EXTENDER SETS AND MULTIDIMENSIONAL SUBSHIFTS

NIC ORMES AND RONNIE PAVLOV

Abstract. In this paper, we consider extender sets, first defined in [4], which are a $\mathbb{Z}^d$ extension of follower sets from one-dimensional symbolic dynamics. As our main result, we show that for any $d \geq 1$ and any $\mathbb{Z}^d$ subshift X, if there exists n so that the number of extender sets of words on a d-dimensional hypercube of side length n is less than or equal to n, then X is sofic. We also give an example of a non-sofic system for which this number of extender sets is $n+1$ for every n.

We prove this theorem in two parts. First we show that if the number of extender sets of words on a d-dimensional hypercube of side length n is less than or equal to n for some n, then there is a uniform bound on the number of extender sets for words on any sufficiently large rectangular prisms; to our knowledge, this result is new even for $d = 1$. We then show that such a uniform bound implies soficity, which extends a well-known result in $d = 1$.

1. Introduction

For any $\mathbb{Z}$ subshift X (i.e. a closed shift-invariant subset of $A^{\mathbb{Z}}$ for a finite set A) and finite word w appearing in some point of X, the **follower set** of w, written $F_X(w)$, is defined as the set of all one-sided infinite sequences s such that the infinite word ws occurs in some point of X. (In some sources, the follower set is defined as the set of all finite words which can legally follow w, but the former definition may be obtained by taking limits of the latter.) It is well-known that for a $\mathbb{Z}$ subshift X, finiteness of $\{F_X(w) \ : \ w \text{ in the language of } X\}$ is equivalent to X being sofic, i.e. the image of a shift of finite type under a continuous shift-commuting map. (For instance, see [5].)

In [4], extender sets were defined and introduced as a natural extension of follower sets to $\mathbb{Z}^d$ subshifts with $d > 1$. The **extender set** of any finite word w in the language of X with shape $S \subset \mathbb{Z}^d$, written $E_X(w)$, is the set of all configurations on $\mathbb{Z}^d \setminus S$ which, when concatenated with w, form a legal point of X. We can no longer speak of a subshift having only finitely many extender sets, since extender sets of patterns with different shapes cannot be compared as in the one-dimensional case. One way to deal with this is examine the growth rate of the number of distinct extender sets for words in X with a given shape S (which we denote by $N_S(X)$), as the size of S approaches infinity. This works nicely in the one-dimensional case; our Lemma 3.4 (which is routine) shows that soficity of a one-dimensional subshift is equivalent to boundedness of the number of extender sets of n-letter words as $n \to \infty$. Interestingly, this sequence need not stabilize; Example 3.5, due to Martin Delacourt ([1]), demonstrates a $\mathbb{Z}$ sofic shift X where $N_{[1,n]}(X)$ oscillates between

2010 *Mathematics Subject Classification.* Primary: 37B50; Secondary: 37B10, 37A15.
Key words and phrases. $\mathbb{Z}^d$; shift of finite type; sofic; multidimensional.

two values as n increases (In this paper, $[a, b]$ for $a, b \in \mathbb{Z}$ will always represent the set $\{a, a+1, \ldots, b\}$).

There are many relations between properties of X and the behavior of $N_S(X)$. For instance, it is easy to see that when X is a nearest-neighbor shift of finite type, the extender set of a word with shape $[1, n]^d$ is determined entirely by the letters on the boundary. This implies that for such X, $N_{[1,n]^d}(X)$ is bounded from above by $|A_X|^{\partial[1,n]^d} \leq |A_X|^{2dn^{d-1}}$, where A_X denotes the alphabet of X. It was conjectured in [4] that X sofic implies that $\frac{\log N_{[1,n]^d}(X)}{n^d} \to 0$, but this remains open.

A partial answer was proven in [4], using an argument basically present in [7]. A finite sequence S_n of sets, $1 \leq n \leq N$, was defined to be a union increasing chain if $S_n \nsubseteq \bigcup_{i=1}^{n-1} S_i$ for all $1 \leq n \leq N$. Theorem 2.3 of [4] states that if there exist union increasing chains of size $e^{\omega(n^{d-1})}$ of extender sets of words with shape $[1, n]^d$, then X is not sofic. These results can, broadly speaking, be thought of as showing that a very fast growth rate for extender sets implies that a subshift is not an SFT or sofic. Our main result is in the opposite direction, namely it demonstrates that a slow enough growth rate implies soficity.

Theorem 1.1. *For any $d \geq 1$ and any $\mathbb{Z}^d$ subshift X, if there exists n so that $N_{[1,n]^d}(X) \leq n$, then X is sofic.*

The proof of Theorem 1.1 is broken into two mostly disjoint parts:

Theorem 1.2. *For any $d \geq 1$ and any $\mathbb{Z}^d$ subshift X, if there exists n so that $N_{[1,n]^d}(X) \leq n$, then there exist K, N such that for any rectangular prism R with dimensions at least K, $N_R(X) \leq N$.*

Theorem 1.3. *For any $d \geq 1$ and any $\mathbb{Z}^d$ subshift X, if there exist K, N so that $N_R(X) \leq N$ for all rectangular prisms R with all dimensions at least K, then X is sofic.*

Theorem 1.3 can be thought of as a generalization of the previously mentioned fact that one-dimensional shifts with only finitely many follower sets are sofic. We also show that the upper bound in Theorem 1.1 cannot be improved.

Theorem 1.4. *For any $d \geq 1$, there exists a nonsofic $\mathbb{Z}^d$ subshift X for which $N_{[1,n]^d}(X) = n + 1$ for all n.*

Our results have similarities to the famous Morse-Hedlund theorem.

Theorem 1.5. ([6]) *If X is a $\mathbb{Z}$ subshift and there exists an n such that the number of words of length n is less than or equal to n, then X consists entirely of periodic points. Equivalently, there is a uniform upper bound on the number of words of length n.*

It is well-known that the bound in Theorem 1.5 is also tight. Sturmian subshifts have no periodic points and have so-called minimal complexity, i.e. any Sturmian subshift has $n + 1$ words of length n for all n. (For an introduction to Sturmian subshifts, see [2].)

There are similarities between Theorems 1.1 and 1.5; in fact Theorem 1.5 is used in our proof of Theorem 1.2. However, there are also some interesting differences. In the usual proof of Theorem 1.5, a key component is that the number of n-letter words is nondecreasing in n. However, Example 3.5 shows that $N_{[1,n]}(X)$ is not necessarily nondecreasing.

2. Definitions and preliminaries

Let A denote a finite set, which we will refer to as our alphabet.

Definition 2.1. A **pattern** over A is a member of A^S for some $S \subset \mathbb{Z}^d$, which is said to have **shape** S. For $d = 1$, patterns are generally called **words**, especially in the case where S is an interval.

For any patterns $v \in A^S$ and $w \in A^T$ with $S \cap T = \varnothing$, we define the concatenation vw to be the pattern in $A^{S \cup T}$ defined by $(vw)|_S = v$ and $(vw)|_T = w$.

Definition 2.2. For any finite alphabet A, the $\mathbb{Z}^d$**-shift action** on $A^{\mathbb{Z}^d}$, denoted by $\{\sigma_t\}_{t \in \mathbb{Z}^d}$, is defined by $(\sigma_t x)(s) = x(s+t)$ for $s, t \in \mathbb{Z}^d$.

We always think of $A^{\mathbb{Z}^d}$ as being endowed with the product discrete topology, with respect to which it is obviously compact.

Definition 2.3. A $\mathbb{Z}^d$ **subshift** is a closed subset of $A^{\mathbb{Z}^d}$ which is invariant under the $\mathbb{Z}^d$-shift action.

Definition 2.4. The **language** of a $\mathbb{Z}^d$ subshift X, denoted by $L(X)$, is the set of all patterns which appear in points of X. For any finite $S \subset \mathbb{Z}^d$, $L_S(X) := L(X) \cap A^S$, the set of patterns in the language of X with shape S.

Any subshift inherits a topology from $A^{\mathbb{Z}^d}$, and is compact. Each σ_t is a homeomorphism on any $\mathbb{Z}^d$ subshift, and so any $\mathbb{Z}^d$ subshift, when paired with the $\mathbb{Z}^d$-shift action, is a topological dynamical system. Any $\mathbb{Z}^d$ subshift can also be defined in terms of disallowed patterns: for any set $\mathcal{F}$ of patterns over A, one can define the set $X(\mathcal{F}) := \{x \in A^{\mathbb{Z}^d} \ : \ x|_S \notin \mathcal{F} \text{ for all finite } S \subset \mathbb{Z}^d\}$. It is well known that any $X(\mathcal{F})$ is a $\mathbb{Z}^d$ subshift, and all $\mathbb{Z}^d$ subshifts are representable in this way. All $\mathbb{Z}^d$ subshifts are assumed to be nonempty in this paper.

Definition 2.5. A $\mathbb{Z}^d$ **shift of finite type (SFT)** is a $\mathbb{Z}^d$ subshift equal to $X(\mathcal{F})$ for some finite $\mathcal{F}$. If $\mathcal{F}$ consists only of patterns consisting of pairs of adjacent letters, then $X(\mathcal{F})$ is called **nearest-neighbor**.

Definition 2.6. A (topological) **factor map** is any continuous shift-commuting map ϕ from a $\mathbb{Z}^d$ subshift X onto a $\mathbb{Z}^d$ subshift Y. A factor map ϕ is **1-block** if $(\phi x)(v)$ depends only on $x(v)$ for $v \in \mathbb{Z}^d$.

Definition 2.7. A $\mathbb{Z}^d$ **sofic shift** is the image of a $\mathbb{Z}^d$ SFT under a factor map. It is well-known that for any $\mathbb{Z}^d$ sofic shift Y, there exists a nearest-neighbor $\mathbb{Z}^d$ SFT X and 1-block factor map ϕ so that $Y = \phi(X)$.

For $d = 1$, any $\mathbb{Z}$ sofic shift can also be defined using graphs; a $\mathbb{Z}$ subshift is sofic if and only if it is the set of labels of bi-infinite paths for some (edge-)labeled graph $\mathcal{G}$ (see [5] for a proof.)

Definition 2.8. For any $\mathbb{Z}^d$ subshift X and rectangular prism $R = \prod_{i=1}^d [0, n_i - 1]$, the R**-higher power shift** of X, denoted X^R, is a $\mathbb{Z}^d$ subshift with alphabet $L_R(X)$ defined by the following rule: $x \in (L_R(X))^{\mathbb{Z}^d} \in X^R$ if and only if the point y defined by concatenating the $x(v)$, viewed themselves as patterns with shape X, is in X. Formally,

$$\forall v \in \mathbb{Z}^d, y(v) := \left(x(\lfloor v_1 n_1^{-1} \rfloor, \ldots, \lfloor v_d n_d^{-1} \rfloor)\right)(v_1 \pmod{n_1}, \ldots, v_d \pmod{n_d}).$$

Definition 2.9. For any $\mathbb{Z}$-subshift X and word $w \in L_{[1,n]}(X)$, the **follower set** of w is $F_X(w) = \{x \in A^{\{n+1,n+2,\dots\}} \ : \ wx \in L(X)\}$. For any n, we use $M_{[1,n]}(X)$ to denote $|\{F_X(w) \ : \ w \in L_{[1,n]}(X)\}|$, the number of distinct follower sets of words of length n.

Definition 2.10. For any $\mathbb{Z}^d$-subshift X and pattern $w \in L_S(X)$, the **extender set** of w is $E_X(w) = \{x \in A^{\mathbb{Z}^d \setminus S} \ : \ wx \in X\}$. For any S, we use $N_S(X)$ to denote $|\{E_X(w) \ : \ w \in L_S(X)\}|$, the number of distinct extender sets of patterns with shape S.

3. Proofs

For the proof of Theorem 1.2 we need the following finite version of the Morse-Hedlund theorem. We include a proof for completeness, though it is essentially the same proof as that of the original theorem.

Lemma 3.1. *For any word $w \in A^N$ and $n \leq \frac{N}{4}$ so that the number of n-letter subwords of w is less than or equal to n, we can write $w = tuv$ where $|t| = |v| = n$ and u is periodic with some period less than or equal to n.*

Proof. Since there are less than or equal to n subwords of w of length n and there are $N-n+1 > n$ values of i for which $w(i)w(i+1)\dots w(i+n-1)$ is a subword of w, there exists an n-letter subword of w which appears twice. In fact, by the pigeonhole principle we may fix indices $i < k \in [1, n+1]$ such that $w(i)w(i+1)\dots w(i+n-1) = w(k)w(k+1)\dots w(k+n-1)$. Similarly, we may fix indices $\ell < j \in [N-n, N]$ such that $w(\ell-n+1)\dots w(\ell-1)w(\ell) = w(j-n+1)\dots w(j-1)w(j)$. Set $w' = w(i)w(i+1)\dots w(j-1)w(j)$.

It suffices to show that w' is periodic of period less than or equal to n; if this is true, then taking $t = w(1)\dots w(n)$, $u = w(n+1)\dots w(N-n)$, and $v = w(N-n+1)\dots w(N)$ completes the proof since u is a subword of w'.

Let us now consider the number of m-letter subwords of w' for values of $m \in [1, n]$. If the number of one-letter subwords of w' is equal to 1, then w' is of the form $ss\dots s$ for some symbol s and we are done. If not, then the number of one-letter subwords of w' is greater than 1, whereas the number of n-letter subwords of w' is less than or equal to n. Therefore, there must be an $m \in [1, n-1]$ for which the number of $(m+1)$-letter subwords of w' is less than or equal to the number of m-letter subwords of w'. Fix m to be this number for the remainder of the proof.

We now claim that for every m-letter subword t of w', there exists $a \in A$ so that ta is a subword of w' as well. For any choice of t aside from the m-letter suffix of w', this is obvious. But it is true for the suffix as well, since by construction of w', if t is a suffix of w' then t is also the suffix of $w(\ell-n+1)\dots w(\ell-1)w(\ell)$ which means $tw(\ell+1)$ is a subword of w'. A similar argument shows that for every m-letter subword t of w', there exists a $b \in A$ so that bt is a subword of w' as well.

Note that because the number of m-letter subwords is less than or equal to the number of $(m+1)$-letter subwords of w', the a and b constructed in the previous paragraph are always unique.

Let $p = k - i$, and note that $w(i)w(i+1)\dots w(i+m-1) = w(i+p)w(i+1+p)\dots w(i+m-1+p)$. Since there is a unique a which extends the word $w(i)w(i+1)\dots w(i+m-1)$ as a subword of w', we get that $w(i+1)w(i+2)\dots w(i+m) = w(i+1+p)w(i+2+p)\dots w(i+m+p)$. Using the same argument and working inductively, we see that $w(i+r)w(i+r)\dots w(i+r) = w(i+r+p)w(i+r+p)\dots w(i+r+p)$ for any $0 \le r \le j-i-p$. In other words, w' is periodic with period $p \le n$. □

We remark that since the word u in the previous lemma is periodic with period less than or equal to n, this clearly implies that u is periodic with period $n!$ (though this may be a meaningless statement if $|u| \le n!$)

Proof of Theorem 1.2. Consider a $\mathbb{Z}^d$ subshift X and n so that $|N_{[1,n]^d}(X)| \le n$. Define an equivalence relation on $L_{[1,n]^d}(X)$ by $w \sim w'$ iff $E_X(w) = E_X(w')$. For each of the $k \le n$ equivalence classes, choose a lexicographically maximal element, and denote the collection of these words by M. Then for every $w \in L_{[1,n]^d}(X)$, there exists $w' \in M$ so that $E_X(w) = E_X(w')$. Equivalently, in any $x \in X$ containing w, w can be replaced by w' to make a new point $x' \in X$.

Now, consider any rectangular prism $R = \prod_{i=1}^d [1, n_i]$ with $n_i > \max 4n, 2n + n!$ for all i, and any finite word $v \in L_R(X)$. Iterate the following procedure: if v contains a subword with shape $[1,n]^d$ which is not in M, then replace it by the element of M in its equivalence class. Since each of these replacements increases the entire word on R in the lexicographic ordering, the procedure will eventually terminate, yielding a word v' in which every subword with shape $[1,n]^d$ is in M. (These replacements could possibly be done in many different ways or orders; simply choose a particular one and call the result v'.) In particular, v' contains less than or equal to n distinct subwords with shape $[1,n]^d$. Since v' was obtained from v by a sequence of replacements with identical extender sets, $E_X(v) = E_X(v')$.

We wish to bound the number of such possible v' for a given R. For any translate of the $(d-1)$-dimensional hypercube $t + [1,n]^{d-1} \subset \prod_{i=2}^d [1, n_i]$, consider the subpattern $v'|_{[1,n_1]\times(t+[1,n]^{d-1})}$. This can be viewed as an n_1-letter word in the x_1-direction, where each "letter" is a cross-section with shape $t + [1,n]^{d-1}$. When viewed in this way, each n-letter subword of $v'|_{[1,n_1]\times(t+[1,n]^{d-1})}$ is a subpattern of v' with shape $[1,n]^d$, and there are less than or equal to n such subpatterns. Therefore, by Lemma 3.1, $v'|_{[n+1,n_1-n]\times(t+[1,n]^{d-1})}$ is periodic with period $n!e_1$. Since $t + [1,n]^{d-1}$ was arbitrary, in fact $v'|_{[n+1,n_1-n]\times\prod_{i=2}^d [1,n_i]}$ is periodic with period $n!e_1$ as well. In other words, if t and $t + n!e_1$ both have first coordinate between $n+1$ and $n_1 - n$ inclusive, then $v'(t) = v'(t + n!e_1)$. A similar proof shows that if t and $t + n!e_i$ both have ith coordinate between $n+1$ and $n_i - n$ inclusive, then $v'(t) = v'(t + n!e_i)$.

The above shows that except for sites within n of the boundary of R, v' is determined by the subpattern occuring within a d-dimensional hypercube of side length $n!$. More specifically, the values of v' on the sites in $\prod_{i=1}^d ([1,n] \cup [n_i - n + 1, n_i] \cup [n+1, n+n!])$ uniquely determine v', and there are $(2n+n!)^d$ such sites. So, regardless of our choice for R, there are less than or equal to $|A_X|^{(2n+n!)^d}$ possible v'. Since $E_X(v) = E_X(v')$ for every v, this shows that $|N_R(X)| \le |A_X|^{(2n+n!)^d}$ for every R with all dimensions at least n, completing the proof for $K = 2n + n!$ and $N = |A_X|^{(2n+n!)^d}$. □

We now need a few lemmas for the proof of Theorem 1.3. The first shows that for the purposes of proving X sofic, we may always without loss of generality pass to a higher power shift.

Lemma 3.2. *For any d, for any $\mathbb{Z}^d$ subshift X and rectangular prism $R \subseteq \mathbb{Z}^d$, X is sofic if and only if the higher power shift $X^{[R]}$ is sofic.*

Proof. $\Longrightarrow$: Suppose that X is sofic. Then there is a 1-block factor map ϕ and $\mathbb{Z}^d$ nearest-neighbor SFT Y so that $X = \phi(Y)$. But then it is easy to check that $X^{[R]} = \phi^{[R]}(Y^{[R]})$, where $\phi^{[R]}$ acts on patterns in A_Y^R via coordinatewise action of ϕ. Since $Y^{[R]}$ is a $\mathbb{Z}^d$ SFT and $\phi^{[R]}$ is a factor map, clearly $X^{[R]}$ is sofic.

$\Longleftarrow$: Suppose that $X^{[R]}$ is sofic, and without loss of generality, write $R = \prod_{i=1}^d [0, n_i - 1]$. Then there is a 1-block factor map ψ and $\mathbb{Z}^d$ nearest-neighbor SFT Z so that $X^{[R]} = \psi(Z)$. Define a $\mathbb{Z}^d$ nearest-neighbor SFT Z' with alphabet $A_Z \times R$ by the following rules:

(1) In the x_i-direction, a letter of the form $(a, (v_1, \ldots, v_d))$ must be followed by a letter of the form $(b, (v_1, \ldots, v_{i_1}, v_i + 1 \pmod{n_i}, v_{i+1}, \ldots, v_d))$.
(2) In rule (1), if $v_i \neq n_i - 1$, then $b = a$.
(3) In rule (1), if $v_i = n_i - 1$, then b must be a legal follower of a in the x_i-direction in the nearest-neighbor SFT Z.

The effect of these rules is that in any point of Z', $\mathbb{Z}^d$ is partitioned into translates of R, each translate of R has a constant "label" from A_Z, and the "labels" of these translates comprise a legal point of Z. We now define a 1-block factor map ϕ' on Z' by the rule $\phi'(a, v) = (\phi(a))(v)$, i.e. the letter of A_Z appearing at location v in $\phi(a)$, which was by definition a pattern in A_Z^R. This has the effect of, in each point of Z', filling every translate of R with the image under ϕ of the letter of A_Z which was its label. Since these labels comprise a point of Z and since $\phi(Z) = X^{[R]}$, the reader may check that $\phi'(Z') = X$, and so X is sofic.

□

Our next lemma shows that an upper bound for $N_R(X)$ over all large finite rectangular prisms R must also be an upper bound for $N_R(X)$ even when we allow R to have some infinite dimensions.

Lemma 3.3. *For any d and any $\mathbb{Z}^d$ subshift X, if there exist K, N so that $N_R(X) \leq N$ for any rectangular prism R with dimensions at least K, then it is also the case that $N_{R'}(X) \leq N$ for any "infinite rectangular prism" of the form $R' = \prod_{i=1}^d I_i$, where each of the I_i is either an interval of integers with length at least K or $\mathbb{Z}$.*

Proof. Consider any K, N, X satisfying the hypotheses of the theorem, and any "infinite rectangular prism" R' with all dimensions either finite and greater than K or infinite. Suppose for a contradiction that there exist $N+1$ distinct configurations $w_1, \ldots, w_{N+1}$ in $L_{R'}(X)$ and that their extender sets $E_X(w_i)$ are distinct. Then, for each pair $i < j \in [1, N+1]$, there exists a pattern $v_{ij} \in L_{R'^c}(X)$ s.t. $v_{ij} w_i \in X$ and $v_{ij} w_j \notin X$ or vice versa. By compactness, for each v_{ij}, there exists n_{ij} so that $v_{ij} w_i|_{[-n_{ij}, n_{ij}]^d \cap R'} \in L(X)$ and $v_{ij} w_i|_{[-n_{ij}, n_{ij}]^d \cap R'} \notin L(X)$, or vice versa. This property is clearly preserved by increasing n_{ij}. Therefore, if we define $M = \max(K, \{n_{ij}\}_{i<j})$, then for every $i < j \in [1, N+1]$, either $v_{ij} w_i|_{[-M,M]^d \cap R'} \in L(X)$ and $v_{ij} w_i|_{[-M,M]^d \cap R'} \notin L(X)$ or vice versa. Put another way, $E_X(w_i|_{[-M,M]^d \cap R'})$ contains a pattern which equals v_{ij} on R'^c, and $E_X(w_j|_{[-M,M]^d \cap R'})$ contains no such pattern, or vice versa. Either way, this shows that $E_X(w_i|_{[-M,M]^d \cap R'}) \neq$

$E_X(w_j|_{[-M,M]^d \cap R'})$ and, since i, j were arbitrary, that all $N+1$ of the extender sets $E_X(w_i|_{[-M,M]^d \cap R'})$, $i \in [1, N+1]$, are distinct. Since $[-M,M]^d \cap R'$ is a finite rectangular prism with all dimensions at least K, this contradicts the hypotheses of the theorem. Our original assumption was therefore wrong, and $N_{R'}(X) \leq N$. □

Our final preliminary lemma shows that for $d = 1$, boundedness of $N_{[1,n]}(X)$ is equivalent to soficity of X.

Lemma 3.4. *For a $\mathbb{Z}$ subshift X, X is sofic if and only if $N_{[1,n]}(X)$ is a bounded sequence.*

Proof. $\Longrightarrow$: If X is sofic, then there is a 1-block map ϕ and nearest-neighbor SFT Y, with alphabet A_Y, so that $X = \phi(Y)$. Then, for any finite word $w \in L_{[1,n]}(Y)$, clearly $E_X(w) = \bigcup_{y \in \phi^{-1}(w)} \phi(E_Y(y))$. Since Y is a nearest-neighbor SFT, this clearly depends only on the set of pairs of first and last letters of ϕ-preimages of w, and there are fewer than $2^{|A_Y|^2}$ such sets. Therefore, $N_{[1,n]}(X) \leq 2^{|A_Y|^2}$ for all n, and so the sequence $N_{[1,n]}(X)$ is bounded.

$\Longleftarrow$: We prove the contrapositive, the proof will be similar to Lemma 3.3. Suppose that X is not sofic. Then there are infinitely many follower sets $F(p)$ for infinite pasts $p \in A_X^{\mathbb{Z}^-}$. For any N, choose N pasts $p_1, \ldots, p_N$ with distinct follower sets. This means that for every $i < j \in [1, N]$, there exists a future $f_{ij} \in A_X^{\mathbb{N}}$ so that either $p_i f_{ij} \in X$ and $p_j f_{ij} \notin X$, or vice versa. By compactness, there exists N_{ij} so that for any $n > N_{ij}$, either $p_i|_{(-n,0)} f_{ij} \in X$ and $p_j|_{(-n,0)} f_{ij} \notin X$ or vice versa. But then if we take $M = \max N_{ij}$, then for every $i < j \in [1, N]$, either $p_i|_{(-M,0)} f_{ij} \in X$ and $p_j|_{(-M,0)} f_{ij} \notin X$ or vice versa, meaning that the N extender sets $E_X(p_i|_{(-M,0)})$, $i \in [1, N]$, are distinct. Therefore, $N_{[1,M]}(X) \geq N$. Since N was arbitrary, $N_{[1,n]}(X)$ is not bounded. □

As an aside, before proving Theorem 1.3 we present the example mentioned in the introduction, of a $\mathbb{Z}$ sofic shift X where $N_{[1,n]}(X)$ is bounded, but does not stabilize. In fact, the number of distinct follower sets of words of length n also fails to stabilize for this shift, which may be of independent interest.

Example 3.5. ([1]) Define X to be the sofic shift consisting of all labels of bi-infinite paths on the labeled graph $\mathcal{G}$ below. Then for all $n > 1$, $M_{[1,2n]}(X) = 14$, $M_{[1,2n+1]}(X) = 13$, $N_{[1,2n]}(X) = 46$, and $N_{[1,2n+1]}(X) = 44$.

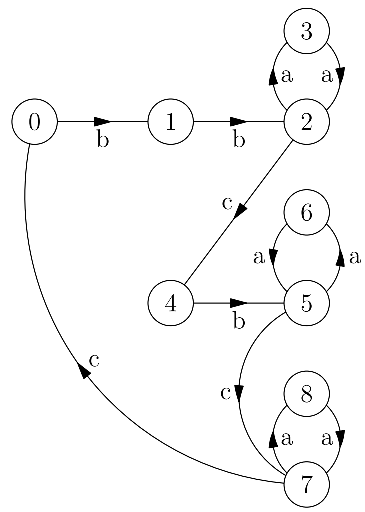

Proof. The reader may check that $\mathcal{G}$ is follower-separated (see [5] for a definition), and so for any $w \in L(X)$, the follower set $F_X(w)$ is determined by the set of terminal vertices for paths in $\mathcal{G}$ with label w, which we'll denote by $T(w)$. We now simply describe the possible sets $T(w)$ for words of even and odd length, with examples of words realizing each set. We use the notation $*$ to indicate that any word may replace the $*$, and n to represent any nonnegative integer.

(even length)

$T(w)$	w
$\{0\}$	$*cc$
$\{1\}$	$*ccb$
$\{2\}$	$*cbb$
$\{3\}$	$*bba$
$\{4\}$	$*bbc$
$\{5\}$	$*bcb$
$\{6\}$	$*cba$
$\{7\}$	$*cbc$
$\{8\}$	$*bca$
$\{1,5\}$	$a^n cb$
$\{3,6\}$	ba^{2n+1}
$\{4,7\}$	$ba^{2n}c$
$\{0,4,7\}$	$a^n c$
$\{2,3,5,6,7,8\}$	$a^{2(n+1)}$

(odd length)

$T(w)$	w
$\{0\}$	$*cc$
$\{1\}$	$*ccb$
$\{2\}$	$*cbb$
$\{3\}$	$*bba$
$\{4\}$	$*bbc$
$\{5\}$	$*bcb$
$\{6\}$	$*cba$
$\{7\}$	$*cbc$
$\{8\}$	$*bca$
$\{1,5\}$	$a^n cb$
$\{2,5\}$	$ba^{2(n+1)}$
$\{0,4,7\}$	$a^n c$
$\{2,3,5,6,7,8\}$	a^{2n+1}

We leave it to the reader to check that there are no follower sets aside from the ones described here, and so $M_{[1,2n]}(X) = 14$ and $M_{[1,2n+1]}(X) = 13$ for all $n > 1$. Informally, the reason that words of even length have an additional follower set is that the word $ba^{2n}c$ has a follower set (given by the set $\{3,6\}$ of terminating states) which can not be recreated by odd length; every cycle has even length, knowledge of at least one letter on each side of the cycle is required to create a new follower set, and knowledge of two letters on either side makes the word synchronizing (meaning there is only a single terminating state.)

Since listing 46 and 44 extender sets similarly (for even and odd lengths respectively) would be quite cumbersome, we will not give a complete list of these, but will give a sketch of how they appear. First, note that $\mathcal{G}$ is also predecessor separated, and so the extender set of a word $w \in L(X)$ is determined entirely by the set $\{v \to v'\}$ of possible pairs of initial and terminal vertices of paths in $\mathcal{G}$ with label w, which we denote by $S(w)$. Note that partitioning the vertices into $\{0, 2, 5, 7\}$ and $\{1, 3, 4, 6, 8\}$ shows that $\mathcal{G}$ is bipartite. The reader may check that every possible singleton $\{v \to v'\}$ for pairs v, v' in the same vertex class occurs as $S(w)$ for a word w of even length, and every possible singleton $\{v \to v'\}$ for pairs v, v' in opposite vertex classes occurs as $S(w)$ for a word w of odd length. This contributes $5^2 + 4^2 = 41$ extender sets to $N_{[1,2n]}(X)$ and $2 \cdot 5 \cdot 4 = 40$ extender sets to $N_{[1,2n+1]}(X)$ for every $n > 1$. The remaining sets $S(w)$, along with w presenting them, appear in the table below. Note that though we can informally pair up the first through fourth types in each case, again the word $ba^{2n-2}c$ creates an extender set with no analogous extender set for a word of odd length.

(even length)

$S(w)$	w
$\{2 \to 2, 3 \to 3, 5 \to 5, 6 \to 6, 7 \to 7, 8 \to 8\}$	a^{2n}
$\{1 \to 3, 4 \to 6\}$	ba^{2n-1}
$\{2 \to 5, 7 \to 1\}$	$a^{2n-2}cb$
$\{3 \to 4, 6 \to 7, 8 \to 0\}$	$a^{2n-1}c$
$\{1 \to 4, 4 \to 7\}$	$ba^{2n-2}c$

(odd length)

$S(w)$	w
$\{2 \to 3, 3 \to 2, 5 \to 6, 6 \to 5, 7 \to 8, 8 \to 7\}$	a^{2n-1}
$\{1 \to 2, 4 \to 5\}$	ba^{2n}
$\{3 \to 5, 8 \to 1\}$	$a^{2n-1}cb$
$\{2 \to 4, 5 \to 7, 7 \to 0\}$	$a^{2n}c$

□

Proof of Theorem 1.3. Our proof proceeds by induction on d. The base case $d = 1$ is precisely Lemma 3.4. We now assume that the result holds for $\mathbb{Z}^{d-1}$ subshifts, and will prove it for $\mathbb{Z}^d$ subshifts. To that end, assume that X is a $\mathbb{Z}^d$ subshift and that there exist K, N so that for any rectangular prism R with dimensions at least K, $N_R(X) \leq N$. By Lemma 3.3, the same is true even if R has some infinite dimensions.

Note that by Lemma 3.2, we may without loss of generality replace X by the higher power shift $X^{[[1,K]^d]}$. Since $N_R(X^{[[1,K]^d]}) \leq N$ for all rectangular prisms, with no restrictions on the dimension, we will assume this property for X in the remainder of the proof.

Define $X' = \{x|_{\mathbb{Z}^{d-1} \times \{0\}} \ : \ x \in X\}$, the set of restrictions of points of X to hyperplanes spanned by the first $d-1$ cardinal directions. By the assumption above, there are fewer than N distinct extender sets for $x \in X'$, and so we define equivalence classes C_i, $i \in [1, M]$, $M \leq N$, for the equivalence relation defined by

$x \sim y$ if $E_X(x) = E_X(y)$. In a slight abuse of notation, we denote by $E_X(C_i)$ the common extender set shared by all $x \in C_i$.

Now, consider any $x \in X'$ and $k \in [1, d-1]$. By the pigeonhole principle, there exist $i < j \in [1, M+1]$ so that $\sigma_{ie_k} x \sim \sigma_{je_k} x$. But then for $y \in L_{\mathbb{Z}^{d-1} \times \{0\}^c}(X)$, $y \in E_X(x) \iff \sigma_{ie_k} y \in E_X(\sigma_{ie_k} x) \iff \sigma_{iek} y \in E_X(\sigma_{je_k} x) \iff y \in E_X(\sigma_{(j-i)e_k} x)$. Therefore, $x \sim \sigma_{(j-i)e_k} x$, and the same logic shows that $\sigma_{(j-i)me_k} x \sim x$ for any $m \in \mathbb{Z}$. Since $j - i \leq M \leq N$, this shows that the C_i containing x is invariant under shifts by $N!e_k$ for $k \in [1, d-1]$. Since $x \in X'$ was arbitrary, this means that every C_i is invariant under shifts by each $N!e_k$. We may then, again by Lemma 3.2, replace X by its higher power shift $X^{[[1,N!]^{d-1} \times \{0\}]}$, which allows us to assume without loss of generality that all of the C_i are shift-invariant subsets of $A_X^{\mathbb{Z}^{d-1}}$. The classes C_i need not, however, be closed. Their closures are $\mathbb{Z}^{d-1}$ subshifts though, and we will show that they in fact must be sofic.

Claim 1: $\overline{C_i}$ is sofic for each i.

It suffices to show that for any rectangular prism $R \subseteq \mathbb{Z}^{d-1}$ and $w, w' \in L_R(\overline{C_i})$, $E_X(w) = E_X(w') \implies E_{\overline{C_i}}(w) = E_{\overline{C_i}}(w')$, since then $N_{R \times \{0\}}(X) \leq N \implies N_R(\overline{C_i}) \leq N$ for all rectangular prisms R, which will imply the desired conclusion by our inductive hypothesis.

So, assume that $E_X(w) = E_X(w')$ for $w, w' \in L_R(\overline{C_i})$. Suppose also that $vw \in \overline{C_i}$ for some $v \in A_X^{\mathbb{Z}^{d-1} \setminus R}$. Then there exists $v_n \in A_X^{\mathbb{Z}^{d-1} \setminus R}$ so that $v_n \to v$ and $v_n w \in C_i$ for all n. Then for any $y \in E_X(C_i)$, $y v_n w \in X$, since all $v_n w$ share the same class C_i. Since $E_X(w) = E_X(w')$, $y v_n w' \in X$ as well. Similarly, for any $y \notin E_X(C_i)$, $y v_n w \notin X$, and so $y v_n w' \notin X$. But then $E_X(v_n w') = E_X(C_i)$, and so $v_n w' \in C_i$. By taking limits, $vw' \in \overline{C_i}$. We've then shown that $vw \in \overline{C_i} \implies vw' \in \overline{C_i}$. The converse is true by the same proof, and so $E_{\overline{C_i}}(w) = E_{\overline{C_i}}(w')$, completing the proof of soficity of $\overline{C_i}$ as described above.

Since all elements of any class C_i are interchangeable in points of X, we can define $V \subseteq [1, M]^{\mathbb{Z}}$ which lists legal sequences of classes (in the e_d-direction) within points in X:

$$V := \{(i_n) \in [1, M]^{\mathbb{Z}} \;:\; \exists x \in X \text{ such that } \forall n \in \mathbb{Z}, x|_{\mathbb{Z}^{d-1} \times \{n\}} \in C_{i_n}\}.$$

It is obvious that V is shift-invariant since X is shift-invariant. However, it is not immediately clear that V is closed since the C_i are not necessarily closed. We will show that V is closed by proving the following claim.

Claim 2: $X = \{x \in A_X^{\mathbb{Z}^d} \;:\; \exists (i_n) \in V \text{ such that } \forall n \in \mathbb{Z}, x|_{\mathbb{Z}^{d-1} \times \{n\}} \in \overline{C_{i_n}}\}$.

In other words, given $v \in V$, not only can you make points of X by substituting in configurations from the classes given by the letters in v, but you may also substitute configurations from the closures of these classes.

$X \subseteq \{x \in A_X^{\mathbb{Z}^d} : \exists (i_n) \in V$ such that $\forall n \in \mathbb{Z}, x|_{\mathbb{Z}^{d-1}\times\{n\}} \in \overline{C_{i_n}}\}$:

First, we note that for any $x \in X$, by definition $x|_{\mathbb{Z}^{d-1}\times\{n\}} \in X'$ for all $n \in \mathbb{Z}$, and so each of these is in some class C_i. Define $v = (i_n) \in [1, M]^{\mathbb{Z}}$ by saying that the $x|_{\mathbb{Z}\times\{n\}}$ of x is in C_{i_n}. Then by definition of V, $v \in V$. This clearly shows the desired containment.

$X \supseteq \{x \in A_X^{\mathbb{Z}^d} : \exists (i_n) \in V$ such that $\forall n \in \mathbb{Z}, x|_{\mathbb{Z}\times\{n\}} \in \overline{C_{i_n}}\}$:

Choose any $x \in A_X^{\mathbb{Z}^d}$ so that there is $v = (i_n) \in V$ with the property that $\forall n \in \mathbb{Z}$, $x|_{\mathbb{Z}^{d-1}\times\{n\}} \in \overline{C_{i_n}}$. Then, for each $n \in \mathbb{Z}$, there exists a sequence $x^{(k,n)} \in C_{i_n}$ so that $x^{(k,n)} \underset{k\to\infty}{\longrightarrow} x|_{\mathbb{Z}^{d-1}\times\{n\}}$ for all n. Also, since $v \in V$, there exists $x' \in X$ so that $\forall n \in \mathbb{Z}$, $x'|_{\mathbb{Z}^{d-1}\times\{n\}} \in C_{i_n}$.

We now define, for every k, the point $x^{(k)} \in A_X^{\mathbb{Z}^d}$ by

$$x^{(k)}|_{\mathbb{Z}^{d-1}\times\{n\}} = \begin{cases} x^{(k,n)} \text{ if } |n| \leq k \\ x'|_{\mathbb{Z}^{d-1}\times\{n\}} \text{ if } k \leq |n| \end{cases}$$

The central $2k+1$ $(d-1)$-dimensional hyperplanes of $x^{(k)}$ are given by the $x^{(k,n)}$, and the remaining $(d-1)$-dimensional hyperplanes are unchanged from x'. We note that $x^{(k)}$ can be obtained from $x' \in X$ by making $2k+1$ consecutive replacements of $x'|_{\mathbb{Z}^{d-1}\times\{n\}}$ by $x^{(k,n)}$. Since these replacements involve configurations in the same class C_{i_n}, each of these replacements preserves being in X, and so $x^{(k)} \in X$ for all k. Finally, we note that $x^{(k)} \to x$, so $x \in X$ as well, showing the desired containment.

Claim 3: V is a sofic subshift.

We first show that V is closed and therefore a subshift. Let $v^{(k)} \in V$ and $v^{(k)} \to v = (i_n)$. By passing to a subsequence, we may assume that for all $k \geq |n|$, $v_n^{(k)} = i_n$. By definition of V, for every $k \in \mathbb{N}$ there exists $x^{(k)} \in X$ where $x^{(k)}|_{\mathbb{Z}^{d-1}\times\{n\}} \in C_{v_n^{(k)}}$ for every n. For all $n \leq k$, since $C_{v_n^{(k)}} = C_{i_n} = C_{v_n^{(n)}}$ we may replace the pattern $x^{(k)}|_{\mathbb{Z}^{d-1}\times\{n\}}$ in $x^{(k)}$ with $x^{(n)}|_{\mathbb{Z}^{d-1}\times\{n\}}$ to form a legal point in X. In such a way we obtain a new sequence of points $y^{(k)} \in X$ where $y^{(k)}|_{\mathbb{Z}^{d-1}\times\{n\}} \in C_{v_n^{(k)}}$, but with the additional property that $y^{(k)}|_{\mathbb{Z}^{d-1}\times\{n\}} = y^{(n)}|_{\mathbb{Z}^{d-1}\times\{n\}}$ for $k \geq |n|$. The sequence $y^{(k)}$ clearly converges to a point $y \in A_X^{\mathbb{Z}^d}$, and $y \in X$ since X is closed. Since $y|_{\mathbb{Z}^{d-1}\times\{n\}} = y^{(n)}|_{\mathbb{Z}^{d-1}\times\{n\}} \in C_{i_n}$, we have $v \in V$.

We claim that $N_{[1,m]}(V) \leq N$ for all $m \in \mathbb{N}$, which will prove the claim by Lemma 3.4. Suppose for a contradiction that there are $N+1$ words $v^{(1)}, \ldots, v^{(N+1)} \in L_m(V)$ s.t. the extender sets $E_V(v^{(i)})$ are all distinct. Then, for every $i < j \in$

$[1, N+1]$, there exists $w^{(ij)} \in L_{[1,m]^c}(V)$ s.t. either $v^{(i)}w^{(ij)} \in V$ and $v^{(j)}w^{(ij)} \notin V$ or vice versa.

For each $v^{(i)} = v_1^{(1)} \dots v_m^{(1)}$, define $S^{(i)} \in A_X^{\mathbb{Z}^{d-1} \times [1,m]}$ by choosing $S^{(i)}|_{\mathbb{Z}^{d-1} \times \{n\}}$ to be any row in $C_{v_n^{(i)}}$. Similarly, for each $w^{(ij)}$, define a pattern $O^{(ij)} \in A_X^{\mathbb{Z}^{d-1} \times [1,m]^c}$ by choosing $O^{(ij)}|_{\mathbb{Z}^{d-1} \times \{b\}}$ to be any row in $C_{w_n^{(ij)}}$. Then, by Claim 2, $S^{(i)}O^{(ij)} \in X$ and $S^{(j)}O^{(ij)} \notin X$ or vice versa, meaning that all extender sets $E_X(S^{(i)})$, $i \in [1, N+1]$, are distinct. This is a contradiction to Lemma 3.3, and so our original assumption was wrong, $N_{[1,m]}(V)$ is a bounded sequence, and V is sofic.

We may now finally construct an SFT cover of X to show that it is sofic. Since V is sofic by Claim 3, we may define a 1-block factor ψ and nearest-neighbor SFT W so that $\psi(W) = V$. For each $a \in A_V$, since $\overline{C_a}$ is sofic by Claim 1, there is a 1-block factor ϕ_a and nearest-neighbor $\mathbb{Z}^{d-1}$ SFT Y_a (whose alphabet we denote by A_a) so that $\phi_a(Y_a) = \overline{C_a}$. Now, define a nearest-neighbor $\mathbb{Z}^d$ SFT Y with alphabet $A_Y := \bigcup_{a \in A_W} (\{a\} \times A_{Y_{\psi(a)}})$ by the following adjacency rules:

(1) Any pair of letters (a, s), (b, t) which are adjacent in one of the first $d-1$ cardinal directions must share the same first coordinate, i.e. $a = b$.
(2) (a, s) may legally precede (a, t) in the e_i-direction for $i \in [1, d-1]$ if and only if s may legally precede t in the same direction in $Y_{\psi(a)}$.
(3) (a, s) may legally precede (b, t) in the e_d-direction if and only if a may legally precede b in W. (There is no restriction on the second coordinates s, t.)

Clearly for any $y \in Y$, these rules force the $(d-1)$-dimensional hyperplanes $y|_{\mathbb{Z}^{d-1} \times \{n\}}$ to have constant first coordinate (say a_n), force the second coordinates to form a point in $Y_{\psi(a_n)}$, and force the sequence (a_n) to be in W. We now define the 1-block factor map ϕ on Y by $\phi(a, s) = \phi_{\psi(a)}(s)$.

Claim 4: $\phi(Y) = X$.

$\phi(Y) \subseteq X$: Take any $y \in Y$, and define a_n to be the first coordinate shared by all letters in $y|_{\mathbb{Z}^{d-1} \times \{n\}}$. Then by definition of Y, $(a_n) \in W$. Also by definition of Y, the second coordinates of the letters in $y|_{\mathbb{Z}^{d-1} \times \{n\}}$ form a point of $Y_{\psi(a_n)}$, call it $b^{(n)}$. Then, $(\psi(a_n)) \in V$, and for every $n \in \mathbb{Z}$, $(\phi(y))|_{\mathbb{Z}^{d-1} \times \{n\}} = \phi_{\psi(a_n)}(b^{(n)})$ is the $\phi_{\psi(a_n)}$-image of a point of $Y_{\psi(a_n)}$, and so is in $\overline{C_{a_n}}$. But then, by Claim 2, $\phi(y) \in X$, and since $y \in Y$ was arbitrary, $\phi(Y) \subseteq X$.

$\phi(Y) \supseteq X$: Choose any $x \in X$. For every $n \in \mathbb{Z}$, $x|_{\mathbb{Z}^{d-1} \times \{n\}}$ is in one of the C_i, and if we define a sequence i_n by $x|_{\mathbb{Z}^{d-1} \times \{n\}} \in C_{i_n}$, then $(i_n) \in V$. Choose any $(a_n) \in W$ s.t. $(\psi(a_n)) = (i_n)$. For each $n \in \mathbb{Z}$, since $x|_{\mathbb{Z}^{d-1} \times \{n\}} \in \overline{C_{i_n}} = \overline{C_{\psi(a_n)}}$, there exists $b^{(n)} \in Y_{\psi(a_n)}$ s.t. $\phi_{\psi(a_n)}(b^{(n)}) = x|_{\mathbb{Z}^{d-1} \times \{n\}}$. Define a point $y \in A_Y^{\mathbb{Z}^2}$ by setting, for all $t = (t_1, \dots, t_d) \in \mathbb{Z}^d$, $y(t) = (a_{t_d}, b^{(n)}(t_1, \dots, t_{d-1}))$. Then $y \in Y$ since (a_n) is in W, and for all $n \in \mathbb{Z}$, the second coordinates in each hyperplane

$y|_{\mathbb{Z}^{d-1}\times\{n\}}$ create a point of $Y_{\psi(a_n)}$. It is also clear that $\phi(y) = x$, and since X was arbitrary, $\phi(Y) \supseteq X$.

Since Y was an SFT and ϕ a factor map, this shows that X is sofic, completing the proof of Theorem 1.3.

□

We would like to say a bit more about shifts satisfying the hypotheses of Theorem 1.3, i.e. those with eventually bounded numbers of extender sets, because in fact they satisfy a much stronger (though technical) condition than just being sofic.

Definition 3.6. We say that a $\mathbb{Z}^d$ nearest-neighbor SFT X is **decouplable** if either $d = 1$ (in which case X is automatically called decouplable) or there exist $i \in [1, d]$, a $\mathbb{Z}$ nearest-neighbor SFT W, and decouplable $\mathbb{Z}^{d-1}$ nearest-neighbor SFTs Y_a for each $a \in A_W$, with disjoint alphabets, so that

$$X = \{x \in A_X^{\mathbb{Z}^d} \ : \ \exists w = (w_n) \in W \text{ s.t. } \forall n \in \mathbb{Z}, x|_{\mathbb{Z}^{i-1}\times\{n\}\times\mathbb{Z}^{d-i}} \in Y_{w_n}\}.$$

(Here, we have made the obvious identification between $\mathbb{Z}^{i-1} \times \{n\} \times \mathbb{Z}^{d-i}$ and $\mathbb{Z}^{d-1}$.)

In other words, X is decouplable if one can construct it by starting from a one-dimensional nearest neighbor SFT and then arbitrarily replacing occurrences of each letter in its alphabet by points from a $\mathbb{Z}^{d-1}$ decouplable nearest-neighbor SFT associated to that letter. This definition is obviously recursive; X is decouplable if its $(d-1)$-dimensional hyperplanes are given by decouplable SFTs, whose $(d-2)$-dimensional hyperplanes are given by decouplable SFTs, and so on. This means that though X is a $\mathbb{Z}^d$ SFT, its behavior is in some sense one-dimensional.

Remark 3.7. In fact the SFT cover in the proof of Theorem 1.3 can always be chosen to be decouplable. By the inductive hypothesis, the covers Y_a can each be chosen to be decouplable $\mathbb{Z}^{d-1}$ SFTs, and then the construction of X from W and all Y_a clearly yields a decouplable $\mathbb{Z}^d$ SFT.

In order to give an application of Theorem 1.3 and to elucidate the idea of decouplable SFTs, we will present a brief example.

Example 3.8. Define X to be the $\mathbb{Z}^2$ subshift on $\{0, 1\}$ consisting of all $x \in \{0, 1\}^{\mathbb{Z}^2}$ with either no 1s, a single 1, or two 1s.

Then it is not hard to see that for any $S \subseteq \mathbb{Z}^2$ with $|S| > 1$, $N_S(X) = 3$. It is easily checked that the three possible extender sets for $w \in L_S(X)$ are:

• If w contains no 1s, then $E_X(w)$ consists of all patterns on S^c with either no 1s, a single 1, or two 1s.
• If w contains a single 1, then $E_X(w)$ consists of all patterns on S^c with either no 1 or a single 1.
• If w contains two 1s, then $E_X(w)$ consists of the single pattern on S^c with no 1s, namely 0^{S^c}.

We will now describe how the proof of Theorem 1.3 yields an SFT cover for X. Using the language of the proof of Theorem 1.3, X' consists of all bi-infinite 0-1 sequences with either no 1s, a single 1, or two 1s. X' is broken into three classes of rows with the same extender sets in X, which are again classified by number of 1s contained:

- $C_0 = \{x \in \{0,1\}^{\mathbb{Z}} \ : \ x$ contains no 1s$\} = \{0^{\mathbb{Z}}\}$.
- $C_1 = \{x \in \{0,1\}^{\mathbb{Z}} \ : \ x$ contains exactly one 1$\}$.
- $C_2 = \{x \in \{0,1\}^{\mathbb{Z}} \ : \ x$ contains exactly two 1s$\}$.

(We've written the C_is with subscripts starting from 0 rather than 1 so that V can be more easily described; clearly this has no effect on the proof.) Note that each of the C_i is shift-invariant, but C_1 and C_2 are not closed. However, each closure $\overline{C_i}$ is sofic. This is easily checked, but we will momentarily explicitly describe SFT covers of the $\overline{C_i}$ anyway.

We now wish to find V, the $\mathbb{Z}$ subshift with alphabet $\{0,1,2\}$ which describes how the rows in various classes can fit together to make points of X. This is not so hard to see: since points of X must have at most two 1s, and since the classes C_i are partitioned by number of 1s,

$$V = \{v \in \{0,1,2\}^{\mathbb{Z}} \ : \ v \text{ has only finitely many nonzero digits, and } \sum v_n \leq 2\}.$$

Points of X are then constructed by beginning with a sequence in V, writing it vertically, and replacing each letter v_n with an arbitrary element of C_{v_n}. So, for instance, one could start with $\ldots 0002000 \ldots \in V$, replace all 0s by the single sequence $\ldots 000000 \ldots \in C_0$, and replace 2 by any sequence in C_2, for instance $\ldots 0001001000 \ldots$. Clearly every point obtained in this way will have at most two 1s and so will be in X. In addition though, as described in Claim 2, one can also replace each v_n by an arbitrary element of the closure $\overline{C_{v_n}}$. For instance, if we chose to replace the 2 in our earlier sequence by $\ldots 0001000 \ldots$, which is not in C_2 (in fact it's in C_1), but is in $\overline{C_2}$, we would still arrive at a legal point of X.

We now note that V is a sofic shift, with nearest-neighbor SFT cover W defined as follows: $A_W = \{A,B,C,D,E,F\}$, and legal adjacent pairs in W are AA, AB, BC, CC, CD, DE, EE, AF, and FE. So,

$$W = \{A^\infty, C^\infty, E^\infty, A^\infty BC^\infty, C^\infty DE^\infty, A^\infty BC^n DE^\infty, A^\infty FE^\infty\}.$$

The factor ψ is defined by $\psi(A) = \psi(C) = \psi(E) = 0$, $\psi(B) = \psi(D) = 1$, and $\psi(F) = 2$, and it is easily checked that $\psi(W) = V$.

Following our proof of Theorem 1.3, the next step is to construct SFT covers of each $\overline{C_i}$, which is straightforward. Define Y_0 to consist of the single fixed point $\{a^\infty\}$, and ϕ_0 by $\phi_0(a) = 0$. Define Y_1 to have alphabet $\{a,b,c\}$ with legal adjacent pairs aa, ab, bc, and cc; then $Y_1 = \{a^\infty, c^\infty, a^\infty bc^\infty\}$. Define ϕ_1 by $\phi_1(a) = \phi_1(c) = 0$ and $\phi_1(b) = 1$. Finally, define Y_2 to have alphabet $\{a,b,c,d,e\}$ with legal adjacent pairs aa, ab, bc, cc, cd, de, and ee; then $Y_2 = \{a^\infty, c^\infty, e^\infty, a^\infty bc^\infty, c^\infty de^\infty, a^\infty bc^n de^\infty\}$. Define ϕ_2 by $\phi_2(a) = \phi_2(c) = \phi_2(e) = 0$ and $\phi_2(b) = \phi_2(d) = 1$. The reader may check that $\phi_i(Y_i) = \overline{C_i}$ for each i.

We may now construct an SFT cover Y for X following our proof. The alphabet $A_Y := \bigcup_{a \in A_W} (\{a\} \times A_{\psi(a)}) = \{(A,a),(B,a),(B,b),(B,c),(C,a),(D,a),(D,b),(D,c),$ $(E,a),(F,a),(F,b),(F,c),(F,d),(F,e)\}$. The adjacency rules are that horizontally adjacent letters have the same first (capital) coordinate Π and second (lowercase) coordinates satisfying the adjacency rules given by $Y_{\psi(\Pi)}$, and that vertically adjacent letters have first (capital) coordinates satisfying the adjacency rules of W. So, for instance, (D,a) cannot appear immediately to the left of (D,c), since ac is not a legal pair in $Y_{\psi(D)} = Y_1$. On the other hand, (A,a) can appear below (F,d), since AF is a legal pair in W. The map ϕ is defined, as before by $\phi(a,b) = \phi_{\psi(a)}(b)$, meaning that $\phi(B,b) = \phi(D,b) = \phi(F,b) = \phi(F,d) = 1$, and all other letters of A_Y have ϕ-image 0. It's easy to see that $\phi(Y) = X$; the rules defining Y mean that there are at most two letters with second coordinate b or d, and these are the only letters of A_Y which have ϕ-image 1.

□

Finally, we will prove Theorem 1.4, but first need the following definition and theorem from [3].

Definition 3.9. ([3]) A $\mathbb{Z}^d$ subshift X is **effective** if there exists a forbidden list (w_n) for X and a Turing machine which, on input n, outputs w_n.

It is easy to see that not every subshift is effective; there are only countably many Turing machines, and so only countably many effective subshifts.

Theorem 3.10. ([3]) *Any $\mathbb{Z}^d$ sofic shift is effective.*

Proof of Theorem 1.4. For any Sturmian $\mathbb{Z}$ subshift S, we can extend S to a $\mathbb{Z}^d$ subshift $\widetilde{S}$ by enforcing constancy along all cardinal directions e_i, $i \in [2,d]$. There are uncountably many Sturmian subshifts, and so there exists one, call it S', s.t. $\widetilde{S'}$ is not effective. (In fact, effectiveness of Sturmian S and/or the shift $\widetilde{S}$ is equivalent to computability of the rotation number defining S, but we will not need this fact.) By Theorem 3.10, $\widetilde{S'}$ is not sofic. In addition, by the minimal complexity definition of Sturmian subshifts, for every $n \in \mathbb{N}$, $|L_{[1,n]^d}(\widetilde{S'})| = |L_{[1,n]}(S')| = n+1$, and so trivially $N_{[1,n]}(\widetilde{S'}) \leq n+1$. Since $\widetilde{S'}$ is not sofic, by Theorem 1.1, in fact $N_{[1,n]}(\widetilde{S'}) = n+1$ for all n.

□

References

[1] Martin Delacourt, personal communication.

[2] N. Pytheas Fogg, *Substitutions in Dynamics, Arithmetics and Combinatorics*, Lecture Notes in Mathematics **1794**, Springer, Berlin, 2002.

[3] M. Hochman, *On the dynamics and recursive properties of multidimensional symbolic systems*, Inventiones Mathematica **176** (2009), no. 1, 131–167.

[4] S. Kass and K. Madden, *A sufficient condition for non-soficness of higher-dimensional subshifts*, Proc. Amer. Math. Soc., to appear.

[5] D. Lind and B. Marcus, *An Introduction to Symbolic Dynamics and Coding*, Cambridge University Press, 1995.

[6] M. Morse and G. A. Hedlund. *Symbolic dynamics II. Sturmian trajectories*, Amer. J. Math. **62** (1940) 1–42.

[7] R. Pavlov, *A class of nonsofic multidimensional shift spaces*, Proc. Amer. Math. Soc. **141** (2013), no. 3, 987–996.

Nic Ormes, Department of Mathematics, University of Denver, 2360 S. Gaylord St., Denver, CO 80208

E-mail address: `normes@du.edu`

URL: `www.math.du.edu/~ormes/`

Ronnie Pavlov, Department of Mathematics, University of Denver, 2360 S. Gaylord St., Denver, CO 80208

E-mail address: `rpavlov@du.edu`

URL: `www.math.du.edu/~rpavlov/`