Playing with Subshifts*

Ville Salo[1] and Ilkka Törmä[2]

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

Abstract. We study the class of sequence-building games, where two players pick letters from a finite alphabet to construct an infinite word. The outcome is determined by whether the resulting word lies in a prescribed subshift (a win for player A) or not (a win for player B). We investigate the relation between the target subshift and the set of turn orders for which A has a winning strategy.

Keywords: subshifts, games

1 Introduction

Subshifts are the central objects of symbolic dynamics, and also have an important role in coding and information theory. A subshift is a set of infinite sequences over a finite alphabet defined by forbidden patterns. A sequence lies in the subshift if and only if no forbidden pattern occurs in it. Subshifts can be viewed as streams of information flowing from a transmitter towards a receiver. The asymptotic rate at which information can be sent is captured by the notion of entropy.

Consider the following scenario: a person (we will call her Alice, or A) is trying to send some information through the subshift. She does this by choosing, one by one, a letter from the alphabet, and constructing an infinite sequence from them. She can pick any letters she fancies, as long as the resulting sequence contains no forbidden pattern. But at some prescribed moments an adversary (who we will call Bob, or B) manages to insert a letter to the end of the finite sequence A has constructed thus far. This may cause a forbidden word to appear, destroying the information channel, but not necessarily: if A knows beforehand the moment at which B will interfere, she can choose her sequence so that no matter which letter B inserts, no forbidden pattern results. The situation can be thought of as a game, where A wins if no forbidden pattern ever occurs.

Games of this form have previously been studied in [3], from the point of view of combinatorics on words. The article in question concentrates on subshifts in which no large enough approximate squares, or patterns approximately of the

* Research supported by the Academy of Finland Grant 131558

form ww for long enough w, occur. The results of the paper state that for certain definitions of 'approximate' and 'long enough', A has a winning strategy even if every choice for A is followed by some t choices for B.

In this article, we take a different view on the situation. Instead of asking whether a given subshift admits a winning strategy for A when a prescribed order for the turns of A and B is used, we fix a subshift X and study the set of turn orders for which A has a winning strategy on it. We call this set the winning shift of X, since it actually turns out to be a subshift over the alphabet $\{A, B\}$. We study the realization of binary subshifts as winning shifts, and the connections between properties of X and its winning shift. In particular, connections between the entropies of X and its winning shift are obtained.

2 Definitions

2.1 Standard Definitions

Let S be a finite set, called the *alphabet*. We denote by S^* the set of finite words over S, and by λ the empty word of length 0. The set $S^{\mathbb{N}}$ of infinite state sequences, or *configurations*, is called the *full shift on* S. If $x \in S^{\mathbb{N}}$ and $i \in \mathbb{N}$, then we denote by x_i the ith 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 *appears in* x, if $w = x_{[i,i+|w|-1]}$ for some i. For words $u, v \in S^*$, the notation uv^∞ has the intuitive meaning. Two elements $x, y \in S^{\mathbb{Z}}$ are *asymptotic* if $x_i = y_i$ for all sufficiently large i. For a letter $s \in S$ and $u \in S^*$, we denote by $|u|_s$ the number of occurrences of s in u.

We define a metric d on the full shift by setting $d(x, y) = 0$ if $x = y$, and by setting $d(x, y) = 2^{-i}$ where $i = \min\{j \mid x_j \neq y_j\}$ if $x \neq y$. The topology defined by d makes $S^{\mathbb{N}}$ a compact metric space. We define the *shift map* $\sigma : S^{\mathbb{N}} \to S^{\mathbb{N}}$ by $\sigma(x)_i = x_{i+1}$. Clearly σ is a continuous surjection from the full shift to itself.

A *subshift* is a closed subset X of the full shift with the property $\sigma(X) \subset X$. Alternatively, a subshift is defined by a set $F \in S^*$ 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* (SFT for short), if F is regular, then $\mathcal{X}_F$ is *sofic*, and if F is recursively enumerable, then $\mathcal{X}_F$ is *effective*. The length of the longest forbidden pattern of an SFT is called its *window size*. We define $\mathcal{B}_k(X) = \{w \in \Sigma^k \mid \exists x \in X : w \sqsubset x\}$ as the set of words of length k appearing in X, and define the *language* of X as $\mathcal{B}(X) = \bigcup_{k \in \mathbb{N}} \mathcal{B}_k(X)$. Since a subshift is uniquely defined by its language [2], we may write $X = \mathcal{B}^{-1}(L)$, if $\mathcal{B}(X)$ is the set of factors (contiguous subwords) of the language $L \subset S^*$.

We say a word $u \in \mathcal{B}(X)$ is *left extendable* if for all $n \in \mathbb{N}$ there exists v with $|v| \geq n$ such that $vu \in \mathcal{B}(X)$, and X is *left extendable* if every $u \in \mathcal{B}(X)$ is. The *left extendable part of* X is the (left extendable) subshift $\mathcal{B}^{-1}(\{u \in \mathcal{B}(X) \mid u \text{ is left extendable}\})$. We say X is *transitive* if for all $u, w \in \mathcal{B}(X)$ there exists $v \in \mathcal{B}(X)$ such that $uvw \in \mathcal{B}(X)$, and *mixing* if the length of v can be chosen arbitrarily, as long as it is sufficiently large (depending on u and w). It is known that for SFTs and sofic shifts, the minimum length of v does not depend on u or w, but is a constant, called the *mixing distance of* X. The *entropy* of a subshift X is defined as $h(X) = \lim_{n \to \infty} \frac{1}{n} \log |\mathcal{B}_n(X)|$.

An *alternating finite automaton* is a sextuple $A = (Q_\exists, Q_\forall, \Sigma, q_1, F, \delta)$, where $Q_\exists$ is the finite set of *existential states*, $Q_\forall$ the finite set of *universal states*, Σ the finite *input alphabet*, $q_1 \in Q = Q_\exists \cup Q_\forall$ the *initial state*, $F \subset Q$ the *final states* and $\delta : (Q \times \Sigma) \to 2^Q$ the *transition function*. For a word $w \in \Sigma^*$ and $q \in Q$, we denote

$$q(w) = \begin{cases} q \in F, & \text{if } w = \lambda \\ \exists r \in \delta(q, w_0) : r(w_{[1,|w|-1]}), & \text{if } w \neq \lambda \text{ and } q \in Q_\exists \\ \forall r \in \delta(q, w_0) : r(w_{[1,|w|-1]}), & \text{if } w \neq \lambda \text{ and } q \in Q_\forall \end{cases}$$

where λ denotes the empty word. We say w is *accepted* by A if $q_1(w)$ holds. Deterministic automata can be seen as alternating automata with $|\delta(q, a)| = 1$ for all $q \in Q, a \in \Sigma$. Alternating automata were first defined in [1] (although slightly more generally), and there it was proved that they accept exactly the regular languages.

2.2 Word Games

We now define games in which two players take turns to pick letters from an alphabet and build a finite or infinite word. A *word game* is a triple (S, n, X), where S is a finite set, $n \in \mathbb{N} \cup \{\mathbb{N}\}$ and $X \subset S^n$ is the *target set*. An *ordered word game* is a tuple (S, n, X, a), where (S, n, X) is a subshift game and $a \in \{A, B\}^n$ is the *turn order*. If $n = \mathbb{N}$ and X is a subshift, we have *subshift games* and *ordered subshift games*, respectively. An ordered word game (S, n, X, a) should be understood as the players A and B building a word in S^n by choosing one coordinate at a time, with the coordinate i being chosen by a_i. If the resulting word lies in X, then A wins, and otherwise B does.

Let $G = (S, n, X, a)$ be an ordered word game. A *strategy for* G is a function $s : S^* \to S$ that specifies the next pick of a player, given the word constructed thus far. The *play* of a pair (s_A, s_B) of strategies for G is the sequence $x = p(G, s_A, s_B) \in S^n$ defined inductively by $x_i = s_{a_i}(x_{[0,i-1]})$. We say that a strategy s is *winning for* A if $p(G, s, s_B) \in X$ for all strategies s_B (A wins the game no matter how B plays), and *winning for* B if $p(G, s_A, s) \notin X$ for all strategies s_A.

A strategy for the word game (S, n, X) is a function $\zeta : \{A, B\}^n \times S^* \to S$ that, given a turn order $a \in \{A, B\}^n$, produces a strategy $\zeta(a, \cdot)$ for the ordered word game (S, n, X, a). Consider a pair (ζ_A, ζ_B) of strategies for (S, n, X). We say that A *should play with* ζ_A *and* B *with* ζ_B, if for all $a \in \{A, B\}^n$ we have $p((S, n, X, a), \zeta_A(a, \cdot), \zeta_B(a, \cdot)) \in X$ iff A has a winning strategy for (S, n, X, a). Intuitively, the strategies ζ_A and ζ_B are always 'at least as good' as any winning strategies. This notion will simplify some of our proofs.

Given a set $X \subset S^n$, where $n \in \mathbb{N} \cup \{\mathbb{N}\}$, we define

$$W(X) = \{a \in \{A, B\}^n \mid A \text{ has a winning strategy for } (S, n, X, a)\}.$$

The alphabet S will always be clear from the context. If $n = \mathbb{N}$ and X is a subshift, $W(X)$ is called the *winning shift* of X. We denote by $\mathcal{W}$ the class of all winning shifts.

In a sense, this notion generalizes one of the problem families studied in [3]: Given a periodic sequence $x \in \{A, B\}^{\mathbb{N}}$ and a natural parametrized class of subshifts $(X_i)_{i \in \mathcal{I}}$, for which parameters $i \in \mathcal{I}$ do we have $x \in W(X_i)$. To showcase our formalism, we state the two results of [3] which are of this form.

Theorem 1 (Theorem 1.4 of [3]). *For all $\epsilon > 0$ and $t \in \mathbb{N}$, we have*

$$(AB^t)^\infty \in W(\mathcal{X}(\{wuw \mid |w| > N_{\epsilon,t}, |u| \leq (2-\epsilon)^{\frac{|w|}{t+1}}\})),$$

for large enough $N_{\epsilon,t} \in \mathbb{N}$.

Theorem 2 (Theorem 1.5 of [3]). *For all $\epsilon > 0$ and $t \in \mathbb{N}$, we have*

$$(AB^t)^\infty \in W(\mathcal{X}(\{uv \mid |u| = |v| > N_{\epsilon,t}, H(u,v) < ((2t+2)^{-1} - \epsilon)|u|\})),$$

for large enough $N_\epsilon \in \mathbb{N}$, where $H(u, v)$ denotes the Hamming distance of u and v, that is, the number of coordinates in which they differ.

We note that in terms of symbolic dynamics, $(AB^t)^\infty$ is a rather trivial object, while the subshifts on the right are quite complicated in that they are not sofic. Our take in this article is, simply put, to make the right side simpler and see what can happen in the left.

3 Results

As stated in the introduction, the winning shift $W(X)$ of a subshift X is a subshift itself. We begin this section with the proof of this fact.

Proposition 1. *Let $X \subset S^{\mathbb{N}}$ be a subshift. Then $W(X)$ is a downward closed subshift with relation to the order $A < B$.*

Proof. We first prove that $W(X)$ is closed. For that, we take $a \notin W(X)$, and prove that there exists an open neighborhood U of a such that $U \cap W(X) = \emptyset$. Since $a \notin W(X)$, there exists a strategy s_B that is winning for B for the game $G = (S, X, a)$. Now, if there would exist plays $p(G, s, s_B)$ with arbitrarily long prefixes free of forbidden words, then by compactness a winning play for A would exist, contradicting the definition of s_B. Thus there exists $N \in \mathbb{N}$ such that the prefix of length N of any game $p(G, s, s_B)$ contains a forbidden word, and we can choose $U = \{b \in \{A, B\}^{\mathbb{N}} \mid \forall i \leq N : a_i = b_i\}$.

It is clear that $a \in W(X)$ implies $\sigma(a) \in W(X)$, since if s is a winning strategy for A for (S, X, a), then $w \mapsto s(s(\lambda)w)$ is winning for A for $(S, X, \sigma(a))$: A can simply pretend that a single letter was played before the game defined by $\sigma(a)$ started. It is likewise clear that $W(X)$ is downward closed, since changing a B into an A in a member of $W(X)$ allows A to play arbitrarily in that coordinate, and still win. □

Note that $W(X)$ is an ideal (a downward closed subshift, which is also closed under coordinatewise maxima of two sequences) iff it is either $\{A^\infty\}$ or $\{A, B\}^{\mathbb{N}}$. However, there do exist interesting downward closed subshifts.

Example 1.

- $\mathcal{B}^{-1}(0^*(10^*20^*)^*)$ is a mixing sofic shift whose (left extendable) winning shift $\mathcal{B}^{-1}(A^*BA^*)$ is nontransitive proper sofic.
- $\mathcal{B}^{-1}(0^*1^*)$ is an SFT whose (left extendable) winning shift $\mathcal{B}^{-1}(A^*BA^*)$ is nontransitive proper sofic.
- $\mathcal{B}^{-1}((01 + 0001)^*)$ is a transitive SFT whose (left extendable) winning shift $\mathcal{B}^{-1}((AB(AA)^+)^*)$ is transitive proper sofic.

We now prove closure properties of some classes of subshifts under the operation W, and make some basic observations about $\mathcal{W}$.

Proposition 2. *Let $X \subset S^{\mathbb{N}}$ be a sofic shift. Then $W(X)$ is sofic.*

Proof. Since X is sofic, there exists a finite automaton M accepting its language $\mathcal{B}(X)$. A word $w \in \{A, B\}^*$ is in $\mathcal{B}(W(X))$ if and only if

$$\Box_0 v_0 \in S : \Box_1 v_1 \in S : \cdots \Box_{|w|-1} v_{w-1} \in S : v = v_0 v_1 \cdots v_{|w|-1} \in \mathcal{B}(X), \quad (1)$$

where each $\Box_i$ is $\exists$ if $w_i = A$, and $\forall$ otherwise. This formula is checkable by an alternating automaton that chooses the letter v_i either existentially or universally, depending on whether it last read A or B from its input, and keeps track of whether M would accept the resulting v. Since alternating automata accept only regular languages, we see that $\mathcal{B}(W(X))$ is regular, and thus $W(X)$ is sofic. $\Box$

Question 1. What is the class of $W(X)$ for X SFT?

Proposition 3. *Let $X \subset S^{\mathbb{N}}$ be an effective subshift. Then $W(X)$ is effective.*

Proof. This is clear from the characterization (1): any word $w \notin \mathcal{B}(W(X))$ will fail to satisfy the formula, and an algorithm enumerating $S^* - \mathcal{B}(X)$ will find the words $v \notin \mathcal{B}(X)$ that cause this to happen. $\Box$

We now give a rather trivial full characterization of subshifts in $\mathcal{W}$, when there is no requirement on the subshift on which the game is being played.

Proposition 4. *If a subshift $X \subset \{A, B\}^{\mathbb{N}}$ is downward closed with relation to $A < B$, then $W(X) = X$.*

Proof. In a downward closed subshift, each player P should always play the symbol P. $\Box$

Corollary 1. *The subshifts in $\mathcal{W}$ are exactly the subshifts over $\{A, B\}$ which are downward closed with relation to $A < B$.*

Proof. Every subshift in $\mathcal{W}$ is downward closed by Proposition 1, and if the subshift $X \subset \{A, B\}^{\mathbb{N}}$ is downward closed, then $W(X) = X \in \mathcal{W}$.

Corollary 2. *The class $\mathcal{W}$ is closed under union and intersection.*

By the previous corollary, also the classes of winning shifts of sofic shifts and effective subshifts are closed under these operations, since both classes are closed under the operation W. However, Corollary 1 is not very useful when considering classes of subshifts which are not closed under W, in particular the SFTs. Therefore, we mention the following standard trick:

Proposition 5. *Let $X \subset S^{\mathbb{N}}$ and $Y \subset R^{\mathbb{N}}$ be subshifts. Then $W(X \times Y) = W(X) \cap W(Y)$.*

Proof. Clearly, both A and B should play using their 'product strategies': if the player P should play ζ_S on (S, X) and ζ_R on (R, Y), then P should play $(u, v) \mapsto (\zeta_S(u), \zeta_R(v))$ on $(S \times R, X \times Y)$. With these choices of strategies, the claim immediately follows. □

Corollary 3. *The class of winning shifts of SFTs is closed under intersection.*

We now turn to the question of how different properties of a subshift X result in $W(X)$ having similar properties.

Proposition 6. *If X is a mixing (transitive) SFT, then the left-extendable part of $W(X)$ is mixing (transitive).*

Proof. Let X be mixing with mixing distance and window size m. Let $u, v \in \mathcal{B}(W(X))$ be left-extendable, and let $n \in \mathbb{N}$ and $k \geq 2m$. We claim that $a = A^n u A^k v A^\infty \in W(X)$, and that the following is a winning strategy for A on a. In the first $n + |u|$ coordinates, play as on $A^n u A^\infty$. In the next m coordinates, play a mixing word so that you can continue as you would play on $A^{k-m} v A^\infty$. The strategy is winning for A, since the symbols played in the first $n + |u| + m$ coordinates have no effect on the game played on the tail vA^∞. This implies that $uA^k v$ is a left-extendable word in $\mathcal{B}(W(X))$, and thus $W(X)$ is mixing. The case of transitive X is similar. □

If $W(X)$ has high entropy, then in a typical point of $W(X)$, B is able to play quite often, but A still wins. Since B can play arbitrarily, the B-coordinates can thus be chosen arbitrarily, with the resulting configuration still in X. This gives a lower bound for the entropy of X, which we approximate in the following.

Lemma 1. *Let $X \subset S^{\mathbb{N}}$ be a subshift. Denoting*

$$d = \limsup_n \max \left\{ \frac{|w|_B}{n} \mid w \in \mathcal{B}_n(W(X)) \right\},$$

we have $h(X) \geq d \log |S|$.

Proof. Let $\epsilon > 0$ and $n \in \mathbb{N}$ be arbitrary, and let $w \in \mathcal{B}(W(X))$ be such that $|w| \geq n$ and $\frac{|w|_B}{|w|} \geq d - \epsilon$. Let s be a winning strategy for A on wA^∞. By enumerating all the possible strategies for B and the prefixes of the resulting games of length $|w|$, we find that $|\mathcal{B}_{|w|}(X)| \geq |S|^{|w|_B}$. This implies that

$$\frac{1}{|w|} \log |\mathcal{B}_{|w|}(X)| \geq \frac{|w|_B}{|w|} \log |S| \geq (d - \epsilon) \log |S|.$$

Since $|w|$ can be chosen arbitrarily large and ϵ arbitrarily small, the claim follows. □

Lemma 2. *Let $2 > k > 1$, and let $\epsilon > 0$ be such that $\left(\frac{e}{\epsilon}\right)^\epsilon < k$. Then we have $\binom{n}{\lfloor \epsilon n \rfloor} \leq k^n$ for all n large enough.*

Proof. For large enough n we have that

$$\binom{n}{\lfloor n\epsilon \rfloor} \leq \left(\frac{n\epsilon}{\lfloor n\epsilon \rfloor}\right)^{n\epsilon} \left(\frac{ne}{n\epsilon}\right)^{n\epsilon} \leq k^n,$$

using the approximations $\binom{n}{m} \leq \left(\frac{ne}{m}\right)^m$ and $\left(\frac{x}{\lfloor x \rfloor}\right)^x \leq \left(1 - \frac{1}{x}\right)^{-x} \longrightarrow e$. □

Proposition 7. *Let $h(W(X)) \geq \log k$, and let $\epsilon > 0$ be such that $\left(\frac{e}{\epsilon}\right)^\epsilon < k$. Then $h(X) \geq \epsilon \log |S|$. In particular, if $h(W(X)) > 0$, then $h(X) > 0$.*

Proof. First, note that $W(X)$ is binary, so $k \leq 2$. If equality holds, then $X = S^{\mathbb{N}}$ and the claim holds. Otherwise, Lemma 2 implies that

$$\limsup_n \max \left\{ \frac{|w|_B}{n} \mid w \in \mathcal{B}_n(W(X)) \right\} \geq \epsilon,$$

and Lemma 1 gives the claim. □

We are not yet able to say much about whether X having entropy very close to $\log |S|$ causes $W(X)$ to have positive entropy as well. However, we conjecture this to be the case.

Conjecture 1. Fix the alphabet S. For all $\epsilon > 0$ there exists $\delta > 0$ such that if $X \subset S^{\mathbb{N}}$ is a subshift with $h(X) > \log |S| - \delta$, then $h(W(X)) > \log 2 - \epsilon$.

In the binary case, the entropy of a subshift and its winning shift are actually equal. This follows from a simple combinatorial lemma.

Proposition 8. *If X is a binary subshift, then $h(X) = h(W(X))$.*

Proof. We first claim that for all target sets $Y \subset \{0, 1\}^n$, where $n \in \mathbb{N}$, we have $|W(Y)| = |Y|$. We prove this by induction, starting with the case $n = 1$, which is easily seen true. Suppose then that $n > 1$, and let $Y_b = \{w \in \{0, 1\}^{n-1} \mid bw \in Y\}$ for $b \in \{0, 1\}$. We clearly have $Y = 0Y_0 \cup 1Y_1$, and by the induction hypothesis, also $|W(Y_b)| = |Y_b|$ holds for $b \in \{0, 1\}$.

Let $a \in W(Y)$, and suppose that $a_0 = A$. If A has a winning strategy in which she starts with $b \in \{0, 1\}$, then $a_{[1,n-1]} \in W(Y_b)$. Conversely, if $a_{[1,n-1]} \in W(Y_b)$, then A has a winning strategy that starts with b, and then follows the strategy for $a_{[1,n-1]}$. Thus

$$|\{a \in W(Y) \mid a_0 = B\}| = |W(Y_0)| + |W(Y_1)| - |W(Y_0) \cap W(Y_1)|.$$

On the other hand, if $a_0 = B$, then $a_{[1,n-1]}$ must be in $W(Y_0) \cap W(Y_1)$ and the converse also holds, so

$$|\{a \in W(Y) \mid a_0 = B\}| = |W(Y_0) \cap W(Y_1)|.$$

All in all, we have that $|W(Y)| = |W(Y_0)| + |W(Y_1)| = |Y_0| + |Y_1| = |Y|$, and the claim is proved.

From the above result it follows that $|\mathcal{B}_n(W(X))| = |W(\mathcal{B}_n(X))| = |\mathcal{B}_n(X)|$ holds for all $n \in \mathbb{N}$, and thus $h(X) = h(W(X))$. □

References

1. Ashok K. Chandra, Dexter C. Kozen, and Larry J. Stockmeyer. Alternation. *J. ACM*, 28:114–133, January 1981.
2. Douglas Lind and Brian Marcus. *An introduction to symbolic dynamics and coding*. Cambridge University Press, Cambridge, 1995.
3. Wesley Pegden. Highly nonrepetitive sequences: winning strategies from the local lemma. *Random Structures Algorithms*, 38(1-2):140–161, 2011.