

On the geometry and regularity of invariant sets of piecewise-affine automorphisms on the Euclidean space

C. SİNAN GÜNTÜRK† and NGUYEN T. THAO‡

† *Courant Institute of Mathematical Sciences, NYU, 251 Mercer Street, New York, NY 10012, USA*

(*e-mail: gunturk@courant.nyu.edu*)

‡ *City College of New York, CUNY, Convent Avenue at 138th Street, New York, NY 10031, USA*

(*e-mail: tnguyen@ccny.cuny.edu*)

(Received 28 April 2016 and accepted in revised form 2 November 2018)

Abstract. In this paper, we derive geometric and analytic properties of invariant sets, including orbit closures, of a large class of piecewise-affine maps T on \mathbb{R}^d . We assume that (i) T consists of finitely many affine maps defined on a Borel measurable partition of \mathbb{R}^d , (ii) there is a lattice $\mathcal{L} \subset \mathbb{R}^d$ that contains all of the mutual differences of the translation vectors of these affine maps, and (iii) all of the affine maps have the same linear part that is an automorphism of \mathcal{L} . We prove that finite-volume invariant sets of such piecewise-affine maps always consist of translational tiles relative to this lattice, up to some multiplicity. When the partition is Jordan measurable, we show that closures of bounded orbits of T are invariant and yield Jordan measurable tiles, again up to some multiplicity. In the latter case, we show that compact T -invariant sets also consist of Jordan measurable tiles. We then utilize these results to quantify the rate of convergence of ergodic averages for T in the case of bounded single tiles.

Key words: arithmetic and algebraic dynamics, low-dimensional dynamics, tilings, topological dynamics

2010 Mathematics Subject Classification: 37Exx (Primary); 37Dxx, 11K38 (Secondary)

1. Introduction

1.1. *Piecewise maps of affine automorphisms.* This paper studies geometric and analytic properties of orbit closures of a class of piecewise-affine maps on the Euclidean space which we call ‘piecewise-affine automorphisms’. Consider a general piecewise-affine

map on \mathbb{R}^d given by

$$T(v) = A_i(v) := L_i v + \tau_i \quad \text{if } v \in \Omega_i, \quad (1)$$

where $\{\Omega_i\}_{i \in \Lambda}$ is a partition of \mathbb{R}^d and (L_i, τ_i) are the linear transformation and translation components of the affine maps A_i , $i \in \Lambda$, associated with this partition. For the class of piecewise-affine maps considered in this paper, we assume that:

- (i) Λ is finite and all the Ω_i are Borel measurable;
- (ii) there is a lattice $\mathcal{L} \subset \mathbb{R}^d$ for which $\{\tau_i - \tau_j : i, j \in \Lambda\} \subset \mathcal{L}$; and
- (iii) $L_i = L$ for all $i \in \Lambda$, where L is an automorphism on \mathcal{L} .

Fixing a basis for \mathcal{L} , we may assume that L is represented by a unimodular matrix, i.e., an element of $\text{GL}(d, \mathbb{Z})$ which we also denote by L . In this basis, the coefficients of all the τ_i would belong to a single coset of \mathbb{Z}^d which we may identify with a point $a \in \mathbb{T}^d := \mathbb{R}^d / \mathbb{Z}^d$. Hence no generality is lost by assuming that $\mathcal{L} = \mathbb{Z}^d$.

With these assumptions, T has a factor on \mathbb{T}^d given by the invertible affine map (also called an *affine automorphism* [Dan00])

$$S(u) := Lu + a, \quad u \in \mathbb{T}^d. \quad (2)$$

More precisely, the canonical projection $\langle \cdot \rangle : \mathbb{R}^d \rightarrow \mathbb{T}^d$ induces a semiconjugacy from (\mathbb{R}^d, T) to (\mathbb{T}^d, S) via the intertwining relation

$$\langle T(v) \rangle = S(\langle v \rangle), \quad v \in \mathbb{R}^d. \quad (3)$$

We will say that T is a piecewise extension of the affine automorphism S (or, in short, a *piecewise-affine automorphism*), and denote the set of all piecewise extensions of S by $\mathcal{P}(S) := \mathcal{P}(L, a)$.

Given an affine automorphism S , we are interested in structural properties of bounded orbits of all $T \in \mathcal{P}(S)$. It should be noted that piecewise-affine maps can exhibit extremely complicated behavior. Indeed, the problem of determining whether all trajectories of a given piecewise-affine map remain bounded is known to be algorithmically undecidable, even when the model class of piecewise-affine maps only incorporate finite partitions that are determined by linear inequalities with rational coefficients [BBKT01, DDBB⁺09]. Hence, it can be said that, generally, the boundedness question for orbits must be established by exploiting special properties of particular maps. Some cases are trivial, such as when each affine piece A_i is contracting. However, in our case with unimodular L , the A_i generically yield unbounded orbits as individual maps on \mathbb{R}^d . Therefore boundedness of orbits of T can only follow from the fine interactions of the A_i with the partition domains Ω_i . This paper will not address the question of deciding when orbits are bounded (sufficient conditions in some special classes of examples that fall in our setting can be found in [DD03, DGWY10, Yil02]), but rather we will consider what bounded orbits (i.e., their closures) generally look like.

Due to (3), orbits of S already give partial information about orbits of T . Indeed, every orbit of T is determined by an orbit of S up to a translation by a sequence in \mathbb{Z}^d . As is well known, the nature of the orbits of S depends significantly on the spectrum of L . For example, when L is unipotent, every orbit of S is dense in a finite union of cosets of some subtorus of \mathbb{T}^d , a result that falls within Ratner's theory [Tao12, §1.1]. On the other hand, when L is hyperbolic, some orbits exhibit complicated behavior, even though

ergodicity with respect to the Haar measure on \mathbb{T}^d implies that almost every orbit is uniformly distributed in \mathbb{T}^d . While most of the general results of this paper will hold without any spectral constraints, only building on a suitable ergodic decomposition of the system (\mathbb{T}^d, S) , some of the more specific ones concerning individual orbits can readily utilize additional information that may follow from spectral properties.

1.2. *Summary of the main results.* We say that a set A is T -invariant if $T(A) = A$ and essentially T -invariant if $T(A) \Delta A$ is null with respect to the measure of interest. (The reason for our choice of convention regarding invariant sets is explained in §3.3.) The theoretical contributions of this paper can then be summarized as follows. Given any piecewise-affine automorphism T satisfying (i)–(iii) above, we show that:

- geometrically, any essentially T -invariant set of finite measure consists of a disjoint union of one or more suitably defined \mathbb{Z}^d -tiles (see paragraph below);
- analytically, closures of bounded orbits of T are essentially T -invariant (and therefore consist of \mathbb{Z}^d -tiles) provided the partition $\{\Omega_i\}_{i \in \Lambda}$ is not too complicated (e.g., Jordan measurable); and
- statistically, in the case of single tiles, the rate of convergence of ergodic averages on any given orbit of T can be controlled by the underlying affine automorphism S and the regularity of the partition $\{\Omega_i\}_{i \in \Lambda}$.

Our results are simplest to state when S is ergodic on \mathbb{T}^d . In this case, the first two contributions listed above are based on the following results proved in this paper.

(T) (*Tiling*) Any essentially T -invariant set A of finite Lebesgue measure is an essential m -tile for some integer $m \geq 0$. This means that there exist disjoint sets $\Gamma_1, \dots, \Gamma_m$ in \mathbb{R}^d such that each Γ_i is a fundamental domain for the lattice \mathbb{Z}^d (i.e., $\{\Gamma_i + k\}_{k \in \mathbb{Z}^d}$ partition \mathbb{R}^d) and

$$\lambda_d(A \Delta \Gamma) = 0,$$

where $\Gamma := \Gamma_1 \cup \dots \cup \Gamma_m$ (Γ is called an exact m -tile) and λ_d is the d -dimensional Lebesgue measure. If the multiplicity m equals one, then T is ergodic on Γ . (For details, see Theorem 3.1.)

(R) (*Regularity of invariant sets and orbit closures*) If the partition $\{\Omega_i\}_{i \in \Lambda}$ is Jordan measurable in the sense that

$$\lambda_d\left(\bigcup_{i \in \Lambda} \partial \Omega_i\right) = 0,$$

then, for almost every v_0 , the closure of the forward orbit $V = \mathcal{O}_T^+(v_0) := \{T^n(v_0) : n \geq 0\}$ is essentially T -invariant and Jordan measurable. More generally, every compact essentially T -invariant set K is Jordan measurable. Moreover, it is possible to find a Jordan measurable exact m -tile Γ that is a measure-equivalent subset of K (or \bar{V}) and can be arranged to contain V in the latter case. (For details, see Theorems 4.2, 5.1, and 5.2.)

For the general case, i.e., without ergodicity of S , we show that a similar picture holds, albeit with some modifications. The starting point is a partitioning of \mathbb{T}^d (and uniquely so) into S -invariant sets (Π_α) such that each element Π of this partition is a finite union of subtoral cosets and S is ergodic (but not necessarily uniquely ergodic) with respect to the

uniform (surface) measure μ_Π on Π . Let $\tilde{\Pi} \subset \mathbb{R}^d$ be the preimage of Π under $\langle \cdot \rangle$, and let $\mu_{\tilde{\Pi}}$ be the corresponding uniform measure on $\tilde{\Pi}$. (One can also view $\mu_{\tilde{\Pi}}$ as the restriction of the k -dimensional Hausdorff measure on the Borel sets of $\tilde{\Pi}$, where $k = \dim(\Pi)$; see §3 for details.) In this context, we establish a generalized notion of tiling for orbits of T within $\tilde{\Pi}$ and show that **(T)** and **(R)** continue to hold for $\mu_{\tilde{\Pi}}$ -almost everywhere (a.e.) bounded orbits of T and, in particular, for every bounded orbit whose toral projection is dense in Π . It should be noted that, in the unipotent case, the above mentioned partition of \mathbb{T}^d is a partition into minimal sets for S . However, in the general setting, it is not possible to match *all* orbit closures of S with finite unions of subtoral cosets in \mathbb{T}^d . Furthermore, there may not be any decomposition of \mathbb{T}^d into minimal sets for S either.

It is natural to ask whether unimodularity of L is necessary, i.e., if the tiling property continues to hold when L is merely an endomorphism of \mathbb{T}^d . While certain partitions may still result in the tiling property, it turns out that generally this is not the case (see Example 7.6).

1.2.1. *Regularity of invariant tiles and convergence rate of ergodic averages.* The most significant implication of the regularity analysis of invariant tiles is that it enables us to derive an ‘effective ergodic theorem’ for T (at least in the single tile case), i.e., a quantitative bound on the convergence rate of ergodic averages along orbits of T based on a suitable quantitative measure of regularity. For simplicity, consider the case when S is ergodic on $\Pi = \mathbb{T}^d$ and Γ (as given by **(R)**) is a Jordan measurable, essentially T -invariant, exact 1-tile for \mathbb{R}^d . In this case, the projection $\langle \cdot \rangle$ (when restricted to Γ) defines a measure-preserving isomorphism between Γ (with the Lebesgue measure restricted to the Borel subsets of Γ) and \mathbb{T}^d (with the Haar measure), and the intertwining relation (3) implies that T is ergodic on Γ . Hence, for any $f \in L^1(\Gamma)$ and for almost every $v_0 \in \Gamma$, the sequence of iterates $\mathbf{v} = (v_n)_0^\infty$, where $v_n := T^n(v_0)$, is contained in Γ and yields

$$D_N(f, \mathbf{v}) := D_N(f, \mathbf{v}, \Gamma) := \left| \frac{1}{N} \sum_{n=0}^{N-1} f(v_n) - \int_\Gamma f(v) \, dv \right| \rightarrow 0 \quad \text{as } N \rightarrow \infty. \quad (4)$$

The quantitative bound that we establish in this paper on the rate of decay of $D_N(f, \mathbf{v})$ incorporates two ingredients: a regularity estimate for the invariant tile Γ and a regularity estimate for the function f . Regarding the first, let us define

$$\rho_\Gamma(\varepsilon) := \lambda_d(\mathcal{N}_\varepsilon(\partial\Gamma)), \quad \varepsilon > 0, \quad (5)$$

where $\mathcal{N}_\varepsilon(B)$ denotes the (open) ε -neighborhood of a set B . Note that Jordan measurability of Γ implies *a priori* that $\lim_{\varepsilon \rightarrow 0^+} \rho_\Gamma(\varepsilon) = 0$. Regarding the second ingredient, let $f : \Gamma \rightarrow \mathbb{C}$ be uniformly continuous and let ω_f denote its modulus of continuity (with respect to the Euclidean metric on Γ). Note that, again, the uniform continuity of f means *a priori* that $\lim_{\varepsilon \rightarrow 0^+} \omega_f(\varepsilon) = 0$. With these two ingredients, we establish the quantitative bound

$$D_N(f, \mathbf{v}) \lesssim_d \omega_f(\sqrt{d} \lfloor D_N(\mathbf{u})^{-1} \rfloor^{-1/d}) + \|f\|_\infty \rho_\Gamma(\sqrt{d} \lfloor D_N(\mathbf{u})^{-1} \rfloor^{-1/d}), \quad (6)$$

where $\mathbf{u} := (\langle v_n \rangle)_0^\infty$ denotes the projection of \mathbf{v} on \mathbb{T}^d (which itself results in an orbit of S), and $D_N(\mathbf{u})$ denotes the standard discrepancy (with respect to axis-parallel rectangles) of

the first N terms of \mathbf{u} in \mathbb{T}^d . (Here the notation $\lfloor \cdot \rfloor$ denotes the integer part, and $A \lesssim_d B$ means $A \leq C B$, where C is a numerical constant that depends on d only.)

Note that this is a general-purpose upper bound. But it is ready to be turned into an effective bound with additional information on f , Γ and \mathbf{u} . In particular, bounds on $D_N(\mathbf{u})$ can be obtained by standard analytic tools, such as the Erdős–Turán inequality (see, e.g., [KN74]), which can exploit the algebraic structure of S .

In the special case when S is uniquely ergodic, note that $D_N(\mathbf{u}) \rightarrow 0$ uniformly over all $u_0 \in \mathbb{T}^d$. Hence we get that, for every f continuous on $\overline{\Gamma}$, $D_N(f, \mathbf{v}) \rightarrow 0$ uniformly over all v_0 for which \mathbf{v} is contained in Γ . This is a close approximation to unique ergodicity for T on Γ , even if it does not hold *per se*.

1.3. *Relation to other work on piecewise maps.* There is a large body of literature on piecewise maps and naturally this paper has relations to many of them. The piecewise-affine automorphisms studied in this paper admit absolutely continuous invariant measures with density equal to the indicator function of their invariant sets. Starting with [LY73, AY84] and [GB89], piecewise *expanding* maps (uniformly or in area) have been studied extensively in terms of the refinements of the sufficient conditions on the set of discontinuities that guarantee existence of absolutely continuous invariant measures. However, our maps are not area-expanding, but, instead, are (locally) area-preserving. In this respect, they match more closely with piecewise isometries such as piecewise translations and rotations, piecewise parabolic maps [AFNZ00, AFL09] and with area-preserving piecewise hyperbolic maps. The work [ZL13] concerning invariant measures with bounded variation density for general piecewise area-preserving maps is of particular relevance to this paper. In one dimension and for partitions consisting of intervals, the maps that we consider simply reduce to interval translation maps [BT03] on their full domain and, furthermore, to interval exchange transformations on their (bounded) invariant sets. In higher dimensions, certainly a much wider variety of spectral possibilities is present.

The tiling property of invariant sets has been known for some time for a restricted class of maps in which L is a skew transformation. This was noted in [DD03, GT04] and was proved in [GT05]. The case of piecewise translations (i.e., when L is the identity transformation) falls under the work of Adler *et al* [AKM⁺05, ANST10, ANS⁺17]. To the best of our knowledge, the present paper is the first work that extends these partial results to all $L \in GL(d, \mathbb{Z})$.

1.4. *Organization of the paper.* In §2, we present the background material concerning dynamics of affine automorphisms on \mathbb{T}^d which will be needed in the subsequent sections. Section 3 contains our general tiling theorem for invariant sets of piecewise-affine automorphisms and is central to this paper. The results of this section are derived based on measure-theoretic principles only. In §4, after deriving some general properties of orbit closures of piecewise homeomorphisms, we show that closures of bounded orbits of piecewise-affine automorphisms have the tiling property whenever the partitions associated with these maps are Jordan measurable. We further this analysis in §5, where we show that any bounded orbit closure and, more generally, any compact essentially invariant set, is Jordan measurable. In §6, building on all the tools and results of this paper,

we derive quantitative bounds on the convergence rate of the ergodic averages associated with T . Section 7 is devoted to selected examples from algorithmic analog-to-digital (A/D) conversion that have motivated this paper and concludes with a list of open problems and challenges. Some of our more technical lemmas (which are non-dynamical) are presented in Appendix A. In addition, Appendix B is concerned with extracting invariant sets of finite-to-one maps.

2. Preliminaries on toral dynamics of affine automorphisms

In this section, we will discuss the structure of generic orbits of affine automorphisms of the torus by means of a recursive, ‘semi-explicit’ ergodic decomposition, the meaning of which will be made precise below. Let S be an affine automorphism of \mathbb{T}^d given by $S(u) := Lu + a$, where L is a unimodular matrix and a is an arbitrary element of \mathbb{T}^d . It is well known (e.g., [Bro76, Theorem 3.3]) that S is ergodic on \mathbb{T}^d (with respect to the Haar measure) if and only if whenever k is in \mathbb{Z}^d :

- (i) $(L^\top)^n k = k$ for $n > 0$ implies $L^\top k = k$; and
- (ii) $L^\top k = k$ and $k \cdot a = 0 \pmod{1}$ implies $k = 0$.

When S is not ergodic, it is natural to seek an ergodic decomposition. However, a ‘general-purpose’ decomposition of the Haar measure on \mathbb{T}^d via Choquet’s theorem does not reveal the nature of the ergodic components explicitly, so we will take a more direct approach that also extracts geometric information on the ergodic components of interest to us. As will be discussed below, these ergodic components turn out to be uniform measures supported on certain lower-dimensional submanifolds Π of the form $G + P$, where G is a subtorus and P is a finite subset of \mathbb{T}^d . In other words, each such Π is a finite union of cosets of a subtorus of \mathbb{T}^d . Note that Π determines G uniquely (since any connected component of Π is a coset of G) but determines P only up to translations by elements of G . The uniform probability measure on Π will be denoted by μ_Π . (In other words, μ_G is the same as the Haar measure on G and μ_Π is the normalized sum of translated copies of μ_G on the cosets that constitute Π .) The main result that we will show in this section is the following proposition.

PROPOSITION 2.1. *There exists a unique partition of \mathbb{T}^d into S -invariant subsets (Π_α) , where each Π_α is a finite union of cosets of some subtorus and S is ergodic on Π_α with respect to μ_{Π_α} .*

This result will follow as a consequence of several basic facts concerning dynamics of affine toral automorphisms. Let us define the following classes of sets that will be used frequently in this paper.

$$\mathcal{C} := \{G + P : G \text{ is a subtorus and } P \text{ is a finite subset of } \mathbb{T}^d\}, \quad (7)$$

$$\mathcal{C}_S := \{\Pi \in \mathcal{C} : \Pi \text{ is } S\text{-invariant}\}, \quad (8)$$

$$\mathcal{C}_S^e := \{\Pi \in \mathcal{C}_S : S \text{ is ergodic on } \Pi \text{ with respect to } \mu_\Pi\}. \quad (9)$$

We now discuss what it means for a set $\Pi = G + P \in \mathcal{C}$ to be S -invariant. The relation

$$S(u + v) = L(u) + S(v),$$

which is valid for all u and v , implies that $S(\Pi) = L(G) + S(P)$. Since any subtorus G is of the form $(V + \mathbb{Z}^d)/\mathbb{Z}^d$, where V is a linear subspace of \mathbb{R}^d defined by rational equations and $L(\mathbb{Z}^d) = \mathbb{Z}^d$, we know that $L(G)$ is also a subtorus. Hence $S(\Pi) \in \mathcal{C}$. Matching the connected components of Π and $S(\Pi)$, we get that $G + P$ is S -invariant if and only if G is L -invariant and P is S -invariant modulo translations by G .

When G is an L -invariant subtorus of \mathbb{T}^d , we have $S(G + p) = G + S(p)$ for all $p \in \mathbb{T}^d$. This relation induces a map $G + p \mapsto G + S(p)$ on \mathbb{T}^d/G , which we may denote by S_G . Let $\pi : \mathbb{T}^d \rightarrow \mathbb{T}^d/G$ be the canonical projection given by $\pi(p) := G + p$. With this notation, we can equivalently say that $G + P$ is S -invariant if and only if G is L -invariant and $\pi(P)$ is S_G -invariant. Since P is finite, the latter holds if and only if $\pi(P)$ decomposes into finitely many distinct periodic S_G -orbits of cosets of G . Note that the coset $S_G(G + p)$ is the same as the set $S(G + p)$, i.e., we have the commutation relation $S_G \circ \pi = \pi \circ S$. Hence periodicity of $(S^n(G + p))$ will mean periodicity of $(S_G^n(\pi(p)))$.

With these elementary observations, we are ready to proceed. The proof of Proposition 2.1 will rely on the following crucial lemma.

LEMMA 2.1. *Suppose G is an L -invariant subtorus of \mathbb{T}^d and $(S^n(G))$ is periodic. If S is not ergodic on the orbit $\mathcal{O}_S(G)$, then there exists an L -invariant proper subtorus G' of G such that $(S^n(G' + p))$ is periodic for every $p \in G$.*

Proof. Let $e_k(u) := e^{2\pi i k \cdot u}$ for $u \in \mathbb{T}^d$, $k \in \mathbb{Z}^d$. Note that the system $\{e_k\}_{k \in \mathbb{Z}^d/G^\perp}$ (with domain restricted to G) forms an orthonormal basis of $L^2(G)$, where, for notational ease, we represent each coset $G^\perp + k$ with a unique element in it. Denote $S^j(0)$ by s_j , $j \in \mathbb{Z}$, so that $S^j(G) = G + s_j$. Define

$$\varphi_{j,k}(u) := 1_{S^j(G)}(u)e_k(u - s_j) \quad \text{for } u \in \mathbb{T}^d, j \in \mathbb{Z}, k \in \mathbb{Z}^d,$$

where $1_{S^j(G)}$ denotes the indicator function of $S^j(G)$.

Let M be the period of $(S^n(G))$ and denote $\mathcal{O}_S(G) = G \cup (G + s_1) \cdots \cup (G + s_{M-1})$ by Π . Then the system

$$\{\varphi_{j,k} : j = 0, \dots, M - 1, k \in \mathbb{Z}^d/G^\perp\}$$

(with domain restricted to Π) is an orthonormal basis of $L^2(\Pi)$. Note that, by periodicity, we have $s_{j+M} - s_j \in G$ for all j . We also have $\varphi_{j+M,k}(u) = e_k(s_j - s_{j+M})\varphi_{j,k}(u)$, i.e., these two functions are equal up to a phase factor. (However the phase factor need not vanish unless k is zero, i.e., in G^\perp .)

If S is not ergodic on Π , then there exists an S -invariant $f \in L^2(\Pi)$ that is not constant μ_Π -a.e. Consider the expansion

$$f = \sum_{j=0}^{M-1} \sum_{k \in \mathbb{Z}^d/G^\perp} \langle f, \varphi_{j,k} \rangle \varphi_{j,k}.$$

Observing that $1_{S^j(G)} \circ S = 1_{S^{j-1}(G)}$ and $e_k(Su - s_j) = e_{L^\top k}(u - s_{j-1})$,

$$\varphi_{j,k} \circ S = \varphi_{j-1, L^\top k} \quad \text{for all } j \in \mathbb{Z} \text{ and } k \in \mathbb{Z}^d,$$

so that

$$f \circ S = \sum_{j=0}^{M-1} \sum_{k \in \mathbb{Z}^d / G^\perp} \langle f, \varphi_{j,k} \rangle \varphi_{j-1, L^\top k}. \tag{10}$$

With L^\top -invariance of G^\perp , we also have

$$f = \sum_{j=0}^{M-1} \sum_{k \in \mathbb{Z}^d / G^\perp} \langle f, \varphi_{j, L^\top k} \rangle \varphi_{j, L^\top k} \tag{11}$$

so that equating the coefficients in (10) and (11), we get that, for all $k \in \mathbb{Z}^d / G^\perp$,

$$\langle f, \varphi_{j, L^\top k} \rangle = \begin{cases} \langle f, \varphi_{j+1, k} \rangle & \text{if } j = 0, \dots, M-2, \\ \langle f, \varphi_{0, k} \rangle e_k (s_{M-1} - s_{-1}) & \text{if } j = M-1. \end{cases} \tag{12}$$

Since f is not constant, there exists a non-zero element k_* in \mathbb{Z}^d / G^\perp for which $\langle f, \varphi_{j, k_*} \rangle \neq 0$ for some $j = 0, \dots, M-1$. The orbit $((L^\top)^n k_*)$ in \mathbb{Z}^d / G^\perp has to be finite (i.e., periodic) or else we would get a violation of $f \in L^2(\Pi)$.

Let N be the smallest positive integer such that $(L^\top)^N k_* = k_*$ and define

$$k_0 := \begin{cases} k_* & \text{if } N = 1, \\ L^\top k_* - k_* & \text{if } N > 1. \end{cases}$$

Note that $k_0 \neq 0$ and $(L^\top)^N k_0 = k_0$. Consider the closed subgroup

$$H := \{u \in G : k_0 \cdot u = (L^\top k_0) \cdot u = \dots = ((L^\top)^{N-1} k_0) \cdot u = 0 \pmod{1}\}.$$

We are going to show that $(S^n(H + p))$ is MN -periodic for any given $p \in G$. Since $(L^\top)^N k_0 = k_0$, H is L -invariant so that, for any coset $H + p$, we have $S^n(H + p) = H + S^n(p)$ for all n . Then $(S^n(H + p))$ is MN -periodic if and only if $S^{MN}(p) - p \in H$. Since $S^M(G) = G$, we already know that $S^{MN}(p) \in G$ and hence $S^{MN}(p) - p \in G$. To show that $S^{MN}(p) - p \in H$, we analyze the cases $N = 1$ and $N > 1$ separately.

Case $N = 1$. We have $L^\top k_0 = k_0$. Note that, for any $u \in \mathbb{T}^d$,

$$k_0 \cdot (S^M(u) - u) = k_0 \cdot ((L^M - I)u + (L^{M-1} + \dots + L + I)a) = M k_0 \cdot a$$

so that $k_0 \cdot (S^M(u) - u)$ is independent of u . Meanwhile, (12), together with the knowledge that $\langle f, \varphi_{j, k_0} \rangle \neq 0$ for some $j = 0, \dots, M-1$, implies that

$$\langle f, \varphi_{0, k_0} \rangle = \langle f, \varphi_{1, k_0} \rangle = \dots = \langle f, \varphi_{M-1, k_0} \rangle = \langle f, \varphi_{0, k_0} \rangle e_{k_0} (s_{M-1} - s_{-1}) \neq 0$$

so that $k_0 \cdot (S^M(s_{-1}) - s_{-1}) = k_0 \cdot (s_{M-1} - s_{-1}) = 0 \pmod{1}$. Thus, we also have $k_0 \cdot (S^M(p) - p) = 0 \pmod{1}$, which proves that $S^M(p) - p \in H$.

Case $N > 1$. For any integer $i \geq 0$,

$$\begin{aligned} ((L^\top)^i k_0) \cdot (S^{MN}(p) - p) &= k_0 \cdot (L^{MN+i} p - L^i p + (L^{MN-1} + \dots + L + I)L^i a) \\ &= ((L^\top)^{MN} k_0 - k_0) \cdot (L^i p) + ((L^\top)^{MN} k_* - k_*) \cdot (L^i a) \\ &= 0. \end{aligned}$$

This proves that $S^{MN}(p) - p \in H$.

Next, H is a proper closed subgroup of G (including the possibility that H is $\{0\}$) since $k_0 \neq 0$. As such, H can be written in the form $G' + F$, where G' is a proper subtorus of G and F is a finite subgroup of G . Since $H \in \mathcal{C}$ and is L -invariant, it follows from our earlier discussion that G' is L -invariant. We then see that $(S^n(G' + p))$ must be periodic because it remains within $\mathcal{O}_S(H + p)$, which can be expressed as a finite union of cosets of G' given by

$$\bigcup_{i=0}^{MN-1} G' + F + S^i(p),$$

due to the fact that $(S^n(H + p))$ is MN -periodic. □

COROLLARY 2.1. *Suppose G is an L -invariant subtorus of \mathbb{T}^d and $(S^n(G + p))$ is periodic for some $p \in \mathbb{T}^d$. If S is not ergodic on $\mathcal{O}_S(G + p)$, then there exists an L -invariant proper subtorus G' of G such that $(S^n(G' + p' + p))$ is periodic for every $p' \in G$.*

Proof. For the given p , consider the affine automorphism $S'(u) := Lu + (a + Lp - p)$. Then the commutation relation

$$S(\cdot + p) = S'(\cdot) + p \tag{13}$$

implies that (i) $(S^n(G + p))$ is periodic if and only if $(S^n(G))$ is periodic, and (ii) S is ergodic on $\Pi \in \mathcal{C}_S$ if and only if S' is ergodic on $\Pi - p \in \mathcal{C}_{S'}$. Hence, if S is not ergodic on $\mathcal{O}_S(G + p)$, then Lemma 2.1 (applied to S' on G) implies that there exists an L -invariant proper subtorus G' of G such that $(S^n(G' + p'))$ is periodic for every $p' \in G$. Therefore, by (13) again, $(S^n(G' + p' + p))$ is also periodic for every $p' \in G$. □

Proof of Proposition 2.1. Starting with $G = \mathbb{T}^d$ and $p = 0$, we apply Corollary 2.1 recursively to every periodic orbit $\mathcal{O}_S(G + p)$ until every branch of the process terminates, i.e., S is ergodic on the resulting periodic orbit. More precisely, each such branch is characterized by a finite sequence of periodic orbits

$$\mathcal{O}_S(G_0 + p_0), \quad \mathcal{O}_S(G_1 + p_1 + p_0), \quad \dots, \quad \mathcal{O}_S(G_r + p_r + \dots + p_0),$$

where $(G_0, p_0) = (\mathbb{T}^d, 0)$ and, for all $i \geq 1$ (provided $r \neq 0$), G_i is a proper subtorus of G_{i-1} and p_i is an arbitrary element of G_{i-1} . The termination (i.e., finiteness of r) is guaranteed because $\dim G_i < \dim G_{i-1}$. Each orbit $\mathcal{O}_S(G_i + p_i + \dots + p_0)$ is periodic, but S is ergodic only on $\mathcal{O}_S(G_r + p_r + \dots + p_0)$. In order to get distinct branches, we only consider p_i that are distinct modulo G_i .

The final result is a partitioning of \mathbb{T}^d into a collection of subsets (Π_α) , where each $\Pi_\alpha \in \mathcal{C}_S^e$.

Finally, we show that there is only one partition with this property. Suppose (Π_α) and (Π_β) are two such partitions of \mathbb{T}^d . The ergodicity and continuity of S implies that it has a dense orbit O_α in any Π_α . We know that O_α intersects with some Π_β , but S -invariance implies that $O_\alpha \subset \Pi_\beta$. Taking the closure, we have $\Pi_\alpha \subset \Pi_\beta$. By symmetry, there must exist $\Pi_{\alpha'}$ such that $\Pi_\beta \subset \Pi_{\alpha'}$. But then $\Pi_\alpha \subset \Pi_{\alpha'}$, which implies that $\alpha = \alpha'$, and therefore $\Pi_\alpha = \Pi_\beta$. Hence the two partitions are identical. □

3. Tiling of T -invariant sets

In this section, we will establish the first core result of this paper, namely, the tiling property of essentially invariant sets of piecewise-affine automorphisms. We start by defining the measures that will be relevant to us.

3.1. *The measures μ_Π and $\mu_{\tilde{\Pi}}$.* Let T be a piecewise-affine automorphism and let S be its toral factor. As we saw in §2, \mathbb{T}^d admits an ergodic decomposition into S -invariant sets (Π_α) , each of which is of the form $\Pi = G + P$ for some subtorus G and a finite set P . Here each Π is equipped with its uniform probability measure μ_Π defined on its Borel sets. Up to a constant factor, μ_Π is equal to the k -dimensional Hausdorff measure on \mathbb{T}^d restricted to the Borel sets of Π , where $k = \dim(G)$.

For any set $A \subset \mathbb{T}^d$, let \tilde{A} denote its preimage under the canonical map $\langle \cdot \rangle$. Note that $(\tilde{\Pi}_\alpha)$ is a partition of \mathbb{R}^d and that $T(\tilde{\Pi}_\alpha) \subset \tilde{\Pi}_\alpha$ for each α because of (3). Hence we may study T on each $\tilde{\Pi}_\alpha$ individually.

In the immediate discussion below, we will work with a general $\Pi = G + P \in \mathcal{C}$, but when the maps S and T are incorporated into our discussion, we will assume that $\Pi \in \mathcal{C}_S$ or, actually, is an element of the partition (Π_α) .

When $G = \mathbb{T}^d$, clearly, $\tilde{\Pi} = \mathbb{R}^d$. Otherwise, when G is a proper subtorus of \mathbb{T}^d (including the case $G = \{0\}$), $\tilde{\Pi}$ is a countable union of distinct cosets of the rational subspace V of \mathbb{R}^d , where $G = (V + \mathbb{Z}^d)/\mathbb{Z}^d$ and, alternatively, is a finite union of cosets of the subgroup $\tilde{G} := V + \mathbb{Z}^d$ of \mathbb{R}^d . We equip $\tilde{\Pi}$ with its own uniform surface measure $\mu_{\tilde{\Pi}}$ which assigns the same weight on every coset. Likewise, and again up to a constant factor, the resulting measure $\mu_{\tilde{\Pi}}$ is equal to the k -dimensional Hausdorff measure on \mathbb{R}^d restricted to the Borel sets of $\tilde{\Pi}$. While the exact normalization of $\mu_{\tilde{\Pi}}$ will not affect the main results of this paper, there is a natural choice that will help to simplify our notation. We would like $\mu_{\tilde{\Pi}}$ and μ_Π to be compatible in the sense that $\mu_{\tilde{\Pi}}(B) = \mu_\Pi(\langle B \rangle)$ for any Borel set $B \subset \tilde{\Pi}$ on which $\langle \cdot \rangle$ is one-to-one. This is a special case of the more general equality

$$\mu_{\tilde{\Pi}}(B) = \int_\Pi N_B d\mu_\Pi, \tag{14}$$

where $N_B : \mathbb{T}^d \rightarrow \mathbb{N} \cup \{\infty\}$ is the *multiplicity* function defined by

$$N_B(u) := |\{v \in B : \langle v \rangle = u\}| = |B \cap \langle \cdot \rangle^{-1}(u)|. \tag{15}$$

(Here $|\cdot|$ denotes cardinality.) Indeed, if $\langle \cdot \rangle$ is one-to-one on B , then $N_B = 1_{\langle B \rangle}$ so that (14) implies that $\mu_{\tilde{\Pi}}(B) = \mu_\Pi(\langle B \rangle)$.

Let us see why (14) is valid with a suitable normalization of $\mu_{\tilde{\Pi}}$. First, note that Borel measurability of N_B (for B Borel) follows from the continuity of the canonical map $\langle \cdot \rangle$. (See, for example, the general discussion in [Fed69, 2.2.13 and 2.10.10]. Alternatively, this fact follows from the explicit representation

$$N_B = \sum_{n \in \mathbb{Z}^d} 1_{\langle B \cap Q_n \rangle},$$

where $Q_n := n + [0, 1)^d$.) Next, let $\nu(B)$ denote the right-hand side of (14) for any Borel set $B \subset \tilde{\Pi}$. It is clear from (15) that $N_\emptyset = 0$, and whenever (B_n) is a disjoint countable

family we have $N_{\cup B_n} = \sum N_{B_n}$. Employing the monotone convergence theorem, we get that ν is a Borel measure. It remains to show that ν is uniform on $\tilde{\Pi}$. The identity $N_{B+x} = N_B(\cdot - \langle x \rangle)$, which is valid for any set B and any $x \in \mathbb{R}^d$, implies, in particular, that $\nu(B+x) = \nu(B)$ whenever $x \in \tilde{G}$ and B is a Borel set in any of the cosets $\tilde{G} + \tilde{p}$ that constitute $\tilde{\Pi}$. Hence ν restricted to $\tilde{G} + \tilde{p}$ is nothing but a translate of the Haar measure on \tilde{G} , which is, of course, the uniform measure and it is unique up to a constant multiple. Since μ_Π has the same weight on each coset $G + p$ that constitutes Π , ν also has the same weight on each $\tilde{G} + \tilde{p}$. This shows that $\mu_{\tilde{\Pi}}$ can be normalized to equal ν and therefore (14) has been established.

Note that, for any Borel set B in $\tilde{\Pi}$, $N_B = 0$ μ_Π -a.e. if and only if $\mu_{\tilde{\Pi}}(B) = 0$. More generally, $N_{B_1} = N_{B_2}$ μ_Π -a.e. if and only if $\mu_{\tilde{\Pi}}(B_1 \Delta B_2) = 0$.

3.2. *Interaction between T and $\mu_{\tilde{\Pi}}$.* Let $\Pi \in \mathcal{C}_S$, i.e., μ_Π remain invariant by S . Note that while $\mu_{\tilde{\Pi}}$ is invariant under each affine component of T , it need not be invariant under T itself when the set of points with multiple preimages is not negligible, which is typically the case. We show below that T preserves the measure of a set if and only if it is ‘essentially one-to-one’ on this set.

For a Borel set $B \subset \tilde{\Pi}$, we say that T is *essentially one-to-one* on B if the set of points $v \in B$ for which $T(v)$ has more than one preimage in B is $\mu_{\tilde{\Pi}}$ -null. To make this notion more precise, let us first define, for an arbitrary non-empty set A in $\tilde{\Pi}$

$$m_A := \min_{v \in A} |T^{-1}(T(v)) \cap A| \quad \text{and} \quad M_A := \max_{v \in A} |T^{-1}(T(v)) \cap A|. \tag{16}$$

Clearly, $|T^{-1}(T(v)) \cap A| \geq 1$ for all $v \in A$ so that we have $M_A \geq m_A \geq 1$. We set $M_\emptyset := m_\emptyset := 1$ (instead of the convention that would set the maximum of a function over the empty set equal to $-\infty$ and the minimum equal to $+\infty$). Hence T is one-to-one on A if and only if $M_A = 1$. Also

$$m_A |T(A)| \leq |A| \leq M_A |T(A)| \tag{17}$$

for all sets. Similarly, for any Borel set B in $\tilde{\Pi}$ of non-zero measure, we define

$$\ell_B := \text{ess inf}_{v \in B} |T^{-1}(T(v)) \cap B| \quad \text{and} \quad L_B := \text{ess sup}_{v \in B} |T^{-1}(T(v)) \cap B|, \tag{18}$$

both with respect to $\mu_{\tilde{\Pi}}$. It is clear that $L_B \geq \ell_B \geq 1$. If $\mu_{\tilde{\Pi}}(B) = 0$, we set $L_B := \ell_B := 1$ (instead of the conventional values of $-\infty$ and $+\infty$). With this convention, it follows that T is essentially one-to-one on B if and only if $L_B = 1$.

Note that

$$1 \leq m_B \leq \ell_B \leq L_B \leq M_B \leq |\Lambda| < \infty,$$

where Λ denotes the index set labeling the partition of the domain of T , as introduced in §1.1. If we define

$$B^* := \{v \in B : \ell_B \leq |T^{-1}(T(v)) \cap B| \leq L_B\}, \tag{19}$$

then we have

$$\mu_{\tilde{\Pi}}(B \setminus B^*) = 0; \quad \ell_B = m_{B^*} \quad \text{and} \quad L_B = M_{B^*}. \tag{20}$$

LEMMA 3.1. Let $\Pi \in \mathcal{C}_S$. Then, for all Borel sets $B \subset \tilde{\Pi}$,

$$\ell_B \mu_{\tilde{\Pi}}(T(B)) \leq \mu_{\tilde{\Pi}}(B) \leq L_B \mu_{\tilde{\Pi}}(T(B)), \tag{21}$$

and the following are equivalent when $\mu_{\tilde{\Pi}}(B) < \infty$.

- (a) $\mu_{\tilde{\Pi}}(T(B)) = \mu_{\tilde{\Pi}}(B)$.
- (b) $N_{T(B)} = N_B \circ S^{-1}$ μ_{Π} -a.e.
- (c) T is essentially one-to-one on B .

The implications (c) \Rightarrow (a) and (b) \Rightarrow (a) are valid when $\mu_{\tilde{\Pi}}(B) = \infty$ as well.

Proof. Let $C \subset \tilde{\Pi}$ be an arbitrary Borel set. For any $u \in \Pi$, let $[u]$ denote $\langle \cdot \rangle^{-1}(u)$. We begin by noting that $T(C \cap [u]) = T(C) \cap [S(u)]$ so that setting $A := C \cap [u]$ in (17) yields

$$m_{C \cap [u]}|T(B) \cap [S(u)]| \leq |C \cap [u]| \leq M_{C \cap [u]}|T(C) \cap [S(u)]|.$$

Observing that $m_C \leq m_{C \cap [u]} \leq M_{C \cap [u]} \leq M_C$ and employing definition (15), we obtain

$$m_C N_{T(C)}(S(u)) \leq N_C(u) \leq M_C N_{T(C)}(S(u)) \quad \text{for all } u \in \Pi. \tag{22}$$

We now integrate all three functions in (22) over Π , which gives us, using the S -invariance of μ_{Π} and the relation (14),

$$m_C \mu_{\tilde{\Pi}}(T(C)) \leq \mu_{\tilde{\Pi}}(C) \leq M_C \mu_{\tilde{\Pi}}(T(C)). \tag{23}$$

This is almost (21) except for the constants of equivalence. We will now tighten these constants. For any Borel $B \subset \tilde{\Pi}$, let B^* be as in (19). It already follows from (23) that $T(C)$ is measure zero if and only if C is. Setting $C := B \setminus B^*$, which is measure zero by (20), we get that $T(B \setminus B^*)$ is also measure zero so that $\mu_{\tilde{\Pi}}(T(B)) = \mu_{\tilde{\Pi}}(T(B^*))$. We now set $C := B^*$ in (23). Invoking (20) again, we obtain

$$\begin{aligned} \ell_B \mu_{\tilde{\Pi}}(T(B)) &= m_{B^*} \mu_{\tilde{\Pi}}(T(B^*)) \\ &\leq \mu_{\tilde{\Pi}}(B^*) = \mu_{\tilde{\Pi}}(B) \leq M_{B^*} \mu_{\tilde{\Pi}}(T(B^*)) = L_B \mu_{\tilde{\Pi}}(T(B)), \end{aligned}$$

which is the desired form of (21).

(c) \Rightarrow (a). This is immediate from (21) since T essentially one-to-one means $L_B = \ell_B = 1$.

(b) \Rightarrow (a). This follows by integrating $N_{T(B)}$ and $N_B \circ S^{-1}$ over Π and using (14) along with the fact that μ_{Π} is invariant under S^{-1} .

(a) \Rightarrow (b) and (c). Assume that $\mu_{\tilde{\Pi}}(T(B)) = \mu_{\tilde{\Pi}}(B) < \infty$. Let

$$B_1 := \{v \in B : |T^{-1}(T(v)) \cap B| = 1\}.$$

Since T is one-to-one on B_1 , we have $M_{B_1} = m_{B_1} = 1$ so that setting $C := B_1$ in (22) and replacing u by $S^{-1}(u)$ already gives $N_{T(B_1)} = N_{B_1} \circ S^{-1}$. We will extend this to B by showing that $B \setminus B_1$ is of measure zero. We clearly have $\mu_{\tilde{\Pi}}(T(B_1)) = \mu_{\tilde{\Pi}}(B_1)$. Meanwhile, for $B \setminus B_1$,

$$2\mu_{\tilde{\Pi}}(T(B \setminus B_1)) \leq \mu_{\tilde{\Pi}}(B \setminus B_1).$$

Indeed, if $B \setminus B_1 \neq \emptyset$, we have $m_{B \setminus B_1} \geq 2$ so that we can invoke (23) with $C := B \setminus B_1$; if $B \setminus B_1 = \emptyset$, the inequality, of course, still holds. This shows that

$$2\mu_{\tilde{\Pi}}(T(B \setminus B_1)) \leq \mu_{\tilde{\Pi}}(B) - \mu_{\tilde{\Pi}}(B_1) = \mu_{\tilde{\Pi}}(T(B)) - \mu_{\tilde{\Pi}}(T(B_1)) = \mu_{\tilde{\Pi}}(T(B \setminus B_1)),$$

where the last equality follows from the fact that $T(B \setminus B_1)$ and $T(B_1)$ are disjoint by definition of B_1 . Since $\mu_{\tilde{\Pi}}(T(B \setminus B_1)) < \infty$, we obtain $\mu_{\tilde{\Pi}}(T(B \setminus B_1)) = \mu_{\tilde{\Pi}}(B \setminus B_1) = 0$. Therefore $\mu_{\tilde{\Pi}}(T(B) \setminus T(B_1)) = 0$. Hence $L_B = 1$, i.e., T is essentially one-to-one. In addition,

$$N_{T(B)} = N_{T(B_1)} = N_{B_1} \circ S^{-1} = N_B \circ S^{-1} \quad \mu_{\tilde{\Pi}}\text{-a.e.},$$

where, in the last equality, we have also used the S -invariance of $\mu_{\tilde{\Pi}}$. □

3.3. *The tiling property of T -invariant sets.* In this subsection, we will prove our first main theorem which states that T -invariant sets comprise disjoint tiles. We start below with a discussion on our convention of T -invariance and its implications.

3.3.1. *T -invariant sets.* We say that a set A is T -invariant if $T(A) = A$. Similarly, for $\Pi \in \mathcal{C}_S$, we say that a Borel set $A \subset \tilde{\Pi}$ is *essentially T -invariant* if $\mu_{\tilde{\Pi}}(A \Delta T(A)) = 0$.

The significance of this definition is the following crucial observation. Suppose A is essentially T -invariant with $\mu_{\tilde{\Pi}}(A) < \infty$. It follows from Lemma 3.1((a) \Rightarrow (c)) that T is essentially one-to-one on A . Moreover, since T is also essentially one-to-one on any Borel subset of A , it again follows from Lemma 3.1(this time, (c) \Rightarrow (a)) that $(A, \mu_{\tilde{\Pi}}, T|_A)$ is an invertible measure-preserving system. Equivalently, 1_A is an invariant density for T on $\tilde{\Pi}$ since, for any Borel subset B of $\tilde{\Pi}$,

$$\mu_{\tilde{\Pi}}(A \cap T^{-1}B) = \mu_{\tilde{\Pi}}(T(A \cap T^{-1}B)) = \mu_{\tilde{\Pi}}(T(A) \cap B) = \mu_{\tilde{\Pi}}(A \cap B).$$

(Here, we have used the set identity $T(A \cap T^{-1}B) = T(A) \cap B$.)

Let us note that, in many examples, no set B of positive and finite measure satisfies the alternative invariance condition $T^{-1}(B) = B$, which, otherwise, would also be of interest.

It is natural to ask how one might encounter T -invariant sets. Theorem B.1 provides a general-purpose formula for the largest invariant set contained in any given set, which is valid for any ‘finite-to-one’ map. However, for some examples of T , there may not be any non-empty invariant sets in the strict sense (see Examples 7.1 and 7.2). Orbit closures are generally not expected to be strictly invariant either, but we will see in the next section that they are essentially invariant provided the partition is Jordan measurable.

Let us also note the following simple observation for future reference.

LEMMA 3.2. *Let $\Pi \in \mathcal{C}_S$ and A, B be two Borel sets in $\tilde{\Pi}$. If B is essentially T -invariant and A is $\mu_{\tilde{\Pi}}$ -equivalent to B , i.e., $\mu_{\tilde{\Pi}}(A \Delta B) = 0$, then A is also essentially T -invariant.*

Proof. Note that (21) implies $\mu_{\tilde{\Pi}}(T(A \Delta B)) = 0$, and since $T(A) \Delta T(B) \subset T(A \Delta B)$, we have $\mu_{\tilde{\Pi}}(T(A) \Delta T(B)) = 0$ as well. The conclusion then follows from the basic relation $A \Delta T(A) \subset (A \Delta B) \cup (B \Delta T(B)) \cup (T(B) \Delta T(A))$. □

3.3.2. *Tiles.* Note that a set $X \subset \mathbb{R}^d$ is \mathbb{Z}^d -invariant (i.e., $X + \mathbb{Z}^d = X$) if and only if it is the preimage under $\langle \cdot \rangle$ of a subset of \mathbb{T}^d . Given such a \mathbb{Z}^d -invariant set X , we say that Γ is an *exact tile* (or simply a *tile*) for X if $\{\Gamma + k\}_{k \in \mathbb{Z}^d}$ is a partition of X . More generally, given a non-negative integer m , we say that Γ is an *m -tile* for X if it is the union of m disjoint tiles for X , with 0-tile meaning the empty set. It is easy to see that Γ is an m -tile for X if and only if $N_{\Gamma} = m$ everywhere on $\langle X \rangle$.

We say that a Borel set $B \subset \tilde{\Pi}$ is an *essential m -tile* for $\tilde{\Pi}$ if it differs from an exact m -tile Γ for $\tilde{\Pi}$ by a $\mu_{\tilde{\Pi}}$ -null set. Equivalently, $N_B = m$ holds $\mu_{\tilde{\Pi}}$ -a.e. on $\tilde{\Pi}$.

THEOREM 3.1. *Let $\Pi \in \mathcal{C}_S^e$. Any essentially T -invariant set $A \subset \tilde{\Pi}$ of finite measure is an essential m -tile for $\tilde{\Pi}$, where it must necessarily be the case that $\mu_{\tilde{\Pi}}(A) = m$. If $m = 1$, then $T|_A$ is ergodic with respect to $\mu_{\tilde{\Pi}}$.*

Proof. The essential T -invariance of A yields $N_A = N_{T(A)}$ $\mu_{\tilde{\Pi}}$ -a.e. and Lemma 3.1((a) \Rightarrow (b)) yields $N_{T(A)} = N_A \circ S^{-1}$ $\mu_{\tilde{\Pi}}$ -a.e. Hence $N_A = N_A \circ S^{-1}$ $\mu_{\tilde{\Pi}}$ -a.e. Since S is ergodic, it follows that N_A is constant $\mu_{\tilde{\Pi}}$ -a.e. By definition of N_A , this constant is a non-negative integer m , and (14) implies that $\mu_{\tilde{\Pi}}(A) = m$.

The ergodicity of $T|_A$ for $m = 1$ follows from the fact that $\langle \cdot \rangle : A \rightarrow \Pi$ is a measure-preserving isomorphism between $(A, \mu_{\tilde{\Pi}})$ and (Π, μ_{Π}) intertwining T and S . Indeed, for any Borel subset B of A such that $T|_A^{-1}B = B$, (3) (combined with the fact that $T|_A$ is essentially one-to-one) implies that $\langle B \rangle$ is essentially S -invariant, so that ergodicity of S implies that $\mu_{\tilde{\Pi}}(B) = \mu_{\Pi}(\langle B \rangle) = 0$ or 1 . □

Remarks.

- $T|_A$ may or may not be ergodic when $m \geq 2$ (Examples 7.3 and 7.4).
- Any Borel exact m -tile Γ that is $\mu_{\tilde{\Pi}}$ -equivalent to A would be essentially T -invariant due to Lemma 3.2, and in the case $m = 1$, $T|_{\Gamma}$ would be ergodic with respect to $\mu_{\tilde{\Pi}}$. Note that neither A nor Γ may be exactly invariant, but working with an exact tile Γ in the case of $m = 1$ comes with the convenience of $\langle \cdot \rangle$ being a bijection between Γ and Π . This feature will be employed in §6 when we discuss convergence rate of ergodic averages.

4. Essential invariance and tiling of orbit closures for T

The results of the previous section were mainly of measure-theoretic nature, involving essentially T -invariant sets. The results of the present section and the next one will combine topology and measure by providing sufficient conditions for essential invariance, tiling and basic regularity properties of orbit closures for T in terms of the partition associated with T .

4.1. Orbit closures for piecewise homeomorphisms: topologically approximate invariance and essential invariance. The main result of this subsection will be a purely topological lemma on a sufficient condition for approximate invariance of orbit closures of piecewise homeomorphisms. Note that, for any map $f : X \rightarrow X$, any forward orbit $V = \mathcal{O}_f^+(v_0)$ is nearly invariant, satisfying

$$V = f(V) \cup \{v_0\}.$$

If f is a continuous map on a Hausdorff topological space X and V has compact closure, then this relation extends to \bar{V} because

$$\bar{V} = \overline{f(V) \cup \{v_0\}} \tag{24}$$

and $\overline{f(V)} = \overline{f(\bar{V})} = f(\bar{V})$, noting that, in Hausdorff spaces, compact sets are closed. However, this nice fact can easily fail for discontinuous maps when orbit closures intersect

with the set of discontinuities. To clarify the topological context of invariance properties of orbit closures for our piecewise automorphisms T , we will consider the larger class of piecewise homeomorphisms. (Some of our results can even be extended to piecewise continuous maps, but this generality will not be needed.)

Let $\{\Sigma_i\}_{i \in \Lambda}$ be a finite partition of a Hausdorff topological space X and let $f : X \rightarrow X$ be a piecewise map defined by

$$f(v) := f_i(v) \quad \text{for all } v \in \Sigma_i, \quad i \in \Lambda,$$

where each $f_i : X \rightarrow X$ is a homeomorphism. In short, we will call f a *piecewise homeomorphism*.

Given any collection of sets $\{B_i\}_{i \in \Lambda}$, we will write $f\{B\}$ for $\{f(B_i)\}_{i \in \Lambda}$ and ∂B for $\bigcup_{i \in \Lambda} \partial B_i$. Then $\partial f\{B\}$ stands for $\bigcup_{i \in \Lambda} \partial f(B_i)$. With this notation, we have

$$\partial f\{\Sigma\} = \bigcup_{i \in \Lambda} f_i(\partial \Sigma_i) \tag{25}$$

since $\partial f(\Sigma_i) = \partial f_i(\Sigma_i) = f_i(\partial \Sigma_i)$.

As in topological dynamics (of continuous maps), we will say that v_0 is (positively) *recurrent* (for (X, f)) if there exists a sequence of indices $n_k \rightarrow \infty$ (as $k \rightarrow \infty$) for which $f^{n_k}(v_0) \rightarrow v_0$. Therefore v_0 is recurrent if and only if $v_0 \in \overline{f(V)}$.

The following result shows that orbit closures of a piecewise homeomorphism f are approximately f -invariant provided the partition associated with f is sufficiently regular in the sense that $\partial f\{\Sigma\}$ is a small set.

LEMMA 4.1. *Any orbit $V = \mathcal{O}_f^+(v_0)$ of a piecewise homeomorphism $f : X \rightarrow X$ associated with a partition Σ satisfies $\overline{V} \Delta f(\overline{V}) \subset \{v_0\} \cup \partial f\{\Sigma\}$. If v_0 is recurrent, then, in fact, we have $\overline{V} \Delta f(\overline{V}) \subset \partial f\{\Sigma\}$.*

Proof. Note that

$$\overline{f(V)} = \bigcup_{i \in \Lambda} \overline{f_i(V \cap \Sigma_i)} = \bigcup_{i \in \Lambda} f_i(\overline{V \cap \Sigma_i})$$

and

$$f(\overline{V}) = \bigcup_{i \in \Lambda} f_i(\overline{V} \cap \Sigma_i)$$

so that

$$\overline{f(V)} \Delta f(\overline{V}) \subset \bigcup_{i \in \Lambda} f_i(\overline{V \cap \Sigma_i}) \Delta f_i(\overline{V} \cap \Sigma_i) \subset \bigcup_{i \in \Lambda} f_i([\overline{V \cap \Sigma_i}] \Delta [\overline{V} \cap \Sigma_i]). \tag{26}$$

In addition, we have the inclusions

$$\overline{V \cap \Sigma_i} \subset \overline{V} \cap \overline{\Sigma_i} \subset [\overline{V} \cap \Sigma_i] \cup \partial \Sigma_i$$

and

$$\overline{V} \cap \Sigma_i \subset [\overline{V \cap \Sigma_i} \cup \overline{\Sigma_i^c}] \cap \Sigma_i \subset \overline{V \cap \Sigma_i} \cup \partial \Sigma_i$$

so that $[\overline{V \cap \Sigma_i}] \Delta [\overline{V} \cap \Sigma_i] \subset \partial \Sigma_i$. The proof is complete once we inject this bound in (26) and employ (24). □

The following corollary is now a trivial measure-theoretic extension of Lemma 4.1.

COROLLARY 4.1. *Let f be a piecewise homeomorphism on X that is equipped with a continuous (atomless) Borel measure μ . If the partition Σ associated with f is such that $\mu(\partial f\{\Sigma\}) = 0$, i.e., $f_i(\Sigma_i)$ is a continuity set of μ for all $i \in \Lambda$, then every orbit closure for f is essentially f -invariant with respect to μ .*

4.2. Essential invariance and tiling of orbit closures for T . Recall the classes of sets defined in (7)–(9). For any $\Pi \in \mathcal{C}_S$, equip $\tilde{\Pi}$ with its relative Euclidean topology and the measure $\mu_{\tilde{\Pi}}$. Consider T as a piecewise homeomorphism defined on $\tilde{\Pi}$ and write $\Omega \cap \tilde{\Pi}$ for the partition of $\tilde{\Pi}$ consisting of the sets $\{\Omega_i \cap \tilde{\Pi}\}_{i \in \Lambda}$.

In analogy with the terminology in the Euclidean space, we will call any continuity set of $\mu_{\tilde{\Pi}}$ *$\tilde{\Pi}$ -Jordan measurable*. We will also say that the partition Ω is $\tilde{\Pi}$ -Jordan measurable if $\Omega_i \cap \tilde{\Pi}$ is $\tilde{\Pi}$ -Jordan measurable for every $i \in \Lambda$: in other words, if $\mu_{\tilde{\Pi}}(\partial(\Omega \cap \tilde{\Pi})) = 0$.

We first record the following result, which is merely an application of Corollary 4.1.

THEOREM 4.1. *Let $\Pi \in \mathcal{C}_S$ and let $V = \mathcal{O}_T^+(v_0)$ be any bounded orbit of T in $\tilde{\Pi}$.*

- (i) *If $\dim(\Pi) = 0$, then V is eventually periodic.*
- (ii) *If $\dim(\Pi) \geq 1$ and Ω is $\tilde{\Pi}$ -Jordan measurable, then \overline{V} is essentially T -invariant.*

Proof. The case $\dim(\Pi) = 0$ implies that Π is a finite set. Therefore every bounded subset of $\tilde{\Pi}$ is finite. Hence V is an eventually periodic orbit.

For the case $\dim(\Pi) \geq 1$, we employ Corollary 4.1 with $X = \tilde{\Pi}$, $f = T$, $\Sigma = \Omega \cap \tilde{\Pi}$ and $\mu = \mu_{\tilde{\Pi}}$ and note that each T_i preserves $\mu_{\tilde{\Pi}}$ so that each $T_i(\Omega_i \cap \tilde{\Pi})$ is also a continuity set of $\mu_{\tilde{\Pi}}$ and $\mu_{\tilde{\Pi}}(\{v_0\}) = 0$. □

The next result concerns some basic observations on the multiplicity function of bounded orbits and their closures.

LEMMA 4.2. *Let $\Pi \in \mathcal{C}_S$ and let $V = \mathcal{O}_T^+(v_0)$ be any bounded orbit of T in $\tilde{\Pi}$.*

- (i) *If $\langle V \rangle$ is not a periodic orbit of S , then $N_V \leq 1$.*
- (ii) *If $\langle V \rangle$ is dense in Π , then $N_{\overline{V}} \geq 1$ on Π .*
- (iii) *If $\dim(\Pi) \geq 1$ and $\langle V \rangle$ is dense in Π , then $N_V \leq 1 \leq N_{\overline{V}}$ on Π .*

Proof. For (i), let $v_n := T^n(v_0)$ and note that if $N_V(u) \geq 2$ for some $u \in \Pi$, then there exist $n_2 > n_1$ such that $\langle v_{n_2} \rangle = \langle v_{n_1} \rangle = u$. But then $\langle V \rangle$ would be a periodic orbit of S since $\langle v_{n_2} \rangle = S^{n_2 - n_1} \langle v_{n_1} \rangle$.

For (ii), note that \overline{V} is compact and therefore $\langle \overline{V} \rangle = \overline{\langle V \rangle} = \Pi$.

For (iii), note that density implies non-periodicity when $\dim(\Pi) \geq 1$ so that we can combine (i) and (ii). □

Note that, when V is bounded, $\langle V \rangle$ is a periodic orbit of S if and only if V is an eventually periodic orbit of T . (The ‘if’ part is obvious. The ‘only if’ part follows from the observation that V has to be a finite set since it is contained in a bounded set whose toral projection is finite.) Therefore (i) could be restated as ‘if V is not eventually periodic, then $N_V \leq 1$ ’, but the stated form of (i) is more robust because its conclusion holds even when V is not bounded. Meanwhile, (ii) and (iii) could fail without the boundedness

assumption on V . For example, any orbit V of the affine map $T(v) := v + a$ on \mathbb{R} , where a is irrational, yields projection $\langle V \rangle$ which is a dense subset of \mathbb{T} ; but V is closed and $\langle V \rangle$ is a countable subset of \mathbb{T} , so $N_{\overline{V}}(u) = N_V(u) = 0$ for uncountably many $u \in \mathbb{T}$.

Our final result in this section shows that ergodicity of S implies tiling of orbit closures for T .

THEOREM 4.2. *Let $\Pi \in \mathcal{C}_S^e$ and let $V = \mathcal{O}_T^+(v_0)$ be any bounded orbit of T in $\tilde{\Pi}$.*

- (i) *If $\dim(\Pi) = 0$, then there exist $l \geq 0$ and $m \geq 1$ such that $\mathcal{O}_T^+(T^l(v_0))$ is an exact m -tile.*
- (ii) *If $\dim(\Pi) \geq 1$ and Ω is $\tilde{\Pi}$ -Jordan measurable, then \overline{V} is an essential m -tile for some $m \geq 0$. If, in addition, $\langle V \rangle$ is dense in Π , then $m \geq 1$.*

Proof. For (i), we know that V is eventually periodic by Theorem 4.1(i). Let $l \geq 0$ be such that $V' = \mathcal{O}_T^+(T^l(v_0))$ is periodic. Then V' is T -invariant and of finite measure, so it is an essential m -tile by Theorem 3.1. Clearly, the only measure zero set in $\tilde{\Pi}$ is the empty set, so V' is an exact m -tile. Furthermore, $m \geq 1$ because V' is non-empty.

For (ii), Theorem 4.1(ii) shows that \overline{V} is essentially T -invariant. Clearly, $\mu_{\tilde{\Pi}}(\overline{V})$ is finite since \overline{V} is compact. Hence Theorem 3.1 implies that \overline{V} is an essential m -tile for some $m \geq 0$. The final claim follows immediately from Lemma 4.2(ii). □

Remarks.

- It is not difficult to construct examples of T that yield orbit closures with tiling multiplicity $m > 1$. Trivial examples follow by scaling up any single tile example by an integer; for a non-trivial example, see Example 7.3.
- We know by Theorem 3.1 that $(T, \overline{V}, \tilde{\mu}_\Pi)$ is ergodic when \overline{V} is a single tile. It is natural to ask whether this is always the case for orbit closures when Ω is $\tilde{\Pi}$ -Jordan measurable. However, this is not true, even if S has additional favorable properties, such as unique ergodicity or topological transitivity (see Example 7.5).
- Note that the boundary $\partial(\Omega \cap \tilde{\Pi})$ that appears in Theorem 4.2 is with respect to the topology of $\tilde{\Pi}$. It can be checked that $\partial(\Omega \cap \tilde{\Pi})$ is contained in $\partial\Omega \cap \tilde{\Pi}$, where $\partial\Omega$ now stands for the boundary of Ω in \mathbb{R}^d . Hence, we may, alternatively, check the stronger condition $\mu_{\tilde{\Pi}}(\partial\Omega \cap \tilde{\Pi}) = 0$ if it is more convenient to do so.

It is safe to claim that all practical examples of maps T satisfy this basic regularity assumption. As a particular case, we note that if P is a polyhedron in \mathbb{R}^d (bounded or not), then, for any affine subspace V of \mathbb{R}^d , the cross-section $P \cap V$ is also a polyhedron (in V). Therefore, if the partition Ω consists entirely of polyhedral sets, then it is automatically $\tilde{\Pi}$ -Jordan measurable for any $\Pi \in \mathcal{C}$.

5. Regularity of orbit closures and invariant sets for T

In this section, we will show that compact invariant sets inherit the basic regularity of the partition associated with T . In particular, we will show that, given $\Pi \in \mathcal{C}_S^e$, if the partition is $\tilde{\Pi}$ -Jordan measurable, then the closure of every bounded orbit and, in general, any compact essentially invariant set in Π is also $\tilde{\Pi}$ -Jordan measurable.

As in §4, we will first consider general piecewise homeomorphisms and then discuss the implications for the piecewise-affine automorphisms T .

5.1. *Orbit closures for piecewise homeomorphisms: the boundary.* Our first result shows that the boundary of an orbit closure for a piecewise homeomorphism is controlled by its image and the boundary of the underlying partition.

LEMMA 5.1. *Any orbit $V = \mathcal{O}_f^+(v_0)$ of a piecewise homeomorphism $f : X \rightarrow X$ associated with a partition Σ satisfies $\partial\bar{V} \setminus f(\partial\bar{V}) \subset \{v_0\} \cup \partial f\{\Sigma\}$. If v_0 is recurrent, then, in fact, we have $\partial\bar{V} \setminus f(\partial\bar{V}) \subset \partial f\{\Sigma\}$.*

Proof. We start by noting that, for any function f and sets A, B, C , where $A = B \setminus C$, we have the inclusion $A \setminus f(A) \subset B \setminus f(B) \cup f(C) \setminus C$. Since $\partial\bar{V} = \bar{V} \setminus \overset{\circ}{V}$,

$$\partial\bar{V} \setminus f(\partial\bar{V}) \subset [\bar{V} \setminus f(\bar{V})] \cup [f(\overset{\circ}{V}) \setminus \overset{\circ}{V}].$$

Lemma 4.1 readily implies that $[\bar{V} \setminus f(\bar{V})] \subset \{v_0\} \cup \partial f\{\Sigma\}$ (with $\{v_0\}$ being unnecessary if v_0 is recurrent, i.e., contained in $\bar{f}(\bar{V})$), so it suffices to consider $[f(\overset{\circ}{V}) \setminus \overset{\circ}{V}]$. Note that

$$f(\overset{\circ}{V}) = \bigcup_{i \in \Lambda} f_i(\overset{\circ}{V} \cap \Sigma_i) \subset \bigcup_{i \in \Lambda} f_i(\overset{\circ}{V} \cap \overset{\circ}{\Sigma}_i) \cup f_i(\partial \Sigma_i) \subset \bigcup_{i \in \Lambda} f_i(\overset{\circ}{V} \cap \overset{\circ}{\Sigma}_i) \cup \partial f\{\Sigma\}.$$

For any given $i \in \Lambda$,

$$\overset{\circ}{V} \cap \overset{\circ}{\Sigma}_i \subset \bar{V} \cap \overset{\circ}{\Sigma}_i \subset (\overline{V \cap \Sigma_i} \cup \overline{\Sigma_i^c}) \cap \overset{\circ}{\Sigma}_i \subset \overline{V \cap \Sigma_i},$$

and therefore $\overset{\circ}{V} \cap \overset{\circ}{\Sigma}_i \subset \overline{V \cap \Sigma_i}$. Since f_i is a homeomorphism,

$$f_i(\overset{\circ}{V} \cap \overset{\circ}{\Sigma}_i) \subset f_i(\overline{V \cap \Sigma_i}) = \overline{f_i(V \cap \Sigma_i)} = \overline{f(V \cap \Sigma_i)} \subset \overline{f(\overset{\circ}{V})} \subset \overset{\circ}{V}.$$

Taking the union over $i \in \Lambda$, we get $f(\overset{\circ}{V}) \subset \overset{\circ}{V} \cup \partial f\{\Sigma\}$ and therefore $f(\overset{\circ}{V}) \setminus \overset{\circ}{V} \subset \partial f\{\Sigma\}$, which completes the proof. □

COROLLARY 5.1. *Let $V = \mathcal{O}_f^+(v_0)$ be any orbit of a piecewise homeomorphism $f : X \rightarrow X$ associated with a partition Σ . Then*

$$\partial\bar{V} \subset \bigcap_{n=1}^{\infty} f^n(\partial\bar{V}) \cup \bigcup_{k=0}^{\infty} f^k(\partial f\{\Sigma\}) \cup V^*, \tag{27}$$

where V^* is defined to be \emptyset if v_0 is recurrent and V otherwise.

Proof. Lemma 5.1 clearly implies that

$$\partial\bar{V} \subset f(\partial\bar{V}) \cup \partial f\{\Sigma\} \cup V^*. \tag{28}$$

Noting that $f(V^*) \subset V^*$ and iteratively applying f to (28), we get that, for any $n \geq 1$,

$$\partial\bar{V} \subset f^n(\partial\bar{V}) \cup \bigcup_{k=0}^{n-1} f^k(\partial f\{\Sigma\}) \cup V^* \subset f^n(\partial\bar{V}) \cup \bigcup_{k=0}^{\infty} f^k(\partial f\{\Sigma\}) \cup V^*.$$

Intersecting these supersets over all $n \geq 1$ completes the proof. □

We will analyze this bound on $\partial\bar{V}$ in the next subsection in the special case of the piecewise automorphism T on $\tilde{\Pi}$.

5.2. *Regularity of orbit closures for T.* The main result of this section is the following theorem.

THEOREM 5.1. *Let $\Pi \in \mathcal{C}_S^e$ and let $V = \mathcal{O}_T^+(v_0)$ be any bounded orbit of T in $\tilde{\Pi}$.*

- (i) *If Ω is $\tilde{\Pi}$ -Jordan measurable, then so is \bar{V} .*
- (ii) *If, in addition, S is uniquely ergodic on Π and v_0 is a recurrent point for $(\tilde{\Pi}, T)$, then*

$$\partial \bar{V} \subset \bigcup_{k=0}^{\infty} T^k(\partial T\{\Omega \cap \tilde{\Pi}\}).$$

To prove this theorem, we will need the following lemma.

LEMMA 5.2. *Let $\Pi \in \mathcal{C}_S^e$ and let $D \subset \tilde{\Pi}$ be any Borel set. If $\langle D \rangle$ is not dense in Π , then $\mu_{\tilde{\Pi}}(\bigcap_{n=1}^{\infty} T^n(D)) = 0$. If, in addition, S is uniquely ergodic on Π , then $\bigcap_{n=1}^{\infty} T^n(D) = \emptyset$.*

Proof. Let $D_T := \bigcap_{n=1}^{\infty} T^n(D)$. Note that

$$\langle D_T \rangle \subset \bigcap_{n=1}^{\infty} \langle T^n(D) \rangle = \bigcap_{n=1}^{\infty} S^n(\langle D \rangle) =: \langle D \rangle_S,$$

so it suffices to show that $\mu_{\Pi}(\langle D \rangle_S) = 0$ and, in the uniquely ergodic case, that $\langle D \rangle_S = \emptyset$.

Due to the bijectivity of S, we have $S^{-1}(\langle D \rangle_S) = \langle D \rangle \cap \langle D \rangle_S$, which shows that $S^{-1}(\langle D \rangle_S)$ and therefore also its homeomorphic image $\langle D \rangle_S$ are not dense in Π .

It is also true that $S^{-1}(\langle D \rangle_S) \subset \langle D \rangle_S$. Therefore, for any $u \in \langle D \rangle_S$, the backward orbit $U := \mathcal{O}_S^-(u)$ which remains in $\langle D \rangle_S$ is not dense in Π . Since μ_{Π} -a.e. orbit of S is dense in Π due to ergodicity of S, we conclude that $\mu_{\Pi}(\langle D \rangle_S) = 0$. If, in addition, S (and therefore S^{-1}) is uniquely ergodic on Π , then every orbit U of S^{-1} would have to be dense, which immediately implies that $\langle D \rangle_S = \emptyset$. □

Proof of Theorem 5.1. We start by inspecting the bound (27) of Corollary 5.1. For (i), we lose no generality by assuming that $\dim(\Pi) \geq 1$ because, when $\dim(\Pi) = 0$, $\tilde{\Pi}$ is equipped with the discrete topology so that every set has empty boundary.

Note that $\partial \bar{V}$ is a compact nowhere dense set in $\tilde{\Pi}$. With Lemma A.3, we know that $\langle \partial \bar{V} \rangle$ is compact nowhere dense, and in particular, not dense, in Π . Hence, the first implication in Lemma 5.2 shows that $\bigcap_{n=1}^{\infty} T^n(\partial \bar{V})$ is $\mu_{\tilde{\Pi}}$ -null. For (i), as in the proof of Theorem 4.1, we know that the assumption that $\mu_{\tilde{\Pi}}(\partial(\Omega \cap \tilde{\Pi})) = 0$ and Lemma 3.1 imply that $\mu_{\tilde{\Pi}}(\partial T\{\Omega \cap \tilde{\Pi}\}) = 0$. Repeatedly applying Lemma 3.1 shows that the T-iterates of this set, $T^k(\partial T\{\Omega \cap \tilde{\Pi}\})$, are also $\mu_{\tilde{\Pi}}$ -null for all $k \geq 0$. Taking the union over k, invoking Corollary 5.1 for T on $\tilde{\Pi}$ and noting that with $\dim(\Pi) \geq 1$ we automatically have $\mu_{\tilde{\Pi}}(V^*) = 0$, we get that $\mu_{\tilde{\Pi}}(\partial \bar{V}) = 0$.

For (ii), it suffices to invoke Corollary 5.1 again, but this time employing the second implication in Lemma 5.2. □

Remark. One can also give an alternative proof of Theorem 5.1(i) using Theorem 3.1 as follows. Lemma 5.1 implies that $\mu_{\tilde{\Pi}}(\partial \bar{V} \setminus T(\partial \bar{V})) = 0$ so that

$$0 \leq \mu_{\tilde{\Pi}}(T(\partial \bar{V}) \setminus \partial \bar{V}) = \mu_{\tilde{\Pi}}(T(\partial \bar{V})) - \mu_{\tilde{\Pi}}(\partial \bar{V}) + \mu_{\tilde{\Pi}}(\partial \bar{V} \setminus T(\partial \bar{V})) \leq 0,$$

and therefore $\mu_{\tilde{\Pi}}(\partial \bar{V} \triangle T(\partial \bar{V})) = 0$. Now Theorem 3.1 shows that $\partial \bar{V}$ is an essential m-tile for some non-negative integer m. If we had $m \geq 1$, Corollary A.1 would imply

that $N_{\partial\bar{V}}(u) \geq 1$ everywhere on Π , yielding $\langle \partial\bar{V} \rangle = \Pi$. However, $\langle \partial\bar{V} \rangle$ is nowhere dense according to Lemma A.3. Hence $m = 0$.

5.3. *Regularity of invariant sets for T.* The main result of this section is the following extension of Theorem 5.1.

THEOREM 5.2. *Let $\Pi \in \mathcal{C}_S^e$, Ω be $\tilde{\Pi}$ -Jordan measurable, and let K be a compact essentially T -invariant subset of $\tilde{\Pi}$.*

- (i) *K is $\tilde{\Pi}$ -Jordan measurable.*
- (ii) *There is a $\tilde{\Pi}$ -Jordan measurable exact m -tile Γ such that $\overset{\circ}{K} \subset \Gamma \subset K$ and $\mu_{\tilde{\Pi}}(K \setminus \Gamma) = 0$.*

We will prove this theorem with the help of the lemma below which will provide a measure-theoretic reduction convenient for our objective. Let us say that a Borel set $A \subset \tilde{\Pi}$ is *trim* if, for every $v \in A$ and every open neighborhood U of v , we have $\mu_{\tilde{\Pi}}(A \cap U) > 0$. Equivalently, the support of $\mu_{\tilde{\Pi}}|_A$ (restriction of $\mu_{\tilde{\Pi}}$ to A) is equal to A . The empty set is automatically trim, but all other trim sets have non-zero measure.

LEMMA 5.3. *Let $\Pi \in \mathcal{C}$. Any compact set K in $\tilde{\Pi}$ can be expressed as the disjoint union of two sets, A and B , where A is either empty (if $\mu_{\tilde{\Pi}}(K) = 0$) or else compact and trim, B is of measure zero and $\mu_{\tilde{\Pi}}(\partial B) \leq \mu_{\tilde{\Pi}}(\partial A)$.*

Proof. Let E be the set of $v \in K$ for which $K \cap U$ is of measure zero for some open neighborhood $U := U_v$ of v . Let $G := \bigcup_{v \in E} U_v$, $A := K \setminus G$ and $B := K \cap G$. A is compact since G is open. Furthermore, $G \supset E$ so that A and E are disjoint, and therefore A is trim (including the case $A = \emptyset$). By Lindelöf’s lemma, G can be reduced to the union of a countable subfamily $(U_{v_n})_{n \geq 1}$, which implies that $B = \bigcup_{n \geq 1} K \cap U_{v_n}$ is of measure zero. Finally, $\bar{B} \subset K \setminus \overset{\circ}{A} = (\partial A) \cup B$ so that $\mu_{\tilde{\Pi}}(\partial B) \leq \mu_{\tilde{\Pi}}(\partial A)$. □

Proof of Theorem 5.2. Let K be a compact essentially T -invariant set in $\tilde{\Pi}$. If K is of measure zero, then it is automatically $\tilde{\Pi}$ -Jordan measurable (because it is closed), so we may assume that K is of non-zero measure.

Let A be the compact trim subset of K and let B be its residual as prescribed in Lemma 5.3. B is of measure zero, so A is essentially T -invariant as well. It now suffices to show that A is $\tilde{\Pi}$ -Jordan measurable since $\partial K \subset \partial A \cup \partial B$ and $\mu_{\tilde{\Pi}}(\partial B) \leq \mu_{\tilde{\Pi}}(\partial A)$.

We may assume that $\dim(\Pi) \geq 1$ because if $\dim(\Pi) = 0$, then the topology of $\tilde{\Pi}$ is discrete so that all subsets of $\tilde{\Pi}$ are automatically $\tilde{\Pi}$ -Jordan measurable.

By Theorem 3.1, we know that A is an essential m -tile. We also know that $m \geq 1$ because the trimness of A implies that $\mu_{\tilde{\Pi}}(A) > 0$. Since A is essentially T -invariant, the set of points in A whose forward orbits remain in A has measure m . Take any such orbit V_1 in A such that $\langle V_1 \rangle$ is dense in Π . By Theorems 4.1(ii) and 4.2(ii), we know that the compact set $A_1 := \bar{V}_1$ is an essentially T -invariant essential k_1 -tile for some $1 \leq k_1 \leq m$, is contained in A and, by Theorem 5.1(i), is $\tilde{\Pi}$ -Jordan measurable. If $k_1 < m$, that is, if $\mu_{\tilde{\Pi}}(A \setminus A_1) > 0$, we repeat the process (since $A \setminus A_1$ is also essentially T -invariant) and extract a compact $\tilde{\Pi}$ -Jordan measurable essentially T -invariant essential k_2 -tile $A_2 = \bar{V}_2$

for some orbit $V_2 \subset A \setminus A_1$ such that $\langle V_2 \rangle$ is dense in Π so that $k_2 \geq 1$. Since A_1 is $\tilde{\Pi}$ -Jordan measurable and $A_1 \cap A_2 \subset \partial A_1$, we get that $\mu_{\tilde{\Pi}}(A_1 \cup A_2) = k_1 + k_2$. This process will terminate after finitely many steps, resulting in compact $\tilde{\Pi}$ -Jordan measurable sets A_1, \dots, A_l such that $\mu_{\tilde{\Pi}}(A \setminus (A_1 \cup \dots \cup A_l)) = 0$. We claim that $A \setminus (A_1 \cup \dots \cup A_l)$ is empty because if it contained a point v , then $U := (A_1 \cup \dots \cup A_l)^c$ would be an open neighborhood of v and the trimness of A would contradict the fact that $\mu_{\tilde{\Pi}}(A \cap U) = 0$. Hence we have $A = A_1 \cup \dots \cup A_l$, which shows that A (and therefore K) is $\tilde{\Pi}$ -Jordan measurable and proves the claim in (i).

For (ii), first note that, by Corollary A.1, we have $N_K^\circ \leq m \leq N_K$ on Π . Next we apply Lemma A.2 to the pair $(\overset{\circ}{K}, K)$ to obtain a Borel measurable exact m -tile Γ such that $\overset{\circ}{K} \subset \Gamma \subset K$, which yields $\partial \Gamma \subset \partial K$ so that Γ , too, is $\tilde{\Pi}$ -Jordan measurable. \square

5.4. *Embedding orbits of T in regular exact tiles.* The main result of this section is the following theorem.

THEOREM 5.3. *Let $\Pi \in \mathcal{C}_S^e$ and let Ω be $\tilde{\Pi}$ -Jordan measurable. Suppose V is a bounded orbit of T in $\tilde{\Pi}$ and $\langle V \rangle$ is dense in Π . Let $m \geq 1$ denote the multiplicity of tiling for \bar{V} . Then there exists a $\tilde{\Pi}$ -Jordan measurable exact m -tile Γ such that $V \subset \Gamma \subset \bar{V}$.*

Note that, in the setting of Theorem 5.3, the conclusion of Theorem 5.1 which says that $\mu_{\tilde{\Pi}}(\partial \bar{V}) = 0$ is in effect. Furthermore, Theorem 5.2 shows that a $\tilde{\Pi}$ -Jordan measurable exact m -tile Γ could be found so that $\overset{\circ}{V} \subset \Gamma \subset \bar{V}$. Since V need not be contained in $\overset{\circ}{V}$, we may wonder whether Γ could be arranged to contain $\overset{\circ}{V} \cup V$ as well. Unfortunately, it is possible that $\sup N_{\overset{\circ}{V} \cup V}^\circ > m$ for some orbits so that no exact m -tile can contain $\overset{\circ}{V} \cup V$. (For example, let $a \in (0, 1)$ be irrational and let $T : \mathbb{R} \rightarrow \mathbb{R}$, $T(v) := v + a - 1_{[1-a, \infty)}(v)$. Then, for any $v_0 \in (1, 2)$, it can be checked that $\overset{\circ}{V} = (0, 1)$ (so that $m = 1$) yet $N_{\overset{\circ}{V} \cup V}^\circ(\langle v_0 \rangle) = 2$.) To circumvent this obstacle, we will implement a careful ‘surgery’ on $\overset{\circ}{V}$ to remove a $\mu_{\tilde{\Pi}}$ -null subset that overlaps with $V + \mathbb{Z}^d \setminus \{0\}$ while maintaining $\tilde{\Pi}$ -Jordan measurability of the resulting set. Our next result, which is essentially topological (not concerning measure theory or dynamics), provides the main tool of this surgery.

LEMMA 5.4. *Let $\Pi \in \mathcal{C}$. For any bounded set $V \subset \tilde{\Pi}$, there exists W such that $V \subset W \subset \bar{V}$, $N_W \leq \max(N_V, N_{\overset{\circ}{V}})$ and $\bar{V} \setminus \overset{\circ}{W} \subset \partial \bar{V} + \mathbb{Z}^d$.*

Proof. Let Z_V denote $\partial \bar{V} + \mathbb{Z}^d$ and set $W := V \cup (\overset{\circ}{V} \setminus Z_V)$. It is clear that $V \subset W \subset \bar{V}$. Next, note that $V \setminus \overset{\circ}{V} \subset \partial \bar{V} \subset Z_V$ so that $V \setminus (\overset{\circ}{V} \setminus Z_V) \subset Z_V$, and therefore

$$W = [V \setminus (\overset{\circ}{V} \setminus Z_V)] \cup (\overset{\circ}{V} \setminus Z_V) \subset Z_V \cup (\overset{\circ}{V} \setminus Z_V).$$

Since Z_V is \mathbb{Z}^d -invariant, $\langle Z_V \rangle$ and $\langle \overset{\circ}{V} \setminus Z_V \rangle$ are disjoint. The same is then true for their respective subsets $\langle V \setminus (\overset{\circ}{V} \setminus Z_V) \rangle$ and $\langle \overset{\circ}{V} \setminus Z_V \rangle$, which together constitute $\langle W \rangle$. Therefore

$$N_W = \max(N_{V \setminus (\overset{\circ}{V} \setminus Z_V)}, N_{\overset{\circ}{V} \setminus Z_V}) \leq \max(N_V, N_{\overset{\circ}{V}}).$$

Note that Z_V is closed since it is the sum of a compact set and a closed set. Therefore $\overset{\circ}{V} \setminus Z_V$ is open, which implies that $\overset{\circ}{W} \supset \overset{\circ}{V} \setminus Z_V$, and so $\overline{V} \setminus \overset{\circ}{W} \subset \overline{V} \setminus (\overset{\circ}{V} \setminus Z_V) \subset Z_V$. \square

Remark. In Lemma 5.4, since both $\overline{V} \setminus \overset{\circ}{W}$ and $\partial \overline{V}$ are bounded sets, we can, in fact, say that $\overline{V} \setminus \overset{\circ}{W} \subset \partial \overline{V} + K$ for some finite set $K \subset \mathbb{Z}^d$.

Proof of Theorem 5.3. First of all, note that the case $\dim(\Pi) = 0$ is readily handled by Theorem 4.2(i), so we assume that $\dim(\Pi) \geq 1$. Let $m \geq 1$ be the multiplicity of (essential) tiling for \overline{V} as implied by Theorem 4.2(ii). With Lemma 4.2(iii), we have $N_V \leq 1$, and with Corollary A.1 applied to $A = \overline{V}$, we have $N_{\overline{V}} \leq m \leq N_{\overline{V}}$ on Π . Lemma 5.4 yields W such that $V \subset W \subset \overline{V}$ and $N_W \leq m$, and Lemma A.2 with $A = W$ and $B = \overline{V}$ yields a Borel exact m -tile Γ such that $V \subset W \subset \Gamma \subset \overline{V}$. With Lemma 5.4 again, we have $\partial \Gamma \subset \overline{V} \setminus \overset{\circ}{W} \subset \partial \overline{V} + \mathbb{Z}^d$ and Theorem 5.1 now yields $\mu_{\overline{\Pi}}(\partial \Gamma) = 0$. \square

6. Convergence rate of ergodic averages

For convenience, the setting of this section will be limited to the case when S is ergodic on $\Pi := \mathbb{T}^d$. We will assume that the partition Ω is Jordan measurable. We start with a bounded, Jordan measurable, exact 1-tile Γ on which T is ergodic. This set may have been obtained using Theorem 5.2 or Theorem 5.3.

Almost every orbit $\mathcal{O}_T^+(v_0)$ originating in Γ will remain in Γ . As we stated earlier in (4), the ergodic theorem implies that, for every $f \in L^1(\Gamma)$, $D_N(f, \mathbf{v}) \rightarrow 0$ for almost every such $v_0 \in \Gamma$. In this section, we will quantify the rate of convergence for functions f that are uniformly continuous on Γ .

The above qualitative result can actually be slightly strengthened via the toolkit of uniform distribution in compact spaces (e.g., as in [KN74]). Let us consider T on the compact set $\overline{\Gamma}$. Since $\partial \Gamma$ is null, $\overline{\Gamma}$ is a Jordan measurable essentially T -invariant essential 1-tile on which T is ergodic. It then follows that, for almost every $v_0 \in \overline{\Gamma}$ (and, consequently, for almost every $v_0 \in \Gamma$), the sequence $\mathbf{v} = (v_n)_0^\infty$, where $v_n := T^n(v_0)$, is uniformly distributed in $\overline{\Gamma}$, which means that, for any such v_0 , $D_N(f, \mathbf{v}) \rightarrow 0$ for every continuous $f : \overline{\Gamma} \rightarrow \mathbb{C}$ or, equivalently, for every uniformly continuous f on Γ . This result has the advantage that the ‘good’ orbits, i.e., those that are uniformly distributed in $\overline{\Gamma}$, result in the convergence of ergodic averages for a whole class of functions at once.

For additional improvements, we consider the toral projection. Given any orbit $\mathcal{O}_T^+(v_0)$ that remains in Γ , let $\mathbf{u} = (u_n)_0^\infty$ in \mathbb{T}^d be defined by $u_n := \langle v_n \rangle$, where $v_n := T^n(v_0)$, so that $u_n = S^n(u_0)$. Denoting the inverse of the restriction of $\langle \cdot \rangle$ to Γ by $\langle \cdot \rangle_\Gamma : \mathbb{T}^d \rightarrow \Gamma$, we have $v_n = \langle u_n \rangle_\Gamma$. Let $g := f \circ \langle \cdot \rangle_\Gamma$, where $f : \Gamma \rightarrow \mathbb{C}$. It then follows by substitution that

$$D_N(f, \mathbf{v}) = D_N(g, \mathbf{u}) := D_N(g, \mathbf{u}, \mathbb{T}^d) := \left| \frac{1}{N} \sum_{n=0}^{N-1} g(u_n) - \int_{\mathbb{T}^d} g(u) du \right|. \tag{29}$$

By the same reasoning as above but this time applied to S on \mathbb{T}^d , ergodicity of S implies that \mathbf{u} is uniformly distributed in \mathbb{T}^d for almost every u_0 : that is, the discrepancy

$$D_N(\mathbf{u}) := \sup_{R \in \mathcal{R}} D_N(1_R, \mathbf{u}, \mathbb{T}^d) \rightarrow 0 \quad \text{as } N \rightarrow \infty,$$

where \mathcal{R} stands for the collection of axis-parallel rectangles in \mathbb{T}^d .

For a uniformly distributed sequence \mathbf{u} , note that $D_N(h, \mathbf{u}, \mathbb{T}^d) \rightarrow 0$ for every Riemann integrable h on \mathbb{T}^d (when viewed as a function on $[0, 1)^d$). This is an important distinction because $g := f \circ \langle \cdot \rangle_\Gamma$ may be (and typically would be) discontinuous whereas the Jordan measurability of Γ implies that $f \circ \langle \cdot \rangle_\Gamma$ is Riemann integrable on \mathbb{T}^d as long as f is bounded and continuous on Γ . (This follows from Lebesgue’s criterion for Riemann integrability, noting that the set of discontinuities of $f \circ \langle \cdot \rangle_\Gamma$ is contained in the set of discontinuities of $\langle \cdot \rangle_\Gamma$, which is a subset of $\langle Z \cup \partial\Gamma \rangle$, where Z stands for the ‘continuous integer grid’, i.e., the measure-zero set of points in \mathbb{R}^d with at least one integer coordinate.) Hence it follows that almost every $v_0 \in \Gamma$ yields $D_N(f, \mathbf{v}) \rightarrow 0$ for every bounded and continuous f on Γ (instead of every uniformly continuous f). While this is yet another improvement of the qualitative convergence result, the reformulation of the problem on \mathbb{T}^d actually paves the path for a quantitative bound, too. For this, we will apply some well-known bounds on (29), which we discuss next.

For a general Riemann integrable h on \mathbb{T}^d , $D_N(h, \mathbf{u})$ can be controlled by a combination of *a priori* knowledge on the regularity of h and the discrepancy $D_N(\mathbf{u})$. The advantage of this route is that quantitative bounds for $D_N(\mathbf{u})$ can be given via analytic tools (e.g., the exponential sums that appear in the Erdős–Turán inequality [KN74]) that can exploit the algebraic structure of S . For example, for unipotent L , one ends up with Weyl sums which are well studied. The most commonly used bound for $D_N(h, \mathbf{u})$ is given by the standard version of the Koksma–Hlawka inequality [KN74], which provides the bound

$$D_N(h, \mathbf{u}) \lesssim_d \text{Var}_{HK}(h) D_N(\mathbf{u}),$$

where $\text{Var}_{HK}(h)$ denotes the variation of h in the sense of Hardy and Krause. Unfortunately, this is a highly restricted definition of bounded variation and $f \circ \langle \cdot \rangle_\Gamma$ can easily fail to be of bounded variation in this sense if $\partial\Gamma$ is not aligned with coordinate axes. An alternative tool (also due to Hlawka) that is applicable in our case enlarges the class of h that can be handled to the class of all Riemann integrable functions at the expense of a weaker bound. Given a Riemann integrable function h on $[0, 1)^d$ and a grid partition \mathcal{P} of $[0, 1)^d$ into half-open intervals, let $s(h, \mathcal{P})$ be the mean oscillation of h over \mathcal{P} , that is, the difference between the upper and lower Darboux sums of h corresponding to \mathcal{P} . For any $t > 0$, let

$$S(h, t) := \sup_{\|\mathcal{P}\| \leq t} s(h, \mathcal{P}),$$

where $\|\mathcal{P}\|$ is defined to be the maximum edge-length for the intervals in \mathcal{P} . Then Hlawka [Hla71] provides the bound

$$D_N(h, \mathbf{u}) \lesssim_d S(h, \lfloor D_N(\mathbf{u})^{-1} \rfloor^{-1/d}). \tag{30}$$

Equipped with Hlawka’s bound, we now proceed to state and prove our main quantitative improvement of (4).

THEOREM 6.1. *Let $\mathbf{v} = (v_n)_0^\infty$ be any sequence in the Jordan measurable exact 1-tile Γ and let $f : \Gamma \rightarrow \mathbb{C}$ be uniformly continuous. Then*

$$D_N(f, \mathbf{v}) \lesssim_d \omega_f(\sqrt{d} \lfloor D_N(\mathbf{u})^{-1} \rfloor^{-1/d}) + \|f\|_{\infty} \rho_\Gamma(\sqrt{d} \lfloor D_N(\mathbf{u})^{-1} \rfloor^{-1/d}), \tag{31}$$

where ω_f is the modulus of continuity of f on Γ , $\mathbf{u} := (\langle v_n \rangle)_0^\infty$ is the projection of \mathbf{v} on \mathbb{T}^d , $D_N(\mathbf{u})$ is its N -term discrepancy and ρ_Γ is defined by (5).

Proof. For consistency of our notation, we first identify \mathbb{T}^d with $[0, 1)^d$. For any $k \in \mathbb{Z}^d$, let $Q_k := k + [0, 1)^d$ and $\Gamma_k := Q_k \cap \Gamma$. Then, for some finite subset γ of \mathbb{Z}^d , Γ is the disjoint union of Γ_k , $k \in \gamma$, and $[0, 1)^d$ is the disjoint union of $\langle \Gamma_k \rangle = \Gamma_k - k$, $k \in \gamma$. Clearly, each Γ_k , and therefore $\langle \Gamma_k \rangle$, is Jordan measurable. Let $g_k := g 1_{\langle \Gamma_k \rangle}$, where $g := f \circ \langle \cdot \rangle_\Gamma$, so that $g = \sum_{k \in \gamma} g_k$. Let \mathcal{P} be a partition of $[0, 1)^d$ with $\|\mathcal{P}\| \leq t$. It is clear that

$$s(g, \mathcal{P}) \leq \sum_{k \in \gamma} s(g_k, \mathcal{P}).$$

For each $k \in \gamma$, let \mathcal{P}_k be the collection of intervals in \mathcal{P} that are fully contained in $\langle \Gamma_k \rangle$, let \mathcal{P}_k^c be those that are fully contained in $\langle \Gamma_k \rangle^c = [0, 1)^d \setminus \langle \Gamma_k \rangle$ and let $\partial \mathcal{P}_k$ be the remaining ones. Note that $g_k(u) = f(u + k) 1_{\langle \Gamma_k \rangle}(u)$. Hence

$$\sum_{I \in \mathcal{P}_k} \left(\sup_I g_k - \inf_I g_k \right) \lambda_d(I) \leq \omega_f(t\sqrt{d}) \sum_{I \in \mathcal{P}_k} \lambda_d(I) \leq \omega_f(t\sqrt{d}) \lambda_d(\langle \Gamma_k \rangle),$$

where the last inequality uses the fact that the intervals in \mathcal{P} are disjoint. Note that the corresponding sum over \mathcal{P}_k^c is identically zero. For each $I \in \partial \mathcal{P}_k$, we have $I \cap \langle \Gamma_k \rangle \neq \emptyset$ and $I \cap \langle \Gamma_k \rangle^c \neq \emptyset$, and therefore

$$I \subset (\langle \Gamma_k \rangle \cap \mathcal{N}_{t\sqrt{d}}(\langle \Gamma_k \rangle^c)) \cup (\langle \Gamma_k \rangle^c \cap \mathcal{N}_{t\sqrt{d}}(\langle \Gamma_k \rangle)) = \mathcal{N}_{t\sqrt{d}}(\partial \langle \Gamma_k \rangle),$$

where the last equality is due to Lemma 6.1 (see below). Note that the boundary of $\langle \Gamma_k \rangle$ is with respect to the Euclidean metric on $[0, 1)^d$. Hence

$$\sum_{I \in \partial \mathcal{P}_k} \left(\sup_I g_k - \inf_I g_k \right) \lambda_d(I) \leq 2\|f\|_\infty \sum_{I \in \partial \mathcal{P}_k} \lambda_d(I) \leq 2\|f\|_\infty \lambda_d(\mathcal{N}_{t\sqrt{d}}(\partial \langle \Gamma_k \rangle)),$$

where we have used the disjointness of $I \in \mathcal{P}$ again. Combining this with the bound for the sum over \mathcal{P}_k , we get

$$s(g_k, \mathcal{P}) \leq \omega_f(t\sqrt{d}) \lambda_d(\langle \Gamma_k \rangle) + 2\|f\|_\infty \lambda_d(\mathcal{N}_{t\sqrt{d}}(\partial \langle \Gamma_k \rangle)).$$

Summing over $k \in \gamma$ and taking the supremum over \mathcal{P} , we now get

$$S(g, t) \leq \omega_f(t\sqrt{d}) + 2\|f\|_\infty \sum_{k \in \gamma} \lambda_d(\mathcal{N}_{t\sqrt{d}}(\partial \langle \Gamma_k \rangle)). \tag{32}$$

We claim that

$$\sum_{k \in \gamma} \lambda_d(\mathcal{N}_{t\sqrt{d}}(\partial \langle \Gamma_k \rangle)) \leq \lambda_d(\mathcal{N}_{t\sqrt{d}}(\partial \Gamma)), \tag{33}$$

where the boundary of Γ is with respect to the Euclidean metric on \mathbb{R}^d . To see this, note that, for any $\epsilon > 0$, $x \in \mathcal{N}_\epsilon(\partial \langle \Gamma_k \rangle)$ means that there exist $x^* \in [0, 1)^d$ such that $|x - x^*| < \epsilon$ and two sequences, (x_n) in $\langle \Gamma_k \rangle$ and (x'_n) in $\langle \Gamma_k \rangle^c$, such that $x^* = \lim x_n = \lim x'_n$. Then the sequences $(x_n + k)$ and $(x'_n + k)$ are, respectively, in $\Gamma_k \subset \Gamma$ and $Q_k \setminus \Gamma_k \subset \Gamma^c$, both having the limit $x^* + k$, which must be in $\partial \Gamma$. This implies that $x + k \in \mathcal{N}_\epsilon(\partial \Gamma)$ and therefore we get $V_k := \mathcal{N}_\epsilon(\partial \langle \Gamma_k \rangle) + k \subset \mathcal{N}_\epsilon(\partial \Gamma)$. Since the V_k , $k \in \mathbb{Z}^d$, are disjoint, by translation invariance of Lebesgue measure and setting $\epsilon = t\sqrt{d}$, we obtain (33) which now yields

$$S(g, t) \leq \omega_f(t\sqrt{d}) + 2\|f\|_\infty \rho_\Gamma(t\sqrt{d}). \tag{34}$$

The proof is completed by setting $t = \lfloor D_N(\mathbf{u})^{-1} \rfloor^{-1/d}$ in this bound and using Hlawka's inequality (30) for $h = g$. □

Remark. There are other generalizations of the Koksma-Hlawka inequality that can be used in our analysis. Among them, we would like to single out the one given in [BCGT13], which is applicable to piecewise smooth functions with singularities along arbitrary Borel sets. In our context, this would mean restricting our analysis to smooth f (at least C^d); however, the resulting bound would also be somewhat stronger, especially if additional structure is available regarding Γ .

What remains to be shown is the following lemma.

LEMMA 6.1. *Let X be a convex subset of \mathbb{R}^d , considered as a metric space equipped with the Euclidean metric. For any $A \subset X$ and $t > 0$,*

$$\mathcal{N}_t(\partial A) = (A \cap \mathcal{N}_t(A^c)) \cup (A^c \cap \mathcal{N}_t(A)).$$

Proof. Note that $\mathcal{N}_t(\partial A) \subset \mathcal{N}_t(\overline{A}) = \mathcal{N}_t(A)$, so that $A^c \cap \mathcal{N}_t(\partial A) \subset A^c \cap \mathcal{N}_t(A)$. Similarly, $A \cap \mathcal{N}_t(\partial A) \subset A \cap \mathcal{N}_t(A^c)$. Taking the union, it follows that $\mathcal{N}_t(\partial A) \subset (A \cap \mathcal{N}_t(A^c)) \cup (A^c \cap \mathcal{N}_t(A))$.

For the reverse inclusion, let $x \in (A \cap \mathcal{N}_t(A^c)) \cup (A^c \cap \mathcal{N}_t(A))$. Without loss of generality, we may assume that $x \notin \partial A$ as, otherwise, $x \in \mathcal{N}_t(\partial A)$ trivially. The first case is $x \in A^\circ \cap \mathcal{N}_t(A^c)$. There is $v \in A^c$ such that $|x - v| < t$. If $v \in \partial A$, then $x \in \mathcal{N}_t(\partial A)$, so assume that $v \in (A^c)^\circ$. Consider the line segment $I := [x, v]$ with its relative topology. Then $I \cap A^\circ$ and $I \cap (A^c)^\circ$ are non-empty open subsets of the connected space I , which implies that their union is not all of I , i.e., there is a point $y^* \in I \setminus (A^\circ \cup (A^c)^\circ) = (x, v) \cap \partial A$, and therefore $x \in \mathcal{N}_t(\partial A)$. The second case $x \in (A^c)^\circ \cap \mathcal{N}_t(A)$ is handled similarly by switching the roles of A and A^c . Hence $(A \cap \mathcal{N}_t(A^c)) \cup (A^c \cap \mathcal{N}_t(A)) \subset \mathcal{N}_t(\partial A)$. \square

Remark. While the first inclusion above is valid in a general metric space, convexity was used critically for the reverse inclusion. A counterexample without convexity is $X = (-1, 1) \setminus \{0\}$ and $A := (0, 1)$.

7. Examples and counterexamples, motivation and applications, open problems and challenges

7.1. *Motivation and applications to algorithmic A/D conversion.* Piecewise affine maps are typically found in applications where switching plays an important role, such as hybrid control systems [DDBB⁺09] and A/D conversion algorithms with feedback (algorithmic converters) [DGWY10, Gün12]. The subfamily of piecewise-affine automorphisms that have motivated this paper appear in the latter setting. Below we give a brief overview of algorithmic A/D converters and go over some important applications of our results.

By an algorithmic A/D converter, we mean a process of encoding signals into discrete valued sequences that is implemented by carrying out an autonomous operation (the algorithm) on some auxiliary state space. More precisely, let \mathcal{X} be a space of input signals, let Λ be a finite index set and let \mathcal{V} be a chosen state space, such as \mathbb{R}^d . The algorithm is implemented using two associated maps; given the input and the current state, the first map $F : \mathcal{X} \times \mathcal{V} \rightarrow \mathcal{V}$ determines the next state, and a second map $Q : \mathcal{X} \times \mathcal{V} \rightarrow \Lambda$ determines the next output index. Q is also called a ‘quantization rule’. Given an initial

state v_0 (which may depend on x), the recursion

$$\begin{cases} v_n = F(x, v_{n-1}) \\ i_n = Q(x, v_{n-1}) \end{cases}, \quad n = 1, 2, \dots \tag{35}$$

defines the overall encoding $(x \mapsto (i_n)_1^\infty)$. It is required that the map $(x \mapsto (i_n)_1^\infty)$ is invertible on \mathcal{X} , and also, given a metric on \mathcal{X} , it is desirable that each $x \in \mathcal{X}$ can be approximately recovered from $(i_n)_1^N$ with high accuracy as N increases.

Note that the map Q can be equivalently described by a partition $\{\Omega_{x,i}^Q\}_{i \in \Lambda}$, where $\Omega_{x,i}^Q := \{v \in \mathcal{V} : Q(x, v) = i\}$. Similarly, F can be seen as a family of maps $\{T_x\}_{x \in \mathcal{X}}$ on \mathcal{V} , where $T_x(v) := F(x, v)$. Then $(i_n)_1^\infty$ becomes the itinerary of the orbit of v_0 with respect to T_x and the partition $\{\Omega_{x,i}^Q\}_{i \in \Lambda}$. In general, F and Q need not be linked, but in most examples of algorithmic converters, T_x is a piecewise-affine map whose defining partition coincides with the partition $\{\Omega_{x,i}^Q\}_{i \in \Lambda}$ induced by Q .

While the algorithmic converter model is general enough to accommodate time-varying signals $x = (x_n)_1^\infty$ as well, in the examples we shall see below, we will only consider constant input signals, i.e., \mathcal{X} will be a subset of \mathbb{R} .

We will present two examples of algorithmic A/D converters that are relevant to this paper: *the golden ratio encoder* and *sigma-delta quantization*.

7.1.1. The golden ratio encoder. This encoder was proposed in [DGWY10] to address some practical considerations regarding electronic circuit implementation of β -expansions. Let us first recall β -expansions. For any $\beta > 1$, let $\Lambda = \{0, \dots, \lceil \beta \rceil - 1\}$. For simplicity, let us assume that $1 < \beta \leq 2$ so that $\Lambda = \{0, 1\}$. Pick any $\eta \in [(\beta - 1)\beta, \beta]$, and define $T_x = T$ on $\mathcal{V} = \mathbb{R}$ by

$$T(v) := \beta v - Q(v) := \beta v - \begin{cases} 0 & \text{if } \eta v < 1, \\ 0 \text{ or } 1 & \text{if } \eta v = 1, \\ 1 & \text{if } \eta v > 1. \end{cases}$$

We set $v_0 = x \in \mathcal{X} = [0, (\beta - 1)^{-1}]$ and $i_n = Q(v_{n-1})$. It is easily checked that T maps $[0, (\beta - 1)^{-1}]$ into itself. If $\eta = \beta$ and $Q(\eta^{-1}) = 1$, one gets the so-called ‘greedy’ β -expansion, whereas if $\eta = \beta(\beta - 1)$ and $Q(\eta^{-1}) = 0$, one gets the ‘lazy’ β -expansion [DK02]. The values of α strictly in between these two extremes correspond to ‘cautious’ β -expansions [DDGV06]. In all cases, x can be recovered via the inversion formula $x = \sum_1^\infty i_n \beta^{-n}$.

Clearly, β -expansions are not associated with a toral automorphism in the above form. However, certain special values of β (namely, the *Pisot units*) can be realized in a toral automorphism. The golden ratio encoder uses $\beta = \phi := (1 + \sqrt{5})/2$.

The starting point of the golden ratio encoder is the ‘multiplier-free’ recursion

$$w_n = w_{n-1} + w_{n-2} - i_n. \tag{36}$$

Noting that $\phi^2 = \phi + 1$, it is straightforward to check that if $\phi^{-1}w_{-1} + w_0 = x$, then any encoding decision (i_n) that yields a bounded solution (w_n) results in $x = \sum_1^\infty i_n \phi^{-n}$.

In order to frame this encoder in our formulation of piecewise-affine automorphisms, let $v_n := (w_n, w_{n-1})$. Then

$$v_n = Lv_{n-1} - i_n e,$$




FIGURE 1. An orbit closure for the golden ratio encoder for $\eta = (1.4, 0.8)$ along with two of its integer translates as a demonstration of its tiling property.

where $L(v^1, v^2) := (v^1 + v^2, v^1)$ and $e := (1, 0)$. L defines a hyperbolic toral automorphism on \mathbb{T}^2 with eigenvalues ϕ and $-1/\phi$. Once a quantization rule $i_n := Q(v_{n-1}) \in \Lambda := \{0, 1\}$ is specified, we obtain a piecewise-affine automorphism $T \in \mathcal{P}(L, 0)$ with $\Omega_i := Q^{-1}(i)$ and $\tau_i := -ie, i \in \Lambda$.

As in the case of classical β -expansions, there is some freedom in the choice of the quantization rule that yields bounded orbits. In [DGWY10], a parametric family of such rules was given that correspond to half-space partitions

$$\Omega_0 = \{v \in \mathbb{R}^2 : \eta \cdot v < 1\}, \quad \Omega_1 = \Omega_0^c.$$

Let us denote the resulting piecewise-affine automorphism by T_η . The simplest case is $\eta = (1, 1)$ when $[0, 1]^2$ is seen to be invariant under $T_{(1,1)}$. Similarly to β -expansions, it turns out that there is an open set U for the parameter η (not containing $(1, 1)$) and an open bounded set V for the state variable v_n such that $T_\eta(V) \subset V$ for all $\eta \in U$, thereby allowing for robust implementation.

As a particular case of our results in this paper, it follows that, inside V , Lebesgue a.e. orbit closure for T is a tile. An example is given in Figure 1.

7.1.2. Sigma-delta quantization. The most basic form of $\Sigma\Delta$ quantization is based on the difference equation

$$(\Delta w)_n := w_n - w_{n-1} = x - i_n. \tag{37}$$

For each x , we are interested in an encoding (i_n) that results in a bounded solution (w_n) . It then follows that

$$x = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N i_n.$$

The rate of convergence above is $O(N^{-1})$ and generally not better. Inherently more efficient encodings are obtained within the unipotent family by using the difference operator Δ^m instead. Here m is a positive integer that stands for the order of the resulting $\Sigma\Delta$ quantization scheme. In this case, for any $\varphi \in \ell^1$ with $\sum \varphi_n = 1$,

$$x - \sum \varphi_n i_n = \sum \varphi_n (\Delta^m w)_n = \sum (\tilde{\Delta}^m \varphi)_n w_n, \tag{38}$$

where $(\tilde{\Delta}\varphi)_n := \varphi_n - \varphi_{n+1}$ is the adjoint of Δ . For an N -term approximation, we require that the support of φ is $\{1, \dots, N\}$. It is not hard to find such a φ with $\sum \varphi_n = 1$ and $\|\tilde{\Delta}^m \varphi\|_1 = O(N^{-m})$ so that

$$\left| x - \sum \varphi_n i_n \right| = O(N^{-m}).$$

The price paid for the higher accuracy of higher-order $\Sigma\Delta$ schemes is that, in order to keep (w_n) bounded, one either has to use a larger index set Λ for the same set \mathcal{X} of inputs or seek increasingly more complicated quantization rules for setting i_n in a small Λ . Which route would be more feasible depends on the constraints of a given application.

In order to formulate $\Sigma\Delta$ quantization as an algorithmic converter and, in particular, as a piecewise-affine automorphism on the state space $\mathcal{V} = \mathbb{R}^m$, we set

$$v_n := ((\Delta^{m-1} w)_n, (\Delta^{m-2} w)_{n-1}, \dots, w_{n-m+1}).$$

Then it follows that

$$v_n = L v_{n-1} + (x - i_n) e,$$

where this time $L(v^1, v^2, \dots, v^m) := (v^1, v^2 + v^1, \dots, v^m + v^{m-1})$ and $e := (1, 0, \dots, 0)$. Note that the map $S(v) := Lv + xe$ defines a generalized skew translation on \mathbb{T}^m , and once $\Lambda \subset \mathbb{Z}$ and $i_n = Q(x, v_{n-1})$ is specified, we get a piecewise-affine automorphism $T_x \in \mathcal{P}(L, xe)$ with $\tau_i = (x - i)e, i \in \Lambda$. Typical quantization rules Q used in practice are ‘linear’ in the sense that $Q(x, v) = \text{round}_\Lambda(\eta \cdot v + \gamma x)$, although more general partitions have been proposed as well (e.g., [DD03, GT04]).

Two examples of tiling orbit closures are illustrated in Figures 2 and 3. In Figure 2, we plot an orbit of a two-dimensional piecewise-affine automorphism associated with a second-order $\Sigma\Delta$ quantization scheme with $\Lambda = \{0, 1\}$ and the partition $\{\Omega_0, \Omega_1\}$ is defined according to the linear rule of the previous paragraph with $\eta = (\frac{3}{2}, 1)$ and $\gamma = 0$. In this example, $x = 1/\sqrt{5}$. It can be seen that the boundary of the orbit closure is significantly more complicated than the linear boundary of the partition. We have observed that the resulting tiles become even more complicated if the slope $s := \eta_1/\eta_2$ is decreased. Nevertheless, as we have shown in this paper, all of these tiles are Jordan measurable. On the other hand, these tiles become more regular as s is increased. In fact, for $s > 2$, they were identified by inspection to be polygonal single tiles [GT04].

In Figure 3, we demonstrate another orbit, using the same L and Λ , but this time with a partition defined by a piecewise linear curve. We have found that the orbit closure shown in (a) yields an invariant 2-tile Γ : the T -invariance of Γ is demonstrated in (b), and the fact that Γ is a 2-tile is demonstrated in (c) and (d), where two single tiles Γ_1 and Γ_2 are identified. This decomposition, however, is certainly not unique.

In the case of single tiles, our results on the regularity of these tiles and the corresponding estimates on the rate of convergence of ergodic averages allow us to better




FIGURE 2. An orbit closure for the second-order $\Sigma\Delta$ quantization with the ‘linear’ rule: $x = 1/\sqrt{5}$, $\Lambda = \{0, 1\}$, $\eta = (\frac{3}{2}, 1)$, $\gamma = 0$.

quantify the rate of convergence of approximations in $\Sigma\Delta$ quantization. Note that (38) represents the error as an average over (w_n) . First, note that $w_n = f(v_n) := l \cdot v_n$, where $l_i = \binom{m-1}{i-1}$, $i = 1, \dots, m$. Define $\mu := \int_{\Gamma} f(v) dv$ and $W_l := \sum_1^l (w_n - \mu)$. Then $(\Delta W)_n = w_n - \mu$ so that

$$x - \sum \varphi_n i_n = \sum (\tilde{\Delta}^m \varphi)_n w_n = \sum (\tilde{\Delta}^{m+1} \varphi)_n W_n, \tag{39}$$

where the last equality uses the fact that $\sum (\tilde{\Delta}^m \varphi)_n = 0$. Noting that $|W_n| = n D_n(f, V)$, it follows that any effective bound on the discrepancy of the form $D_N(f, V) \lesssim N^{-\delta}$ for some $\delta \in (0, 1]$ implies that

$$\left| x - \sum \varphi_n i_n \right| \lesssim \|\tilde{\Delta}^{m+1} \varphi\|_1 N^{1-\delta}, \tag{40}$$

which yields an improved bound of $O(N^{-m-\delta})$ using a suitable φ with $\|\tilde{\Delta}^{m+1} \varphi\|_1 = O(N^{-m-1})$, e.g., by means of a discrete B -spline of degree $m + 1$. To achieve this via Theorem 6.1, one would need to analyze the invariant sets Γ of the specific examples further and estimate the function ρ_{Γ} .

7.2. *Further examples and counterexamples.* In this paper, we have alluded to examples and counterexamples of piecewise-affine transformations with various special properties, which we now provide. Straightforward details are omitted.

Example 7.1. (T with no bounded orbit and no invariant set) Let $0 < a < 1$, and let $T : \mathbb{R} \rightarrow \mathbb{R}$ be given by

$$T(v) := v + a - \mathbf{1}_{(-\infty, 0)}(v).$$

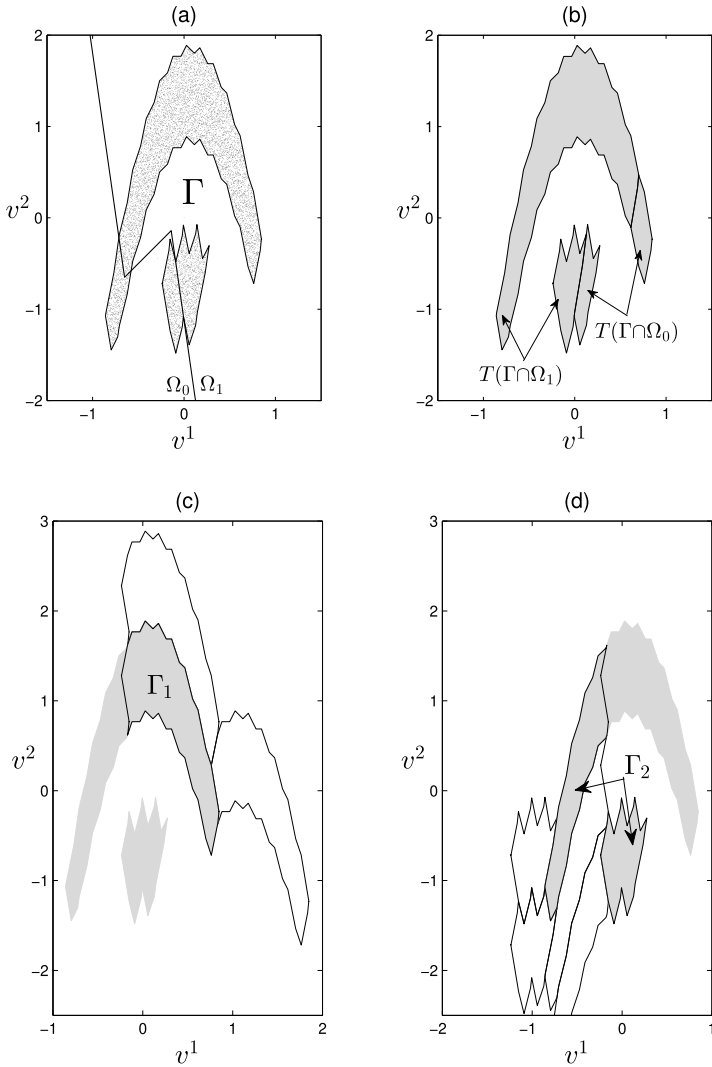


FIGURE 3. An orbit closure yielding a 2-tile.

Example 7.2. (T with a bounded forward orbit but with no strictly invariant set) Let $0 < a < 1$ be any irrational, let $A := \{na \bmod 1 : n \geq 0\}$ and let $T : \mathbb{R} \rightarrow \mathbb{R}$ be given by

$$T(v) := v + a - (1_{(-\infty, 0)} + 1_{A \cap [1-a, 1)} - 1_{[0, \infty) \setminus A})(v).$$

Example 7.3. (T ergodic on a non-trivial invariant 2-tile) Consider

$$T(v) := v + a + 1_{[1-a, 3/2)}(v) - 3 \cdot 1_{[5/2, \infty)}(v),$$

where $0 < a < 3/4$ is an irrational number. Then $\Gamma := [a - \frac{1}{2}, 1) \cup [2, a + \frac{5}{2})$ is a T -invariant 2-tile on which T is ergodic.

Example 7.4. (T not ergodic on an invariant 2-tile, despite ergodic S) Let $0 < a < 1$ and let $T : \mathbb{R} \rightarrow \mathbb{R}$ be given by

$$T(v) := v + a - 1_{[1-a, 1) \cup [2-a, \infty)}(v).$$

$[0, 1)$ and $[1, 2)$ are both T -invariant.

Example 7.5. (T not ergodic on a 2-tile orbit closure, despite ergodic S and Jordan measurable Ω) Let α be an irrational number in $(0, 3/4)$. With $m_0 := 0$, let us define $(n_k)_1^\infty$ and $(m_k)_1^\infty$ recursively by

$$\begin{aligned} n_k &:= \min\{n : n > m_{k-1} + k \text{ and } 2^{-2k-1} < \langle n\alpha \rangle < 2^{-2k}\}, \quad k \geq 1, \\ m_k &:= \min\{m : m > n_k + k \text{ and } 2^{-2k-2} < \langle m\alpha \rangle < 2^{-2k-1}\}, \quad k \geq 1. \end{aligned}$$

With these two sequences, we define $T : [0, 2] \rightarrow [0, 2]$ by

$$T(v) := v + \alpha - 1_{[1-\alpha, 1) \cup [2-\alpha, 2]}(v) + 1_{\mathcal{N}}(v) - 1_{\mathcal{M}}(v),$$

where $\mathcal{N} := \{\langle n_k\alpha \rangle : k \geq 1\}$ and $\mathcal{M} := \{1 + \langle m_k\alpha \rangle : k \geq 1\}$. Note that $\mathcal{N} \subset (0, 1/4)$ and $\mathcal{M} \subset (1, 5/4)$ so that $[1 - \alpha, 1) \cap \mathcal{N} = \emptyset$ and $[2 - \alpha, 2] \cap \mathcal{M} = \emptyset$.

Then:

- $\partial\Omega = \{0, 1, 2, 1 - \alpha, 2 - \alpha\} \cup \mathcal{M} \cup \mathcal{N}$, so Ω is Jordan measurable;
- $[0, 1]$ and $[1, 2]$ are essentially T -invariant, so T is not ergodic on $[0, 2]$; and
- the orbit $V = \mathcal{O}_T^+(0)$ is dense in $[0, 2]$.

Remark. In relation to Example 7.5, we learned from one of the anonymous referees the following related (and more interesting) example due to Veech ([Vee69], [MT02, p. 1035]). Let θ and α be two numbers in $(0, 1)$. Take two copies of the unit circle (with unit circumference) and mark off the segment $[0, \alpha]$ on each of them. A map on the union of these two circles is defined as follows. Every point is rotated by θ on the circle on which it sits, but if the point lands in the segment $[0, \alpha]$, then it is moved to the same position on the other circle. Veech showed that if θ is an irrational number with unbounded partial quotients, then there are irrational α for which this system is minimal, yet not ergodic with respect to the Lebesgue measure. Clearly, this map can be implemented as a piecewise-affine automorphism on, say, $[0, 2]$.

Example 7.6. (Piecewise extension of a toral endomorphism with a non-tiling invariant set) Consider the map

$$T(v) := 2v - 1_{[3/4, 1)}(v) - 2 \cdot 1_{(1, 3/2]}(v).$$

The interval $[0, \frac{3}{2}]$ is the largest bounded set that is invariant under T . T is ergodic on this interval with respect to the measure with density ρ given by

$$\rho(v) := \frac{1}{2} \cdot 1_{[0, 1/2)}(v) + 1_{[1/2, 1)}(v) + \frac{1}{2} \cdot 1_{[1, 3/2]}(v).$$

(The invariance of ρ is straightforward; the ergodicity follows, for example, by [BG97, ‘Folklore Theorem’ 6.1.1].) Since this measure is equivalent to the Lebesgue measure on $[0, 3/2]$, it follows that the orbit closure of almost every initial point is equal to $[0, 3/2]$. Note, however, that the relation $\sum_{n \in \mathbb{Z}} \rho(v + n) \equiv 2$ may be considered to be a general form of tiling for the invariant density. This is generally true.

7.3. *Some open problems and challenges.* We close this section with a selection of open problems and challenges.

- (1) *Stability.* As we stated in the introduction, the problem of determining whether all trajectories of a general piecewise-affine map remain bounded is out of reach even for relatively mild classes of maps. However, it would still be useful to have partial information for applications. For example, are there useful criteria that can be applied to these maps to determine a ‘region of stability,’ i.e., a set of starting points resulting in bounded orbits?
- (2) *Single vs multiple tiles.* What are the general mechanisms behind the generation of single tiles vs multiple tiles as invariant sets?
- (3) *Shapes of tiles.* What is the mapping that takes a piecewise-affine automorphism to its invariant set(s)? In particular, knowing the linear part L and the translations τ_i and, say, in the case of a single tile invariant set Γ , how is the shape of Γ related to the shape of the partition Ω ?
- (4) *Regularity of tiles.* If we know the regularity properties of Ω , what can we say about the regularity of Γ , e.g., in terms of the ρ_Γ function, or simply in terms of the Hausdorff dimension of $\partial\Gamma$?
- (5) *Approximating tiles by orbits and orbit segments.* When is it possible for a piecewise-affine automorphism considered in this paper to have the property that the forward orbit of every non-empty open ball in an invariant tile covers the tile in finite time?
- (6) *Tiles of infinite area?* Our results in this paper are based on invariant sets of finite measure. However, it is possible to have non-trivial examples of unbounded invariant sets of infinite area. But an m -tile for $m = \infty$ is somewhat ambiguous. Is there any analog of tiling in this setting? Also, along these lines, what do infinite-area orbit closures look like?

Acknowledgement. The authors gratefully acknowledge the useful comments and the suggestions made by the anonymous referees.

A. Appendix. Technical lemmas

Let X be a topological space, $f : X \rightarrow \mathbb{R}$. Recall that f is lower semi-continuous (l.s.c.) on X if and only if $\{f > \alpha\}$ is open for all $\alpha \in \mathbb{R}$ and f is upper semi-continuous (u.s.c.) on X if and only if $\{f < \alpha\}$ is open for all $\alpha \in \mathbb{R}$. For $f : X \rightarrow \mathbb{N}$, it follows that f is l.s.c. if and only if $\{f \geq n\}$ is open for all $n \geq 0$ and f is u.s.c. if and only if $\{f \leq n\}$ is open for all $n \geq 0$. In particular, the indicator function 1_A is l.s.c. if and only if A is open and 1_A is u.s.c. if and only if A is closed. Semi-continuity is preserved under addition and multiplication by a non-negative constant.

LEMMA A.1. *Let $\Pi \in \mathcal{C}$ and let A be a bounded subset of $\tilde{\Pi}$.*

- (i) *If A is open in $\tilde{\Pi}$, then N_A is l.s.c. on Π .*
- (ii) *If A is closed in $\tilde{\Pi}$, then N_A is u.s.c. on Π .*

Proof. To show that N_A is l.s.c. (u.s.c.) on Π , it suffices to show that, for any $u \in \Pi$, there is a neighborhood \mathcal{N} of u in Π such that $1_{\mathcal{N}}N_A$ is l.s.c. (u.s.c.) on Π . Let us denote the projection $\langle \cdot \rangle$ by p .

(i) Let A be open in $\tilde{\Pi}$. For any $u \in \Pi$, let U be an open neighborhood of u that is also the homeomorphic image of a bounded open set V in $\tilde{\Pi}$ such that $k + V$ is disjoint from V for all $k \in \mathbb{Z}^d \setminus \{0\}$. Let $V_k := k + V, k \in \mathbb{Z}^d$. Then each V_k is also homeomorphic to U and $\{V_k : k \in \mathbb{Z}^d\}$ is a partition of $p^{-1}(U)$ into open sets. Since A is bounded, there exists a finite index set I such that $\{A \cap V_k : k \in I\}$ partitions $A \cap p^{-1}(U)$, again into open sets. Hence

$$1_U N_A = 1_U \sum_{k \in I} 1_{A \cap V_k}(p^{-1}(\cdot)) = \sum_{k \in I} 1_{\langle A \cap V_k \rangle},$$

where, in the last step, we have used the fact that $\langle A \cap V_k \rangle \subset U$. Since each $\langle A \cap V_k \rangle$ is open, it follows that $1_{\langle A \cap V_k \rangle}$ and therefore $1_U N_A$ is l.s.c. on Π .

(ii) The proof will be similar to part (i). Let A be closed in $\tilde{\Pi}$. For any $u \in \Pi$, let U, V, V_k be the same as above. Pick a closed set $F \subset U$ with $u \in \overset{\circ}{F}$, and let $W_k := V_k \cap p^{-1}(F)$ so that W_k is homeomorphically mapped to F by p for each $k \in \mathbb{Z}^d$. As before, we have that, for some finite index set $I, \{A \cap W_k : k \in I\}$ is a partition of $A \cap p^{-1}(F)$ into closed sets. Hence

$$1_F N_A = 1_F \sum_{k \in I} 1_{A \cap W_k}(p^{-1}(\cdot)) = \sum_{k \in I} 1_{\langle A \cap W_k \rangle}.$$

Since each $\langle A \cap W_k \rangle$ is closed, it follows that $1_F N_A$ is u.s.c. on Π . □

COROLLARY A.1. *Let $\Pi \in \mathcal{C}$. If A is a bounded essential m -tile for $\tilde{\Pi}$, then $N_A^\circ \leq m \leq N_{\bar{A}}$ on Π .*

Proof. The desired relation $N_A^\circ \leq m \leq N_{\bar{A}}$ readily holds μ_Π -a.e., since $N_A^\circ \leq N_A \leq N_{\bar{A}}$ holds everywhere and $N_A = m$ μ_Π -a.e. Meanwhile, Lemma A.1 implies that N_A° is l.s.c. and $N_{\bar{A}}$ is u.s.c., so the sets $\{N_A^\circ > m\}$ and $\{N_{\bar{A}} < m\}$ are simultaneously μ_Π -null and open, and therefore empty. □

LEMMA A.2. *Let $\Pi \in \mathcal{C}$. Suppose A and B are two subsets of $\tilde{\Pi}$ such that $A \subset B$ and $N_A \leq m \leq N_B$ on Π for some integer $m \geq 0$. Then there exists an exact m -tile Γ for $\tilde{\Pi}$ such that $A \subset \Gamma \subset B$. If A and B are Borel measurable, then Γ can be chosen to be Borel measurable as well.*

Proof. We will construct Γ using the following ‘greedy water-filling’ procedure. Given any exact tile Q for $\tilde{\Pi}$, let $Q_k := Q + k, k \in \mathbb{Z}^d$. Let $(k_n)_{n=1}^\infty$ be any enumeration of \mathbb{Z}^d . Let $A_0 := A$. For $n = 1, 2, \dots$, define

$$I_n := \{u \in \Pi : N_{A_{n-1}}(u) < m\} \tag{41}$$

and

$$A_n := A_{n-1} \cup E_n, \tag{42}$$

where

$$E_n := p^{-1}(I_n) \cap Q_{k_n} \cap (B \setminus A).$$

Here $p(v) := \langle v \rangle$.

Note that since $(Q_{k_n})_1^\infty$ is a disjoint family, $(E_n)_1^\infty$ is also a disjoint family. Moreover, all the E_n are disjoint from $A_0 = A$. Therefore, for any $n \geq 1$, (42) is a disjoint union. Since E_n is subset of a tile, we have $N_{E_n} = 1_{\langle E_n \rangle}$, and therefore

$$N_{A_n} = N_{A_{n-1}} + 1_{\langle E_n \rangle}.$$

Furthermore, p is bijective when restricted to any tile, so

$$\langle E_n \rangle = \langle p^{-1}(I_n) \cap Q_{k_n} \rangle \cap \langle Q_{k_n} \cap (B \setminus A) \rangle = I_n \cap \langle Q_{k_n} \cap (B \setminus A) \rangle.$$

This relation shows that if $N_{A_{n-1}}(u) = m$ for $u \in \Pi$ and $n \geq 1$, then $u \notin I_n$ and therefore $u \notin \langle E_n \rangle$ so that $N_{A_l}(u) = m$ for all $l \geq n$. Since $N_{A_0} \leq m$, we get that $N_{A_n} \leq m$ for all n . Also, since $A_0 \subset B$ and $E_n \subset B$, it follows from (42) that $A_n \subset B$ for all n . We define

$$\Gamma := \bigcup_{n=0}^\infty A_n = A \cup \bigcup_{n=1}^\infty E_n,$$

the latter expression being a disjoint union. Hence

$$N_\Gamma = \lim_{n \rightarrow \infty} N_{A_n} = N_A + \sum_{n=1}^\infty N_{E_n} = N_A + \sum_{n=1}^\infty 1_{I_n} 1_{\langle Q_{k_n} \cap (B \setminus A) \rangle}. \tag{43}$$

The first relation above shows that $N_\Gamma \leq m$. Hence, the infinite sums above reduce to a finite sum for every point in Π .

Clearly, $A \subset \Gamma \subset B$. We claim that $N_\Gamma = m$ everywhere. Let $u \in \Pi$ be arbitrary and suppose that $N_\Gamma(u) < m$. Then $N_{A_{n-1}}(u) \leq N_\Gamma(u)$ implies that $u \in I_n$ for all n . Therefore it follows from the last equality in (43) that $N_\Gamma(u) = N_A(u) + N_{B \setminus A}(u) = N_B(u)$. Since $N_B(u) \geq m$, we get a contradiction. Hence $N_\Gamma(u) = m$.

Finally, if A and B are Borel sets, then choosing Q to be Borel implies recursively that all E_n , A_n , and therefore Γ , are also Borel. □

LEMMA A.3. *Let $\Pi \in \mathcal{C}$ and let D be a compact nowhere dense subset of $\tilde{\Pi}$. Then $\langle D \rangle$ is a compact nowhere dense subset of Π .*

Proof. It is clear that $\langle D \rangle$ is compact. Suppose $\langle D \rangle$ contains an open ball B . Without loss of generality, we may assume that the radius of B is strictly less than $1/2$. Hence there exists an open ball \tilde{B} in $\tilde{\Pi}$ such that $\{\tilde{B} + k\}_{k \in \mathbb{Z}^d}$ is a partition of the preimage of B under the canonical projection. Since D is bounded and $B \subset \langle D \rangle$, there exists a finite set $K \subset \mathbb{Z}^d$ such that $B = \bigcup_{k \in K} \langle D \cap (\tilde{B} + k) \rangle$. Note that $D \cap (\tilde{B} + k) \subset D$ and D having empty interior implies that $D \cap (\tilde{B} + k)$ is nowhere dense in $\tilde{\Pi}$. Since $\langle \cdot \rangle$ is a homeomorphism between \tilde{B} and B , it follows that $\langle D \cap (\tilde{B} + k) \rangle$ is nowhere dense in Π . This contradicts the fact that a finite union of nowhere dense sets is nowhere dense. □

B. Appendix. Invariant sets of finite-to-one maps

Let X be any set and let $f : X \rightarrow X$ be any map, not necessarily onto. As before, we say that a set A is f -invariant if $f(A) = A$. It is easy to see that the set $X_f := \bigcap_{n=0}^\infty f^n(X)$

is positively f -invariant, i.e., $f(X_f) \subset X_f$. However, X_f need not be f -invariant in general. (For example, let $X := \{(x, y) \in \mathbb{N}^2 : y \leq x\} \cup \{a, b\}$, where a and b are two new points (say, in \mathbb{Z}^2) and define $f : X \rightarrow X$ via $f(x, y) := (x, y - 1)$ if $y \geq 1$, $f(x, 0) := a$, $f(a) := b$ and $f(b) = b$. Then $X_f = \{a, b\}$ and $f(X_f) = \{b\}$.) It turns out, as we will show below, that X_f is f -invariant if $\text{card}(f^{-1}(x)) < \infty$ for all $x \in X$. In this case, we will say that f has *finite preimages* or f is *finite-to-one* [PH15]. Clearly, the piecewise-affine automorphisms that are considered in this paper have this property since they are constructed by piecing together finitely many injective maps.

The following theorem is a more general form of the claim made above.

THEOREM B.1. *Let $f : X \rightarrow X$ be finite-to-one and let W be any subset of X . Then the largest f -invariant subset of W is equal to $W_f := \bigcap_{n=0}^{\infty} f^n(W_0)$, where $W_0 := \bigcap_{n=0}^{\infty} f^{-n}(W)$.*

Proof. We need to show that W_f contains any f -invariant subset of W and is f -invariant.

For the first claim, let $A \subset W$ and $f(A) = A$. We then have $A \subset f^{-n}(A) \subset f^{-n}(W)$ for all $n \geq 0$ so that $A \subset \bigcap_{n=0}^{\infty} f^{-n}(W) = W_0$. Therefore $A = \bigcap_{n=0}^{\infty} f^n(A) \subset \bigcap_{n=0}^{\infty} f^n(W_0) = W_f$.

For the second claim, note that $f(W_0) \subset \bigcap_{n=1}^{\infty} f(f^{-n}(W)) \subset \bigcap_{n=1}^{\infty} f^{-(n-1)}(W) = W_0$, i.e., W_0 is positively f -invariant. If we let $W_n := f^n(W_0)$ for $n \geq 0$, then, clearly, $W_{n+1} = f(W_n) \subset W_n$ so that $f(W_f) = f(\bigcap_{n=0}^{\infty} W_n) \subset \bigcap_{n=0}^{\infty} f(W_n) = \bigcap_{n=1}^{\infty} W_n = W_f$. Hence it remains to show that $W_f \subset f(W_f)$. This is the only place where we will use the assumption that f is finite-to-one. There is nothing to prove if W_f is empty. For any $x \in W_f$, let $Y_n := f^{-1}(x) \cap W_n$, $n \geq 0$. Then $(Y_n)_0^{\infty}$ forms a decreasing sequence of finite sets. Furthermore, each Y_n is non-empty since $x \in W_{n+1} = f(W_n)$. Hence $f^{-1}(x) \cap W_f = \bigcap_{n=0}^{\infty} Y_n \neq \emptyset$, i.e., $x \in f(W_f)$. \square

REFERENCES

- [AFL09] P. Ashwin, X. Fu and C. Lin. On planar piecewise and two-torus parabolic maps. *Internat. J. Bifur. Chaos* **19**(07) (2009), 2383–2390.
- [AFNZ00] P. Ashwin, X.-C. Fu, T. Nishikawa and K. Zyczkowski. Invariant sets for discontinuous parabolic area-preserving torus maps. *Nonlinearity* **13**(3) (2000), 819.
- [AKM⁺05] R. Adler, B. Kitchens, M. Martens, C. Pugh, M. Shub and C. Tresser. Convex dynamics and applications. *Ergod. Th. & Dynam. Sys.* **25**(4) (2005), 321–352.
- [ANS⁺17] R. Adler, T. Nowicki, G. Świrszcz, C. Tresser and S. Winograd. Error diffusion on acute simplices: invariant tiles. *Israel J. Math.* **221** (2017), 445–469.
- [ANST10] R. Adler, T. Nowicki, G. Świrszcz and C. Tresser. Convex dynamics with constant input. *Ergod. Th. & Dynam. Sys.* **30**(8) (2010), 957–972.
- [AY84] J. C. Alexander and J. A. Yorke. Fat baker's transformations. *Ergod. Th. & Dynam. Sys.* **4**(3) (1984), 1–23.
- [BBKT01] V. D. Blondel, O. Bournez, P. Koiran and J. N. Tsitsiklis. The stability of saturated linear dynamical systems is undecidable. *J. Comput. System Sci.* **62**(3) (2001), 442–462.
- [BCGT13] L. Brandolini, L. Colzani, G. Gigante and G. Travaglini. On the Koksma–Hlawka inequality. *J. Complexity* **29**(2) (2013), 158–172.
- [BG97] A. Boyarsky and P. Góra. *Laws of Chaos (Probability and its Applications)*. Birkhäuser Boston, Boston, MA, 1997.
- [Bro76] J. R. Brown. *Ergodic Theory and Topological Dynamics (Pure and Applied Mathematics, 70)*. Academic Press, New York, 1976.

- [BT03] H. Bruin and S. Troubetzkoy. The Gauss map on a class of interval translation mappings. *Israel J. Math.* **137**(1) (2003), 125–148.
- [Dan00] S. G. Dani. Dynamical systems on homogeneous spaces. *Dynamical Systems, Ergodic Theory and Applications (Encyclopaedia of Mathematical Sciences, 100)*, revised edn. Ed. Y. G. Sinai. Springer, Berlin, 2000, pp. 264–359.
- [DD03] I. Daubechies and R. DeVore. Approximating a bandlimited function using very coarsely quantized data: a family of stable sigma-delta modulators of arbitrary order. *Ann. of Math. (2)* **158**(2) (2003), 679–710.
- [DDBB⁺09] J. Daafouz, M. D. Di Benedetto, V. D. Blondel, G. Ferrari-Trecate, L. Hetel, M. Johansson, A. L. Juloski, S. Paoletti, G. Pola, E. De Santis and R. Vidal. Switched and piecewise affine systems. *Handbook of Hybrid Systems Control*. Cambridge University Press, Cambridge, 2009, pp. 87–137.
- [DDGV06] I. Daubechies, R. A. DeVore, C. Sinan Güntürk and V. A. Vaishampayan. A/D conversion with imperfect quantizers. *IEEE Trans. Inform. Theory* **52**(3) (2006), 874–885.
- [DGWY10] I. Daubechies, S. Güntürk, Y. Wang and Ö. Yılmaz. The golden ratio encoder. *IEEE Trans. Inform. Theory* **56**(10) (2010), 5097–5110.
- [DK02] K. Dajani and C. Kraaikamp. From greedy to lazy expansions and their driving dynamics. *Expo. Math.* **20**(4) (2002), 315–327.
- [Fed69] H. Federer. *Geometric Measure Theory (Die Grundlehren der mathematischen Wissenschaften, Band 153)*. Springer, New York, 1969.
- [GB89] P. Góra and A. Boyarsky. Absolutely continuous invariant measures for piecewise expanding C^2 transformation in \mathbf{R}^N . *Israel J. Math.* **67**(3) (1989), 272–286.
- [GT04] C. Sinan Güntürk and N. T. Thao. Refined error analysis in second-order $\Sigma\Delta$ modulation with constant inputs. *IEEE Trans. Inform. Theory* **50**(5) (2004), 839–860.
- [GT05] C. Sinan Güntürk and N. T. Thao. Ergodic dynamics in sigma-delta quantization: tiling invariant sets and spectral analysis of error. *Adv. Appl. Math.* **34**(3) (2005), 523–560.
- [Gün12] C. Sinan Güntürk. Mathematics of analog-to-digital conversion. *Comm. Pure Appl. Math.* **65**(12) (2012), 1671–1696.
- [Hla71] E. Hlawka. Discrepancy and Riemann integration. *Studies in Pure Mathematics (Presented to Richard Rado)*. Academic Press, London, 1971, pp. 121–129.
- [KN74] L. Kuipers and H. Niederreiter. *Uniform Distribution of Sequences (Pure and Applied Mathematics)*. Wiley-Interscience [John Wiley & Sons], New York, 1974.
- [LY73] A. Lasota and J. A. Yorke. On the existence of invariant measures for piecewise monotonic transformations. *Trans. Amer. Math. Soc.* **186**(1974) (1973), 481–488.
- [MT02] H. Masur and S. Tabachnikov. Rational billiards and flat structures. *Handbook Dyn. Sys.* **1** (2002), 1015–1089.
- [PH15] Krzysztof Plotka. On lineability and additivity of real functions with finite preimages. *J. Math. Anal. Appl.* **421**(2) (2015), 1396–1404.
- [Tao12] T. Tao. *Higher Order Fourier Analysis (Graduate Studies in Mathematics, 142)*. American Mathematical Society, Providence, RI, 2012.
- [Vee69] W. A. Veech. Strict ergodicity in zero dimensional dynamical systems and the Kronecker–Weyl theorem mod 2. *Trans. Amer. Math. Soc.* **140** (1969), 1–33.
- [Yıl02] Ö. Yılmaz. Stability analysis for several second-order sigma-delta methods of coarse quantization of bandlimited functions. *Constr. Approx.* **18**(4) (2002), 599–623.
- [ZL13] Y. Zhang and C. Lin. Invariant measures with bounded variation densities for piecewise area preserving maps. *Ergod. Th. & Dynam. Sys.* **33**(4) (2013), 624–642.