Ergod. Th. & Dynam. Sys. (2020), **40**, 1788–1804
doi:10.1017/etds.2018.126

The automorphism group of a shift of slow growth is amenable

VAN CYR† and BRYNA KRA‡

† *Bucknell University, Lewisburg, PA 17837, USA*
(e-mail: van.cyr@bucknell.edu)
‡ *Northwestern University, Evanston, IL 60208, USA*
(e-mail: kra@math.northwestern.edu)

(*Received* 21 *August* 2017 *and accepted in revised form* 2 *November* 2018)

Abstract. Suppose (X, σ) is a subshift, $P_X(n)$ is the word complexity function of X, and $\mathrm{Aut}(X)$ is the group of automorphisms of X. We show that if $P_X(n) = o(n^2/\log^2 n)$, then $\mathrm{Aut}(X)$ is amenable (as a countable, discrete group). We further show that if $P_X(n) = o(n^2)$, then $\mathrm{Aut}(X)$ can never contain a non-abelian free monoid (and, in particular, can never contain a non-abelian free subgroup). This is in contrast to recent examples, due to Salo and Schraudner, of subshifts with quadratic complexity that do contain such a monoid.

Key words: subshift, automorphism, block complexity, amenable
2010 Mathematics Subject Classification: 37B50 (Primary); 43A07, 68R15 (Secondary)

1. *Amenability and the automorphism group*

For a subshift (X, σ) over a finite alphabet, let $\mathrm{Aut}(X) = \mathrm{Aut}(X, \sigma)$ denote the group of all automorphisms of the system, meaning the collection of all homeomorphisms $\phi : X \to X$ such that $\phi \circ \sigma = \sigma \circ \phi$. The automorphism group of many subshifts with positive entropy, including the full shift and more generally any mixing shift of finite type, is a countable group that contains many structures, including isomorphic copies of any finite group, countably many copies of $\mathbb{Z}$, and the free group on countably many generators (see [**1**, **8**]). In particular, when given the discrete topology, these automorphism groups are never amenable. This behavior is in contrast to what happens in minimal shifts of zero entropy: if the complexity function $P_X(n)$, which counts the number of words in the language of the shift, satisfies $\limsup_{n\to\infty} (\log(P_X(n))/n^\beta) = 0$ for some $\beta < 1/2$, then the automorphism group $\mathrm{Aut}(X)$ is amenable; furthermore, every finitely generated torsion-free subgroup of the automorphism group has subexponential growth [**5**]. For lower complexities, one can sometimes carry out a more detailed analysis

of the automorphism group, and this is done for polynomial growth in [**5**], and with extra assumptions on the dynamics, sometimes one can give a complete description of the automorphism group (see [**3**, **4**, **6**]).

We continue the systematic study of automorphism groups here, focusing on subshifts with zero entropy. These automorphism groups are constrained by the subexponential growth rate of words in the language of the shift, and it seems plausible that for any subshift (X, σ) of zero entropy, we may have the following dichotomy: either $\mathrm{Aut}(X)$ contains a non-abelian free subgroup or $\mathrm{Aut}(X)$ is amenable. Somewhat surprisingly, we cannot rule out that such an alternative holds for any shift, even without an assumption on the entropy. For example, this dichotomy holds for any mixing subshift of finite type [**1**], since the automorphism group contains the free group on two generators, and it holds for other classes of subshifts with positive entropy, such as Toeplitz systems, where the automorphism group is abelian [**7**]. In a different vein, the automorphism group of a zero-entropy subshift can never contain logarithmically distorted elements [**2**].

Furthermore, a stronger result is plausible, namely that for any zero-entropy subshift, the automorphism group is amenable. Numerous results support this statement: the automorphism group is amenable for any minimal subshift whose complexity is stretched exponential with exponent less than $1/2$, for all subshifts with linear complexity, and for several other classes of low complexity shifts (see [**3**–**6**]).

To address these questions and conjectures, we give a detailed analysis of the algebraic properties of $\mathrm{Aut}(X)$ for shifts whose complexity is at most quadratic. In [**3**], we showed that for a transitive shift with subquadratic growth, after quotienting the automorphism group by the subgroup generated by the shift, we are left with a periodic group. This left open a stronger description of this automorphism group, as well as what happens without an assumption of transitivity. As a first step in addressing this, we show (see §2 for precise definitions) the following theorem.

THEOREM 1. *Assume that* (X, σ) *is a subshift whose complexity function satisfies* $P_X(n) = o(n^2/\log^2 n)$. *Then* $\mathrm{Aut}(X)$ *is amenable (as a countable discrete group).*

In particular, the automorphism group of any shift whose complexity is $o(n^{2-\varepsilon})$, for some $\varepsilon > 0$, is amenable.

The techniques to prove Theorem 1 follow a basic strategy developed in [**5**], but deducing the theorem without the assumption that X is minimal adds significant technical difficulties. One of the new ideas used is the construction of a descending chain of subshifts for which each term retains some of the properties that make minimal shifts easier to study. We believe that this technique should prove to be applicable in other settings.

Unfortunately, our methods do not easily extend to a shift whose complexity is $o(n^2)$, but in this setting we are able to prove a morally weaker result that holds for this larger class of shifts.

THEOREM 2. *Assume that* (X, σ) *is a subshift whose complexity function satisfies* $\liminf_{\to\infty} (P_X(n)/n^2) = 0$. *Then* $\mathrm{Aut}(X)$ *does not contain an isomorphic copy of the free monoid on two generators.*

In particular, such an automorphism group cannot contain any non-abelian free subgroup. The interest in this theorem is the contrast with a recent result of Salo and Schraudner [**11**]. They constructed a subshift (X, σ) whose complexity function satisfies $P_X(n) = (n+1)^2$ and whose automorphism group contains a free monoid on two generators, and the automorphism group is still amenable. This highlights the subtle issues that arise in addressing questions and conjectures on amenability of $\mathrm{Aut}(X)$ when (X, σ) is not minimal, and the difficulty in passing beyond quadratic complexity.

2. *Background*

2.1. *Subshifts.* Let $\mathcal{A}$ be a finite alphabet and endow $\mathcal{A}^{\mathbb{Z}}$ with the topology induced by the metric

$$d(x, y) := 2^{-\inf\{|i| : x_i \neq y_i\}}.$$

For $x \in \mathcal{A}^{\mathbb{Z}}$, we denote the ith coordinate of x by x_i. For $n \in \mathbb{N}$, an element $w = (w_0, \ldots, w_{n-1}) \in \mathcal{A}^n$ is called a *word of length* n. If w is a word of length n, then the set

$$[w]_0^+ := \{x \in \mathcal{A}^{\mathbb{Z}} : x_i = w_i \text{ for all } 0 \leq i < n\}$$

is the *cylinder set* determined by w. The *(left) shift* $\sigma : \mathcal{A}^{\mathbb{Z}} \to \mathcal{A}^{\mathbb{Z}}$ is the map $x \mapsto \sigma x$ given by $(\sigma x)_i := x_{i+1}$ for all $i \in \mathbb{Z}$, and it is a homeomorphism of $\mathcal{A}^{\mathbb{Z}}$. If $x \in \mathcal{A}^{\mathbb{Z}}$ and there exists $p > 0$ such that $\sigma^p x = x$, then x is *periodic* of period p. If no such p exists, x is *aperiodic*.

A closed, σ-invariant subset $X \subset \mathcal{A}^{\mathbb{Z}}$ together with the shift $\sigma : X \to X$ is called a *subshift*. If X is a subshift, we define the *language* $\mathcal{L}(X)$ of X to be

$$\mathcal{L}(X) := \left\{ w \in \bigcup_{n=1}^{\infty} \mathcal{A}^n : [w]_0^+ \cap X \neq \varnothing \right\}.$$

For $n \in \mathbb{N}$, the set $\mathcal{L}_n(X) := \mathcal{L}(X) \cap \mathcal{A}^n$ denotes the set of *words of length* n in the language of X, and we denote the length of word $w \in \mathcal{L}(X)$ by $|w|$.

2.2. *Complexity.* The *complexity function of* X is the function $P_X : \mathbb{N} \to \mathbb{N}$ defined by $P_X(n) := |\mathcal{L}_n(X)|$. If $x \in \mathcal{A}^{\mathbb{Z}}$, then the *orbit closure* $\mathcal{O}(x)$ *of* x under the shift

$$\overline{\mathcal{O}(x)} := \overline{\{\sigma^i x : i \in \mathbb{Z}\}}$$

is also a subshift. We commit a slight abuse of terminology and refer to $P_{\overline{\mathcal{O}(x)}}(n)$ as the *complexity function of* x. To avoid confusion, we use the lower-case letter x to refer to an element of $\mathcal{A}^{\mathbb{Z}}$ and the upper-case letter X to refer to a subshift of $\mathcal{A}^{\mathbb{Z}}$. The basic result relating dynamical properties of x to its complexity is the Morse–Hedlund theorem [**9**]: an element $x \in \mathcal{A}^{\mathbb{Z}}$ is aperiodic if and only if its complexity function is bounded below by $n+1$ for all n.

Suppose $w = (w_0, \ldots, w_{n-1}) \in \mathcal{L}_n(X)$ and $L \in \mathbb{N}$ is fixed. We say that w *extends uniquely* L times to the right and left (in the language of X) if there is a unique $u = (u_0, \ldots, u_{n+2L-1}) \in \mathcal{L}_{n+2L}(X)$ such that $w_i = u_{i+L}$ for $0 \leq i < n$. If $w \in \mathcal{L}(X)$ and $u \in \mathcal{L}_n(X)$ for some $n \geq |w|$, we say that w is a *subword* of u if there exists $0 \leq i < n - |w|$ such that $u_j = w_j$ for $i \leq j < |w|$. Thus if w extends uniquely L times in X, then if

$x \in X$ and if $w = (x_j, \dots, x_{j+n-1})$ for some $j \in \mathbb{Z}$, then $u = (x_{j-L}, \dots, x_{j+n+L-1})$. Rephrasing this, whenever w is a word in x, then u is also a word in x and w is a subword of u.

2.3. *The automorphism group.* If (X, σ) is a subshift and $\mathrm{Hom}(X)$ is the group of all homeomorphisms from X to itself, then the *group of automorphisms* of X, denoted $\mathrm{Aut}(X)$, is the centralizer of σ in $\mathrm{Hom}(X)$. (Strictly speaking, we should write $\mathrm{Aut}(X, \sigma)$ instead of $\mathrm{Aut}(X)$, but we assume that the subshift is endowed with the shift and omit explicit mention of σ from most of our notation.) A function $\varphi : X \to X$ is called a *sliding block code* if there exists $R \in \mathbb{N} \cup \{0\}$ such that $(\varphi x)_0$ is a function of $(x_{-R}, \dots, x_R)$ for all $x \in X$. In this case, R is called a *range* of φ. The classical Curtis–Hedlund–Lyndon theorem [**8**] states that every automorphism of X is a sliding block code. In particular, this means that for any subshift (X, σ), the automorphism group $\mathrm{Aut}(X)$ is countable.

For $R \in \mathbb{N} \cup \{0\}$, define $\mathrm{Aut}_R(X)$ to be the set of all $\phi \in \mathrm{Aut}(X)$ such that both ϕ and its inverse are given by sliding block codes of range R. Since any block code of range R is also a block code of range $R + 1$, we have

$$\mathrm{Aut}_0(X) \subset \mathrm{Aut}_1(X) \subset \mathrm{Aut}_2(X) \subset \mathrm{Aut}_3(X) \subset \cdots$$

and $\mathrm{Aut}(X) = \bigcup_{R=0}^{\infty} \mathrm{Aut}_R(X)$. If $\varphi \in \mathrm{Aut}_R(X)$ and if $w \in \mathcal{L}(X)$ is such that $|w| \geq 2R + 1$, then we define $\varphi(w)$ to be the word of length $|w| - 2R$ obtained by applying the block code (of range R) defining φ to w. Note that this definition is not intrinsic to φ but rather to φ together with a range R. Whenever we apply an automorphism $\varphi \in \mathrm{Aut}_R(X)$ to a word w, we are implicitly choosing R to be the range of φ.

For each $w \in \mathcal{L}(X)$ and $n \in \mathbb{N}$, define the function $\mathcal{W}_n^w : ([w]_0^+ \cap X) \to \mathcal{L}_{|w|+2n}(X)$ by

$$\mathcal{W}_n^w(x) := (x_{-n}, x_{-n+1}, \dots, x_0, x_1, \dots, x_{|w|+n-2}, x_{|w|+n-1}).$$

With this notation, w extends uniquely n times to the right and left (in the language of X) if $\mathcal{W}_n^w(x)$ is independent of $x \in [w]_0^+ \cap X$.

Suppose $u_1, \dots, u_k, v \in \mathcal{L}(X)$. An automorphism $\phi \in \mathrm{Aut}(X)$ *preserves occurrences of* v if $\phi([v]_0^+ \cap X) \subset [v]_0^+ \cap X$ and $\phi^{-1}([v]_0^+ \cap X) \subset [v]_0^+ \cap X$. If $D \in \mathbb{N}$, then ϕ *preserves occurrences of* v *when it is* D *units from* $u_1, \dots, u_k$ if, for any $x \in [v]_0^+ \cap X$ such that $\mathcal{W}_D^v(x)$ does not contain u_i as a subword for any $1 \leq i \leq k$, we have $\phi(x) \in [v]_0^+ \cap X$ and $\phi^{-1}(x) \in [v]_0^+ \cap X$.

To illustrate the usefulness of this notion, we note the following lemma.

LEMMA 3. *Let* $R \in \mathbb{N} \cup \{0\}$ *be fixed and suppose* $\varphi, \psi \in \mathrm{Aut}_R(X)$. *Suppose* $w \in \mathcal{L}(X)$ *extends uniquely* $2R$ *times to the right and left and let* $\tilde{w} \in \mathcal{L}(X)$ *be the unique word obtained by this extension. If* $\varphi(\tilde{w}) = \psi(\tilde{w})$, *then* $\varphi^{-1} \circ \psi$ *preserves occurrences of* $\tilde{w}$.

Proof. Since w extends uniquely $2R$ times to both sides (to $\tilde{w}$), it suffices to show that $\varphi^{-1} \circ \psi$ preserves occurrences of w. Let $x \in [w]_0^+ \cap X$. By assumption, $\mathcal{W}_{2R}^w(x) = \tilde{w}$. Since φ and ψ are block codes of range R,

$$((\varphi x)_{-R}, \dots, (\varphi x)_{|w|+R-1}) = \varphi(\tilde{w}) = \psi(\tilde{w}) = ((\psi x)_{-R}, \dots, (\psi x)_{|w|+R-1}).$$

Since φ^{-1} is a block code of range R and $\varphi^{-1}(\varphi x) = x \in [w]_0^+$, we have

$$((\varphi^{-1} y)_0, \dots, (\varphi^{-1} y)_{|w|}) = w$$

for any $y \in X$ satisfying

$$(y_{-R}, \ldots, y_{|w|+R-1}) = \varphi(\tilde{w}).$$

Since $y := \psi x$ satisfies this condition, it follows that $(\varphi^{-1} \circ \psi)(x) \in [w]_0^+$. A similar argument, with the roles of φ and ψ reversed, shows that $(\psi^{-1} \circ \varphi)(x) \in [w]_0^+$.

Therefore we have $(\varphi^{-1} \circ \psi)([w]_0^+ \cap X) \subseteq [w]_0^+ \cap X$ and $(\varphi^{-1} \circ \psi)^{-1}([w]_0^+ \cap X) \subseteq [w]_0^+ \cap X$. □

Let (X, σ) be a subshift and suppose $w \in \mathcal{L}(X)$. Let $X(w) \subset X$ denote the (possibly empty) subshift of X obtained by forbidding the word w:

$$X(w) := \{x \in X : \sigma^j x \notin [w]_0^+ \text{ for all } j \in \mathbb{Z}\}.$$

2.4. *Amenability.* If G is a group and $F \subset G$, let $|F|$ denote the cardinality of the set F, and for $g \in G$ the set gF is defined to be the set $\{gf : f \in F\}$. A discrete, countable group G is *amenable* if there exists a sequence $(F_k)_{k\in\mathbb{N}}$ of finite subsets of G such that every $g \in G$ is contained in all but finitely many F_k and such that

$$\lim_{k\to\infty} \frac{|F_k \Delta g F_k|}{|F_k|} = 0$$

for all $g \in G$. In this case, the sequence $(F_k)_{k\in\mathbb{N}}$ is called a *Følner sequence for G*.

3. *Technical lemmas*

We start with a bound on the complexity for the subshift obtained by forbidding the occurrences of some word.

LEMMA 4. *Suppose (X, σ) is a subshift and $w \in \mathcal{L}(X)$. If the cylinder set $[w]_0^+ \cap X$ contains at least one aperiodic point, then, for all $n \geq |w|$, we have*

$$P_{X(w)}(n) \leq P_X(n) - (n - |w| + 1).$$

Proof. Let $x \in [w]_0^+ \cap X$ be aperiodic and let $n \geq |w|$ be fixed. There are two cases to consider.

Case 1. Assume that every word of length n that occurs in x contains w as a subword. By the Morse–Hedlund theorem and aperiodicity of x, there are at least $n + 1$ distinct words of length n that occur in x. Since none of these words are in the language of $X(w)$, we have $P_{X(w)}(n) \leq P_X(n) - (n + 1)$, which gives the inequality in the statement.

Case 2. Assume there is a word of length n that occurs in x that does not contain w as a subword. Without loss of generality, we can assume that $x \in [w]_0^+$, and that either $(x_1, \ldots, x_n)$ or $(x_{|w|-n-1}, \ldots, x_{|w|-2})$ does not contain w as a subword (otherwise we replace x with an appropriate shift of itself). First suppose that $(x_1, \ldots, x_n)$ does not contain w as a subword and, for each $y \in X$, let $\mathcal{W}(y)$ denote the unique $v \in \mathcal{L}_n(X)$ such that $y \in [v]_0^+$. Then our assumption is that w is the leftmost subword (of length $|w|$) in $\mathcal{W}(x)$ and w is not a subword of $\mathcal{W}(\sigma x)$. Therefore each of the words $\mathcal{W}(x)$, $\mathcal{W}(\sigma^{-1}x)$, $\mathcal{W}(\sigma^{-2}x)$, $\ldots$, $\mathcal{W}(\sigma^{|w|-n}x)$ contains w as a subword, and the rightmost occurrence of w as a subword of $\mathcal{W}(\sigma^{-i}x)$ begins at the ith letter, for each $0 \leq i < n - |w| + 1$. It follows

that these words are all distinct and none of them are words in the language of $X(w)$. Therefore $P_{X(w)}(n) \leq P_X(n) - (n - |w| + 1)$ in this case. On the other hand, if w is not a subword of $(x_{|w|-n-1}, \ldots, x_{|w|-2})$, the argument is similar with the roles played by left and right reversed. □

We use this to bound the maximal length of a descending chain of subshifts obtained by forbidding a word at each step.

LEMMA 5. *Suppose (X, σ) is a subshift and $L \in \mathbb{N}$ is fixed. Let*

$$X =: X_0 \supset X_1 \supset X_2 \supset \cdots \supset X_k$$

be a descending chain of non-empty subshifts, where for $0 \leq i < k$ we have $X_{i+1} = X_i(w_i)$ for some word $w_i \in \mathcal{L}(X_i)$ with $|w_i| \leq L$. Further, suppose that for each $i = 0, 1, \ldots, k-1$ the cylinder set $[w_i]_0^+ \cap X_i$ contains at least one aperiodic point. Then

$$k < P_X(2L-1)/L.$$

Proof. By inductively applying Lemma 4 and using the fact that $|w_i| \leq L$ for all i, we have $P_{X_k}(2L-1) \leq P_X(2L-1) - kL$. If $k \geq P_X(2L-1)/L$, then $P_{X_k}(2L-1) \leq 0$. However, this is impossible, as a non-empty subshift has at least one word of every length. □

The proof of the following lemma is a straightforward modification of the proof of Lemma 3.1 in [**5**].

LEMMA 6. *Let $d, N \in \mathbb{N}$ with $d > 1$ and $N > 2d/\log 2$. Suppose (X, σ) is a subshift such that $P_X(n) \leq n^d$ for all $n \geq N$. Define k_n to be*

$$\min\{k \in \mathbb{N} : \textit{no word } w \in \mathcal{L}_n(X) \textit{ extends uniquely } k \textit{ times to the right and left}\}.$$

Then, for all integers $n \geq N$, there exists $m \leq n \log n$ satisfying $k_m \geq Cn$ where $C = \log(2)/2d$.

Proof. Suppose $P_X(n) \leq n^d$ for all $n \geq N$. For contradiction, suppose that there exists $n \geq N$ such that $k_m < Cn$ for all $m \leq n \log n$. Then since every word of length s can be extended in at least two distinct ways to a word of length $s + 2k_s$ (by adding k_s letters to each side), we have that

$$P_X(s + 2k_s) \geq 2P_X(s).$$

Therefore the assumption that $k_m < Cn$ for all $m \leq n \log n$ and the fact that $n \geq N > 2d/\log 2$ imply that

$$P_X(\lfloor n \log n \rfloor) \geq 2^{\lfloor n \log n \rfloor / Cn} > 2^{(n \log n - 1)/Cn} = \frac{n^{(\log 2)/C}}{2^{1/Cn}} = \frac{n^{(\log 2)/C}}{2^{2d/n \log 2}} > \frac{n^{2d}}{2}.$$

But this contradicts the fact that $P_X(n \log n) \leq (n \log n)^d < n^{2d}/2$. □

Suppose $Y \subseteq X$ are two subshifts where Y is obtained by forbidding a finite number of words from the language of X. If $u \in \mathcal{L}(Y)$ and u extends uniquely T times to the right and left (as a word in the language of Y), it might not extend uniquely T times to the right

and left when thought of as a word in the language of X. The following lemma resolves this issue, showing that if u appears in some element of X and is sufficiently far from any occurrence of the forbidden words, then u behaves as though it occurs in the language of Y.

LEMMA 7. *Let (X, σ) be a subshift and let $w_1, \ldots, w_{k-1} \in \mathcal{L}(X)$. Suppose*

$$Y = \{x \in X : \sigma^i(x) \notin [w_j]_0^+ \text{ for any } i \in \mathbb{Z} \text{ and } j \in \{1, 2, \ldots, k-1\}\}$$

and suppose $u \in \mathcal{L}(Y)$ is a word which extends uniquely (in $\mathcal{L}(Y)$) at least T times to the right and left. Let $v \in \mathcal{L}_{|u|+2T}(Y)$ be the unique word such that u is obtained by removing the rightmost and leftmost T letters from v. Then there exists $D \in \mathbb{N}$ such that, for any $x \in [u]_0^+ \cap X$, if $\sigma^i x \notin [w_j]_0^+$ for any $-D \le i \le D$ and $1 \le j < k$, then $\sigma^{-T} x \in [v]_0^+$.

Proof. For contradiction, suppose not. For each $D \in \mathbb{N}$, choose $x_D \in [u]_0^+ \cap X$ such that $\sigma^{-T} x_D \notin [v]_0^+$ and $\sigma^i x_D \notin [w_j]_0^+$ for any $-D \le i \le D$ and $1 \le j < k$. Let $x \in X$ be a limit point of $\{x_D : D \in \mathbb{N}\}$. Then $x \in [u]_0^+$ and $\sigma^i x \notin [w_j]_0^+$ for any $i \in \mathbb{Z}$ and $1 \le j < k$. Thus $x \in Y$, and since u extends uniquely T times to the right and left (as a subword of Y) we have $\sigma^{-T} x \in [v]_0^+$. This contradicts the fact that $\sigma^{-T} x_D \notin [v]_0^+$ for all $D \in \mathbb{N}$ and x is a limit point of $\{x_D : D \in \mathbb{N}\}$. □

We now generalize Lemma 3 to a form which is more useful in our setting.

LEMMA 8. *Let $R \in \mathbb{N}$ be fixed and suppose (X, σ) is a subshift. Let*

$$X =: X_0 \supset X_1 \supset X_2 \supset \cdots \supset X_k$$

be a descending chain of non-empty subshifts, where for $0 \le i < k$ there exists $w_i \in \mathcal{L}(X_i)$ such that w_i extends uniquely at least $2R$ times to the right and left (as a word in $\mathcal{L}(X_i)$), and $X_{i+1} = X_i(w_i)$. Suppose, further, that there exists $w_k \in \mathcal{L}(X_k)$ that extends uniquely at least $2R$ times to the right and left but for which $X_k(w_k) = \varnothing$. Let $\tilde{w}_i \in \mathcal{L}(X_i)$ be the unique word of length $|w_i| + 4R$ obtained by extending w_i by $2R$ letters on each side. If $\varphi, \psi \in \mathrm{Aut}_R(X)$ are such that $\varphi(\tilde{w}_i) = \psi(\tilde{w}_i)$ for all $i = 0, 1, \ldots, k$, then there exists D such that for all $i = 0, 1, 2, \ldots, k$ the automorphism $(\varphi^{-1} \circ \psi)$ preserves all occurrences of $\tilde{w}_0$ and preserves all occurrences of $\tilde{w}_i$ that occur at least D units from $\tilde{w}_0, \tilde{w}_1, \ldots, \tilde{w}_{i-1}$.

Proof. Fix $\varphi, \psi \in \mathrm{Aut}_R(X)$. First we show that $\varphi^{-1} \circ \psi$ preserves occurrences of $\tilde{w}_0$. Let $x \in [\tilde{w}_0]_0^+ \cap X$. Since $\psi(\tilde{w}_0) = \varphi(\tilde{w}_0)$, it follows that $\psi(x) \in [\varphi(\tilde{w}_0)]_0^+ \cap X$. Note that $\varphi(\tilde{w}_0)$ is $2R$ letters shorter than $\tilde{w}_0$, since the block code defining φ has range R. Since $\varphi^{-1}([\varphi(\tilde{w}_0)]_0^+ \cap X) \subset [w_0]_0^+ \cap X$, we have that $(\varphi^{-1} \circ \psi)(x) \in [w_0]_0^+ \cap X$. But $[w_0]_0^+ \cap X = \sigma^{-R}([\tilde{w}_0]_0^+) \cap X$ since w_0 extends uniquely $2R$ times to the right and left. Since $x \in [\tilde{w}_0]_0^+ \cap X$ was arbitrary, we have $(\varphi^{-1} \circ \psi)([\tilde{w}_0]_0^+ \cap X) \subset [\tilde{w}_0]_0^+ \cap X$. A similar argument, with the roles of φ and ψ reversed, shows that $(\psi^{-1} \circ \varphi)([\tilde{w}_0]_0^+ \cap X) \subset [\tilde{w}_0]_0^+ \cap X$. Therefore $\varphi^{-1} \circ \psi$ preserves occurrences of $\tilde{w}_0$.

For the second statement, we proceed by induction. Assume that we have shown that $\varphi^{-1} \circ \psi$ and $\psi^{-1} \circ \varphi$ preserve occurrences of $\tilde{w}_0$ and there exists D_k such that, for all $i < k$, the automorphisms $\varphi^{-1} \circ \psi$ and $\psi^{-1} \circ \varphi$ preserve occurrences of $\tilde{w}_i$ that occur

at least D_k units from $\tilde{w}_0, \tilde{w}_1, \ldots, \tilde{w}_{i-1}$. We show that there exists $D_{k+1} \geq D_k$ such that $\varphi^{-1} \circ \psi$ and $\psi^{-1} \circ \varphi$ also preserve occurrences of $\tilde{w}_k$ that occur at least D_{k+1} units from $\tilde{w}_0, \ldots, \tilde{w}_{k-1}$. By Lemma 7 applied to the subshifts X and $Y := X_k$, there exists $\tilde{D}_{k+1} \geq D_k$ such that, for any $x \in [\tilde{w}_k]_0^+ \cap X$, if $\sigma^i x \notin [\tilde{w}_j]_0^+$ for any $-\tilde{D}_{k+1} \leq i \leq \tilde{D}_{k+1}$ and $1 \leq j < k$, then $\sigma^{-2R} x \in [\tilde{w}_k]_0^+$. Define $D_{k+1} := (k+1)\tilde{D}_{k+1}$ and let $x \in [\tilde{w}_k]_0^+ \cap X$ be such that $\sigma^i x \notin [\tilde{w}_j]_0^+$ for any $-D_{k+1} \leq i \leq D_{k+1}$ and $1 \leq j < k$. Define

$$\mathcal{P}_0 := \{p \in \mathbb{Z} : \sigma^p x \in [\tilde{w}_0]_0^+\}.$$

Since $\varphi^{-1} \circ \psi$ and $\psi^{-1} \circ \varphi$ both preserve occurrences of $\tilde{w}_0$, observe that $\mathcal{P}_0$ is equal to the set $\{p \in \mathbb{Z} : \sigma^p(\varphi^{-1} \circ \psi)x \in [\tilde{w}_0]_0^+\}$ (in other words, occurrences of $\tilde{w}_0$ can neither be created nor destroyed by applying $\varphi^{-1} \circ \psi$ to x). Next define

$$\mathcal{P}_1 := \{p \in \mathbb{Z} : \sigma^p x \in [\tilde{w}_1]_0^+\}.$$

Since $\varphi^{-1} \circ \psi$ and $\psi^{-1} \circ \varphi$ both preserve occurrences of $\tilde{w}_1$ when they occur at least D_k units from $\tilde{w}_0$, it follows that for any $t \in \mathbb{N}$ any element of

$$\mathcal{P}_1 \triangle \{p \in \mathbb{Z} : \sigma^p(\varphi^{-1} \circ \psi)^t(x) \in [\tilde{w}_1]_0^+\}$$

is within distance D_k of an element of $\mathcal{P}_0$. Further, defining for each $1 < i < k$ the set

$$\mathcal{P}_i := \{p \in \mathbb{Z} : \sigma^p x \in [\tilde{w}_i]_0^+\},$$

it follows by induction that, for any $t \in \mathbb{N}$, any element of

$$\mathcal{P}_i \triangle \{p \in \mathbb{Z} : \sigma^p(\varphi^{-1} \circ \psi)^t(x) \in [\tilde{w}_i]_0^+\}$$

lies either within distance iD_k of an element of $\mathcal{P}_0$, within distance $(i-1)D_k$ of an element of $\mathcal{P}_1, \ldots,$ or within distance D_k of an element of $\mathcal{P}_{k-1}$. Recall that, by assumption, the set

$$\left\{p \in \mathbb{Z} : \sigma^p x \in \bigcup_{i=0}^{k-1} [\tilde{w}_i]_0^+\right\}$$

does not contain any element within distance D_{k+1} of the origin. But $\tilde{D}_{k+1} \leq D_{k+1} - kD_k$, and so, for any $t \in \mathbb{N}$, the set

$$\left\{p \in \mathbb{Z} : \sigma^p(\varphi^{-1} \circ \psi)^t(x) \in \bigcup_{i=0}^{k-1} [\tilde{w}_i]_0^+\right\}$$

does not contain any element within distance $\tilde{D}_{k+1}$ of the origin. However, $x \in [\tilde{w}_k]_0^+$ and so, as previously, we have $(\varphi^{-1} \circ \psi)x$, $(\psi^{-1} \circ \varphi)x \in [w_k]_0^+$. Since this occurrence of w_k (in the element x) is at least $\tilde{D}_{k+1}$ units from any occurrence of $\tilde{w}_0, \ldots, \tilde{w}_{k-1}$, we have $(\varphi^{-1} \circ \psi)x$, $(\psi^{-1} \circ \varphi)x \in [\tilde{w}_k]_0^+$.

Finally, note that, although the parameter $\tilde{D}_k$ may depend on the automorphisms φ and ψ, any parameter $D \geq \tilde{D}_k$ also suffices to reach the conclusion on the lemma for these automorphisms. Since $\mathrm{Aut}_R(X)$ is finite, we can take D to be the largest of the parameters $\tilde{D}_k$ that are constructed as $\varphi, \psi \in \mathrm{Aut}_R(X)$ vary over all possible combinations. □

The following lemma allows us to adapt techniques from [**5**] which relied on the fact that in a minimal shift all words occur syndetically.

LEMMA 9. *Suppose (X, σ) is a subshift. Let*

$$X =: X_0 \supset X_1 \supset X_2 \supset \cdots \supset X_k$$

be a descending chain of non-empty subshifts where for each $0 \le i < k$ there exists $w_i \in \mathcal{L}(X_i)$ such that w_i extends uniquely at least T times to the right and left (as a word in $\mathcal{L}(X_i)$), and $X_{i+1} = X_i(w_i)$. Suppose, further, that there exists $w_k \in \mathcal{L}(X_k)$ that extends uniquely at least T times to the right and left but for which $X_k(w_k) = \varnothing$. Let $\tilde{w}_i \in \mathcal{L}(X_i)$ be the unique word of length $|w_i| + 2T$ obtained by extending the word w_i by T letters on each side. Finally, for each $i = 1, 2, \ldots, k$, let $D \in \mathbb{N}$ be the constant obtained from Lemma 8. Then there exists $G \in \mathbb{N}$ such that, for any $x \in X$, the set

$$\begin{aligned}\mathcal{S}_x := \{j \in \mathbb{Z} : \sigma^j x \in [\tilde{w}_0]_0^+\} \cup \{j \in \mathbb{Z} : \sigma^j x \in [\tilde{w}_i]_0^+ \text{ for some } 1 \le i \le k \\ \text{and } \sigma^s x \notin [\tilde{w}_t]_0^+ \text{ for any } t < i \text{ and any } j - kD \le s \le j + kD + |\tilde{w}_i| - 1\}\end{aligned}$$

is syndetic with gap at most G.

Proof. If not, then for each $G \in \mathbb{N}$ there exists $x_G \in X$ such that $\mathcal{S}_x$ is not syndetic with gap less than G. Without loss (shifting x if necessary), we can assume that $\{-\lfloor G/2 \rfloor, \ldots, 0, \ldots, \lfloor G \rfloor/2\} \cap \mathcal{S}_x = \varnothing$. Since X is compact, we can pass to a subsequence of $(x_G)_{G \in \mathbb{N}}$ converging to some $y \in X$. Then $\mathcal{S}_y = \varnothing$. Therefore $y \in X_k(w_k)$, contradicting the assumption that $X_k(w_k) = \varnothing$. □

We use this to describe the set of automorphisms preserving occurrences of the sequence of words.

LEMMA 10. *Assume (X, σ) is a subshift and $D \in \mathbb{N}$. Let $\mathcal{W} = \{\tilde{w}_0, \tilde{w}_1, \ldots, \tilde{w}_k\} \subset \mathcal{L}(X)$ be a finite set of words for which there exists $G \in \mathbb{N}$ such that, for any $x \in X$, the set*

$$\begin{aligned}\mathcal{S}_x := \{j \in \mathbb{Z} : \sigma^j x \in [\tilde{w}_0]_0^+\} \cup \{j \in \mathbb{Z} : \sigma^j x \in [\tilde{w}_i]_0^+ \text{ for some } 1 \le i \le k \\ \text{but } \sigma^s x \notin [\tilde{w}_t]_0^+ \text{ for any } t < i \text{ and any } j - kD \le s \le j + kD\}\end{aligned}$$

is syndetic with gap at most K. Let $R \in \mathbb{N}$ be such that $|\tilde{w}_i| > 2R$ for all $\tilde{w}_i \in \mathcal{W}$ and define

$$\begin{aligned}\mathcal{H} := \{\varphi \in \mathrm{Aut}_R(X) : \varphi \text{ preserves occurrences of } \tilde{w}_0 \text{ and occurrences of} \\ \tilde{w}_i \text{ that occur at least } D \text{ units from } \tilde{w}_0, \tilde{w}_1, \ldots, \tilde{w}_{i-1} \text{ for all } 1 \le i \le k\}.\end{aligned}$$

Then $\langle \mathcal{H} \rangle$, the subgroup of $\mathrm{Aut}(X)$ generated by $\mathcal{H}$, is finite.

Proof. For any $x \in X$, define

$$\mathcal{S}_0(x) := \{j \in \mathbb{Z} : \sigma^j x \in [\tilde{w}_0]_0^+\}.$$

Note that, since any element of $\mathcal{H}$ preserves occurrences of $\tilde{w}_0$, we have

$$\mathcal{S}_0(x) = \mathcal{S}_0(\varphi x) \quad \text{for any } x \in X \text{ and } \varphi \in \mathcal{H}.$$

Since every element of $\langle \mathcal{H} \rangle$ is a product of elements of $\mathcal{H}$, it follows that $\mathcal{S}_0(x) = \mathcal{S}_0(\psi x)$ for any $x \in X$ and $\psi \in \langle \mathcal{H} \rangle$.

For any $x \in X$ and $1 \le i \le k$, define

$$\mathcal{S}_i(x) := \{j \in \mathbb{Z} : \sigma^j x \in [\tilde{w}_i]_0^+ \text{ but for any } t < i \text{ and any } j - (i-t)D \le s \le j + (i-t)D \text{ we have } \sigma^s x \notin [\tilde{w}_t]_0^+\}.$$

We claim that $\mathcal{S}_i(x) = \mathcal{S}_i(\varphi x)$ for any $x \in X$, $\varphi \in \mathcal{H}$, and $i \le k$. We proceed by induction on i (the case $i = 0$ having already been shown). Suppose $i < k$ and we have $\mathcal{S}_m(x) = \mathcal{S}_m(\varphi x)$ for any $x \in X$, $\varphi \in \mathcal{H}$, and $0 \le m \le i$. Fix $x \in X$ and $\varphi \in \mathcal{H}$. We first show that $\mathcal{S}_{i+1}(x) \subseteq \mathcal{S}_{i+1}(\varphi x)$. If $j \in \mathcal{S}_{i+1}(x)$, then $\sigma^j x \in [\tilde{w}_{i+1}]_0^+$ and, for any $j - D \le s \le j + D$ and $t \le i$, we have $\sigma^s x \notin [\tilde{w}_t]_0^+$. Since φ preserves occurrences of $\tilde{w}_{i+1}$ that occur at least D units from $\tilde{w}_0, \ldots, \tilde{w}_i$, we have $\sigma^j(\varphi x) \in [\tilde{w}_{i+1}]_0^+$. Since φ preserves occurrences of $\tilde{w}_0$, we also have $\sigma^s(\varphi x) \notin [\tilde{w}_0]_0^+$ for any $j - (i+1)D \le s \le j + (i+1)D$, by definition of $\mathcal{S}_{i+1}(x)$. We want to show that $j \in \mathcal{S}_{i+1}(\varphi x)$. For contradiction, suppose there were some $1 \le t \le i$ and some $j - (i+1-t)D \le s \le j + (i+1-t)D$ for which $\sigma^s(\varphi x) \in [\tilde{w}_t]_0^+$. Then $\sigma^s(x) \notin [\tilde{w}_t]_0^+$ because $j \in \mathcal{S}_{i+1}(x)$. But since φ preserves all occurrences of $\tilde{w}_t$ that occur at least D units from $\tilde{w}_0, \ldots, \tilde{w}_{t-1}$, there must be some $0 \le u < t$ such that s is within D units of an occurrence in x of $\tilde{w}_u$ (otherwise $\tilde{w}_t$ could not appear at location s in φx). Since s is within $(i+1-t)D$ units of j, this occurrence of $\tilde{w}_u$ in x must be within $(i+1-t+1)D \le (i+1-u)D$ units of j. But this contradicts the fact that $j \in \mathcal{S}_{i+1}(x)$. Therefore no such t exists and so $\sigma^s(\varphi x) \notin [\tilde{w}_t]_0^+$ for all $t \le i$ and all $j - (i+1-t)D \le s \le j + (i+1-t)D$. This means that $j \in \mathcal{S}_{i+1}(\varphi x)$. It follows that $\mathcal{S}_{i+1}(\varphi x) \subseteq \mathcal{S}_{i+1}(x)$. The reverse containment is proven similarly using the point $y = \varphi x$ and the automorphism $\varphi^{-1} \in \mathcal{H}$. Thus $\mathcal{S}_{i+1}(x) = \mathcal{S}_{i+1}(\varphi x)$. By induction, we have $\mathcal{S}_i(x) = \mathcal{S}_i(\varphi x)$ for any $x \in X$, $\varphi \in \mathcal{H}$, and $i \le k$, as claimed.

Now for $x \in X$, define

$$\mathcal{T}_x := \bigcup_{i=0}^{k} \mathcal{S}_i(x).$$

We have shown that for any $\varphi \in \mathcal{H}$ we have $\mathcal{T}_{\varphi x} = \mathcal{T}_x$. Since any element of $\langle \mathcal{H} \rangle$ is a product of elements of $\mathcal{H}$, this result holds for any $\varphi \in \langle \mathcal{H} \rangle$ as well. Furthermore, note that $\mathcal{S}_x \subseteq \mathcal{T}_x$ and so by assumption there exists $K \ge 1$ such that $\mathcal{T}_x$ is syndetic with gap at most K for all $x \in X$. Fix $x \in X$ and order the elements of $\mathcal{T}_x$ as

$$\cdots < j_{-2} < j_{-1} < j_0 < j_1 < j_2 < \cdots .$$

For any $n \in \mathbb{Z}$, let $m_1(n), m_2(n) \in \{0, 1, \ldots, k\}$ be such that $\sigma^{j_n} x \in [\tilde{w}_{m_1(n)}]_0^+$ and $\sigma^{j_{n+1}} x \in [\tilde{w}_{m_2(n)}]_0^+$. Let $u \in \mathcal{L}$ be a word of length at most K such that $\sigma^{j_n} x \in [\tilde{w}_{m_1(n)} u \tilde{w}_{m_2(n)}]_0^+$. If $\varphi \in \mathcal{H}$, then since $\mathcal{T}_{\varphi x} = \mathcal{T}_x$ we have $\sigma^{j_n}(\varphi x) \in [\tilde{w}_{m_1(n)} v \tilde{w}_{m_2(n)}]_0^+$ for some $v \in \mathcal{L}(X)$ such that $|v| = |u|$. Moreover, since $|\tilde{w}_i| > 2R$ and $\varphi \in \mathrm{Aut}_R(X)$, the block code defining φ has the property that, for any $y \in X$, if $m_1(n), m_2(n) \in \mathcal{T}_y$ (and are consecutive in $\mathcal{T}_y$) and if $\sigma^{j_n} y \in [\tilde{w}_{m_1(n)} u \tilde{w}_{m_2(n)}]_0^+$, then $\sigma^{j_n}(\varphi y) \in [\tilde{w}_{m_1(n)} v \tilde{w}_{m_2(n)}]_0^+$ (in other words, the fact that the word u turned into the word v upon application of φ is independent of y). Since any element of $\langle \mathcal{H} \rangle$ is a product of elements of $\mathcal{H}$, this same result holds for any $\varphi \in \langle \mathcal{H} \rangle$.

By assumption, any element $x \in X$ can be decomposed as

$$x = \cdots \tilde{w}_{-2} u_{-2} \tilde{w}_{-1} u_{-1} \tilde{w}_0 u_0 \tilde{w}_1 u_1 \tilde{w}_2 u_2 \cdots ,$$

where $|u_i| \le K$ and $\tilde{w}_i$ begins at a location in $\mathcal{T}_x$ for all $i \in \mathbb{Z}$. Therefore if $\varphi \in \langle \mathcal{H} \rangle$, then φ is determined entirely by which word v_i it turns u_i into (recall $|v_i| = |u_i| \le K$ and the mapping $u_i \mapsto v_i$ is independent of x). As there are only finitely many permutations of the set of words of length at most K in $\mathcal{L}(X)$, $\langle \mathcal{H} \rangle$ has only finite many elements. □

Our final technical lemma quantifies a property of functions which grow subexponentially.

LEMMA 11. *Let $g : \mathbb{N} \to \mathbb{N}$ be such that $\log g(n) = o(n)$. For any $k \in \mathbb{N}$ and all sufficiently small $\varepsilon > 0$, there exists $M \in \mathbb{N}$ such that if $N \ge M$ and if $f : \{0, 1, \ldots, N\} \to \mathbb{N}$ is a non-decreasing function satisfying $f(N) \le g(N)$, then there exists $x \in \{0, 1, \ldots, N - k\}$ such that*

$$\frac{f(x+k) - f(x)}{f(x)} < \varepsilon.$$

Proof. Let $k \in \mathbb{N}$ and $\varepsilon > 0$ be fixed. Find $M \in \mathbb{N}$ such that for all $N \ge M$ we have

$$\frac{\log g(N)}{N} < \varepsilon/4k. \tag{1}$$

Without loss of generality, we can assume that $M > k$. Let $N \ge M$ and let $f : \{0, 1, \ldots, N\} \to \mathbb{N}$ be non-decreasing. Suppose that for all $0 \le x \le N - k$ we have

$$\frac{f(x+k) - f(x)}{f(x)} \ge \varepsilon.$$

Then by induction, $f(nk) \ge (1+\varepsilon)^n f(0)$ for all $0 \le n < \lfloor N/k \rfloor$. In particular, since f is non-decreasing,

$$f(N) \ge (1+\varepsilon)^n f(0)$$

where $n = \lfloor N/k \rfloor$. Therefore

$$\log g(N) \ge \log f(N) \ge n \log(1+\varepsilon) + \log f(0) \ge \frac{N}{2k} \log(1+\varepsilon) > \frac{N\varepsilon}{3k}$$

for all sufficiently small ε, a contradiction of (1). □

4. *Amenability of* Aut(X)

Our goal in this section is to prove Theorem 1. We do this first with the added assumption that (X, σ) has dense aperiodic points.

THEOREM 12. *Let X be a subshift with dense aperiodic points and suppose $P_X(n) = o(n^2/\log^2 n)$. Then* Aut(X) *is amenable (as a countable discrete group).*

Proof. Let C be as in Lemma 6. Fix $R \in \mathbb{N}$ sufficiently large such that $P_X(n) \le n^2$ for all $n \ge 2R/C$. Define $X_0 = X$ and, by Lemma 6, choose a word $w_0 \in \mathcal{L}(X_0)$ such that $|w_0| \le (2R/C)\log(2R/C)$ and w_0 extends uniquely at least $2R$ times to the right and to the left. Define $\tilde{w}_0$ to be the (unique) extension of w_0 exactly $2R$ times to each side and set $X_1 := X_0(\tilde{w}_0)$. Continue this process inductively: once we have constructed the non-empty subshift X_i, apply Lemma 6 to find a word $w_i \in \mathcal{L}(X_i)$ such that $|w_i| \le (2R/C)\log(2R/C)$ and such that w_i extends at least $2R$ times to the right and

to the left (note that $P_{X_i}(n) \leq P_{X_j}(n)$ for all n, whenever $i > j$, and so the parameter N appearing in Lemma 6 can be taken to be $2R/C$ for all of the shifts we construct). Define $\tilde{w}_i$ to be the (unique) extension (in $\mathcal{L}(X_i)$) of w_i exactly $2R$ times to each side and set $X_{i+1} := X_i(\tilde{w}_i)$. If X_{i+1} is empty, the process ends. By Lemma 5 this process ends after at most s_R steps, where $s_R \leq P_X(2L-1)/L$ and $L = \lfloor 2R/C \log 2R/C \rfloor$. It follows that, as a function of R, $s_R = o(R/\log R)$. To summarize, for any $R \in \mathbb{N}$ we have constructed a sequence of non-empty subshifts

$$X := X_0 \supset X_1 \supset X_2 \supset \cdots \supset X_{s_R}$$

such that for each $i = 0, 1, \ldots, k_R - 1$ there exists $w_i \in \mathcal{L}(X_i)$ that extends uniquely (in $\mathcal{L}(X_i)$) at least $2R$ times to each side and is such that $|w_i| \leq (2R/C)\log(2R/C)$ and $X_{i+1} = X_i(\tilde{w}_i)$ (where $\tilde{w}_i$ is the extended version of w_i). Note that for all sufficiently large R, $|\tilde{w}_i| \leq R^2$. For fixed $R \in \mathbb{N}$, let $\mathcal{W}_R = \{\tilde{w}_0, \ldots, \tilde{w}_{s_R}\}$ and let D be the parameter coming from Lemma 8. Let $\mathcal{G}_R \subset \mathrm{Aut}(X)$ be the subgroup of automorphisms generated by

$$\{\varphi \in \mathrm{Aut}_R(X) : \varphi \text{ preserves occurrences of } \tilde{w}_0 \text{ and occurrences of } \tilde{w}_i \text{ that occur at least } D \text{ units from } \tilde{w}_0, \tilde{w}_1, \ldots, \tilde{w}_{i-1} \text{ for all } 1 \leq i \leq k\}.$$

By Lemma 9, there exists $G \in \mathbb{N}$ such that for any $x \in X$ the set

$$\mathcal{S}_x := \{j \in \mathbb{Z} : \sigma^j x \in [\tilde{w}_0]_0^+\} \cup \{j \in \mathbb{Z} : \sigma^j x \in [\tilde{w}_i]_0^+ \text{ for some } 1 \leq i \leq k \text{ but } \sigma^s x \notin [\tilde{w}_t]_0^+ \text{ for any } t < i \text{ and any } j - D \leq s \leq j + D + |\tilde{w}_i| - 1\}$$

is syndetic with gap at most G. By Lemma 10, $\mathcal{G}_R$ is finite. Furthermore, by the pigeonhole principle, if $\mathcal{S} \subset \mathrm{Aut}_R(X)$ is any set satisfying

$$|\mathcal{S}| > P_X(R^2)^{s_R},$$

then there exist $\varphi, \psi \in \mathcal{S}$ such that $\varphi(\tilde{w}_i) = \psi(\tilde{w}_i)$ for all $i = 0, 1, 2, \ldots, s_R$. In other words, if $\mathcal{S} \subset \mathrm{Aut}_R(X)$ is any set satisfying $|\mathcal{S}| > P_X(R^2)^{s_R}$, then there exist $\varphi, \psi \in \mathcal{S}$ such that $(\varphi^{-1} \circ \psi) \in \mathcal{G}_R$.

Since $P_X(R^2) \leq R^4/\log^2(R^2) \leq R^4$ for all sufficiently large R, we have

$$P_X(R^2)^{s_R} \leq (R^4)^{o(R/\log(R))} = e^{o(R)},$$

meaning that this grows subexponentially in R. Define $g : \mathbb{N} \to \mathbb{N}$ by $g(R) := P_X(R^2)^{s_R}$. Then, by Lemma 11, for any $k \in \mathbb{N}$ and any sufficiently small $\varepsilon > 0$, there exists M such for any $N \geq M$ and any non-decreasing function $f : \{0, 1, \ldots, N\} \to \mathbb{N}$, which satisfies $f(N) \leq g(N)$, there exists $0 \leq x \leq N - k$ such that

$$\frac{f(x+k) - f(x)}{f(x)} < \varepsilon.$$

We are now ready to prove that $\mathrm{Aut}(X)$ is amenable. Let $k \in \mathbb{N}$ be fixed. Choose $\varepsilon < 1/k$ sufficiently small such that Lemma 11 applies and let M be the constant obtained from this lemma. Choose $R > \max\{k, M\}$ large enough such that

$$\frac{g(R+k) - g(R)}{g(R)} < \frac{\varepsilon}{4k}.$$

Let $f : \{0, 1, 2 \ldots, R\} \to \mathbb{N}$ be the function

$$f(n) := |\{(\varphi(\tilde{w}_0), \varphi(\tilde{w}_1), \ldots, \varphi(\tilde{w}_{k_R})) : \varphi \in \mathrm{Aut}_n(X) \hookrightarrow \mathrm{Aut}_R(X)\}|.$$

Here, for $n \le R$, the notation $\mathrm{Aut}_n(X) \hookrightarrow \mathrm{Aut}_R(X)$ refers to the embedded image of $\mathrm{Aut}_n(X)$ in $\mathrm{Aut}_R(X)$ obtained by using the natural identification of a block code of range n as a block code of range R. It follows that if $\mathcal{S} \subset \mathrm{Aut}_n(X)$ is any set containing more than $f(n)$ elements, then there exist $\varphi, \psi \in \mathrm{Aut}_n(X)$ such that $(\varphi^{-1} \circ \psi) \in \mathcal{G}_R$. In other words, $\mathrm{Aut}_n(X)$ can be covered by $f(n)$ cosets of $\mathcal{G}_R$.

By Lemma 11, there exists $0 \le n \le R - k$ such that

$$\frac{f(n+k) - f(n)}{f(n)} < \varepsilon/2.$$

Fix such an n and let $\varphi_1, \ldots, \varphi_{f(n)} \in \mathrm{Aut}_n(X)$ be representatives of distinct cosets of $\mathcal{G}_R$ and such that $\mathrm{Aut}_n(X)$ is contained in

$$F_k := \bigcup_{i=0}^{f(n)} \varphi_i \cdot \mathcal{G}_R.$$

Observe that F_k is finite and contains $\mathrm{Aut}_n(X)$. Now let $\varphi_{f(n)+1}, \ldots, \varphi_{f(n+k)} \in \mathrm{Aut}_{n+k}(X)$ be $f(n+k) - f(n)$ additional representatives of distinct cosets of $\mathcal{G}_R$ and such that $\mathrm{Aut}_{n+k}(X)$ is contained in

$$\tilde{F}_k := \bigcup_{i=0}^{f(n+k)} \varphi_i \cdot \mathcal{G}_R.$$

Observe that if $\psi \in \mathrm{Aut}_k(X)$ then for any $i = 0, 1, \ldots, f(n)$ we have $(\psi \circ \varphi_i) \in \mathrm{Aut}_{n+k}(X)$. Therefore

$$\psi \circ F_k \subset \tilde{F}_k$$

and since $F_k \subset \tilde{F}_k$, we have

$$\frac{|F_k \triangle (\psi \circ F_k)|}{|F_k|} \le \frac{2|\tilde{F}_k \setminus F_k|}{|F_k|} = \frac{2(f(n+k) - f(n))}{f(n)} < \varepsilon \le \frac{1}{k}.$$

Let $n_k \in \mathbb{N}$ be the constant n constructed above. Observe that $n_k \to \infty$ as $k \to \infty$.

Construct the set F_k for each $k \in \mathbb{N}$. We claim that $(F_k)_{k \in \mathbb{N}}$ is a Følner sequence in $\mathrm{Aut}(X)$. By construction, F_k is finite for each k and we have shown that

$$\frac{|F_k \triangle (\psi \circ F_k)|}{|F_k|} < \frac{1}{k}$$

for each $k \in \mathbb{N}$. Finally, since $n_k \to \infty$ as $k \to \infty$, we have that if $\psi \in \mathrm{Aut}(X)$ then $\{k : \psi \notin F_k\}$ is finite. Thus we have constructed a Følner sequence for $\mathrm{Aut}(X)$ and so it is amenable. □

We use this to complete the proof of Theorem 1.

Proof of Theorem 1. Let $Y \subset X$ be the closure of the aperiodic points in X. By Theorem 12, $\mathrm{Aut}(Y)$ is amenable. For any $\varphi \in \mathrm{Aut}(X)$, observe that $x \in X$ is aperiodic if and only if $\varphi(x)$ is aperiodic. Therefore for any $\varphi \in \mathrm{Aut}(X)$, we have $\varphi(Y) = Y$ and the

map $h : \mathrm{Aut}(X) \to \mathrm{Aut}(Y)$ given by $h(\varphi) := \varphi|_Y$ is a homomorphism. Since the image of h is a (closed) subgroup of $\mathrm{Aut}(Y)$, it is amenable. Thus to check that $\mathrm{Aut}(X)$ is amenable, it suffices to check that $\ker(h)$ is amenable.

To show this, it suffices to show that any finitely generated subgroup of $\ker(h)$ is amenable. Let $\varphi_1, \ldots, \varphi_m \in \ker(h)$. We claim that the set

$$\mathcal{S} := \{x \in X : \varphi_i(x) \neq x \text{ for at least one element of } \langle \varphi_1, \ldots, \varphi_m \rangle\}$$

is finite. For contradiction, suppose $\mathcal{S}$ is infinite. Choose $R \in \mathbb{N}$ such that $\varphi_1, \ldots, \varphi_m \in \mathrm{Aut}_R(X)$. By construction, if $w \in \mathcal{L}(Y)$ is a word of length R and if $\mathrm{Id} \in \mathrm{Aut}(Y)$ denotes the identity, then $\varphi_i(w) = \mathrm{Id}(w)$ (as a block map). Therefore if $x \in \mathcal{S}$, there exists $j \in \mathbb{Z}$ such that $\sigma^j x \in [u]_0^+$ for some word $u \in \mathcal{L}(X) \setminus \mathcal{L}(Y)$ of length R (otherwise x is comprised entirely of words of length R on which φ acts as the identity). For each $x \in \mathcal{S}$ choose a word $u_x \in \mathcal{L}_R(X) \setminus \mathcal{L}_R(Y)$ such that $\sigma^j x \in [u_x]_0^+$ for some $j \in \mathbb{Z}$. Since $\mathcal{L}_R(X) \setminus \mathcal{L}_R(Y)$ is finite, there exists $u \in \mathcal{L}_R(X) \setminus \mathcal{L}_R(Y)$ such that $u_x = u$ for infinitely many $x \in \mathcal{S}$. For each such x, let $j_x \in \mathbb{Z}$ be such that $\sigma^{j_x} x \in [u]_0^+$. Every infinite collection of points has an aperiodic limit point, and so there is some aperiodic $y \in [u]_0^+$ (see, for example, [**4**, Lemma 4.1]). This contradicts the fact that $u \notin \mathcal{L}(Y)$. Therefore $\mathcal{S}$ is finite, proving the claim.

Since the set $\mathcal{S}$ is finite, it follows that $\langle \varphi_1, \ldots, \varphi_m \rangle$ is finite (and hence amenable). □

COROLLARY 13. *Let X be a subshift and suppose there exists $\varepsilon > 0$ such that $P_X(n) = O(n^{2-\varepsilon})$. Then* $\mathrm{Aut}(X)$ *is amenable (as a countable discrete group).*

Proof. Any function which is $O(n^{2-\varepsilon})$ is also $o(n^2/\log^2 n)$, and so this follows immediately from Theorem 1. □

5. *Shifts of subquadratic growth*

Recall that if G is a group and $g_1, g_2 \in G$ then

$$\langle g_1, g_2 \rangle^+ := \{g_{i_1} g_{i_2} \cdots g_{i_j} : i_1, \ldots, i_j \in \{1, 2\}, k \in \mathbb{N} \cup \{0\}\}$$

is the *monoid generated by* g_1 *and* g_2 in G (where we interpret the case $k = 0$ to give the identity). This monoid is *free* if whenever

$$g_{i_1} g_{i_2} \cdots g_{i_j} = g_{k_1} g_{k_2} \cdots g_{k_l}$$

we have $j = l$ and $g_{i_n} = g_{k_n}$ for all $n = 1, \ldots, j$. If this monoid is free, it is said to have *rank* 2 because it is generated by two elements of G.

Let (X, σ) be a subshift. The *full group* of (X, σ), denoted $[\sigma]$, is the group of all maps $\varphi : X \to X$ such that there exists $k : X \to \mathbb{Z}$ such that $\varphi(x) = \sigma^{k(x)}(x)$ for all $x \in X$. The group $[\sigma] \cap \mathrm{Aut}(X)$ is the group of all *orbit-preserving* automorphisms. This is a normal, abelian subgroup of $\mathrm{Aut}(X)$.

We recall the statement of Theorem 2.

THEOREM 2. *Let (X, σ) be a subshift such that* $\liminf P_X(n)/n^2 = 0$. *Then* $\mathrm{Aut}(X)$ *does not contain a free monoid of rank* 2.

Our main tool to prove this theorem is the following rephrasing given in [**3**] of a result of Quas and Zamboni.

LEMMA 14. (Quas and Zamboni [**10**]) *Let $n, k \in \mathbb{N}$. Then there exists a finite set $F \subset \mathbb{Z}^2 \setminus \{(0, 0)\}$ (which depends on n and k) such that, for every $\eta \in \mathcal{A}^{\mathbb{Z}^2}$ satisfying $P_\eta(n, k) \leq nk/16$, there exists a vector $v \in F$ such that $\eta(x + v) = \eta(x)$ for all $x \in \mathbb{Z}^2$.*

We now adapt the technique developed in [**3**] to prove Theorem 2.

Proof of Theorem 2. Let $X \subset \mathcal{A}^{\mathbb{Z}}$ be a subshift and suppose $P_X(n) = o(n^2)$. For each $\varphi \in \mathrm{Aut}(X)$ and each $x \in X$, define $\eta_{\varphi,x} \in \mathcal{A}^{\mathbb{Z}^2}$ by setting $\eta_{\varphi,x}(i, j) := (\varphi^j \sigma^i)(x)$ (this is the space time of the system). Finally, let

$$Y_{\varphi,x} \subset \mathcal{A}^{\mathbb{Z}^2} := \overline{\{\eta_{\varphi,x} \circ S^j T^i : (i, j) \in \mathbb{Z}^2\}}$$

where $S, T : \mathbb{Z}^2 \to \mathbb{Z}^2$ are the vertical and horizontal shifts, respectively: $S(i, j) := (i, j + 1)$ and $T(i, j) := (i + 1, j)$. Since $P_X(n) = o(n^2)$, it follows from [**3**, Lemma 2.1] that $P_{Y_{\varphi,x}}(n, n) = o(n^2)$ (however, the rate at which $P_X(n)/n^2$ tends to zero depends on the range of φ).

For contradiction, suppose $\varphi, \psi \in \mathrm{Aut}(X)$ generate a free monoid. Let $Z \subset X$ be the closure of the aperiodic elements of X and, as already noted, Z is $\mathrm{Aut}(X)$-invariant. By Lemma 14, there exists a finite set $F \subset \mathbb{Z}^2 \setminus \{(0, 0)\}$ such that for any $x \in Z$ the maps $\eta_{\varphi,x}$ and $\eta_{\psi,x}$ are both periodic with some period vector in F. Note that if x is aperiodic, then $\eta_{\varphi,x}$ and $\eta_{\psi,x}$ cannot be horizontally periodic. Therefore there exists $M \in \mathbb{N}$ such that $\eta_{\varphi,x}$ and $\eta_{\psi,x}$ both have period vectors with y-coordinate M, for all aperiodic $x \in Z$. It follows that $\varphi^M x$ and $\psi^M x$ are both shifts of x. As this holds for all aperiodic $x \in X$, the restrictions of φ^M and ψ^M to Z are both elements of the abelian group $\mathrm{Aut}(Z) \cap [\sigma]$. In particular, the restriction of the commutator $[\varphi^M, \psi^M]$ to Z is the identity.

Suppose the range of $[\varphi^M, \psi^M]$ is R. If $w \in \mathcal{L}_R(X)$ is such that $[w]_0^+$ contains an aperiodic point, then $[\varphi^M, \psi^M]$ acts like the identity map (when thought of as a block code of range R) on w. It follows that if $[\varphi^M, \psi^M]$ does not act like the identity map on $[w]_0^+$, then $[w]_0^+ \cap X$ does not contain any aperiodic points. Furthermore, this means that $[w]_0^+ \cap X$ cannot contain periodic points of arbitrarily large period. Therefore, $X \setminus Z$ is finite and so there exists $k \in \mathbb{N}$ such that $[\varphi^{kM}, \psi^{kM}]$ is the identity on X, contradicting the fact that φ^{kM} and ψ^{kM} do not commute (since φ and ψ generate a free monoid). □

Remark 15. Note that our proof shows a slightly stronger result, namely that if $\varphi, \psi \in \mathrm{Aut}(X)$ then there exists $n > 0$ such that φ^n commutes with ψ^n.

While Theorem 2 applies to a larger class of shifts than Theorem 1, it does not conclude that $\mathrm{Aut}(X)$ amenable. Nevertheless, Theorem 2 does give algebraic information about $\mathrm{Aut}(X)$ (in particular, it cannot contain a non-abelian free subgroup) and a recent result of Salo and Schraudner shows that it is essentially optimal.

THEOREM 16. (Salo and Schraudner [**11**]) *There exists a subshift (X, σ) such that*

$$P_X(n) = (n + 1)^2$$

and is such that $\mathrm{Aut}(X)$ is amenable and contains a free monoid of rank 2.

The example that they construct is the Cartesian product of two copies of the subshift X on the alphabet $\{0, 1\}$ where each $x \in X$ contains at most one occurrence of 1.

We present here a second example (different from that of [**11**]) of a shift of quadratic growth whose automorphism group contains a free monoid of rank 2.

Let $\mathcal{A}$ be the eight-letter alphabet $\{0, 1, a, b, p, 1p, ap, bp\}$. We consider the subshift $X \subset \mathcal{A}^{\mathbb{Z}}$ consisting of the following:

(i) the coloring of all 0s;

(ii) any coloring which is all 0s except at a single location where it is one of 1, a, b, p, $1p$, ap, or bp;

(iii) any coloring which is all 0s except at two locations, one of which is p and the other of which is one of 1, a, or b.

We leave it to the reader to check that this does indeed form a (closed) subshift and that its complexity function grows quadratically. To show that this contains a free monoid, we define two automorphisms of X which we call φ_a and φ_b. These are block codes of range 1 and we claim they generate a free monoid of rank 2. Rather than define them on each block, we find the following description helpful. A person, p, walks down a bi-infinite hallway. At one location in the hallway there is a nail on the wall and there are two possible pictures, a and b, than can be hung on the nail. When the nail is unoccupied, its location is denoted 1. When it holds picture a, its location is a. When it holds picture b, its location is b. If the person is standing in front of the nail, the person/nail is denoted $1p$, ap, or bp (depending on the state of the nail). Now we can define our automorphisms. When φ_a is applied to an element of X it moves the person one space to the right. If this causes the person to be standing in front of the nail they take one of three actions:

- if the nail is unoccupied, the person hangs picture a on it;
- if the nail holds picture a, the person removes it and leaves the nail unoccupied;
- if the nail holds picture b, the person leaves the picture undisturbed.

We claim that these rules can be implemented by a block code of range 1 and that φ_a is invertible. Similarly, when φ_b is applied to an element of X it moves the person one space to the right. If this causes the person to be standing in front of the nail, the analogous rules (with the roles of a and b reversed) apply. This is also invertible and can be implemented by a block code of range 1. Note that φ_a and φ_b carry elements of X to elements of X. Finally, suppose $w = (w_1, w_2, \ldots, w_k) \in \{a, b\}^k$ and let $g \in \langle \varphi_a, \varphi_b \rangle^+$ be

$$g = g_1 g_2 \cdots g_k,$$

where $g_i \in \{\varphi_a, \varphi_b\}$ for each $1 \leq i \leq k$ is the automorphism corresponding to letter w_i. First we show how to find k by observing the action of g on X. For each $i \in \mathbb{N}$, let $x_i \in X$ be the configuration which has a 1 at the origin, a p exactly i spaces to the left of the origin, and 0s elsewhere. Note that gx_k places the person at the origin and gx_i places the person off the origin for all $i \neq k$. Consequently, the length of a minimal presentation of g by φ_a and φ_b can be deduced from this information and all representations of g as a product of φ_a and φ_b (but not their inverses) have the same length. Now fix $1 \leq i \leq k$. Then gx_{k-i+1} is a configuration with the letter representing g_i (a or b) at the origin. Therefore the natural surjection from $\{a, b\}^*$ to $\langle \varphi_a, \varphi_b \rangle^+$ is an injection and so this is the free monoid of rank 2.

Acknowledgements. We thank Michael Schraudner for sharing his preprint and for pointing out that their example gives a subshift with zero entropy whose automorphism group is amenable but not virtually abelian. We also thank the anonymous referee for numerous helpful comments. The second author was partially supported by NSF grant 1500670.

REFERENCES

[**1**] 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.
[**2**] V. Cyr, J. Franks, S. Petite and B. Kra. Distortion and the automorphism group of a shift. *J. Mod. Dyn.*, to appear.
[**3**] V. Cyr and B. Kra. The automorphism group of a shift of subquadratic growth. *Proc. Amer. Math. Soc.* **144**(2) (2016), 613–621.
[**4**] V. Cyr and B. Kra. The automorphism group of a shift of linear growth: beyond transitivity. *Forum Math. Sigma* **3** (2015), e5.
[**5**] V. Cyr and B. Kra. The automorphism group of a minimal shift of stretched exponential growth. *J. Mod. Dyn.* **10** (2016), 483–495.
[**6**] S. Donoso, F. Durand, A. Maass and S. Petite. On automorphism groups of low complexity subshifts. *Ergod. Th. & Dynam. Sys.* **36**(1) (2016), 64–95.
[**7**] S. Donoso, F. Durand, A. Maass and S. Petite. On automorphism groups of Toeplitz subshifts. *Discrete Anal.* (2017), Paper No. 11, 19 pp.
[**8**] G. A. Hedlund. Endomorphisms and automorphisms of the shift dynamical system. *Math. Systems Theory.* **3** (1969), 320–375.
[**9**] M. Morse and G. A. Hedlund. Symbolic dynamics II. Sturmian trajectories. *Amer. J. Math.* **62** (1940), 1–42.
[**10**] A. Quas and L. Zamboni. Periodicity and local complexity. *Theoret. Comput. Sci.* **319** (2004), 229–240.
[**11**] V. Salo and M. Schraudner. Automorphism groups of subshifts through group extensions. *Preprint*.