On Derivatives and Subpattern Orders of Countable Subshifts

Ville Salo[1] and Ilkka Törmä[2]

[1] TUCS – Turku Centre for Computer Science, Finland,
University of Turku, Finland,
vosalo@utu.fi
[2] University of Turku, Finland,
iatorm@utu.fi

Abstract. We study the computational and structural aspects of countable two-dimensional SFTs and other subshifts. Our main focus is on the topological derivatives and subpattern posets of these objects, and our main results are constructions of two-dimensional countable subshifts with interesting properties. We present an SFT whose iterated derivatives are maximally complex from the computational point of view, a sofic shift whose subpattern poset contains an infinite descending chain, a family of SFTs whose finite subpattern posets contain arbitrary finite posets, and a natural example of an SFT with infinite Cantor-Bendixon rank.

Keywords: countable subshifts, Cantor-Bendixon rank, subpattern poset

1 Introduction

In this article, we study the computational and structural aspects of countable two-dimensional SFTs, with an emphasis on properties of the topological derivative and the so-called subpattern order. Our approach is mainly constructive, in that our main results are examples of subshifts with interesting properties. The study of computational aspects of tilings started with the observation of Wang that from a seed tile, one can simulate a Turing machine by simply drawing its run on the set of tilings. It was conjectured by Wang that without a seed tile, computation cannot be forced, and even that there in fact exists a periodic tiling. The tile sets of [2] and [7] provide counterexamples, and further show that a Turing machine can be forced to run on every tiling by using a self-similar tile construction. In the case that there exist only countably many tilings, the situation is different: since every dynamical system contains a minimal subsystem (which is then the orbit closure of a single point), a countable subshift must contain a periodic point. However, Turing machines can still be run on such tilings, so that configurations not containing a seed tile form a small recursive set. From this observation, a wealth of interesting behavior emerges [4].

In [1], the notion of topological derivative is shown to be very useful for studying structural properties of countable two-dimensional SFTs. Namely, it is straightforward that such an SFT will eventually become empty when the derivative is iterated transfinitely, and at the very end of this process, simple structure must emerge: Unless the SFT in question is in fact finite, the second-to-last nonempty derivative must contain a point with exactly one vector of periodicity, and the last level must contain at least two fully periodic points.

The main open problem in [1] about topological derivatives asks for a countable two-dimensional SFT with infinite rank. This has been completely solved in [4], and we give an independent weaker solution to this problem, hopefully showing a more natural example of why this type of behavior may occur in a countable two-dimensional SFT.

Another notion studied in [1] is the partial order induced by subpattern inclusion: we say $x \succcurlyeq y$ if and only if all patterns seen in y are also seen somewhere in x (that is, x has more patterns than y). Not much is known about this partial order in the class of countable two-dimensional SFTs, in particular (as far as we are aware) it is still open whether there exists such an SFT with an infinite descending chain. We solve this in the positive for two-dimensional sofic shifts, by finding such a chain in a countable sofic shift. While we cannot solve the descending chain problem for SFTs, we find rich structure in this class of partial orders by order-embedding every finite partially ordered set in the partial order of a countable SFT.

In addition to addressing the questions of [1], we study the computational complexity of the derivative of a countable two-dimensional SFT. The extension problem of such an SFT (that is, solving whether a given pattern P occurs in a valid configuration) is easily seen to be Π^0_1, and the complexity of this problem may increase by at most two levels in the arithmetical hierarchy when the derivative is taken. We show a converse to this: there exists a countable two-dimensional SFT X, whose kth derivative is Π^0_{2k+1}-complete for all k. This implies that while the very last levels of the derivation process have simple structure, during the process the complexity may rise arbitrarily high and only decrease in a limit ordinal.

The structure of our paper is as follows. In Section 2, we give the relevant definitions and notation used in the rest of the article. In Section 3, we give our results about ranks attainable from subshifts in one and two dimensions. We solve the one-dimensional sofic case, which is drastically different from its two-dimensional counterpart. For the case of countable two-dimensional SFTs, we obtain transfinite rank with a natural subshift, which does not directly involve computation of any kind. In Section 4, we show that derivatives of countable two-dimensional SFTs can climb arbitrarily high in the arithmetical hierarchy. In Section 5, we order-embed an arbitrary finite poset in the subpattern poset of a countable two-dimensional SFT, and find a countable sofic shift with an infinite descending chain.

2 Definitions and Notation

Let S be a finite set of *symbols*, called the *alphabet*, endowed with the discrete topology. For an integer dimension $d \geq 1$, the set $S^{\mathbb{Z}^d}$, equipped with the product topology, as called the *d-dimensional full shift on S*. Elements x of $S^{\mathbb{Z}^d}$ are called *configurations*. A *pattern over S* is a pair (D, s), where $\mathbf{0} \in D \subset \mathbb{Z}^d$ is a finite *domain*, and $s : D \to S$ gives the arrangement of symbols in D. A pattern $P = (D, s)$ *occurs* in a configuration x, denoted $P \sqsubset x$, if we have $x_{D+\boldsymbol{n}} = P$ for some $\boldsymbol{n} \in \mathbb{Z}^d$. For all $k \in [1, d]$, we define the *shift map* $\sigma_k : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ by $\sigma_k(x)_{\boldsymbol{n}} = x_{\boldsymbol{n}+e_k}$, where $\{e_1, \ldots, e_d\}$ is the natural generator set of $\mathbb{Z}^d$.

A *d-dimensional subshift over S* is a closed subset $X \subset S^{\mathbb{Z}^d}$ satisfying $\sigma_k^{-1}(X) = X$ for all $k \in [1, d]$. Alternatively, all subshifts X can be defined by a set $\mathcal{F}$ of *forbidden patterns* as $X = \{x \in S^{\mathbb{Z}^d} \mid \forall P \in \mathcal{F} : P \not\sqsubset x\}$. If $\mathcal{F}$ is finite, then X is said to be *of finite type* (SFT for short). Given a finite domain $D \subset \mathbb{Z}^d$, the set of patterns occurring in the points of a subshift X with domain D is denoted $\mathcal{B}_D(X)$, the set of all patterns of X is $\bigcup_D \mathcal{B}_D(X) = \mathcal{B}(X)$, and the set of symbols occurring in X is denoted $\mathcal{A}(X)$. A *block map* is a continuous mapping $f : X \to Y$, where X and Y are d-dimensional subshifts (possibly over different alphabets), which intertwines the shift maps of X and Y: $f \circ \sigma_k = \sigma_k \circ f$ for all $k \in [1, d]$. Alternatively, a block map f can be defined by a *local function* $F : \mathcal{B}_D(X) \to \mathcal{A}(Y)$ by $f(x)_{\boldsymbol{n}} = F(x)_{D+\boldsymbol{n}}$ for all $x \in X$ and $\boldsymbol{n} \in \mathbb{Z}^d$, where D is a finite domain, called the *neighborhood* of f. An image of a subshift under a block map is a subshift, and images of SFT's are called *sofic*.

A *partially ordered set* (poset for short) is a tuple $(S, \geq)$, where $\geq$ is a binary relation on S which is reflexive ($x \geq x$ holds for all x), antisymmetric ($x \geq y$ and $y \geq x$ imply $x = y$) and transitive ($x \geq y$ and $y \geq z$ imply $x \geq z$). A *preorder* is a partial order which is not necessarily antisymmetric. An *order-embedding* between two partially ordered sets $(S, \geq)$ and $(T, \succcurlyeq)$ is a function $f : S \to T$ such that for all $x, y \in S$ we have $x \geq y$ iff $f(x) \succcurlyeq f(y)$.

We define a preorder, called the *subpattern order*, on the configurations of $S^{\mathbb{Z}^d}$ by stating that $x \succcurlyeq y$ holds iff $P \sqsubset y$ implies $P \sqsubset x$ for all patterns P, meaning that x has more patterns than y. If $x \succcurlyeq y$ and $x \preccurlyeq y$, we denote $x \approx y$, and if $x \succcurlyeq y$ and $x \not\approx y$, we denote $x \succ y$. The *subpattern poset* of a subshift $X \subset S^{\mathbb{Z}^d}$ is the poset $(X/{\approx}, \succcurlyeq)$, where $\approx$-equivalent elements of X are identified.

Given a topological space X, the *Cantor-Bendixon derivative of X* is defined as $X' = \{x \in X \mid x \in \overline{X - \{x\}}\}$. Thus X' consists of the nonisolated points of X (which might be isolated in X'). We inductively define the λth derivative of X for all ordinals λ. First, $X^{(0)} = X$. If $\lambda = \alpha + 1$, then $X^{(\lambda)} = (X^{(\alpha)})'$, and if λ is a limit ordinal, then $X^{(\lambda)} = \bigcap_{\alpha<\lambda} X^{(\alpha)}$. The *Cantor-Bendixon rank* of an element $x \in X$ is the least ordinal λ such that $x \notin X^{(\lambda)}$, if it exists. Also, the *rank* of X is the least ordinal λ such that $X^{(\lambda)} = X^{(\lambda+1)}$.

Let ϕ be a formula in first-order arithmetic. If ϕ contains only bounded quantifiers, then we say ϕ is Σ^0_0 and Π^0_0. For all $n > 0$, we say ϕ is Σ^0_n if it is equivalent to a formula of the form $\exists k : \psi$ where ψ is Π^0_{n-1}, and ϕ is Π^0_n, if it is equivalent to a formula of the form $\forall k : \psi$ where ψ is Σ^0_{n-1}. This classification is called the *arithmetical hierarchy*. A subset X of $\mathbb{N}$ is Σ^0_n or Π^0_n, if $X = \{x \in \mathbb{N} \mid \phi(x)\}$ for some ϕ with the corresponding classification. The nonstandard quantifier $\exists^\infty n : \phi(n)$ has the meaning 'there exist infinitely many n such that $\phi(n)$.'

A subset $X \subset \mathbb{N}$ is *many-one reducible* (or simply *reducible*) to another set $Y \subset \mathbb{N}$, if there exists a computable function $f : \mathbb{N} \to \mathbb{N}$ such that $x \in X$ iff $f(x) \in Y$. If every set in a class $\mathcal{C}$ is reducible to X, then X is said to be $\mathcal{C}$*-hard*. If, in addition, X is in $\mathcal{C}$, then X is $\mathcal{C}$*-complete*.

In the proof of one of our results, we utilize *counter machines*, which we define here informally. A counter machine M consists of a finite state set Σ and a finite set of counters, each of which holds a value in $\mathbb{N}$. On a single step, the machine increments or decrements some of its counters by 1 and goes to a new state $s \in \Sigma$, depending on its previous state and which of its counters contained the value 0. This action may also be nondeterministic. A counter machine can be used to simulate a Turing machine, and thus to execute any algorithm.

3 Ranks of Subshifts

Ranks of subshifts have usually been studied in the countable case, where we have the following basic result:

Lemma 1 ([1]). *A subshift X has $X^{(\lambda)} = \emptyset$ for some ordinal λ if and only if it is countable.*

In this restricted setting, the rank of a subshift has usually been defined as the smallest such λ.

If a configuration x is isolated in a subshift X, then there exists a pattern $P = (D, s) \in \mathcal{B}(X)$ such that x is the only element of X with $x_D = P$. In this case, we say that P *isolates* x. We also remark here that the derivatives of subshifts are subshifts.

First, we look at one-dimensional sofic shifts. We use the following well-known result:

Definition 1. *A* context *of a word $v \in S^*$ in a subshift $X \subset S^{\mathbb{Z}}$ is $C_X(v) = \{(w, w') \in (S^*)^2 \mid wvw' \sqsubset X\}$.*

Lemma 2 ([6]). *A subshift $X \subset S^{\mathbb{Z}}$ is sofic if and only if it has a finite number of contexts.*

The following lemma relates the contexts of words in a subshift and its derivative.

Lemma 3. *For a subshift $X \subset S^{\mathbb{Z}}$,*

$$C_X(u) = C_X(v) \implies C_{X^{(1)}}(u) = C_{X^{(1)}}(v).$$

Proof. Let $C_X(u) = C_X(v)$, and suppose that $(w, w') \in C_{X^{(1)}}(u) - C_{X^{(1)}}(v)$. Then $wvw' \not\sqsubset X^{(1)}$, so the set of points $x \in X$ with $x_{[0,|wvw'|-1]} = wvw'$ is finite. But since $C_X(u) = C_X(v)$, these are in a bijective correspondence with the points y such that $y_{[0,|wuw'|-1]} = wuw'$, which implies that $wuw' \not\sqsubset X^{(1)}$, a contradiction. □

Corollary 1. *The derivative of a one-dimensional sofic shift is sofic.*

Proposition 1. *All one-dimensional sofic shifts have finite rank.*

Proof. Let X be sofic, and let k be the number of different contexts in X. If $X \neq X^{(1)}$, necessarily $X \neq S^{\mathbb{Z}}$, so we may choose $u \notin \mathcal{B}(X)$. Further, choose $v \in \mathcal{B}(X) - \mathcal{B}(X^{(1)})$. Now, $C_X(u) \neq C_X(v)$, but $C_{X^{(1)}}(u) = \emptyset = C_{X^{(1)}}(v)$. Let then $C_X(w) = C_X(v)$. By Lemma 3, we have $C_{X^{(1)}}(w) = \emptyset$, so $X^{(1)}$ has at most $k-1$ different contexts. It is then clear that $X^{(i)} = X^{(i+1)}$ for some $i \leq k$. □

The situation here is in stark contrast to the two-dimensional case, where very high ranks can be obtained even for SFTs [4]. Note that a one-dimensional sofic shift has zero topological entropy if and only if it is countable.

It was asked in [1] whether the rank of a countable two-dimensional SFT can be infinite. This problem was completely solved in [4], where the possible ranks were proven to be exactly those of Π^0_1 subshifts. However, to our knowledge, there does not exist an example where an infinite rank arises 'naturally'. We prove a weaker version of [4, Theorem 4.4] in Example 1 with such a natural example, using no direct computation.

Example 1. For all $k \in \mathbb{N}$ there exists a countable two-dimensional SFT X of rank at least $k\omega$.

Proof. We first assume $k = 1$, and then generalize the construction for all $k \in \mathbb{N}$.

The SFT X contains one infinite horizontal *dedicated line*, and on the line may be situated several *dedicated points*. The top and bottom halves are colored differently. Around a dedicated point one must have two (perhaps infinite) *diamonds*, a red and a blue one, whose left and right corners must hit the dedicated line. Conversely, exactly one point must be found inside a diamond (this is established by sending signals along the dedicated line). Two distinct diamonds may not overlap, unless one is completely inside the other (including a complete overlap). The insides of the diamonds are colored differently from their outsides.

From the top (bottom) corner of every red (blue) diamond, a *decrement signal* is sent to the right (left, respectively). Also, the top (bottom) corner of every red (blue) diamond must absorb one decrement signal traveling one tile above (below) it. The area between the line and a signal is colored differently from its complement. See Figure 1 for a clarifying picture.

We first show that X is countable. Indeed, for each $(n, m) \in \mathbb{N}^2$, if a configuration x of X contains a dedicated point with diamonds of sizes n and m, then there are at most $n+m-1$ dedicated points in x, since the size of the red (blue) diamonds decreases to the right (left). The number of ways to arrange these points and the surrounding diamonds is countable. One can also check that the

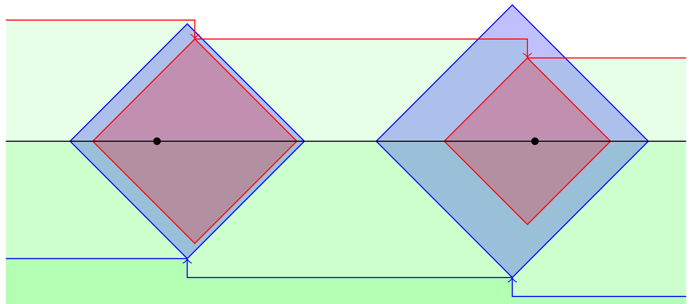

Fig. 1. The diamond construction.

number of exceptional points (ones containing, say, an infinite diamond or just signals) is countable.

Next, we show that $X^{(\omega)}$ is a nonempty set of finite rank. First, the isolated points of X are exactly those that contain at least one dedicated point with finite diamonds, and whose rightmost (leftmost) dedicated point is surrounded by a red (blue) diamond of size 1. In general, if $x \in X$ contains a dedicated point with finite diamonds, we say that x *has type* (n, m) if the blue (red) diamond of the rightmost (leftmost) dedicated point is of size n (m, respectively). It is then easy to see that for all $k \in \mathbb{N}$, the set $X^{(k)}$ will contain all of X, except for the points of type (n, m) with $n + m - 1 \leq k$. Then $X^{(\omega)}$ is nonempty, but will consist of only the exceptional points, and clearly $X^{(\omega+k)} = \emptyset$ for some finite k.

Let $k \in \mathbb{N}$. We generalize the above construction by allowing each configuration of X to contain k dedicated lines. We claim that X now has rank at least $k\omega$. Namely, the finite derivatives $X^{(n)}$ will remove the points where all k lines are present, and their types satisfy an inequality analogous to the above. Inductively, derivatives of the form $X^{(m\omega+n)}$ will remove points containing $k-m$ lines, and so $X^{(k\omega)}$ again consists of exceptional points only. □

4 Computability Aspects of Derivatives

In this section, we study the computational power of the kth derivative of a countable two-dimensional SFT, which turns out to possibly climb very high in the arithmetical hierarchy. We start with an upper bound, which we then reach with a construction.

Lemma 4. *Given a two-dimensional* Π^0_k *subshift* X *and a pattern* P*, it is* Π^0_{k+2} *whether* $P \sqsubset X^{(1)}$*.*

Proof. Given X and P, we have $P \sqsubset X^{(1)}$ iff $P \sqsubset X$ and for all $n \in \mathbb{N}$, there exist two distinct equal-sized extensions $Q_1, Q_2 \sqsubset X$ of P that agree on the $n \times n$-square around P. This is clearly Π^0_{k+2}. □

The following construction shows that the bound given by Lemma 4 on the complexity of kth derivatives of Π^0_1 subshifts is strict, and can be attained by a single countable SFT. In particular, it implies that Corollary 1 fails miserably in higher dimensions, since two-dimensional sofic shifts are Π^0_1. The rank of the subshift we build will be $\omega+k$ for some finite k, and with slight modifications, we could guarantee its ωth derivative to be recursive (with this exact construction, it is probably already recursive, but we are unable to perfectly verify this). This is interesting, since the complexity of the derivatives first rises without bound, but then collapses in the limit. We start with a definition, and a classical computability lemma.

Definition 2. *For $k \in \mathbb{N}$, denote by Φ_k the set of first-order arithmetical formulas with k free variables and only bounded quantifiers. For $k, l \in \mathbb{N}$, denote by ϕ^k_l the lth formula in Φ_k, ordered alphabetically.*

Lemma 5 (Lemma 2 in [5]). *Let $k \in \mathbb{N}$ and $\phi \in \Phi_{2k+1}$. Then there exists $\psi \in \Phi_{k+1}$, effectively computable from ϕ, such that*

$$\forall n_1 : \exists n_2 : \cdots \forall n_{2k-1} : \exists n_{2k} : \forall n_{2k+1} : \phi(n_1, \ldots, n_{2k+1})$$

is equivalent to

$$\exists^\infty n_1 : \exists^\infty n_2 : \cdots \exists^\infty n_k : \forall n_{k+1} : \psi(n_1, \ldots, n_{k+1}).$$

We denote $\psi = I(\phi)$ in the above lemma. With this result, we can transform alternating quantifiers into infinitary ones, and the application to derivatives is rather straightforward.

Theorem 1. *There exists a countable two-dimensional SFT X for which the problem whether $P \sqsubset X^{(k)}$ for a given pattern P is Π^0_{2k+1}-complete, for all $k \in \mathbb{N}$.*

Proof. We explicitly construct the subshift X. A typical configuration $x \in X$ consists of the *input*, a segment of the form $1^l 2^k$ extending to the right from the origin, and the *computation area*, a filled cone extending upwards from the input. The input also extends upwards in order to be accessible in the computation area. The rest of x is filled with 0's. See Figure 2 for a visualization.

Inside the computation area, a *ball* bounces between the walls of the cone, and in every sweep, one step of a counter machine M is simulated. This is simply to ensure that X is countable, as every nontrivial computation now has a starting point. The machine operates inside the cone, where the values of its counters are stored as the distances of special *counter symbols* from the vertical line going through the origin. A counter c with value n is represented by a length-n sequence of a symbol $\langle c \rangle$ extending right from the central column. Configurations of X thus correspond to computation histories of M in a concrete way.

Given the input (l, k), the machine M sequentally guesses k natural numbers $n_1, \ldots, n_k$ and then checks in an infinite loop that $I(\phi^{2k+1}_l)(n_1, \ldots, n_k, n_{k+1})$ holds for all $n_{k+1} \in \mathbb{N}$. If the check fails at some point, a tiling error is produced.

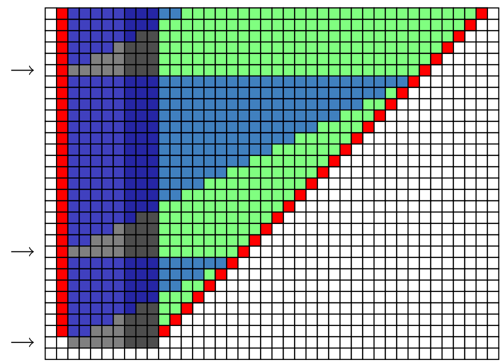

Fig. 2. A typical point of X. The transparent blue zone represents the bouncing ball. Counters and states, whose values are updated at the lines indicated by an arrow, are not shown here.

The guesses are also made using loops, so that for every $i \in [1, k]$, M starts enumerating all $n \in \mathbb{N}$, and at some point decides that n_i gets the value n. Thus larger guesses for the numbers take more time to compute, and since the computation is visible in the subshift, the input pattern 01^l2^k0 occurs in $X^{(k)}$ if and only if

$$\exists^\infty n_1 : \cdots \exists^\infty n_k : \forall n_{k+1} : I(\phi_l^{2k+1})(n_1, \ldots, n_{k+1})$$

is true. But by Lemma 5, this is equivalent to

$$\forall n_1 : \exists n_2 : \cdots \forall n_{2k-1} : \exists n_{2k} : \forall n_{2k+1} : \phi_l^{2k+1}(n_1, \ldots, n_{2k+1}),$$

and thus the subshift $X^{(k)}$ is Π^0_{2k+1}-hard in the sense of the claim. Since it reaches the upper bound given by Lemma 4, it is actually Π^0_{2k+1}-complete.

The only thing left to prove is the countability of X. For each input pattern 01^l2^k0, there are only a countable number of ways to complete it into a configuration, since the computation structure is forced, and only k nondeterministic moves are made by M. Configurations which do not contain input, or which contain an infinite input, are degenerate, and a simple case analysis shows that they, too, form a countable set. □

5 Subpattern Order

We now focus on the subpattern posets of countable multidimensional subshifts.

Proposition 2. *A countable d-dimensional subshift does not contain an infinite upward chain with relation to $\succ$.*

Proof. Suppose that a subshift $X \subset S^{\mathbb{Z}^d}$ contains a chain $(x^i)_{i \in \mathbb{N}}$ with $x^i \succ x^{i-1}$ properly for all $i \geq 1$. We show that X is uncountable by constructing an injective map $f : \{0,1\}^{\mathbb{N}} \to X$. First, note that if some x^i with $i \geq 1$ were periodic, then the relation $x^i \succ x^{i-1}$ could not be proper, and so all x^i with $i \geq 1$ are aperiodic.

Define $k_0 = 0$, and consider the one-cell pattern $P(\epsilon) = x^0_0$, which must occur somewhere in x^1. Since x^1 is aperiodic, there exist two distinct patterns $P(0)$ and $P(1)$ of size $(2k_1+1)^d$ for some k_1 occurring in x^1 with $P(\epsilon)$ in the center. In general, for all $n \in \mathbb{N}$, there exists $k_n \in \mathbb{N}$ such that for all words $w \in \{0,1\}^{n-1}$, we have two distinct patterns $P(w0)$ and $P(w1)$ of size $(2k_n+1)^d$ occurring in x^n and containing $P(w)$ in their centers.

For all $w \in \{0,1\}^{\mathbb{N}}$, we define $f(w) = \lim_{n \longrightarrow \infty} P(w_{[0,n-1]})$. Then f is a well-defined injection from $\{0,1\}^{\mathbb{N}}$ to X, and the claim is proved. □

This is a generalization of [1, Theorem 3.7], which states the result for two-dimensional SFTs. The above proof is more direct, but their method also directly generalizes to all countable subshifts. For antichains we have the following example.

Example 2. There exists a countable two-dimensional SFT with an infinite number of periodic points and an infinite antichain in its subpattern poset: take the SFT where horizontal and vertical lines form an infinite grid, and every rectangle is forced to be a square using a diagonal signal.

While no countable subshift contains an infinite ascending chain, and a simple countable SFT with an infinite antichain exists, the problem of descending chains turns out to be much more involved. We repeat the following, yet unsolved, conjecture from [1].

Conjecture 1. There is no countable two-dimensional SFT with an infinite downward chain for $\succ$.

We will not prove this conjecture, but provide a counterexample in the sofic case. For this result, we use a lemma for simplicity's sake, although the subshift we construct with it also has a direct implementation using signals.

Lemma 6 ([3]). *Let $X \subset S^{\mathbb{Z}}$ be a one-dimensional subshift such that $\mathcal{B}(X)$ is co-RE. Then $\{x \in S^{\mathbb{Z}^2} \mid \exists y \in X : \forall i,j \in \mathbb{Z} : x_{(i,j)} = y_j\}$ is a two-dimensional sofic shift.*

Theorem 2. *There exists a countable two-dimensional sofic shift with an infinite decreasing chain with relation to $\succ$.*

Proof. We will explicitly construct said sofic shift by defining several layers, each of which is a sofic shift, placing them over each other with some constraints and finally applying a block map that forgets almost all of the data. We start with the *powers of* 2 *shift*, the orbit closure of the point containing a horizontal line of 1's at the coordinate 2^n for all $n \geq 0$, and 0's everywhere else. It is sofic by Lemma 6. We may also put an arbitrary number of *dedicated points* on top of the lines.

We then present the *powers of* 2 *gadget*, which will control the way in which the dedicated points appear. For the time being, concentrate for simplicity's sake on the bottom line. If two points A, B lie on the line with distance d, we want a point C to appear $2d$ steps to the right of B, and if $d > 1$, another point D $\frac{d}{2}$ steps to the left of A. This is achieved with six signals emitted by each point, presented in Figure 3. Each point emits slope-1 and slope-$\frac{1}{2}$ signals to the left, a signal up and down, and slope-1 and slope-2 signals to the right. The downward signal is destroyed when the other signals hit it, and they must be correctly matched on each side. If $d = 1$ in the above situation, then A only emits the rightward signals, and creates a *forbidden zone* on its left. The zone continues infinitely upward and to the left, and no dedicated points may be situated inside it. Two signals of different types (dashed and dotted lines in the figure) may always cross each other and the horizontal lines, and a dashed line may cross a downward signal that has already encountered another dashed line. The grey area in the figure is another forbidden zone. It consists of the rightmost $\frac{2}{3}$ of the space between two points and continues infinitely upwards. Also, no forbidden zone may appear below the lowest line.

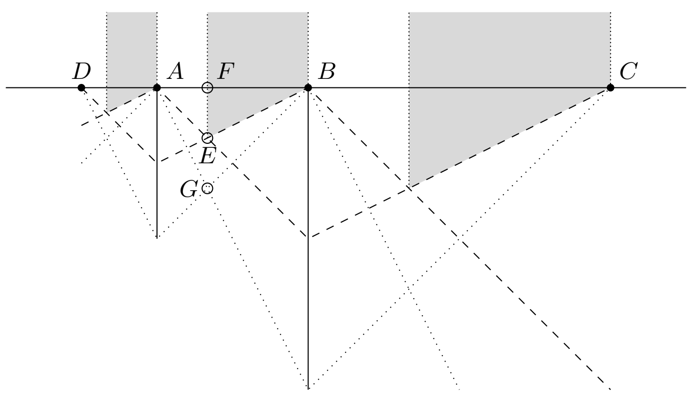

Fig. 3. The signals of the powers of 2 gadget.

We now drop the single line hypothesis, and require that all signals attempting to enter a forbidden zone are destroyed. On lines other than the first one, when an upward signal emitted by a point reaches the first horizontal line, it checks with a slope-1 *check signal* (line segment AB in Figure 4) whether the (perhaps infinite) rectangle formed by the two lines and the left border of the next borbidden zone is at least as wide as it is high. If this is the case, it forces two consecutive points to appear one step to the right, and if not, the whole interval gets forbidden infinitely upwards. Figure 4 shows an interval in which the check succeeds, and the two points are forced to appear in the upper line.

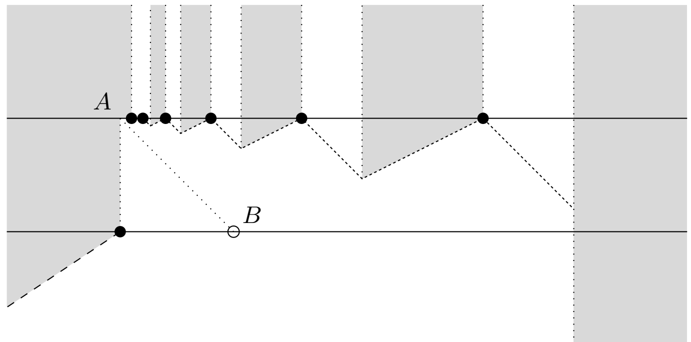

Fig. 4. A successful check, which forces two points to appear on the upper line. The powers of 2 gadget, in turn, forces more points to appear, until the interval ends. Not all signals are shown.

We then claim that signals emitted by points from different lines do not destructively interact with each other. Namely, if the distance of A and B in Figure 3 is d, then the maximum distance between two points on the line AF is at most $\frac{1}{2}\frac{1}{3}d = \frac{d}{6}$. Then the length of the downward signals emitted by these points is at most $\frac{d}{3}$, which is the length of the segment EF. Thus no signals emitted by these points propagate below the point E. Also, all emitted leftward signals are caught by downward ones. Finally, only the dotted signals and the downward signals that have already encountered a dashed signal are able to reach the dashed line AE, but these will just intersect without interaction, and no signal can reach the dotted line AG.

Consider then the sofic shift $X \subset \{0,1\}^{\mathbb{Z}^2}$ obtained by mapping every dedicated point to 1, and the rest to 0. We first prove that it is countable. Suppose first that a configuration of X contains points in only one horizontal line. There might be forbidden zones extending infinitely downward from the line. On intervals between two such forbidden zones, no points are ever seen, since the check signals always fail. On an infinite interval free of forbidden zones, we either have an infinite sequence of points determined by the position of the leftmost one (which always exists), or just one lone point.

If points are seen on two lines, we know exactly the positions of all lines. Consider the leftmost point on the lowest occupied line. Because of the forbidden zone it creates on its left, it is the overall leftmost point. If there is another point on the lowest line, then all points on the line are determined by the gadget. Next, all points of the line above are determined by the upward signals and the gadget. By induction, all the points of the configuration are determined. If there are no other points on the lowest line, the next line is determined by the gadget, and we repeat the above argument. So all in all, a configuration is determined by the position of the lowest and leftmost point, and whether there are other points on the same line or some other line, and X is countable.

Finally, we show that X contains an infinite downward chain with relation to $\succ$. Consider the point x_1 containing infinitely many points on the lowest line. The gadgets force x_1 to consist of vertical 'stripes' of exponential width that begin with a 1 on the lowest line, followed by a prefix of the lowest line on the second one (and the patterns on the higher lines determined by it), and a forbidden zone. See Figure 4 to better visualize this. The prefixes become arbitrarily long as the width of the stripes increases. In the orbit closure of x_1, we thus find the point x_2 containing a 1 in the lowest line followed by the stripe pattern on the first line of x_1 lifted to the second one, which creates a similar prefix pattern on the third line. We repeat the argument to find x_3, on which the striped pattern lies on the third line, and inductively we obtain the chain $(x_i)_{i \geq 1}$. Clearly $x_i \succ x_{i+1}$ holds properly for all i, since x_{i+1} was chosen from the orbit closure of x_i, and the claim is proved. □

We now consider the structure of finite subpattern posets. We first show that they do not capture the whole class of finite posets: in fact, no nontrivial meet-semilattice can occur as a subpattern poset. This is an easy consequence of the results in [1].

Proposition 3. *A nontrivial countable two-dimensional SFT contains at least two periodic points.*

Proof. Let X be a countable two-dimensional SFT with exactly one periodic point $x \in X$, which must be unary: $x_{\boldsymbol{n}} = 0$ for all $\boldsymbol{n} \in \mathbb{Z}^2$. If X was nontrivial, [1] would imply the existence of a point $y \in X$ with exactly one direction of periodicity. We have two possibilities.

1. There exists $n \in \mathbb{N}$ such that every $(n \times n)$-square of y contains a non-0 symbol. But then x is not the only minimal point, and since X contains only periodic minimal points, this is a contradiction.
2. For all $n \in \mathbb{N}$, an $(n \times n)$-square of 0's occurs in y. Since y is periodic in one direction, this implies that y contains arbitrarily thick stripes of 0's, and since $y \neq x$, there are non-0 symbols between some of them. But if the thickness of the stripes is larger that the window size of X, we can build an uncountable number of points using the stripes, a contradiction.

Thus X must be trivial. □

Corollary 2. *The subpattern poset of a countable two-dimensional SFT cannot be a nontrivial meet-semilattice, that is, a poset in which every two elements have a greatest lower bound.*

However, we have the following embedding result, for which we do not know an essentially simpler proof.

Proposition 4. *All finite posets can be order-embedded in the subpattern poset of some countable two-dimensional SFT. Furthermore, the subpattern poset itself can be made finite.*

Proof. Let $(S, \geq)$ be a finite poset, and for all $x \in S$, define

$$r(x) = \max\{n \in \mathbb{N} \mid \exists y_1, \ldots, y_n \in S : x = y_1 > \ldots > y_n\} - 1.$$

This is one less than the maximal length of a descending chain beginning from x. Define $B = r^{-1}(0)$, the set of minimal elements of S. Define also

$$p(x) = \{y \in S \mid x > y, \nexists z \in S : x > z > y\},$$

the set of immediate predecessors of x. We also inductively define $k(x) = 1$ for all $x \in B$, and $k(x) = 1 + \sum_{y \in p(x)} k(y)$ for $x \notin B$. This is an auxiliary 'height' function we need in our construction.

We build a two-dimensional SFT X in whose subpattern poset S can be order-embedded via $f : S \to X$. First, for each $x \in B$, X contains a unary point $f(x)$. Let then $x \in S - B$. We assume that $f(y)$ has already been defined for all $y \in S$ with $r(y) < r(x)$, using the construction we are about to present if $y \notin B$.

The point $f(x)$ contains a horizontal *dedicated half-line*, starting from the origin and extending right. Below the line there is a vertical sequence of *ruler rectangles*, starting with one of size 1×1 below the origin. The nth rectangle, in general, has size $\phi(n, r(x)) \times n$, where $\phi(n, 1) = n$ and $\phi(n, r) = \sum_{i=1}^{n} \phi(i, r-1)$. If $r(x) = 1$, this is achieved with a diagonal signal forcing the rectangles to be squares, and in general by stacking a sequence of lower-rank rectangles inside the large ones. See Figure 5 for a visualization, and note that $\phi(n, r)$ is a polynomial of degree r in n.

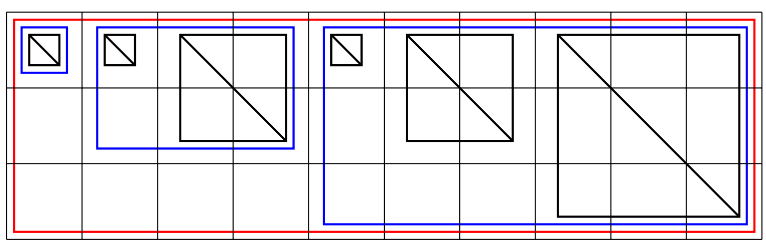

Fig. 5. A ruler rectangle of size $\phi(3, 3) \times 3$.

Above the dedicated half-line we put $|p(x)|$ sequences of *data rectangles*, one for each $y \in p(x)$, stacked on top of each other. The data rectangles of y contain patterns from the configuration $f(y)$. The left and right ends of the data rectangles are forced to align with those of the ruler rectangles, and the heights increase by the respective $k(y)$ every step. If $y \in B$, the height of the nth rectangle of the sequence is n, and it will be filled with the unary pattern of $f(y)$. If $y \notin B$, it is created using this construction, and consists of a finite number of horizontal sequences of rectangles whose total width increases by a constant $k(y)$ every step. For each $n \in \mathbb{N}$, the nth data rectangle of the sequence corresponding to y has height $n \cdot k(y)$, and it is forced to contain a pattern of $f(y)$ aligned with the right border of the rectangle as in Figure 6. The linear growth is easily forced by SFT rules. In the construction, each $f(x)$ will extend the alphabet with completely new symbols (apart from the ones used to simulate the $f(y)$), and each region in the construction will have a different unary background symbol to differentiate them from each other.

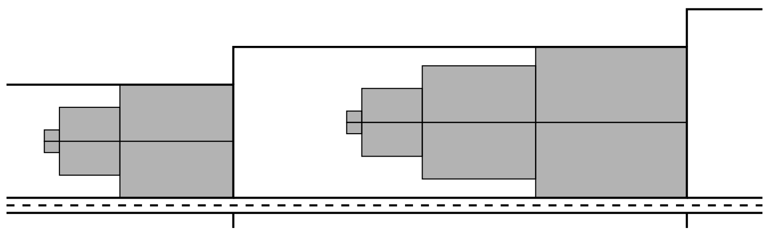

Fig. 6. The alignment of $f(y)$ in the data rectangles. The ruler rectangles lie below the horizontal line.

That the construction can be done using only SFT rules is clear, as is the fact that f becomes an order-embedding of $(S, \geq)$ into $(X, \succcurlyeq)$. Furthermore, a simple case analysis shows that X is countable and its subpattern poset is finite. □

6 Conclusions

In this paper, we have presented several constructions related to the computational and topological structure of countable sofic and finite type subshifts. In Theorem 1, we presented a single countable SFT whose kth derivative is Π^0_{2k+1}-complete for all k, the highest possible among the kth derivatives of all Π^0_1 subshifts. We also studied the subpattern posets of countable subshifts, our main result being Theorem 2. The theorem is much more interesting in conjunction with Conjecture 1 than in itself, since if the conjecture is true, the sofic counterexample may be helpful in finding a proof for it.

We have achieved the exact maximal computational strength of the kth derivative of a countable SFT, but it would be interesting to see what happens in the first limit ordinal and beyond.

Question 1. Let λ be any computable ordinal. What is the maximal computational power of the λth derivative of a countable SFT?

In particular, does there exist a countable SFT whose ωth derivative is Π^0_k-hard for all k? Can we reach some levels of the analytical hierarchy?

We also have some interesting open questions regarding the subpattern posets of subshifts. We have shown here that all finite posets can be order-embedded in the subpattern poset of some countable SFT. For which infinite posets does this hold? As a more concrete question, let $\mathcal{F} = \{A \subset \mathbb{N} \mid |A| < \infty\}$, and consider the partial order $<$ defined on $\mathcal{F}$ by $A < B$ iff $B \subset A$.

Question 2. Let $\mathcal{F}' \subset \mathcal{F}$ be computable. Can the induced poset $(\mathcal{F}', <)$ be order-embedded in the subpattern poset of some countable SFT or sofic shift?

Example 2 shows that the poset $(\{\{n\} \mid n \in \mathbb{N}\}, <)$ can be order-embedded in a countable SFT. Also, Theorem 2 shows that the poset $(\{\{1, \dots, n\} \mid n \in \mathbb{N}\}, <)$ can be order-embedded in a countable sofic shift, while Conjecture 1 claims that this is impossible for an SFT. Note that the above question does not clash with Proposition 2, since we have inverted the subset relation.

References

1. 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.
2. Robert Berger. The undecidability of the domino problem. *Mem. Amer. Math. Soc. No.*, 66:72, 1966.
3. 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.
4. Emmanuel Jeandel and Pascal Vanier. Π_1^0 sets and tilings. In *Theory and Applications of Models of Computation (TAMC)*, volume 6648 of *Lecture Notes in Computer Science*, pages 230–239, 2011.
5. G. Kreisel, J. Shoenfield, and Hao Wang. Number theoretic concepts and recursive well-orderings. *Arch. Math. Logik Grundlagenforsch.*, 5:42–64, 1960.
6. Douglas Lind and Brian Marcus. *An introduction to symbolic dynamics and coding.* Cambridge University Press, Cambridge, 1995.
7. Raphael M. Robinson. Undecidability and nonperiodicity for tilings of the plane. *Invent. Math.*, 12:177–209, 1971.