ON THE AUTOMORPHISM GROUPS OF MULTIDIMENSIONAL SHIFTS OF FINITE TYPE

MICHAEL HOCHMAN

Abstract. We investigate algebraic properties of the automorphism group of multidimensional shifts of finite type (SFTs). We show that positive entropy implies that the automorphism group contains every finite group, and together with transitivity implies that the center of the automorphism group is trivial (i.e. consists only of the shift action). We also show that positive entropy and dense minimal points (in particular, dense periodic points) imply that the automorphism group of X contains a copy of the automorphism group of the one-dimensional full shift, and hence contains nontrivial elements of infinite order. On the other hand we construct a mixing, positive entropy SFT which, modulo the shift action, is the union of finite groups.

1. Introduction

Let Σ be a finite alphabet, T the shift map on the full shift $\Sigma^{\mathbb{Z}}$, and $X \subseteq \Sigma^{\mathbb{Z}}$ a mixing shift of finite type. An automorphism of X is a shift-commuting homeomorphism of X, and the group of automorphisms is denoted $\operatorname{aut} X$. The dynamics of individual automorphisms have been much studied as discrete models for evolution of complex physical systems. The automorphism group as a whole is of interest as well, and has been studied both from the dynamical perspective and algebraically. It is the last point of view that we adopt here.

The classical work on this topic is due to Hedlund [3], who showed for the full shift $X = \Sigma^{\mathbb{Z}}$ that $\operatorname{aut} X$ is countable but very large: it contains isomorphic copies of every finite group, free products of finite groups, and free abelian groups of every rank. Another measure of its largeness, due to Ryan [10], is that the center of $\operatorname{aut} X$ consists precisely of the shift action. Similar results for the automorphism group of general mixing one-dimensional SFTs were proved by Boyle, Lind and Rudolph [2]. In the general case one can also rely on Hedlund's results in conjunction with a later theorem of Kim and Roush, who showed that for every n and every mixing one-dimensional SFT X, the automorphism group of the full shift on n symbols is algebraically embedded in $\operatorname{aut} X$ [6].

In the present work we consider the automorphism groups of multidimensional SFTs, and in particular are interested in when these groups are "large". This is not a straightforward generalization of the one-dimensional theory; multidimensional SFTs can differ rather drastically from their one-dimensional relatives, and one cannot apply the same combinatorial constructions as one can there. Under the hypothesis of strong irreducibility, T. Ward [12] has shown that the automorphism group contains every finite group. He also gave an example of a multidimensional mixing zero entropy SFT whose automorphism group is trivial (consisting only of powers of the shifts).

These results point to the fact that mixing is perhaps not the right condition to study in connection with the automorphism group, and in this work we show that positive entropy plays an important role. Our first result is similar to Ward's, but under the assumption of positive entropy (a much weaker condition than strong specification).

Theorem 1. *Let X be a positive-entropy $\mathbb{Z}^d$-SFT. Then* $\operatorname{aut} X$ *contains copies of every finite group.*

It should be noted that positive entropy is not necessary for the automorphism group of a two-dimensional SFT to be large. Consider the two-dimensional SFT $X \subseteq \{0,1\}^{\mathbb{Z}^2}$ in which each vertical column of symbols is constant. This is essentially the one dimensional full shift $\{0,1\}^{\mathbb{Z}}$ in which we formally extend the $\mathbb{Z}$-action generated by the shift to a $\mathbb{Z}^2$-action generated by the shift and the identity. Clearly, the automorphism groups are the same.

Under the additional hypothesis of transitivity, we get the following analogue of Ryan's theorem:

Theorem 2. *If $X \subseteq \Sigma^{\mathbb{Z}^d}$ is a topologically transitive SFT with positive entropy, then the shift action is the center of* $\operatorname{aut} X$.

Finally, we obtain an analogue of Kim and Roush's result, for which we need a slightly stronger dynamical hypothesis:

Theorem 3. *Let $X \subseteq \Sigma^{\mathbb{Z}^d}$ be a positive entropy SFT. If the minimal points are dense in X (in particular, if the periodic points are dense), then* $\operatorname{aut} X$ *contains isomorphic copies of the automorphism group of the one-dimensional full shift on any number of symbols.*

Theorem 3 cannot be significantly improved, as shown by the following two examples. Recall that a countable group G is locally finite if it is the increasing union of finite subgroups. A group is virtually simple if it has a finite index simple subgroup.

Example 4. There exists a positive-entropy $\mathbb{Z}^2$ SFT X such that $\operatorname{aut} X = \mathbb{Z}^2 \oplus G$, where $\mathbb{Z}^2$ is the shift action and G is locally finite and virtually simple.

Example 5. There exists a strongly-mixing positive-entropy $\mathbb{Z}^2$ SFT X such that $\operatorname{aut} X = \mathbb{Z}^2 \oplus G$, where $\mathbb{Z}^2$ is the shift action and G is a locally finite group which has as a homomorphic image an infinite virtually simple group.

Recall that if X is a mixing one-dimensional SFT then $\operatorname{aut} X$ acts on the (finite) set of periodic points of period n, and since the periodic points are dense these actions separate points in $\operatorname{aut} X$, implying that $\operatorname{aut} X$ is residually finite. In the examples above $\operatorname{aut} X$ is not residually finite. This can be seen as strengthening the fact that periodic points need not be dense in a multidimensional SFT.

Examples 4 and 5 show that the analogue of theorem 3 is false in general for multidimensional SFTs, since the automorphism group of the full two-dimensional shift contains a copy of $\mathbb{Z}^3$ but the group in our example does not. It also shows that in general there is no multidimensional analogue of Kim and Roush's result, since one cannot embed the automorphism group from our example in the automorphism group of the full two-dimensional shift, which is residually finite (this is again due to its actions on sets of periodic points, which are dense in the full shift).

As it happens, these constructions also provide an example of interesting behavior of a different sort. Let X be a $\mathbb{Z}^2$-subshift. If $x, x' \in X$ and $H^+, H^- \subseteq \mathbb{R}^2$ are closed half-spaces with common boundary ("complementary half-spaces"), we say that x, x' are (H^+, H^-, r)-compatible if $x(u) = x'(u)$ for all $u \in \mathbb{Z}^d$ within distance r of the boundary $\partial H^+ = \partial H^-$. We say that X is *half-space Markov* (with step r) if, for every such pair of complementary half-spaces H^+, H^- and every pair of (H^+, H^-, r)-compatible points $x, x' \in X$, the point

$$y(u) = \begin{cases} x(u) & u \in H^+ \\ x'(u) & u \in H^- \end{cases}$$

is in X. Clearly every r-step SFT is half-space Markov. M. Boyle raised the question of whether the converse holds. Using example 4 we show:

Example 6. There is a half-space Markov subshift which is not an SFT.

After some preliminaries in section 2, the subsequent sections (3, 4 and 5) are devoted to proving theorems 1, 2 and 3 respectively. In section 6 we construct the examples 4 and 5. In section 7 we discuss half-space Markov systems, and in section 8 we list some open problems.

2. Preliminaries

See [8, 7] for general background on shifts of finite type and symbolic dynamics in one dimension.

We fix the dimension d and an alphabet Σ. For a point $x \in \{1,\ldots,k\}^{\mathbb{Z}^d}$ and $u \in \mathbb{Z}^d$ let $T^u x$ denote the translate of x by u, i.e. $(T^u x)(v) = x(u+v)$. A symbolic system, or subshift, is a non-empty closed subset of $\{1,\ldots,k\}^{\mathbb{Z}^d}$ which is invariant under the $\mathbb{Z}^d$ action $\{T^u\}_{u\in\mathbb{Z}^d}$.

For $F \subseteq \mathbb{Z}^d$ we refer to a Σ-coloring $a \in \Sigma^F$ of F as a *pattern* or F-pattern. If $X \subseteq \Sigma^{\mathbb{Z}^d}$ and $x \in X$ we write $x|_F$ for the pattern obtained by restricting x to F. We will say a pattern $a \in \Sigma^F$ appears in $x \in \Sigma^{\mathbb{Z}^d}$ at position $u \in \mathbb{Z}^d$ if $a(v) = x(u+v)$ for $v \in F$, i.e. if $(T^u x)|_F = a$. We say a appears in a point $x \in \Sigma^{\mathbb{Z}^d}$ if it appears at some position in it, and if $X \subseteq \Sigma^{\mathbb{Z}^d}$ is a subshift we say a appears in X is it appears in some $x \in X$. Patterns $a \in \Sigma^E$ and $b \in \Sigma^F$ are *congruent* if F is a translate of E, say $F = E + u$, and for all $v \in E$ we have $a(v) = b(u+v)$. Note that if $u \neq 0$ then formally $a \neq b$ because they have different domains, but one should think of them as translates of the same pattern.

A shift of finite type (SFT) is a subshift defined by a finite set of allowed patterns, i.e. X is an SFT if there is a finite set $F \subseteq \mathbb{Z}^d$ and a family of words $L \subseteq \Sigma^F$ such that $x \in X$ if and only if $(T^u x)|_{F+u}$ is congruent to a pattern in L for every $u \in \mathbb{Z}^d$. Equivalently one can define an SFT by a finite set of disallowed patterns.

We write $[i;j] = [i,j] \cap \mathbb{Z}$ for the segment of integers between i, j. For $r \in \mathbb{N}$ and $u \in \mathbb{Z}^d$ let

$$B_r(u) = \{v \in \mathbb{Z}^d \,:\, \|u - v\|_\infty \leq r\}$$

So $B_0(u) = \{u\}$. We write $B_r = B_r(0)$. Note that B_r is a box of side $2r+1$, and that $|B_r(u)| = |B_r| = (2r+1)^d$.

If X is an SFT defined by a set of words $L \subseteq \Sigma^{B_n}$ we say that X is an n-step SFT. It is classical that every SFT is isomorphic to an SFT which is 1-step, i.e. one in which the only restrictions are between a symbol and its immediate neighbors. In what follows we will assume that our SFTs are 1-step (except in section 6 where it will be convenient for the construction to allow longer-range dependencies).

For a subshift $X \subseteq \Sigma^{\mathbb{Z}}$ let $N(k)$ denote the number of B_k-patterns appearing in X. The topological entropy of X is the limit

$$h_{\mathrm{top}}(X) = \lim_{k\to\infty} \frac{1}{(2k+1)^d} \log N(k) = \inf_k \frac{1}{(2k+1)^d} \log N(k)$$

For finite $F \subseteq \mathbb{Z}^d$ we define the k-boundary of F by

$$\partial_k F = \{u \in F \,:\, \exists v \in (\mathbb{Z}^d \setminus F)\ \|u - v\|_\infty \leq k\}$$

In particular, $B_r \setminus \partial_k B_r = B_{r-k}$. For brevity we write $\partial = \partial_1$.

Let $X \subseteq \Sigma^{\mathbb{Z}^d}$ be a subshift. A point x is transitive if its orbit $\{T^u x \,:\, u \in Z^d\}$ is dense in X, and X is transitive if it contains a transitive point, or equivalently if for every pair of nonempty open sets $U, V \subseteq X$, there is a $u \in \mathbb{Z}^d$ with $T^u U \cap V \neq \emptyset$. Symbolically this means that if $E, F \subseteq \mathbb{Z}^d$ are finite and $a \in \Sigma^E$, $b \in \Sigma^F$ appear in X, then there is an $x \in X$ in which both a and b appear.

A point $x \in X$ is minimal if for every finite $F \subseteq \mathbb{Z}^d$, the set $\{u \in Z^d \,:\, x|_{F+u} = x|_F\}$ is syndetic, i.e. has bounded gaps. This is equivalent to the property that for every y in the orbit closure of x, also x is in the orbit closure of y; or to the property that the orbit closure of x contains no nontrivial subsystems.

A subshift X is mixing if for all nonempty open sets $U, V \subseteq X$ it holds that $T^u U \cap V \neq \emptyset$ for all but finitely many $u \in \mathbb{Z}^d$. Equivalently, for every two finite patterns $a \in \Sigma^E$ and $b \in \Sigma^F$ which appear in X, and for all but finitely many $u \in \mathbb{Z}^d$, there exists $x \in X$ such that $x|_E = a$ and $x|_{E+u} = b$.

For a homeomorphism $S : X \to X$ of a compact metric space, a point $x \in X$ is recurrent for the $\mathbb{Z}$-action generated by S if $S^{n(k)} x \to x$ for some sequence $|n(k)| \to \infty$.

We will use the following classical fact: if $X, Y \subseteq \Sigma^{\mathbb{Z}^d}$ are subshifts and $\varphi : X \to Y$ is a continuous shift-commuting map (and in particular if $X = Y$ and φ is an automorphism), then there is an n and a map $\varphi_0 : \Sigma^{B_n} \to \Sigma$ such that $\varphi(x)(0) = \varphi_0(x|_{B_n})$ for every $x \in X$; in other words, $\varphi(x)(0)$ depends only on $x|_{B_n}$. The minimal n for which this holds will be called the *window width* of φ, and denoted $R(\varphi)$.

3. Markers and the automorphism group

3.1. **Embedding Finite Groups in** $\mathrm{aut}(X)$**.** The key to constructing automorphisms of SFT is the so-called "marker method". We must first define multidimensional markers. Note that the following definition differs from that in [12].

Definition 7. Let $m > k \geq m/2$ and $F = \partial_k B_m$ be the "square annulus" of thickness k, surrounding a cube of side $2m + 1 - 2k \leq m + 1$. A marker (for F) is a word $a \in \Sigma^F$ such that if $x \in \Sigma^{\mathbb{Z}^d}$ and a appears in x at u' and u'' then either $u' = u''$ or $\|u' - u''\| > m + k$.

Let $X \subseteq \Sigma^{\mathbb{Z}^d}$ be a subshift, $F = \partial_k B_m$ and $a \in \Sigma^F$ an admissible word. A word $b \in B_{m-k}$ is an *admissible completion* of a if $a \cup b$ appears in X.

Let a be a marker for $\partial_k B_m$ and b', b'' admissible completions of a. If $x \in \Sigma^{\mathbb{Z}^d}$ and $a \cup b', a \cup b''$ appear at distinct locations in x, then the marker condition implies that the corresponding occurrences of b', b'' do not intersect each other or the surrounding a's.

For example, consider the following 2-dimensional words $a, b \in \{0,1\}^{\partial_1 B_1}$:

$$a = \begin{matrix} 1 & 1 & 1 \\ 0 & & 1 \\ 0 & 0 & 0 \end{matrix} \quad , \quad b = \begin{matrix} 1 & 1 & 1 \\ 0 & & 0 \\ 0 & 0 & 0 \end{matrix}$$

One easily checks that a is a marker, but b isn't because in the pattern

$$\begin{matrix} 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{matrix}$$

b occurs here twice, at displacement 1.

Markers can be used in a standard way to construct automorphisms:

Proposition 8. *Let $a \in \Sigma^{\partial_k B_m}$ be an admissible marker for an SFT X, and let $b_0, b_1 \in \Sigma^{B_{m-k}}$ be distinct completions of a. Then the map $\varphi : X \to X$ which exchanges every occurrence of $a \cup b_i$ with $a \cup b_{1-i}$ is well defined and is an automorphism of X. If a has n distinct completions $b_0, \ldots, b_{n-1}$ then the group S_n of permutations of $\{1, \ldots, n\}$ is embedded in $\mathrm{aut}(X)$ via $\pi \mapsto \varphi_\pi$, where $\varphi_\pi(x)$ replaces each occurrence of $a \cup b_i$ in x with $a \cup b_{\pi(i)}$.*

Proof. Suppose $a \cup b_i$ appears in x at u and $a \cup b_j$ appears in x at v. Then by the marker property either $u = v$ (and hence $i = j$) or else the patterns b_i, b_j at u, v respectively do not intersect each other nor do they intersect the a's surrounding them, or any other occurrence of a. Therefore φ is well defined and $\varphi(x)$ contains a's in the same places that x did. φ is clearly continuous and shift-commuting, and is its own inverse, so $\varphi \in \mathrm{aut}\, X$.

The second statement follows similarly: φ_π is an automorphism and clearly $\varphi_\pi \circ \varphi_\tau = \varphi_{\pi\tau}$, and $\pi \mapsto \varphi$ is an injection because each $a \cup b_i$ appears in X. □

3.2. **The Existence of Markers in Positive Entropy SFTs.** We are interested in showing that there exist markers with many completions in as wide a class of SFTs as possible. To our knowledge, in the literature markers have always been constructed explicitly, both in the one-dimensional case and the multidimensional case. These constructions rely on specification-type properties of the subshifts which allow one to glue together patterns in convenient ways, and while these conditions are usually satisfied automatically in one-dimension, in the multidimensional case they must be assumed explicitly, as in [12]. It turns out that to guarantee the existence of markers what is needed is not a mixing condition but rather positive entropy.

We will show this in two stages. First, we show that there are words of the form $a \in \Sigma^{\partial B_m}$ which are not markers but which have multiple completions. Then we show that these words can be "extended" to be completions of markers.

Lemma 9. *Let $X \subseteq \Sigma^{\mathbb{Z}^d}$ have positive entropy. Then for any integer $M > 0$, there is a $k \in \mathbb{N}$ and a word $a \in \Sigma^{\partial B_k}$ which has M compatible completions.*

Proof. By the definition of topological entropy, the number of B_k-words occurring in X is at least $2^{h|B_k|}$. On the other hand the number of ∂B_k-words occurring in X is at most $|\Sigma|^{|\partial B_k|}$. Since $|\partial B_k|/|B_k| \to 0$, the result follows from the pigeonhole principle. □

Let X be an SFT. Let $a \in \Sigma^{\partial B_{k+1}}$ and for some $n > k$ suppose that $a' \in \Sigma^{\partial_{n-k} B_n}$ occurs in X and $a'|_{\partial B_k} = a$. Then every completion of a is a completion of a' (here we use the assumption that X is one-step). In particular if a'' is a marker and has a completion in which a appears, then the number of completions of a'' is at least that of a. It follows that for a positive entropy SFT, if we want to have a marker with M completions it suffices to find a k and a ∂B_k-word with M completions (these exist by the previous lemma), and then show that this word is a subword of the completion of some marker.

For a subshift $X \subseteq \Sigma^{\mathbb{Z}^d}$, a finite set $F \subseteq \mathbb{Z}^d$ and a pattern $a \in \Sigma^F$ appearing in X, let $N(k|a)$ denote the number of admissible elements of Σ^{B_k} which extend a. More precisely, for k such that $F \subseteq B_k$ we define

$$N(k|a) = \#\{x|_{B_k} \,:\, x \in X\,,\, x|_F = a\}$$

The following is a reformulation of part of the Shannon-McMillan theorem:

Lemma 10. *Let $X \subseteq \Sigma^{\mathbb{Z}^d}$ be a dynamical system with a globally supported ergodic measure of entropy $h > 0$. Let $a \in \Sigma^F$ be a word appearing in X. Then for every $\varepsilon > 0$ and large enough k,*

$$N(k|a) \geq 2^{(h-\varepsilon)\cdot|B_k|}$$

The following lemma is purely geometrical.

Lemma 11. *Let $k \in \mathbb{N}$ and $F = \partial_{9k} B_{10k}$. If $u', u'' \in \mathbb{Z}^d$ and $0 < \|u' - u''\|_\infty \leq 19k$ then there is a $v \in \mathbb{Z}^d$ such that*

$$B_k + v \subseteq (F + u') \cap (F + u'')$$

Proof. This is clearly true for $d = 1$ and follows from this case by projecting on each coordinate. □

Let $b \in \Sigma^{B_r}$ be an admissible word for a system $X \subseteq \Sigma^{Z^d}$. For $k \geq r$ let $M(k|b)$ be the number of markers $a \in \Sigma^{\partial_{9k} B_{10k}}$ such that $a \cup b$ appears in X. We next deduce the existence of markers in positive-entropy symbolic systems.

Proposition 12. *Let $X \subseteq \Sigma^{\mathbb{Z}^d}$ be a subshift with a globally supported measure of entropy $h > 0$ and let $b \in \Sigma^{B_r}$ be admissible for X. Then for every $\varepsilon > 0$ and for all large enough k we have*

$$M(k|b) \geq 2^{(h-\varepsilon)|\partial_{9k} B_{10k}|}$$

Proof. We prove the lemma for $d = 2$, the general case being completely analogous.

Let $N(i,j)$ be the number of $[1;i] \times [1;j]$-patterns appearing in X, so $N(k) = N(k,k)$. Fix n so that $h \leq \frac{1}{n} \log N(n) \leq h + \frac{1}{2}\varepsilon$. For i,j large enough, every rectangle $F = [1;i] \times [1;j]$ can be tiled by squares of side n except a region of size at most $n \cdot (i+j)$, and $\frac{n(i+j)}{ij} \to 0$ as $i,j \to \infty$. It follows that when i,j are large enough, $h \leq \frac{1}{i \cdot j} \log N(i,j) < h + \varepsilon$ (see also [1]).

Fix a large k and write $F = \partial_{9k} B_{10k}$. Suppose that $a \in \Sigma^F$ is such that $a \cup b$ appears in X and a is not a marker. Then this is because there exists a $u = u(a) \in \mathbb{Z}^2$ with $\|u\| \leq 19k$ and a $x \in \Sigma^{\mathbb{Z}^d}$ with a appearing in x at positions 0 and u. It follows from lemma 11 that $F \cap (F + u)$ contains a translate $E = B_k + v$ of B_k by some $v \in \mathbb{Z}^d$. This means that if $w \in E$ then $w - v \in F$ and $a(w) = a(w - v)$. Note that E, v, w may be chosen to depend only on u.

We proceed to estimate, for each $u \in B_{19k}$, the number of such patterns a (admissible for X) for which the above holds with this u. The key observation is that a is determined completely by the pattern $a|_{F \setminus E}$, since given $a|_{F \setminus E}$ we can reconstruct $a|_E$ as follows: for $w \in E$ if $w - v \in F \setminus E$ then $a(w - v)$ is known and $a(w) = a(w - v)$; in case $w - v \in E$ we iterate this procedure and look at $w - 2v$, and so on. Eventually we will land in $F \setminus E$, and the symbol there is equal to $a(w)$.

Thus the number of patterns a corresponding to a given u is bounded by the number of $F \setminus E$-patterns in X. Assuming k is large enough we can decompose $F \setminus E$ into large rectangles $R_1, \ldots, R_m$ such that for each i, the number of R_i-patterns appearing in

X is $< 2^{(h+\varepsilon)|R_i|}$. Thus the number of possibilities for $a|_{F\setminus E}$ is bounded above by $2^{(h+\varepsilon)|F\setminus E|}$, and the total number of patterns $a \in \Sigma^F$ with $a \cup b$ admissible which are *not* markers is at most

$$\#\{\text{choices of } u\} \cdot 2^{(h+\varepsilon)|F\setminus E|} \quad \leq \quad (2 \cdot 19k+1)^2 2^{(h+\varepsilon)|F\setminus E|}$$

So the number of $a \in B_{10k}$ extending b such that $a|_{\partial_{9k}B_{10k}}$ is not a marker is at most

$$\begin{aligned} N(k|b) \cdot (38k+1)^2 2^{(h+\varepsilon)|F\setminus E|} &\leq 2^{(h+\varepsilon)|B_k|}(38k+1)^2 2^{(h+\varepsilon)|F\setminus E|} \\ &= (38k+1)^2 2^{(h+\varepsilon)(|B_{10k}|-|B_k|)} \end{aligned}$$

because $N(k|b) \leq N(k) \leq 2^{(h+\varepsilon)|B_k|}$ and $|F \setminus E| = |B_{10k}| - |B_k|$.

But by lemma 10 we know that if k is large enough then $N(10k|b) \geq 2^{(h-\varepsilon/2)|B_{10k}|}$, so

$$\begin{aligned} &\#\{a \in \Sigma^{B_{10k}} \,:\, a|_{B_r} = b\} - \#\{a \in \Sigma^{B_{10k}} \,:\, a|_F \text{ is not a marker}\} \\ &\quad \geq \quad 2^{(h-\varepsilon/2)|B_{10k}|} - (38k+1)^2 2^{(h+\varepsilon)(|B_{10k}|-|B_k|)} \end{aligned}$$

But $|B_{10k}| - |B_k| = (20k+1)^d - (2k+1)^d \leq (1-\delta)|B_{10k}|^d$ for some $\delta > 0$ independent of k, and we may assume that ε is small enough so that $h - \varepsilon/2 > (h+\varepsilon)(1-\delta)$. The result follows. □

Corollary 13. *Let $X \subseteq \Sigma^{\mathbb{Z}^d}$ be an SFT with positive entropy, and let $M > 1$. Then there is a marker $a \in \Sigma^{\partial_{9k}B_{10k}}$ occurring in X with at least M completions in X.*

Proof. Let μ be an invariant measure on X with positive entropy. Let $X_0 = \operatorname{supp}\mu \subseteq X$; then X_0 also has positive entropy (though may no longer be an SFT). Applying lemma 9 we can find a word $b \in \Sigma^{\partial B_k}$ admissible for X_0 with M completions in X_0 (and hence b is admissible and has M completions in X as well). Applying the proposition, for a large enough k we can find a marker $a \in \Sigma^{\partial_{9k}B_{10k}}$ in X_0 with $b \cup a$ appearing in X_0. Thus a is a marker with M completions in X. □

This allows us to prove theorem 1:

Proof. (of theorem 1) We have seen that for every n the group $\operatorname{aut} X$ contains the group of permutations on n symbols, and every finite group is embeddable in one of these. □

4. The center of a positive-entropy SFT is trivial

In this section we generalize to the multidimensional setting a theorem of Ryan [10] stating that if X is a one-dimensional mixing SFT then the center of $\operatorname{aut} X$ consists only of the shift action.

Throughout this section we fix a positive entropy SFT $X \subseteq \Sigma^{\mathbb{Z}^d}$ and an automorphism φ in the center of $\operatorname{aut} X$, i.e. $\varphi\psi = \psi\varphi$ for all $\psi \in \operatorname{aut} X$. Let $R = \max\{R(\varphi), R(\varphi^{-1})\}$ be the larger of the window widths of φ,φ^{-1}. We can find a $k > R$ and a pattern $a \in \Sigma^{\partial_{9k} B_{10k}}$ which is an admissible marker for X with exactly $N \geq 3$ admissible completions, which we denote by $b_1, \ldots, b_N \in \Sigma^{B_k}$.

Lemma 14. *Suppose $x \in X$ and a appears in x at u. Then there is a unique $v \in \mathbb{Z}^d$ with $\|v\|_\infty \leq R + 2k$ such that a appears in $\varphi(x)$ at $u + v$.*

Proof. Suppose there were no such v. Suppose that b_1 is the completion of a in x. Let ψ be the marker automorphism which exchanges $a \cup b_1$ and $a \cup b_2$ and has no other effect. Thus $x|_{B_k(u)} = b_1$ and $\psi(x)|_{B_k(u)} = b_2$. Since φ is an automorphism, $\varphi(x)|_{B_{k+R}(u)} \neq \varphi(\psi(x))|_{B_{k+R}(u)}$. On the other hand, by assumption no part of the square $B_{k+R}(u)$ intersects the interior of a marker a in $\varphi(x)$, so $\psi(\varphi(x))|_{B_{k+R}(u)} = \varphi(x)|_{B_{k+R}(u)}$ and hence $\psi(\varphi(x))|_{B_{k+R}(u)} \neq \varphi(\psi(x))|_{B_{k+R}(u)}$, and $\varphi\psi \neq \psi\varphi$, contrary to the hypothesis. Uniqueness of v follows from the marker property. □

Lemma 15. *Under the conditions of the previous lemma, the displacement v does not depend on the completion $x|_{B_k(u)}$ of the marker a at u.*

Proof. Let x_i be the word obtained from x by setting $x|_{B_k(u)} = b_i$. Applying the previous lemma to x_i we get a vector $v_i \in \mathbb{Z}^d$ with $\|v_i\|_\infty \leq R+2k$ such that a appears in $\varphi(x_i)$ at $u + v_i$. Suppose $v_i \neq v_j$ for some i, j. Then if ψ is the automorphism that exchanges $a \cup b_i$ and $a \cup b_j$, we see that $\psi(\varphi(x_i))$ has an occurrence of a at $u + v_i$ while $\varphi(\psi(x_i))$ has an occurrence of a at $u + v_j$, so $\varphi\psi \neq \psi\varphi$, which is impossible (here we used the marker property of a and the fact that $R < k$ to deduce that these two possibilities are mutually exclusive). □

Lemma 16. *Under the same conditions as the previous lemma, the completion of a at u in x is the same as the completion of the displaced a in $\varphi(x)$.*

Proof. Let $k(i)$ be the index such that if a at u in x is completed by b_i then the a at $u + v$ in $\varphi(x)$ is completed by $b_{k(i)}$. Then k is $1 - 1$ because φ is an automorphism and $k > R$; hence $k(\cdot)$ is a permutation of $\{1, \ldots, N\}$. If $k \neq \mathrm{id}$ let π be another permutation which does not commute with k, and ψ the automorphism of X which acts on the completions of a like π. Then ψ does not commute with φ. $\square$

Lemma 17. *Suppose that X is transitive. Under the circumstances described in the previous lemma, the displacement v does not depend on $x|_{\mathbb{Z}^d \setminus B_{10k}(u)}$.*

Proof. Suppose there are points in X with occurrences of a whose displacement by φ is $v' \neq v''$ with $\|v'\|, \|v''\| \leq R + 2k$. By transitivity we can find a point $x \in X$ with occurrences of a at two places u', u'' such that a appears in $\varphi(x)$ at $u'+v', u''+v''$. Using lemma 15 we may assume that both the occurrences of a in x (at u', u'') are completed by the same completion b_1 and that every other marker within a $100(\|u' - u''\| + k)$ radius of u' and u'' is completed with some word b_N (note that changing the completions of markers other than the ones at u', u'' does not affect the displacement of those markers because the displacement is determined by $x|_{B_{12k+R}(u')}$ and $x|_{B_{12k+R}(u'')}$ and by the marker property and the assumption that $k > R$, all changes to other markers do not affect these patterns).

Define an automorphism ψ of X as follows. For $y \in X$, when ψ sees a pair of occurrences of a at some position u and at $u + (u'' - u')$ in y, it checks three things: (a) that both are completed by b_1 or both are completed by b_2, (b) that every other occurrence of a within a $100(\|u' - u''\|_\infty + k)$ radius of u or $u + (u'' - u')$ is completed with a b_N, and (c) that applying φ to y will translate the occurrences of a at u and $u + (u'' - u')$ by a different amount. If the answer to all three is yes then ψ applies the permutation $b_1 \leftrightarrow b_2$ to the completions of these a's. Clearly ψ is well defined and is an automorphism because it is its own inverse.

Finally, consider the points $\varphi(\psi(y))$ and $\psi(\varphi(y))$. In the first of these, first the b_1's are changed to b_2 and then the markers are moved by different amounts. In the second case, they markers are first moved, so when we apply ψ their offset from each other are no longer $u_1 - u_2$; so applying ψ has no effect on their completions, which remain b_1. Thus $\varphi(\psi(y)) \neq \psi(\varphi(y))$, which is impossible. $\square$

In summary, we conclude that φ displaces every occurrence of a by v.

The proof of theorem 2 is another application of the same technique:

Proof. (of theorem 2) Let $\varphi\psi = \psi\varphi$ for all $\psi \in \operatorname{aut} X$, let a be a marker and v as above, so φ acts on every occurrence of a as the translation T^v.

Suppose that $\varphi \neq T^v$. We can then find a point $y \in X$ such that $\varphi(y) \neq T^v(y)$, so for a large enough n we have $\varphi(y)|_{B_{n-R}} \neq T^v(y)|_{B_{n-R}}$. Consider all points $y' \in X$ such that $y'|_{B_n} = y|_{B_n}$, and choose such a y' which has the maximal number of a's intersecting

∂B_n. Let $E \subseteq \mathbb{Z}^d$ be the union of B_n and the domains of all these a's and their interiors. Let c be the word obtained from $y|_E$ by replacing every completion of a in $y|_E$ with a fixed admissible completion b_0. Since φ acts on a's and their completions like T^v, it follows that whenever $x \in X$ with $x|_E = c$, then $\varphi(x)|_E \neq T^v(x)|_E$. Furthermore, suppose that $x|_E = c$ and $\psi \in \operatorname{aut} X$ acts by permuting the completions of a but fixes $a \cup b_0$. Then ψ preserves c, for otherwise there would be an occurrence of $a \cup b$ for some $b \neq b_0$ with b overlapping c; but it cannot overlap any of the a in c by the marker property, so a must intersect ∂B_n, contradicting the definition of y' and E.

Let $b_1 \neq b_2$ be an admissible completions of a distinct from b_0, and use the transitivity of X to find a point $x \in X$ containing both c and $a \cup b_1$ at some displacement u. Define $\psi \in \operatorname{aut} X$ as the involution which exchanges $a \cup b_1$ and $a \cup b_1$ if and only if it sees c at displacement u from the a. Then φ translates $a \cup b_1$ by v but does not do this for c, so $\varphi\psi$ and $\psi\varphi$ act differently on x. Thus ψ and φ do not commute. □

5. SFTs with dense minimal points

In this section we show that if X is a positive entropy $\mathbb{Z}^d$-SFT with dense minimal points (e.g. with dense periodic points) then the automorphism group of the one-dimensional full shift on any number of symbols is embedded in $\operatorname{aut} X$ (this generalizes a result of Kim and Roush for dimension 1 [6]). This implies that the automorphism group is "at least as large" as in the one-dimensional case. In the following sections we will show that positive entropy alone is not enough to ensure this.

5.1. **The One-Dimensional case.** For simplicity we first explain how to do this in the one-dimensional case, following [7]. Fix M. We want to show that

$$\operatorname{aut}\{1, \ldots, M\}^{\mathbb{Z}} \hookrightarrow \operatorname{aut} X$$

for $X \subseteq \Sigma^{\mathbb{Z}}$ a positive entropy SFT. We may assume X is one-step. Pick a marker a for X with M^2 different admissible completions which we label $\alpha_{1,1}, \alpha_{1,2}, \ldots, \alpha_{M,M}$.

Fix a large R. Let $x \in X$. We define a *coded stretch* in x to be a subword $y = x_{[s;t]}$ of x which is maximal with respect to the property that occurrences of words of the form $b \cup \alpha_{i,j}$ appear in y in an arithmetic progression of gap R, and y begins and ends with occurrences of these patterns. A coded stretch may be finite, one-sided- or two-sided-infinite. Every $x \in X$ can be written uniquely as

$$\ldots z_{-2} y_{-1} z_{-1} y_0 z_0 y_1 z_1 \ldots$$

where y_i are coded stretches and z_i are (possibly empty) words in which no $a \cup \alpha_{ij}$ occurs. From the mixing properties of one-dimensional SFTs it follows that if R is chosen large enough then there will exist arbitrarily long (and infinitely long) coded stretches.

In order to construct a group embedding $\text{aut}\{1,\dots,M\}^Z \hookrightarrow \text{aut}\, X$, the idea is to map $\varphi \in \text{aut}\{1,\dots,M\}^{\mathbb{Z}}$ to an automorphism $\overline{\varphi} \in \text{aut}(X)$ which acts on coded stretches in X as though they were sequences of symbols in $\{1,\dots,M\}^{\mathbb{Z}}$, the symbols being represented by the sequence of α_{ij}'s appearing in the stretch, and acting on this data as φ would. As a first attempt one would try to have each α_{ij} code a single letter from the alphabet $\{1,\dots,M\}$. However then one runs into trouble because it is not clear how to define $\overline{\varphi}$ at the beginning or end of a finite coded stretch.

To overcome this problem we will use a more elaborate coding scheme. We identify the symbol $\alpha_{i,j}$ with the pair of indices (i,j), $1 \le i,j \le M$, which we write as $\begin{bmatrix} i \\ j \end{bmatrix}$. Thus a coded stretch corresponds to a sequence

$$\cdots \begin{matrix} i(-1) & i(0) & i(1) & i(2) \\ j(-1) & j(0) & j(1) & j(2) \end{matrix} \cdots$$

which we think of as a pair of sequences of symbols, with the top row going right and the bottom one going left, thus:

$$\cdots \begin{matrix} i(-1) & \rightarrow & i(0) & \rightarrow & i(1) & \rightarrow & i(2) \\ j(-1) & \leftarrow & j(0) & \leftarrow & j(1) & \leftarrow & j(2) \end{matrix} \cdots$$

If the stretch ends we imagine the two sequences joined, thus:

$$\begin{matrix} i(-1) & \rightarrow & i(0) & \rightarrow & i(1) & \rightarrow & i(2) & \\ \uparrow & & & & & & & \cdots \\ j(-1) & \leftarrow & j(0) & \leftarrow & j(1) & \leftarrow & j(2) & \end{matrix}$$

In this way a finite coded stretch corresponds to a single cycle (since the top and bottom sequences close up at each end), a one-sided infinite coded stretch represents a single infinite sequence, and a two-sided infinite coded stretch represents two infinite sequences.

We now define $\overline{\varphi}$ to act on coded stretches as φ would. To be precise, suppose $\varphi(x)(0)$ depends on the word $x|_{[-I,I]}$. Then $\overline{\varphi}(x)$ will depend on the word $x|_{[-RI,RI]}$. The only effect $\overline{\varphi}$ has is on occurrences of $a \cup \alpha_{i,j}$. Such an occurrence is transformed by $\overline{\varphi}$ into $a \cup \alpha_{i'j'}$, as follows. To determine i' it looks right and left at the neighboring IR symbols of x in every direction. This suffices for $\overline{\varphi}$ to follow the top row of the coded stretch which it is considering and recover I symbols in each direction; it may wrap around in case the coded stretch ends. It then applies $\varphi$to these $2I+1$ symbols to obtain the symbol i'. The symbol j' is obtained similarly by following the bottom sequence of the coded stretch and applying φ to $2I+1$ symbols from it.

It is straightforward to check that $\varphi \mapsto \overline{\varphi}$ is a group homomorphism $\mathrm{aut}\{1,\ldots,M\} \to \mathrm{aut}(X)$, and it is an injection because to every point in the full shift there corresponds at least one coded stretch representing the same sequence. Note that the action on finite coded stretches corresponds to the action of φ on periodic points of the full shift. Also note that under this embedding the shift on $\{1,\ldots,M\}^{\mathbb{Z}}$ is not mapped to the shift on X, but to the automorphism that "shifts" the data in coded stretches.

5.2. **The Multidimensional Case.** Let $X \subseteq \Sigma^{\mathbb{Z}^d}$ be a positive entropy SFT. In order to adapt this construction to the multidimensional setting we would like to define coded stretches in two-dimensional words. However, we cannot guarantee that there exist periodic occurrences of a pattern, or even "almost periodic" occurrences. We therefore require explicitly that the minimal points are dense in X. This is a natural requirement from a dynamical point of view, and is satisfied, for instance, if the periodic points are dense. It is equivalent to the condition that for every finite $F \subseteq \mathbb{Z}^d$ and every $a \in \Sigma^F$, there is an $x \in X$ such that for some R there is an occurrence of a in $(T^u x)|_{B_R}$ for every u.

Fix $M \in \mathbb{N}$; we will show how to embed $\mathrm{aut}\{1,\ldots,M\}^{\mathbb{Z}}$ in $\mathrm{aut}(X)$. Choose a marker a with (at least) $1+M^2$ completions, the first M^2 of which we enumerate as before as $\{\alpha_{i,j}\}_{1\leq i,j\leq M}$. It will be convenient to think of the other completion(s) of a as "blank", i.e. not containing any data. We will call an occurrence of a word of the form $a \cup \alpha_{ij}$ an *occupied cell*. Let $R > 0$ be such that a appears R-syndetically in some word in X.

Let $B_k^+, B_k^- \subseteq \mathbb{Z}^d$ denote the right and left half-cubes of side k, defined by $B_k^+ = \{u \in B_k \,:\, u_1 > 0\}$ and $B_k^- = \{u \in B_k \,:\, u_1 < 0\}$. Let $x \in X$ and suppose there is an occupied cell in x at position u. We say that an occupied cell in x at position v in x is the *next* occupied cell (following the occupied cell at u) if (i) $v \in B_{2R}^+(u)$, (ii) this is the only occupied cell satisfying (i), and (iii) there is no other occupied cell in x for which (i),(ii) hold except the one at u. We say that an occupied cell of located in x at v is the *previously* occupied cell (before the one at u) if the one at u is the next occupied cell after the one at v. Note that the relations defined above is local in the sense that it is determined by the pattern induced by x on a finite neighborhood of u.

For $x \in X$ we now define a coded stretch as a maximal sequence of occupied cells in x occurring at locations $u(n) \in \mathbb{Z}^d$, such that for each n, the cell at $u(n+1)$ is the next occupied cell after the one at x. As before, we allow finite, one-sided and two-sided sequences. Note that the first coordinate of the locations of these words is a strictly increasing sequence, so there is no danger of loops forming.

For $\varphi \in \mathrm{aut}\{1, \ldots, M\}^{\mathbb{Z}}$ define $\overline{\varphi} \in \mathrm{aut}(X)$ to act on coded stretches as before; the only difference is that now the geometry of a coded stretch is a little looser, no longer consisting of periodically occurring occupied cells; nonetheless the next occupied cell is uniquely defined and everything goes through.

6. SFTs with small automorphism group

In this section we will construct some $\mathbb{Z}^2$-SFTs with small automorphism group. We give two examples. The first is an SFT with positive entropy whose automorphism group is of the form $\mathbb{Z}^2 \oplus G$, where G is a locally finite virtually simple group. The second example, which is slightly more elaborate, produces a strongly mixing SFT whose automorphism group is of the form $\mathbb{Z}^2 \oplus H$ where H is locally finite and factors onto a virtually simple group.

Both constructions are modifications of a certain zero entropy SFT X which we describe first. The points of X typically display a hierarchical nested pattern of squares, and in this regard X is somewhat reminiscent of Robinson's minimal SFT [9]. However the systems are quite different; indeed, a key feature of our example is that it is not minimal.

6.1. **A zero-entropy system.**

It will be convenient to think of the symbols of the language as square tiles of unit side length, which may be centered at points in $\mathbb{Z}^d$. The adjacency rules will restrict which tiles can be neighbors: we now say that $u, v \in \mathbb{Z}^2$ are neighbors if $\|u - v\|_\infty = 1$ (so our SFT is defined by a set of allowable 2×2 patterns). Thus each tile t has eight neighbors, four of which share an edge with it (the tiles directly above, below, to the right and to the left of t), and four sharing a corner with it. The latter four we call the NorthWest (NW), NorthEast (NE), SouthEast (SE) and SouthWest (SW) neighbors according to their position.

The alphabet Σ consists of 33 tiles. The first of these is a blank tile (which is represented by a black square in our diagrams). The remaining 32 tiles have arrows drawn on them with the head and tail meeting (different) sides of the square at their midpoint and perpendicular to the side. There are two kinds of arrow tiles: straight arrows where the arrow leaves an edge and goes straight to the opposite edge, and corner tiles where the arrow leaves an edge and turns clockwise 90°. See figure 6.1. The straight arrows are referred to by the direction of the arrow: up, down, left, right.

The corner tiles are referred to as NW,NE,SE,SW according to their position is the square they can form.

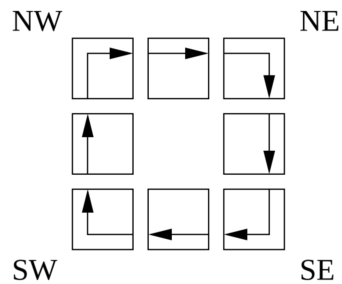

FIGURE 6.1. The arrow tiles (uncolored)

The first rule is that tiles must be arranged so that the head of every arrow meets the tail of an arrow on one of the neighboring tiles. For an occurrence of an arrow tile, define the successor and predecessor as the tile bordering on the head or tail of the arrow respectively. We define a path in a word $x \in \Sigma^{\mathbb{Z}^2}$ to be an injection $p : \mathbb{Z} \to \mathbb{Z}^2$ or $p : \mathbb{Z}/n\mathbb{Z} \to Z^2$ for some n, such that the tile at $p(i)$ is an arrow tile for all i and is the successor of the tile at $p(i-1)$. We will generally identify a path with its image, i.e. the collection of tiles in the path. Paths may be infinite, as in the first case, or finite, as in the second; paths of the second type are called *closed*. Since all the arrows turn counter-clockwise, every closed path is rectangular.

Note that we use the term *rectangle* for a set of the form $[i;j] \times [i',j'] \subseteq \mathbb{Z}^2$, but use the term *rectangular path* to mean a path which forms the boundary of a rectangle. We will also sometimes use the term rectangle to refer to a rectangular pattern. A similar remark holds for the terms *square* and *square path*.

In addition to the arrow marks, the arrow tiles each occurs in one of four colors; the colors too are called NW,NE,SE,SW. Thus for instance, there is a NW corner in each of the colors NW,NE,SE an SW. We represent the colors by different shades of gray; NW the lightest, and proceeding clockwise to SW, which is the darkest.

The allowed 2×2 patterns are precisely those which appear in figure 6.2, in which each allowed 2×2 configuration appears. In more detail, the adjacency rules are as follows:

(1) An arrow tile cannot be the successor of an arrow tile of a different color (hence every path is monochromatic).
(2) Every blank tile is surrounded by a 3×3 square path (which may be of any color).

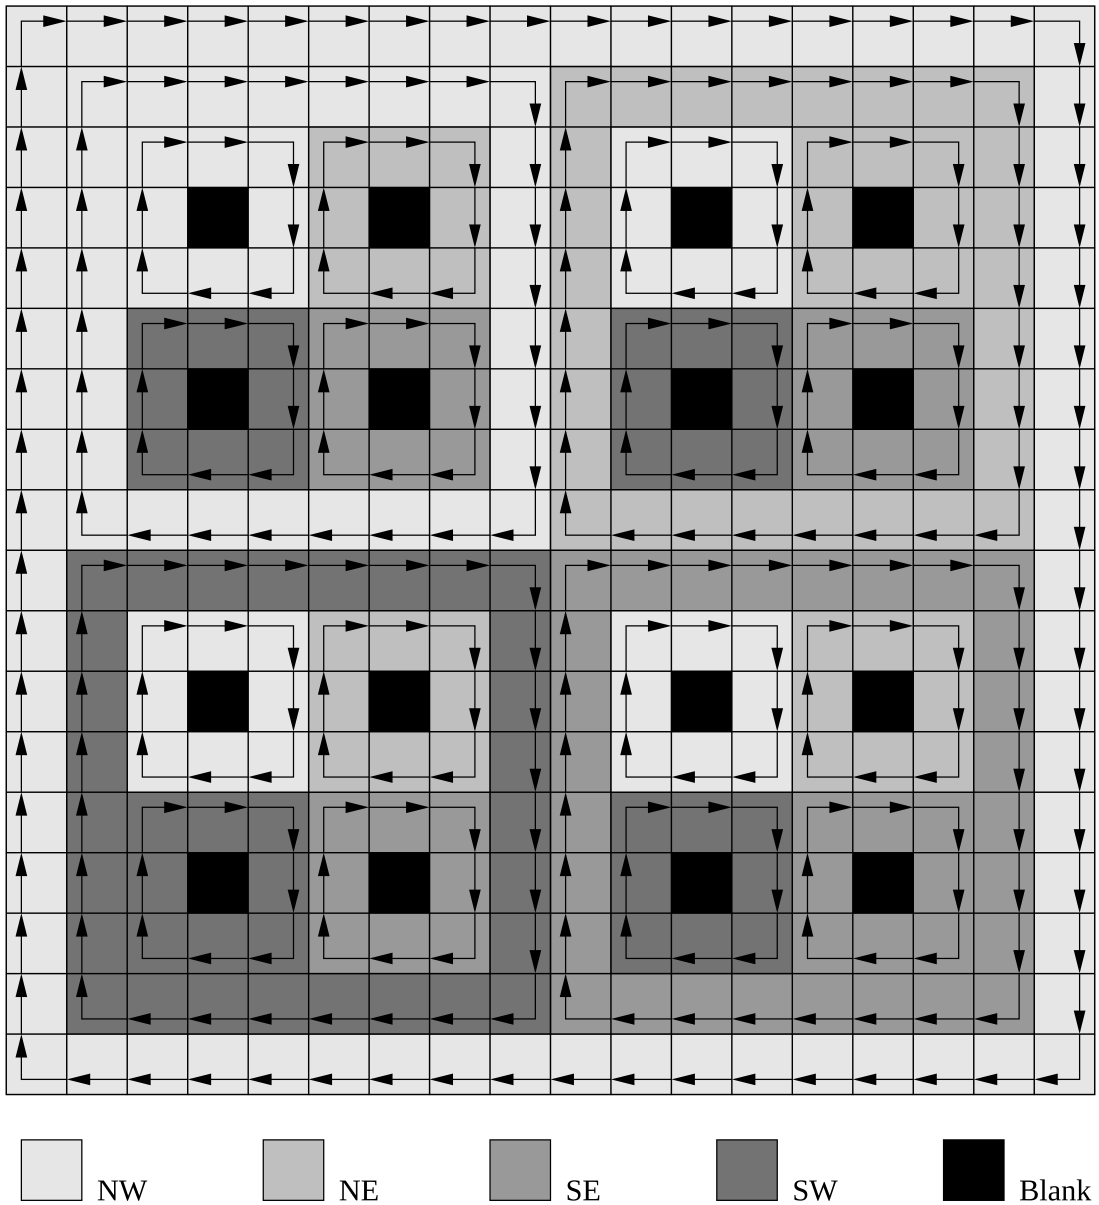

Figure 6.2. The level-3 square with NW colored boundary

(3) Corners may not appear consecutively in a path.
(4) When two corner tiles are horizontally adjacent the left one must be colored NW or SW and the right one NE or SE, respectively. When two corner tiles are vertically adjacent, the top one must be colored NW or NE and the bottom

on SW or SE, respectively. This determines which closed paths can abut on each other.

(5) The only tiles which may be the SE neighbor of a NW corner tile are (a) a NW corner tile colored NW, or (b) a blank tile. A corresponding condition holds for other corner types.

(6) Every NW corner tile which is colored NW has a NW corner tile as its NW neighbor (the neighbor's color is not restricted). A similar statement holds for other corner types.

(7) There are restrictions on the adjacency of straight arrows to each other (the allowed combinations are those appearing in figure 6.2).

Let $X \subseteq \Sigma^{\mathbb{Z}^2}$ be the SFT defined by these rules. We next describe the structure of configurations in X.

For $n = 1, 2, \ldots$ set $k(n) = 2^{n+1} + 2^{n-1} - 2$. For $\theta \in \{$NW,NE,SE,SW$\}$ define patterns $a_n^\theta \in \Sigma^{B_{k(n)}}$ inductively as follows: a_1^θ is the 3×3 square with a blank in the center, surrounded by a square path colored θ. For the induction step, define a_{n+1}^θ by constructing the square pattern b by arranging the squares $a_n^{\mathrm{NW}}, a_n^{\mathrm{NE}}, a_n^{\mathrm{SE}}, a_n^{\mathrm{SW}}$ is a square of side $2k(n)$, with a_n^ω in position ω within the large square. Next surround the composite square with a square path of color θ; it may be verified that this construction is consistent with the adjacency rules. Thus the side length of a_{n+1}^θ is $2k(n) + 2 = k(n+1)$. Note that the four patterns described at the n-th stage differ only in the color of their boundary path. The words a_n^θ for $n \geq 1$ are called level-n squares. The four level n squares which appear as subpatterns of a level-$(n+1)$ square will be referred to as siblings, and the level-$(n+1)$ pattern containing them their parent. We will also refer to them by their relative position inside the level-$(n+1)$ square, i.e. as the NW-subsquare, the SE-subsquare, etc. It will be convenient to refer to blank tiles as level-0 squares.

Let $\omega = (\omega_1, \omega_2, \ldots)$ be a sequence with $\omega_i \in \{$NW,NE,SE,SW$\}$. For each n the sequence $\omega_n, \omega_{n-1}, \ldots, \omega_1$ defines a decreasing sequence of level-i subsquares, $i = n, n-1, \ldots, 1$, of $a_n^{\omega_n}$, namely start with $a_n^{\omega_n}$, and given an level-i subsquare choose the subsquare one level down whose border is ω_{i-1} (and is also in relative position ω_{i-1}). Let $x_{\omega,n}$ be the pattern obtained by translating $a_n^{\omega_n}$ so that the level-1 subsquare determined by $\omega_n, \omega_{n-1} \ldots \omega_1$ is at the origin. The sequence $x_{\omega,n}$ is seen to be consistent and its union covers $\mathbb{Z}^2$ (note that $x_{\omega,n}$ covers B_n for every ω). We have thus defined a unique point $x_\omega \in X$ by $x_\omega = \cup_{n=1}^\infty x_{\omega,n}$.

Given $\omega, \omega' \in \{\text{NW,NE,SE,SW}\}^{\mathbb{N}}$, both $x_{\omega,n}$ and $x_{\omega',n}$ are translates of a level-n square, and each level-n square contains copies of all four level-$(n-1)$ squares. It follows that the orbit closure of x_ω and $x_{\omega'}$ are the same. It will develop that these are precisely the transitive points of X.

Lemma 18. *Suppose $F \subseteq \mathbb{Z}^2$ is a finite set consisting of all points within a closed path. Suppose $w \in \Sigma^F$ appears in X and that the pattern $w|_{\partial F}$ is a closed path (each point contains an arrow pointing to the next point). Then F is a translate of $B_{k(n)}$ and w is congruent to a_n^θ for some n and $\theta \in \{NW,NE,SE,SW\}$.*

Proof. By induction on $|F|$. Since every closed path is rectangular F is a rectangle, so we may write $F = [i;j] \times [i';j']$. It is clearly not possible that $i = j$ or $i' = j'$, nor that $j = i+1$ or $j' = i'+1$ since then we would have two corners consecutive in the path $w|_{\partial F}$. Assume that the color of $w|_{\partial F}$ is NW; the other cases are treated similarly.

Let $E = F \setminus \partial F = [i+1; j-1] \times [i'+1, j'-1]$. If E is a 1×1 square then $w|_E$ must be a blank tile and $w = a_1^{\text{NW}}$. If this is not the case, then no blank tile can appear in ∂E (because a blank tile must be surrounded by a closed path). Thus every tile in ∂E is an arrow tile and hence belongs to a closed path, and if $E_1, \ldots, E_m \subseteq E$ are the rectangles surrounded by these paths, then $E_1, \ldots, E_m$ are disjoint and $\partial E \subseteq \cup E_i$. If E_i contains the NW corner of E then $w|_{\partial E_i}$ must be colored NW, because the NW corner of E_i is the SE neighbor of the NW corner of F. Similarly if E_j contains some other corner of E then it is colored according to the position of that corner. Thus each corner of E is contained in a different one of the E_i's. Let $E_{i(1)}, E_{i(2)}, E_{i(3)}, E_{i(4)}$ be the E_i's containing the NW,NE,SE,SW corners of E, respectively. Consider $E_{i(1)}$; since it does not contain the NE corner of E there is some E_j whose NW corner is the tile to the right of the NE corner of $E_{i(1)}$. By the color restrictions on neighboring corners, E_j must be colored NE. For this reason, the NE corner tile of E_j cannot have a corner tile to its right, so $E_j = E_{i(2)}$. We have shown that $E_{i(1)}$ shares an edge with $E_{i(2)}$. Similarly, $E_{i(2)}$ and $E_{i(3)}$ share an edge, as do $E_{i(3)}, E_{i(4)}$ and $E_{i(4)}, E_{i(1)}$. Thus $\partial E \subseteq E_{i(1)} \cup E_{i(2)} \cup E_{i(3)} \cup E_{i(4)}$. Now by the induction hypothesis each $E_{i(j)}$ is of the form $a_{n(i)}^{\theta(i)}$ and it is simple to verify that the only way to arrange four squares Q_1, Q_2, Q_3, Q_4 so that each consecutive pair share an edge and the outer perimeter is a rectangle is when all four squares are of the same size and they are arranged in a square. It follows that $w|_E$ is the union of four level-n squares arranged in a square and their boundaries are colored as in the construction of a_{n+1}^{NW}, so $w = a_{n+1}^{\text{NW}}$. □

Similar reasoning gives the following:

Lemma 19. *There is no legal tiling of a set $F = [0;\infty) \times [i;j]$ such that ∂F is a path.*

This give us the following important property of X:

Lemma 20. *If $x \in X$ contains a level-n square for some $n \geq 1$, then this square is a subsquare of some level-$(n+1)$ square in X.*

Proof. Let $x \in X$ and suppose a_n^θ appears in x at u. Then the θ-corner of this level-n square is colored θ so its θ-neighbor tile, located at v, is also a θ-corner colored θ. Consider the path that this tile belongs to. If this path is finite then it must be the boundary of a level-$(n+1)$ square because one of its subsquares is a level-n square. Otherwise it would have to be L-shaped; but this is ruled out by reasoning as in lemma 19. □

Corollary 21. *Let $x \in X$ and suppose x contains a blank tile. Then x is a translate of a x_ω for some $\omega \in \{NW,NE,SE,SW\}^{\mathbb{N}}$.*

Proof. Every blank tile is surrounded by a path and so x contains a level-1 square, and from the previous lemma we get an increasing sequence E_n of $k(n) \times k(n)$ squares such that $x|_{E_n} = a_n^{\omega_n}$ for some $\omega_n \in \{\text{NW,NE,SE,SW}\}$. Thus x is a translate of x_ω. □

We have shown that if X is transitive then every transitive point is of the form x_ω. Our next goal is to show that the x_ω are indeed transitive points for X. To this end we give some dynamical characterizations of the points x_ω which will be useful later.

Lemma 22. *Every point x_ω is recurrent under the $\mathbb{Z}$-actions in directions e_1, e_2 and is not fixed by T^{e_1} or T^{e_2}.*

Proof. The recurrence of x_ω is clear from the construction. For the second statement note that corner tiles with the same orientation cannot share a side, so a pattern containing corner tiles cannot be fixed by T^{e_1} or T^{e_2}; and every x_ω contains corner tiles. □

For the next lemma the reader may refer to figure 6.3, which schematically depicts portions of four types of exceptional points. Additional points may be obtained by coloring and rotating the four below.

Lemma 23. *If $x \in X$ is not a translate of one of the x_ω, then either x is not recurrent under the $\mathbb{Z}$-actions in one of the directions e_1, e_2, or else it is fixed by one of the transformations T^{e_1} or T^{e_2}. Every such point is in the orbit closure of the x_ω's.*

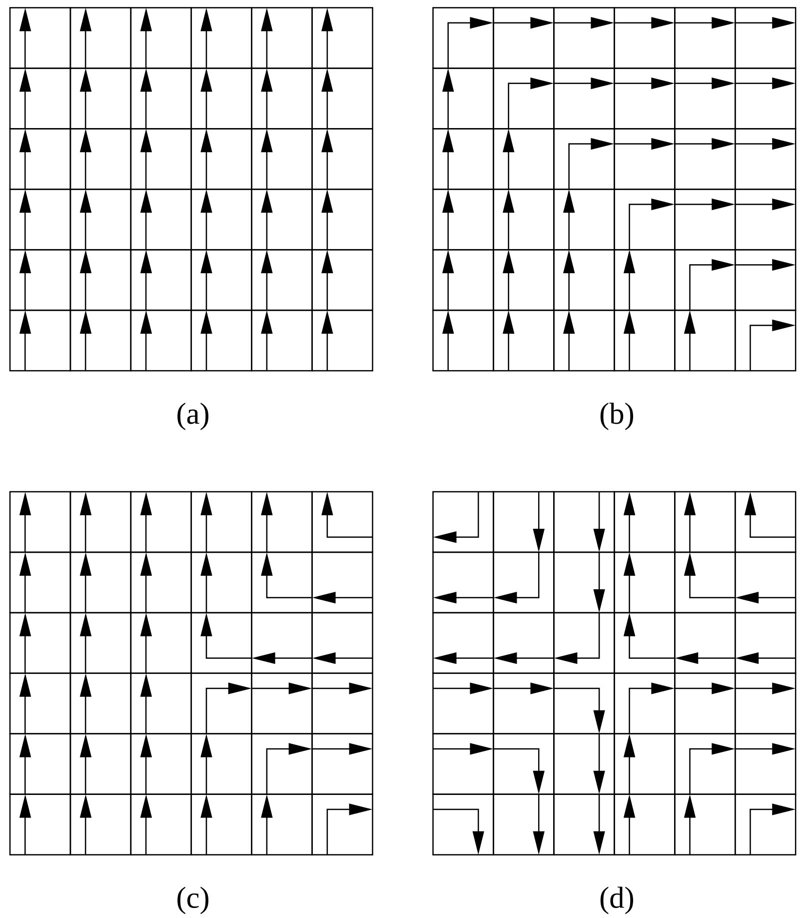

FIGURE 6.3. Exceptional points (uncolored)

Proof. Suppose $x \neq x_\omega$ for every $\omega \in \{\text{NW,NE,SE,SW}\}^{\mathbb{N}}$. Then by lemma 22 x does not contain a blank, and by lemmas 18 and 19, every path in x is infinite and contains at most one corner tile.

Suppose x contains no corner tiles. Then every path is a straight line, and these must all be parallel (since they cannot intersect). It follows that x is the union of parallel lines, and each line is fixed by translation in the direction it is pointing (or the opposite direction); thus x is fixed by T^{e_1}or T^{e_2} and is of the form (a) in figure 6.3.

By lemmas 18 and 19, an infinite path containing a corner tile must be L-shaped, and divides $\mathbb{Z}^2$ into two regions, one of which is a quarter space (i.e. a translate of the intersection of two of the half-spaces $u + (\mathbb{Z} \times \pm\mathbb{N})$ or $v + (\pm\mathbb{N} \times \mathbb{Z})$). Such a region can contain no infinite straight lines, and one sees that in this region every path must be L-shaped with the same orientation as the original corner tile, and that the union of these must fill the region.

Suppose all corner tiles in x have orientation θ. Let Q be the union of all paths containing a corner tile; this is either all of $\mathbb{Z}^2$ or a quarter-space. But it cannot be a quarter-space because the complement of a quarter space cannot be covered with non-intersecting straight lines. Thus x is of the form (b) in figure 6.3. x is not fixed by T^{e_1}or T^{e_2}, but its one-sided orbit under either of these or their inverses converges to a fixed point of that transformation, x is not recurrent in either of these directions.

Suppose the corner tiles appearing in x are of two orientations, θ an θ'. Let Q be the quarter-space which is the union of paths containing θ-corners, and Q' the quarter-space which is the union of paths containing θ'-corners. Since $\mathbb{Z}^2$ is not the union of two quarter-spaces, x also contains straight lines; let L be their union. One can check that then L must be a half space, and x is of the form (c) in figure 6.3. Once again x is not fixed by T^{e_1}or T^{e_2} but the one-sided action of one of these or one of their inverses send x to a fixed point of that transformation, so x is not recurrent under T^{e_1}or T^{e_2}.

If x contains corner tiles in exactly three orientations, then we would have to cover the complement of three quarter-spaces with straight lines, and this is impossible; so this case does not arise.

Suppose that x contains corners of all four orientations. Let Q^θ be the union of paths which contain θ-corners. Since the complement of the union of four quarter-spaces cannot be covered by straight lines unless it is empty, $\mathbb{Z}^2 = \cup_\theta Q^\theta$ and x must have the form (d) in figure 6.3. As before x is not fixed by T^{e_1} or T^{e_2} nor is it recurrent under the action of any one of them.

Finally, the descriptions we have given of x are still somewhat variable in that the paths may bear various colors. These do not affect any of the statements made so far. We leave it to the reader to work out which colorings/directions can be combined and to show that all these points are in the closure of the x_ω's. □

Corollary 24. *The transitive points in X are precisely the translates of the x_ω's.*

Next, we show that the recurrence properties of a point x_ω essentially determines ω. Let us say that the NW,SW level-n subsquares of a level-$(n+1)$ square are on the *left side* of the level-$(n+1)$ square, and similarly that the NW,NE subsquares are the top subsquares, etc.

Lemma 25. *Let $x \in X$, let a be a level-n square appearing in x. Let $u \in \mathbb{Z}^d$ be the position of some sub-pattern p in the interior of a which contains a corner tile. Then a*

is on the left side of b *if and only if* p *repeats at* $u + k(n)e_1$*, but fails to repeat at least at one of the locations* $u - k(n)e_1$ *or* $u + 2k(n)e_1$*.*

Proof. Let b be the parent rectangle of a, so b is a level-$(n+1)$ square, and let c be the level-$(n+2)$ square which is the parent of b.

Suppose first that a is on the left side of b. Then there is another level-n square immediately to the right of a, i.e. one of a's three siblings. Since $k(n)$ is the width of level-n squares, p repeats at $u + k(n)e_1$.

We now distinguish two cases according to the position of b in c. Suppose that b is on the right side of c. Then immediately to the left of b there is another level-$(n+1)$ square b', and so to the left of a we find the boundary paths of b and b', followed by a level-n square. Thus p repeats at $u - (k(n) - 2)e_1$. Since no corner tile can appear in X twice with displacement 2 and p contains a corner tile, we conclude that p does not repeat at $u - k(n)e_1$.

Otherwise b is on the left side of c. Then there is a level-$(n+1)$ square to the right of b and an analysis similar to the above shows that p repeats at offset $u + (2k(n) + 2)e_1$, so it does not repeat at $u + 2k(n)e_1$.

In the converse direction, suppose that a is on the right side of b and we will show that the condition in the lemma fails. For this purpose we may assume that p is a single corner tile. An analysis similar to the above shows that the tile at u repeats at $u - k(n)e_1$. We will show that either the tile at u fails to repeat at $u + k(n)e_1$ or else it repeats also at $u + 2k(n)e_1$, which gives the desired contradiction.

Indeed if the tile at u repeats at $u + k(n)e_1$ then b is not on the left side of c, for if it were then the tile would repeat at $u + (k(n) + 2)e_1$. So b is on the right side of c. Let c' be the level-$(n+2)$ square to the right of c; the location $u + k(n)e_1$ falls within c', and actually it must fall inside the left level-$(n+1)$ subsquare of c', and on the left side of that. Hence the level-n square to which $u + k(n)e_1$ belongs is on the left of its surrounding level-$(n+1)$ square, so there is another level-n square immediately to its right. But then the tile at $u + k(n)e_1$ repeats at $u + k(n)e_1 + k(n)e_1$, as claimed. □

Similar statements characterize when a square is on the top, right or bottom of its surrounding square. It follows that the sequence ω of a point x_ω is determined by its recurrence properties.

Theorem 26. *The automorphism group of* X *consists only of the shift action.*

Proof. Let $\varphi \in \mathrm{aut}(X)$ and $x = x_\omega$. Since x is transitive its image is transitive, so $\varphi(x) = T^v x_{\omega'}$ for some $v \in \mathbb{Z}^2$ and $\omega' \in \{\mathrm{NW},\mathrm{NE},\mathrm{SE},\mathrm{SW}\}^{\mathbb{N}}$. We will show that $\omega_n = \omega'_n$

for every large enough n. This shows that $x_{\omega'} = T^u x_\omega$, and hence $\varphi(x) = T^{u+v}x$. Since x is transitive this shows that $\varphi = T^{u+v}$.

Let $R = R(\varphi)$ be the window radius of φ. Since both $x, \varphi(x)$ are transitive, we may choose N large enough that both $a = x|_{B_N}$ and $a' = (\varphi(x))|_{B_{N-R}}$ contain a level-1 square. Note that if a appears at v in x then a' appears at v in $\varphi(x)$.

Let b_n denote the level-n square in x containing $x|_{B_N}$; this is defined for every large enough n. Similarly let b'_n denote the level-n square in $\varphi(x)$ containing a'. For all n note that ω_n is the position of b_n in b_{n+1} and ω'_n is the position of b'_n in b'_{n+1}, and we must show that these coincide.

Fix a large enough n so that b_n is defined and fix a corner tile in a'. By the previous lemma, for n large enough we can determine ω_n in terms of the distances between b_n and the level-n tiles to its left and right. But if one of these distances is r, say, then the said corner tile in a' will also repeat at displacement r in the same direction, because a' in fact repeats with this displacement. Hence by the lemma, $\omega'_n = \omega_n$ for all n large enough, as claimed. □

Corollary 27. $h(X) = 0$.

Proof. We have seen (theorem 1) that positive entropy implies the existence of non-trivial automorphisms. □

Also note that since the distance between NW corners of level-2 squares in X is always even, X is not weakly mixing. We will modify the construction in the coming sections in order to achieve these properties.

6.2. **Adding entropy.** To raise the entropy of X we will resort to two tricks. The first, which we analyze in this section, is very simple: we simply add a second blank tile to Σ, which is completely interchangeable with the original blank tile. Let us call the first blank tile 0 and the new blank tile 1. Denote the new alphabet $\widehat{\Sigma}$ and the new system $\widehat{X}$. Alternatively, $\widehat{X}$ can be described as a factor of the system $X \times \{0,1\}^{\mathbb{Z}^2}$ via the map $\rho : X \times \{0,1\}^{\mathbb{Z}^2} \to \widehat{\Sigma}^{\mathbb{Z}^2}$ defined by

$$\rho(x,y)(0) = \begin{cases} x(0) & \text{if } x(0) \text{ is not a blank} \\ y(0) & \text{otherwise} \end{cases}$$

Since $\{0,1\}^{\mathbb{Z}^2}$ is strongly mixing and X is transitive, we conclude that so is $X \times \{0,1\}^{\mathbb{Z}^2}$ and this is inherited by $\widehat{X}$.

We also have a factor map $\pi : \widehat{X} \to X$ which forgets the color of blank tiles. The preimage of a point $x_\omega \in X$ under π in uncountable and consists of all $\{0,1\}$-colorings of its blank tiles, whereas each $x \in X$ which is not of the form x_ω has a single preimage under π and $\pi(x) = x$. A simple calculation shows that for any $x_\omega \in X$ the blanks appear with positive density, and each of these may independently be replaced with the new blank tile, we conclude that $\widehat{X}$ has positive entropy (in fact its entropy is easily seen to be equal to $\frac{1}{25}\log 2$, since $\frac{1}{25}$ is the maximal density of blank tiles in X).

Define a colored level-n square to be any level-n square whose blanks have been colored arbitrarily by $0, 1$. Clearly every point $x \in \widehat{X}$, except for the exceptional points of figure 6.3, is the increasing union of colored level-n squares.

Lemma 28. *Let $x \in \widehat{X}$, let a be a level-n colored square appearing in x. Let $u \in \mathbb{Z}^d$ be the position of some sub-pattern p in the interior of a which does not contain any blanks and contains a corner tile. If p repeats at $u + k(n)e_1$, but fails to repeat at least at one of the locations $u - k(n)e_1$ or $u + 2k(n)e_1$, then p is on the left side of a.*

This lemma is one direction of lemma 25 adapted to $\widehat{X}$; the proof is exactly the same as the original. Using it we get

Theorem 29. *Let $\varphi \in \operatorname{aut} \widehat{X}$. Then there exists a unique $u \in \mathbb{Z}^2$ such that for every $x \in X$ if $x(0)$ is not one of the blank tiles then $(T^u\varphi)(x)(0) = x(0)$.*

Proof. Fix $x_\omega \in X$ (so $x \in \widehat{X}$ also). Using the recurrence properties of x_ω under T^{e_1} and T^{e_2} again we see that $\varphi(x_\omega)$ is not one of the exceptional points in figure 6.3, and hence neither is $\varphi(x_\omega)$, so $\varphi(x_\omega)$ contains blanks tiles. We claim that there is a $u \in \mathbb{Z}^2$ such that $\pi(T^u\varphi(x_\omega)) = x_\omega$; the proof of this uses the same argument of theorem 26, only substituting lemma 28 for one direction of 25.

Now if $R = R(\varphi)$ is the window width of φ, note that for $n > R$ and $x \in \widehat{X}$, every point on the boundary of a (colored) level-n square in x is surrounded by non-blank tiles within a distance of n in every direction, and hence $T^u\varphi$ acts identically on the symbols on the boundaries of level-n squares in x. But this determines uniquely the position and value of all arrow tiles in $T^u\varphi(x_\omega)$, and since the level-n boundary tiles agree with those of x_ω we see that $T^u\varphi(x)(v) = x(v)$ whenever $x_\omega(v)$ is not a blank tile.

The case where x is an exceptional point of X follows by approximation. □

Define now $G_n < \operatorname{aut} \widehat{X}$ to be the group of automorphism which act by permuting the patterns on the interiors of colored level-n squares, independently of the pattern in their complement. G_n is clearly a group and $G_n < G_{n+1}$. Let $G = \cup G_n$. If $\varphi \in \operatorname{aut} \widehat{X}$, $R = R(\varphi)$ its window width and $u = u(\varphi) \in \mathbb{Z}^2$ as in the above theorem, then for $x \in \widehat{X}$

if $T^u\varphi(x)(v)$ contains a blank tile then it is determined only by the interior configuration of the colored level-$(R+\|u\|_\infty+1)$-square it belongs to (because $R(T^u\varphi) = R(\varphi)+\|u\|_\infty$, and every blank tile in a level-n square is at distance at least n from the boundary of the square). Also, $T^u\varphi$ acts identically on arrow tiles. Thus $T^u\varphi \in G_{R+\|u\|_\infty+1}$. If we identify $\mathbb{Z}^2$ with the shift action in $\operatorname{aut}\widehat{X}$ then $\operatorname{aut}\widehat{X} = \mathbb{Z}^2 + G$. In fact this is a direct sum, since the embedding of $\mathbb{Z}^2$ in $\operatorname{aut}\widehat{X}$ splits via the homomorphism $\varphi \mapsto u(\varphi)$. That this is a homomorphism follows from the uniqueness of u.

Finally, it is easy to see that G_n is just the full permutation group on the set of completions of a $k(n) \times k(n)$ square path (we simply construct an automorphism for any given permutation). Thus G_n is finite. Let $A_n < G_n$ be the simple subgroup of index 2; let $A = \cup A_n$. Then A is a simple subgroup of index 2 of G. We have proved:

Theorem 30. $\operatorname{aut}\widehat{X} = \mathbb{Z}^2 \oplus G$ *where* $\mathbb{Z}^2$ *is the shift action and* G *is locally finite and contains a simple subgroup of index* 2.

Unfortunately, $\widehat{X}$ factors onto X (by forgetting the tile coloring), so it is not weakly mixing. We correct this in the next section.

6.3. **Mixing things up.** In order to get mixing one needs to make a more fundamental change to the system. We retain the alphabet Σ of X, but now instead of the rule that blank tiles may not be adjacent to each other we now allow blank tiles to appear in rectangular configurations of dimensions 1×1, 1×2 , 2×1 and 2×2. As before, each such configuration is required to be surrounded by a closed path. All other adjacency rules are the same as before. We denote the resulting SFT by X_*. Note that $X \subseteq X_*$. Figure 6.4 describes part of a typical pattern in X_*.

It is no longer true that every closed path in X_* is a square, but a typical point in X_* is still made up of nested closed rectangular paths, as described below. First note that, as before, every closed path is a rectangle and its interior is the union of four rectangles, each of which has precisely one corner adjacent to the corner of the enclosing path; the proof of this is the same as of lemma 18.

Define a level-0 rectangle to be one of the four allowed rectangular patterns consisting of blanks, and define a level-1 rectangle to be a rectangular pattern whose boundary is a closed path surrounding a level-0 rectangle. For $n > 1$, define a rectangular pattern to be a level-n rectangle if the pattern along its boundary is a closed path and its interior is the union of four rectangular patterns, each of which is a level-m rectangle for some for $m < n$, and at least one of them is of level $n-1$. Thus the level-n squares from the previous section are level-n rectangles in the sense just defined. Every rectangular

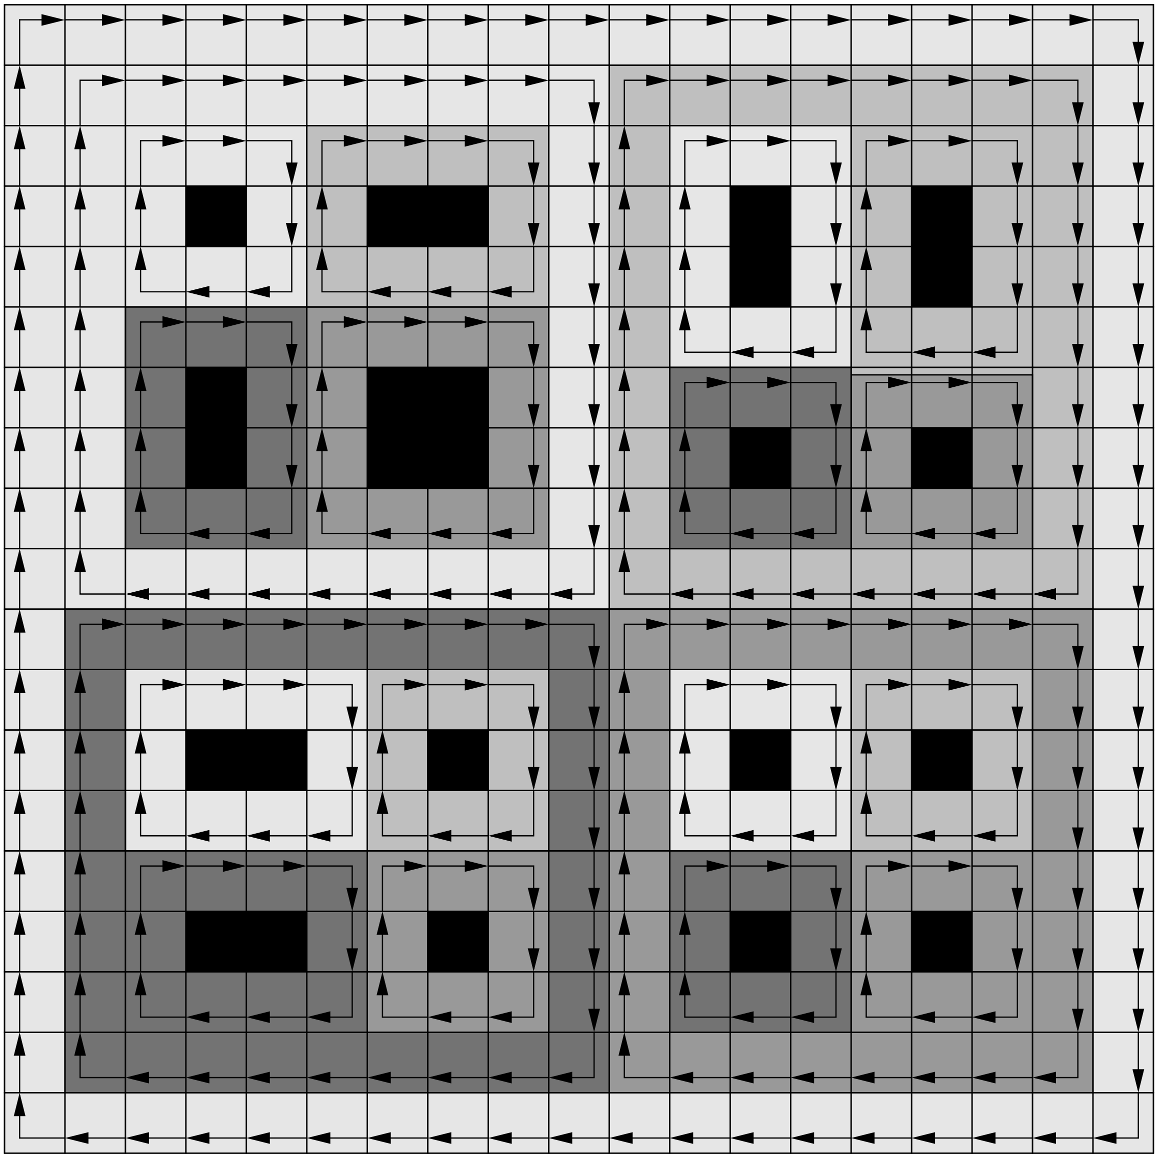

FIGURE 6.4. Part of a level-3 rectangle (the NW corner or a_3^{NW})

pattern whose boundary is a closed path is a level-n rectangle for some n. Note also that if w is a level-n rectangle then we may change the color of the border of w arbitrarily (keeping it monochromatic), and the result is an admissible pattern which is also a level-n rectangle.

For $n \geq 1$ let $k^-(n), k^+(n)$ denote, respectively, the minimal and maximal lengths of a side of a level-n rectangle. Let $k(n)$ be as before.

Lemma 31. *$k^-(n) = k(n) = 2^{n+1} + 2^{n-1} - 2$ and $k^+(n) = 2^{n+1} + 2^n - 2$ for every $n \geq 1$, and for $n \geq 2$ each of the four subrectangles making up a level n rectangle are level $n-1$ rectangles.*

Proof. The proof is by induction. For $n = 1$ we see that the smallest level-1 rectangle is of dimensions 3×3 and the largest is of dimensions 4×4, so $k^-(1) = 3 = k(1)$ and $k^+(1) = 4$, which is consistent with the claim.

Let w be a level-n rectangle, $n \geq 2$. By definition w consists of its boundary (a closed path) and the union of four rectangle of a lower level, with one of them being level-$(n-1)$. Suppose this is the NW subrectangle, for instance. Each of its sides is therefore of length at least $k^-(n-1) = 2^n + 2^{n-2} - 2$. Since this is strictly greater than $k^+(m)$ for $m < n-1$, and since the SW and NE subrectangles of w meet the NW rectangle along one of its edges, we see that they must be level-$(n-1)$ rectangles as well; they share an edge with the SE subrectangle so it too is level-$(n-1)$. It also follows that the minimal and maximal side lengths of w, denoted i, j respectively, satisfy

$$\begin{aligned} i &\geq 2k^-(n-1) + 2 \\ j &\leq 2k^+(n-1) + 2 \end{aligned}$$

so

$$\begin{aligned} k^-(n) &\geq 2k^-(n-1) + 2 \\ k^+(n) &\leq 2k^+(n-1) + 2 \end{aligned}$$

substituting the value of $k^{\pm}(n-1)$ gives $k^-(n) \geq 2^{n+1} + 2^{n-1} - 2$ and $k^+(n) \leq 2^{n+1} + 2^n - 2$; one checks by construction that these are actually equalities. $\square$

Another induction proves:

Lemma 32. *Let w be a level-n rectangle. Then w is a level-n rectangle if and only if every level-0 square in w is an isolated blank. A point $x \in X_*$ is in X is and only if all level-0 rectangles are isolated blanks.*

The following lemma, and its proof, are analogs of lemma 21.

Lemma 33. *If a point $x \in X_*$ contains a level-0 rectangle, then there is an increasing sequence of rectangles F_n with $\cup F_n = \mathbb{Z}^2$ such that $x|_{F_n}$ is a level n rectangle. Otherwise*

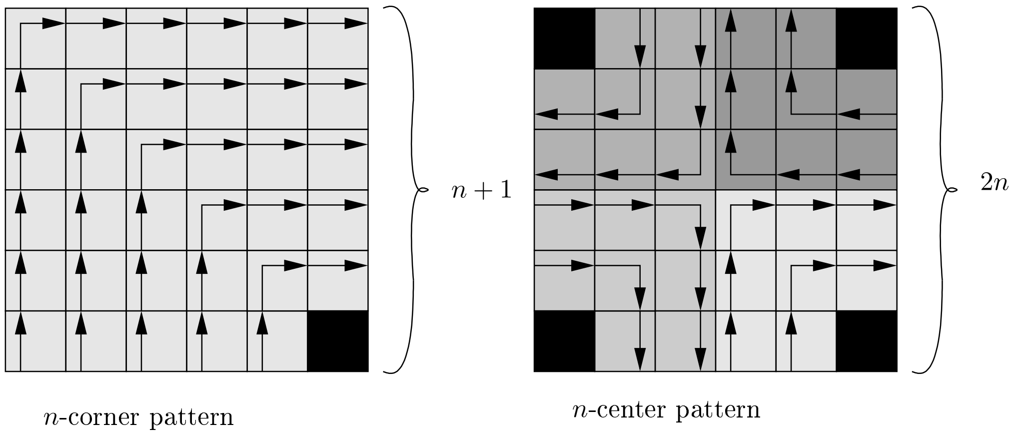

Figure 6.5.

(if x does not contain a level-0 rectangle) then it is one of the non-transitive points of X, as in figure 6.3.

We shall use that fact that for every $k^-(n) \leq i, j \leq k^+(n)$ there exists a level-n rectangle of dimensions $i \times j$. It follows that

Proposition 34. *X_* is transitive.*

Proof. As in the case of X, every block appears in some level-n rectangle, so it suffices to show that if a, b are level-n rectangles then they appear together in some point in X_*. We may embed a, b in level-$(n+1)$ rectangles a', b' and since we may change the boundary coloring of a', b' freely, we may assume that the boundary of a' is colored NW and of b' is colored SE. We may also find admissible rectangles c', d'so that a', b', c', d' may be arranged in a level-$(n+2)$ rectangle with a' the NW sub-rectangle and b' the SE rectangle. Since each level-$(n+2)$ square appears in X_* we are done. □

Our next goal is to show that every automorphism φ of X_* preserves X, and thus acts on X by translation.

For $n \geq 1$, define words $b_n^{\mathrm{NW}}, b_n^{\mathrm{NE}}, b_n^{\mathrm{SW}}, b_n^{\mathrm{SE}} \in \Sigma^{B_{n+1}}$ and $c \in \Sigma^{B_{2n+2}}$ as follows: b_n^{NW} is the $(n+1) \times (n+1)$ square pattern at the NW corner of a_n^{NW}, and c is the central $2n \times 2n$ square pattern of a_n^{NW} (the pattern c does not depend on the superscript NW, which was chosen arbitrarily). See figure 6.5. We call the b_n^{θ}'s n-corner patterns and call c the n-center pattern.

We use the following simple fact: whenever b_n^{NW} appears in a point $x \in X_*$, it is located in the NW corner of the level-n rectangle to which it belongs, and similarly when NW replaced with the other directions. Similarly, when c_n appears in a point $x \in X$ it is at the center of an instance of a level-n rectangle.

The following lemma is an immediate consequence of the size bounds on level-n rectangles:

Lemma 35. *Let $R > 0$ be an integer. For every sufficiently large n, the following holds. Let $x \in X_*$ and let $E, F \subseteq \mathbb{Z}^2$ be rectangles. Suppose $x|_F$ is a level-n rectangle and that E is a $(k(n) - R) \times (k(n) - R)$ square such that $x|_{F \cap E}$ contains a blank tile. Then $x|_E$ contains either an n-corner or the n-center of $x|_F$.*

Lemma 36. *Suppose $x \in X$ is a transitive point and $\varphi \in \operatorname{aut} X_*$. Then $\varphi(x)$ contains a blank tile.*

Proof. Since x is transitive for X, it is recurrent for the $\mathbb{Z}$-action determined by each of the directions $\pm e_1, \pm e_2$ but not fixed by translation in either of these directions. Therefore this is true also of $\varphi(x)$. Since it is not true for any of the exceptional points from lemma 22, by lemma 33 we are done. □

Proposition 37. *Suppose $x \in X$ and $\varphi \in \operatorname{aut}(X_*)$. Then $\varphi(x) \in X$.*

Proof. We prove this for the case that x is a transitive point of X; the general case follows.

Write $y = \varphi(x)$. By the previous lemma, y contains a blank tile, and if every level-0 rectangle in y is of dimensions 1×1 then by lemma 32 $y \in X$; thus we may assume that y contains a pair adjacent blank tiles.

Let $E_n \subseteq \mathbb{Z}^2$ be a sequence of $k(n) \times k(n)$ squares such that $x|_{E_n}$ is a level-n square and $\cup E_n = \mathbb{Z}^2$. Let $R = R(\varphi)$ be the window radius of φ and write $E'_n = E_n \setminus \partial E_n$ and $E''_n = E_n \setminus \partial_{R+1} E_n$. Then $x|_{E'_n}$ determines $y|_{E''_n}$; since $\cup E_n = \mathbb{Z}^2$ we also have $\cup E''_n = \mathbb{Z}^2$, so for large enough n the word $y|_{E''_n}$ contains a pair of adjacent blank tiles. Let $u \in E''_n$ be the location of this pair, and let F be the rectangle containing u such that $y|_F$ is a level-n rectangle. By lemma 35, either the n-center pattern of $y|_F$ or on of its n-corner patterns is in $y|_{E''_n}$.

Suppose for instance that $y|_{E''_n}$ contains the NW n-corner of $y|_F$ and that the two blanks at u are horizontally adjacent. Let $D \subseteq \mathbb{Z}^2$ be a $k(n) \times k(n)$ square such $x|_D$ is a level-n square which is a sibling of $x|_{E_n}$ and shares a vertical edge with it. Defining $D' = \partial D$ and $D'' = \partial_{R+1} D$ as before, we note that $x|_{D'} = x|_{E'_n}$ so $y|_{D''} = x|_{E''_n}$, and so $y|_{D''}$ also contains the NW n-corner pattern of a level-n rectangle in y which contains two horizontally adjacent blank tiles. However, the displacement between the n-corner patterns in $y|_{E''_n}$ and $y|_{D''}$ is exactly $k(n)$, so at least one of the level-n rectangles must be of width $k(n)$; by lemma 32 this is impossible because both contain a pair of horizontally adjacent blank tiles.

The other cases are treated similarly. In the case that $y|_{E''_n}$ contains an n-center pattern we must also take into account the relative position of u (the location of the two blanks) with respect the n-center patters. We omit the details. □

Theorem 38. *Let $\varphi \in \operatorname{aut} X_*$. Then there is a unique $u \in \mathbb{Z}^2$ and an N such that $T^{-u}\varphi$ fixes the boundary of every level-n rectangle for $n > N$.*

Proof. Let $R = R(\varphi)$ be the window width of φ. Let u be such that $T^{-u}\varphi|_X = \mathrm{id}_X$. If $x \in X_*$ is one of the exceptional points of lemma 22 then $\varphi(x) = T^u(x)$ because $x \in X$, so the claim certainly holds. So assume x is the union of its level-n rectangles. Note that for $n > R$, if $v \in \mathbb{Z}^2$ is the position in x of a point on the boundary of a level-n rectangle then $x|_{B_n(v)}$ is a legal B_n-pattern in X. Hence $T^{-u}\varphi(x)(v) = x(v)$. □

Define $G_n < \operatorname{aut} X_*$ as before to be the group of automorphisms which act by permuting the interiors of level-n rectangles independently of their context or the color of the boundary, so $G_n < G_{n+1}$. If $i \times j$ is a permissible size of a level-n rectangle, let $S_{i\times j}$ denote the group of permutations on admissible interiors of an $i \times j$ level-n rectangle. Let $D_n \subseteq \mathbb{N} \times \mathbb{N}$ be the set of pairs (i, j) such that $i \times j$ are the dimensions of a level-n rectangle, and let $S_{i,j}$ be the group of automorphisms of X_* which act by permuting the interiors of level-n rectangle of dimensions $i \times j$ independently of their exteriors. It is easy to verify that as $n \to \infty$ the number of completions of the boundary of an $i \times j$ level-n rectangle with $i \neq k^{\pm}(n)$ or $j \neq k^{\pm}(n)$ is unbounded (i.e. if $k^{\pm(n)} \neq i(n), j(n) \to \infty$ and $(i(n), j(n)) \in D_n$ then the number of completions of the boundary of an $i(n) \times j(n)$ level-n rectangle tends to infinity with n). Clearly,

$$G_n = \bigoplus_{(i,j)\in D_n} S_{i,j}$$

As before, let $G = \cup G_n$; from theorem 38 it follows that $\operatorname{aut} X_* = \mathbb{Z}^2 \oplus G$, where $\mathbb{Z}^2$ is the shift action.

We claim that there is a surjective homomorphism $G \to H$ with H is a virtually simple infinite group. We describe the kernel $K < G$ of this homomorphism explicitly. For $(i, j) \in D_n$ let

$$\pi_{i,j} : G_n \to S_{i,j}$$

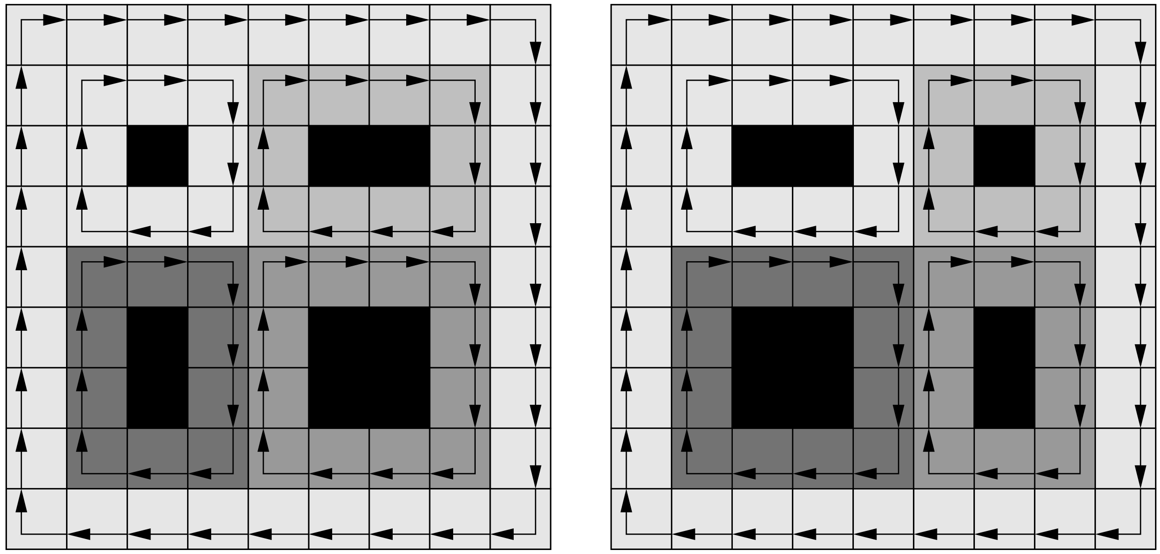

FIGURE 6.6. Two admissible completions of a single boundary

be the canonical projection (note that (i,j) determine n) . Define

$$K = \{\varphi \in G \,:\, \pi_{k(n),(k(n)+1)}(\varphi) = 1 \text{ for all large enough } n$$

This is clearly a normal subgroup of G since if $\varphi \in K$ and $\psi \in G$ then whenever n is such that $\pi_{k(n),(k(n)+1)}(\varphi) = 1$ and $\psi \in G_n$, we have

$$\pi_{k(n)\times k(n+1)}(\psi^{-1}\varphi\psi) = \pi_{k(n)\times k(n+1)}(\psi)^{-1} \cdot 1 \cdot \pi_{k(n)\times k(n+1)}(\psi) = 1$$

Let $H = G/K$. We claim that H is virtually simple. In fact, H is the increasing union of permutation groups, because one easily verifies that the restriction of $\pi_{k(n),k(n+1)}$ to $S_{k(m)\times(k(m)-1)}$ is one-to-one whenever $m \leq n$, and so the projection of $S_{k(m),(k(m)+1)}$ to G/K is one-to-one. Clearly G/K is the increasing union over m of the images of these groups. Thus we have proved

Theorem 39. $\operatorname{aut} X_* = \mathbb{Z}^2 \oplus G$, *where G is a locally finite group factoring onto an infinite virtually simple group.*

It remains to verify our claims about the dynamics of X_*.

Theorem 40. *X_* has positive entropy*

Proof. Consider the two level-2 rectangles a, a' in figure 6.6. Note that they are both of dimensions 9×9. For $n \geq 2$ set $r(n) = 11 \cdot 2^{n-2} - 2$. One can construct a level-n

rectangle all of whose level-2 subrectangles have dimensions 9×9; such a rectangle can be seen inductively to have dimensions $r(n)\times r(n)$. Since such a level-n rectangle has 4^{n-2} level-2 subrectangles, and each of these can be filled in independently as either a or a', the number of level-n rectangles in X_* of dimensions $r(n)\times r(n)$ is at least $2^{4^{n-2}}$. Thus

$$\frac{1}{r(n)^2}\log N(r(n)) \geq \frac{4^{n-2}}{(11\cdot 2^{n-2}-2)^2} \geq \frac{1}{121}$$

and hence $h(X_*)=\lim \frac{1}{n}\log N(n) \geq \frac{1}{200} > 0$. □

Finally, we prove:

Theorem 41. *X_* is strongly mixing.*

Proof. Fix $N>n$ and let $\overline{w}=w_1,\ldots,w_{2^{N-n}}$ and $\overline{h}=h_1,\ldots,h_{2^{N-n}}$ be sequences of integers such that $k^-(n)\leq w_i,h_j\leq k^+(n)$. For each $1\leq i,j\leq 2^{N-n}$ let a_{ij} be a level-n rectangle of dimensions $w_i\times h_j$. From this data one can construct a level-N rectangle $b=b(\overline{w},\overline{h})$ whose level-n subrectangles are divided into 2^{N-n} columns and rows, and the level-n rectangle in row i and column j is a_{ij}. We adopt the convention that the numbering of rows and columns increases from the NW corner to the SE corner of the level-N rectangle. The level-n rectangle in column i and row j of b will be identified as $b_{\langle i,j\rangle}$.

Fix a level-n rectangle a of dimensions $w\times h$. Let N satisfy $2^{N-n}>2N^2$ and let S be the family of level-N rectangles constructed as above from data $\overline{w}=w_1,\ldots,w_{2^{N-n}}$, $\overline{h}=h_1,\ldots,h_{2^{N-n}}$, and a_{ij}, under the constraint that

$$w_i = w \qquad , \qquad h_j = h \qquad , \qquad a_{i,j}=a$$

for $N^2<i,j\leq 2^{N-n}-N^2$, and

$$w_i+w_{i+2^{N-n}-N} = h_i+h_{i+2^{N-n}-N} = k^-(n)+k^+(n)$$

and for $1\leq i,j\leq N^2$ (the rectangles $a_{i,j}$ for $1\leq i,j\leq N^2$ and $N^2+1\leq i,j\leq 2^{N-n}$ may be chosen arbitrarily subject to the restrictions on their dimensions).

The level-n rectangles of any $b\in S$ consist of $(2^{N-n}-2N^2)^2$ copies of a arranged in a grid, which we refer to as the central grid; they are surrounded by N^2 more level-n rectangles on each side. Because the widths of the boundary level-n rectangles on the left and right borders balance, as do the top and bottom ones, we see that all rectangles in S are of the same dimensions $w'\times h'$. The position of the central grid varies between members of S, but is completely determined by the parameters $w_1,\ldots,w_{N^2},h_1,\ldots,h_{N^2}$ from which b is constructed.

For $b \in S$, any two neighboring central copies of a in the central grid differ in position by a vector $u \in \mathbb{Z}^2$ with $\|u\|_\infty \leq k^+(n) + 2N$. The first term bounds the width and height of a (we could use $\max\{w, h\}$ instead) and the second term bounds the number of paths that separate the level-n rectangles; these consist of the boundaries of level-k rectangles for $n+1 \leq k \leq N-1$. Similar considerations show that the $\|\cdot\|_\infty$-distance between a point on the boundary of b and the nearest central copy of a is $\leq N^2k^+(n) + (N^2+1)2N$.

Consider four rectangles $b_1, b_2, b_3, b_4 \in S$. Since they are all of dimensions $w' \times h'$ they can be arranged as siblings in a level $N+1$ rectangle c. Let $U_N \subseteq \mathbb{Z}^2$ be the set of displacements between copies of a's in c as we vary the b_i's. We claim that there is a constant C independent of N (but depending on n) so that $B_{2^N} \setminus B_{CN^3} \subseteq U_N$. Indeed, this follows from the preceding remarks and the fact that in each b the copies of a within the central grid appear syndetically, with gaps bounded by CN, within a square of side $w' - CN^2$, but varying the data of a $b \in S$ (the numbers $w_1, \ldots, w_{N^2}, h_1, \ldots, h_{N^2}$) allows us to displace the central grid by all integer vectors of $\|\cdot\|_\infty$-norm $\leq (k^+(n) - k^-(n))N^2$.

It follows that $\mathbb{Z}^2 \setminus (\cup_N U_N)$ is finite. Since every finite admissible pattern appears in level-n rectangle for some n, this shows that for any open set $\emptyset \neq V \subseteq X_*$, the set $\{u \in \mathbb{Z}^2 : T^u V \cap V \neq \emptyset\}$ is cofinite in $\mathbb{Z}^2$. Now the fact that X_* is transitive implies that X_* is strongly mixing. □

7. Half-space Markov systems

We next describe a half-space Markov system that is not an SFT. Our example will be a 1-step half space Markov systems. The relevant definitions are given in the introduction.

As motivation we first give a construction that is half space Markov in most directions but fails in two exceptional ones. Let X be the system constructed in section 6 and let Y be the factor of X obtained by erasing the colors of the tiles and retaining only the arrows.

Let ℓ be a line not parallel to the axes. We shall assume that the slope of ℓ is > 1 in absolute value; the analysis in the case < 1 is completely analogous, and in the case $= 1$ is simpler. Let $E = \{u : d(u, \ell) \leq 1\}$. We examine the information that $y|_E$ contains about y; often it determines it completely.

In our analysis we shall refer to level-n squares in Y; these are defined as the images of level-n squares in X, which makes sense since the factor map is a 1-block code.

If ℓ intersects the boundary of a level-n square in y, then it does so either at two points or at a corner (although E may spend more time at each boundary, because it is a "thickened" line). If it intersects at two points, necessarily on different sides of the square, then that the arrow directions at the points of intersection allow us to know which side of the arrow is "inside" and which is outside the square. In the case that the intersection occurs at a corner we may treat it as a collapsed case of entering and leaving the square at the same point. Thus we may be read $y|_E$ as a balanced sequence of parentheses (in the case of a corner, matching parentheses are collapsed into a single symbol), and use this to partition the symbols of $y|_E$ into finite sets consisting of symbols belonging to the same square, and deduce which of these squares are contained in which others.

There are two ways that ℓ can intersect a square. The first alternative is that ℓ enters and leaves a square via adjacent sides, in which case we shall identify the sides by the location of their common corner (e.g. NW, etc.). In this case from the symbols at the points of intersection with the square's boundary we can deduce the location of a corner of the square, simply by following the arrows there to the intersection of the line they determine.

The other possibility is that ℓ enter and leave via parallel (non-adjacent) sides. When this happens it is easily deduced from the structure of squares that while the line is inside the square it either passes within distance 1 of some level-0 square, or else it intersects a vertical edge of some sub-square. Since the slope is > 1, it must then intersect the sub-square also in a horizontal edge or corner. Thus there is a sub-square (which may be at level-0) for which we can determine the location of some corner.

Suppose now that ℓ intersects a level-n square Q and that $Q', Q'' \subseteq Q$ are two subsquares of Q that ware not nested in each other and are also intersected by ℓ, at levels k', k'' respectively. By passing to subsquares of Q', Q'' if necessary, we can assume that we know the position of a corner of each. We do not assume that we know n, k', k'', but since ℓ will have to cross every intermediate square between Q, Q' and between Q, Q'', we can see by examining $y|_E$ how far each subsquare is nested inside Q, i.e. we can compute $n - k', n - k''$. One can then show, based on the sequence $k(n)$ of square sizes and on the structure of squares in X, that knowing the position of a corner of each of Q', Q'' and their types (NE,NW etc.) is enough to deduce n, k', k'' and the size and location of Q', Q'' and thus of Q. This argument is similar to that in lemma 25, and we omit it.

It follows from this discussion that if ℓ intersects infinitely many pairs of squares that are not mutually nested, then we can deduce y from $y|_E$. This shows that in the case above, gluing along $y|_E$ is permitted because two points with the same restriction to E coincide.

Now, it is easy to see that if a line passes through a square Q, then it will intersect two non-nested subsquares unless there is some $\omega \in \{NW, NE, SW, SE\}$ so that, starting at $Q_0 = Q$, for each n the line enters Q_ℓ at points that are closer to the ω-corner of Q_n than any other and pass exclusively through the ω-subsquare Q_{n+1} of Q_n. For this to occur in every square that ℓ passes through, the hierarchy of squares must collapse to one that is linearly ordered by inclusion. This can occur only for points y that are images of points x_ω for $\omega \in \{NW, NE, SW, SE\}^{\mathbb{N}}$ that are eventually constant. The picture that we see in $y|_E$ is then similar to the intersection of a line, going from SW to NE, with figure 6.3(b), and the same pattern $y|_E$ occurs when y is exceptional of this form. One may verify that in all these cases, though, gluing produces an admissible point.

One can similarly verify that in the case of other exceptional points and in the case of a line of some other orientation intersecting (b) of figure 6.3, gluing along E produces an admissible point.

This establishes the half space Markov property for lines not parallel to the axes. However, this is not the case for the two remaining lines. In fact, one can obtain identical patterns on vertical lines by choosing a non-exceptional point $x \in X$ and two vertical lines which pass through blank tiles. The width-one strips will then coincide (to get wider ones we must take care that the position of the blanks in is the same for the first few levels of the nested structure). One cannot glue along these lines, since it would produce closed paths which are not squares.

To correct this, we modify X as follows. We replace the arrow tiles with two types of arrow tiles, type A and type B. The rules governing these tiles, besides those inherited from X, are:

- All corner tiles are type A. If the corner is in a level-1 square (i.e. it is diagonally adjacent to a blank) it leads into a type-B arrow; otherwise it leads into a type-A arrow.
- A type A arrow leads into another type A arrow unless there is a corner tile adjacent to it on its right (i.e. at $90°$ from the direction of the arrow). In this case, the next arrow along the path is a type B arrow.

- A type B arrow leads into another type B arrow, or into a corner tile (which is type A).

Let us denote by $\widehat{X}$ the SFT thus defined. Proceeding clockwise around the boundary of a level-n square, beginning at any corner, one sees alternating type A and type B tiles, with one alternation per side. The transition from type B to type A occurs on entering the corner tiles, and from type A to type B the transition occurs at e the middle of each side (in level-1 squares this occurs upon leaving the corner tile). The rules for exceptional points are derived from these.

Let $\widehat{Y}$ be the sofic shift obtained by forgetting all colors, but retaining arrows and types. Each level-n square can be extended uniquely to a colored one, so a non-exceptional point in Y can be extended uniquely to one in $\widehat{Y}$. Thus, as before, for lines not parallel to the axes it is easy to check that $y|_\ell$ determines y when $y \in \widehat{Y}$ is non-exceptional. One can verify directly that for exceptional x the Markov property is satisfied, even though $y|_\ell$ may not determine y (this occurs, for example, with when ℓ is the line of slope -1 and y is a point of the type (b) in figure 6.3.

For half-spaces determined by vertical lines, we note that the problem that occurred in Y does not occur, because when a vertical line intersects a level-n square, the type of tiles indicates which side of the square it is on, and one can check that this provides enough information to reconstruct the point, and similarly for horizontal lines.

It remains to demonstrate that $\widehat{Y}$ is not an SFT. Notice that in $\widehat{Y}$ all the level-n squares are the same pattern (in X they differed only by the color of their boundaries but in $\widehat{Y}$ this information has now been pushed "farther out" to the boundary of level-$(n+1)$ squares). If $\widehat{Y}$ were an SFT with interactions of length $\leq n$, then the periodic point obtained by tiling the plane regularly with level-n squares would be admissible for $\widehat{Y}$, because it is locally admissible; indeed, the patterns along the "seams" all appear in the pattern at the center of a level-$(n+1)$ square. Thus, since no point in $\widehat{Y}$ is both periodic and contains a square (of any level), $\widehat{Y}$ cannot be an SFT.

8. Open problems

In this section we discuss some open questions about automorphism groups of higher dimensional SFTs.

A fundamental problem which we do not resolve is whether for each d and $i \neq j$ one can embed the automorphism groups of d-dimensional full shifts on i and j symbols in each other.

Another basic question is whether the automorphism group of the full shift on two symbols in dimension $d > 1$ can be embedded in any of the automorphism groups of a lower-dimensional SFT.

Other questions about $\operatorname{aut} X$ arise from the recent connections observed between SFTs and recursion theory [5, 11, 4]. In [2] the authors noted that for 1-dimensional SFTs one can always determine if an automorphism, given by its block code representation, is the identity. This is because we may apply the block code to the set of sufficiently long admissible words and observe whether it acts as the identity. It follows that any finitely presented subgroup of $\operatorname{aut} X$ has decidable word problem. In higher dimensions this argument fails, for one cannot in general identify all admissible finite patterns of a given shape for a given SFT. It is true that if φ is the identity one can demonstrate this, since by compactness it must act as the identity near the origin for all sufficiently large locally admissible patterns; but it seems that the converse is false. One is therefore led to the question: can every finitely presented group be embedded in $\operatorname{aut} X$ for some SFT X? More generally, suppose that G is a countable group such that the product function $R : G \times G \to G$ is recursive. Can G be embedded in, or realized as, the automorphism group of an SFT modulo the shift?

References

[1] Paul Balister, Béla Bollobás, and Anthony Quas. Entropy along convex shapes, random tilings and shifts of finite type. *Illinois J. Math.*, 46(3):781–795, 2002.

[2] Mike Boyle, Douglas Lind, and Daniel Rudolph. The automorphism group of a shift of finite type. *Trans. Amer. Math. Soc.*, 306(1):71–114, 1988.

[3] G. A. Hedlund. Endormorphisms and automorphisms of the shift dynamical system. *Math. Systems Theory*, 3:320–375, 1969.

[4] Michael Hochman. On the dynamics and recursive properties of multidimensional symbolic systems. *preprint*, 2007.

[5] Michael Hochman and Tom Meyerovitch. A characterization of the entropies of multidimensional shifts of finite type. *Annals of Mathematics*, to appear, 2007.

[6] K. H. Kim and F. W. Roush. On the automorphism groups of subshifts. *Pure Math. Appl. Ser. B*, 1(4):203–230 (1991), 1990.

[7] Bruce P. Kitchens. *Symbolic dynamics*. Universitext. Springer-Verlag, Berlin, 1998. One-sided, two-sided and countable state Markov shifts.

[8] Douglas Lind and Brian Marcus. *An introduction to symbolic dynamics and coding*. Cambridge University Press, Cambridge, 1995.

[9] Raphael M. Robinson. Undecidability and nonperiodicity for tilings of the plane. *Invent. Math.*, 12:177–209, 1971.

[10] J. Patrick Ryan. The shift and commutativity. *Math. Systems Theory*, 6:82–85, 1972.

[11] Steve Simpson. Medvedev degrees of 2-dimensional subshifts of finite type. *Ergodic Theorey Dynamical Systems*, to appear, 2007.

[12] Thomas Ward. Automorphisms of $\mathbf{Z}^d$-subshifts of finite type. *Indag. Math. (N.S.)*, 5(4):495–504, 1994.

Current address: Fine Hall, Washington Rd., Princeton, NJ 08544

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