

ANNALS OF MATHEMATICS

A characterization of the entropies of multidimensional shifts of finite type

By MICHAEL HOCHMAN and TOM MEYEROVITCH



SECOND SERIES, VOL. 171, NO. 3

May, 2010

ANMAAH

A characterization of the entropies of multidimensional shifts of finite type

By MICHAEL HOCHMAN and TOM MEYEROVITCH

Abstract

We show that the values of entropies of multidimensional shifts of finite type (SFTs) are characterized by a certain computation-theoretic property: a real number $h \geq 0$ is the entropy of such an SFT if and only if it is right recursively enumerable, i.e. there is a computable sequence of rational numbers converging to h from above. The same characterization holds for the entropies of sofic shifts. On the other hand, the entropy of strongly irreducible SFTs is computable.

1. Introduction

A shift of finite type (SFT) is an ensemble of colorings of \mathbb{Z} (a one-dimensional SFT) or \mathbb{Z}^d for $d > 1$ (a multidimensional SFT) defined by local rules. SFTs are one of the fundamental objects of study in symbolic dynamics, and their most significant invariant is their (topological) entropy, which measures the asymptotic growth of the number of legal colorings of finite regions (see §2 for definitions). Besides having been studied extensively from a dynamical perspective as topological analogs of Markov chains [21], [28], [27], SFTs appear naturally in a wide range of other disciplines. In information theory, SFTs were used by Shannon as models for discrete communication channels [24], for which entropy describes the capacity; similarly, SFTs model “two-dimensional” channels [8]. SFTs have been used to study the dynamics of geodesic flows and have played an important role in the classification of the dynamics of Anosov and Axiom A diffeomorphisms [1], [3], where entropy is again a fundamental invariant. In mathematical physics SFTs are often called hard-core models, and are used to model a wide variety of physical systems; this is the thermodynamic formalism [23]. In this setup it is of central importance to understand the equilibrium states of the system, which are the invariant measures of maximal entropy.

It is well known that in the one-dimensional case the entropy of an SFT may be effectively calculated, since it is the logarithm of the spectral radius of a certain

positive integer matrix which is derived from the combinatorial description of the system. D. Lind [16] has given an algebraic characterization of the numbers which arise as entropies of one-dimensional SFTs. A Perron number is a real algebraic integer greater than 1 and greater than the modulus of its algebraic conjugates. The entropies of one-dimensional SFTs are precisely the nonnegative rational multiples of logarithms of Perron numbers.

In higher dimensions the problem becomes much more difficult. The dynamics of multidimensional SFTs is vastly more complicated than their one-dimensional counterparts. For instance, strongly irreducible multidimensional SFTs may have more than one measure of maximal entropy [4], and zero entropy can coexist with rather complex dynamics [20]. In general it is undecidable whether a given set of rules define a nonempty SFT [2], [22]. Regarding the entropy, even when the rules defining an SFT enjoy good symmetry properties, calculating the entropy is usually beyond current technology. As a result numerical methods have been developed to approximate the entropy (e.g. [10]), but these usually apply to restricted class systems.

One should note that for certain \mathbb{Z}^d -actions which arise as automorphisms of compact groups (a class which includes some SFTs), explicit expressions for the entropy have been obtained by D. Lind, K. Schmidt and T. Ward [18]. We note however that while these expressions are explicit they do not provide much information on the properties of the entropies, e.g. whether they are algebraic, well approximable, etc.

In this paper we characterize those real numbers which can occur as entropies of multidimensional SFTs in terms of their computation-theoretic properties. It is natural to say that a real number h is *computable* if it can be calculated to any desired accuracy. More precisely, h is computable if there is an algorithm which, given input $n \in \mathbb{N}$, produces a rational number $r(n)$ with $|h - r(n)| < 1/n$. For example, every algebraic number is computable (since there are numerical methods for computing the roots of an integer polynomial), and so are e , π , since they can be written as power series with computable coefficients and rate of convergence.

A weaker notion is the following. A real number h is *right recursively enumerable* (sometimes called *upper semi recursive*) if there exists a Turing machine which, given n , computes a rational number $r(n) \geq h$ such that $r(n) \rightarrow h$ (equivalently, the right Dedekind cut $\{q \in \mathbb{Q} : q > h\}$ is a recursively enumerable set of rationals).

The class of right recursively enumerable numbers is countable since algorithms may be put in one-to-one correspondence with finite 0, 1-valued sequences, and hence there are only countably many of them. If h is computable then there is an algorithm computing $r(n)$ with $|h - r(n)| < \frac{1}{n}$, so the computable sequence $r(n) + \frac{1}{n}$ converges to h from above. This shows that the class of right recursively

enumerable numbers contains the computable numbers, and it can be shown to be strictly larger. For example, if we choose a recursive enumeration T_1, T_2, \dots of all Turing machines, and let $b_n = 0$ if T_n halts and $= 1$ otherwise, then $x = 0.b_1b_2b_3\dots$ can be seen to be the decreasing limit of a recursive sequence, but one cannot compute b_n as a function of n , since this would solve the halting problem. For more information, see [15].

THEOREM 1.1. *For $d \geq 2$ the class of entropies of d -dimensional SFTs is the class of nonnegative right recursively enumerable numbers.*

The property of right recursive enumerability is a necessary condition for a number to be the entropy of an SFT because the naive approximation algorithm, which counts locally admissible patterns on cubes, converges from above to the entropy. This follows from the work of Friedland [9]; we provide a different proof below. The main novelty here is the sufficiency of the condition.

A sofic system is a subshift factor of an SFT, i.e. an ensemble of colorings of \mathbb{Z}^d obtained from a fixed SFT X by applying a local transformation to each coloring in X (for a definition, see §2). In the one dimensional case, Coven and Paul [5] showed that every sofic system Y can be extended to an SFT X with the same entropy as Y . In particular, this implies that the class of entropies of sofic shifts is the same as that of SFTs. Whether the covering theorem is true in the multidimensional case is still open and seems quite hard (see [7] for a partial result). However some circumstantial evidence in favor of the covering theorem is provided by the following:

THEOREM 1.2. *For $d \geq 2$, the class of entropies of d -dimensional sofic shifts is the same as that of d -dimensional SFTs.*

This is a consequence of the fact that the entropy of sofic shifts is right recursively enumerable (Corollary 3.3 below), and the fact that an SFT is in particular a sofic system.

It is worth emphasizing that since there are noncomputable numbers which are right recursively enumerable, it follows from Theorem 1.1 that there are SFTs whose entropy cannot be computed effectively. (It is known that, for cellular automata [13] and general subshifts [25], entropy cannot be computed from the description of the system.) However, if one assumes strong enough mixing properties of the system the situation improves. Recall that an SFT is *strongly irreducible* if any two admissible patterns far enough apart may be extended admissibly to the whole lattice (see §2).

THEOREM 1.3. *The entropy of a strongly irreducible SFT is computable.*

We do not know if computability characterizes the entropies of strongly irreducible SFTs.

The rest of this paper is organized as follows. In the next section we introduce notation and background. In Section 3 we prove that the entropy of any SFT or sofic shift is right recursively enumerable, and that of a strongly irreducible SFT is computable. In Section 4 we outline the construction which constitutes the proof of the other direction of Theorem 1.1. Sections 6–8 give the details of the construction. In Section 9 we discuss some open problems.

Acknowledgement. This work was done during the authors' graduate studies, and we would like to thank our advisors, Benjamin Weiss and Jon Aaronson, for their support and advice. We also thank Mike Boyle for his comments.

2. Preliminaries

In this section we provide some background from symbolic dynamics and define SFTs and entropy. See [17], [14] for more information on these subjects.

2.1. *Subshifts and SFTs.* For a finite alphabet Σ let $\Sigma^{\mathbb{Z}^d}$ be the space of Σ -colorings of \mathbb{Z}^d (this is called the full shift on Σ). For a subset $F \subseteq \mathbb{Z}^d$ we refer to a function $a \in \Sigma^F$ as a coloring of F or an F -pattern. We say that patterns $a \in \Sigma^F$ and $b \in \Sigma^{F+u}$ are *congruent* if $a(v) = b(v+u)$ for every $v \in F$. We say that a pattern $a \in \Sigma^F$ appears at u in a pattern $b \in \Sigma^E$ if $b|_{F+u}$ and a are congruent.

If $E \subseteq F$ and $a \in \Sigma^F$ then a induces a coloring of E by restriction, namely $a|_E$. For a finite set $F \subseteq \mathbb{Z}^d$ and pattern $a \in \Sigma^F$ the cylinder set defined by a is

$$[a] = \{x \in \Sigma^{\mathbb{Z}^d} : x|_F = a\}.$$

We endow $\Sigma^{\mathbb{Z}^d}$ with the product topology, which is generated by the cylinder sets and makes $\Sigma^{\mathbb{Z}^d}$ into a compact metrizable space.

For $u \in \mathbb{Z}^d$ let $\sigma^u : \Sigma^{\mathbb{Z}^d} \rightarrow \Sigma^{\mathbb{Z}^d}$ be the homeomorphisms

$$(\sigma^u(x))(v) = x(v+u), \quad v \in \mathbb{Z}^d.$$

This gives an action of \mathbb{Z}^d on $\Sigma^{\mathbb{Z}^d}$ called the shift action. A subset $X \subseteq \Sigma^{\mathbb{Z}^d}$ is invariant under the shift action if $\sigma^u(X) = X$ for every $u \in \mathbb{Z}^d$. A closed invariant set $X \subseteq \Sigma^{\mathbb{Z}^d}$ is called a \mathbb{Z}^d -subshift.

A d -dimensional subshift of finite type (SFT) is defined by a finite alphabet Σ , a finite set $F \subseteq \mathbb{Z}^d$, and a collection $L \subseteq \Sigma^F$ of Σ -colorings of F , called the *syntax*. A Σ -coloring $x \in \Sigma^{\mathbb{Z}^d}$ of \mathbb{Z}^d is *admissible* for L if the pattern induced by x on every translate of F is congruent to a pattern in L . The SFT defined by L is the set $X \subseteq \Sigma^{\mathbb{Z}^d}$ of all admissible x . From the definition it is clear that an SFT is closed and shift-invariant.

Given an SFT X defined by a syntax $L \subseteq \Sigma^F$, we say that a finite pattern is *globally admissible* for X if it appears in X . In contrast we say that a pattern

$a \in \Sigma^E$ is *locally admissible* if $a|_{F+u}$ is congruent to a pattern in L whenever $F+u \subseteq E$. A globally admissible pattern is locally admissible, but the latter is not true in general.

An SFT $X \subseteq \Sigma^{\mathbb{Z}^d}$ is *strongly irreducible* if there is a constant $r > 0$, called a *gap*, such that for every $A, B \subseteq \mathbb{Z}^d$ satisfying $\|u-v\|_\infty \geq r$ for $u \in A, v \in B$, and for every pair of globally admissible $a \in \Sigma^A$ and $b \in \Sigma^B$, there is a point $x \in X$ with $x|_A = a$ and $x|_B = b$ (in other words, $a \cup b$ is globally admissible).

2.2. Topological entropy of subshifts. For a subshift $X \subseteq \Sigma^{\mathbb{Z}^d}$ and $F \subseteq \mathbb{Z}^d$ we say that a pattern $a \in \Sigma^F$ appears in X if $a = x|_F$ for some $x \in X$. For a set F let $N_X(F)$ denote the number of distinct Σ -colorings of F which appear in X . Let

$$F_n = \{1, \dots, n\}^d$$

denote the discrete d -dimensional cube of side n . The (topological) entropy $h(X)$ of X is defined by

$$h(X) = \lim_{n \rightarrow \infty} \frac{1}{|F_n|} \log N_X(F_n).$$

By convention the logarithm is to base 2. The limit above exists, and is in fact equal to $\inf_{n \in \mathbb{N}} \frac{1}{|F_n|} \log N_X(F_n)$.

2.3. Products, factors and isomorphism. Let $X \subseteq \Sigma^{\mathbb{Z}^d}$ and $Y \subseteq \Delta^{\mathbb{Z}^d}$ be two \mathbb{Z}^d -subshifts. The product system $X \times Y \subseteq (\Sigma \times \Delta)^{\mathbb{Z}^d}$ is then a symbolic \mathbb{Z}^d -system also, and satisfies $h(X \times Y) = h(X) + h(Y)$.

A continuous-onto map $\varphi : X \rightarrow Y$ is called a *factor map* if it commutes with the action; i.e. $\sigma^u \circ \varphi = \varphi \circ \sigma^u$ for all $u \in \mathbb{Z}^d$. An *isomorphism* is an invertible factor map. Both entropy and the property of being an SFT are invariants of isomorphism (although isomorphic SFTs are usually not defined by the same syntax), as is irreducibility.

A factor of an SFT is called a *sofic system*. In general a sofic system is not an SFT.

Every factor map $\varphi : X \rightarrow Y$ arises from a so-called *block code*, which means the following: There exist a finite set $F \subseteq \mathbb{Z}^d$ and a function $\varphi_0 : \Sigma^F \rightarrow \Delta$ such that

$$(\varphi(x))(u) = \varphi_0((\sigma^u x)|_F).$$

Conversely, given such a φ_0 we can define φ by this formula, and then φ is a factor map from X onto its image.

A factor map $\pi : X \rightarrow Y$ of symbolic systems $X \subseteq \Sigma^{\mathbb{Z}^d}$ and $Y \subseteq \Delta^{\mathbb{Z}^d}$ is called a *one-block map* if it is determined by a single symbol, i.e. it is induced by a map $\varphi_0 : \Sigma \rightarrow \Delta$. We will always assume our factor maps are one-block maps. There is no loss of generality in this since given a factor map $\varphi : X \rightarrow Y$ there is

a system X' isomorphic to X via an isomorphism $\psi : X' \rightarrow X$ so that the factor map $\psi\varphi : X' \rightarrow Y$ is a one-block map.

Similarly, an SFT is called one-step if it is defined by a syntax $L \subseteq \Sigma^{\{0,1\}^d}$. Every SFT is isomorphic to a one-step SFT. Note that for a one-step SFT one can splice a pattern into a larger pattern as long as the boundary of the subpattern agrees with the super-pattern. To be precise, for a subset $F \subseteq \mathbb{Z}^d$ let

$$\partial F = \{u \in F : \|u - v\|_\infty = 1 \text{ for some } v \in \mathbb{Z}^d \setminus F\}.$$

Then, for a one-step SFT X , given $x \in X$ and $a \in \Sigma^F$ such that $a|_{\partial F} = x|_{\partial F}$, the point y obtained by setting $y(u) = a(u)$ for $u \in F$ and $y(u) = x(u)$ otherwise, is globally admissible.

2.4. Invariant measures and entropy. Given a symbolic system X , a Borel measure μ on X is invariant under the shift action if $\mu(\sigma^u(A)) = \mu(A)$ for every Borel set $A \subseteq X$ and every $u \in \mathbb{Z}^d$. We denote the set of invariant Borel probability measures by $\mathcal{M}(X)$. The weak-* topology on $\mathcal{M}(X)$ is the topology in which $\mu_n \rightarrow \mu$ if $\int f d\mu_n \rightarrow \int f d\mu$ for every continuous function f on X . This makes $\mathcal{M}(X)$ into a compact metrizable space.

For $\mu \in \mathcal{M}(X)$ we denote its measure-theoretic entropy with respect to the shift action by $h(\mu)$. We recall the following facts:

- (1) The entropy function $h : \mathcal{M}(X) \rightarrow \mathbb{R}^+$ is upper semi-continuous.
- (2) The variational principle: $h(X) = \max_{\mu \in \mathcal{M}(X)} h(\mu)$.

See [6] for definitions, proofs and a detailed discussion of the one-dimensional case, or [19] for a proof of the variational principle in the multidimensional case.

3. Computability of entropies

In this section we show that the entropy of an SFT is right recursively enumerable. This follows from the work of Friedland [9], but for completeness we give a short alternative proof and extend the result to sofic systems. We also prove that the entropy of a strongly irreducible SFT is computable.

Let the syntax $L \subseteq \Sigma^F$ define a (possibly empty) SFT X . The definition of entropy provides us with the sequence $N_n = N_X(F_n)$ such that $\frac{1}{n^d} \log N_n$ converges to $h(X)$ from above, and if N_n is computable this sequence shows that h is upper-semi recursive. However, N_n is not computable in general. Indeed, determining whether $N_n > 0$ is equivalent to deciding if the SFT defined by L is nonempty, and this is in general undecidable [22], [2].

Let us say that a finite pattern $a \in \Sigma^{F_n}$ is locally admissible if $a|_{F+u}$ is congruent to a pattern in L whenever $F + u \subseteq F_n$. Instead of N_n , consider the sequence

$$\tilde{N}_n = \#\{\text{locally admissible } F_n\text{-patterns}\}.$$

Clearly \tilde{N}_n is computable. If $x \in X$ then $x|_{F_n}$ is one of the patterns counted by \tilde{N}_n , so $\tilde{N}_n \geq N_n$. The inequality can be strict, because not all locally admissible F_n -patterns need arise in this way: there can be locally admissible finite patterns which do not extend to globally admissible coloring of \mathbb{Z}^d . Nonetheless,

THEOREM 3.1. *For L, X and \tilde{N}_n as above, $\frac{1}{n^d} \log \tilde{N}_n \rightarrow h(X)$ from above. Consequently, $h(X)$ is right recursively enumerable.*

Proof. Denote

$$\tilde{h} = \limsup_{n \rightarrow \infty} \frac{1}{n^d} \log \tilde{N}_n.$$

Since $\tilde{N}_n \geq N_n$ and $\frac{1}{n^d} \log N_n \geq h(X)$, we have $\frac{1}{n^d} \log \tilde{N}_n \geq h(X)$, so that it suffices to show $\tilde{h} \leq h(X)$.

Define a sequence of measures ν_n on $\Sigma^{\mathbb{Z}^d}$ as follows. Let $W_n \subseteq \Sigma^{F_n}$ be the set of locally admissible colorings of F_n . Let μ_n denote the probability measure obtained by coloring each translate $F_n + u$ for $u \in n\mathbb{Z}^d$ independently and uniformly with patterns from W_n . Let $\nu_n = \frac{1}{|F_n|} \sum_{u \in F_n} \sigma^u \mu_n$. Then ν_n is an invariant probability measure and its entropy is easily shown to be

$$h(\nu_n) = \frac{1}{n^d} \log \tilde{N}_n.$$

Let $\nu_{n(k)}$ be a subsequence such that $h(\nu_{n(k)}) \rightarrow \tilde{h}$ and let ν be a weak-* accumulation point $\nu_{n(k)}$; we may assume $\nu_{n(k)} \rightarrow \nu$. Since entropy is upper semi-continuous in the weak-* topology, we have

$$h(\nu) \geq \tilde{h}.$$

On the other hand we claim that $\nu(X) = 1$, and so ν can be regarded as an invariant probability measure on X . To show this, we prove that $\nu([a]) = 0$ for any $a \in \Sigma^F \setminus L$, where $[a]$ is the cylinder set defined by a . Indeed, for every k and $u \in F_k$, if $(F + u) \subseteq F_k + v$ for some $v \in k\mathbb{Z}^d$ then $\mu_k(\sigma^{-u}([a])) = 0$ and so

$$\nu_k([a]) \leq \frac{1}{k^d} \#\{u \in F_k : (F + u) \cap k\mathbb{Z}^d \neq \emptyset\} \leq \frac{k^d - (k - 2 \text{diam } F)^d}{k^d}$$

where $\text{diam } F$ is the diameter of F with respect to the norm

$$\|u\|_\infty = \max_{i=1 \dots d} |u_i|.$$

Therefore,

$$\nu([a]) = \lim_{k \rightarrow \infty} \nu_{n(k)}([a]) = 0.$$

Finally, the variational principle implies that $h(\nu) \leq h(X)$, and the theorem follows. □

With the same notation as above, let $Y \subseteq \Delta^{\mathbb{Z}^d}$ be a symbolic factor of X arising from a one-block map $\varphi_0 : \Sigma \rightarrow \Delta$ and its componentwise extension $\varphi : X \rightarrow Y$. Write

$$\tilde{M}_n = \#\{\varphi(a) \in \Delta^{F_n} : a \in W_n\}$$

where as before, W_n is the set of locally admissible F_n -patterns for L .

THEOREM 3.2. *With the above notation, $\frac{1}{|F_n|} \log \tilde{M}_n \rightarrow h(Y)$ from above. Consequently $h(Y)$ is right recursively enumerable.*

Proof. Denote $\tilde{h}(Y) = \limsup \frac{1}{n^d} \log \tilde{M}_n$. Since φ is onto we have $\tilde{M}_n \geq N_Y(F_n)$, so that $\frac{1}{n^d} \log \tilde{M}_n \geq h(Y)$. Thus we only need to show $\tilde{h}(Y) \leq h(Y)$.

Let θ_k be measures on $\Delta^{\mathbb{Z}^d}$ defined by coloring each translate $F_k + u$ for $u \in k\mathbb{Z}^d$ with patterns drawn uniformly from $\{\varphi(a) : a \in W_k\}$. Then $\eta_k = \frac{1}{k^d} \sum_{u \in F_k} \sigma^u \theta_k$ is an invariant measure on $\Delta^{\mathbb{Z}^d}$, and $h(\eta_k) = \frac{1}{k^d} \log \tilde{M}_k$. Let μ_k be measures on $\Sigma^{\mathbb{Z}^d}$ such that the pattern on $F_k + u$ for $u \in k\mathbb{Z}^d$ is drawn from W_k according to a distribution which projects under φ to the uniform distribution on $\{\varphi(a) : a \in W_k\}$. Thus $\theta_k = \varphi(\mu_k)$. Let $\nu_k = \frac{1}{k^d} \sum_{u \in F_k} \sigma^u \mu_k$, so that $\eta_k = \varphi(\nu_k)$. Choose a subsequence $\eta_{n(k)}$ so that there is a measure η on Y with $\eta_{n(k)} \rightarrow \eta$ and there is a measure ν on X with $\nu_{n(k)} \rightarrow \nu$; now, $\eta = \varphi(\nu)$ satisfies $\eta_{n(k)} \rightarrow \eta$. By upper semi-continuity, $h(\eta) \geq \lim h(\eta_{n(k)}) = \tilde{h}$, and so we will be done if we show that η is supported on Y . For this it is enough to show that ν is supported on X , i.e. that $\nu([b]) = 0$ whenever $b \in \Sigma^F \setminus L$. The proof of this is identical to the proof of the same statement at the end of Theorem 3.1. \square

COROLLARY 3.3. *The entropy of every sofic shift is right recursively enumerable.*

Proof. As noted in Section 2, every sofic shift is a one-block factor of some SFT. \square

We turn now to strongly irreducible SFTs and the proof of Theorem 1.3. Let $Q_n = \{-n, -n + 1, \dots, n, n + 1\}^d$ denote the symmetric cube. Recall that the boundary of the cube Q_n is $\partial Q_n = Q_n \setminus Q_{n-1}$. For $n > k$ and patterns $a \in \Sigma^{Q_k}$ and $b \in \Sigma^{\partial Q_n}$ we say that a, b are compatible on Q_n if there is a locally admissible pattern $c \in \Sigma^{Q_n}$ with $c|_{Q_k} = a$ and $c|_{\partial Q_n} = b$. Such a c is called a completion of a, b .

We state the following for 1-step SFTs, though it is easily adapted to the general case.

LEMMA 3.4. *Let $X \subseteq \Sigma^{\mathbb{Z}^d}$ be a nonempty strongly irreducible 1-step SFT and $a \in \Sigma^{Q_k}$. Given N , consider the following conditions:*

- (1) $a \neq b|_{Q_k}$ for every locally admissible $b \in \Sigma^{Q_N}$.
- (2) a and $b|_{\partial Q_N}$ are compatible on Q_N for every locally admissible $b \in \Sigma^{Q_N}$.

Then

- (a) a fails to be globally admissible if and only if (1) holds for some N (equivalently, all sufficiently large N).
- (b) a is globally admissible if and only if (2) holds for some N (equivalently, all sufficiently large N).

Proof. (a) is clear: If a appears in an infinite configuration then it appears at the center of locally admissible patterns on cubes of all sizes. The converse follows by compactness.

We turn to the proof of (b). First, note that if (2) holds for N then it holds for all $M > N$. Indeed, if $b \in \Sigma^{\mathcal{Q}^M}$ is locally admissible then so is $b|_{\mathcal{Q}_N}$. Hence $a, b|_{\partial\mathcal{Q}_N}$ are compatible on \mathcal{Q}_N . Let c be a completion for them, so that $c|_{\partial\mathcal{Q}_N} = b|_{\partial\mathcal{Q}_N}$. Since the SFT is 1-step we can define $c' \in \Sigma^{\mathcal{Q}^M}$ by $c'(u) = c(u)$ for $u \in \mathcal{Q}_N$ and $c'(u) = c(u)$ otherwise, and we get a locally admissible pattern c' with $c'|_{\mathcal{Q}_k} = a$ and $c'|_{\partial\mathcal{Q}_M} = b|_{\partial\mathcal{Q}_M}$. Hence $a, b|_{\partial\mathcal{Q}_M}$ are compatible on \mathcal{Q}_M .

Suppose (2) holds for some N , so that $N > k$. Choose any $x \in X$ and let $b = x|_{\mathcal{Q}_N}$. Then b is locally admissible, and so by assumption we can choose a completion c for a, b . Since $c|_{\partial\mathcal{Q}_N} = b|_{\partial\mathcal{Q}_N} = x|_{\partial\mathcal{Q}_N}$, we can redefine $x(u) = c(u)$ for $u \in \mathcal{Q}_N$ and obtain a globally admissible pattern y with $y|_{\mathcal{Q}_N} = a$. Thus a is globally admissible.

Conversely, assume that a is globally admissible; we will establish (2) for all sufficiently large N . Let r be a gap for X , so that for $x \in X$ we have $a, x|_{\partial\mathcal{Q}_{k+r+1}}$ are compatible on \mathcal{Q}_{k+r+1} . Thus, for every locally admissible $b \in \Sigma^{\mathcal{Q}_{k+r+1}}$, if $a, b|_{\partial\mathcal{Q}_{k+r+1}}$ are not compatible on \mathcal{Q}_{k+r+1} then b is not globally admissible. By (a) we see that for sufficiently large N every locally admissible $b' \in \Sigma^{\mathcal{Q}_N}$ satisfies $b'|_{\mathcal{Q}_{k+r+1}} \neq b$. Since there are finitely many such b 's, it follows that if N is large enough, then for every locally admissible $b' \in \Sigma^{\mathcal{Q}_N}$ the patterns $a, b'|_{\partial\mathcal{Q}_{k+r+1}}$ are compatible on \mathcal{Q}_{k+r+1} . For such an N , this implies (by the splicing argument at the beginning of the proof of (b)) that $a, b'|_{\partial\mathcal{Q}_N}$ are compatible on \mathcal{Q}_N , as required. \square

COROLLARY 3.5. *For a nonempty strongly irreducible SFT X it is decidable whether a finite pattern a is globally admissible.*

Proof. To decide if a is globally admissible, iterate over $N = 1, 2, 3, \dots$ and find the first N for which condition (1) or (2) in the lemma holds (for each N the conditions are finitely checkable). If (2) holds then a is globally admissible; if (1) holds it is not. The lemma guarantees this algorithm halts in finite time and gives the correct answer. Note that to apply the lemma one does not need to know the gap. \square

We remark that an old argument of Wang gives an algorithm that decides global admissibility of a pattern in an SFT, assuming the SFT has dense periodic points; see [22]. It is known that strong irreducibility implies dense periodic points in dimension 2, so that in this case the corollary above does not give new information. Whether strong irreducibility implies dense periodic points for $d \geq 3$ seems to be open.

Returning to the question of computability of entropy, we say that a number h is left recursively enumerable if there is an algorithm which, given n , produces a rational number $s(n)$ with $s(n) \rightarrow h$ and $s(n) \leq h$. If h is both right and left recursively enumerable then it is computable. To see this let $r(n), s(n)$ be computable sequences with $s(n) \leq h \leq r(n)$ and $r(n), s(n) \rightarrow h$. Now given n , we can calculate $r(k), s(k)$ for $k = 1, 2, 3 \dots$ until such a k is reached that $r(k) - s(k) < \frac{1}{n}$. Then $r(k)$ satisfies $|r(k) - h| < \frac{1}{n}$. This algorithm shows that h is computable.

We can now prove Theorem 1.3, which we repeat here for convenience:

THEOREM. *The entropy of a strongly irreducible SFT is computable.*

Proof. Let X be a strongly irreducible SFT, and we may assume it is nonempty. We already know that $h(X)$ is right recursively enumerable, so it suffices to show that it is left recursively enumerable, i.e. to exhibit an algorithm which given $n \in \mathbb{N}$ returns a rational number $s(n)$ such that $s(n) \rightarrow h(X)$ and $s(n) \leq h$.

The algorithm is as follows. First, identify all the globally admissible patterns $a_1, \dots, a_{k(n)} \in \Sigma^{Q_n}$ (this is computable by the corollary above). With this notation we have $k(n) = N_X(Q_n) = N_X(F_{(2n+1)^d})$, and $\frac{1}{|Q_n|} \log k(n) \rightarrow h(X)$. Next, find the smallest number r' so that each globally admissible pattern $b \in \Sigma^{Q_{n+r'+1}}$ is compatible on $Q_{n+r'+1}$ with a_i for $i = 1, \dots, k(n)$. Set

$$s(n) = \frac{1}{|Q_{n+r'}|} \log k(n).$$

Note that $r' \leq r$, where r is the gap for X . Hence

$$s(n) \geq \frac{|Q_n|}{|Q_{n+r}|} \cdot \frac{1}{|Q_n|} \log k(n) \rightarrow h(X).$$

On the other hand, consider a large Q_m , and consider the collection of translates of Q_n by elements of the lattice $2(n+r')\mathbb{Z}^d$ which fall inside Q_m . By choice of r' we can color each of these translates in an arbitrary globally admissible way and complete it to a globally admissible Q_m pattern. Since the number of translates is $\frac{1}{|Q_{n+r'}|} |Q_m|$ (for convenience assume that m is a multiple of $2(n+r')$, we see that

$$k(m) \geq k(n) \frac{1}{|Q_{n+r'}|} |Q_m|.$$

Thus, letting $m \rightarrow \infty$, we have

$$s(n) = \frac{1}{|Q_{n+r'}|} \log k(n) \leq \frac{1}{|Q_m|} \log k(m) \rightarrow h(X);$$

hence $s(n) \leq h(X)$, and also $s(n) \rightarrow h(X)$, as desired. \square

Note that the algorithm given in the proof does not require prior knowledge of a gap for the X . It may of course be applied to a nonirreducible SFT, but in that case may not halt on some inputs, and even if it does, the sequence $s(n)$ will not necessarily behave as above.

4. Outline of the main construction

Let h be a right recursively enumerable number. To prove the remaining direction of Theorem 1.1, we must construct for every $d \geq 2$ a d -dimensional SFT with entropy h . We first make some simplifying assumptions. We may restrict ourselves to dimension 2, since given an SFT $X \subseteq \Sigma^{\mathbb{Z}^d}$ the system $X' \subseteq \Sigma^{\mathbb{Z}^{d+1}}$ defined by

$$X' = \{x' \in \Sigma^{\mathbb{Z}^{d+1}} : \forall j \in \mathbb{Z} \exists x \in X \forall u \in \mathbb{Z}^d x'(u, j) = x(u)\}$$

is easily seen to be a $d + 1$ -dimensional SFT and $h(X') = h(X)$. Furthermore, since (a) the product of SFTs is an SFT, (b) $h(X \times Y) = h(X) + h(Y)$ and (c) n is the entropy of the full shift on 2^n symbols, it suffices to prove the statement under the assumption that $h \in [0, 1]$.

Our construction has three main steps:

Step 1: Constructing the base (§6). We construct an SFT X some of whose symbols are marked 0, 1, and such that the density of 1's in each point of X is very uniform. It will be possible to estimate this density by observing any sufficiently large and well-distributed set of coordinates.

Step 2: Pruning (§7). In this step we “kill” all points $x \in X$ such that the frequency of 1's in x is strictly greater than h . In this way we obtain an SFT Y such that the symbol 1 appears in each $y \in Y$ with frequency at most h , and for some points the frequency is h . Furthermore, Y will still have zero entropy. We achieve this by superimposing another layer on top of X which represents calculations of a certain Turing machine, using as input the underlying patterns from X . This machine halts when it detects a density of 1's greater than h . The result is that a point $x \in X$ with density of 1's greater than h cannot be extended to a pattern in Y ; otherwise, it can be.

Step 3: Adding “Random” bits (§8). We extend Y to an SFT Z by allowing two new symbols, say “ α ” and “ β ”, to appear independently over every occurrence of a 1 in Y . This system Z has entropy h .

For Steps 1 and 2 we utilize certain SFTs with special geometric and arithmetic properties. The existence of such systems, and their use in representing Turing machines in SFTs, appears first in Robinson’s paper [22]. However, we will not refer directly to Robinson’s construction, which would in any case require some modification to suit our needs. Instead we rely on a theorem of Mozes [20] about the realization of substitution systems by SFTs. This theorem, which allows us to easily construct variants of Robinson’s system, is presented in the next section together with another technical definition. Following that, we give the details of Steps 1, 2 and 3.

Before moving on, we note that our arguments give the following result, which may be of independent interest:

THEOREM 4.1. *A real number $r \geq 0$ is right recursively enumerable if and only if there is an alphabet Σ , a symbol $a \in \Sigma$ and an SFT $X \subseteq \Sigma^{\mathbb{Z}^d}$ such that*

$$\sup_{x \in X} \lim_{n \rightarrow \infty} \frac{1}{|F_n|} \#\{u \in F_n : x(u) = a\} = r$$

(and in particular the limit above exists for every $x \in X$). Furthermore if r is computable then one can find Σ , a and an SFT X so that $\lim_{n \rightarrow \infty} \frac{1}{|F_n|} \#\{u \in F_n : x(u) = a\} = r$ for every $x \in X$.

5. Substitutions and superpositions

In this section we describe two technical devices for constructing SFTs.

5.1. Subshifts defined by substitution. Given a finite alphabet Σ , a *substitution rule* is a map $s : \Sigma \rightarrow \Sigma^{F_k}$ for some integer $k > 1$, where $F_k = \{1, \dots, k\} \times \{1, \dots, k\}$ (in the terminology of [20], this is a deterministic $k \times k$ substitution system with property A). The map s extends naturally to a map $s_n : \Sigma^{F_n} \rightarrow \Sigma^{F_{n \cdot k}}$ by identifying $\Sigma^{F_{n \cdot k}}$ with $(\Sigma^{F_k})^{F_n}$.

Starting from a single symbol located at $(1, 1) \in \mathbb{Z}^2$ and iterating the substitution map, we obtain a sequence of colorings of F_{k^n} for $n = 0, 1, 2, \dots$. Such patterns are called s -blocks. A point $x \in \Sigma^{\mathbb{Z}^2}$ is admissible for s if every finite subpattern of x appears in some s -block. The subshift $W \subseteq \Sigma^{\mathbb{Z}^2}$ associated with s is the set of admissible patterns; this is seen to be closed and shift invariant.

Define $s_\infty : W \rightarrow W$ by applying s to each symbol of x ; more precisely, $s_\infty(x)(u) = s(x(u'))(u'')$, where $u' \in \mathbb{Z}^2$ and $u'' \in F_k$ are the unique vectors such that $u = ku' + u''$. Clearly s_∞ maps W into W . We say that x is derived from y if

$T^v x = s_\infty(y)$ for some $v \in F_k$. It is not hard to show that each $x \in W$ is derived from some $y \in \tilde{W}$; if this y is unique, we say that s has *unique derivation*.

THEOREM 5.1 (Theorem 4.5 of [20]). *Let $s : \Sigma \rightarrow \Sigma^{F_k}$ be a substitution rule with unique derivation and let W be the associated dynamical system. Then there exists an alphabet Δ , an SFT $\tilde{W} \subseteq \Delta^{\mathbb{Z}^2}$, and a one-block factor map $\varphi : \tilde{W} \rightarrow W$. Furthermore φ is an injection on a set having full measure with respect to every invariant measure on \tilde{W} .*

Note that Theorem 5.1 is false in dimension $d = 1$.

PROPOSITION 5.2. *If s , W and \tilde{W} are as in Theorem 5.1, then $h(\tilde{W}) = 0$.*

Proof. For any μ invariant on \tilde{W} , the map φ is an isomorphism of dynamical systems between (\tilde{W}, μ) and $(W, \varphi\mu)$ where $\varphi\mu$ is the push-forward of μ to \tilde{W} . Hence it suffices to show that the latter system has zero measure-theoretic entropy. By the variational principle it suffices to show that $h(W) = 0$. Fix m . Since every large enough s -block is composed of an array of smaller s -blocks of dimension $k^m \times k^m$ arranged in a square, it follows that for $n > k^m$ an admissible F_n -pattern can decompose F_n into $(\lfloor \frac{n}{k^m} \rfloor - 2)^2$ disjoint s -blocks of dimension $k^m \times k^m$ together with a “small” remaining region near the boundary. Thus the number of F_n patterns is at most

$$N_{F_n}(W) \leq \#\{k \times k \text{ } s\text{-blocks}\}^{(\lfloor n/k^m \rfloor - 2)^2} \cdot |\Sigma|^{4nk^m}$$

where the second term on the right-hand side is the number of ways to fill in the region near the boundary of F_n not covered by the s -blocks. Since there are only $|\Sigma|$ different s -blocks of dimension $k^m \times k^m$ (because each is derived from one of the original symbols), for all large enough n we have

$$\frac{1}{n^2} \log N_{F_n}(W) \leq \frac{(\lfloor \frac{n}{k^m} \rfloor - 2)^2 \log |\Sigma|}{n^2} + \frac{4k^m \log |\Sigma|}{n} \rightarrow \frac{1}{k^m};$$

as m was arbitrary, $h(W) = 0$. □

We use Theorem 5.1, which is due to Mozes, to construct systems similar in many respects to Robinson’s system from [22]. We remark that although this allows a more economical exposition the gain is cosmetic. Indeed, the proof of Theorem 5.1 relies on an elaborate extension of Robinson’s techniques. There has recently been a revival of interest in substitutions and their realization using local rules; see e.g. [11].

5.2. Superposition. Given an SFT X defined by a syntax L , *superposition* is a syntactic process which gives an SFT X' which factors onto a subshift of X . Informally, this is done by adding data to each symbol of X and enriching the syntax with rules relating to the new data.

More precisely, suppose X is an SFT defined by a syntax $L \subseteq \Sigma^F$. A system Y is superimposed over X if it is obtained by the following process. (a) For a finite set Δ , we replace each symbol of $\sigma \in \Sigma$ with one or more symbols of the form $(\sigma, \delta) \in \Sigma \times \Delta$. Let Σ' be the set of these pairs. For the new symbol $(\sigma, \delta) \in \Sigma'$, we say that δ is superimposed over σ ; we also frequently refer to this pair as the symbol σ marked with δ . (b) We extend each pattern $a \in L \subseteq \Sigma^F$ to one or more patterns $a' \in (\Sigma')^F$ by superimposing new symbols over each symbol of a . Call the new syntax L' . The SFT X' defined by L' has the property that every pattern appearing in X' consists of a Δ -pattern superimposed over a Σ -pattern, and the Σ -pattern is admissible for X .

Note that the map $\pi : X' \rightarrow \Sigma^{\mathbb{Z}^2}$ which erases the superimposed layer of data maps X' into a subsystem of X . We say that $x \in X$ is represented in X' if one can turn x into a point of X' by superimposing a suitable Δ -pattern over x ; i.e., if $x = \pi(x')$ for some $x' \in X'$.

6. Step 1: Constructing the base

In this section we construct a two-dimensional SFT X whose symbols are marked with the symbols 0, 1. The symbol 1 may appear with any density in points of x , but for each fixed $x \in X$ the density of 1's will be extremely uniform.

6.1. *An almost periodic SFT.* Consider the substitution on the alphabet $\{\circ, \bullet\}$ defined the rule

$$\bullet \mapsto \begin{pmatrix} \circ & \bullet \\ \bullet & \circ \end{pmatrix} \quad \circ \mapsto \begin{pmatrix} \circ & \circ \\ \bullet & \circ \end{pmatrix}.$$

Let W denote the dynamical system defined by these rules. See Figure 6.1.

We say that a set $E \subseteq \mathbb{Z}^2$ is a 2-net if $E = \cup_{n=1}^{\infty} I_n \times J_n$ where each I_n and J_n are translates of $2^n\mathbb{Z}$, the I_n 's are pairwise disjoint, and the J_n 's are pairwise disjoint. We refer to $I_n \times J_n$ as the n -th level associated with E . Note that if u belongs to some level of E then the row and column to which u belongs do not intersect any other level.




Figure 6.1. Three iterations of the substitution.

PROPOSITION 6.1. *Let $w \in W$ and*

$$E(w) = \{u \in \mathbb{Z}^2 : w(u) = \bullet\}.$$

Then $E = E(w)$ is a 2-net.

Proof. For $n = 0, 1, 2, \dots$ let a_n be the sequence of F_{2^n} square patterns obtained by applying the substitution rule to the initial symbol $a_0 = \bullet$. It is sufficient to show that there is a 2-net $E = \cup_{n=1}^{\infty} I_n \times J_n$ such that $\{u \in F_{2^n} : a_n(u) = \bullet\} = E \cap F_{2^n}$. To verify this, one proves by induction that the above holds for

$$I_n = J_n = 2^n \mathbb{Z} + 2^{n-1}. \quad \square$$

We remark that the system W supports a unique invariant probability measure and as a measure-preserving system this is an odometer, i.e. is isomorphic to a zero-dimensional abelian group along with a free minimal \mathbb{Z}^2 action generated by translation by two elements of the group.

This substitution rule has a unique derivation, since one may check that there is a unique way to derive the central 6×6 square of the 8×8 pattern in Figure 6.1 from a 4×4 pattern.

Let \widetilde{W} be the SFT associated to W by Theorem 5.1. Then to each point in \widetilde{W} there is associated, via a one-block map, a $\{\circ, \bullet\}$ pattern defining a 2-net.

6.2. *Marking the columns of \widetilde{W} .* We now superimpose another layer on top of \widetilde{W} . Begin by superimposing the symbols 0, 1 on top of the \widetilde{W} with the constraint that the symbols 0, 1 cannot be placed vertically adjacent to each other. This forces each column in the resulting system to be marked either entirely with 0's or entirely with 1's.

For a point $w \in \widetilde{W}$, the new coloring induces a $\{0, 1\}$ -coloring of each level $I \times J$ in the decomposition given by the proposition. This coloring is constant on the intersection of $I \times J$ with columns; we now force it to be constant on the intersection of the grid with rows. For this, superimpose two new symbols " \longleftrightarrow ," " \longleftrightarrow " on top of the existing ones. We think of \longleftrightarrow as transmitting a "0" signal, and of \longleftrightarrow as transmitting a "1" signal. The rules are that over a symbol marked \bullet , the symbol \longleftrightarrow appears always together with the symbol 0, and \longleftrightarrow appears always together with the symbol 1. We also require that \longleftrightarrow and \longleftrightarrow cannot appear as horizontal neighbors, so the arrow type is constant on rows.

Call the resulting system X (it is of course an SFT) and let $x \in X$ be superimposed over a point $w \in \widetilde{W}_0$. Let $I \times J$ be some level of the 2-net induced by w , and suppose that $w(u)$ is marked 0 for some $u \in I \times J$. Since it is also marked \bullet , it bears the symbol \longleftrightarrow (and not \longleftrightarrow); this forces the entire row to which u belongs to be marked with \longleftrightarrow . Every other $v \in I \times J$ belonging to the same row

is thus marked \bullet and \longleftrightarrow , and so it must be marked 0. A similar analysis holds if $w(u)$ is marked 1.

In short, the 0, 1-coloring of each grid $I \times J$ is constant on rows and columns, and thus is completely constant. If $I_n \times J_n$ are the levels of the 2-net induced by a point $x \in X$ then each $I_n \times J_n$ determines a collection of columns which is 2^n -periodic in the horizontal direction, and all these columns bear the same symbol 0 or 1.

For $x \in X$, let $\delta(x)$ be the upper density of 1's in x , i.e.

$$\delta(x) = \limsup_{n \rightarrow \infty} \frac{|\{u \in F_n : x(u) = 1\}|}{|F_n|}$$

where as usual $F_n = \{1, \dots, n\}^2$. If $I_n \times J_n$ are the levels of the 2-net induced by x ; then a simple calculation shows that

$$\delta(x) = \sum_{n=1}^{\infty} \rho_n \cdot 2^{-n}$$

where ρ_n is 0 or 1 according to the coloring x induces on $I_n \times J_n$. Since the I_n 's and J_n 's are pairwise disjoint the arrows transmitting information between the points of each grid occupy different rows, and hence do not interact. Therefore, we are free to color each level 0 or 1 independently of the coloring of the other levels. Consequently, any sequence $\rho_n \in \{0, 1\}$ may arise, and so there are points $x \in X$ with $\delta(x)$ taking on any value in the range $[0, 1]$.

We will call a point in X exceptional if it is superimposed over an exceptional point of \tilde{W} . For an exceptional point $x \in X$ there are complementary half-spaces and/or quarter-spaces such that the restriction of x to each of them looks like a nonexceptional point. Thus the above analysis applies to each of these regions separately. This is not to say that we can glue admissible half- and quarter-spaces together arbitrarily, and indeed for exceptional points the arrows from different parts can interact; but this will not matter to us.

Finally, we claim that X has zero entropy. Indeed, W has zero entropy, and it is simple to check that if a is a square pattern admissible for W then every extension of a to a pattern b admissible for X is determined by the symbols of b on the boundary of the square. It follows that X has entropy 0.

7. Step 2: Pruning

Let h be a fixed, right, recursively enumerable number. Let X be the system constructed in the previous section. Our goal in this section is to construct an SFT Y superimposed over X which “kills” points with density of 1's greater than h .

More precisely, we will want

$$\sup\{\delta(y) : y \in Y\} = h$$

(here δ is the natural extension of δ from X to Y) and that the supremum will be achieved.

7.1. *Boards.* We define a substitution system over the alphabet

$$\Sigma = \{ |, -, \ulcorner, \urcorner, \llcorner, \lrcorner, \top, \perp, \vdash, \dashv, +, \blacksquare, \square \}.$$

The substitution rules are described in Figure 7.1 together the symmetric rules obtained by rotating by multiples of 90° . Let us denote by b_n the $5^n \times 5^n$ -pattern obtained by applying the substitution rule n times to the symbol \blacksquare . It is not hard to show that \blacksquare appears with period 5 in every b_n . Given k and $n > k$, since \blacksquare appear in b_{n-k} with period 5 we see that b_i appears in b_{n-k+i} with period 5^i , so that b_k appears in b_n with period 5^k .

As can be seen from Figure 7.1, this substitution rule produces patterns which induce certain grid-like shapes on \mathbb{Z}^2 . More precisely, define finite sets $I_n \subseteq \mathbb{N}$ inductively by $I_1 = \{1, 2, 4, 5\}$ and

$$I_{n+1} = I_n \cup (I_n + 5^n) \cup (I_n + 3 \cdot 5^n) \cup (I_n + 4 \cdot 5^n).$$

One sees by induction that $\min I_n = 1$ and $\max I_n = 5^n$, so that the union above is disjoint, and hence $|I_n| = 4^n$ for each n . Let

$$B_n = (I_n \times \{1, 2, \dots, 5^n\}) \cup (\{1, 2, \dots, 5^n\} \times I_n).$$

This is the set obtained by “filling in” the rows and columns between points of $I_n \times I_n$. The set B_n is called an n -board. An infinite board is any set $B \subseteq \mathbb{Z}^2$ which is the limit of a sequence of translates $B_n + u_n$ for some $u_n \in \mathbb{Z}^d$, where by limit we mean that $u \in B$ if and only if eventually $u \in B_n + u_n$. It is simple to check that every infinite board B has density zero; i.e. for every $\varepsilon > 0$ there is an N so that $\frac{|B \cap (F_N + u)|}{|F_N|} < \varepsilon$ for every $u \in \mathbb{Z}^2$. This follows from the recursion formula for I_n .

Consider the patterns $b_n|_{B_n}$. One shows by induction that these patterns do not contain the symbols \blacksquare, \square . Also, for $u \in B_n$ the points $v \in B_n$ which are adjacent to u – i.e., which differ from u by $\pm e_1$ or $\pm e_2$ – are determined by $b_n(u)$ by interpreting the symbol $b_n(u)$ as a collection of lines pointing to the neighbors of u in B_n . Thus, \perp indicates that there are neighbors left, right and above the current symbol; $-$ indicates neighbors to the left and right of it, etc. One can show that if $u \in F_{5^n} = \{1, \dots, 5^n\} \times \{1, \dots, 5^n\}$ and $b_n(u) \notin \{\blacksquare, \square\}$, then there are a unique k and translate A of B_k so that $u \in A$ and $b_n|_A$ is congruent to $b_k|_{B_k}$. In the large square in Figure 7.1 there are two boards visible; a 1-board in the center, and a 2-board surrounding it. If we iterate the substitution one more step, each \blacksquare

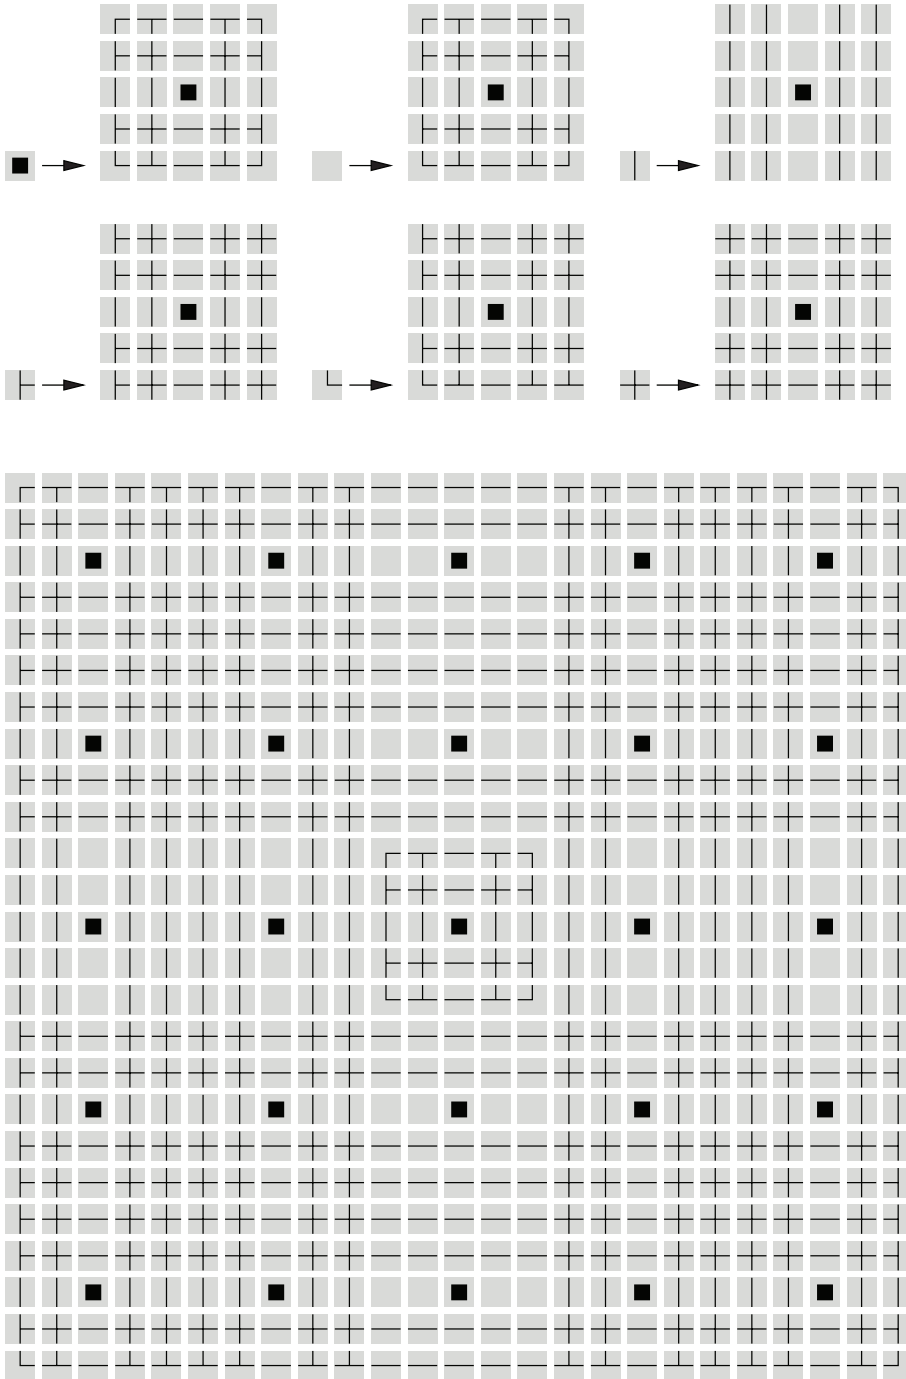


Figure 7.1. The substitution rules, up to rotation. The symbol \square is represented as an empty square. The large 25×25 pattern is obtained by applying the substitution rules twice to \blacksquare .

will turn into a 1-board plus \blacksquare 's; the 1-board will turn into a 2-board plus \blacksquare 's, and the 2-board will turn into a 3-board, plus \blacksquare 's.

Let R denote the dynamical system defined by these rules. From the remarks above it follows that each $r \in R$ determines a pairwise disjoint collection of boards, with n -boards appearing periodically with period 5^n ; and if $u \in r$ and $r(u) \notin \{\blacksquare, \square\}$, then u belongs to one of these boards and the neighbors of u in this board can be determined from symbols $r(u)$. By compactness, there will exist points $r \in R$ and infinite boards B so that $r|_B$ is marked similarly to a finite board. Since infinite boards cannot overlap and each occupies at least some quarter-space, there can be at most four infinite boards in r , and since each has density 0, the density of points belonging to infinite boards in r is zero.

R has unique derivation; indeed, the locations of the corner tiles determine the derivation of a point. We denote by \tilde{R} the SFT associated to R by Mozes' theorem. We identify points in \tilde{R} with the point in R they are mapped to by the given one-block map; in general this identification is many-to-one.

7.2. Turing machines and their representations in SFTs. A Turing machine is an automaton with a finite number of internal states which reads and writes data on a one-sided infinite array of *cells* indexed by \mathbb{N} , called the *tape*. Each cell contains one symbol from the data alphabet (so in our model the input is an infinite sequence). The computation begins with the machine located at the 0-th (leftmost) cell and in a special initial state, and the tape contains some data which is the input to the computation. The state of the data tape along with the location and internal state of the machine are called a *configuration*; a configuration uniquely determines all future configurations. The computation proceeds in discrete time steps. At each iteration the machine is located at some cell, reads the symbol written there and based on this data and on its internal state, performs three actions: (a) it replaces the current data symbol with a new one, (b) it moves one cell to the left or to the right, and (c) it updates its internal state. The computation may halt after a finite number of steps if the machine either moves off the tape (steps left at cell 0) or enters a designated state, called the halting state. Barring these occurrences, the computation continues forever.

Although a very simple model, any algorithm written in a modern computer programming language can be implemented as a Turing machine, and it is generally accepted that any effective computation can be performed by a Turing machine; this is Church's thesis. For background and basic facts on this subject, see [12].

Let X be the SFT constructed in Section 6, let \tilde{R} be the SFT described above and let T be a Turing machine whose data alphabet includes symbols 0, 1. We construct an SFT Y_T superimposed over $X \times \tilde{R}$ such that when a point $y \in Y$ is superimposed over $(x, r) \in X \times \tilde{R}$, each board induced by r has superimposed over it a pattern representing the run of T on the input given by the sequence of

0, 1's appearing in x along the columns of the board. This construction, which we describe next, is similar to the one used by Robinson in [22], except that Robinson's machines always ran on an "empty" input.

Let ξ, ρ be symbols in the alphabets of X, \tilde{R} respectively. We superimpose new symbols over (ξ, ρ) only if ρ represents a point in a board (i.e. $\rho \neq \square, \blacksquare$), and the adjacency rules for the new symbols will only restrict pairs of neighbors which belong to the same board (note that this can be determined locally). Thus $(x, r) \in X \times \tilde{R}$ will be represented in Y_T if and only if for each (finite or infinite) board B induced by r there exists a locally admissible pattern superimposed over $(x, r)|_B$.

For a board $B_n + u$ let us call the points $I_n \times I_n + u$ the *nodes* of the board. Note that (ξ, ρ) represents a node if and only if $\rho \in \{\ulcorner, \lrcorner, \llcorner, \lrcorner, \top, \perp, \vdash, \dashv, +\}$. The data superimposed over a node will include a combination of data symbol (from the machine's data alphabet) and possibly also a machine state; this information may be represented by the alphabet $\Delta_1 \cup (\Delta_1 \times \Delta_2)$ where the union is disjoint, Δ_1 is the machine's data alphabet and Δ_2 its state space.

Each row of nodes in a board is to represent a finite portion of the configuration of the machine. More precisely, each node will contain either a data symbol or a data symbol and a machine state; this is called the cell's configuration. Suppose $x \in X$ and $r \in \tilde{R}$ induces a board B . We can arrange things so that

- (1) The data symbols in the nodes of the bottom row are the symbol 0 or 1 induced by x on that node.
- (2) The node at the lower left corner of B contains the initial state of the machine, and no other node in the bottom row contains a machine state.
- (3) Each row of nodes except the bottom one represents the configuration obtained by iterating the computation one step from the configuration given in the row below it. In particular, no row can appear admissibly above a row containing a halting state.

Properties (1) and (2) are easily implemented by restricting the types of symbols which may be superimposed over (ξ, ρ) when $\rho \in \{\lrcorner, \dashv, \perp, \lrcorner\}$.

Implementing (3) with local rules requires a little more effort. First, note that in the course of the operation of a Turing machine T , the configuration of a cell i at a time $t > 1$ is a function of the configurations of the cells $i - 1, i, i + 1$ at time $t - 1$; indeed the data on the cell is determined by the configuration at i , and the presence and state of the machine depend on the configurations of the cells at $i - 1, i + 1$ (in case $i = 0$, the dependence is on the cells at $i, i + 1$ only). We write $T(u, v, w)$ for the state of i at time $t + 1$ given that at time t cells $i - 1, i, i + 1$ were in states u, v, w respectively (we allow $u = \text{"null"}$ in case $i = 0$). If we forget the geometry of the boards and imagine configurations of the machine represented as sequences of cell configurations stacked one on top of the other, this transition is

“local” and can be enforced by a local rule that every pattern of the form $\begin{array}{c} v' \\ u \quad v \quad w \end{array}$ must satisfy $v' = T(u, v, w)$.

However, when we represent cell configurations in nodes of a board the transition from row to row is no longer local, since in a board the nodes representing successive cells are spread out in space and may be arbitrarily far apart. We can overcome this by using the rows and columns between nodes (which belong to the board, and therefore do not overlap for distinct boards) to “transmit information”. In this way we can guarantee that the symbol superimposed over the immediate neighbors of each node indicate the cell configuration at each of the neighboring nodes. This can be implemented in a manner similar to the way in which we synchronized the coloring of 2-nets in X in Section 6. Briefly, over each grid point marked $-$ we superimpose a pair of symbols (u, v) where u, v are node configurations. We require that each pair of horizontally adjacent $-$'s are marked with the same pair, so all members of an uninterrupted horizontal sequence of $-$'s carry the same pair. When a pair $+ -$ appears and $+$ has configuration u we require that over $-$ there is a pair (u, v) for some v ; and similarly for pairs $\perp -$ and $\top -$. The symmetric condition is imposed for $- +$, $- \perp$ and $- \top$. The result is that every uninterrupted horizontal sequence of $-$'s carries the pair (u, v) where u is the configuration of the node at which the sequence ends on the left, and v the configuration of the node ending the sequence on the right.

Next, over each symbol $|$ we superimpose a triple (u, v, w) , where u, v, w are cell configurations and u or w may also be “blank”. As for $-$'s, we require that the marking is constant for each uninterrupted vertical sequence of $|$'s. The markings are determined as follows. If a $|$ is located immediately above a node with configuration v , and the nodes to the left and right of that node have configurations u, w respectively, then $|$ carries (u, v, w) ; u or w is “blank” in the case there is no node to the left or right of the node below $|$ (i.e. if it is at the edge of the board). Note that by the previous discussion, u, v, w may be determined by looking at the immediate neighbors of the node below the $|$. Thus the column of $|$'s above each node represents the configuration of that node and its neighbors. Finally, we require that when a node $+, \top, \ulcorner$ or \lrcorner in state v' appears vertically above a $|$ marked (u, v, w) , then $v' = T(u, v, w)$; this is the point where we encode the transition rules of the Turing machine in the rules of the SFT. These conditions can be seen to force property (3).

We summarize this construction and its properties in the following proposition:

PROPOSITION 7.1. *Given the systems X, \tilde{R} from Sections 6 and 7.1 respectively, and given a Turing machine T , there exists an SFT Y_T superimposed over $X \times \tilde{R}$ such that the following are equivalent:*

- (1) $(x, r) \in X \times \tilde{R}$ is represented in Y_T .
- (2) For each finite or infinite board B induced by r and containing the symbol \perp , when T is run on the sequence of 0, 1-s induced by x on the columns of B the number of steps it runs without halting is at least equal to the number of rows in B .

Furthermore, $h(Y_T) = 0$.

Proof. The equivalence follows from the discussion preceding the theorem. We only note that if a board B induced by r does not contain the symbol \perp then it can always be extended, e.g. by a pattern in which all rows are the same and contain only data. Note that in general, there may be infinitely many ways to superimpose a pattern over an infinite board which does not contain \perp . Thus the projection from Y_T into $X \times \tilde{R}$ is not an injection.

It remains to check that $h(Y_T) = 0$. Given an $N \times N$ pattern a appearing in $X \times \tilde{R}$, if $B_n + u$ is a board induced by \tilde{R} and contained in F_N then there is a unique way to extend a to a locally admissible Y_T pattern. This is true also for symbols in a which do not lie in any board. Given $\varepsilon > 0$, a simple estimate shows that if N is large enough these points make up all but an ε -fraction of the points in F_N , the remaining points coming from “boards” which intersect the boundary of F_N or infinite boards, all of which have density tending to zero as $N \rightarrow \infty$. Hence a can be completed in at most $2^{\varepsilon(N) \cdot N^2}$ ways with $\varepsilon(N) \rightarrow 0$. It now follows that

$$N_{Y_T}(F_n) \leq N_{X \times \tilde{R}}(F_n) \cdot 2^{\varepsilon(n)n^2}.$$

Therefore,

$$h(Y_T) \leq \frac{1}{n^2} \lim_{n \rightarrow \infty} N_{X \times \tilde{R}}(F_n) + \lim_{n \rightarrow \infty} \varepsilon(n) \leq h(X) + h(\tilde{R}) = 0$$

as claimed. □

7.3. Pruning. Our aim now is to find a Turing machine T so that $(x, r) \in X \times \tilde{R}$ is represented in Y_T if and only if $\delta(x) \leq h$.

Recall that this machine T will receive as its input sequences of 0, 1’s induced by points $x \in X$ on translates of I_n . Write $I = \cup_{n=1}^{\infty} I_n$, and enumerate the elements of I as $I = \{i_1 < i_2 < \dots\}$, where $i_1 = 1$. Note that the first 4^n elements of this sequence are precisely the elements of I_n ; this follows easily from the recursion relation defining the I_n ’s. If $(x, r) \in X \times \tilde{R}$ and $B_n + u$ is an n -board induced by r , then the 0, 1-coloring induced by x on $I_n \times \{1\} + u$ is the sequence $(x_j)_{j=1}^{4^n}$ such that x_j is the symbol 0 or 1 appearing on the i_j -th column in $T^u x$. It follows that for any $k \leq n$, the first 4^k symbols of this sequence correspond to a pattern induced by x on some translate of I_k .

LEMMA 7.2. *There exists a sequence of finite sets $M_n \subseteq I_{m \cdot 2^{n^2}}$ such that $\{i_m : m \in M_n\}$ is a complete set of residue classes modulo 2^n ; i.e., for every $0 \leq j < 2^n$ there exists a unique $m \in M_n$ such that $i_m \equiv j \pmod{2^n}$.*

Proof. By the recursion formula for I_k given, Section 7.1 and the fact that I_k is an increasing sequence, for any $k \leq q$ we have

$$I_k + 5^q \subseteq I_q + 5^q \subseteq I_{q+1}.$$

In particular, since $1 \in I_1$, we may show by induction that for any r and t ,

$$1 + 5^q + 5^{2q} + \dots + 5^{tq} \subseteq I_{q+2q\dots+tq+1}.$$

Given n , since $\gcd(2^n, 5) = 1$ we may choose $q \leq 2^n$ so that $5^q \equiv 1 \pmod{2^n}$. Since the set

$$\{1 + 5^q + \dots + 5^{tq}\}_{t=1}^{2^n}$$

is a complete set of residues modulo 2^n and is contained in $I_{2^{3n}}$, the existence of M_n follows. □

It is clearly possible to compute a sequence of sets $M_n \subseteq I_{2^{3n}}$ with the above properties. The proof above gives an algorithm for doing so, since the identity $5^q = 1 \pmod{2^n}$ is solved by $q = \phi(2^n)$ (here ϕ is Euler's function).

Let $r(n)$ be a computable sequence and $h \leq r(n) \rightarrow h$. We can now describe our algorithm:

ALGORITHM 7.3. *Input:* $(x_n)_{n \in \mathbb{Z}} \in \{0, 1\}^{\mathbb{Z}}$.

For $N = 1, 2, 3 \dots$ *do*

- (1) Calculate $r(N)$.
- (2) Calculate the relative frequency δ_N of 1's in the sequence $(x_m : m \in M_N)$, i.e.

$$\delta_N = \frac{1}{2^N} \#\{m \in M_N : x_m = 1\}.$$

- (3) If $\delta_N > r(N) + 2^{-N}$ then halt.

PROPOSITION 7.4. *Let* $x \in X$ *and let* $(x_n)_{n=1}^\infty$ *be the 0, 1-valued sequence with* x_n *equal to the color of the* i_n -*th column of* x . *Then Algorithm 7.3 halts on the input* (x_n) *if and only if* $\delta(x) > h$, *and if it halts, the number of steps it runs before halting depends only on* $\delta(x)$ *(not on* x *).*

Proof. It suffices to show that $\delta(x) - 2^{-N} \leq \delta_N \leq \delta(x) + 2^{-N}$ for every N . Indeed, if $\delta(x) = h + \varepsilon$ for some $\varepsilon > 0$ then $\delta_N \geq \delta(x) - 2^{-N}$ implies that $\delta_N > h - \varepsilon/2$ for large enough N , and since $r(N) \rightarrow h$ for large enough N we will have $\delta_N > r(N) + 2^{-N}$ and the algorithm will halt. On the other hand if $\delta(x) \leq h$ then $\delta_N \leq \delta(x) + 2^{-N}$ implies that $\delta_N \leq h + 2^{-N} \leq r(N) + 2^{-N}$, so that the algorithm will run forever.

Fix $N \geq 1$ and let $E = \cup_{n=1}^\infty U_n \times V_n$ be the 2-net induced by x . Note that

$$\delta(x) = \sum_{n=1}^\infty \rho_n 2^{-n}$$

where $\rho_n \in \{0, 1\}$ is the symbol induced by x on the grid $U_n \times V_n$.

Note that $|M_n| = 2^n$. Let $J_n = \{j \in M_N : i_j \in U_n\}$. Since $\{i_j : j \in M_N\}$ is a complete set of residues modulo 2^N , for each $n \leq N$ we have

$$|J_n| = 2^{N-n}$$

and since the U_n 's are pairwise disjoint so are the J_n 's, so that

$$|M_N \setminus \bigcup_{n=1}^N J_n| = 1.$$

Let $M_N \setminus \bigcup_{n=1}^N J_n = \{i\}$ and let $\rho' \in \{0, 1\}$ be the symbol induced on the i -th column of x . Then

$$\begin{aligned} \delta_N &= \frac{1}{2^N} \sum_{j \in M_N} 1_{\{x_j=1\}} = \frac{1}{2^N} \left(1_{\{x_i=1\}} + \sum_{n=1}^N \sum_{j \in J_n} 1_{\{x_j=1\}} \right) \\ &= \frac{\rho'}{2^N} + \sum_{n=1}^N \frac{\rho_n |J(n)|}{2^N} = \frac{\rho'}{2^N} + \sum_{n=1}^N \rho_n 2^{-n}. \end{aligned}$$

The desired inequality follows.

Regarding the number of steps the algorithm runs before halting, we see that this depends only on N and $\delta(x)$. □

Let T be a Turing machine implementing this algorithm and whose input is the sequence of 0's and 1's which is the input to the algorithm. We make one important assumption about the implementation, namely that there are integers t_N such that the machine performs the first N iterations of the loop in at most t_N steps (or halts before that), independent of the input. Such an implementation does not present any difficulty. Another thing to note is that as we have defined it, the entire tape is taken up by input data. In order to provide the machine with space to store its intermediate calculations one can allow it to superimpose another layer of symbols over the input alphabet. Formally, this can be done by setting the machine's alphabet to be $\{0, 1\} \times \{0', 1'\}$, with the input represented by the first coordinate and the machines modifying the second coordinate as it pleases.

Let $Y = Y_T$; this is the system Y whose construction was the goal of the second step in the outline given in Section 4.

PROPOSITION 7.5. *If $(x, r) \in X \times \tilde{R}$ then (x, r) is represented in Y_T if and only if $\delta(x) \leq h$.*

Proof. By 7.1 it suffices to show that the condition $\delta(x) \leq h$ is equivalent to the fact that for any finite or infinite board B induced by r representing an $N \times N$ grid ($N \in \mathbb{N} \cup \{\infty\}$), if $r|_B$ contains the symbol \perp then the algorithm does not halt after N steps when run on the input induced by x on the columns of B . The

proposition now follows easily from Proposition 7.4 and the fact that r induces boards of arbitrarily large size. \square

Finally, we note that the topological entropy of $Y = Y_T$ is zero by Proposition 7.1.

8. Step 3: Adding and calculating entropy

Let Y be the system constructed in the previous section. Let Z be the SFT superimposed over Y by adding one of the symbols α, β over each occurrence of the symbol 1. We place no other restrictions on the configurations of α, β 's which may appear. In this section we estimate the entropy of Z and show that it is indeed equal to $h = \sup\{\delta(y) : y \in Y\}$.

Write $F_n = \{1, \dots, n\}^2$ again and for $y \in Y$ denote

$$f_n(y) = \frac{1}{|F_n|} \#\{u \in \mathbb{Z}^2 : y(u) \text{ is marked "1"}\}$$

so that $\delta(y) = \limsup_{n \rightarrow \infty} f_n(y)$. Since $\delta(y) \leq h$ for every $y \in Y$ there is a sequence $\varepsilon_n \rightarrow 0$ such that

$$\sup_{y \in Y} f_n(y) < h + \varepsilon_n$$

(such a sequence ε_n exists by general considerations, but in our case by the proof of Proposition 7.4 one can choose $\varepsilon_n = 2^{-n+1}$).

We now estimate the number of patterns induced by Z on the box $F_n = \{1, \dots, n\}^2$. For each pattern induced on F_n by $y \in Y$, the number of ways to superimpose the symbols α, β and get an admissible pattern for Z is $2^{f_n(y)|F_n|} = 2^{f_n(y)n^2}$. Summing over all patterns induced on F_n by Y and using the fact that $f_n(y) \leq h + \varepsilon_n$ we have

$$N_Z(F_n) \leq N_Y(F_n) \cdot 2^{f_n(y)n^2} \leq N_Y(F_n) \cdot 2^{n^2(h+\varepsilon_n)}$$

so that

$$\limsup_{n \rightarrow \infty} \frac{1}{|F_n|} \log N_Z(F_n) \leq \limsup_{n \rightarrow \infty} \frac{1}{|F_n|} \log N_Y(F_n) + \limsup_{n \rightarrow \infty} (h + \varepsilon_n) = h$$

because $h(Y) = \limsup_{n \rightarrow \infty} \frac{1}{|F_n|} \log N_Y(F_n) = 0$.

On the other hand, if $y_n \in Y$ satisfy $f_n(y_n) \rightarrow h$ then clearly the number of ways to extend $y_n|_{F_n}$ to a pattern in Z is $2^{f_n(y_n)|F_n|}$ and so

$$\liminf_{n \rightarrow \infty} \frac{1}{|F_n|} \log N_Z(F_n) \geq \limsup_{n \rightarrow \infty} f_n(y_n) = h.$$

The entropy estimate $h(Z) = h$ follows.

This completes the proof of Theorem 1.1.

9. Concluding remarks

Many questions remain about the relation between the dynamics, SFTs and their entropies. Let us take a closer look at the system Z constructed above. We can write Z as a disjoint union $Z = \cup_{0 \leq r \leq h} Z_r$ where Z_r is the (nonempty) set of points $z \in Z$ with $\delta(z) = r$; each Z_r is a closed, shift-invariant set, so every orbit closure in Z lies in some Z_r . Hence Z is not transitive. Z also does not have periodic points, since it factors onto the infinite, uniquely ergodic system W .

We remark that if h is computable instead of merely right recursively enumerable, then one can modify Algorithm 7.3 so as to also kill points whose density of 1's is less than h (computability implies both right and left recursive enumerability). For this algorithm the resulting system is essentially the system Z_h above. However, it is still not transitive, since there are many ways to extend an infinite board which does not contain a bottom row; this does not affect entropy, since infinite boards have density zero, but means that Z_h has a transient part.

Problem 9.1. Is every right recursively enumerable number h the entropy of a transitive SFT?

Conversely, we have seen that the entropy of strongly irreducible SFTs is computable. This raises the following:

Problem 9.2. What is the class of entropies of multidimensional strongly irreducible SFTs?

Another mechanism which may be related to entropy is the presence of periodic points. For a two-dimensional SFT X denote by P_n the number of $n \times n$ patterns which can be repeated to produce an admissible tiling of the lattice with period k in both directions. Clearly P_n is computable, and in certain situations one can show that $\lim_{n \rightarrow \infty} \frac{1}{n^2} \log P_n \rightarrow h$ and $\frac{1}{n^2} \log P_n \leq h + \varepsilon_n$ for a sequence ε_n which decays to 0 at a computable rate. This implies that the entropy is computable, because for \tilde{N}_n as in Section 3 we have $h \in (\frac{1}{n} \log P_n - \varepsilon_n, \frac{1}{n} \log \tilde{N}_n)$, and so given n we can examine the difference $\frac{1}{n} \log \tilde{N}_n - (\frac{1}{n} \log P_n - \varepsilon_n)$ for $n = 1, 2, 3 \dots$, stop the first time it is less than $\frac{1}{n}$, and give $\frac{1}{n} \log P_n - \varepsilon_n$ as our estimate.

Friedland [9] used this observation to deduce that if the syntax of an SFT enjoys a certain spatial symmetry then the entropy is computable. We note also that strongly irreducible SFTs in two dimensions have dense periodic points, but whether this is so in higher dimensions seems to be open [26].

Problem 9.3. Do dense periodic points for an SFT imply that the entropy is computable?

Finally, we repeat here an old question which we mentioned in the introduction:

Problem 9.4. Is every sofic shift a factor of an SFT with the same entropy?

References

- [1] R. L. ADLER and B. WEISS, *Similarity of Automorphisms of the Torus*, *Mem. Amer. Math. Soc.* **98**, Amer. Math. Soc., Providence, R.I., 1970. MR 41 #1966 Zbl 0195.06104
- [2] R. BERGER, *The Undecidability of the Domino Problem*, *Mem. Amer. Math. Soc. No. 66*, Amer. Math. Soc., Providence, RI, 1966. MR 36 #49 Zbl 0199.30802
- [3] R. BOWEN, *On Axiom A Diffeomorphisms*, *Reg. Conf. Series in Math.*, No. 35, Amer. Math. Soc., Providence, R.I., 1978. MR 58 #2888 Zbl 0383.58010
- [4] R. BURTON and J. E. STEIF, Non-uniqueness of measures of maximal entropy for subshifts of finite type, *Ergodic Theory Dynam. Systems* **14** (1994), 213–235. MR 95f:28023 Zbl 0807.58023
- [5] E. M. COVEN and M. E. PAUL, Sofic systems, *Israel J. Math.* **20** (1975), 165–177. MR 52 #4260 Zbl 0307.28015
- [6] M. DENKER, C. GRILLENBERGER, and K. SIGMUND, *Ergodic Theory on Compact Spaces*, *Lecture Notes in Math.* **527**, Springer-Verlag, New York, 1976. MR 56 #15879 Zbl 0328.28008
- [7] A. DESAI, Subsystem entropy for \mathbb{Z}^d sofic shifts, *Indag. Math.* **17** (2006), 353–359. MR 2321105 Zbl 1104.37006
- [8] S. FORCHHAMMER and J. JUSTESSEN, Entropy bounds for constrained two-dimensional random fields, *IEEE Trans. Inform. Theory* **45** (1999), 118–127. MR 99k:94023 Zbl 0946.94011
- [9] S. FRIEDLAND, On the entropy of \mathbf{Z}^d subshifts of finite type, *Linear Algebra Appl.* **252** (1997), 199–220. MR 98b:54052 Zbl 0893.54031
- [10] S. FRIEDLAND and U. N. PELED, Theory of computation of multidimensional entropy with an application to the monomer-dimer problem, *Adv. in Applied Math.* **34** (2005), 486–522. MR 2005m:82020 Zbl 1076.82009
- [11] C. GOODMAN-STRAUSS, Matching rules and substitution tilings, *Ann. of Math.* **147** (1998), 181–223. MR 99m:52027 Zbl 0941.52018
- [12] J. E. HOPCROFT and J. D. ULLMAN, *Introduction to Automata Theory, Languages, and Computation*, *Addison-Wesley Series in Computer Science*, Addison-Wesley Publ. Co., Reading, MA, 1979. MR 83j:68002 Zbl 0426.68001
- [13] L. P. HURD, J. KARI, and K. CULIK, The topological entropy of cellular automata is uncomputable, *Ergodic Theory Dynam. Systems* **12** (1992), 255–265. MR 93f:58124 Zbl 0770.58017
- [14] B. P. KITCHENS, *Symbolic Dynamics, One-Sided, Two-Sided and Countable State Markov Shifts*, *Universitext*, Springer-Verlag, New York, 1998. MR 98k:58079 Zbl 0892.58020
- [15] K.-I. KO, *Complexity Theory of Real Functions*, *Prog. Theoret. Comput. Sci.*, Birkhäuser, Boston, 1991. MR 93i:03057 Zbl 0791.03019
- [16] D. A. LIND, The entropies of topological Markov shifts and a related class of algebraic integers, *Ergodic Theory Dynam. Systems* **4** (1984), 283–300. MR 86c:58092 Zbl 0546.58035
- [17] D. LIND and B. MARCUS, *An Introduction to Symbolic Dynamics and Coding*, Cambridge Univ. Press, Cambridge, 1995. MR 97a:58050 Zbl 1106.37301
- [18] D. LIND, K. SCHMIDT, and T. WARD, Mahler measure and entropy for commuting automorphisms of compact groups, *Invent. Math.* **101** (1990), 593–629. MR 92j:22013 Zbl 0774.22002
- [19] M. MISIUREWICZ, A short proof of the variational principle for a \mathbf{Z}_+^N action on a compact space, in *International Conference on Dynamical Systems in Mathematical Physics* (Rennes, 1975), *Astérisque* **40**, Soc. Math. France, Paris, 1976, pp. 147–157. MR 56 #3250 Zbl 0368.54013

- [20] S. MOZES, Tilings, substitution systems and dynamical systems generated by them, *J. Analyse Math.* **53** (1989), 139–186. MR 91h:58038 Zbl 0745.52013
- [21] W. PARRY, Intrinsic Markov chains, *Trans. Amer. Math. Soc.* **112** (1964), 55–66. MR 28 #4579 Zbl 0127.35301
- [22] R. M. ROBINSON, Undecidability and nonperiodicity for tilings of the plane, *Invent. Math.* **12** (1971), 177–209. MR 45 #6626 Zbl 0197.46801
- [23] D. RUELLE, *Thermodynamic Formalism, The Mathematical Structures of Equilibrium Statistical Mechanics*, second ed., *Cambridge Math. Lib.*, Cambridge Univ. Press, Cambridge, 2004. MR 2006a:82008 Zbl 1062.82001
- [24] C. E. SHANNON, A mathematical theory of communication, *Bell System Tech. J.* **27** (1948), 379–423, 623–656. MR 10,133e Zbl 1154.94303
- [25] J. G. SIMONSEN, On the computability of the topological entropy of subshifts, *Discrete Math. Theor. Comput. Sci.* **8** (2006), 83–95. MR 2008b:37025 Zbl 1153.37320
- [26] T. WARD, Automorphisms of \mathbf{Z}^d -subshifts of finite type, *Indag. Math.* **5** (1994), 495–504. MR 97a:28014 Zbl 0823.28007
- [27] B. WEISS, Subshifts of finite type and sofic systems, *Monatsh. Math.* **77** (1973), 462–474. MR 49 #5308 Zbl 0285.28021
- [28] R. F. WILLIAMS, Classification of subshifts of finite type, *Ann. of Math.* **98** (1973), 120–153; errata, *ibid.* **99** (1974), 380–381. MR 48 #9769 Zbl 0282.58008

(Received March 8, 2007)

E-mail address: hochman@math.princeton.edu

DEPARTMENT OF MATHEMATICS, PRINCETON UNIVERSITY, FINE HALL - WASHINGTON RD,
PRINCETON, NJ, UNITED STATES

and

SCHOOL OF MATHEMATICS, INSTITUTE FOR ADVANCED STUDY, EINSTEIN DR,
PRINCETON, NJ 08540, UNITED STATES
<http://www.math.princeton.edu/~hochman/>

E-mail address: tomm@post.tau.ac.il

SCHOOL OF MATHEMATICAL SCIENCES, TEL AVIV UNIVERSITY, RAMAT AVIV,
TEL AVIV 69978, ISRAEL

<http://www.math.tau.ac.il/~tomm/>