

Computational Aspects of Cellular Automata on Countable Sofic Shifts^{*}

Ville Salo¹ and Ilkka Törmä²

¹ University of Turku, Finland,
Turku Centre for Computer Science, Finland,
`vosalo@utu.fi`

² University of Turku, Finland,
`iatorm@utu.fi`

Abstract. We investigate the computational properties of cellular automata on countable (equivalently, zero entropy) sofic shifts with an emphasis on nilpotency, periodicity, and asymptotic behavior. As a tool for proving decidability results, we prove the Starfleet Lemma, which is of independent interest. We present computational results including the decidability of nilpotency and periodicity, the undecidability of stability of the limit set, and the existence of a Π_1^0 -complete limit set and a Σ_3^0 -complete asymptotic set.

Keywords: cellular automata, subshifts, nontransitivity, undecidability

1 Introduction

Symbolic dynamics is often viewed from the perspective of coding, and then infinite words represent information flowing to the left in the orbit given by the shift action. From this perspective, positive dynamical entropy is crucial, and transitivity of the subshift is a very natural assumption. In the case of a sofic shift, transitivity essentially means that the behavior of the medium stays the same no matter what has been sent sofar. However, from a purely mathematical perspective, zero entropy, or countable, sofic shifts are quite interesting, consisting of a simple periodic background pattern broken by a finite number of local disturbances. They lie in the other extremity of the spectrum of sofic shifts, where almost no information can be sent. Still, as far as we know, the conjugacy problem is still open even in the case of countable SFTs.

Our point of view is that of investigating the computational capabilities of cellular automata running on a countable sofic shift. Such a cellular automaton can be thought of as a kind of counter machine, with the distances between the local disturbances representing the counter values. As a tool for proving decidability results for such systems, we prove the Starfleet Lemma, which is of independent interest.

^{*} Research supported by the Academy of Finland Grant 131558

We present a variety of computational results: We prove the decidability of nilpotency and periodicity and that limit sets and asymptotic sets always lie in Π_1^0 and Σ_3^0 in the arithmetical hierarchy, respectively. Further, we prove that limit sets and asymptotic sets in fact capture the respective classes by finding complete sets for them. The examples proving completeness are based on a rather direct simulation of counter machines by CA, and we also obtain many undecidability results using this idea.

2 Definitions

Let S be a finite set, called the *state set* or *alphabet*. The set $S^{\mathbb{Z}}$ of bi-infinite state sequences, or *configurations*, is called the *full shift on S* . If $x \in S^{\mathbb{Z}}$, then we denote by x_i the i th coordinate of x , and we adopt the shorthand notation $x_{[i,j]} = x_i x_{i+1} \dots x_j$. If $w \in S^*$, we denote $w \sqsubset x$, and say that w occurs in x , if $w = x_{[i, i+|w|-1]}$ for some i . For words $t, u, v, w \in S^*$, the notation ${}^\infty t u v w {}^\infty$ has the intuitive meaning, with the word v starting at coordinate 0. The dot can be omitted if the position of the words is irrelevant. Two elements $x, y \in S^{\mathbb{Z}}$ are *right (left) asymptotic* if $x_i = y_i$ for all sufficiently large (small) i .

We define a metric d on the full shift by $d(x, y) = \inf\{2^{-n} \mid x_{[-n,n]} = y_{[-n,n]}\}$. The topology defined by d makes $S^{\mathbb{Z}}$ a compact metric space. We define the *shift map* $\sigma : S^{\mathbb{Z}} \rightarrow S^{\mathbb{Z}}$ by $\sigma(x)_i = x_{i+1}$. Clearly σ is a homeomorphism from the full shift to itself.

A *subshift* is a closed subset X of the full shift with the property $\sigma(X) = X$. Alternatively, a subshift is defined by a set $F \subset S^*$ of *forbidden words* as the set $\mathcal{X}_F = \{x \in S^{\mathbb{Z}} \mid \forall w \in F : w \not\sqsubset x\}$. If F is finite, then \mathcal{X}_F is of *finite type*, and if F is regular, \mathcal{X}_F is *sofic*. We define $\mathcal{B}_k(X) = \{w \in S^k \mid \exists x \in X : w \sqsubset x\}$ as the set of words of length k occurring in X , and define $\mathcal{B}(X) = \bigcup_{k \in \mathbb{N}} \mathcal{B}_k(X)$. The language of a subshift determines it [1], so we may also denote $X = \mathcal{B}^{-1}(L)$, if $\mathcal{B}(X)$ is the set of factors of L . We denote by σ_X the restriction of σ to X .

Given a directed graph $G = (V, E)$, we define its *edge shift*, a subshift of $E^{\mathbb{Z}}$, as the set of bi-infinite paths on its edges. Every sofic shift X is equal to the set of labels of bi-infinite paths on a directed graph with labeled edges, and the edge shift of the (essentially unique) deterministic graph with the least number of edges is called the *minimal Shannon cover* of X .

Let $X \subset S^{\mathbb{Z}}$ be a subshift. A *cellular automaton* is a continuous function $c : X \rightarrow X$ with the property $c \circ \sigma_X = \sigma_X \circ c$. Equivalently, cellular automata can be defined by *local functions* $C : \mathcal{B}_{2r+1}(X) \rightarrow \mathcal{B}_1(X)$ for some $r \in \mathbb{N}$ such that $c(x)_i = C(x_{[i-r, i+r]})$. The smallest such r is called the *radius* of c . The *limit set* of c is $\bigcap_{n \in \mathbb{N}} c^n(X)$, and c is *stable* if the limit set is equal to some $c^n(X)$. A state 0 is *quiescent* for c if $c({}^\infty 0 {}^\infty) = {}^\infty 0 {}^\infty$. A CA c on X is *weakly nilpotent* if for all $x \in X$ we have an n such that $c^n(x) = {}^\infty 0 {}^\infty$, and *nilpotent* if the n are uniformly bounded. Also, c is *weakly periodic* if for all $x \in X$ we have an n such that $c^n(x) = x$, and *periodic* if the n are uniformly bounded. A *spaceship* of c is a configuration $x \in X$ such that $c^n(x) = \sigma^i(x)$ for some $n \in \mathbb{N}$ and $i \in \mathbb{Z}$. We say the spaceship is *nontrivial* if it is not spatially periodic.

Let $k \in \mathbb{N}$. A k -counter machine is a quintuple $M = (\Sigma, k, \delta, q_0, q_f)$, where Σ is a finite state set, $q_0, q_f \in \Sigma$ the initial and final states and $\delta : (\Sigma - \{q_f\}) \rightarrow [1, k] \times \Sigma^2 \cup [1, k] \times \{\uparrow, \downarrow\} \times \Sigma$ the transition function. A configuration of M is an element of $\mathbb{N}^k \times \Sigma$, with the interpretation of (n_1, \dots, n_k, s) being that the machine is in state s with counter values n_1, \dots, n_k . The machine operates in steps as directed by δ . The case $\delta(s) = (i, r, t) \in [1, k] \times \Sigma^2$ means that if M is in state s , then it will change its state to r if $n_i = 0$, and otherwise to t . The case $\delta(s) = (i, d, r) \in [1, k] \times \{\uparrow, \downarrow\} \times \Sigma$ means that M increments or decrements n_i by 1, depending on d (but never decrementing it below 0), and then goes to state r . The machine M is *reversible* if every configuration can be reached in one step from at most one other configuration.

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_n^0 if it is equivalent to a formula of the form $\exists k : \psi$ where ψ is Π_{n-1}^0 , and ϕ is Π_n^0 , if it is equivalent to a formula of the form $\forall k : \psi$ where ψ is Σ_{n-1}^0 . This classification is called the *arithmetical hierarchy*. A subset X of \mathbb{N} is Σ_n^0 or Π_n^0 , if $X = \{x \in \mathbb{N} \mid \phi(x)\}$ for some ϕ with the corresponding classification. It is known that the class of sets that are both Σ_1^0 and Π_1^0 coincides with the class of recursive sets.

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 $f : \mathbb{N} \rightarrow \mathbb{N}$ 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.

3 Basics

We remark here, and it would not be hard to prove either, that a sofic shift is countable if and only if the cycles of its Shannon cover are disjoint. Thus the configurations of a countable sofic shift consist of arbitrarily long periodic patterns (corresponding to the cycles of its Shannon cover), separated by finite period-breaking patterns (the transitions between cycles). In particular, there are a finite number of periodic configurations and a global upper bound on the number of transitions in a single configuration. Furthermore, a sofic shift is easily seen to be countable if and only if it has zero topological entropy with respect to the shift map, but we will not elaborate on this here.

We now prove a fundamental result concerning the dynamics of cellular automata on countable sofic shifts, which is very useful for decidability results. For this, we need a canonical representation for configurations of the shift, given by the following lemma.

Lemma 1. *Let X be a countable sofic shift. Then there exists a finite set T of tuples of words in $\mathcal{B}(X)$ such that every point $x \in X$ is representable as $x = {}^\infty u_0 v_1 u_1^{n_1} \cdots u_{m-1}^{n_{m-1}} v_m u_m^\infty$ for a unique $t = (u_0, \dots, u_m, v_1, \dots, v_m) \in T$.*

Proof. Let p be a common period for all the periodic points of X , and let $U = \{u \in \mathcal{B}_p(X) \mid {}^\infty u \in X\}$. For all $x \in X$ we add a tuple $T(x) \in T$ corresponding

to x as follows. First, there is a unique $u_0 \in U$ such that x is of the form ${}^\infty u_0 a y$ for some $a \in S - \{(u_0)_1\}$ and $y \in S^{\mathbb{N}}$. Let the coordinate of a be i_1 . We inductively define i_k as the first coordinate greater than i_{k-1} such that $u_k := x_{[i_k-p, i_k-1]} \in U$ and $x_{i_k} \neq (u_k)_1$. When such a coordinate can no longer be found, we are in the right periodic tail of x . Finally, the last u_m is chosen as we would choose u_0 for x^R , the reversal of x . The v_k are then the words occurring in between the u_k -periodic patterns, and we set $T(x) = (u_0, \dots, u_m, v_1, \dots, v_m)$.

This representation is now unique, since the parsing process is deterministic from left to right. \square

We continue to write $T(x)$ for the unique tuple $(u_0, \dots, u_m, v_1, \dots, v_m)$ given by the previous lemma for $x \in X$, and $M(x)$ for the tuple (n_1, \dots, n_{m-1}) .

Lemma 2 (Starfleet Lemma). *If c is a cellular automaton on a countable sofic shift X , then for all $x \in X$ there exists a tuple $t = (r_0, \dots, r_l, s_1, \dots, s_l) \in \mathcal{B}(X)^{2l+1}$ (not necessarily in T) such that*

- for all $N \in \mathbb{N}$, there exist $n \in \mathbb{N}$ and $n_1, \dots, n_{l-1} \geq N$ such that

$$c^n(x) = {}^\infty r_0 s_1 r_1^{n_1} s_2 \cdots s_{l-1} r_{l-1}^{n_{l-1}} s_l r_l^\infty$$

modulo a power of the shift, and

- for all $i \in [1, l]$, the configuration ${}^\infty r_{i-1} s_i r_i^\infty$ is a nontrivial spaceship for c .

We call a tuple t as above a *starfleet* for x .

Proof. Let $x \in X$. Define $U = \{(T(c^n(x)), M(c^n(x))) \mid n \in \mathbb{N}\}$, and for all $t \in T$, let $M_x(t)$ be the set $\{t' \mid (t, t') \in U\}$. Let $t = (u_0, \dots, u_m, v_1, \dots, v_m) \in T$, and let $C_x(t)$ be the closure of the set $M_x(t)$ in $\bar{\mathbb{N}}^{m-1}$, where $\bar{\mathbb{N}} = \mathbb{N} \cup \{\infty\}$ is the one-point compactification of \mathbb{N} .

If $C_x(t)$ is finite for all $t \in T$, then x must evolve into a spaceship, and the requirements of the lemma clearly hold. If this is not the case, let $t = (u_0, \dots, u_m, v_1, \dots, v_m) \in T$ and $t' = (n_1, \dots, n_{m-1}) \in C_x(t)$ be such that t' contains a maximal number of infinite coordinates. Let (k_1, \dots, k_{l-1}) be the infinite coordinates of t' , let $k_0 = 0$ and $k_l = m$, and consider the words

$$w_i = v_{k_{i-1}+1} u_{k_{i-1}+1}^{n_{k_{i-1}+1}} \cdots u_{k_i-1}^{n_{k_i-1}} v_{k_i},$$

for $i = 1, \dots, l$.

Clearly, by the maximality of l , each ${}^\infty u_{k_{i-1}} w_i u_{k_i}^\infty$ will evolve into a nontrivial spaceship ${}^\infty r_{i-1} s_i r_i^\infty$ in some p steps. Let x_1, x_2, \dots be the subsequence of the orbit of x such that $T(x_i) = t$ and $M(x_i)$ converges to t' . Then, the subsequence $x'_i = c^p(x_i)$ shows that $(r_0, \dots, r_l, s_1, \dots, s_l)$ has the required properties. \square

For CA on a mixing sofic shift, it is well-known that injectivity implies surjectivity [1]. However, in general, surjectivity and injectivity do not imply each other even in a countable SFT. The following example shows this, and will also turn out useful later on.

Example 1. Let $X = \mathcal{B}^{-1}(0^*1^*2^*)$ and let f be a cellular automaton that expands the pattern of 1's to both directions by one. Clearly, such a CA is injective on X , but it is not surjective. Similarly, the CA g which removes a single 1 from each end is surjective, but not injective (in fact, not even preinjective, that is, two left and right asymptotic configurations can have the same image).

Of course, a CA that is surjective and injective is still reversible. Furthermore, the usual decidability results hold also in the general case: Surjectivity, injectivity and reversibility of cellular automata are decidable properties on all sofic shifts, which is easily seen using standard techniques of symbolic dynamics.

4 Nilpotency and Periodicity

In this section, we prove that nilpotency and periodicity are decidable properties for cellular automata on countable sofic shifts. The decidability of nilpotency is a rather direct application of the Starfleet Lemma and the following result.

Lemma 3. *Let X be a subshift. Then a cellular automaton c on X is nilpotent if and only if it is weakly nilpotent.*

Proposition 1. *For cellular automata on countable sofic shifts, nilpotency is decidable.*

Proof. We will prove that one of the following cases holds for a CA c on a countable sofic shift X : either c is nilpotent, it is not nilpotent on periodic configurations, or X contains a nontrivial spaceship for c . Since the latter cases imply non-nilpotency and all three are semi-decidable, this proves the proposition.

Assume that c is not nilpotent, but is nilpotent on the finitely many periodic configurations of X , and let ${}^\infty 0^\infty$ be their limit. We will show that X contains a nontrivial spaceship for c . By Lemma 3, c is not even weakly nilpotent, so there exists a configuration $x \in X$ with $c^n(x) \neq {}^\infty 0^\infty$ for all n . Let $(u_0, \dots, u_m, v_1, \dots, v_m) \in \mathcal{B}(X)^{2m+1}$ be a starfleet for x , which must be nontrivial: $v_i \notin 0^*$ for some i . But now ${}^\infty u_{i-1} v_i u_i^\infty$ is a nontrivial spaceship for c . \square

Note that this is not true in general: It is proved in [2] that on the full shift, nilpotency is undecidable for one-dimensional cellular automata. In fact, there is an analogue of Rice's theorem on the limit sets of CA on the full shift [3], implying that all nontrivial properties of the limit set are undecidable. An interesting question is whether some weaker form of Rice's theorem holds on countable sofic shifts, since Theorem 2 implies that individual limit sets can have very complicated structure.

The decidability of periodicity also follows from the Starfleet Lemma.

Lemma 4. *A cellular automaton c on an arbitrary subshift $X \subset S^{\mathbb{Z}}$ is periodic iff it is weakly periodic.*

Proof. We only need to prove that weak periodicity implies periodicity, so let c be weakly periodic. Then its *two-way trace subshift*

$$\{y \in S^{\mathbb{Z}} \mid \exists (x^i)_{i \in \mathbb{Z}} \in X^{\mathbb{Z}} : \forall i \in \mathbb{Z} : x^{i+1} = c(x^i) \wedge y_i = x_0^i\}$$

contains only periodic points, and thus must be finite [4]. Then c is periodic with period the least common multiple of this finite set of periods. \square

Proposition 2. *For cellular automata on countable sofic shifts, periodicity is decidable.*

Proof. We will prove that one of the following cases holds for a CA c on a countable sofic shift X : either c is periodic, c is noninjective, or X contains a nonperiodic spaceship for c . Since the latter cases imply nonperiodicity and all three are semi-decidable, this proves the proposition.

Suppose on the contrary that c is nonperiodic (hence not weakly periodic by Lemma 4) and injective, and all spaceships are periodic. Now let $x \in X$ be arbitrary, and let $(u_0, \dots, u_m, v_1, \dots, v_m)$ be a starfleet for it. Since the configurations ${}^\infty u_{i-1} v_i u_i^\infty$ are spaceships, they must be periodic, so $m = 1$. Now $c^n(x)$ is a spaceship for large enough n , and by injectivity of c , x is a spaceship itself, hence periodic. This is a contradiction, since c was not weakly periodic. \square

We remark that the decidability proofs of nilpotency and periodicity both involve three semi-decidable properties, of which every cellular automaton must possess at least one. In addition to proving the propositions, these triplets are interesting in themselves, shedding light on the possible dynamics of cellular automata on countable sofic shifts.

5 Computation and Limit Sets

In this section, we focus on the limit sets of cellular automata and their computational properties. We start by defining a natural tool for proving undecidability results, namely, a cellular automaton which simulates a counter machine. Since the limit set is exactly the set of configurations with an infinite chain of preimages, it is natural to run a counter machine in reverse.

Theorem 1 ([5]). *For any k -counter machine $M = (\Sigma, k, \delta, q_0, q_f)$ there exists a reversible $(2k + 2)$ -counter machine $M' = (\Sigma', 2k + 2, \delta', p_0, p_f)$ such that*

$$(m_1, \dots, m_k, q_0) \Rightarrow_{M'}^* (n_1, \dots, n_k, q_f)$$

if and only if

$$(m_1, \dots, m_k, 0, 0, 0, \dots, 0, p_0) \Rightarrow_{M'}^* (m_1, \dots, m_k, 0, 0, n_1, \dots, n_k, p_f).$$

Let $M = (\Sigma, k, \delta, q_0, q_f)$ be a counter machine. We construct a countable SFT X_M and a CA c_M on X_M simulating M in a specific way.

Denote $\Sigma' = \Sigma \times \{\leftarrow, \rightarrow\}$. The subshift X_M is defined as the subset of

$$\mathcal{B}^{-1}(a^* \Sigma' b^*) \times \prod_{i \in [1, k] \cup \{\#\}} \mathcal{B}^{-1}(a^* i b^*)$$

where none of the symbols $[1, k]$ can occur to the left of $\#$, and not all symbols $[1, k] \cup \#$ can occur on the same side of $s \in \Sigma'$. A configuration of X_M is *good* if it contains some $s' \in \Sigma'$, the symbol $\#$ and all symbols $i \in [1, k]$.

The automaton c_M simulates M as follows. The value of a counter i is coded as its distance from the symbol $\#$. Each state $(s, \rightarrow) \in \Sigma'$ starts on $\#$, where the automaton can immediately check whether any counter has value 0. If needed, the state travels to the right until it encounters the counter symbol. It shifts the symbol one step to the desired direction, flips its own arrow, returns to the $\#$ and enters a new state as determined by δ . If x is a configuration of M , we denote by $X_M(x)$ the corresponding configuration of X_M .

If M is reversible, we may add a *direction bit* to each state and have c_M simulate M either forward or backward, flipping the bit whenever an image or preimage does not exist. We denote this augmented automaton by \bar{c}_M . Note that this is possible since a CA can, in one step, check which counters have value 0, when the state lies on the $\#$ -symbol, even if the reverse function cannot actually be computed by a counter machine. The notation $X_M^{\rightarrow}(x)$ stands for the forward and $X_M^{\leftarrow}(x)$ for the backward configuration.

When this basic construction is used in proofs, it will be modified as needed.

It is easy to see that the language of a limit set of a CA is always Π_1^0 , and the following complementary theorem holds.

Theorem 2. *There exists CA $c : X \rightarrow X$ on a countable SFT X such that the limit set of c is Π_1^0 -complete.*

Proof. Let M be a reversible counter machine whose language is Σ_1^0 -complete with the property that the initial configuration $(n, 0, \dots, 0, q_0)$ has no preimage for any n , and if $n \notin L(M)$, then M does not halt when started on this configuration. Such a machine can be constructed using Theorem 1 and its proof. We modify \bar{c}_M to start filling the space with a new spreading state when an accepting state is reached in a forward state.

Now we claim that the forward CA configuration corresponding to

$$x_n = (n, 0, \dots, 0, q_0)$$

is in the limit set of \bar{c}_M if and only if $n \notin L(M)$. It is clear that if $n \notin L(M)$, then an infinite chain of preimages is obtained by running M backwards towards the initial state. Consider the case $n \in L(M)$. Since M is reversible and \bar{c}_M simulates it on good configurations, the only possible preimage chain for $X_M^{\rightarrow}(x_n)$ under \bar{c}_M could be obtained by running M back and forth between the initial and accepting states. But this is not possible, since we introduced a spreading state, and adding a spreading state clearly cannot add any new preimages for good configurations. \square

Theorem 3. *It is undecidable whether a CA on a countable SFT is stable.*

Proof. Given a reversible counter machine M , it easily follows from Theorem 1 that it is undecidable whether M halts on $x = (0, \dots, 0, q_0)$. Modify \bar{c}_M so that $X_M^{\leftarrow}(x)$ becomes a fixed point, but no other configuration is affected. We now claim that the modified \bar{c}_M is stable iff M halts from x .

Suppose first that M does not halt from x . Then every point $\bar{c}_M^n(X_M^{\rightarrow}(x))$ has a single preimage chain of length n ending in $X_M^{\rightarrow}(x)$. Thus \bar{c}_M is unstable.

Suppose then that M halts from x . Then $X_M^{\rightarrow}(x)$ eventually evolves into a final state, from which \bar{c}_M will run backwards to $X_M^{\leftarrow}(x)$, which is a fixed point. Thus there are only a finite number of points of the form $\bar{c}_M^n(X_M^{\rightarrow}(x))$ without an infinite preimage chain. Furthermore, since every point except $X_M^{\rightarrow}(x)$ has a preimage, every point of X_M not of the form $\bar{c}_M^n(X_M^{\rightarrow}(x))$ has an infinite chain of preimages. Thus \bar{c}_M is stable. \square

6 Asymptotic Sets

The limit set is not the only notion corresponding to ‘where points eventually go’. Another such concept, studied at least in [6], is the asymptotic set. We show that this idea makes sense also in the countable sofic case, and is in some sense stronger than the concept of limit set.

Definition 1. *The asymptotic set of CA $c : X \rightarrow X$ is*

$$\bigcup_{x \in X} \bigcap_{J \in \mathbb{N}} \overline{\{c^n(x) \mid n \geq J\}}$$

A configuration $x \in X$ lies in the asymptotic set iff there exists another configuration $y \in X$ and a subsequence of the orbit $(c^n(y))_{n \in \mathbb{N}}$ which converges to x . Note that the asymptotic set contains all temporally periodic points of c , but not necessarily all the spaceships, unlike the limit set. Asymptotic sets have much stronger computational capabilities than limit sets.

Lemma 5. *The asymptotic set Y of a CA c on a countable sofic shift X is Σ_3^0 .*

Proof. Given a word w , it is clearly (by its form) in Σ_3^0 to check that

$$\exists x \in X : \forall n : \exists m > n : c^m(x)_{[1, |w|]} = w,$$

which is equivalent to $w \sqsubset Y$. Note that the values x of the first quantifier can easily be enumerated by a Turing machine. \square

Theorem 4. *There exists a countable SFT X and a CA $f : X \rightarrow X$ such that the asymptotic set of f is Σ_3^0 -complete.*

Proof. It is easy to see that there exists a recursive set L such that solving

$$\exists k : \forall l : \exists m : (k, l, m, w) \in L$$

for given $w \in \mathbb{N}$ is Σ_3^0 -complete. We say w is a solution to L . We will many-one reduce any such set to the language of the asymptotic set of a CA.

Let M be an always halting counter machine for L which never re-enters the initial state. We construct a counter machine M' with at least states s_1, s_2 , counters C_w, C_k, C_l, C'_l, C_m , and a suitable set of auxiliary counters, having the property that started from configuration $(s_1, w, k, i_1, \dots, i_k)$, M' enters the state s_1 infinitely many times if and only if w is a solution to L for that choice of k , and the i_j are arbitrary. The counter C_w will always contain the value w and is never modified. The counters C_k, C_l, C'_l, C_m in M' will play the role of quantifiers, with C_k containing the guess k . The idea of the construction is that C'_l will loop through the values $[1, n]$ for larger and larger n , which is stored in C_l , and at every value l , C_m is used to search for an m such that $(k, l, m, w) \in L$. Such an event is followed by setting all counters to 0, except C_w, C_k, C_l and C'_l , and visiting the state s_1 . It follows that the state s_1 is visited infinitely many times if and only if k was a correct guess for w .

We wish to have the central pattern of the configuration corresponding to $(s_1, w, \infty, \dots, \infty)$ in the asymptotic set of the corresponding CA if and only if w is a solution to L . Infinite values in counters during the simulation are handled by simply making the machine head check that all counter values are finite before entering s_1 . We also have the problem that while entering s_1 infinitely many times indeed characterizes the solutions of L , the values of other counters might be visible in the asymptotic set. We thus modify all counters except C_w to store their data in the form $2^n \cdot n'$, where n' is odd and n is the useful piece of data, and multiply them by 3 every time M enters the state s_1 . \square

Corollary 1. *There exists a CA c on the full shift such that the asymptotic set of c is Σ_3^0 -complete.*

Proof. Let X and c be given by Theorem 4. We embed X in a full shift and add a spreading state t for c which appears whenever a neighborhood is forbidden in X . Then we may use the same reduction as before, since points outside X only contribute words in t^* to the language of the asymptotic set. \square

From a fixed SFT, asymptotic sets do not necessarily form a larger class than limit sets:

Proposition 3. *There exists a countable sofic shift X and a CA $c : X \rightarrow X$ such that the limit set of c is not the asymptotic set of any CA on X .*

Proof. Let

$$X = \mathcal{B}^{-1}(0^*l0^*\#0^*r0^* + 0^*l'0^*\#0^*r'0^*)$$

and c the CA on X which moves l and r towards $\#$, and moves l' and r' away from $\#$, changing $l\#r$ to $l'\#r'$, but $l\#0$ and $0\#r$ to $0\#0$. It is clear that the limit set is

$$Y = \mathcal{B}^{-1}\left(\bigcup_{n \in \mathbb{N}} (0^*l0^*\#0^*r0^* + 0^*l'0^n\#0^n r'0^*)\right).$$

We claim that no CA e on X has Y as its asymptotic set. Assume on the contrary that this is possible, and let e have radius R . In general, it is clearly true for asymptotic sets that for all N , the $e^n(x)_{[-N,N]}$ must be a word of Y for all x and sufficiently large n . In particular, words of Y must map to words of Y . Also, points in Y have preimages in Y .

Since configurations of the form $y(M, M') = {}^\infty 0l0^M \# 0^{M'} r0^\infty$ appear in the asymptotic set with no restriction on M and M' , but the point $z(M, M') = {}^\infty 0l0^M \# 0^{M'} r0^\infty$ only appears for $M = M'$, a simple case analysis shows that for large M and M' , such points map to points of the same form.

Since $z(M, M)$ is isolated in X , it must be e -periodic, and since it is in the asymptotic set, it must map to a point of the form $z(N, N)$. If for some M , a point of the form $z(N, N)$ with $N \leq 2R - 1$ never appeared in the orbit, we would also find an e -periodic point of the form $z(M, M')$ with $M \neq M'$. Therefore, there exists a point $z(N, N)$ with $N \leq 2R - 1$ such that $z(M, M)$ appears in its orbit for arbitrarily large M . But this is a contradiction, since $z(N, N)$ must be e -periodic.

However, if the subshift is not fixed, we have the following realization theorem. Note that this is not stronger than Theorem 4, since it only implies the existence of a sofic shift, and not an SFT.

Theorem 5. *Let X be a countable sofic shift and let Y be a Σ_3^0 subshift of X . Then there exists a countable sofic shift $Z \supset X$ and a CA $c : Z \rightarrow Z$ such that the asymptotic set of c from Z is exactly Y .*

Proof. The subshift Z has the same periodic points as Y and the cellular automaton c having Y as an asymptotic set from Z behaves as the identity map on the periodic points. Let p be a common period for the periodic points of Z .

If x and y are periodic, and there is a unique point z that is left asymptotic to x and right asymptotic to y , then c also behaves as the identity map on z . For all pairs of periodic points without this property, we construct a subshift $Z(x, y)$ such that started from $Z(x, y)$, the automaton c produces exactly the points of Y left and right asymptotic to x and y , respectively, in the asymptotic set.

Let x be u -periodic and y be v -periodic for $|u| = |v| = p$. The general form of a ‘good’ point in $Z(x, y)$ is

$${}^\infty uM_1u^*wv^*M_2v^\infty,$$

where w is called a *message* moving either left or right, such that ${}^\infty uww^\infty$ appears in X , and u^*wv^* is called the *periodic medium*. The words M_1 and M_2 encode counter machines. Since there are at least two points left and right asymptotic to x and y , respectively, we may choose a nonperiodic word w_{ok} such that ${}^\infty uw_{ok}v^\infty$ is an isolated point in X . Between a pattern of u ’s and v ’s, w_{ok} is shifted a multiple of p steps to the right, and all other words are shifted some multiple of p steps to the left. Since there are only finitely many different choices of x and y , it is easy to handle the finite amount of words w_{ok} separately and move them in a different direction than all others (since these islands are maximal).

By the assumption on Y , the good set of words w that occur between ${}^\infty u$ and v^∞ must be given by the solutions w to

$$\exists k : \forall l : \exists m : (k, l, m, w) \in L$$

for a recursive language L .

The machines M_1 and M_2 have to be encoded slightly differently than usually, so that large values in counters do not change the set of periodic points of Z from those of Y . The finitely many counter markers, however, can be over a separate alphabet altogether. In the evolution of c , the machines M_1 and M_2 move slowly away from each other, so that they cannot be seen in the asymptotic set. The ‘base points’ of counters (where counters become zero) of M_1 and M_2 are on the side of the periodic medium between them. This makes the sending of messages more intuitive, and makes it clearer that the machines actually disappear from the asymptotic set.

The basic idea is that M_2 will have the word w encoded in its counters, and a guess k that it keeps constant. It iterates through all choices of l (checking all values infinitely many times) just like in the proof of Theorem 4, and for each l , it then starts iterating through all m . When the case $(k, l, m, w) \in L$ occurs, it sends w to the left. We omit the details of outputting w at a sufficient speed. After outputting w , M_2 waits for M_1 to respond in order not to send multiple messages at once (which would add entropy to Z). The sofic rule of Z makes sure M_2 is in a waiting state if a message is on its way. When the message w_{ok} is received from M_1 , M_2 continues its computation. The message w is sent infinitely many times iff it is a solution to L and the guess k was correct.

The machine M_1 behaves similarly. It waits for a message from M_2 , and when one is received, it sends the message w_{ok} to the right. Note that there are only finitely many different forms of points in X (and in particular Y), so the rule of the sofic shift can share this information about w between M_1 and M_2 .

Now it is clear that if the automaton is started from a good point, we obtain only words of Y in the asymptotic set. Also, for any $w \sqsubset Y$, a good point puts it in the asymptotic set if M_2 uses w as a message. As for bad points, we note that if only one counter machine occurs in a point, or a counter has an infinite value, then the worst that can happen is that the flow of messages stops at one point, leaving only the periodic medium – which is in Y – in the asymptotic set. \square

In particular, all limit sets are asymptotic sets in the following sense:

Corollary 2. *Let X be a countable sofic shift and $f : X \rightarrow X$ a CA with limit set Y . Then there exists a countable sofic shift $Z \supset X$ and a CA $g : Z \rightarrow Z$ such that the asymptotic set of g is Y .*

Proposition 3 shows that Z above cannot be chosen equal to X in general.

7 Conclusions and Future Work

In this paper, we have studied cellular automata on sofic shifts mainly from the computational point of view. We proved a useful result about the long-term

evolution of points under such automata (the Starfleet Lemma), and showed the decidability of some dynamical properties (nilpotency and periodicity) and undecidability of others (stability). It seems that those properties that depend on the exact behavior of spaceships are likely to be decidable. The Starfleet Lemma is, in these cases, a very powerful tool. We also studied the computational capacity and complexity of asymptotic sets, which turned out to be quite high.

We have already proved several results about the dynamics of cellular automata on countable sofic shifts which did not make it into this article due to lack of space. Future work might involve studying the relations between dynamical systems defined in this way. In particular, the conjugacy problem of countable sofic shifts, which seems more combinatorial than the general conjugacy problem, but challenging nevertheless, is still unsolved.

Acknowledgements

We'd like to thank Jarkko Kari for pointing out that Lemma 3 holds even in the case of a general subshift, Pierre Guillon for his useful comments, in particular on the last section, and the anonymous referees of ICALP 2012 for their suggestions on the overall presentation and structure of the article.

References

1. Lind, D., Marcus, B.: An introduction to symbolic dynamics and coding. Cambridge University Press, Cambridge (1995)
2. Kari, J.: The nilpotency problem of one-dimensional cellular automata. *SIAM J. Comput.* **21**(3) (1992) 571–586
3. Kari, J.: Rice's theorem for the limit sets of cellular automata. *Theoret. Comput. Sci.* **127**(2) (1994) 229–254
4. Ballier, A., Durand, B., Jeandel, E.: Structural aspects of tilings. In Susanne Albers, P.W., ed.: Proceedings of the 25th Annual Symposium on the Theoretical Aspects of Computer Science, Bordeaux, France, IBFI Schloss Dagstuhl (February 2008) 61–72 11 pages.
5. Morita, K.: Universality of a reversible two-counter machine. *Theoretical Computer Science* (1996)
6. Guillon, P., Richard, G.: Asymptotic behavior of dynamical systems and cellular automata. ArXiv e-prints (April 2010)

Appendix

Proof of Lemma 3:

Proof. Suppose on the contrary that the CA $c : X \rightarrow X$ is weakly nilpotent but not nilpotent for some subshift $X \subset S^{\mathbb{Z}}$. Then there exists a configuration x in the limit set of c with $x_0 \neq 0$, for the symbol 0 such that all points reach ${}^\infty 0^\infty$ in finitely many steps. Since x is in the limit set, it has an infinite chain $(x^i)_{i \in \mathbb{N}}$ of preimages by a compactness argument. If r is the radius of c , then since 0 is a quiescent state, there must exist a sequence of integers $(k_i)_{i \in \mathbb{N}}$ such that $|k_i - k_{i+1}| \leq r$ and $\sigma^{k_i}(x^i)_0 \neq 0$ for all $i \in \mathbb{N}$.

Let y be a limit of a converging subsequence of $(\sigma^{k_i}(x^i))_{i \in \mathbb{N}}$ in the product topology of $S^{\mathbb{Z}}$. Since c is weakly nilpotent, there exists $n \in \mathbb{N}$ such that $c^n(y) = {}^\infty 0^\infty$. Let $i \in \mathbb{N}$ be such that

$$\sigma^{k_{i+n}}(x^{i+n})_{[-2rn, 2rn]} = y_{[-2rn, 2rn]}. \quad (1)$$

By definition of the x^i and k_i , we have that $c^n(x^{i+n}) = x^i$ and

$$\sigma^{k_{i+n}}(x^i)_{[-rn, rn]} \neq 0^{2rn+1}.$$

But (1) implies the contrary, and we have a contradiction. \square