



THE AUTOMORPHISM GROUP OF A SHIFT OF LINEAR GROWTH: BEYOND TRANSITIVITY

VAN CYR¹ and BRYNA KRA²

¹ Bucknell University, Lewisburg, PA 17837, USA;
email: van.cyr@bucknell.edu

² Northwestern University, Evanston, IL 60208, USA;
email: kra@math.northwestern.edu

Received 3 November 2014; accepted 12 February 2015

Abstract

For a finite alphabet \mathcal{A} and shift $X \subseteq \mathcal{A}^{\mathbb{Z}}$ whose factor complexity function grows at most linearly, we study the algebraic properties of the automorphism group $\text{Aut}(X)$. For such systems, we show that every finitely generated subgroup of $\text{Aut}(X)$ is virtually \mathbb{Z}^d , in contrast to the behavior when the complexity function grows more quickly. With additional dynamical assumptions we show more: if X is transitive, then $\text{Aut}(X)$ is virtually \mathbb{Z} ; if X has dense aperiodic points, then $\text{Aut}(X)$ is virtually \mathbb{Z}^d . We also classify all finite groups that arise as the automorphism group of a shift.

2010 Mathematics Subject Classification: 37B50 (primary); 68R15, 37B10 (secondary)

1. Introduction

Given a finite alphabet \mathcal{A} , a shift system (X, σ) is a closed set $X \subseteq \mathcal{A}^{\mathbb{Z}}$ that is invariant under the left shift $\sigma : \mathcal{A}^{\mathbb{Z}} \rightarrow \mathcal{A}^{\mathbb{Z}}$, and its automorphism group $\text{Aut}(X)$ is the group of homeomorphisms of X that commute with σ (these notions are made precise in Section 2). For general shift systems, while $\text{Aut}(X)$ is countable, it can be quite complicated: for the full shift [8] or for mixing shifts of finite type [3], $\text{Aut}(X)$ is not finitely generated and is not amenable (see also [2, 7, 9, 10, 16]). The assumption of topological mixing can be used to construct a rich collection of subgroups of the automorphism group. For example, the automorphism group of such a shift contains isomorphic copies of all finite groups, the direct sum of countably many copies of \mathbb{Z} , and the free group on two generators. In these examples, the topological entropy is positive, and the complexity function $P_X(n)$,

© The Author(s) 2015. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/3.0/>), which permits unrestricted re-use, distribution, and reproduction in any medium, provided the original work is properly cited.

which counts the number of nonempty cylinder sets of length n taken over all elements $x \in X$, grows quickly.

When the complexity function of a shift system grows slowly, the automorphism group is often much simpler, and the main goal of this paper is to study the algebraic properties of $\text{Aut}(X)$ in this setting. In contrast to mixing shifts, we study general shifts of low complexity, without an assumption of minimality or transitivity. We show that the automorphism group of any shift of low complexity is amenable, yet its behavior can still be quite complicated.

As $P_X(n)$ is nondecreasing, boundedness is the slowest possible growth property that $P_X(n)$ can have. As expected, this case is simple: the Morse–Hedlund theorem [11] implies that if there exists $n \in \mathbb{N}$ such that $P_X(n) \leq n$, then X is comprised entirely of periodic points. Thus $\text{Aut}(X)$ is a finite group (and we classify all finite groups that arise in this way in Section 7). It follows that if (X, σ) is a shift for which $P_X(n)/n \xrightarrow{n \rightarrow \infty} 0$, then $|\text{Aut}(X)| < \infty$.

It is thus natural to study shifts for which $P_X(n) > n$ for all $n \in \mathbb{N}$. The smallest nontrivial growth rate that such a system can have is linear, by which we mean

$$0 < \limsup_{n \rightarrow \infty} \frac{P_X(n)}{n} < \infty.$$

The class of shifts of linear growth arises naturally in several contexts. For example, this class includes the Sturmian shifts, and more generally the natural coding of any minimal interval exchange transformation. In previous work [5], we studied the algebraic properties of $\text{Aut}(X)$ for transitive shifts of subquadratic growth, and showed that $\text{Aut}(X)/\langle \sigma \rangle$ is a periodic group, where $\langle \sigma \rangle$ denotes the subgroup of $\text{Aut}(X)$ generated by σ . In particular, this holds for transitive shifts of linear growth. Periodic groups, however, can be quite complicated: for example, a periodic group need not be finitely generated, and there are finitely generated nonamenable periodic groups. In this paper, we study $\text{Aut}(X)$ for general (not necessarily transitive) shifts of linear growth. In the transitive case, we prove a stronger result than is implied by [5], showing that $\text{Aut}(X)/\langle \sigma \rangle$ is finite. However, the main novelty of this work is that our techniques remain valid even without the assumption of transitivity.

Depending on dynamical assumptions on the system, shift systems with linear growth exhibit different behavior. Our most general result is the following.

THEOREM 1.1. *Suppose that (X, σ) is a shift system for which there exists $k \in \mathbb{N}$ such that*

$$\limsup_{n \rightarrow \infty} P_X(n)/n < k.$$

Then every finitely generated subgroup of $\text{Aut}(X)$ is virtually \mathbb{Z}^d for some $d < k$.

Let $[\sigma]$ denote the full group of a shift (X, σ) (see Section 2.4 for the definition). With the additional assumption that (X, σ) has a dense set of aperiodic points, we have the following.

THEOREM 1.2. *Suppose that (X, σ) is a shift system for which there exists $k \in \mathbb{N}$ such that*

$$\limsup_{n \rightarrow \infty} P_X(n)/n < k.$$

If X has a dense set of aperiodic points, then $\text{Aut}(X) \cap [\sigma] \cong \mathbb{Z}^d$ for some $d < k$, and $\text{Aut}(X)/(\text{Aut}(X) \cap [\sigma])$ is finite. In particular, $\text{Aut}(X)$ is virtually \mathbb{Z}^d .

With the additional assumption that (X, σ) is topologically transitive, meaning that there exists a point whose orbit is dense in X , we show the following.

THEOREM 1.3. *Suppose that (X, σ) is a transitive shift system for which*

$$0 < \limsup_{n \rightarrow \infty} P_X(n)/n < \infty.$$

Then $\text{Aut}(X)/\langle \sigma \rangle$ is finite. In particular, $\text{Aut}(X)$ is virtually \mathbb{Z} .

For minimal shifts, meaning shifts such that every point has dense orbit, we show the following (note the growth condition on the complexity only assumes \liminf instead of \limsup).

THEOREM 1.4. *Suppose that (X, σ) is a minimal shift for which there exists $k \in \mathbb{N}$ satisfying*

$$\liminf_{n \rightarrow \infty} P_X(n)/n < k.$$

Then $\text{Aut}(X)/\langle \sigma \rangle$ is finite, and $|\text{Aut}(X)/\langle \sigma \rangle| < k$.

For periodic minimal shifts, it is easy to see that $\text{Aut}(X) \cong \mathbb{Z}/n\mathbb{Z}$, where n is the minimal period. Salo and Törmä [14] asked if the automorphism group of any linearly recurrent shift is virtually \mathbb{Z} . Linearly recurrent shifts are minimal, and the factor complexity function grows at most linearly, and so Theorem 1.4 gives an affirmative answer to their question.

Roughly speaking, the proof of Theorem 1.1 splits into two parts. We start by studying shifts with a dense set of aperiodic points in Section 3.2, showing that the automorphism group is locally a group of polynomial growth, with the polynomial growth rate depending on the linear complexity assumption on the shift. We sharpen this result to understand transitive shifts of linear growth, leading to the

proof of Theorem 1.3 in Section 3.3. We then combine this with information on existence of aperiodic points, completing the proof of Theorem 1.1 in Section 4. The proof of Theorem 1.4 in Section 5 proceeds in a different manner, relying on a version of a lemma of Boshernitzan used to bound the number of ergodic probability measures on a shift with linear growth, which we use to bound the number of words in the language of the system that have multiple extensions.

For some of these results, we are able to give examples showing that they are sharp. These examples are included in Section 6.

While writing up these results, we became aware of related work by Donoso et al. [6]. While some of the results obtained are the same, the methods are different, and each method leads to new open directions. In particular, their work includes examples of automorphism groups for certain shifts of polynomial growth. The general classification of automorphism groups for shifts of polynomial growth remains open.

2. Background and notation

2.1. Shift systems. We assume throughout that \mathcal{A} is a fixed finite set endowed with the discrete topology. If $x \in \mathcal{A}^{\mathbb{Z}}$, we denote the value of x at $n \in \mathbb{Z}$ by $x(n)$. The metric $d(x, y) := 2^{-\inf\{|n|:x(n) \neq y(n)\}}$ generates the product topology on $\mathcal{A}^{\mathbb{Z}}$, and endowed with this metric, $\mathcal{A}^{\mathbb{Z}}$ is a compact metric space; henceforth, we assume this metric structure on $\mathcal{A}^{\mathbb{Z}}$.

The *left shift* $\sigma : \mathcal{A}^{\mathbb{Z}} \rightarrow \mathcal{A}^{\mathbb{Z}}$ is the map defined by $(\sigma x)(n) := x(n + 1)$ and is a homeomorphism from $\mathcal{A}^{\mathbb{Z}}$ to itself. If $X \subseteq \mathcal{A}^{\mathbb{Z}}$ is a closed σ -invariant subset, then the pair (X, σ) is called a *subshift of $\mathcal{A}^{\mathbb{Z}}$* , or just a *shift of $\mathcal{A}^{\mathbb{Z}}$* . If the alphabet \mathcal{A} is clear from the context, we refer to (X, σ) as just a *shift*.

The set

$$\mathcal{O}(x) := \{\sigma^n x : n \in \mathbb{N}\}$$

is the *orbit of x* , and we use $\overline{\mathcal{O}(x)}$ to denote its closure. The shift (X, σ) is *transitive* if there exists some $x \in X$ such that $\overline{\mathcal{O}(x)} = X$, and it is *minimal* if $\overline{\mathcal{O}(x)} = X$ for all $x \in X$. A point $x \in X$ is *periodic* if there exists some $n \in \mathbb{N}$ such that $\sigma^n x = x$, and otherwise it is said to be *aperiodic*.

2.2. Complexity of shifts. For a shift (X, σ) and $w = (a_{-m+1}, \dots, a_{-1}, a_0, a_1, \dots, a_{m-1}) \in \mathcal{A}^{2m+1}$, the *central cylinder set $[w]_0$ determined by w* is defined to be

$$[w]_0 := \{x \in X : x(n) = a_n \text{ for all } -m < n < m\}.$$

The collection of central cylinder sets forms a basis for the topology of X . If $w = (a_0, \dots, a_{m-1}) \in \mathcal{A}^m$, then the *one sided cylinder set $[w]_0^+$ determined by w*

is given by

$$[w]_0^+ := \{x \in X : x(n) = a_n \text{ for all } 0 \leq n < m\}.$$

For $m \in \mathbb{N}$, define the set of words $\mathcal{L}_m(X)$ of length m in X by

$$\mathcal{L}_m(X) := \{w \in \mathcal{A}^m : [w]_0^+ \neq \emptyset\},$$

and define the language $\mathcal{L}(X)$ of X to be $\mathcal{L}(X) := \bigcup_{m=1}^{\infty} \mathcal{L}_m(X)$. For $w \in \mathcal{L}(X)$, we denote the length of w by $|w|$. A word in $x \in X$ is also referred to as a *factor* of x .

A measure of the complexity of X is the (*factor*) *complexity function* $P_X : X \rightarrow \mathbb{N}$, which counts the number of words of length n in the language of X :

$$P_X(n) := |\mathcal{L}_n(X)|.$$

If $P_x(n)$ is the complexity function of a fixed $x \in X$, meaning that it is the number of configurations in a block of size n in x , then $P_X(n) \geq \sup_{x \in X} P_x(n)$, with equality holding for all n when X is a transitive shift.

2.3. The automorphism group of a shift. Let $\text{Hom}(X)$ denote the group of homeomorphisms from X to itself. If $h_1, \dots, h_n \in \text{Hom}(X)$, then $\langle h_1, \dots, h_n \rangle$ denotes the subgroup of $\text{Hom}(X)$ generated by h_1, \dots, h_n . Thus the shift $\sigma \in \text{Hom}(X)$ and its centralizer in $\text{Hom}(X)$ is called the *automorphism group* of (X, σ) . We denote the automorphism group of (X, σ) by $\text{Aut}(X)$, and endow it with the discrete topology.

A map $\varphi : X \rightarrow X$ is a *sliding block code* if there exists $R \in \mathbb{N}$ such that for any $w \in \mathcal{L}_{2R+1}(X)$ and any $x, y \in [w]_0$, we have $(\varphi x)(0) = (\varphi y)(0)$. Any number $R \in \mathbb{N} \cup \{0\}$ for which this property holds is called a *range* for φ . The *minimal range* of φ is its smallest range.

If $\varphi : X \rightarrow X$ is a sliding block code of range R , there is a natural map (which, by abuse of notation, we also denote by φ) taking $\bigcup_{m=2R+1}^{\infty} \mathcal{L}_m(X)$ to $\mathcal{L}(X)$. To define this extension of φ , let $m > 2R$, and let $w = (a_0, \dots, a_{m-1}) \in \mathcal{A}^m$. For $0 \leq i < m - 2R$, choose $x_i \in [(a_i, \dots, a_{i+2R})]_0$, and define

$$\varphi(w) := ((\varphi x_0)(0), (\varphi x_1)(0), \dots, (\varphi x_{m-2R-1})(0)).$$

Therefore if w is a word of length at least $2R + 1$, then $\varphi(w)$ is a word of length $|w| - 2R$.

The elements of $\text{Aut}(X)$ have a concrete characterization.

THEOREM 2.1 (Curtis–Hedlund–Lyndon Theorem [8]). *If (X, σ) is a shift, then any element of $\text{Aut}(X)$ is a sliding block code.*

For $R \in \mathbb{N} \cup \{0\}$, we let $\text{Aut}_R(X) \subseteq \text{Aut}(X)$ denote the automorphisms of (X, σ) for which R is a (not necessarily minimal) range. Thus $\text{Aut}(X) = \bigcup_{R=0}^{\infty} \text{Aut}_R(X)$. We observe that if $\varphi_1 \in \text{Aut}_{R_1}(X)$ and $\varphi_2 \in \text{Aut}_{R_2}(X)$, then $\varphi_1 \circ \varphi_2 \in \text{Aut}_{R_1+R_2}(X)$.

In general, the automorphism group of a shift can be complicated, but Theorem 2.1 implies that $\text{Aut}(X)$ is always countable.

2.4. Automorphisms and the full group. The *full group* $[\sigma]$ of a shift (X, σ) is the subgroup of $\text{Hom}(X)$ comprised of the orbit-preserving homeomorphisms:

$$[\sigma] := \{\psi \in \text{Hom}(X) : \psi(x) \in \mathcal{O}(x) \text{ for all } x \in X\}.$$

Thus if $\psi \in [\sigma]$, then there is a function $k_\psi : X \rightarrow \mathbb{Z}$ such that $\psi(x) = \sigma^{k_\psi(x)}(x)$ for all $x \in X$.

It follows from the definitions that the group $\text{Aut}(X) \cap [\sigma]$ is the centralizer of σ in $[\sigma]$. We note two basic facts about $\text{Aut}(X) \cap [\sigma]$ which we will need in order to study $\text{Aut}(X)/(\text{Aut}(X) \cap [\sigma])$ in Section 3.7.

LEMMA 2.2. *If (X, σ) is a shift, then $\text{Aut}(X) \cap [\sigma]$ is normal in $\text{Aut}(X)$.*

Proof. Let $\varphi \in \text{Aut}(X)$, and suppose that $\psi \in \text{Aut}(X) \cap [\sigma]$. Let $k_\varphi : X \rightarrow \mathbb{Z}$ be a function such that $\varphi(x) = \sigma^{k_\varphi(x)}(x)$ for all $x \in X$. Fix $x \in X$, and observe that, since φ and σ commute,

$$\varphi \circ \psi \circ \varphi^{-1}(x) = \varphi \circ \sigma^{k_\psi(\varphi^{-1}(x))} \circ \varphi^{-1}(x) = \sigma^{k_\psi(\varphi^{-1}(x))}(\varphi(x)).$$

As this holds for any $x \in X$, it follows that $\varphi \circ \psi \circ \varphi^{-1} \in \text{Aut}(X) \cap [\sigma]$. Since $\varphi \in \text{Aut}(X)$ and $\psi \in \text{Aut}(X) \cap [\sigma]$ are arbitrary, we have

$$\text{Aut}(X) \cap [\sigma] = \varphi \cdot (\text{Aut}(X) \cap [\sigma]) \cdot \varphi^{-1}$$

for all $\varphi \in \text{Aut}(X)$. So $\text{Aut}(X) \cap [\sigma]$ is normal in $\text{Aut}(X)$. \square

LEMMA 2.3. *If (X, σ) is a shift, then $\text{Aut}(X) \cap [\sigma]$ is abelian.*

Proof. Suppose that $\varphi_1, \varphi_2 \in \text{Aut}(X) \cap [\sigma]$. For $i = 1, 2$, let $k_{\varphi_i} : X \rightarrow \mathbb{Z}$ be functions such that $\varphi_i(x) = \sigma^{k_{\varphi_i}(x)}(x)$ for all $x \in X$. For any $x \in X$,

$$\begin{aligned} \varphi_1 \circ \varphi_2(x) &= \varphi_1 \circ \sigma^{k_{\varphi_2}(x)}(x) = \sigma^{k_{\varphi_2}(x)} \circ \varphi_1(x) \\ &= \sigma^{k_{\varphi_2}(x)} \circ \sigma^{k_{\varphi_1}(x)}(x) = \sigma^{k_{\varphi_1}(x)} \circ \sigma^{k_{\varphi_2}(x)}(x) \\ &= \sigma^{k_{\varphi_1}(x)} \circ \varphi_2(x) = \varphi_2 \circ \sigma^{k_{\varphi_1}(x)}(x) \\ &= \varphi_2 \circ \varphi_1(x). \end{aligned}$$

Therefore $\varphi_1 \circ \varphi_2 = \varphi_2 \circ \varphi_1$. \square

2.5. Summary of group theoretic terminology. For convenience, we summarize the algebraic properties that we prove $\text{Aut}(X)$ may have. We say that a group G is *locally P* if every finitely generated subgroup of G has property P . The group G is *virtually H* if G contains H as a subgroup of finite index. The group G is *K-by-L* if there exists a normal subgroup H of G which is K and such that the quotient G/H is L .

3. Shifts of linear growth with a dense set of aperiodic points

3.1. Cassaigne's characterization of linear growth. Linear growth can be characterized in terms of the (first) difference of the complexity function:

THEOREM 3.1 [4]. *A transitive shift (X, σ) satisfies $p_X(n) = O(n)$ if and only if the difference function $p_X(n+1) - p_X(n)$ is bounded.*

DEFINITION 3.2. Let $w = (a_0, \dots, a_{|w|-1}) \in \mathcal{L}_{|w|}(X)$. For fixed $m \in \mathbb{N}$, we say that w *extends uniquely m times to the right* if there is exactly one word $\tilde{w} = (b_0, \dots, b_{|w|+m-1}) \in \mathcal{L}_{|w|+m}(X)$ such that $a_i = b_i$ for all $0 \leq i < |w|$.

COROLLARY 3.3. *Assume that (X, σ) can be written as the union of finitely many transitive shifts, and satisfies $p_X(n) = O(n)$. Then for any $m, n \in \mathbb{N}$, the number of words of length n that do not extend uniquely m times to the right is at most Bm , where $B = \max_{n \in \mathbb{N}}(p_X(n+1) - p_X(n))$.*

Note that it follows from Cassaigne's theorem that B is finite, as every word that appears in $\mathcal{L}(X)$ also appears in the language of one of the transitive subshifts.

Proof. For any $N \in \mathbb{N}$, the quantity $p_X(N+1) - p_X(N)$ is an upper bound on the number of words of length N that do not extend uniquely to the right. For any word w of length n which does not extend uniquely m times to the right, there exists $0 \leq k < m$ such that w extends uniquely k times to the right, but not $k+1$ times. For fixed k , the number of words for which this is the case is at most the number of words of length $n+k$ that do not extend uniquely to the right. So the number of words of length n that fail to extend uniquely m times to the right is at most

$$\sum_{k=1}^m (p_X(n+k) - p_X(n+k-1)) \leq Bm. \quad \square$$

3.2. Assuming a dense set of aperiodic points. We start by considering shifts with a dense set of aperiodic points. This assumption holds in particular when the

shift has no isolated points: if X has no isolated points, then, for any fixed period, the set of periodic points with that period has empty interior. Then the Baire category theorem implies that the set of all periodic points has empty interior. In particular, the set of aperiodic points is dense. The two assumptions are equivalent if the set of aperiodic points is nonempty.

LEMMA 3.4. *Suppose that (X, σ) is a shift with a dense set of aperiodic points, and that there exists $k \in \mathbb{N}$ such that*

$$\limsup_{n \rightarrow \infty} \frac{P_X(n)}{n} < k.$$

Then there exist $x_1, \dots, x_{k-1} \in X$ such that

$$X = \overline{\mathcal{O}}(x_1) \cup \overline{\mathcal{O}}(x_2) \cup \dots \cup \overline{\mathcal{O}}(x_{k-1}).$$

Proof. Suppose not, and let $x_1 \in X$. Since $\overline{\mathcal{O}}(x_1) \neq X$, there is a word $w_1 \in \mathcal{L}(X)$ such that $[w_1]_0^+ \cap \overline{\mathcal{O}}(x_1) = \emptyset$. Choose $x_2 \in X$ with $x_2 \in [w_1]_0^+$. Let $i < k$, and suppose that we have constructed $x_1, \dots, x_i \in X$ and $w_1, \dots, w_{i-1} \in \mathcal{L}(X)$ such that $[w_{j_i}]_0 \cap \overline{\mathcal{O}}(x_{j_2}) = \emptyset$ whenever $j_2 \leq j_1$. Since $\overline{\mathcal{O}}(x_1) \cup \dots \cup \overline{\mathcal{O}}(x_i) \neq X$, there is a word $w_i \in \mathcal{L}(X)$ such that $[w_i]_0^+ \cap \overline{\mathcal{O}}(x_1) \cup \dots \cup \overline{\mathcal{O}}(x_i) = \emptyset$. Let $x_{i+1} \in [w_i]_0^+$, and we continue this construction until $i = k$.

Let $N > \max_{1 \leq i < k} |w_i|$ be a fixed large integer (to be specified later). Since x_1 is aperiodic, there are at least $N + 1$ distinct factors of length N in $\overline{\mathcal{O}}(x_1)$. Therefore there are at least $N + 1$ distinct factors of length N in X which do not contain the words w_1, \dots, w_{k-1} . We claim that, for $1 \leq i < k$, there are at least $N - |w_i|$ distinct factors in $\overline{\mathcal{O}}(x_{i+1})$ which contain the word w_i but do not contain any of the words $w_{i+1}, w_{i+2}, \dots, w_{k-1}$. Assuming this claim, then for any sufficiently large N we have $p_X(N) \geq kN - \sum_{i=1}^k |w_i|$, a contradiction of the complexity assumption.

We are left with proving the claim. Let $1 \leq i < k$ be fixed. By construction, the word w_i appears in $\mathcal{O}(x_{i+1})$ but $[w_j]_0^+ \cap \overline{\mathcal{O}}(x_{i+1}) = \emptyset$ for any $j > i$. If w_i appears syndetically in x_{i+1} , then so long as N is sufficiently large, every factor of x_{i+1} of length N contains the word w_i . In this case, since x_{i+1} is aperiodic, there are at least $N + 1$ distinct factors in $\overline{\mathcal{O}}(x_{i+1})$ which contain w_i but not w_j for any $j > i$. Otherwise w_i does not appear syndetically in x_{i+1} , and so there are arbitrarily long factors in x_{i+1} which do not contain w_i . Since w_i appears at least once in x_{i+1} , it follows that there are arbitrarily long words which appear in x_{i+1} which contain exactly one occurrence of w_i , and we can assume that w_i occurs as either the rightmost or leftmost factor. Without loss of generality, we assume that there exists a word w of length N which contains w_i as its rightmost factor and

has no other occurrences of w_i . Choose $j \in \mathbb{Z}$ such that

$$w = (x_{i+1}(j), x_{i+1}(j + 1), \dots, x_{i+1}(j + |w| - 1)).$$

By construction, if $0 \leq s < |w| - |w_i|$, then the word

$$w^{(s)} := (x_{i+1}(j + s), x_{i+1}(j + s + 1), \dots, x_{i+1}(j + s + |w| - 1))$$

is a word of length N for which the smallest $t \in \{0, \dots, |w| - |w_i|\}$ such that

$$w_i = (x_{i+1}(j + t), x_{i+1}(j + t + 1), \dots, x_{i+1}(j + t + |w_i| - 1))$$

is $t = |w| - |w_i| - s$. Therefore, the words $w^{(s)}$ are pairwise distinct, and each contains w_i as a factor. By construction, they do not contain w_j for any $j > i$, thus establishing the claim. \square

PROPOSITION 3.5. *Suppose that (X, σ) is a shift with a dense set of aperiodic points, and that there exists $k \in \mathbb{N}$ such that*

$$\limsup_{n \rightarrow \infty} \frac{P_X(n)}{n} < k.$$

Then $\text{Aut}(X)$ is locally a group of polynomial growth with polynomial growth rate at most $k - 1$. Moreover, if $q \in \mathbb{N}$ is the smallest cardinality of a set $x_1, \dots, x_q \in X$ such that $\mathcal{O}(x_1) \cup \mathcal{O}(x_2) \cup \dots \cup \mathcal{O}(x_q)$ is dense in X , then the polynomial growth rate of any finitely generated subgroup of $\text{Aut}(X)$ is at most q .

In Section 6.1, we give an example showing that the growth rate given in this proposition is optimal.

Proof. By Lemma 3.4, there exist $y_1, \dots, y_{k-1} \in X$ such that the union of the orbits $\mathcal{O}(y_1) \cup \mathcal{O}(y_2) \cup \dots \cup \mathcal{O}(y_{k-1})$ is dense in X . Let $x_1, \dots, x_q \in X$ be a set of minimum cardinality for which $\mathcal{O}(x_1) \cup \mathcal{O}(x_2) \cup \dots \cup \mathcal{O}(x_q)$ is dense.

For $1 \leq i \leq q$, define the constant

$$C_i := \inf\{|w| : w \text{ is a factor of } x_i \text{ and } [w]_0^+ \text{ contains precisely one element}\},$$

and define $C_i := 0$ if no such factor exists. Define

$$C := \max_{1 \leq i \leq q} C_i. \tag{1}$$

Fix $R \in \mathbb{N}$. For $i = 1, \dots, q$, let \tilde{w}_i be a factor of x_i such that the following hold.

- (a) $|\tilde{w}_i| \geq 3R + 1$.
- (b) For all $u \in \mathcal{L}_{2R+1}(X)$, there exists i such that u is a factor of \tilde{w}_i .

Note that (b) is possible since $\mathcal{O}(x_1) \cup \mathcal{O}(x_2) \cup \dots \cup \mathcal{O}(x_q)$ is dense.

Without loss of generality, we can assume that there exists $M_1 \geq 0$ such that $[\tilde{w}_i]_0^+$ contains precisely one element for all $i \leq M_1$ and contains at least two elements for all $i > M_1$ (otherwise, reorder x_1, \dots, x_{k-1}). For each $i > M_1$, either there exists $a \geq 0$ such that \tilde{w}_i extends uniquely to the right a times but not $a + 1$ times, or there exists $a \geq 0$ such that \tilde{w}_i extends uniquely to the left a times but not $a + 1$ times. Again, reordering if necessary, we can assume that there exists $M_2 \geq M_1$ such that the former occurs for all $M_1 < i \leq M_2$ and the latter occurs when $i > M_2$. For $i = 1, \dots, q$, we define words w_1, \dots, w_q as follows.

- (i) For $i = 1, \dots, M_1$, the set $[\tilde{w}_i]_0^+$ contains precisely one element. This must be a shift of x_i , and without loss of generality we can assume that it is x_i itself. In this case, we define u_i to be the shortest factor of x_i with the property that $[u_i]_0^+$ contains precisely one element, and define w_i to be the (unique) extension $2R + 2$ times both to the right and to the left of u_i . Observe that if $\varphi, \varphi^{-1} \in \text{Aut}_R(X)$, then $\varphi^{-1}(\varphi(w_i)) = u_i$. Since φ^{-1} is injective and sends every element of $[\varphi(w_i)]_0^+$ to the one point set $[u_i]_0^+$, it follows that $[\varphi(w_i)]_0^+$ contains precisely one element, and that the word $\varphi(w_i)$ uniquely determines the word $\varphi(\tilde{w}_i)$. Moreover, $|u_i| \leq C$, where C is the constant in (1), and so $|w_i| \leq C + 4R + 4$.
- (ii) For $i = M_1 + 1, \dots, M_2$, there exists $a_i \geq 0$ such that \tilde{w}_i extends uniquely to the right a_i times but not $a_i + 1$ times. Define w_i to be the (unique) word of length $|\tilde{w}_i| + a_i$ which has \tilde{w}_i as its leftmost factor. By choice of the ordering, w_i does not extend uniquely to its right.
- (iii) For $i = M_2 + 1, \dots, q$, there exists $a_i \geq 0$ such that \tilde{w}_i extends uniquely to the left a_i times but not $a_i + 1$ times. Define w_i to be the (unique) word of length $|\tilde{w}_i| + a_i$ which has \tilde{w}_i as its rightmost factor. By choice of the ordering, w_i does not extend uniquely to its left.

For $\varphi \in \text{Aut}_R(X)$, we have that $\varphi(w_i)$ determines the word $\varphi(\tilde{w}_i)$, and so, by property (b) in the definition of \tilde{w}_i , the block code determines what φ does to every word in $\mathcal{L}_{2R+1}(X)$. Thus the map $\Phi : \text{Aut}_R(X) \rightarrow \mathcal{L}_{|w_1|-2R}(X) \times \mathcal{L}_{|w_2|-2R}(X) \times \dots \times \mathcal{L}_{|w_q|-2R}(X)$ defined by

$$\Phi(\varphi) = (\varphi(w_1), \varphi(w_2), \dots, \varphi(w_q))$$

is injective. We claim that for $1 \leq i \leq q$, we have

$$|\{\varphi(w_i) : \varphi, \varphi^{-1} \in \text{Aut}_R(X)\}| \leq Bk(C + 4)(R + 1), \tag{2}$$

where B is the constant appearing in Corollary 3.3 and C is the constant in (1).

Before proving the claim, we show how to deduce the proposition from this estimate. It follows from (2) that $|\{\Phi(\varphi) : \varphi, \varphi^{-1} \in \text{Aut}_R(X)\}| \leq (Bk(C+4))^q (R+1)^q$. Since Φ is injective, it follows that we have the bound

$$|\{\varphi \in \text{Aut}_R(X) : \varphi^{-1} \in \text{Aut}_R(X)\}| \leq (Bk(C+4))^q (R+1)^q. \quad (3)$$

Given $\varphi_1, \dots, \varphi_m \in \text{Aut}(X)$, choose $R \in \mathbb{N}$ such that $\varphi_1, \dots, \varphi_m, \varphi_1^{-1}, \dots, \varphi_m^{-1} \in \text{Aut}_R(X)$. Then for any $n \in \mathbb{N}$, any $e_1, \dots, e_n \in \{-1, 1\}$, and any $f_1, \dots, f_n \in \{1, \dots, m\}$, we have

$$\varphi_{f_1}^{e_1} \circ \varphi_{f_2}^{e_2} \circ \dots \circ \varphi_{f_n}^{e_n} \in \text{Aut}_{nR}(X).$$

In particular, if $\mathcal{S} := \{\varphi_1, \dots, \varphi_m, \varphi_1^{-1}, \dots, \varphi_m^{-1}\}$ is a (symmetric) generating set for $\langle \varphi_1, \dots, \varphi_m \rangle$, then any reduced word of length n (with respect to \mathcal{S}) is an element of $\{\varphi \in \text{Aut}_{nR}(X) : \varphi^{-1} \in \text{Aut}_{nR}(X)\}$. By (3), there are at most $(Bk(C+4))^q (nR+1)^q$ such words. Therefore $\langle \varphi_1, \dots, \varphi_m \rangle$ is a group of polynomial growth, and its polynomial growth rate is at most q . This holds for any finitely generated subgroup of $\text{Aut}(X)$ (where the parameter R depends on the subgroup and choice of generating set, but B, C, k , and q depend only on the shift (X, σ)). As $q \leq k-1$, the proposition follows.

We are left with showing that (2) holds. There are three cases to consider, depending on the interval in which i lies.

- (i) Suppose that $1 \leq i \leq M_1$. Then $|w_i| \leq C+4R+4$, and so $\varphi(w_i)$ is a word of length $|w_i| - 2R \leq C+2R+4$. Therefore, there are at most

$$p_X(|w_i| - 2R) \leq p_X(C+2R+4) \leq k \cdot (C+2R+4) \leq k(C+4)(R+1)$$

possibilities for the word $\varphi(w_i)$.

- (ii) Suppose that $M_1 < i \leq M_2$. Then w_i does not extend uniquely to its right. If $\varphi \in \text{Aut}(X)$ is such that $\varphi, \varphi^{-1} \in \text{Aut}_R(X)$, then the word $\varphi(w_i) \in \mathcal{L}_{|w_i|-2R}(X)$ cannot extend uniquely $R+1$ times to its right (as otherwise this extended word would have length $2R+1$, and applying φ^{-1} to it would show that there is only one possible extension of w_i to its right). By Corollary 3.3, there are at most $B(R+1)$ such words. Therefore $\{\varphi(w_i) : \varphi, \varphi^{-1} \in \text{Aut}_R(X)\}$ has at most $B(R+1)$ elements.

- (iii) Suppose that $i > M_2$. Then w_i does not extend uniquely to its left. As in Case (ii), if $\varphi \in \text{Aut}(X)$ is such that $\varphi, \varphi^{-1} \in \text{Aut}_R(X)$, then $\varphi(w_i)$ cannot extend uniquely $R+1$ times to its left. By Corollary 3.3, there are at most $B(R+1)$ such words. Therefore $\{\varphi(w_i) : \varphi, \varphi^{-1} \in \text{Aut}_R(X)\}$ has at most $B(R+1)$ elements.

This establishes (2), and thus the proposition. \square

3.3. The automorphism group of a transitive shift.

LEMMA 3.6. *Suppose that (X, σ) is a transitive shift, and that there exists $k \in \mathbb{N}$ such that*

$$\limsup_{n \rightarrow \infty} \frac{P_X(n)}{n} < k.$$

If $x_0 \in X$ has a dense orbit, then the set

$$\{\varphi(x_0) : \varphi \in \text{Aut}(X)\}$$

is contained in the union of finitely many distinct orbits.

Proof. If X consists entirely of periodic points, then the result follows easily. Thus we assume that X contains at least one aperiodic point. We proceed by contradiction, and suppose that the result does not hold.

Let $\varphi_1, \varphi_2, \dots \in \text{Aut}(X)$ be such that $\varphi_i(x_0) \notin \mathcal{O}(\varphi_j(x_0))$ whenever $i \neq j$. For $N \in \mathbb{N}$, let $R(N)$ be the smallest integer such that we have $\varphi_1, \dots, \varphi_N, \varphi_1^{-1}, \dots, \varphi_N^{-1} \in \text{Aut}_{R(N)}(X)$. For $1 \leq i \leq N$, $m \in \mathbb{N}$, and $-n \leq j \leq n$, we have $\varphi_i^{\pm 1} \circ \sigma^j \in \text{Aut}_{R(N)+n}(X)$. As automorphisms take aperiodic points to aperiodic points, for fixed i , the set

$$\{\varphi_i \circ \sigma^j : -n \leq j \leq n\}$$

contains $2n + 1$ elements. If $i_1 \neq i_2$ and $-n \leq j_1, j_2 \leq n$, then $\varphi_{i_1} \circ \sigma^{j_1}(x_0) \notin \mathcal{O}(\varphi_{i_2} \circ \sigma^{j_2}(x_0))$. Thus the set

$$\{\varphi_i \circ \sigma^j : 1 \leq i \leq N \text{ and } -n \leq j \leq n\}$$

contains $2Nn + N$ elements. Therefore,

$$|\{\varphi \in \text{Aut}_{R(N)+n}(X) : \varphi^{-1} \in \text{Aut}_{R(N)+n}(X)\}| \geq 2Nn + N.$$

It follows that

$$\limsup_{R \rightarrow \infty} \frac{|\{\varphi \in \text{Aut}_R(X) : \varphi^{-1} \in \text{Aut}_R(X)\}|}{R} \geq 2N.$$

Since $N \in \mathbb{N}$ was arbitrary, we have

$$\limsup_{R \rightarrow \infty} \frac{|\{\varphi \in \text{Aut}_R(X) : \varphi^{-1} \in \text{Aut}_R(X)\}|}{R} = \infty. \quad (4)$$

On the other hand, since (X, σ) is transitive, the parameter q in the conclusion of Proposition 3.5 is 1. Then, by (3), we have

$$|\{\varphi \in \text{Aut}_R(X) : \varphi^{-1} \in \text{Aut}_R(X)\}| \leq Bk(C + 2)(R + 1),$$

where B, k, C are as in Proposition 3.5; they depend only on the shift (X, σ) and not on R . This estimate holds for any $R \in \mathbb{N}$, a contradiction of (4). \square

We use this to complete the proof of Theorem 1.3, characterizing the automorphism group of transitive shifts of linear growth.

Proof of Theorem 1.3. Assume that (X, σ) is a transitive shift satisfying

$$\limsup_{n \rightarrow \infty} P_X(n)/n < k$$

for some $k \in \mathbb{N}$. An automorphism in a transitive shift is determined by the image of a point whose orbit is dense, and so Lemma 3.6 implies that the group $\text{Aut}(X)/(\text{Aut}(X) \cap [\sigma])$ is finite (Lemma 2.2 implies that $\text{Aut}(X) \cap [\sigma]$ is normal in $\text{Aut}(X)$). However, the only orbit-preserving automorphisms in a transitive shift are elements of $\langle \sigma \rangle$, since such an automorphism acts like a power of the shift on a point whose orbit is dense. \square

Theorem 1.3 shows that if (X, σ) is transitive and has low enough complexity, then $\text{Aut}(X)$ is highly constrained. One might hope to have a converse to this theorem: if (X, σ) is transitive and is above some ‘complexity threshold’ then $\text{Aut}(X)$ is nontrivial. In Section 6.2, we give an example showing that no such converse holds.

3.4. The automorphism group of a shift with dense aperiodic points.

LEMMA 3.7. *Suppose that (X, σ) has a dense set of aperiodic points, and that there exists $k \in \mathbb{N}$ such that*

$$\limsup_{n \rightarrow \infty} \frac{P_X(n)}{n} < k.$$

Let $x_1, \dots, x_q \in X$ be a set (of minimal cardinality) such that $\mathcal{O}(x_1) \cup \dots \cup \mathcal{O}(x_q)$ is dense in X . Then, for each $1 \leq i \leq q$, the set

$$\{\varphi(x_i) : \varphi \in \text{Aut}(X)\}$$

is contained in the union of finitely many distinct orbits.

Proof. By minimality of the set $\{x_1, \dots, x_q\}$, we have

$$x_i \notin \bigcup_{j \neq i} \overline{\mathcal{O}(x_j)}$$

for any $1 \leq i \leq q$. Therefore there exists $w_i \in \mathcal{L}(X)$ such that $[w_i]_0^+ \cap \mathcal{O}(x_i) \neq \emptyset$ but $[w_i]_0^+ \cap \bigcup_{j \neq i} \overline{\mathcal{O}(x_j)} = \emptyset$. This implies that $[w_i]_0^+ \subseteq \overline{\mathcal{O}(x_i)}$.

Let $\varphi \in \text{Aut}(X)$, and note that φ is determined by $\varphi(x_1), \dots, \varphi(x_q)$. If for some $1 \leq i \leq q$ we have $\mathcal{O}(\varphi(x_j)) \cap [w_i]_0^+ = \emptyset$ for all j , then $\varphi(X) \cap [w_i]_0^+ = \emptyset$ and φ is not surjective, a contradiction. Therefore, for each i there exists $1 \leq j_i \leq q$ such that $\mathcal{O}(\varphi(x_{j_i})) \cap [w_i]_0^+ \neq \emptyset$. By construction, if $\mathcal{O}(\varphi(x_{j_i})) \cap [w_i]_0^+ \neq \emptyset$, then $\varphi(x_{j_i}) \in \overline{\mathcal{O}}(x_i)$, and so $\mathcal{O}(\varphi(x_{j_i})) \cap [w_k]_0^+ = \emptyset$ for any $k \neq i$. That is, the map $i \mapsto j_i$ is a permutation on the set $\{1, 2, \dots, q\}$. Let $\pi_\varphi \in S_q$, where S_q is the symmetric group on q letters, denote this permutation.

Let

$$H := \{\pi_\varphi : \varphi \in \text{Aut}(X)\} \subseteq S_q.$$

For each $h \in H$, choose $\varphi_h \in \text{Aut}(X)$ such that $h = \pi_{\varphi_h}$. Then if $\varphi \in \text{Aut}(X)$ and $h = \pi_\varphi$, the permutation induced by $\varphi_h^{-1} \circ \varphi$ is the identity. It follows that $\varphi_h^{-1} \circ \varphi$ preserves each of the sets $\overline{\mathcal{O}}(x_1), \dots, \overline{\mathcal{O}}(x_q)$. Consequently, for each $1 \leq i \leq q$, the restriction of $\varphi_h^{-1} \circ \varphi$ to $\overline{\mathcal{O}}(x_i)$ is an automorphism of the (transitive) subsystem $(\overline{\mathcal{O}}(x_i), \sigma)$. By Lemma 3.6, the set $\{\psi(x_i) : \psi \in \text{Aut}(\overline{\mathcal{O}}(x_i))\}$ is contained in the union of finitely many distinct orbits. Therefore, the set $\{\varphi_h^{-1} \circ \varphi(x_i) : \varphi \in \text{Aut}(X)\}$ is contained in the union of finitely many distinct orbits. Since

$$\{\varphi(x_i) : \varphi \in \text{Aut}(X)\} \subseteq \bigcup_{h \in H} \varphi_h(\{\varphi_{\pi_\varphi}^{-1} \circ \varphi(x_i) : \varphi \in \text{Aut}(X)\})$$

and automorphisms take orbits to orbits, it follows that $\{\varphi(x_i) : \varphi \in \text{Aut}(X)\}$ is contained in the union of finitely many distinct orbits. \square

LEMMA 3.8. *Let (X, σ) be a shift with a dense set of aperiodic points, and assume that there exists $k \in \mathbb{N}$ such that*

$$\limsup_{n \rightarrow \infty} \frac{P_X(n)}{n} < k.$$

Then $\text{Aut}(X) \cap [\sigma] \cong \mathbb{Z}^d$ for some $d < k$.

Proof. By Lemmas 2.2 and 2.3, $\text{Aut}(X) \cap [\sigma]$ is abelian and normal in $\text{Aut}(X)$. By Lemma 3.4, there exist points $x_1, \dots, x_{k-1} \in X$ such that $\mathcal{O}(x_1) \cup \dots \cup \mathcal{O}(x_{k-1})$ is dense in X . If $\varphi \in \text{Aut}(X) \cap [\sigma]$, then there exist $e_1(\varphi), \dots, e_{k-1}(\varphi) \in \mathbb{Z}$ such that $\varphi(x_i) = \sigma^{e_i(\varphi)}(x_i)$ for all $1 \leq i \leq q$. As an automorphism is determined by the images of x_1, \dots, x_{k-1} , the map $\varphi \mapsto (e_1(\varphi), \dots, e_{k-1}(\varphi))$ is an injective homomorphism from $\text{Aut}(X) \cap [\sigma]$ to \mathbb{Z}^{k-1} . \square

Proof of Theorem 1.2. By Lemma 3.4, there exist $x_1, \dots, x_{k-1} \in X$ such that $\mathcal{O}(x_1) \cup \dots \cup \mathcal{O}(x_{k-1})$ is dense in X . If $\varphi \in \text{Aut}(X)$, then φ is determined by the values of $\varphi(x_1), \dots, \varphi(x_{k-1})$. By Lemma 3.7, the set $\{\varphi(x_i) : \varphi \in \text{Aut}(X)\}$

is contained in the union of finitely many distinct orbits in X . Therefore, modulo orbit-preserving automorphisms, there are only finitely many choices for $\varphi(x_1), \dots, \varphi(x_{k-1})$. It follows that the group $\text{Aut}(X)/(\text{Aut}(X) \cap [\sigma])$ is finite. By Lemma 3.8, $\text{Aut}(X) \cong \mathbb{Z}^d$ for some $d < k$. \square

4. General shifts of linear growth

LEMMA 4.1. *Suppose that (X, σ) is a shift, and that $w \in \mathcal{L}(X)$ is such that $[w]_0^+$ is infinite. Then there exists an aperiodic $x_w \in X$ such that $x_w \in [w]_0^+$.*

Proof. Either w occurs syndetically in every element of $[w]_0^+$ with a uniform bound on the gap, or there exists a sequence y_w of elements of $[w]_0^+$ along which the gaps between occurrences of w in y_w grow.

In the first case, the subsystem

$$\overline{\{\sigma^i x : x \in [w]_0^+, i \in \mathbb{Z}\}}$$

is infinite and so contains an aperiodic point x_w . Since w occurs syndetically with the same bound in every element of $[w]_0^+$, it also occurs syndetically in any limit taken along elements of $[w]_0^+$, and in particular in x_w .

In the second case, there is an element of $x_w \in \overline{\mathcal{O}(y_w)} \cap [w]_0^+$ for which either w occurs only finitely many times or infinitely many times with gaps tending to infinity in the semiinfinite word $\{x(n) : n \geq 0\}$, or the same behavior occurs in the semiinfinite word $\{x(n) : n \leq 0\}$. In either case, x_w is aperiodic. \square

We use this to complete the proof of Theorem 1.1, characterizing the finitely generated subgroups of a shift of linear growth.

Proof of Theorem 1.1. Let (X, σ) be a shift, and assume that there exists $k \in \mathbb{N}$ such that

$$\limsup_{n \rightarrow \infty} \frac{P_X(n)}{n} < k.$$

Let

$$X_{NP} := \overline{\{x \in X : \sigma^i(x) \neq x \text{ for all } i \neq 0\}}$$

be the closure of the set of aperiodic points in X . As automorphisms take aperiodic points to aperiodic points, every element of $\text{Aut}(X)$ preserves X_{NP} . Consequently, restriction to X_{NP} defines a natural homomorphism $h : \text{Aut}(X) \rightarrow \text{Aut}(X_{NP})$.

Let $\varphi_1, \dots, \varphi_N \in \text{Aut}(X)$, and choose $R \in \mathbb{N}$ such that $\varphi_1, \dots, \varphi_N, \varphi_1^{-1}, \dots, \varphi_N^{-1} \in \text{Aut}_R(X)$. By Lemma 3.4, there exists a set $x_1, \dots, x_{k-1} \in X_{NP}$ such that

$$\mathcal{O}(x_1) \cup \dots \cup \mathcal{O}(x_{k-1})$$

is dense in X_{NP} . Let $\{x_1, \dots, x_q\} \subseteq X_{NP}$ be a set of minimal cardinality with the property that

$$\mathcal{O}(x_1) \cup \dots \cup \mathcal{O}(x_q)$$

is dense in X_{NP} . Then, for any $\varphi \in \langle \varphi_1, \dots, \varphi_N \rangle$, the restriction of φ to X_{NP} is determined by $\varphi(x_1), \dots, \varphi(x_q)$. By Lemma 3.7, for each $1 \leq j \leq q$, the set

$$\{\varphi(x_j) : \varphi \in \langle \varphi_1, \dots, \varphi_N \rangle\}$$

is contained in the union of finitely many distinct orbits. Therefore there exists a finite collection of automorphisms $\psi_1, \dots, \psi_M \in \langle \varphi_1, \dots, \varphi_N \rangle$ such that, for any $\varphi \in \langle \varphi_1, \dots, \varphi_N \rangle$, there exists $1 \leq t(\varphi) \leq M$ such that, for all $1 \leq j \leq q$, we have

$$\varphi(x_j) \in \mathcal{O}(\psi_{t(\varphi)}(x_j)).$$

Thus the restriction of $\psi_{t(\varphi)}^{-1} \circ \varphi$ to X_{NP} is orbit preserving. Let

$$K := \{\varphi \in \langle \varphi_1, \dots, \varphi_N \rangle : \text{the restriction of } \varphi \text{ to } X_{NP} \text{ is orbit preserving}\}.$$

Clearly K is a subgroup of $\langle \varphi_1, \dots, \varphi_N \rangle$.

For each $1 \leq i \leq N$, we have that φ_i is a block code of range R . Let

$$\mathcal{W}_R := \{w \in \mathcal{L}_{2R+1}(X) : [w]_0^+ \cap X_{NP} = \emptyset\}.$$

Then by Lemma 4.1, the set

$$Y := \bigcup_{w \in \mathcal{W}_R} [w]_0^+$$

is finite. Since every element of Y is periodic and automorphisms preserve the minimal period of periodic points, the $(\langle \varphi_1, \dots, \varphi_N \rangle$ -invariant) set

$$\begin{aligned} Z := \{ & \varphi_{i_1}^{e_1} \circ \dots \circ \varphi_{i_S}^{e_S}(y) : i_1, \dots, i_S \in \{1, \dots, N\}, \\ & e_1, \dots, e_S \in \{-1, 1\}, S \in \mathbb{N}, y \in Y\} \end{aligned}$$

is finite. For any $1 \leq i \leq N$, the restriction of φ_i to X_{NP} uniquely determines the restriction of φ_i to $X \setminus Z$ (since, by definition of \mathcal{W}_R , all words of length $2R + 1$ that occur in elements of $X \setminus Z$ also occur in X_{NP}). Since $\varphi_1, \dots, \varphi_N$ are automorphisms that preserve Z , they take elements of $X \setminus Z$ to elements of $X \setminus Z$. Thus for any $\varphi \in \langle \varphi_1, \dots, \varphi_N \rangle$, the restriction of φ to X_{NP} uniquely determines the restriction of φ to $X \setminus Z$. In particular, this holds for all $\varphi \in K$. Since Z is finite, there exists a finite collection of automorphisms $\alpha_1, \dots, \alpha_T \in K$ such that, for all $\varphi \in K$, there is an integer $1 \leq s(\varphi) \leq T$ such that $\alpha_{s(\varphi)}^{-1} \circ \varphi$ acts trivially on Z .

With the functions $t(\varphi)$ and $s(\varphi)$ defined as above, we have that, for any $\varphi \in \langle \varphi_1, \dots, \varphi_N \rangle$, the automorphism

$$\alpha_{s(\psi_{t(\varphi)}^{-1} \circ \varphi)}^{-1} \circ \psi_{t(\varphi)}^{-1} \circ \varphi$$

acts trivially on Z , and its restriction to X_{NP} is orbit preserving. Define $H \subseteq \langle \varphi_1, \dots, \varphi_N \rangle$ to be the subgroup of elements $\varphi \in \langle \varphi_1, \dots, \varphi_N \rangle$ such that φ acts trivially on Z and the restriction of φ to X_{NP} is orbit preserving. Every element of H is uniquely determined by its restriction to X_{NP} , and so H is isomorphic to a subgroup of $\text{Aut}(X_{NP}) \cap [\sigma]$. By Lemma 3.8, this subgroup is isomorphic to \mathbb{Z}^d for some $d < k$. On the other hand, for any $\varphi \in \langle \varphi_1, \dots, \varphi_N \rangle$, there exist $1 \leq t \leq M$ and $1 \leq s \leq T$ such that $\alpha_s^{-1} \circ \psi_t^{-1} \circ \varphi \in H$. Therefore H has finite index in $\langle \varphi_1, \dots, \varphi_N \rangle$.

Finally, if $\varphi \in H$, then there is a function $k : X_{NP} \rightarrow \mathbb{Z}$ such that for all $x \in X_{NP}$ we have $\varphi(x) = \sigma^{k(x)}(x)$. Thus if $\psi \in \langle \varphi_1, \dots, \varphi_N \rangle$ and $x \in X_{NP}$, we have

$$\psi \circ \varphi \circ \psi^{-1}(x) = \psi \circ \sigma^{k(\psi^{-1}(x))} \circ \psi^{-1}(x) = \sigma^{k(\psi^{-1}(x))}(x),$$

and if $x \in Z$, we have

$$\psi \circ \varphi \circ \psi^{-1}(z) = z.$$

Therefore $\psi \circ \varphi \circ \psi^{-1} \in H$, and so H is a normal subgroup of $\text{Aut}(X)$.

It follows that $\langle \varphi_1, \dots, \varphi_N \rangle$ is virtually \mathbb{Z}^d . This argument holds for any finitely generated subgroup of $\text{Aut}(X)$, and so every finitely generated subgroup of $\text{Aut}(X)$ is virtually \mathbb{Z}^d for some $d < k$. \square

REMARK 4.2. We note that the proof of Theorem 1.1 actually gives a slightly stronger result: for every finitely generated subgroup of $\text{Aut}(X)$, there exists $d < k$ such that this subgroup is \mathbb{Z}^d -by-finite (that is, the finite-index subgroup isomorphic to \mathbb{Z}^d is normal).

5. Minimal shifts of linear growth

For minimal shifts, we need more information on the words that are uniquely extendable.

DEFINITION 5.1. For $x \in X$, define

$$x_R := \{y \in X : y(i) = x(i) \text{ for all } i \geq 0\}.$$

For $x, y \in X$, we write $x \sim_R y$ if $x_R = y_R$, and define $X_R := X / \sim_R$ to be X modulo this relation.

It is easy to check that \sim_R is an equivalence relation on X , and so X_R is well defined. We view (X_R, σ) as a one-sided shift. If $\varphi \in \text{Aut}(X)$, then φ is a block code (say of range N) and so determines an endomorphism on (X_R, σ) as follows: if $y \in x_R$ and $N \in \mathbb{N}$ is the minimal range of φ , then $\varphi(x_R) := (\sigma^N \circ \varphi(y))_R$. It is easy to check that $\varphi(x_R)$ is well defined.

DEFINITION 5.2. For $x \in X$, we say that x_R is *uniquely left extendable* if it has a unique preimage under the shift σ in X_R , and *nonuniquely left extendable* otherwise.

If $w \in \mathcal{L}_n(X)$ is a word of length n in the language of X , we say that w is *uniquely left extendable* if there is a unique $\hat{w} \in \mathcal{L}_{n+1}(X)$ that ends with w .

Boshernitzan [1] showed that if (X, σ) is minimal and there exists $k \in \mathbb{N}$ such that

$$\liminf_{n \rightarrow \infty} P_X(n) - kn = -\infty,$$

then the number of ergodic probability measures on (X, σ) is finite. In his proof, he makes use of a counting lemma, and we use an infinite version of this lemma to study minimal shifts of linear growth.

LEMMA 5.3 (Infinite version of Boshernitzan's lemma). *Let (X, σ) be a shift for which there exists $k \in \mathbb{N}$ such that*

$$\liminf_{n \rightarrow \infty} P_X(n) - kn = -\infty. \quad (5)$$

Then there are at most $k - 1$ distinct elements of (X_R, σ) which are nonuniquely left extendable.

Proof. We first claim that, for infinitely many n , the number of words of length n that are nonuniquely left extendable is at most $k - 1$. If not, let L_n be the number of words of length n that do not extend uniquely to their left. Then by assumption there exists $N \in \mathbb{N}$ such that for all $n \geq N$ we have $L_n \geq k$. However,

$$P_X(n+1) \geq P_X(n) + L_n,$$

and so $P_X(n) \geq P_X(N) + k \cdot (n - N)$ for all $n \geq N$. This contradicts (5), and the claim follows.

We use this to show that there are at most $k - 1$ elements in X_R which are nonuniquely left extendable. If not, there exist distinct elements $x_1, \dots, x_k \in X_R$ which are all nonuniquely left extendable. Choose $M \in \mathbb{N}$ such that, for any $1 \leq i < j \leq k$, there exists $0 \leq m < M$ such that $x_i(m) \neq x_j(m)$. By the first

claim, there exists $n > M$ such that there are at most k words of length n that are nonuniquely left extendable. For all $1 \leq i \leq k$, the word

$$(x_i(0), x_i(1), \dots, x_i(n-2), x_i(n-1))$$

is a word of length n that is nonuniquely left extendable, and these words are pairwise distinct since $n > M$, leading to a contradiction. Thus the number of elements of (X_R, σ) that are nonuniquely left extendable is at most $k - 1$. \square

NOTATION 5.4. We write $\mathcal{Y}_0 \subseteq X_R$ for the collection of nonuniquely left extendable points in X_R . For $m \in \mathbb{N}$, we write $\mathcal{Y}_m := \sigma^m(\mathcal{Y}_0)$ for the collection of elements of X_R whose preimage under m iterates of σ contains more than one point.

LEMMA 5.5. *Suppose that X is infinite, and that X_R, \mathcal{Y}_0 are as in Notation 5.4. Then $\mathcal{Y}_0 \neq \emptyset$.*

Proof. Since X is infinite, X_R is also infinite. The Morse–Hedlund theorem implies that $P_{X_R}(n)$ is unbounded, and so there are infinitely many $n \in \mathbb{N}$ such that $P_{X_R}(n+1) > P_{X_R}(n)$; let (n_i) be a sequence of such n . If every word of length n in $\mathcal{L}(X)$ extended (left) uniquely to a word of length $n+1$, then $P_{X_R}(n+1) = P_{X_R}(n)$. So for all i , there exists a word $w_i \in \mathcal{L}_{n_i}(X_R)$, and there exist two words $u_i, v_i \in \mathcal{L}_{n_i+1}(X_R)$ such that $u_i \neq v_i$ but the rightmost factor (of length n_i) of both u_i and v_i is w_i . For each $i \in \mathbb{N}$, let $x_i \in [u_i]_0^+$, and let $y_i \in [v_i]_0^+$. By compactness of X_R , we can pass to a subsequence x_{i_j} that converges to some $x_\infty \in X_R$. By passing to a further subsequence if necessary, we can assume that y_{i_j} also converges to some $y_\infty \in X_R$. By construction, $\sigma(x_\infty) = \sigma(y_\infty)$ but $x_\infty \neq y_\infty$. Thus $\sigma(x_\infty) \in \mathcal{Y}_0$. \square

LEMMA 5.6. *If $y \in X_R \setminus \bigcup_{m=0}^\infty \mathcal{Y}_m$, then there is a unique $z \in X$ for which $y = z_R$.*

Proof. If not, there exist distinct $z_1, z_2 \in X$ and $y = (z_1)_R = (z_2)_R$. Thus there exists $i \in \mathbb{N}$ such that $z_1(-i) \neq z_2(-i)$. Set i_0 to be the minimal such i . Then $\sigma^{-i_0+1}y = (\sigma^{-i_0+1}z_1)_R = (\sigma^{-i_0+1}z_2)_R$, but $(\sigma^{-i_0}z_1)_R \neq (\sigma^{-i_0}z_2)_R$. Thus $\sigma^{-i_0+1}y \in \mathcal{Y}_0$, which implies that $y \in \mathcal{Y}_{-i_0+1}$, a contradiction. \square

LEMMA 5.7. *If (X, σ) is a shift, $\varphi \in \text{Aut}(X)$, and $y \in \mathcal{Y}_0$, then there exists $m \geq 0$ such that $\varphi(y) \in \mathcal{Y}_m$.*

Proof. If not, then $\varphi(y) \in X_R \setminus \bigcup_{m=0}^\infty \mathcal{Y}_m$, and so Lemma 5.6 implies that there is a unique $z \in X$ such that $\varphi(y) = z_R$. Since φ is an automorphism, it follows that $\varphi^{-1}(z)$ is the only solution to the equation $y = x_R$, a contradiction of $y \in \mathcal{Y}_0$. \square

We use this to complete the characterization of the automorphism group for minimal aperiodic shifts with linear growth.

Proof of Theorem 1.4. Assume that (X, σ) is an aperiodic minimal shift such that there exists $k \in \mathbb{N}$ with $\liminf_{n \rightarrow \infty} P_X(n)/n < k$.

Fix $y \in \mathcal{Y}_0$ (which is nonempty by Lemma 5.5), and let $\varphi \in \text{Aut}(X)$. By Lemma 5.7, there exists $m \in \mathbb{N}$ such that $\varphi(y) \in \mathcal{Y}_m$. Let $m_\varphi \geq 0$ be the smallest nonnegative integer for which $\varphi(y) \in \mathcal{Y}_{m_\varphi}$. Then there exists $z_\varphi \in \mathcal{Y}_0$ such that $\sigma^{m_\varphi}(z_\varphi) = \varphi(y)$.

Now suppose that $\varphi_1, \varphi_2 \in \text{Aut}(X)$ and $z_{\varphi_1} = z_{\varphi_2}$. We claim that φ_1 and φ_2 project to the same element in $\text{Aut}(X)/\langle \sigma \rangle$. Without loss of generality, suppose that $m_{\varphi_1} \leq m_{\varphi_2}$. Then

$$\varphi_2(y) = \sigma^{m_{\varphi_2}}(z_{\varphi_2}) = \sigma^{(m_{\varphi_2} - m_{\varphi_1})} \circ \sigma^{m_{\varphi_1}}(z_{\varphi_1}) = \sigma^{(m_{\varphi_2} - m_{\varphi_1})} \circ \varphi_1(y).$$

By minimality, every word in $\mathcal{L}(X)$ of every length occurs syndetically in every element of (X, σ) . It follows that all words occur syndetically in every element of (X_R, σ) , and in particular, all words occur syndetically in y . Both φ_2 and $\sigma^{(m_{\varphi_2} - m_{\varphi_1})} \circ \varphi_1$ are sliding block codes. Since $\varphi_2(y) = \sigma^{(m_{\varphi_2} - m_{\varphi_1})} \circ \varphi_1(y)$, it follows that φ_2 and $\sigma^{(m_{\varphi_2} - m_{\varphi_1})} \circ \varphi_1$ have the same image on every word, meaning that they define the same block code. In other words, φ_1 and φ_2 project to the same element in $\text{Aut}(X)/\langle \sigma \rangle$, proving the claim.

Since $|\mathcal{Y}_0| \leq k - 1$, Lemma 5.3 implies that there can be at most $k - 1$ distinct elements of (X_R, σ) that arise as z_φ for $\varphi \in \text{Aut}(X)$. Therefore, there are at most $k - 1$ distinct elements of $\text{Aut}(X)/\langle \sigma \rangle$. \square

This can be used to characterize the automorphism groups for particular systems. We note the simplest case of a Sturmian shift for later use (see [12, Example 4.1]).

COROLLARY 5.8. *If (X, σ) is a Sturmian shift, then $\text{Aut}(X) = \langle \sigma \rangle$.*

Proof. For a Sturmian shift, (X, σ) is minimal, aperiodic, and $P_X(n) = n + 1$ for all $n \in \mathbb{N}$. Applying Theorem 1.4 with $k = 2$, we have that $|\text{Aut}(X)/\langle \sigma \rangle| = 1$. \square

More generally, we have the following.

COROLLARY 5.9. *If (X, σ) is aperiodic, minimal, and there exists $k \in \mathbb{N}$ such that*

$$\liminf_{n \rightarrow \infty} P_X(n) - kn = -\infty,$$

then $\text{Aut}(X)$ is the semidirect product of a finite group and \mathbb{Z} .

Proof. By Theorem 1.4, $\text{Aut}(X)/\langle\sigma\rangle$ is finite. Since $\langle\sigma\rangle$ has infinite order and is contained in the center of $\text{Aut}(X)$, it follows from the classification of virtually cyclic groups (see [15]) that $\text{Aut}(X)$ is the semidirect product of a finite group and \mathbb{Z} . \square

6. Examples

6.1. Automorphism group with large polynomial growth. Proposition 3.5 shows that if (X, σ) is a shift satisfying

$$\limsup_{n \rightarrow \infty} \frac{P_X(n)}{n} < k,$$

then $\text{Aut}(X)$ is locally a group of polynomial growth, with polynomial growth rate at most $k - 1$. The following proposition shows that this estimate of the polynomial growth rate of $\text{Aut}(X)$ is optimal.

PROPOSITION 6.1. *Let $k \in \mathbb{N}$ be fixed, and let $\mathcal{A} = \{0, 1\} \times \{1, \dots, k\}$. There is a shift $X \subseteq \mathcal{A}^{\mathbb{Z}}$ with a dense set of aperiodic points such that $P_X(n) = kn + k$ and $\text{Aut}(X) \cong \mathbb{Z}^k$.*

Proof. Recall that a Sturmian shift is an aperiodic minimal shift of $\{0, 1\}^{\mathbb{Z}}$ whose complexity function satisfies $P_X(n) = n + 1$ for all n . There are uncountably many Sturmian shifts, and any particular Sturmian shift only factors onto countably many other Sturmian shifts (since the factor map must be a sliding block code, of which there are only countably many). Therefore there exist k Sturmian shifts X_1, X_2, \dots, X_k such that there exists a sliding block code taking X_i to X_j if and only if $i = j$. We identify X_i with in a natural way with a shift of $\mathcal{A}^{\mathbb{Z}}$ by writing the elements of X_i with the letters $(0, i)$ and $(1, i)$, and will abuse notation by also referring to this shift as X_i . Let $X := X_1 \cup \dots \cup X_k$ (which is clearly shift invariant, and is closed because the minimum distance between a point in X_i and X_j is 1 whenever $i \neq j$).

Let $\varphi \in \text{Aut}(X)$. As φ is given by a sliding block code, φ must preserve the sets X_1, \dots, X_k . Therefore $\text{Aut}(X) \cong \text{Aut}(X_1) \times \dots \times \text{Aut}(X_k)$. By Corollary 5.8, we have $\text{Aut}(X_i) = \langle\sigma\rangle \cong \mathbb{Z}$ for $i = 1, \dots, k$. So $\text{Aut}(X) \cong \mathbb{Z}^k$. \square

6.2. Quickly growing transitive shifts with trivial automorphism group.

Next, we describe a general process which takes a minimal shift of arbitrary growth rate and produces a transitive shift with essentially the same growth, but whose automorphism group consists only of powers of the shift. This shows that there is no ‘complexity threshold’ above which the automorphism group of a transitive shift must be nontrivial.

LEMMA 6.2. *If (X, σ) is a transitive shift with precisely one dense orbit, then $\text{Aut}(X) = \langle \sigma \rangle$.*

Proof. Suppose that there exists $x_0 \in X$ such that

$$\{y \in X : y \text{ has a dense orbit}\} = \mathcal{O}(x_0).$$

If $\varphi \in \text{Aut}(X)$, then $\varphi(x_0)$ has a dense orbit, and so there exists $k \geq 0$ such that $\varphi(x_0) = \sigma^k(x_0)$. It follows that φ and σ^k agree on the (dense) orbit of x_0 . Since both functions are continuous, they agree everywhere. \square

EXAMPLE 6.3. Let $\mathcal{A} = \{0, 1, 2, \dots, d - 1\}$, and let $X \subseteq \mathcal{A}^{\mathbb{Z}}$ be a minimal shift. Let $\tilde{\mathcal{A}} = \mathcal{A} \cup \{d\}$, where we add the symbol d to the alphabet, and $d \notin \mathcal{A}$. Fix $x_0 \in X$, and define $\tilde{x}_0 \in \tilde{\mathcal{A}}^{\mathbb{Z}}$ by

$$\tilde{x}_0(i) = \begin{cases} x_0(i) & \text{if } i \neq 0; \\ d & \text{if } i = 0. \end{cases}$$

Let $\tilde{X} \subseteq \tilde{\mathcal{A}}^{\mathbb{Z}}$ be the orbit closure of \tilde{x}_0 . Then $\tilde{X} = X \cup \mathcal{O}(\tilde{x}_0)$, (\tilde{X}, σ) is transitive, $p_{\tilde{X}}(n) = p_X(n) + n$ for all $n \in \mathbb{N}$ (for $n \in \mathbb{N}$, $\mathcal{L}_n(X) \subseteq \mathcal{L}_n(\tilde{X})$ by minimality but $\mathcal{L}_n(\tilde{X})$ has exactly n additional words that contain the symbol d), and \tilde{X} has precisely one dense orbit. By Lemma 6.2, $\text{Aut}(\tilde{X}) = \langle \sigma \rangle$.

6.3. $\text{Aut}(X)$ and $\text{Aut}(X)/(\text{Aut}(X) \cap [\sigma])$ are not always finitely generated.

Theorem 1.1 shows that every finitely generated subgroup of $\text{Aut}(X)$ is virtually \mathbb{Z}^d . When X has a dense set of aperiodic points, Theorem 1.2 shows that $\text{Aut}(X)/(\text{Aut}(X) \cap [\sigma])$ is finite. In this section, we show that the result of Theorem 1.2 cannot be extended to the general case, and the words ‘every finitely generated subgroup of’ cannot be removed from the statement of Theorem 1.1. We begin with an example to set up our construction.

EXAMPLE 6.4. Let $\mathcal{A} = \{0, 1\}$, and for $n \in \mathbb{N}$, let $x_n \in \mathcal{A}^{\mathbb{Z}}$ be the periodic point

$$x_n(i) = \begin{cases} 1 & \text{if } i \equiv 0 \pmod{2^n}; \\ 0 & \text{otherwise.} \end{cases}$$

Let X be the closure of the set $\{x_n : n \in \mathbb{N}\}$ under σ . If we define

$$x_{\infty}(i) = \begin{cases} 1 & \text{if } i = 0; \\ 0 & \text{otherwise,} \end{cases}$$

and $\mathbf{0}$ to be the \mathcal{A} -coloring of all zeros, then we have

$$X = \{\mathbf{0}\} \cup \mathcal{O}(x_\infty) \cup \bigcup_{n=1}^\infty \mathcal{O}(x_n).$$

Suppose that $R \in \mathbb{N}$ is fixed, and that $\varphi \in \text{Aut}_R(X)$. Since φ preserves the period of periodic points, $\varphi(\mathbf{0}) = \mathbf{0}$. In particular, the block code φ takes the block consisting of all zeros to $\mathbf{0}$. It follows that there exists $k \in [-R, R]$ such that $\varphi(x_\infty) = \sigma^k(x_\infty)$. For any $m > 2R + 1$, the blocks of length $2R + 1$ occurring in x_m are identical to those appearing in x_∞ , and so $\varphi(x_m) = \sigma^k(x_m)$ for all such m .

Now let $\varphi_1, \dots, \varphi_n \in \text{Aut}(X)$, and find $R \in \mathbb{N}$ such that $\varphi_1, \dots, \varphi_n, \varphi_1^{-1}, \dots, \varphi_n^{-1} \in \text{Aut}_R(X)$. For $1 \leq i \leq n$, let $k_i \in [-R, R]$ be such that for all $m > 2R + 1$ we have $\varphi_i(x_m) = \sigma^{k_i}(x_m)$. Then for $N \in \mathbb{N}$, any $e_1, \dots, e_N \in \{1, \dots, n\}$, any $\epsilon_1, \dots, \epsilon_N \in \{-1, 1\}$, and any $m > 2R + 1$, we have

$$(\varphi_{e_1}^{\epsilon_1} \circ \varphi_{e_2}^{\epsilon_2} \circ \dots \circ \varphi_{e_N}^{\epsilon_N})(x_m) = \sigma^{(\epsilon_1 \cdot k_{e_1} + \epsilon_2 \cdot k_{e_2} + \dots + \epsilon_N \cdot k_{e_N})}(x_m).$$

Then if $\varphi \in \text{Aut}(X)$ is the automorphism that acts like σ on $\mathcal{O}(x_{2R+2})$ and acts trivially on $X \setminus \mathcal{O}(x_{2R+2})$ (this map is continuous because x_{2R+2} is isolated), then $\varphi \notin \langle \varphi_1, \dots, \varphi_n \rangle$. Therefore $\langle \varphi_1, \dots, \varphi_n \rangle \neq \text{Aut}(X)$. Since $\varphi_1, \dots, \varphi_n \in \text{Aut}(X)$ were general, it follows that $\text{Aut}(X)$ is not finitely generated.

On the other hand, we can compute the value of $P_X(n)$ as follows. There are exactly $n + 1$ words in $\mathcal{L}_n(X)$ that contain at most one 1. For all $m \leq \lfloor \log_2(n) \rfloor - 1$, there are exactly 2^m words in $\mathcal{L}_n(X)$ that arise as factors of x_m (all such words contain at least two 1s, and their period can be read off as the spacing between the 1s). In addition, there are exactly $n - 2^{\lfloor \log_2(n) \rfloor}$ words in $x_{\lfloor \log_2(n) \rfloor}$ that contain at least two 1s. Therefore

$$\begin{aligned} P_X(N) &= (n + 1) + \sum_{m=1}^{\lfloor \log_2(n) \rfloor - 1} 2^m + (n - \lfloor \log_2(n) \rfloor) \\ &= 2n + 2^{\lfloor \log_2(n) \rfloor} - \lfloor \log_2(n) \rfloor - 1 < 4n \end{aligned}$$

for all n , so $P_X(n)$ grows linearly. We also remark that $\text{Aut}(X) = \text{Aut}(X) \cap [\sigma]$ for this shift.

PROPOSITION 6.5. *There exists a shift (X, σ) of linear growth that has a dense set of periodic points and is such that none of the groups $\text{Aut}(X)$, $\text{Aut}(X) \cap [\sigma]$, and $\text{Aut}(X)/(\text{Aut}(X) \cap [\sigma])$ are finitely generated.*

Proof. Let X_1 be the shift of $\{0, 1\}^{\mathbb{Z}}$ constructed in the previous example. Let X_2 be the same shift, constructed over the alphabet $\{2, 3\}$ (by identifying 0 with

2 and 1 with 3). Let $X = X_1 \cup X_2$, and observe that $d(X_1, X_2) = 1$. Since $\text{Aut}(X_i) \cap [\sigma] = \text{Aut}(X_i)$ for $i = 1, 2$, we have $\text{Aut}(X) \cap [\sigma] \cong \text{Aut}(X_1) \times \text{Aut}(X_2)$. Therefore $\text{Aut}(X) \cap [\sigma]$ is not finitely generated. On the other hand,

$$P_X(n) = P_{X_1}(n) + P_{X_2}(n) = 2 \cdot P_{X_1}(n) < 8n,$$

so X is a shift of linear growth (and has a dense set of periodic points).

We claim that $\text{Aut}(X)/(\text{Aut}(X) \cap [\sigma])$ is not finitely generated. Define $\delta \in \text{Aut}(X)$ to be the involution of range 0 (or, equivalently, the 1-block code) that exchanges 0 with 2 and 1 with 3. For each $m \in \mathbb{N}$, let $\delta_m \in \text{Aut}(X)$ be the involution of range $(2^{m+1} + 1)$ which exchanges the (unique) orbit of period 2^m in X_1 with the (unique) orbit of period 2^m in X_2 by exchanging 0 with 2 and 1 with 3 in these orbits only (and fixing the remainder of X). For $i \in \mathbb{N}$, let $\tilde{\delta}_i$ be the projection of δ_i to $\text{Aut}(X)/(\text{Aut}(X) \cap [\sigma])$, and let $\tilde{\delta}$ be the projection of δ . These involutions commute pairwise, and one can check that the set $\{\tilde{\delta}_i : i \in \mathbb{N}\} \cup \{\tilde{\delta}\}$ generates a subgroup of $\text{Aut}(X)/(\text{Aut}(X) \cap [\sigma])$ that is isomorphic to $\prod_{i=1}^{\infty} (\mathbb{Z}/2\mathbb{Z})$.

Now let $\mathbf{x} \in X_1$ be the point $\mathbf{x}(i) = 1$ if and only if $i = 0$, and let $\mathbf{y} \in X_2$ be the point $\mathbf{y}(i) = 3$ if and only if $i = 0$. Let $\varphi \in \text{Aut}(X)$ be fixed, and observe that either $\varphi(\mathbf{x}) \in \mathcal{O}(\mathbf{x})$ or $\varphi(\mathbf{x}) \in \mathcal{O}(\mathbf{y})$. In the former case, define $\epsilon := 0$, and in the latter case, define $\epsilon := 1$, so that $\varphi \circ \delta^\epsilon$ preserves the orbit of \mathbf{x} (hence also the orbit of \mathbf{y}). As $\varphi \circ \delta^\epsilon$ is given by a block code which carries the block of all 0s to 0, there are at most finitely many m such that $\varphi \circ \delta^\epsilon$ does not preserve the orbit of the (unique) periodic orbit of period 2^m in X_1 . Let $m_1 < \dots < m_n$ be the set of m for which it does not preserve the orbit. Then

$$\varphi \circ \delta^\epsilon \circ \delta_{m_1} \circ \dots \circ \delta_{m_n} \in \text{Aut}(X) \cap [\sigma].$$

Therefore $\text{Aut}(X)/(\text{Aut}(X) \cap [\sigma])$ is the group generated by $\tilde{\delta}, \tilde{\delta}_1, \tilde{\delta}_2, \tilde{\delta}_3, \dots$. This group is isomorphic to $\prod_{i=1}^{\infty} (\mathbb{Z}/2\mathbb{Z})$, so $\text{Aut}(X)/(\text{Aut}(X) \cap [\sigma])$ is not finitely generated. Finally, as $\text{Aut}(X)$ factors onto a group that it not finitely generated, it is not finitely generated either. □

7. Automorphisms of periodic shifts

We characterize which finite groups arise as automorphism groups of shifts.

DEFINITION 7.1. For $n > 1$ and $m \in \mathbb{N}$, let

$$\mathbb{Z}_n^m := \underbrace{\mathbb{Z}_n \times \mathbb{Z}_n \times \dots \times \mathbb{Z}_n}_{m \text{ times}}$$

where \mathbb{Z}_n denotes $\mathbb{Z}/n\mathbb{Z}$. Let S_m denote the symmetric group on m letters and define a homomorphism $\psi : S_m \rightarrow \text{Aut}(\mathbb{Z}_n^m)$ by

$$(\psi(\pi))(i_1, \dots, i_m) := (i_{\pi(1)}, \dots, i_{\pi(m)}).$$

Then the *generalized symmetric group* is defined as in [13] to be

$$S(n, m) := \mathbb{Z}_n^m \rtimes_{\psi} S_m.$$

Equivalently, $S(n, m)$ is the wreath product $\mathbb{Z}_n \wr S_m$.

THEOREM 7.2. *Suppose that G is a finite group. There exists a shift (X, σ) for which $\text{Aut}(X) \cong G$ if and only if there exist $s \in \mathbb{N}$, $n_1 < n_2 < \dots < n_s$, and $m_1, m_2, \dots, m_s \in \mathbb{N}$ such that*

$$G \cong S(n_1, m_1) \times S(n_2, m_2) \times \dots \times S(n_s, m_s).$$

Proof. Suppose that (X, σ) is a shift for which $\text{Aut}(X)$ is finite. Since $\sigma \in \text{Aut}(X)$, there exists $k \in \mathbb{N}$ such that $\sigma^k(x) = x$ for all $x \in X$. That is, X is comprised entirely of periodic points such that the minimal period of each point is a divisor of k . Since a shift can have only finitely many such points, X is finite. Let $x_1, \dots, x_N \in X$ be representatives of the orbits in X , meaning that $\mathcal{O}(x_i) \cap \mathcal{O}(x_j) = \emptyset$ whenever $i \neq j$, and for all $x \in X$ there exist $i, k \in \mathbb{N}$ such that $x = \sigma^k(x_i)$. For $i = 1, \dots, N$, let p_i be the minimal period of x_i , and, without loss of generality, assume that $p_1 \leq p_2 \leq \dots \leq p_N$. Define $n_1 := p_1$, and inductively define n_2, n_3, \dots, n_s by

$$n_{i+1} := \min\{p_j : p_j > n_i\},$$

where s is the number of steps before the construction terminates. Define

$$m_i := |\{j : p_j = n_i\}|.$$

Let $\varphi \in \text{Aut}(X)$. Then, for $1 \leq i \leq s$, φ induces a permutation on the set of periodic points of minimal period n_i . More precisely, for fixed $1 \leq i \leq s$, we can define $\pi_i^\varphi \in S_{m_i}$ to be the permutation that sends $j \in \{1, 2, \dots, m_i\}$ to the unique integer $k \in \{1, 2, \dots, m_i\}$ such that $\varphi(x_{m_1+\dots+m_{i-1}+j}) \in \mathcal{O}(x_{m_1+\dots+m_{i-1}+k})$. For $1 \leq j \leq m_i$, choose $k_j^i \in \mathbb{Z}_{n_i}$ such that

$$\varphi(x_{m_1+\dots+m_{i-1}+j}) = \sigma^{k_j^i}(x_{m_1+\dots+m_{i-1}+\pi_i^\varphi(j)}).$$

Then the map Φ given by

$$\Phi(\varphi) = (k_1^1, k_2^1, \dots, k_{m_1}^1, \pi_1^\varphi, k_1^2, \dots, k_{m_2}^2, \pi_2^\varphi, \dots, k_1^s, \dots, k_{m_s}^s, \pi_s^\varphi)$$

is a homomorphism from $\text{Aut}(X)$ to $S(n_1, m_1) \times \cdots \times S(n_s, m_s)$. The kernel of this map is trivial, since any such automorphism fixes x_i for all i (and so also fixes every element of X). To check that it is surjective, if π_1, \dots, π_s are permutations ($\pi_i \in S_{m_i}$ for all i), then define

$$\varphi_{\pi_1, \dots, \pi_s}(x_{m_1 + \dots + m_{i-1} + j}) = x_{m_1 + \dots + m_{i-1} + \pi_i(j)},$$

and extend this to an automorphism of (X, σ) . Similarly, for $1 \leq i \leq N$, define

$$\varphi_i(\sigma^k(x_j)) := \sigma^{k + \delta_{i,j}}(x_j),$$

where $\delta_{i,j}$ is the Kronecker delta. Note that each of these maps is given by a block code, where the range is the smallest R such that $\sigma^R(x) = x$ for all $x \in X$. Taken together, this shows that the map Φ is surjective, and thus is an isomorphism.

Conversely, suppose that $n_1 < \cdots < n_s$, and that $m_1, \dots, m_s \in \mathbb{N}$ are given. For $1 \leq i \leq s$ and $1 \leq j \leq m_i$, define

$$x_{i,j}(k) := \begin{cases} j & \text{if } k \equiv 0 \pmod{n_i}; \\ 0 & \text{otherwise.} \end{cases}$$

Let

$$X' = \bigcup_{i=1}^s \bigcup_{j=1}^{m_i} x_{i,j},$$

and let X be the closure of X' under σ . Then X consists of periodic points, with precisely m_i distinct orbits of minimal period n_i , for $1 \leq i \leq s$. The reader can check that the automorphism group of a shift that consists of exactly m distinct orbits of minimal period n is isomorphic to $S(n, m)$. Thus we have

$$\text{Aut}(X) \cong S(n_1, m_1) \times \cdots \times S(n_s, m_s). \quad \square$$

Acknowledgements

We thank J. Davis for pointing us to reference [15], P. Brooksbank for helpful conversations, and the referee for numerous helpful comments. The second author was partially supported by NSF grant DMS-1200971.

References

- [1] M. Boshernitzan, 'A unique ergodicity of minimal symbolic flows with linear block growth', *J. Anal. Math.* **44**(1) (1984), 77–96.
- [2] M. Boyle and W. Krieger, 'Periodic points and automorphisms of the shift', *Trans. Amer. Math. Soc.* **302** (1987), 125–149.
- [3] M. Boyle, D. Lind and D. Rudolph, 'The automorphism group of a shift of finite type', *Trans. Amer. Math. Soc.* **306**(1) (1988), 71–114.

- [4] J. Cassaigne, ‘Special factors of sequences with linear subword complexity’, in *Developments in Language Theory, II (Magdeburg, 1995)* (World Science Publications, River Edge, NJ, 1996), 25–34.
- [5] V. Cyr and B. Kra, ‘The automorphism group of a shift of subquadratic growth’, *Proc. Amer. Math. Soc.* (to appear), Preprint, [arXiv:1403.0238](https://arxiv.org/abs/1403.0238).
- [6] S. Donoso, F. Durand, A. Maass and S. Petite, ‘On automorphism groups of low complexity minimal subshifts’, Preprint, [arXiv:1501.00510](https://arxiv.org/abs/1501.00510).
- [7] D. Fiebig and U. Fiebig, ‘The automorphism group of a coded system’, *Trans. Amer. Math. Soc.* **348**(8) (1996), 3173–3191.
- [8] G. A. Hedlund, ‘Endomorphisms and automorphisms of the shift dynamical system’, *Math. Systems Theory* **3** (1969), 320–375.
- [9] M. Hochman, ‘On the automorphism groups of multidimensional shifts of finite type’, *Ergodic Theory Dynam. Systems* **30**(3) (2010), 809–840.
- [10] K. H. Kim, F. W. Roush and J. B. Wagoner, ‘Automorphisms of the dimension group and gyration numbers’, *J. Amer. Math. Soc.* **5** (1993), 191–212.
- [11] M. Morse and G. A. Hedlund, ‘Symbolic dynamics II. Sturmian trajectories’, *Amer. J. Math.* **62** (1940), 1–42.
- [12] J. Olli, ‘Endomorphisms of Sturmian systems and the discrete chair substitution tiling system’, *Discrete Contin. Dyn. Syst.* **33**(9) (2013), 4173–4186.
- [13] M. Osima, ‘Some remarks on the characters of the symmetric group II’, *Canad. J. Math.* **6** (1954), 511–521.
- [14] V. Salo and I. Törmä, ‘Block maps between primitive uniform and Pisot substitutions’, *Ergodic Theory Dynam. Systems* (to appear), Preprint, [arXiv:1306.3777](https://arxiv.org/abs/1306.3777).
- [15] P. Scott and T. Wall, ‘Topological methods in group theory’, in *Homological Group Theory (Proc. Sympos., Durham, 1977)*, London Mathematical Society Lecture Note Series, 36 (Cambridge University Press, Cambridge–New York, 1979), 137–203.
- [16] T. Ward, ‘Automorphisms of Z^d -subshifts of finite type’, *Indag. Math. (N.S.)* **5**(4) (1994), 495–504.