.
- [19] Pierre Guillon and Gaétan Richard. Asymptotic behavior of dynamical systems and cellular automata. *ArXiv e-prints*, April 2010.
- [20] Donald D. Hearn, M. Pauline Baker, and Warren Carithers. *Computer Graphics with Open GL*. Prentice Hall Press, Upper Saddle River, NJ, USA, 4th edition, 2010.
- [21] 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.
- [22] Michael Hochman. On the dynamics and recursive properties of multidimensional symbolic systems. *Invent. Math.*, 176(1):131–167, 2009.
- [23] Michael Hochman. Non-expansive directions for \mathbb{Z}^2 actions. *Ergodic Theory and Dynamical Systems*, 31:91–112, 1 2011.
- [24] 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.
- [25] Jarkko Kari. The nilpotency problem of one-dimensional cellular automata. *SIAM J. Comput.*, 21(3):571–586, 1992.
- [26] Jarkko Kari. Theory of cellular automata: a survey. *Theoret. Comput. Sci.*, 334(1-3):3–33, 2005.
- [27] 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.
- [28] Jarkko Kari. Decidability and undecidability in cellular automata. *Int. J. General Systems*, 41(6):539–554, 2012.

- [29] Zenon Kulpa. On the properties of discrete circles, rings, and disks. *Computer Graphics and Image Processing*, 10(4):348 – 365, 1979.
- [30] Petr Kůrka. *Topological and symbolic dynamics*, volume 11 of *Cours Spécialisés [Specialized Courses]*. Société Mathématique de France, Paris, 2003.
- [31] Douglas Lind and Brian Marcus. *An introduction to symbolic dynamics and coding*. Cambridge University Press, Cambridge, 1995.
- [32] M. Lothaire. *Algebraic combinatorics on words*, volume 90 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, 2002.
- [33] James Maynard. Small gaps between primes. *Annals of Mathematics*, 181(1):383–413, 2015.
- [34] Thierry Monteil. Another definition for digital tangents. In Isabelle Debled-Rennesson, Eric Domenjoud, Bertrand Kerautret, and Philippe Even, editors, *Discrete Geometry for Computer Imagery*, pages 95–103, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.
- [35] Shahar Mozes. Tilings, substitution systems and dynamical systems generated by them. *J. Analyse Math.*, 53:139–186, 1989.
- [36] Nic Ormes and Ronnie Pavlov. Extender sets and multidimensional subshifts. *Ergodic Theory and Dynamical Systems*, 36:908–923, 5 2016.
- [37] 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.
- [38] Raphael M. Robinson. Undecidability and nonperiodicity for tilings of the plane. *Invent. Math.*, 12:177–209, 1971.
- [39] Ville Salo. Decidability and Universality of Quasiminimal Subshifts. *ArXiv e-prints*, November 2014.
- [40] Ville Salo. *Subshifts with Simple Cellular Automata*. PhD thesis, University of Turku, 2014.
- [41] Ville Salo and Ilkka Törmä. Computational aspects of cellular automata on countable sofic shifts. *Mathematical Foundations of Computer Science 2012*, pages 777–788, 2012.
- [42] Ville Salo and Ilkka Törmä. Computational aspects of cellular automata on countable sofic shifts. *Mathematical Foundations of Computer Science 2012*, pages 777–788, 2012.
- [43] Ville Salo and Ilkka Törmä. Constructions with countable subshifts of finite type. *Fundam. Inf.*, 126(2-3):263–300, April 2013.
- [44] Klaus Schmidt. *Dynamical systems of algebraic origin*, volume 128 of *Progress in Mathematics*. Birkhäuser Verlag, Basel, 1995.

- [45] Helge Tverberg. A proof of the jordan curve theorem. *Bulletin of the London Mathematical Society*, 12(1):34–38, 1980.
- [46] Yitang Zhang. Bounded gaps between primes. *Annals of Mathematics*, 179(3):1121–1174, 2014.
- [47] C. Zinoviadis. Hierarchy and Expansiveness in Two-Dimensional Subshifts of Finite Type. *ArXiv e-prints*, March 2016.