Normal amenable subgroups of the automorphism group of sofic shifts

KITTY YANG

Department of Mathematics, Northwestern University, Evanston, IL 60208, USA
(e-mail: kyang@math.northwestern.edu)

Abstract. Let (X, σ) be a transitive sofic shift and let $\mathrm{Aut}(X)$ denote its automorphism group. We generalize a result of Frisch, Schlank, and Tamuz to show that any normal amenable subgroup of $\mathrm{Aut}(X)$ must be contained in the subgroup generated by the shift. We also show that the result does not extend to higher dimensions by giving an example of a two-dimensional mixing shift of finite type due to Hochman whose automorphism group is amenable and not generated by the shift maps.

1. *Introduction*

1.1. *The center of the group of automorphisms of sofic shifts.* Fix a finite alphabet $\mathcal{A}$. We define a subshift to be a closed set $X \subset \mathcal{A}^{\mathbb{Z}}$ that is invariant under the shift map $\sigma :$ $\mathcal{A}^{\mathbb{Z}} \to \mathcal{A}^{\mathbb{Z}}$, where $(\sigma x)_n = x_{n+1}$, for $(x_n : n \in \mathbb{Z}) \in X$. A well-studied class are shifts of finite type (SFTs), which can be described by a finite set of forbidden words (these notions are defined precisely in §2). Closing SFTs under passage to factors, we obtain the sofic shifts, and this is the context of the result in this paper. There has been much interest in the algebraic properties of the automorphism group of the full shift and mixing SFTs (e.g. [**4**, **10**, **12**, **16**]), and many of these results generalize to sofic shifts. The automorphism group $\mathrm{Aut}(X)$ is the group of homeomorphisms of X to itself that commute with σ, under composition. For any shift X, $\mathrm{Aut}(X)$ trivially contains $\langle \sigma \rangle$, the subgroup generated by the shift. By the Curtis–Hedlund–Lyndon theorem [**10**], any automorphism is a block code; as a corollary, $\mathrm{Aut}(X)$ is always a countable group.

Many results on automorphism groups of SFTs are obtained by manipulating marker automorphisms, initially introduced in [**10**] as a class of finite-order automorphisms of the full shift. Using marker automorphisms, Hedlund proved that the automorphism group of the full shift contains, among others, isomorphic copies of every finite group and the

free group on two generators. Complications arise for mixing SFTs with the existence of forbidden patterns, but Boyle, Lind, and Rudolph [4] showed that marker automorphisms can also be defined for SFTs, and generalized Hedlund's results to SFTs.

Also using marker automorphisms, Kim and Roush [12] embedded the automorphism group of the full shift into the automorphism group of any mixing SFT, using markers to encode data words to act as symbols in a full shift. As a corollary, the automorphism groups of the full 2-shift and the full 3-shift contain the same subgroups up to isomorphism, but it remains an open question whether these automorphism groups are isomorphic. We note that the Kim and Roush construction does not embed the automorphism group of a mixing SFT into the automorphism group of another mixing SFT, as the embedding relies heavily on the lack of forbidden words in the full shift.

On the other hand, a method to distinguish automorphism groups follows from Ryan's theorem [16, 17]. By definition, $\langle \sigma \rangle$ is contained in the center of $\mathrm{Aut}(X, \sigma)$. An SFT is transitive if it contains a point whose orbit is dense. Ryan showed that for any transitive SFT, the center is the subgroup generated by the shift. In the full 4-shift, the shift map has a root, while the shift map for the full 2-shift does not. Using Ryan's theorem, Boyle, Lind, and Rudolph [4] observed that these automorphism groups cannot be isomorphic, as the automorphism group of the full 4-shift contains an element not in the center whose square is in the center, while the automorphism group of the full 2-shift has no such element.

A topological factor of an SFT is not necessarily an SFT. To address this, Weiss [18] introduced the notion of sofic shifts, which form the smallest class of subshifts, closed under taking factors, that contain SFTs. It is natural to ask whether results about automorphism groups of SFTs also apply to sofic shifts. As the definition of marker automorphism relies on the fact that all sufficiently long words are synchronizing and transitive sofic shifts contain an abundance of synchronizing words, many of the results about subgroups of automorphism groups of SFTs also carry over to sofic shifts.

More recently, Frisch, Schlank, and Tamuz [6] generalized Ryan's theorem to show that any normal amenable subgroup of the automorphism group of the full shift must be contained in the subgroup generated by the shift (see §2.4 for precise definitions). In this paper we extend their result to any transitive sofic shift.

THEOREM 1.1. *Let (X, σ) be a transitive sofic shift. Any normal amenable subgroup of* $\mathrm{Aut}(X)$ *is contained in* $\langle \sigma \rangle$.

This generalizes Ryan's theorem, as subgroups of the center are always normal amenable subgroups, and in the case of a transitive sofic shift, our result says the converse also holds.

Complications arise when working in sofic shifts, as there can be arbitrarily long non-synchronizing words. Such words are necessary for a marker automorphism to actually be an automorphism of X. To overcome this problem, we first construct marker automorphisms for left-periodic points composed of synchronizing words as in the definition of extreme proximality (see §2.4). We then extend the result to non-synchronizing left-periodic points that are in the $\mathrm{Aut}(X)$-orbit of the synchronizing points.

1.2. *Methods of proof.* As in [**6**], we use a characterization of the topological boundary due to Furman [**7**] to prove Theorem 1.1 (see §2 for precise definitions).

For any shift X, Aut(X) acts on the set of left-periodic points. In the case of a transitive sofic shift, we show that an invariant subset of this space is in fact a topological boundary for Aut(X), and the kernel of the action is $\langle \sigma \rangle$. By Furman's characterization, we can conclude that any normal amenable subgroup of Aut(X) must be contained in $\langle \sigma \rangle$.

A key ingredient of the paper is Proposition 3.5: given any two cylinder sets of the boundary, we construct an automorphism of X that maps one into the other. This is not needed in the case of the full shift, as [**6**] explicitly constructs a sequence of automorphisms which map proper closed subsets of the boundary to a particular point.

In the proof of Theorem 1.1, we show that the action of Aut(X) on the boundary is extremely proximal. This gives an alternate proof that for an infinite transitive sofic shift, the automorphism group contains a copy of the free group on two generators, and more generally, the free group on any number of generators.

1.3. *Higher dimensions.* In §4 we highlight the obstructions that arise when adapting these methods to higher-dimensional shifts. Hochman [**11**] proves a higher-dimensional analogue of Ryan's theorem: for a transitive $\mathbb{Z}^d$ SFT with positive entropy, the center of the automorphism group is the subgroup generated by the shifts, which can be naturally identified with $\mathbb{Z}^d$. While [**6**] also shows that for the full $\mathbb{Z}^d$-shift, any normal amenable subgroup of the automorphism group is contained in the subgroup generated by the shift maps, $\mathbb{Z}^d$, there is an example due to Hochman of a topological mixing $\mathbb{Z}^2$ SFT with positive entropy whose automorphism group is amenable but strictly larger than $\mathbb{Z}^2$.

2. *Preliminaries*

2.1. *Shift systems.* Let $\mathcal{A}$ be a finite alphabet endowed with the discrete topology and equip $\mathcal{A}^{\mathbb{Z}}$ with the product topology. For $x = (x_n : n \in \mathbb{Z}) \in \mathcal{A}^{\mathbb{Z}}$, let $x_n \in \mathcal{A}$ denote the value of x at $n \in \mathbb{Z}$. Define the shift map $\sigma : \mathcal{A}^{\mathbb{Z}} \to \mathcal{A}^{\mathbb{Z}}$ by setting $(\sigma x)_n = x_{n+1}$ for any $x \in \mathcal{A}^{\mathbb{Z}}$. If $X \subset \mathcal{A}^{\mathbb{Z}}$ is a closed and shift-invariant subset, we call (X, σ) a *shift*. To avoid trivial cases, we assume X is infinite. Given $x \in X$, let

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

be the orbit of x under the shift, and $\overline{\mathcal{O}(x)}$ denote its closure in X.

Given an interval $[i, i + n - 1] \subset \mathbb{Z}$, let $x_{[i,i+n-1]}$ be the word w in $\mathcal{A}^n$ given by $w_j = x_{i+j}$ for $j = 0, 1, \ldots, n - 1$. A word w in $\mathcal{A}^n$ is *allowable in X* if there exist $x \in X$ and $i \in \mathbb{Z}$ such that $w = x_{[i,i+n-1]}$; we say that w *occurs in x at i*. For any word $w \in \mathcal{A}^n$, let $|w| = n$ be the length of the word. We denote the collection of allowable words of length n in X by $\mathcal{L}_n(X)$, and the *language of X*, $\mathcal{L}(X) = \bigcup_{n \in \mathbb{N}} \mathcal{L}_n(X)$, is the set of all finite words that occur in X. Given two words u and w, uw is the word in $\mathcal{A}^{|u|+|w|}$ obtained by concatenating u and w; when we concatenate a word with itself, we simplify notation by writing w^2.

Given a word $w \in \mathcal{L}(X)$, define the cylinder set $[w] \subset X$ to be

$$[w] := \{x \in X : x_n = w_n \text{ for } 0 \leq n < |w|\}.$$

Such cylinder sets are clopen and, together with their translates, form a basis for the subspace topology on X. Thus, we can describe topological properties of a shift in terms of its language. A shift (X, σ) is *transitive* if for any pair of words u and $w \in \mathcal{L}(X)$, there is some word v such that $uvw \in \mathcal{L}(X)$; a shift (X, σ) is *mixing* if for any $u, w \in \mathcal{L}(X)$, there exists an N such that for any $n \geq N$, there is a word $v \in \mathcal{L}_n(X)$ such that uvw is again allowable. Note that mixing implies transitivity.

A word $w \in \mathcal{L}(X)$ is *synchronizing* if, whenever uw and $wv \in \mathcal{L}(X)$, uwv is again allowable in X. It follows that if w is synchronizing, then any word that contains w must also be synchronizing.

We say that a word u is a *root* of w if we can write $w = u^k$ for some $k \in \mathbb{N}$. If w has no proper roots, then w is *primitive*.

Two words u, u' are *conjugate* if we can write $u = wv$ and $u' = vw$ for some v, w (possibly empty) words.

A point $x \in X$ is *periodic* if there exists $k \in \mathbb{N}$ such that $x_n = x_{n-k}$ for all $n \in \mathbb{Z}$. If x is periodic with minimal period k, then any word of length k which appears in x is primitive, and any two subwords of length k are conjugate.

We say $x \in X$ is *left-periodic up to* $N \in \mathbb{Z}$ if there exists $k \in \mathbb{N}$ such that $x_n = x_{n-k}$ for all $n < N$ and $x_N \neq x_{N-k}$. Note that a left-periodic point is not periodic, and that the index N is independent of the choice of k. When the periodic index k is minimal, we say that x is (left-)k-periodic. Denote the set of points in X of period k by Per_k. If X contains a left-periodic point, then it must also contain a periodic point, but the converse is not true.

A shift is of *finite type*, or an SFT, if it can be described by a finite set of forbidden words; that is, X is an SFT if there exists a finite set of words $\mathcal{F}$ such that $x \in X$ if and only if any word that occurs in x is not an element of $\mathcal{F}$. An SFT is *j-step* if $\mathcal{F}$ consists of words of length $j + 1$. We can also characterize SFTs using synchronizing words: a shift X is an SFT if all words of sufficient length are synchronizing.

When (X, σ) is a transitive SFT, it is a classical result that X can be decomposed into disjoint mixing components which are cyclically permuted. More precisely, there exist a period p and a partition $\{E_i\}_{1 \leq i \leq p}$ of X such that each (E_i, σ^p) is mixing, and $\sigma E_i = E_{i+1 \bmod p}$ (see [1, §3], [2, p. 543]). Here, p refers to the greatest common divisor of all k with $\mathrm{Per}_k \neq \emptyset$. This extra structure says that in the case of a transitive SFT, given w and u, we can extend u on the left to $\widetilde{u}$ and, for sufficiently large n, there is a word v of length np with $wv\widetilde{u} \in X$. Transitivity also implies that periodic points are dense for SFTs.

The class of SFTs is not closed under factors, and we can consider the larger natural class which is. Given a compact metric space X and a homeomorphism T, a *topological factor* is a system (Y, S) with a continuous surjective map $\pi : X \to Y$ satisfying $\pi \circ T = S \circ \pi$. We say that a shift is *sofic* if it is a topological factor of an SFT. There are many equivalent definitions of sofic shifts, and we refer the reader to [14, Theorem 3.2.1] for more details. Sofic shifts are the smallest class of shifts which are closed under taking factors and contain SFTs. Note that transitivity and mixing are each preserved under factors, and a transitive (cf. mixing) sofic shift is a factor of a transitive (cf. mixing) SFT. It follows that in a transitive sofic shift, periodic points are also dense. By recoding, we can assume that this factor map is a 0-block map. This presentation is convenient as it allows us to lift words in the sofic shift to words of the same length in the SFT. It follows that in

transitive sofic shifts, as with transitive SFTs, between any two allowable words, we can insert arbitrary spacer words whose lengths form an arithmetic progression.

2.2. *Automorphisms of a shift.*   An *automorphism of* (X, σ) is a homeomorphism from X to itself that commutes with the shift map.

By the Curtis–Hedlund–Lyndon theorem [10], any automorphism is defined by a *block code*: given an automorphism $g \in \mathrm{Aut}(X)$, there exist an $R \in \mathbb{N}$ and a map $\hat{g} : \mathcal{L}_{2R+1}(X) \to \mathcal{A}$ such that $gx_i = \hat{g}(x_{i-R,\ldots,i+R})$. We say that R is a *range* for φ.

The set of automorphisms of X under composition forms a group $\mathrm{Aut}(X, \sigma)$, or simply $\mathrm{Aut}(X)$ when σ is clear from context. Since only finitely many automorphisms can have a given range, $\mathrm{Aut}(X)$ is countable. Given two automorphisms $g_1, g_2 \in \mathrm{Aut}(X)$, let $g_1 g_2$ denote the composition $g_1 \circ g_2$.

In general, it is difficult to construct automorphisms of an arbitrary shift; however, if a shift contains synchronizing words, there are finite-order automorphisms called *marker automorphisms* originally defined by Hedlund [10] for full shifts, and later for SFTs by Boyle, Lind, and Rudolph [4]. We now define marker automorphisms more generally, making slight modifications to conventions introduced in [6].

We say that two words w and u *overlap* if we can write $w = w'v$ and $u = vu'$ (or vice versa). When needed, we specify the length of overlap, and we say that w and u overlap with length i, where $i = |v|$.

Let (X, σ) be a shift and M_ℓ and $M_r \in \mathcal{L}(X)$ be synchronizing words. Let $\mathcal{D} \subset \mathcal{L}_n$ be a set of words of length n appearing in X such that words of the form $M_\ell d M_r$ are allowable for all $d \in \mathcal{D}$. Suppose these words satisfy the following overlap condition: for any d and $d' \in \mathcal{D}$, if $M_\ell d M_r$ and $M_\ell d' M_r$ overlap non-trivially with length i, then $i \le \min(|M_\ell|, |M_r|)$. Then any permutation τ of $\mathcal{D}$ induces an automorphism g_τ on X by sending words of the form $M_\ell d M_r$ to $M_\ell \tau(d) M_r$ and leaving other words unchanged. Such an automorphism is called a *marker automorphism*, and we refer to M_ℓ and M_r as the *left and right markers*, respectively, and $d \in \mathcal{D}$ as *data words*. We note that as originally defined for a j-step SFT, the length of marker words have to be greater than j. The key is that such words are synchronizing, which is the necessary condition to ensure that applying the map does not introduce forbidden words.

Example 2.1. Let $X \subset \{0, 1\}^{\mathbb{Z}}$ be the one-step SFT defined by the forbidden word 11. (X, σ) is the golden mean shift.

Let $M_\ell = 100$ and $M_r = 0101$ be start and end markers, and $\mathcal{D} = \{0, 1\}$ be data words. Given special blocks of the form $M_\ell d M_r$ and $M_\ell d' M_r$, for $d, d' \in \mathcal{D}$, $M_\ell d M_r$ and $M_\ell d' M_r$ can only overlap non-trivially by length 1.

Let $g \in \mathrm{Aut}(X)$ be the marker automorphism induced by the non-trivial permutation on $\mathcal{D}$. It permutes blocks of the form

$$\cdots 100\mathbf{1}0101 \cdots$$
$$\cdots 100\mathbf{0}0101 \cdots$$

and leaves other blocks unchanged.

Example 2.2. Let $(Y, \sigma) \subset \{0, 1\}^{\mathbb{Z}}$ be the even shift, consisting of bi-infinite sequences with only an even number of consecutive 1s. Y is a factor of the golden mean shift in the previous example, so it is sofic, but it is not an SFT. The word 1^{2i+1} for any $i \in \mathbb{N}$ is allowable (as a subword of $01^{2i+2}0$) and 01^{2i+1}, $1^{2i+1}0$ are also allowable. However, $01^{2i+1}0$ is not in the language of Y, so 1^{2i+1} is not a synchronizing word. We have generated arbitrarily long non-synchronizing words; thus, Y cannot be an SFT. We call such shifts *strictly sofic*.

We note that any word which contains 0 is a synchronizing word, so we can define marker automorphisms with markers that contain 0.

2.3. *Generalized Ryan's theorem for sofic shifts.* By generalizing the definition of marker automorphisms, we can adapt the proof of Ryan's theorem to show that for a transitive sofic shift, the center of the automorphism group must be the subgroup generated by the shift. We show that for a transitive sofic shift, the automorphism group contains enough markers so that Ryan's theorem still holds. The key proposition needed, which we state without proof, is that a transitive shift contains infinitely many synchronizing words.

PROPOSITION 2.3. [**14**, Proposition 3.3.16] *Suppose* (X, σ) *is a transitive sofic shift. Then any word* $w \in \mathcal{L}(X)$ *can be extended on the right to a synchronizing word* wu.

Remark 2.4. We note that in [**14**], sofic shifts are defined as the set of all bi-infinite paths on a labeled graph. The definition of synchronizing word in [**14**] is dependent on the graph, while they use the term *intrinsically synchronizing* to denote words we call synchronizing. However, if one chooses the minimal graph presentation for the sofic shift, these definitions coincide.

To prove Ryan's theorem for transitive sofic shifts, it suffices to show that there exist infinitely many synchronizing words which do not overlap themselves. To prove this lemma, we appeal to the following theorem in the field of combinatorics on words, translated into the language of symbolic dynamics.

THEOREM 2.5. ([**5**, Theorem 3.1], [**15**]) *Let* $w \in \mathcal{A}^n$ *be a primitive word. Then there exists* w' *conjugate to* w *which does not overlap itself.*

LEMMA 2.6. *Let* (X, σ) *be a transitive sofic shift. Then for any* $n \in \mathbb{N}$, *there is a synchronizing word* M *of at least length* n *which does not overlap itself non-trivially.*

Proof. Let w be a synchronizing word of at least length n. By transitivity, there exists $u \in \mathcal{L}(X)$ such that $wuw \in \mathcal{L}(X)$ and u is not a power of w. Note that wuw is again a synchronizing word.

Since periodic points are dense, there is a periodic point x of minimal period $k \geq |wuw|$ such that wuw appears in x. By Theorem 2.5, x must contain a subword M of length k which does not overlap itself. Since $|wuw| \leq k$, w must appear in M, and thus M is a synchronizing word. $\square$

For completeness, we give a proof of the generalized Ryan's theorem, and we follow the proof for SFTs, due to Kitchens [**13**, Theorem 3.3.22].

THEOREM 2.7. *Let (X, σ) be a transitive sofic shift. The center of $\mathrm{Aut}(X)$ is $\langle \sigma \rangle$.*

Proof. Let (X, σ) be a transitive sofic shift, and let $\varphi \in \mathrm{Aut}(X)$ commute with all automorphisms. Suppose φ has range R. Recall that for transitive sofic shifts, between any two words we can always insert spacers of lengths that form an arithmetic progression, where the difference is p, the period of X. Using these spacers and the sufficiently long markers produced by Lemma 2.6, we can find a synchronizing $M \in \mathcal{L}(X)$ and $n \in \mathbb{N}$, with $2R + 1 \leq n \leq |M|$, such that for

$$\mathcal{D}(M, n) := \{d \in \mathcal{L}_n(X) : MdM \in \mathcal{L}(X)\},$$

every word of length $2R + 1$ appears as a subword of some element of $\mathcal{D}(M, n)$. This can be done by applying the transitive property simultaneously to M and words of length $2R + 1$ so that the spacers are of the same length. If necessary we can extend M to the left. Repeat the process on the right to get words of the form MdM. For any permutation $\tau \in \mathrm{Sym}(\mathcal{D}(M, n))$, let g_τ denote the marker automorphism induced by τ.

Consider the periodic points of period $|M| + n$ obtained by concatenating Md with itself, for any $d \in \mathcal{D}(M, n)$. We denote such points $\mathrm{Per}(M, n) \subset \mathrm{Per}_{|M|+n}$. Let $\mathrm{Orb}(M, n)$ be the set of distinct σ-orbits in $\mathrm{Per}(M, n)$. Note that $|\mathrm{Orb}(M, n)| \geq 2$, as each word of length $2R + 1$ appears in some $d \in \mathcal{D}(M, n)$.

For any permutation of $\mathrm{Orb}(M, n)$, there is a g_τ, for some $\tau \in \mathrm{Sym}(\mathcal{D}(M, n))$ whose action on $\mathrm{Orb}(M, n)$ coincides with the given permutation. In addition, g_τ acts as the identity on periodic points of period $|M| + n$ which are not in $\mathrm{Per}(M, n)$.

We claim that φ acts on $\mathrm{Orb}(M, n)$. Suppose not. Then φ maps some $x \in \mathrm{Per}(M, n)$ to a periodic point y not in $\mathrm{Per}(M, n)$. Since φ commutes with all g_τ, this means that all points in $\mathrm{Per}(M, n)$ are mapped to the σ-orbit of y, which contradicts that φ permutes the periodic points of each period.

Now we show that φ acts as the identity permutation. Let x and $y \in \mathrm{Per}(M, n)$ be in distinct σ-orbits. We note that there exists $\varphi(x) = \sigma^j(x)$ for some $-R \leq j \leq R$. This equality holds for all points in the σ-orbit of x, and we show that it holds for y as well. Let g_τ be a permutation that takes $\mathcal{O}(x)$ to $\mathcal{O}(y)$. As φ commutes with g_τ, we have

$$\varphi(y) = g_\tau^{-1} \circ \varphi \circ g_\tau(y) = \sigma^j y.$$

As every block of length $2R + 1$ appears in some $d \in \mathcal{D}(M, n)$, we conclude $\varphi = \sigma^j$. $\square$

2.4. *Topological boundaries.* Throughout this section let G be a locally compact group and let Ω be a compact metric space with a continuous G action $G \times \Omega \to \Omega$: for any $g \in G$ and $\omega \in \Omega$,

$$(g, \omega) = g \cdot \omega.$$

We call Ω a *G-space*. Given $\omega \in \Omega$, let $G\omega$ denote the G-orbit of ω,

$$G\omega = \{g \cdot \omega : g \in G\} \subset \Omega,$$

and $\overline{G\omega}$ its closure in Ω.

Let $\mathrm{Prob}(\Omega)$ be the set of Borel probability measures on Ω, equipped with the weak-* topology. Since Ω is compact, $\mathrm{Prob}(\Omega)$ is also a compact metric space. Given $\omega \in \Omega$, let

δ_ω denote the Dirac measure concentrated at ω. The mapping $\omega \mapsto \delta_\omega$ gives an embedding of Ω into $\mathrm{Prob}(\Omega)$.

The G-action on Ω induces an action on $\mathrm{Prob}(\Omega)$ by viewing elements of G as self-homeomorphisms of Ω: for any $g \in G$ and $\omega \in \Omega$,

$$g \cdot \mu = \mu \circ g^{-1}.$$

We say that the G-action on Ω is *minimal* if for any $\omega \in \Omega$, the G-orbit closure $\overline{G\omega} = \Omega$. The G-action on Ω is *strongly proximal* if for all $\mu \in \mathrm{Prob}(\Omega)$, the G-orbit closure $\overline{G\mu} \subset \mathrm{Prob}(\Omega)$ contains a Dirac measure δ_ω for some $\omega \in \Omega$. A G-space Ω is a *topological boundary* if the G-action on Ω is minimal and strongly proximal.

The G-action on Ω is *extremely proximal* if $|\Omega| \geq 2$ and for any proper closed set $C \subsetneq \Omega$, and any open set $U \subset \Omega$ there is some $g \in G$ with $gC \subset U$.

It is known that extreme proximality implies strong proximality [**8**, §3] and the product of strongly proximal actions is again strongly proximal [**9**, §3].

A group G is *amenable* if for every compact G-space Ω, the G-action on $\mathrm{Prob}(\Omega)$ has a fixed point. Examples of amenable groups include abelian groups and finite groups, while the free group is not amenable.

3. Topological boundaries of the automorphism group of transitive sofic shifts

3.1. The action of the automorphism group on left-periodic points.
Recall from §2 that left-periodic points are not periodic. For any shift (X, σ), we define a compact space equipped with an $\mathrm{Aut}(X)$ action.

LEMMA 3.1. *Let (X, σ) be a shift and $k \in \mathbb{N}$. Suppose X contains a left-k-periodic point. We denote the set of left-k-periodic points up to k by Q_k. Then Q_k is an $\mathrm{Aut}(X)$-space, and σ acts trivially on Q_k. If X contains a left-k-periodic point which is transitive, then the kernel of the action is $\langle \sigma \rangle$.*

Proof. Since any automorphism is a block map, the set of left-k-periodic points is invariant under $\mathrm{Aut}(X)$. The set of all left-k-periodic points is precisely $\bigcup_{i \in \mathbb{Z}} \sigma^i Q_k$. Thus, for any $g \in \mathrm{Aut}(X)$ and $x \in Q_k$,

$$gx \in \sigma^i Q_k, \tag{1}$$

for some unique i, since the shifts of Q_k are pairwise disjoint. Define a cocycle $\alpha : \mathrm{Aut}(X) \times Q_k \to \mathbb{Z}$ to be

$$\alpha(g, x) = -i, \tag{2}$$

where i is obtained from equation (1). The cocycle condition ensures that the induced $\mathrm{Aut}(X)$-action on Q_k is well defined, where for $g \in \mathrm{Aut}(X)$, $x \in Q_k$,

$$g \cdot x = \sigma^{\alpha(g,x)} \circ gx. \tag{3}$$

We note here that the action of $\mathrm{Aut}(X)$ on X is different from the action on Q_k, so we use different notation to make clear which action we are referencing. For any $x \in Q_k$, $\alpha(\sigma, x) = -1$, so $\sigma \cdot x = x$.

Suppose in addition that $x \in X$ is a transitive left-k-periodic point. Let $g \notin \langle \sigma \rangle$. For each $n \in \mathbb{N}$, let R_n denote the maximum of n and the range of g. Let $\hat{g}_n$ and $\hat{\sigma}^n$ denote the

block codes of width R_n that induce g and σ^n, respectively. Then there exists a word w_n of length $2R_n + 1$ such that $\hat{g}_n(w_n) \neq \hat{\sigma}^n(w_n)$. Since every w_n appears in x, $g \cdot x \neq x$. $\quad\square$

The set of k-periodic points Per_k is invariant under $\mathrm{Aut}(X)$. We can decompose Per_k into a disjoint union of distinct σ-orbits:

$$\mathrm{Per}_k = \mathcal{O}(x^1) \sqcup \cdots \sqcup \mathcal{O}(x^j). \tag{4}$$

Thus, the action of $\mathrm{Aut}(X)$ on X descends to an action on $\mathrm{Per}_k / \langle \sigma \rangle$.

LEMMA 3.2. *Let (X, σ) be a shift that contains a left-k-periodic point, and Q_k be the set of left-k-periodic points up to k. There is a map*

$$\pi : Q_k \rightarrow \mathrm{Per}_k / \langle \sigma \rangle \tag{5}$$

which is $\mathrm{Aut}(X)$-*equivariant. When (X, σ) is a transitive sofic shift, π is a projection.*

Proof. Given $x \in Q_k$, there exists a unique k-periodic point such that $y_n = x_n$ for all $n < k$. Define the projection π which sends x to the σ-orbit in Per_k containing y. Since any automorphism is a block code, for any $x \in Q_k$ and $g \in \mathrm{Aut}(X)$, $\pi(gx) = g\pi(x)$.

Suppose (X, σ) is a transitive sofic shift, and let $y \in \mathrm{Per}_k$. Consider the word $y_{[k-N,k)}$, which we can extend to the right to be a synchronizing word w. Since X is transitive, there is a word $u \in \mathcal{L}(X)$ of the form $u = wsa$, where $a \neq y_{k+p-1}$ (where p refers to the greatest common divisor of the periods of all periodic points; see §2). Then we can extend u to a left-k periodic point in Q_k (after shifting appropriately so it lands in Q_k) that projects to the σ-orbit containing y. $\quad\square$

A map $s : \mathrm{Per}_k / \langle \sigma \rangle \rightarrow Q_k$ is a section of the projection π if s is a right inverse of π. Let Ω be the collection of all sections $s : \mathrm{Per}_k / \langle \sigma \rangle \rightarrow Q_k$ of the projection π. Since π is equivariant, the action of $\mathrm{Aut}(X)$ on Q_k induces an action on Ω: for any $g \in G$ and $s \in \Omega$,

$$gs = g \cdot (s \circ g^{-1})$$

where we view g^{-1} as the permutation of $\mathrm{Per}_k / \langle \sigma \rangle$ induced by g^{-1}.

Let $\{\Omega^m\}$ denote the fibers of $\pi : Q_k \rightarrow \mathrm{Per}_k / \langle \sigma \rangle$. Given a periodic orbit $\mathcal{O}(x^m) \in \mathrm{Per}_k / \langle \sigma \rangle$,

$$\Omega^m = \{x \in Q_k : \exists i \in \mathbb{Z} \text{ with } x_n = x_{n-i}^m \text{ for all } n < k\}. \tag{6}$$

Let $N \lhd \mathrm{Aut}(X)$ be the normal subgroup given by the kernel of π. Since N preserves the fibers $\{\Omega^m\}$, the restriction of the action on Ω to N is isomorphic to the diagonal action of N on the product of the fibers $\prod_{m=1}^{j} \Omega^m$, where j is the number of distinct k-periodic orbits defined in (4).

3.2. *Extreme proximality.* Given a bi-infinite sequence $x \in X$, we say that x is *synchronizing* if all sufficiently long words that appear in x are synchronizing. If x is periodic and some synchronizing word appears in x, then x itself is synchronizing. Since the definition only depends on words that appear in x, a sequence is synchronizing if and only if any point in its orbit closure is synchronizing.

Every periodic orbit $\mathcal{O}(x)$ is exactly one of the three following types:

(1) x is synchronizing;

(2) x is not synchronizing, but there exists an automorphism $h \in \mathrm{Aut}(X)$ such that hx is synchronizing;

(3) x is not synchronizing, and for all automorphisms $h \in \mathrm{Aut}(X)$, hx is not synchronizing.

Remark 3.3. When X is an SFT, all sequences are synchronizing and Per_k consists only of synchronizing points for large enough k.

In general, the action of $\mathrm{Aut}(X)$ on Per_k may not be transitive. In the case of SFTs, however, for sufficiently large k, $\mathrm{Aut}(X)$ does act on Per_k transitively (see [**4**]). The proof constructs a composition of marker automorphisms which permute periodic points with disjoint orbits, building on work by Boyle and Krieger [**3**] for the full shift. The same proof shows that for a transitive sofic shift, $\mathrm{Aut}(X)$ acts transitively on the synchronizing points in Per_k. However, non-synchronizing points do not contain any synchronizing subwords, so they are fixed by all marker automorphisms.

Let $\mathrm{Syn}_k \subset \mathrm{Per}_k$ be the subset of periodic points of type (1) and type (2). Then Syn_k is the $\mathrm{Aut}(X)$-orbit under $\mathrm{Aut}(X)$ of any synchronizing point, so $\mathrm{Aut}(X)$ must act transitively on Syn_k.

Remark 3.4. It is possible that no periodic points of type (2) exist. We do not know if there exist automorphisms which do not fix non-synchronizing points.

Define $\widetilde{\Omega}$ analogously to Ω in §3.1, but for the restricted action. For any section $s \in \Omega$, where $s : \mathrm{Per}_k /\langle\sigma\rangle \to Q_k$, let $s|_{\mathrm{Syn}_k /\langle\sigma\rangle} : \mathrm{Syn}_k /\langle\sigma\rangle \to Q_k$ be the corresponding element in $\widetilde{\Omega}$. Note that this is not injective and many $s \in \Omega$ project to the same map in $\widetilde{\Omega}$.

We now show that $\widetilde{\Omega}$ is a topological boundary for $\mathrm{Aut}(X)$. Recall that N, the kernel of the action of Per_k, acts on each Ω^m. If we consider the action of N on $\widetilde{\Omega}$,

$$\widetilde{\Omega} \cong \prod_m \Omega^m$$

where the product is taken over values of m where x^m is of type (1) or type (2).

We show that the action of N on Ω^m is extremely proximal, and use this to prove that the full action of $\mathrm{Aut}(X)$ on $\widetilde{\Omega}$ is a topological boundary. The key step is constructing marker automorphisms in Ω^m, where x^m is a synchronizing point. Then for non-synchronizing points, we exploit the fact that such points are in the $\mathrm{Aut}(X)$-orbit of some synchronizing point to achieve the same result.

The set Ω^m is closed in X, so cylinder sets of the form

$$[w]^m := [w] \cap \Omega^m \quad \text{where } w \in \mathcal{L}(X) \tag{7}$$

form a subbase that generates the subspace topology on Ω^m.

PROPOSITION 3.5. *Let (X, σ) be a transitive sofic shift and $k \in \mathbb{N}$. Suppose X contains a left-k-periodic point. Fix $x^m \in \mathrm{Syn}_k$, and define Ω^m as in (6). Let w, u be words in $\mathcal{L}(X)$ such that the corresponding cylinder sets $[w]^m$ and $[u]^m$ in Ω^m are non-empty and proper. Then there is an involution $g \in \mathrm{Aut}(X)$ which acts as the identity on $\mathrm{Per}_k /\langle\sigma\rangle$ and satisfies $g \cdot [w]^m \subset [u]^m$.*

Proof. We fix $x^m \in \mathrm{Syn}_k$, so x^m is of either type (1) or type (2).

Case 1: x^m is synchronizing.
There exists $c \in \mathbb{N}$ such that all subwords of length at least c are synchronizing.

Since points in Ω^m are left-k-periodic up to k, if $|w| \leq k$, we can replace it with a word $\widetilde{w}$ of length $k + 1$ such that $[w]^m = [\widetilde{w}]^m$. So we may assume that $|w|, |u| > k$. Write $a = w_{[0,\ldots,k-1]}$. If $w[0] \neq u[0]$, let $\widetilde{u}$ be the unique extension of u so that w and $\widetilde{u}$ begin with the same letter.

We can assume that $\widetilde{u}$ does not appear as the initial word of w, otherwise a power of the shift will map $[w]^m$ into $[u]^m$ and the identity automorphism satisfies the conclusion of the lemma. We first deal with the case where w does not appear as the initial word of $\widetilde{u}$.

Let a^r be the word obtained by concatenating r copies of a. Since it is a word of at least length c which appears in x^m, it must be synchronizing. As $[w]^m \neq \emptyset$, $a^r w$ is allowable, and must also be synchronizing. Since X is transitive, choose $v \in \mathcal{L}(X)$ such that $\widetilde{u} v a^r$ is an allowable word. Since a is the initial word of $\widetilde{u}$, we can choose v with $|\widetilde{u} v| = rpk$ for some $r \geq \max\{|w|, c\}$, where p is the greatest common divisor of k with $\mathrm{Per}_k \neq \emptyset$. Set a^r to be the left marker, $a^r w$ to be the right marker, and $\mathcal{D} = \{a^{rp}, \widetilde{u} v\}$ to be data words.

To show these markers induce a well-defined marker automorphism, it suffices to check that special words of the form $a^r d a^r w$, for $d \in \mathcal{D}$, satisfy the overlap condition given in the definition of marker automorphism. The word a is a primitive word of length k which appears in x^m, a k-periodic point. While a may overlap itself non-trivially, a^r can only overlap itself by multiples of k, as k is the least period of x^m.

The initial word of w and $\widetilde{u}$ is a, while a does not occur at position k in w or $\widetilde{u}$. Thus, if the special words $a^r a^{rp} a^r w$ and $a^r \widetilde{u} v a^r w$ overlap non-trivially, the length of the overlap must be either less than $|w|$, or exactly $rk + |w|$. In the second case, the special blocks would overlap by a^r. However, this overlap would force $\widetilde{u}$ to be the initial word of w, which contradicts the assumption.

Because special words begin with a^r, a similar argument shows that a special word can only overlap with itself non-trivially by at most $|w|$. By the choice of r, $|a^r| \geq |w|$, so the marker automorphism g determined by the non-trivial permutation on $\mathcal{D}$ is well defined, and g is an involution.

Let $x^i \in \mathrm{Per}_k$. Since no special words appear, g acts as the identity on x^i, so g is in the kernel of the action on $\mathrm{Per}_k /\langle \sigma \rangle$.

Lastly, we show that $g \cdot [w]^m \subset [u]^m$. Let $y \in [w]^m$. Since y is left-k-periodic, the first occurrence of a special word in y is $a^r a^{rp} a^r w$ at $-(2r + rp)k$. Thus,

$$gy = \cdots a\widetilde{u}v.w \cdots .$$

Applying the cocycle α gives

$$g \cdot y = \cdots \widetilde{a}.uvw \cdots \in [u]^m$$

where $\widetilde{a}$ is the initial k-block of u.

Suppose now w is the initial word of $\widetilde{u}$. We can partition $[w]^m$ by the finitely many allowable extensions of w given by wb, where each wb is of length $|\widetilde{u}|$. Applying the process above gives marker automorphisms g_b for each extension wb. Since special words

for each g_b have the same length and are distinct, the involutions g_b all commute. The composition of $\{g_b\}$ is well defined, and is an involution that maps $[w]^m$ into $[u]^m$.

Case 2: x^m is not synchronizing, and there is an $h \in \text{Aut}(X)$ such that $h x^m$ is synchronizing.

Let $h \in \text{Aut}(X)$ where $h x^m = x^i$ is synchronizing, and let $[w]^m, [u]^m$ satisfy the hypothesis. Consider the sets

$$h[w]^m \quad \text{and} \quad h[u]^m.$$

We can partition them into finitely many cylinder sets, so by the previous construction, there is some $g \in N$ that maps

$$g \cdot h[w]^m \subset h[u]^m.$$

Then

$$h^{-1} g h \cdot [w]^m \subset [u]^m$$

with $h^{-1} g h \in N$, as desired. □

Recall that $N \lhd \text{Aut}(X)$ is the kernel of $\pi : \Omega \to \text{Per}_k / \langle \sigma \rangle$, so the automorphism produced above is contained in N.

COROLLARY 3.6. *Let (X, σ) be a transitive sofic shift and $k \in \mathbb{N}$. Suppose that X contains a left-k-periodic point. Then the following statements hold.*
(1) *The action of N on Ω^m, as defined in (6), is minimal.*
(2) *The action of $\text{Aut}(X)$ on Ω, as defined in §3.1, is minimal.*
(3) *The N action on Ω^m is extremely proximal.*
(4) *The $\text{Aut}(X)$ action on Ω is strongly proximal.*

Proof. (1) It suffices to show that the N-orbits of any non-empty open subset U cover all of Ω^m. Let $[u]^m \subset U$ be a non-empty cylinder set with $|u| \geq k$ and $[w]^m$ be an non-empty cylinder with $|w| \geq k$. By Proposition 3.5, there exists an involution $g \in N$ such that $g \cdot [w]^m \subset [u]^m$. Since $g = g^{-1}$, we have $[w]^m \subset g \cdot [u]^m$. As $[u]^m$ was arbitrary, this shows that $\bigcup_{g \in N} g \cdot U = \Omega^m$.

(2) Let $U \subset \Omega$ be an open set. If the intersections $U \cap \Omega^m$ are all non-empty, then by part (1), the N-orbit of U covers Ω. Suppose U is contained in some Ω^m. The action of $\text{Aut}(X)$ on Syn_k is transitive, so there exists $g \in \text{Aut}(X)$ such that $gU \cap \Omega_n$ is non-empty for any Ω_n. By part (1), the $\text{Aut}(X)$-orbit of U covers Ω.

(3) Each Ω^m contains more than two points, and cylinder sets form a subbase that generates the topology on Ω^m. In addition, Ω^m is compact, so any closed set is covered by finitely many cylinder sets. By Proposition 3.5, the N action on each Ω^m is extremely proximal.

(4) As extremely proximal actions are also strongly proximal, by part (3), the N action on each Ω^m is strongly proximal. Thus, the product action of N on $\prod_{m=1}^{j} \Omega^m$ is also strongly proximal. Since the diagonal action of N on the product space is isomorphic (as continuous group actions) on Ω, it follows that the action of $\text{Aut}(X)$, which contains N, on Ω is also strongly proximal. □

We use a proposition of Furman which relates the kernel of boundary actions and normal amenable subgroups.

PROPOSITION 3.7. (Furman [**7**]) *Let G be a discrete group, and consider the following subgroups of G:*

(1) $N = \bigcap_{i \in I} \text{Ker}(G \to \text{Homeo}(X_i))$, *where I is the set of isomorphism classes of boundary actions on the set of G-spaces;*

(2) $\sqrt{G}$, *the group generated by all closed normal amenable subgroups in G.*

Then $N = \sqrt{G}$.

In particular, the kernel of any boundary action contains any normal amenable subgroup. We have now assembled the ingredients to prove Theorem 1.1.

Proof of Theorem 1.1. Let (X, σ) be a transitive sofic shift. As periodic points are dense, there exists some k such that X contains k-periodic points and $\text{Syn}_k \neq \emptyset$. Since X is not finite, X also contains left-k-periodic points.

Corollaries 3.6(2) and 3.6(4) show that Ω is an $\text{Aut}(X)$-boundary. By Proposition 3.7, any normal amenable subgroup of $\text{Aut}(X)$ is contained in the kernel of a boundary action.

An element in Ω is a section of the projection $\pi : \Omega \to \text{Per}_k /\langle \sigma \rangle$, so the kernel of $\text{Aut}(X)$ acting on Ω must be contained in the kernel of $\text{Aut}(X)$ acting on Q_k, the set of left-k-periodic points up to k. Thus, it follows from Lemma 3.1 that the kernel of the $\text{Aut}(X)$ action on Ω is precisely $\langle \sigma \rangle$, and we obtain the desired result.   □

4. *Higher dimensions*

We show that the direct analogue of Theorem 1.1 in higher dimensions fails by giving a counterexample and explain why the methods of proof do not generalize even with stronger hypotheses. Consistent with the definition of one-dimensional shifts given in §2, we define a $\mathbb{Z}^d$-shift to be a closed, translation-invariant subset of $\mathcal{A}^{\mathbb{Z}^d}$. A $\mathbb{Z}^d$-shift is an SFT if it can be described by forbidden patterns in $\mathcal{A}^{\mathcal{F}}$, for some finite set $\mathcal{F} \subset \mathbb{Z}^d$, and a $\mathbb{Z}^d$ sofic shift is a topological factor of a $\mathbb{Z}^d$ SFT. The automorphism group consists of self-homeomorphisms of the shift that commute with the shift maps, which can be identified with $\mathbb{Z}^d$.

Hochman [**11**] constructs a two-dimensional SFT $X \subset \mathcal{A}^{\mathbb{Z}^2}$, which is topologically mixing and has positive entropy. Hochman explicitly computes the automorphism group to be $\mathbb{Z}^2 \oplus \bigcup S_{i,j}$, where $\mathbb{Z}^2$ is generated by the shift maps and $\bigcup S_{i,j}$ is a directed union of infinitely many finite groups, arising from higher-dimensional marker automorphisms. Amenability is closed under taking direct limits and sums; thus, the automorphism group is amenable. In higher dimensions, Ryan's theorem holds [**11**], and the center is the subgroup generated by the shifts, $\mathbb{Z}^2$. In particular, $\text{Aut}(X)$ has normal amenable subgroups that are not contained in the center. While this shift is topologically mixing, the set of periodic points is not dense, which suggests this may not be the right condition to impose.

There are various notions of uniform mixing in higher dimensions (for example, strongly irreducible, uniform filling, and block gluing), each of which imply that periodic points are dense. In each case, if two allowable patterns are sufficiently far apart, there is another allowable pattern which agrees with the original patterns; the distinct notions of

uniform mixing depend on the shape of patterns we consider. In contrast, for $d = 1$, these definitions of uniform mixing are equivalent to topological mixing.

However, even with dense periodic points, we cannot construct a topological boundary for uniformly mixing $\mathbb{Z}^d$ SFTs as we did in the one-dimensional case. Because there are now more directions of periodicity, we cannot construct a space on which the automorphism group acts in the same manner. More specifically, we cannot define a $\mathbb{Z}^d$ cocycle as we did in equation (1).

In the case of the higher-dimensional full shift, Frisch, Schlank, and Tamuz [6] show that any normal amenable subgroup must be contained in the subgroup generated by the shifts; unfortunately, their methods do not generalize to uniformly mixing SFTs. They construct a class of automorphisms of $\mathcal{A}^{\mathbb{Z}^d}$, induced by automorphisms of $\mathcal{A}^{\mathbb{Z}}$, which act independently on bi-infinite sequences of a configuration $x \in \mathcal{A}^{\mathbb{Z}^d}$. This relies strongly on the fact that in the full shift, there are no forbidden blocks. In a more general $\mathbb{Z}^d$-SFT, acting independently on lower-dimensional subspaces may produce forbidden patterns. We note that higher-dimensional marker automorphisms cannot arise from such a construction.

REFERENCES

[1] R. L. Adler and B. Marcus. Topological entropy and equivalence of dynamical systems. *Mem. Amer. Math. Soc.* **20**(219) (1979).

[2] M. Boyle. Lower entropy factors of sofic systems. *Ergod. Th. & Dynam. Sys.* **3**(4) (1983), 541–557.

[3] M. Boyle and W. Krieger. Periodic points and automorphisms of the shift. *Trans. Amer. Math. Soc.* **302**(1) (1987), 125–149.

[4] 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.

[5] A. Ehrenfeucht and D. M. Silberger. Periodicity and unbordered segments of words. *Discrete Math.* **26**(2) (1979), 101–109.

[6] J. Frisch, T. Schlank and O. Tamuz. Normal amenable subgroups of the automorphism group of the full shift. *Ergod. Th. & Dynam. Sys.* **39**(5) (2019), 1290–1298.

[7] A. Furman. On minimal strongly proximal actions of locally compact groups. *Israel J. Math.* **136** (2003), 173–187.

[8] S. Glasner. Topological dynamics and group theory. *Trans. Amer. Math. Soc.* **187** (1974), 327–334.

[9] S. Glasner. *Proximal Flows (Lecture Notes in Mathematics, 517)*. Springer, New York, 1976.

[10] G. A. Hedlund. Endomorphisms and automorphisms of the shift dynamical system. *Math. Systems Theory* **3** (1969), 320–375.

[11] M. Hochman. On the automorphism groups of multidimensional shifts of finite type. *Ergod. Th. & Dynam. Sys.* **30**(3) (2010), 809–840.

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

[13] B. P. Kitchens. *Symbolic Dynamics: One-sided, Two-sided and Countable State Markov Shifts*. Springer, Berlin, 1998.

[14] D. Lind and B. Marcus. *An Introduction to Symbolic Dynamics and Coding*. Cambridge University Press, Cambridge, 1995.

[15] R. C. Lyndon. On Burnside's problem. *Trans. Amer. Math. Soc.* **77** (1954), 202–215.

[16] J. P. Ryan. The shift and commutativity. *Math. Systems Theory* **6** (1972), 82–85.

[17] J. P. Ryan. The shift and commutativity II. *Math. Systems Theory* **8**(3) (1974), 249–250.

[18] B. Weiss. Subshifts of finite type and sofic systems. *Monatsh. Math.* **77** (1973), 462–474.