

Interval exchange transformations over algebraic number fields: the cubic Arnoux-Yoccoz model

J. H. Lowenstein[†], G. Poggiaspalla, and F. Vivaldi

[†]*Dept. of Physics, New York University, 2 Washington Place, New York, NY 10003, USA*
School of Mathematical Sciences, Queen Mary, University of London, London E1 4NS, UK

Abstract

We apply methods developed for two-dimensional piecewise isometries to the study of renormalizable interval exchange transformations over an algebraic number field $\mathbb{Q}(\lambda)$, which lead to dynamics on lattices. We consider the \mathbb{Z} -module \mathcal{M} generated by the translations of the map. On it, we define an infinite family of discrete vector fields, representing the action of the map over the cosets $\mathbb{Q}(\lambda)/\mathcal{M}$, which together form an invariant partition of the field $\mathbb{Q}(\lambda)$. We define a recursive symbolic dynamics, with the property that the eventually periodic sequences coincide with the field elements. We apply this approach to the study of a model introduced by Arnoux and Yoccoz, for which λ^{-1} is a cubic Pisot number. We show that all cosets of \mathcal{M} decompose in a highly non-trivial manner into the union of finitely many orbits.

September 15, 2006

1 Introduction

An important class of interval exchange transformations (IET) consists of those which possess a discrete invariance with respect to scaling. For such maps, the dynamical self-similarity, or renormalizability, can often be associated with the arithmetic properties of the scale factor. Such a connection has been established by Boshernitzan and Carroll: if λ is a quadratic irrational, then any IET defined over an algebraic number field $\mathbb{Q}(\lambda)$ can be renormalized, meaning that, up to rescaling, the number of maps induced on sub-intervals is finite [4]. At present, no generalization of this result to piecewise isometries in higher dimensions has been established. However, for two-dimensional polygonal exchange transformations defined over quadratic number fields, where the rotational part has a rational rotation number, the finiteness of the number of induced maps has been

established rigorously in many specific cases; and the same time, no counterexample has been found.

Going beyond quadratic irrational parameters, the relation between the presence of dynamical self-similarity and arithmetic properties of parameters appears to be much more tenuous. A noteworthy advance in this direction is the family of IETs of Arnoux and Yoccoz, which exhibit renormalizability with a scale factor ω which is the inverse of a Pisot number of degree n , for any integer n . A Pisot number is a real algebraic integer larger than 1, whose algebraic conjugates all lie strictly within the unit circle in the complex plane.

Our objectives in the current article are twofold. The first is to import into a one-dimensional context with interesting scaling behaviour (i.e., the cubic AY model) the recursive tiling formalism which has proved successful for a complete computer-assisted analysis of the orbits, both periodic and aperiodic, of certain two-dimensional piecewise isometries. Our second objective is to translate into the language of lattice maps some questions about the AY model relating dynamics to the arithmetical properties of the scale factor.

Before describing our main results, let us examine briefly the algebraic setting. An IET ϕ permutes the sub-intervals $\Omega_i = [b_i, b_{i+1})$ of an N -element partition of the unit interval $\Omega = [0, 1)$, preserving orientation. It acts as

$$\phi(x) = x + d_i, \quad x \in \Omega_i$$

where the discontinuity points b_i are assumed to be algebraic numbers; they generate a finite algebraic extension of the rationals

$$\mathbb{Q}(b_1, \dots, b_N) = \mathbb{Q}(\lambda)$$

for a suitable algebraic number λ satisfying a polynomial equation of degree n . Note that $d_i \in \mathbb{Q}(\lambda)$, since the translations can be expressed as linear integral combinations of the discontinuity points. Every point of $\mathbb{Q}(\lambda)$ can be written as

$$x = x_0 + x_1\lambda + \dots + x_{n-1}\lambda^{n-1} \quad x_k \in \mathbb{Q} \quad (1)$$

or, in vector notation,

$$x = \Lambda \cdot \mathbf{x}, \quad \mathbf{x} = (x_0, \dots, x_{n-1}), \quad \Lambda = (1, \lambda, \dots, \lambda^{n-1}). \quad (2)$$

Of importance to us will be the \mathbb{Z} -modules \mathcal{M} and \mathcal{N} formed by taking all integral linear combinations of d_1, \dots, d_N and b_1, \dots, b_N , respectively. The module \mathcal{M} will be called the *translation module* of the map. From here on we shall place a “roof” $\widehat{}$ over such sets to denote the restriction to Ω , which corresponds to taking the quotients of the given \mathbb{Z} -module with \mathbb{Z} . We note the inclusions

$$\mathcal{M} \subset \mathcal{N} \subset \mathbb{Q}(\lambda) \quad \Rightarrow \quad \widehat{\mathcal{M}} \subset \widehat{\mathcal{N}} \subset \widehat{\mathbb{Q}(\lambda)}.$$

Since the dynamics is generated by translations, the orbit of a point $\xi \in \widehat{\mathbb{Q}(\lambda)}$ is contained in the set $\widehat{\xi} + \widehat{\mathcal{M}}$, the restriction to Ω of a coset of the translation module \mathcal{M} . The




Figure 1: The interval exchange transformations ϕ and ρ on the interval $[0, 1)$. The scale-invariant map ρ is the first-return map on the scaling domain of ϕ , rescaled to the full unit interval (indicated by dashed lines).

collection of the restricted cosets constitutes a natural invariant partition of the discrete phase space $\widehat{\mathbb{Q}(\lambda)}$, which is itself a \mathbb{Z} -module. The specific choice of set Ξ of representatives of each coset turns out to be important; we shall define Ξ below (equation (18), section 3).

From (2), we see that each restricted coset $\widehat{\xi + \mathcal{M}}$ can be represented as the restriction to the slab $\Lambda \cdot x \in \Omega$ of an n -dimensional integer lattice. The restricted coset can then be represented as the union of finitely many $(n - 1)$ -dimensional lattices $\mathcal{L}_\xi = \cup_i \mathcal{L}_\xi^{(i)}$. Each point of this discrete set is assigned one of finitely many elementary translation vectors, which represent the original interval exchange dynamics. Together, these vectors define a ξ -dependent discrete vector field on \mathcal{L}_ξ . Orbits of the IET are then walks on lattices, with successive steps specified by the vector field.

This work is concerned with the cubic AY map ϕ (figure 1), which will be defined in section 2, equation (4). Our focus will be on the scale-invariant map ρ induced by ϕ on a sub-interval, which contains all relevant dynamical information. Its main feature is renormalizability by a Pisot number scale factor, which places strong constraints on the lattice orbits. Exploiting self-similarity, we shall construct a hierarchical code for all points of the interval Ω , which is a special case of a construct introduced by Vershik [15]. This code has the property that the multiplicative operation of rescaling is (essentially) represented by a right-shift in the symbol sequence (sections 2.2–2.3).

Let us denote by Γ the *discontinuity set* of the map, namely the set of all forward and backward images of the discontinuity points. The dynamical and module-theoretic significance of the hierarchical code are illustrated by the following result, proved in sections 2.3 and 3.4.

Theorem 1 *Let x be a point of Ω , and let $\sigma(x)$ be its hierarchical code. Then*

- (i) $x \in \widehat{\mathbb{Q}(\lambda)}$ if and only if $\sigma(x)$ is eventually periodic;

(ii) $x \in \Gamma$ if and only if $\sigma(x)$ is eventually fixed; furthermore

$$\Gamma = \widehat{\mathcal{N}} = \widehat{\mathcal{M}} \cup \widehat{\eta + \mathcal{M}} \quad \eta = \frac{1}{2} + \frac{\lambda^2}{2}.$$

A simplifying feature specific to the scale-invariant AY model derives from the fact that its translation module \mathcal{M} is the ring $\mathbb{Z}[\lambda]$, the set of all integral linear combinations of 1, λ , and λ^2 . Each point of $\widehat{\xi + \mathcal{M}}$ is of the form $\xi + m_0 + m_1\lambda + m_2\lambda^2$, with m_0 uniquely determined by the remaining coefficients. As a result, for each ξ , there is just one lattice $\mathcal{L}_\xi \simeq \mathbb{Z}^2$, equipped with a vector field consisting of seven elementary vectors. By studying the orbits on these lattices, we shall establish the following *finite decomposition property* :

Theorem 2 *For any $\xi \in \Xi$, the lattice \mathcal{L}_ξ decomposes into a finite number of orbits, each of which passes through at least one of M lattice points contained in a ξ -independent core region. As a consequence, the set $\widehat{\xi + \mathcal{M}}$ consists of at most M orbits of the IET ρ .*

This result is proved in section 3, based primarily on the Pisot character of the scale factor ω^{-1} . An explicit calculation will show that the core region visited by all orbits is contained in a square \mathcal{S} of side 15 centred at the origin, giving the bound $M \leq 15^2$. With additional input, we will sharpen this result below. From theorem 2, it is plain that the number of orbits in any subset of $\widehat{\mathbb{Q}(\lambda)}$ which is intersected by a finite number of cosets is bounded by the constant M multiplied by the number of cosets. A particularly simple case of orbit counting concerns the module $\Gamma = \widehat{\mathcal{N}}$. Computation reveals that it consists of just 6 distinct orbits, and each of these contains at least one of the seven discontinuity points b_i of the IET. Each component of $\widehat{\mathcal{N}}$ decomposes into the union of three orbits, which transit through the central square \mathcal{S} (see figures 3 and 6).

To describe the topology of the orbits on the lattices \mathcal{L}_ξ , we introduce the concept of *neighbours*, which are lattice points whose euclidean distance is smaller than 2. The main results concerning the arrangement of orbits and their boundaries, distilled from the detailed treatment of section 4, are summarized in the following theorem

Theorem 3 *Every orbit \mathcal{O} on the lattice \mathcal{L}_ξ fills a simply connected region; the polygonal line obtained by connecting successive points of the orbit does not cross itself transversally. If $\mathcal{O} \neq \mathcal{L}_\xi$, then the boundary of \mathcal{O} consists of most two boundary curves, which are polygonal lines whose vertices are lattice neighbours. The boundary curves of all the orbits are mutually non-intersecting and self-avoiding, and they all visit the boundary of \mathcal{S} .*

This result combines lemmas 14 and 15, together with some statements appearing at the beginning of section 4.4. As a consequence, the total number of orbits in any \mathcal{L}_ξ cannot exceed the perimeter of the central square \mathcal{S} , and hence the uniform bound M in theorem 2 can be taken to be 56. This improved bound turns out to be still far from optimal. A systematic numerical search of millions of cosets (section 4.6) reveals that all but an exceedingly small fraction of the lattices \mathcal{L}_ξ contain just a single orbit, and there

are no cases of more than three orbits covering the lattice. Moreover, there is no evident pattern governing the occurrence of exceptional cases.

Theorem 3 explains the salient features of explicit orbit plots such as figures 3 and 6. Other properties of the orbits will be discussed in section 4, including their self-similar spiralling outward, and the structure of their boundary curves. To the eye, the boundaries of the orbits appear to be approaching under renormalization a curve with fractal dimension just above unity, which is reminiscent of the Rauzy fractal boundary [13]. Although we have not established a precise connection between the Rauzy boundary and that of our lattice orbits, this similarity is hardly surprising, due to the existence of a semi-conjugacy between the AY map and the two-dimensional map which generates the Rauzy fractal [12]. Additional connections exist between the Rauzy fractal and Pisot units [1]. What we did find experimentally, is a candidate exact substitution rule for the lattice orbits involving 42 symbols, which, when iterated, does produce a scaling dimension equal to the Hausdorff dimension of the Rauzy boundary [7, 11].

Extensions of these results to other examples of algebraic IET may be possible, though not straightforward. It can be shown that, generically, an algebraic interval exchange map does not have the finite decomposition property described in this paper. Indeed it turns out that only the IETs with a zero *Sah-Arnoux-Fathi invariant* (cf. [14]) can feature it. Few such IETs are known. These include the Arnoux-Yoccoz family (of degree n) described here, as well as a notable family of degree 3 proposed by Arnoux and Rauzy [2]. We have conducted preliminary numerical computations which show that at least some members of the Arnoux-Rauzy family have the desired behaviour on lattices.

The careful comments of one referee helped us to improve the presentation and the bibliography, and to correct some imprecisions. This research was supported by EPSRC grant No GR/S62802/01.

2 The Arnoux-Yoccoz map

A family of renormalizable IETs, defined over algebraic number fields of arbitrarily large degree, was discovered by Arnoux and Yoccoz, during their search for pseudo-Anosov diffeomorphisms [3]. The fields are generated by the unique positive root of the polynomials

$$f_n(x) = x^n + x^{n-1} + \dots + x - 1, \quad n \geq 2. \quad (3)$$

The polynomials $f_n(x)$ are irreducible for all n [3, lemma 3], and λ^{-1} is a Pisot number. Because $\lambda < 1$, the unit interval can be partitioned into n intervals of decreasing length $\lambda, \lambda^2, \dots, \lambda^n$. Explicitly

$$\Omega = \bigcup_{k=1}^n \Omega_k \quad \Omega_k = [w_{k-1}, w_k)$$

where

$$w_0 = 0 \quad w_k = \sum_{j=1}^k \lambda^j, \quad k > 0.$$

Let E_s be the map which interchanges the left and right halves of the interval $s = [a, b]$ while preserving orientation, i.e.,

$$E_s(x) = \begin{cases} x + (a+b)/2 & x \in [a, (a+b)/2) \\ x - (a+b)/2 & x \in [(a+b)/2, b) \\ x & x \notin [a, b). \end{cases}$$

The Arnoux-Yoccoz (AY) map is then defined as

$$\phi = E_\Omega \circ E_{\Omega_1} \circ \cdots \circ E_{\Omega_n}, \quad (4)$$

which is an interval exchange transformation on $2n + 1$ intervals. The authors show that ϕ is conjugate to the map which it induces on the sub-interval Ω_1 .

For $n = 2$, the map features two invariant intervals of length $1/2$. Each of them supports a 2-interval exchange, which is a rotation by $\lambda/2$, half the reciprocal of the golden mean. Such a rotation is well-known; in particular, the map induced in the first interval is conjugate to the original map via a contraction with ratio $\omega = \lambda^2/4$. Such self-similar dynamics is conjugate to the shift map on the dynamical system given by the so-called Fibonacci substitution

$$\begin{aligned} \sigma : 0 &\longrightarrow 01 \\ &1 \longrightarrow 1 \end{aligned}$$

whose incidence matrix has λ as eigenvalue. (For an introduction to the theory of substitutions, see [12].)

In the cubic case ($n = 3$), the number λ has very specific properties, studied in the classic work of Rauzy [13]. These are related to the substitution

$$\begin{aligned} \sigma : 1 &\longrightarrow 12 \\ &2 \longrightarrow 13 \\ &3 \longrightarrow 1 \end{aligned} \quad (5)$$

by the Rauzy theorem: let (u_n) be the fixed point of the above substitution. Then there exist three disjoint open sets $\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_3$ of the two-torus whose union is dense and such that $(k\lambda, k\lambda^2) \in \mathcal{R}_{u_k}$. The sets \mathcal{R}_i are images of each other by linear transformations: they are known as the *Rauzy fractals*. Moreover, the itineraries (in the sense of the above intervals $\Omega_0, \Omega_1, \Omega_2$) of the orbits are given by the substitution (5). Conjugacy with a substitution dynamical system ensures unique ergodicity.

2.1 Domain map and scaling sequences

We now reformulate the AY map using the terminology and formalism used in [8, 10] to analyze the recursive return-map structure of a class of two-dimensional piecewise isometries.

The first step is to write the AY map ϕ as a *domain map* on the interval $\Omega = [0, 1)$. The action of the map is to translate each of the sub-domains (half-open intervals) Ω_i , $i = 0, 1, \dots, 6$ as prescribed in table 1.

Interval	Map	Translation
Ω_0	$[0, \frac{1}{2} - \frac{\lambda}{2}) \mapsto [\frac{1}{2} + \frac{\lambda}{2}, 1)$	$\frac{1}{2} + \frac{\lambda}{2}$
Ω_1	$[\frac{1}{2} - \frac{\lambda}{2}, \frac{\lambda}{2}) \mapsto [0, -\frac{1}{2} + \lambda)$	$-\frac{1}{2} + \frac{\lambda}{2}$
Ω_2	$[\frac{\lambda}{2}, \lambda) \mapsto [\frac{1}{2}, \frac{1}{2} + \frac{\lambda}{2})$	$\frac{1}{2} - \frac{\lambda}{2}$
Ω_3	$[\lambda, \lambda + \frac{\lambda^2}{2}) \mapsto [-\frac{1}{2} + \lambda + \frac{\lambda^2}{2}, -\frac{1}{2} + \lambda + \lambda^2)$	$-\frac{1}{2} + \frac{\lambda^2}{2}$
Ω_4	$[\lambda + \frac{\lambda^2}{2}, \lambda + \lambda^2) \mapsto [-\frac{1}{2} + \lambda, -\frac{1}{2} + \lambda + \frac{\lambda^2}{2})$	$-\frac{1}{2} - \frac{\lambda^2}{2}$
Ω_5	$[\lambda + \lambda^2, \frac{1}{2} + \frac{\lambda}{2} + \frac{\lambda^2}{2}) \mapsto [\frac{\lambda}{2} + \frac{\lambda^2}{2}, \frac{1}{2})$	$-\frac{\lambda}{2} - \frac{\lambda^2}{2}$
Ω_6	$[\frac{1}{2} + \frac{\lambda}{2} + \frac{\lambda^2}{2}, 1) \mapsto [-\frac{1}{2} + \lambda + \lambda^2, \frac{\lambda}{2} + \frac{\lambda^2}{2})$	$-1 + \frac{\lambda}{2} + \frac{\lambda^2}{2}$

Table 1: Arnoux-Yoccoz interval exchange map (cubic case)

Next, we identify an infinite *scaling sequence* of nested sub-domains of Ω , namely

$$\tilde{\mathcal{D}}(L) = [0, \omega^L \alpha), \quad \omega = \lambda^3, \quad \alpha = \frac{\lambda^2}{2}(1 - \lambda), \quad L = 1, 2, \dots$$

whose induced first-return maps $\tilde{\rho}(L)$ are related by the operation of scaling by ω :

$$\tilde{\rho}(L) = \omega \tilde{\rho}(L - 1) \omega^{-1}.$$

Since we will be dealing below mainly with the scaling sequence and its first-return maps, it will be convenient to rescale these to the unit interval, introducing

$$\mathcal{D}(L) = \alpha^{-1} \tilde{\mathcal{D}}(L) = [0, \omega^L), \quad \rho(L) = \alpha^{-1} \tilde{\rho}(L) \alpha.$$

The data for the domain map

$$\rho(0) : \mathcal{D}(0) \rightarrow \mathcal{D}(0) \quad \mathcal{D}(0) = \Omega = \bigcup_j \mathcal{D}_j(0)$$

are displayed in table 2. For economy of notation, we set $\rho \stackrel{\text{def}}{=} \rho(0)$, since the scale-invariant map ρ , will be the focus of our attention for most of the remainder of this article.

The level- L return map $\rho(L)$ is related to $\rho(L - 1)$ by

$$\rho(L) = \omega \rho(L - 1) \omega^{-1}. \quad (6)$$

If we write $\rho_j(L)$ for ρ restricted to $\mathcal{D}_j(L)$, then

$$\rho(L) \mathcal{D}_j(L) = \rho_{p(j, \nu_j - 1)}(L - 1) \circ \dots \circ \rho_{p(j, 0)}(L - 1) \mathcal{D}_j(L) \subset \mathcal{D}(L),$$

where

$$p(j) \stackrel{\text{def}}{=} (p(j, 0), \dots, p(j, \nu_j - 1)), \quad p(j, t) \in \{0, 1, \dots, 6\},$$

is the *return path* for $\mathcal{D}_j(L)$, for all $L > 0$, while $p(j, t)$ will be called the *path function*. That is, under the action of $\rho(L - 1)$ the return orbit of $\mathcal{D}_j(L)$ visits in order the level- $(L - 1)$ domains $\mathcal{D}_{p(j, t)}(L - 1)$, $t = 0, \dots, \nu_j - 1$, before returning to $\mathcal{D}(L)$. Furthermore, since $\mathcal{D}(L) = \mathcal{D}_0(L - 1)$, for $L \geq 1$, we have that $p(j, t) = 0$ if and only if $t = 0$.

Domain	Map ρ	Translation	Vector field
$\mathcal{D}_0(0)$	$[0, 1 - \lambda - \lambda^2) \mapsto [\lambda + \lambda^2, 1)$	$\lambda + \lambda^2$	$(1, 1)$
$\mathcal{D}_1(0)$	$[1 - \lambda - \lambda^2, 1 - 2\lambda + \lambda^2) \mapsto [2\lambda - \lambda^2, \lambda + \lambda^2)$	$-1 + 3\lambda$	$(3, 0)$
$\mathcal{D}_2(0)$	$[1 - 2\lambda + \lambda^2, \frac{3}{2} - 2\lambda - \frac{\lambda^2}{2}) \mapsto [1 - \lambda, \frac{3}{2} - \lambda - \frac{3\lambda^2}{2})$	$\lambda - \lambda^2$	$(1, -1)$
$\mathcal{D}_3(0)$	$[\frac{3}{2} - 2\lambda - \frac{\lambda^2}{2}, \frac{1}{2} - \frac{\lambda^2}{2}) \mapsto [\frac{1}{2} + \frac{\lambda^2}{2}, -\frac{1}{2} + 2\lambda + \frac{\lambda^2}{2})$	$-1 + 2\lambda + \lambda^2$	$(2, 1)$
$\mathcal{D}_4(0)$	$[\frac{1}{2} - \frac{\lambda^2}{2}, -\frac{1}{2} + \lambda + \frac{3\lambda^2}{2}) \mapsto [\frac{3}{2} - \lambda - \frac{3\lambda^2}{2}, \frac{1}{2} + \frac{\lambda^2}{2})$	$1 - \lambda - \lambda^2$	$(-1, -1)$
$\mathcal{D}_5(0)$	$[-\frac{1}{2} + \lambda + \frac{3\lambda^2}{2}, \lambda) \mapsto [-\frac{1}{2} + 2\lambda + \frac{\lambda^2}{2}, 2\lambda - \lambda^2)$	$\lambda - \lambda^2$	$(1, -1)$
$\mathcal{D}_6(0)$	$[\lambda, 1) \mapsto [0, 1 - \lambda)$	$-\lambda$	$(-1, 0)$

Table 2: The scale-invariant return map $\rho = \rho(0)$. The final column refers to the lattice map of section 3.

The return paths and times for $\rho(L)$, as well as the return times $\nu_{0,j}$ for level-zero sub-domains (with respect to the original map ϕ) are shown in table 3.

From table 3 we can read off the number of times that the first-return orbit of $\mathcal{D}_j(L)$ visits the level- $(L-1)$ domain $\mathcal{D}_i(L-1)$, i.e.,

$$A_{j,i} = \sum_{t=0}^{\nu_j-1} \delta_{i,p(j,t)},$$

where $\delta_{i,k}$ is the Kronecker delta. These constitute the *incidence matrix*

$$A = \begin{pmatrix} 1 & 0 & 0 & 1 & 0 & 0 & 2 \\ 1 & 2 & 1 & 2 & 0 & 1 & 6 \\ 1 & 1 & 2 & 1 & 1 & 1 & 5 \\ 1 & 1 & 0 & 2 & 0 & 0 & 4 \\ 1 & 0 & 1 & 0 & 2 & 1 & 3 \\ 1 & 1 & 1 & 1 & 1 & 2 & 5 \\ 1 & 0 & 0 & 0 & 1 & 0 & 2 \end{pmatrix}$$

with characteristic polynomial $(x-1)(x^3-5x^2+7x-1)(x^3-7x^2+5x-1)$, whose real roots are 1, $\lambda^3 = \omega$, and λ^{-3} , respectively. In terms of iterations of ρ , the level- L first-return

j	$\nu_{0,j}$	ν_j	$p(j)$
0	6	4	$(0, 6, 3, 6)$
1	37	13	$(0, 6, 3, 6, 1, 6, 2, 5, 6, 1, 6, 3, 6)$
2	37	12	$(0, 6, 3, 6, 1, 6, 2, 5, 6, 2, 4, 6)$
3	17	8	$(0, 6, 3, 6, 1, 6, 3, 6)$
4	24	8	$(0, 6, 4, 5, 6, 2, 4, 6)$
5	37	12	$(0, 6, 4, 5, 6, 2, 5, 6, 1, 6, 3, 6)$
6	7	4	$(0, 6, 4, 6)$

Table 3: Level-to-level return times and paths. Here ν_j is the length of the return path $p(j)$ for $\mathcal{D}_j(L)$, $L > 0$, while $\nu_{0,j}$ is the length of the return path for $\tilde{\mathcal{D}}_j(0)$ with respect to iterations of the original map ϕ .

times may be calculated as

$$T_j(L) = \sum_{k=0}^6 A_{j,k}^L \quad L = 1, 2, \dots,$$

with $T_j(1) = \nu_j$. In terms of iterations of the original AY map, they are

$$T_{0,j}(L) = \sum_{k=0}^6 A_{j,k}^L \nu_{0,k} \quad L = 1, 2, \dots$$

The asymptotic (large- L) behaviour of $T_j(L)$ is dominated by the L th power of the largest eigenvalue of A , $\lambda^{-3} = 6.22226\dots$

Later in this article we will want to exploit the fact that, in the allowed paths listed in table 3, the possible two- and three-step sequences are highly constrained. For example, a step of type 0 can only be preceded and followed by a step of type 6. The full set of such constraints is given in table 4.

j	Predecessors	Predecessor pairs	Successors	Successor pairs
0	6	(3, 6), (4, 6)	6	(6, 3), (6, 4)
1	6	(3, 6), (5, 6)	6	(6, 2), (6, 3)
2	6	(1, 6), (5, 6)	4, 5	(4, 6), (5, 6)
3	6	(0, 6), (1, 6)	6	(6, 0), (6, 1)
4	2, 6	(0, 6), (6, 2)	5, 6	(5, 6), (6, 0)
5	2, 4	(6, 2), (6, 4)	6	(6, 1), (6, 2)
6	0, 1, 3, 4, 5	(4, 5), (2, 4), (2, 5), (6, 0), (6, 1), (6, 3), (6, 4)	0, 1, 2, 3, 4	(0, 6), (1, 6), (2, 4), (2, 5), (3, 6), (4, 5), (4, 6)

Table 4: List of possible predecessors and successors, considered singly or in pairs.

2.2 Recursive tiling, and Vershik symbolic dynamics

It is straightforward to verify that $\mathcal{D}(0) = \Omega$ is completely tiled by the images under $\rho = \rho(0)$ of the sub-domains $\mathcal{D}_j(1)$ of $\mathcal{D}(1)$. The total number of tiles is $\sum_j \nu_j = 61$, namely $\nu_0 = 4$ translates of $\mathcal{D}_0(1)$, $\nu_1 = 13$ translates of $\mathcal{D}_1(1)$, etc. Thanks to the scaling relations, the same partitioning, rescaled by a factor ω^L , occurs for $\mathcal{D}(L)$, and is propagated to all of $\mathcal{D}(0)$ by repeated application of ρ . Thus, at level L , we have a complete tiling of Ω by translates of the seven sub-domains $\mathcal{D}_j(L+1)$, $j = 0, 1, \dots, 6$. For an arbitrary tile, we introduce the notation

$$\mathcal{D}_{jL}^{(t_1, \dots, t_L)}(L) \stackrel{\text{def}}{=} \rho(0)^{t_1} \circ \dots \circ \rho(L-1)^{t_L} \mathcal{D}_{jL}(L),$$

and associate a symbolic code

$$((j_1, t_1), (j_2, t_2), \dots, (j_L, t_L)),$$

where all but the last of the j_k are determined by j_L and the path constraints

$$j_k = p(j_{k+1}, t_{k+1}), \quad k = 1, \dots, L-1.$$

The tiles satisfy the following inclusion relations

$$\mathcal{D}_{j_L}^{(t_1, \dots, t_L)}(L) \subset \mathcal{D}_{j_{L-1}}^{(t_1, \dots, t_{L-1})}(L-1) \subset \dots \subset \mathcal{D}_{j_1}^{t_1}(1) \subset D_{j_0}(0),$$

where $j_0 = p(j_1, t_1)$. Note that j_0 is not written explicitly in the code.

Every point $x \in \Omega$ lies in the intersection of a unique infinite nested sequence of tiles, and hence can be labelled by an infinite code sequence $\sigma(x) = ((j_1, t_1), (j_2, t_2), \dots)$, with the j_k constrained by the path conditions. Note that there may exist a properly linked code sequence corresponding to a nested set of tiles which do not contain x , but which have x as a limit point. These spurious code sequences are excluded from consideration, since for such a point x there exists another code sequence corresponding to nested tiles containing x —cf. proof of proposition 7. (This situation is analogous to the non-uniqueness of the decimal representation of certain rational numbers.)

Suppose x corresponds to code $\sigma(x) = ((j_1, t_1), (j_2, t_2), \dots, (j_L, t_L))$. What is $\sigma(\rho(x))$? Since the first code symbol represents ρ^{t_1} acting on a subset of $\mathcal{D}_{j_1}(1)$, acting once more with ρ simply increments t_1 , unless t_1 is at its maximal value $\nu_{j_1} - 1$. In the latter case, an additional ρ completes the first-return orbit for $\rho(1)$, and so it is t_2 which must be incremented, while we reset t_1 to zero and j_1 to the value dictated by the path constraint. Of course, t_2 (and possibly t_3, t_4 , etc.) may also be maximal, and so the incrementing may take place further to the right in the code sequence. The full prescription, a special case of the scheme introduced by Vershik [15], is the following:

If $t_i < \nu_{j_i} - 1$ and i is the smallest index with this property, then

$$((j_1, t_1), (j_2, t_2), \dots) \mapsto ((j'_1, 0), \dots, (j'_{i-1}, 0), (j_i, t_i + 1), (j_{i+1}, t_{i+1}), \dots),$$

where the j'_k are determined by the path constraints. In the one exceptional case where all t_k are maximal, namely $\sigma(\lambda) = (6, 3)^\infty$, the successor code corresponding to $0 = \rho(\lambda)$ is $\sigma(0) = (0, 0)^\infty$.

We note that all points on the forward orbit of 0 have a code with infinite tail $(0, 0)^\infty$, while those on the backward orbit have tail $(6, 3)^\infty$. All other orbits are characterized by codes with unique tails, i.e., if x and y are on the same orbit, there exists an n such that the symbols of $\sigma(x)$ and $\sigma(y)$ agree from position $n + 1$ onward:

$$\sigma(x) = \dots (j_n, t_n), (j_{n+1}, t_{n+1}), \dots, \quad \sigma(y) = \dots (j'_n, t'_n), (j_{n+1}, t_{n+1}), \dots \quad (7)$$

Here $j'_n = j_n = p(j_{n+1}, t_{n+1})$ and the relative positions of x and y on the orbit is determined by the sign of $t_n - t'_n$. As required by minimality, the Vershik updating forbids the existence of periodic orbits, since every symbolic code is distinct from all of its predecessors on an orbit, and the correspondence of legal codes with points of Ω is one-to-one.

We summarize the above considerations in the following

Lemma 4 *With the exception of the orbit through 0, two points belong to the same orbit if and only if their codes eventually agree, in the sense of equation (7).*

Note that the updating scheme is more complicated than an odometer, even one whose registers are not identical. In the Vershik scheme, the identity of the k th register (not just its reading) depends on the identity and reading of the $(k + 1)$ st register.

We remark that the coding presented here is equivalent to the graph-directed representation of the underlying substitution of the interval-exchange, described by Holton and Zamboni in [6]. Namely, if G is this graph, then the coding described above consists of all the possible paths in a graph isomorphic to the *derived graph* $D(G)$ of G .¹ We also remark that such a graph is equivalent to the prefix-suffix automaton of the substitution as described by Canterini and Siegel in [5]. In this reference, a Vershik-like dynamics on the automaton is also described.

2.3 Scaling operators as symbolic shifts

Multiplication by the scale factor ω and its inverse are represented symbolically by modified right and left symbol shifts, respectively.

Lemma 5 *If the code of $x \in \Omega$ is*

$$\sigma(x) = ((j_1, t_1), (j_2, t_2), \dots), \quad (8)$$

then

$$\sigma(\omega x) = ((p(j_1, t_1), 0), (j_1, t_1), (j_2, t_2), \dots), \quad (9)$$

and, if $t_1 = 0$,

$$\sigma(\omega^{-1}x) = ((j_2, t_2), (j_3, t_3), \dots). \quad (10)$$

Proof. If x has code $\sigma(x)$ given by (8), and n is any positive integer, we have

$$x \in \rho(0)^{t_1} \rho(1)^{t_2} \dots \rho(n-1)^{t_n} \mathcal{D}_{j_n}(n),$$

so that

$$\begin{aligned} \omega x &\in (\omega \rho(0)^{t_1} \omega^{-1}) (\omega \rho(1)^{t_2} \omega^{-1}) \dots (\omega \rho(n-1)^{t_n} \omega^{-1}) (\omega \mathcal{D}_{j_n}(n)) \\ &= \rho(1)^{t_1} \rho(2)^{t_2} \dots \rho(n)^{t_n} \mathcal{D}_{j_n}(n+1) \end{aligned}$$

where we have used (6) Similarly,

$$\begin{aligned} \omega^{-1}x &\in (\omega^{-1} \rho(1)^{t_1} \omega) (\omega^{-1} \rho(2)^{t_2} \omega) \dots (\omega^{-1} \rho(n-1)^{t_n} \omega) (\omega^{-1} \mathcal{D}_{j_n}(n-1)) \\ &= \rho(2)^{t_1} \rho(3)^{t_2} \dots \rho(n)^{t_n} \mathcal{D}_{j_n}(n-1). \end{aligned}$$

¹If G is a directed graph G with vertex set V and edge set E , the derived graph $D(G)$ has vertex set E , and there is an edge from $e \in E$ to $f \in E$ if and only if e and f are consecutive in G , i.e., if the target vertex of e is the source of f in G .

Returning to the symbolic representation and letting n tend to infinity, we get the results. \square

Let x have code (8). The restriction on t_1 can be removed to give a true left shift by introducing the function

$$\tau(x) = t_1 = \min\{t \geq 0 : \rho^{-t}x \in \mathcal{D}_0(0)\} \quad (11)$$

and then defining

$$\gamma(x) = \omega^{-1}\rho^{-\tau(x)}(x).$$

The operator $\rho^{-\tau}$ resets t_1 to 0, and hence by (10),

$$\sigma(\gamma(x)) = ((j_2, t_2), (j_3, t_3), \dots). \quad (12)$$

If repeated application of γ is accompanied by recording j_1 and t_1 at each step, we obtain a convenient encoding algorithm. Here we make use of the coding function $J(x)$ defined by

$$J(x) = j \Leftrightarrow x \in \mathcal{D}_j(0). \quad (13)$$

Algorithm. *Input:* $x \in \Omega$, $n \in \mathbb{N}$; *Output:* the first n symbols of $\sigma(x)$.

- (i) Set $k \leftarrow 0$; $\sigma \leftarrow \emptyset$.
- (ii) Set $k \leftarrow k + 1$.
- (iii) Set $t_k \leftarrow \tau(x)$.
- (iv) Set $x \leftarrow \gamma(x)$.
- (v) Set $j_k = J(x)$; append (j_k, t_k) to σ .
- (vi) If $k = n$, return σ and stop; otherwise go to step 2.

2.4 Eventually periodic codes

Any point $x \in \Omega$ with an eventually periodic code

$$\sigma(x) = ((j'_1, t'_1), \dots, (j'_m, t'_m), ((j_1, t_1), \dots, (j_n, t_n))^\infty)$$

can be computed explicitly from the code. In particular,

$$x = \rho(0)^{t'_1} \dots \rho(m-1)^{t'_m} \omega^m y,$$

where $y \in \mathcal{D}_{j'_m}(0)$ has strictly periodic code

$$\sigma(y) = ((j_1, t_1), \dots, (j_n, t_n))^\infty.$$

Defining

$$\delta = \sum_{\substack{i=1 \\ t_i > 0}}^n \sum_{s_i=0}^{t_i-1} d_{p(j_i, s_i)}(i-1), \quad (14)$$

where $d_j(k)$ is the level- k translation associated with the domain $\mathcal{D}_j(k)$, we obtain

$$y = \rho(0)^{t_1} \rho(1)^{t_2} \cdots \rho(n-1)^{t_n} \omega^n y = \omega^n y + \delta$$

and hence

$$y = (1 - \omega^n)^{-1} \delta. \quad (15)$$

We note that δ and ω are both in $\mathbb{Z}[\lambda]$, and since $x \in \Omega$

$$\sigma(x) \text{ eventually periodic} \quad \Rightarrow \quad x \in \widehat{\mathbb{Q}(\lambda)}. \quad (16)$$

Later we will prove the converse of this statement to complete the proof of theorem 1.

An important property of $x \in \Omega$ with eventually periodic $\sigma(x)$ is expressed by the following proposition.

Proposition 6 *If $\sigma(x)$ is eventually periodic with period n , then the points $\omega^{Nn}x$, $N = 0, 1, 2, \dots$ lie on the same orbit. As $N \rightarrow \infty$, the integer T such that $\omega^{(N+1)n}x = \rho^T(\omega^{Nn}x)$, admits the estimate $T = O(\omega^{-Nn})$.*

Proof. For $x = 0$, the result is trivial; we assume $x \neq 0$. Since multiplication by ω corresponds to a right symbol shift (with appropriate padding), the codes for x and $\omega^n x$ coincide from some position $m+1$ onward. From lemma 4 it then follows that x and $\omega^n x$ belong to the same orbit. Suppose now that $\omega^{Nn}x$ lies on the forward orbit of $\omega^{(N+1)n}x$, i.e.,

$$\omega^{Nn}x = \rho(Nn)^{t_1} \cdots \rho(Nn+n-1)^{t_n} \omega^{(N+1)n}x.$$

Then the number of ρ iterations needed to connect the points is

$$T = \sum_{i=0}^{n-1} \sum_{s_i=0}^{t_i-1} T_{p(j_i, s_i)}(Nn+i) = \sum_{i=0}^{n-1} \sum_{s_i=0}^{t_i-1} \sum_{k=0}^6 A_{p(j_i, s_i), k}^{Nn+i}.$$

The sum is dominated for $N \rightarrow \infty$ by terms proportional to the largest eigenvalue of the positive matrix A^{Nn} , i.e. $T = O(\omega^{-Nn})$. The backward-orbit case, with T large and negative, is handled similarly. \square

The simplest codes belong to the seven discontinuity points of the map ρ , as we see in the following proposition.

Proposition 7 *A point in Ω is a discontinuity point for ρ if and only if its code is periodic with period 1.*

Proof. It is not difficult to verify that the seven discontinuity points (including 0) for ρ are fixed points of the mapping γ with corresponding period-1 codes

$$(0,0)^\infty, \quad (1,4)^\infty, \quad (2,9)^\infty, \quad (3,6)^\infty, \quad (4,2)^\infty, \quad (5,6)^\infty, \quad (6,3)^\infty.$$

There are exactly six additional period-1 codes compatible with the return-path data of table 3, but one can check that they are inadmissible: they correspond to domain sequences with empty intersections. \square

We note that the inadmissible sequences converge to six discontinuity points from the left. The right endpoint of $D_0(0)$ is exceptional, in that the corresponding inadmissible sequence is eventually periodic rather than periodic, namely $(6,0)(6,1)^\infty$.

3 Representations of the scale-invariant AY map

The restriction of an IET to an algebraic number field admits several representations, each of which have certain merits. In what follows we specialize to the case of the scale-invariant cubic AY map; however, some of the constructs apply to any cubic IETs with Pisot scaling, as we shall indicate.

Letting $n = 3$ in equations (1) and (2), we find that each point

$$x = x_0 + x_1\lambda + x_2\lambda^2 \in \widehat{\mathbb{Q}(\lambda)}$$

corresponds to a rational 3-vector

$$\mathbf{x} = (x_0, x_1, x_2) \in U$$

where U is the “rational unit slab”

$$U = \{\mathbf{x} \in \mathbb{Q}^3 : \Lambda \cdot \mathbf{x} \in \Omega\}, \quad \Lambda = (1, \lambda, \lambda^2). \quad (17)$$

We recall further that the map ρ has translation module $\mathcal{M} = \mathbb{Z}[\lambda]$, so that $\mathbb{Q}(\lambda)$ decomposes into cosets $\xi + \mathcal{M}$, with $\xi \in \Xi$, where Ξ is a suitable set of representatives (one from each coset). Our choice is

$$\Xi = \Lambda \cdot \widehat{\mathbb{Q}}^3, \quad (18)$$

where $\widehat{\mathbb{Q}}^3$ is the rational unit cube. We note that each point in $\widehat{\xi + \mathcal{M}}$ may be written as

$$\xi + \Lambda \cdot \mathbf{m}, \quad \mathbf{m} = (m_0, m_1, m_2) \in \mathbb{Z}^3,$$

with m_0 determined by m_1, m_2 via

$$m_0 = -\lfloor \xi + m_1\lambda + m_2\lambda^2 \rfloor,$$

and so the points of $\widehat{\xi + \mathcal{M}}$ correspond to points (m_1, m_2) of an integer lattice $\mathcal{L}_\xi \cong \mathbb{Z}^2$. The above relationships are represented in the commutative diagram

$$\begin{array}{ccc}
 \widehat{\mathbb{Q}(\lambda)} = \widehat{\xi + \mathcal{M}} & \xrightarrow{\varphi_1} & U \\
 \searrow \varphi_3 & & \swarrow \varphi_2 \\
 & & \widehat{\mathbb{Q}^3} \times \mathbb{Z}^2
 \end{array} \tag{19}$$

Here the maps are all bijections.

In the following sub-sections, we discuss the representations of the dynamical map ρ , the scaling map ω , and the shift map γ , in each of the listed contexts. The scaling map corresponds to multiplication by $\omega = \lambda^3$, and we shall use the same symbol for the map and the scale factor. The results will not depend on the specific values of the parameter λ or the scale factor ω ; what is important is that λ and ω are real cubic algebraic integers, and that ω is a *unit* (an invertible element) in $\mathbb{Z}[\lambda]$, with the property that $|\omega| < 1$, while its algebraic conjugates are complex and lie outside the unit circle. Thus ω is the reciprocal of a Pisot number.

Formally, reduction to Ω corresponds to taking the quotient by \mathbb{Z} , and therefore we have $\xi + \mathcal{M} \in \mathbb{Q}(\lambda)/\mathcal{M}$ and $\widehat{\xi + \mathcal{M}} \in (\mathbb{Q}(\lambda)/\mathcal{M})/\mathbb{Z}$, which are quotients of \mathbb{Z} -modules. For the AY map we have $\mathcal{M} \supset \mathbb{Z}$, which accounts for the \mathbb{Z} -module isomorphism $\mathcal{M} \simeq \mathbb{Z}^2$, and hence the representation at the bottom of the diagram (19). The general situation is a bit more complicated: it turns out that, for a full-rank module \mathcal{M} ,

$$\mathcal{M} \cong (\mathbb{Z}/b\mathbb{Z}) \oplus \mathbb{Z}^{n-1} \tag{20}$$

where n is the degree of λ , and b is a certain integer which depends on the arithmetical properties of \mathcal{M} . We shall not pursue this matter here.

In order to avoid notational clutter, we have adopted the following conventions. The corresponding elements of the various sets are labelled by the same letter, but with different font styles: italic for algebraic numbers in $\widehat{\mathbb{Q}(\lambda)}$, sans serif for 3-vectors in U , and boldface for 2-vectors in \mathcal{L}_ξ . For example,

$$x \in \widehat{\xi + \mathcal{M}} \cong \mathbf{x} \in U \cong \mathbf{x} \in \mathcal{L}_\xi.$$

Furthermore, the maps ρ, ω , and γ are denoted by the same letters (font style unchanged) in all three contexts.²

3.1 One-dimensional representation: coset partition

We represent the map in $\widehat{\mathbb{Q}(\lambda)}$, the restriction of the field $\mathbb{Q}(\lambda)$ to the unit interval Ω .

²With apologies to Gertrude Stein, a ρ is a ρ is a ρ .

The *dynamical map* ρ is defined as

$$\rho : \widehat{\mathbb{Q}(\lambda)} \rightarrow \widehat{\mathbb{Q}(\lambda)} \quad \rho(x) = x + d_{J(x)}(0)$$

where $d_j(0)$ is the level-0 translation associated with the domain $\mathcal{D}_j(0)$, while the coding function J was defined in equation (13). Since the translation $d_j(0) \in \mathcal{M}$, each restricted coset $\widehat{\xi + \mathcal{M}}$, $\xi \in \Xi$, is ρ -invariant.

The *scaling map* ω is defined as multiplication by λ^3 in $\widehat{\mathbb{Q}(\lambda)}$. In contrast to the dynamical map, the scaling map does not leave the typical coset of \mathcal{M} invariant. Rather we have

$$\omega(\widehat{\xi + \mathcal{M}}) \subset \widehat{\xi' + \mathcal{M}} \quad \xi' \equiv \omega(\xi) \pmod{\mathcal{M}}$$

and $\xi' \in \Xi$ is determined uniquely by the above congruence. However, for every $n \in \mathbb{N}$

$$\frac{1}{n}\mathcal{M} = \bigcup_{\xi: n\xi \in \mathcal{M}} (\xi + \mathcal{M})$$

is invariant. Repeated application of ω eventually leads back to $\xi + \mathcal{M}$, since $\Xi \cap \frac{1}{n}\mathcal{M}$ contains precisely n^3 elements, so that there exists a well-defined function of ξ ,

$$\mu(\xi) = \min\{t \geq 1 : \omega^t(\xi) \equiv \xi \pmod{\mathcal{M}}\}, \quad (21)$$

which we call the *order* of ξ . Clearly, $\mu(\omega(\xi)) = \mu(\xi)$ and $\mu(\xi) \leq n^3$. A sharper bound for $\mu(\xi)$ can be obtained by considering the ideal factorization of n in $\mathbb{Z}[\lambda]$. We shall not pursue this matter here.

While $\omega(x)$ is well-defined for all $x \in \widehat{\mathbb{Q}(\lambda)}$, the same cannot be said for $\omega^{-1}(x) = \lambda^{-3}x$, which is meaningful only for $x \in \mathcal{D}_0(0) = [0, \omega)$.

The *left shift map* γ is defined as

$$\gamma : \widehat{\mathbb{Q}(\lambda)} \rightarrow \widehat{\mathbb{Q}(\lambda)} \quad \gamma(x) = \omega^{-1} \circ \rho^{-\tau(x)}(x)$$

where the function τ was defined in (11). Note that γ is defined on all of $\widehat{\mathbb{Q}(\lambda)}$, because, by construction, $\tau(x) \in \mathcal{D}_0(0)$. The transformation of cosets with respect to γ parallels that of ω : if $x \in \xi + \mathcal{M}$, then $\gamma(x) \in \xi' + \mathcal{M}$, where $\xi' \in \Xi$, $\xi' \equiv \omega^{-1}(\xi) \pmod{\mathcal{M}}$, and

$$\mu(\xi) = \min\{n \geq 1 : \gamma^n(\xi) \equiv \xi \pmod{\mathcal{M}}\}.$$

3.2 Three-dimensional representation: Pisot decomposition

We represent the map over the rational unit slab U , see (17), which is embedded in euclidean \mathbb{R}^3 .

The *dynamical map* ρ now acts as a translation by a 3-vector with integer components:

$$\rho : U \rightarrow U \quad \rho(x) = x + \mathbf{d}_{J(\Lambda \cdot x)}(0).$$

The coding function J and the vector Λ are defined in (13) and (17), respectively, while the values of $d_j(0)$ are listed in table 2.

The *scaling map* ω is defined as

$$\omega : U \rightarrow U \quad \omega(\mathbf{x}) = \mathbf{W} \cdot \mathbf{x}, \quad (22)$$

where \mathbf{W} is the matrix over \mathbb{Z} defined by the equation

$$\Lambda \cdot \mathbf{W} = \omega \Lambda. \quad (23)$$

The matrix \mathbf{W} is the matrix of multiplication by ω in $\mathbb{Q}(\lambda)$, with respect to the basis Λ , given by

$$W = \begin{pmatrix} 1 & -1 & 0 \\ -1 & 2 & -1 \\ -1 & 0 & 2 \end{pmatrix}. \quad (24)$$

Hence its characteristic polynomial is the minimal polynomial of ω , with roots ω and $\omega^{-\frac{1}{2}}e^{\pm i\theta}$, where

$$\cos \theta = \frac{1}{2}\omega^{\frac{1}{2}}(\text{Tr } \mathbf{W} - \omega).$$

The corresponding eigenvectors are \mathbf{e}_0 and $\mathbf{e}_{\pm} = \mathbf{e}_1 \mp i\mathbf{e}_2$, given by

$$\begin{aligned} \mathbf{e}_0 &= (1 + \lambda + \lambda^2, 1 + \lambda, 1) \\ \mathbf{e}_1 &= \left(\frac{\lambda^3-1}{2}, \frac{1-\lambda}{2}, 1\right) \\ \mathbf{e}_2 &= \frac{1}{2}\sqrt{\lambda^2 + 4\lambda - 1} (1, -\lambda^{-1}, 0). \end{aligned} \quad (25)$$

Applying both sides of (23) on \mathbf{e}_{\pm} , we obtain

$$\Lambda \cdot \mathbf{e}_{\pm} = 0,$$

which implies that the plane spanned by \mathbf{e}_1 and \mathbf{e}_2 is orthogonal to Λ . Note that there is no reason to assume any orthogonality relations among the 3 linearly independent real basis vectors \mathbf{e}_k . Introducing new coordinates by means of the matrix $\mathbf{T} = \mathbf{T}_{ij} = (\mathbf{e}_j)_i$, we have

$$\mathbf{T}^{-1} \cdot \mathbf{W} \cdot \mathbf{T} = \begin{pmatrix} \omega & 0 & 0 \\ 0 & \omega^{-\frac{1}{2}} \cos \theta & -\omega^{-\frac{1}{2}} \sin \theta \\ 0 & \omega^{-\frac{1}{2}} \sin \theta & \omega^{-\frac{1}{2}} \cos \theta \end{pmatrix}.$$

Since $0 < \omega < 1 < 1/\sqrt{\omega}$, repeated application of the scaling map is asymptotically expanding in directions perpendicular to Λ , with an accompanying linear transformation conjugate to a rotation through angle θ .

Repeated application of the inverse map $\omega^{-1}(\mathbf{x})$, where it is well-defined, is dominated by expansion along the \mathbf{e}_0 direction, with eigenvalue ω^{-1} . Since \mathbf{e}_0 has a non-zero projection onto planes perpendicular to Λ , the expansion will be present in such planes as well as parallel to Λ .

The *left shift map* γ

$$\gamma : U \rightarrow U \quad \gamma(\mathbf{x}) = \omega^{-1} \circ \rho^{-\tau(\mathbf{x})}(\mathbf{x})$$

combines an \mathcal{M} -displacement with application of the matrix W^{-1} . Since $\Lambda \cdot \rho^{-\tau(\mathbf{x})}(\mathbf{x}) \in [0, \omega)$, $\gamma(\mathbf{x})$ is always well-defined. Moreover, as we shall see, the component of $\rho^{-\tau(\mathbf{x})}$ along \mathbf{e}_0 is uniformly bounded, so that for sufficiently large vectors in a plane perpendicular to Λ , the contraction (with rotation) due to ω^{-1} in fact dominates under repetition of γ .

The above contraction property underpins the main result of this section. In preparation for it, we establish some notation. We represent $\mathbf{x} \in U$ as

$$\mathbf{x} = (x_0, x_1, x_2) = \alpha_0(\mathbf{x})\mathbf{e}_0 + \alpha_1(\mathbf{x})\mathbf{e}_1 + \alpha_2(\mathbf{x})\mathbf{e}_2$$

where the vectors \mathbf{e}_k are given in (25) and

$$\alpha_k(\mathbf{x}) = \sum_{j=0}^2 (\mathbb{T}^{-1})_{kj} x_j, \quad k = 0, 1, 2.$$

Then we define

$$x_{\perp} = \sqrt{x_1^2 + x_2^2} \quad \alpha_{\perp}(\mathbf{x}) = \sqrt{\alpha_1(\mathbf{x})^2 + \alpha_2(\mathbf{x})^2} \quad (26)$$

(the quantity x_{\perp} will be used in the next section) and

$$R = \sup \left\{ \alpha_{\perp} \left(\sum_{s=0}^t d_{p(j,s)}(0) \right) : 0 \leq j \leq N-1, 0 \leq t \leq \nu_j - 1 \right\}. \quad (27)$$

Lemma 8 *Let $B = \{\mathbf{x} \in U : \alpha_{\perp}(\mathbf{x}) \leq R_0\}$, where $R_0 = R\sqrt{\omega}/(1 - \sqrt{\omega})$. Then $\gamma(B) \subset B$.*

Proof. For $k = 1, 2$, we have

$$\alpha_k(\omega^{-1}(\mathbf{x})) = \sum_{j=0}^2 (\mathbb{T}^{-1} \cdot W \cdot \mathbb{T})_{kj} \alpha_j(\mathbf{x}) = \sqrt{\omega} \sum_{j=0}^2 R_0(\theta)_{kj} \alpha_j(\mathbf{x}),$$

where $R_0(\theta)$ is the matrix representing a rotation by angle θ about the α_0 axis. From the rotational invariance of $\alpha_{\perp}(\mathbf{x})$ and the triangle inequality for α_{\perp} , we have

$$\alpha_{\perp}(\gamma(\mathbf{x})) = \sqrt{\omega} \alpha_{\perp}(\rho^{-\tau(\mathbf{x})}\mathbf{x}) \leq \sqrt{\omega} \left(\alpha_{\perp}(\mathbf{x}) + \sup_{\mathbf{x}} |\alpha_{\perp}(\rho^{-\tau(\mathbf{x})}\mathbf{x}) - \alpha_{\perp}(\mathbf{x})| \right). \quad (28)$$

Now

$$\begin{aligned} & \sup_{\mathbf{x}} \left| \alpha_{\perp}(\rho^{-\tau(\mathbf{x})}\mathbf{x}) - \alpha_{\perp}(\mathbf{x}) \right| \\ &= \sup \{ |\alpha_{\perp}(\rho^t(\mathbf{y})) - \alpha_{\perp}(\mathbf{y})| : 0 \leq j \leq 6, 0 \leq t \leq \nu_j - 1, \Lambda \cdot \mathbf{y} \in \mathcal{D}_j(1) \} \\ &\leq \sup \{ \alpha_{\perp}(\rho^t(\mathbf{y}) - \mathbf{y}) : 0 \leq j \leq 6, 0 \leq t \leq \nu_j - 1, \Lambda \cdot \mathbf{y} \in \mathcal{D}_j(1) \} \\ &= R, \end{aligned}$$

However, for all j, t and \mathbf{y} as specified above, we also have

$$\rho^t(\mathbf{y}) - \mathbf{y} = \sum_{s=0}^t d_{p(j,s)}(0).$$

This, together with (28), gives

$$\alpha_{\perp}(\gamma(x)) \leq \sqrt{\omega}(\alpha_{\perp}(x) + R).$$

If $x \in B$, the above inequality gives

$$\alpha_{\perp}(\gamma(x)) \leq \sqrt{\omega}(\alpha_{\perp}(x) + R) \leq \sqrt{\omega}(R_0 + R) = R_0, \quad (29)$$

so that $\gamma(B) \subset B$, as desired. \square

3.3 Two-dimensional representation: discrete vector fields

In this section we represent the map over the lattice $\mathcal{L}_{\xi} \cong \mathbb{Z}^2$; each lattice supports a ξ -dependent discrete vector field.

For every $\xi \in \Xi$, we represent the dynamical map ρ as

$$\rho : \mathcal{L}_{\xi} \rightarrow \mathcal{L}_{\xi} \quad \rho(\mathbf{m}) = \mathbf{m} + \mathbf{d}_{J(\Lambda \cdot \mathbf{m})}(0),$$

where J is the coding function (13). The vector $\mathbf{m} \in U$ depends on $\mathbf{m} = (m_1, m_2)$ and ξ via

$$\mathbf{m} = (-\lfloor \xi + m_1\lambda + m_2\lambda^2 \rfloor, m_1, m_2)$$

while the vector field values

$$\mathbf{d}_j(0) = (d_j(0)_1, d_j(0)_2) \quad j = 0, \dots, 6$$

can be read off table 2, and are shown in figure 2.




Figure 2: Lattice translation vectors corresponding to the action of $\rho(0)$ on the seven domains $\mathcal{D}(0)$. The integer components are listed in table 2.

With reference to the commutative diagram (19), the action of the scaling map ω on each \mathcal{L}_{ξ} can be inferred from that on U (cf. (22)) as $\varphi_2 \circ \omega \circ \varphi_2^{-1}$. This is done by

restricting to the last two components of $\mathbf{m} = (m_0, m_1, m_2) \in U$, and taking into account that the image point will be in $\mathcal{L}_{\xi'}$, for some $\xi' \in \Xi$ uniquely determined by the congruence $\xi' \equiv \omega(\xi) \pmod{\mathcal{M}}$. Explicitly

$$\omega(\mathbf{m}) = \left(\sum_{j=0}^2 W_{1j} m_j + \left\lfloor \sum_{j=0}^2 W_{1j} \xi_j \right\rfloor, \sum_{j=0}^2 W_{2j} m_j + \left\lfloor \sum_{j=0}^2 W_{2j} \xi_j \right\rfloor \right),$$

where $m_0 = -\lfloor \xi + m_1 \lambda + m_2 \lambda^2 \rfloor$.

The action of the left shift map γ is again our main concern. The trapping of the γ -orbit of an arbitrary point in U in a bounded set B (lemma 8), is mirrored by orbit-trapping in each \mathcal{L}_ξ under iteration of $\gamma^{\mu(\xi)}$. The following lemma is the \mathcal{L}_ξ equivalent of lemma 8; in addition, it establishes that the trapping regions B_ξ for γ -orbits in the various \mathcal{L}_ξ are uniformly bounded.

Lemma 9 *For every $\xi \in \Xi$, the set*

$$B_\xi = \{(m_1, m_2) : (\xi_0 + m_0, \xi_1 + m_1, \xi_2 + m_2) \in B\}, \quad m_0 = -\lfloor \xi + m_1 \lambda + m_2 \lambda^2 \rfloor,$$

is a finite subset of \mathcal{L}_ξ contained in a closed disk of uniformly bounded radius

$$r_\xi = \max\{x_\perp : x \in B\} + \xi_\perp < r_0 + \sqrt{2}.$$

If $k = \mu(\xi)$ is the order of ξ , then $\gamma^k(B_\xi) \subset B_\xi$. Moreover, for arbitrary $\mathbf{m} \in \mathcal{L}_\xi$ and sufficiently large integer t , the points $\gamma^{tk}(\mathbf{m})$ belong to a limit cycle in B_ξ .

Proof. Clearly $|\mathbf{m}| \leq x_\perp + \xi_\perp$, from which the first statement follows. Since $\{(\xi_0, \xi_1, \xi_2) : \xi \in \Xi\}$ is a subset of the unit cube, ξ_\perp is bounded above by $\sqrt{2}$. To prove the final statement, we note that by lemma 8 the points $\gamma^{tk}(\mathbf{m})$ are restricted to a uniformly bounded portion of the lattice. Thus, the sequence must eventually terminate in a limit cycle. Each point of the limit cycle is a fixed point for an iterate of γ^k . Such a fixed point, by lemma 8, cannot exist outside B_ξ . \square

Since every choice of ξ and $\mathbf{m} \in \mathcal{L}_\xi$ corresponds to a unique $x \in U$, we immediately get from lemma 9 the following corollary.

Corollary 10 *For all $x \in U$, there exists a positive integer t_0 such that $\gamma^t(x) \in B$ for all $t > t_0$.*

3.4 Proofs of main results

We have now laid all the groundwork for proving theorems 1 and 2 for the scale-invariant AY map ρ , together with a related result for the full AY map ϕ (proposition 13 below).

We shall also establish the following additional results.

Proposition 11 *Every ρ -orbit in \mathcal{L}_ξ contains a unique periodic-code point, which belongs to B_ξ , with one exception. The exceptional case is an orbit in \mathcal{L}_0 which contains two periodic-code points, namely $(0,0)$ and $(1,0)$.*

Proposition 12 *Let M be the bound of theorem 2. Then for every $n \in \mathbb{N}$, the number of orbits in $\widehat{\frac{1}{n}\mathcal{M}}$ is less than Mn^3 ; the code period for any $x \in \widehat{\frac{1}{n}\mathcal{M}}$ is also less than Mn^3 . Furthermore, we have the estimates*

$$M \leq 225 \quad r \leq \frac{15}{2}$$

with r as in lemma 9.

Proof of theorem 1(i) and proposition 11. Let $x \in \Omega$ be given, and let $\sigma(x)$ be the code of x . If $\sigma(x)$ is eventually periodic, then $x \in \widehat{\mathbb{Q}(\lambda)}$, from equation (16) of section 2.4. Conversely, let $x \in \widehat{\mathbb{Q}(\lambda)}$. Then $x \in \widehat{\xi + \mathcal{M}}$ for some $\xi \in \Xi$; let \mathbf{m} be the corresponding point of \mathcal{L}_ξ . By lemma 9, repeated application of $\gamma^{\mu(\xi)}$ on \mathbf{m} , eventually terminates in a limit cycle, each point of which has a periodic code. This establishes theorem 1(i).

If n is the code period for points of the limit cycle, there exists a t such that $\gamma^{nt}(\mathbf{m})$ is a periodic-code point \mathbf{p} with the same code tail as \mathbf{m} , and hence \mathbf{p} and \mathbf{m} share the same ρ -orbit. Due to the rules of Vershik updating, there is no other point with periodic code on the same orbit, except in the case $\xi = 0$, where the orbit containing the origin (code $(0,0)^\infty$) also contains $(1,0)$ (code $(6,3)^\infty$). This exception corresponds to the fact that there is a unique maximal code which gets mapped into the unique minimal code. This completes the proof of proposition 11. \square

Proof of theorem 2. From lemma 9 the number of points in the core region B_ξ is less than or equal to the number M of lattice points in a square of side $2[r] + 1$. The number of orbits, by proposition 11, does not exceed the number of periodic-code points in B_ξ , and hence is bounded above by M . \square

Proof of proposition 12. There are n^3 cosets $\frac{k_0}{n} + \frac{k_1}{n}\lambda + \frac{k_2}{n}\lambda^2 + \mathcal{M}$, $k_i = 0, \dots, n-1$, and each one restricted to Ω has at most M orbits.

To bound the code period, we consider without loss of generality a point in some $\widehat{\xi + \mathcal{M}}$ with strictly periodic code. The corresponding lattice point \mathbf{m} is one of at most M points in the core subset of \mathcal{L}_ξ . Moreover, for all $t \in \mathbb{N}$, $\gamma^{t\mu(\xi)}(\mathbf{m})$ is also in the core, and so there exists a $t \leq M$ such that $\mathbf{m} = \gamma^{t\mu(\xi)}(\mathbf{m})$. Thus the period is bounded by $M\mu(\xi) \leq Mn^3$.

To estimate M and r , we calculate explicitly $\mathbf{d} = \sum_{s=0}^{t-1} \mathbf{d}_{p(j,s)}(0)$ and $\alpha_\perp(\mathbf{d})$ for all 54 of the combinations $j = 0, \dots, 6$, $t = 1, \dots, \nu_j - 1$. The lattice displacement vector \mathbf{d} maximizing $\alpha_\perp(\mathbf{d})$ is found to be $(-4, 8, 1)$, for $j = 1, t = 12$, giving

$$R = \frac{1}{11}(311 + 268\lambda + 155\lambda^2) = 6.759\dots$$

and

$$R_0 = \frac{R}{(\lambda^{-\frac{3}{2}} - 1)} = 4.522\dots$$

We now turn to lemma 9. We need to bound x_\perp , where $x = \alpha_0 e_0 + \alpha_1 e_1 + \alpha_2 e_2$. The condition $\Lambda \cdot x \in [0, 1)$, with $\Lambda \cdot e_1 = \Lambda \cdot e_2 = 0$ places a bound on α_0 :

$$\alpha_0 < \frac{1}{\Lambda \cdot e_0} = \frac{1}{1 + 2\lambda + 3\lambda^2} = 0.336\dots$$

Further, the triangle inequality requires

$$x_\perp \leq \alpha_0 (e_0)_\perp + (\alpha_1 e_1 + \alpha_2 e_2)_\perp.$$

We note that the second term on the right-hand side achieves its maximum value on the boundary of the circle $\alpha_1^2 + \alpha_2^2 = R_0^2$, so that

$$r \leq \alpha_0 (e_0)_\perp + \max_{|\alpha_1| \leq R_0} \left(\alpha_1 e_1 + \sqrt{R_0^2 - \alpha_1^2} e_2 \right)_\perp + \sqrt{2}.$$

The function of α_1 on the right-hand side is found by elementary methods to be maximized for

$$\alpha_1 = -R_0 \left(\frac{1}{2} - \frac{1}{2} \sqrt{\frac{\lambda - \lambda^2}{2}} \right)^{\frac{1}{2}} = -2.573\dots$$

Evaluating the bound on r using interval arithmetic, using $\lambda \in (543/1000, 544/1000)$, we obtain the rigorous bound

$$r < \frac{15}{2}.$$

By lemma 9, the set B_ξ is contained in a disk of radius $\frac{15}{2}$, hence in a square of side 15, containing 225 lattice points. \square

Proof of theorem 1(ii) We already know (proposition 7) that the discontinuity points are precisely the points with fixed code. The points of Γ are obtained from the discontinuity points via iteration, and the forward and backward dynamics do not change the tail of the code (apart from one exception mentioned above, where, however, the length of the tail remains unchanged). Thus if $x \in \Gamma$, then $\sigma(x)$ is eventually fixed. Conversely, if $\sigma(x)$ is eventually fixed, then the orbit of x contains a point with fixed code, which shows that $x \in \Gamma$.

From table 2 we see that the \mathbb{Z} -module generated by the translations and discontinuity points of ρ is just

$$\mathcal{N} = \mathcal{M} \cup (\eta + \mathcal{M}), \quad \eta = \frac{1}{2} + \frac{\lambda^2}{2}.$$

The restriction $\widehat{\mathcal{N}}$ of the above module, corresponds to the pair of lattices \mathcal{L}_0 and \mathcal{L}_η . For each lattice, we explicitly verify that a 15×15 square centred at the origin lies on the orbit of one of the following points in the square:

$$\begin{aligned} L_0 : & \quad (0, 0), (1, 0), (-1, -1), (-2, 1) \\ L_\eta : & \quad (-2, -1), (0, -1), (1, 1). \end{aligned}$$

These are precisely the lattice counterparts of the 7 discontinuity points. By theorem 2 there are no other orbits, and hence $\widehat{\mathcal{N}} = \Gamma$. \square

The results we have obtained for the map ρ can easily be extended to the full AY map ϕ . Recall that there is a ratio $\alpha = -\frac{1}{2} + \frac{\lambda}{2} + \lambda^2 = (9 + 8\lambda + 5\lambda^2)^{-1}$ between the defining intervals of the 2 maps. For the AY map, the \mathbb{Z} -modules \mathcal{M} and \mathcal{N} are, attaching superscripts to distinguish the contexts,

$$\mathcal{M}^{(\phi)} = \alpha\mathcal{M}^{(\rho)} = \{y = y_0 + y_1\lambda + y_2\lambda^2 : y_i \in \frac{1}{2}\mathbb{Z}, y_0 + y_1 + y_2 \in \mathbb{Z}\} \subset \mathcal{N}^{(\phi)},$$

$$\mathcal{N}^{(\phi)} = \alpha\mathcal{N}^{(\rho)} = \frac{1}{2}\mathcal{M}.$$

Similarly, for $\xi \in \Xi$,

$$\alpha(\xi + \mathcal{M}^{(\rho)}) = \alpha\xi + \mathcal{M}^{(\phi)}.$$

These correspondences can readily be verified by explicit multiplication by α . In addition, we have

Proposition 13 *The discontinuity set $\Gamma^{(\phi)}$ is equal to $\mathcal{N}^{(\phi)}$. It decomposes into 6 disjoint orbits, whose restrictions to $[0, \alpha)$ are just the 6 constituent orbits of $\Gamma^{(\rho)}$, rescaled.*

Proof. For each discontinuity point w for ϕ , there exists a non-negative integer t such that $\alpha^{-1}\phi^{-t}(w)$ is a discontinuity point for ρ . We have verified this by explicit iteration of ϕ^{-1} , with the details listed in table 5. \square

w	t	$\alpha^{-1}\phi^{-t}(w)$
0	0	0
$\frac{1}{2} - \frac{\lambda}{2}$	6	λ
$\frac{\lambda}{2}$	23	$-\frac{1}{2} + \lambda + \frac{3\lambda^2}{2}$
λ	31	$1 - 2\lambda + \lambda^2$
$\lambda + \frac{\lambda^2}{2}$	16	$\frac{3}{2} - 2\lambda - \frac{\lambda^2}{2}$
$\lambda + \lambda^2$	7	$1 - \lambda - \lambda^2$
$\frac{1}{2} + \frac{\lambda}{2} + \frac{\lambda^2}{2}$	3	$\frac{1}{2} - \frac{\lambda^2}{2}$

Table 5: ρ -discontinuity points as rescaled pre-images of ϕ -discontinuity points.

4 Lattice orbits

Thus far we have determined that each lattice \mathcal{L}_ξ is composed of a finite number of disjoint walks, each of which passes through a periodic-code point in a core region surrounding the

origin. Exactly how the orbits coexist on each lattice has not yet been made clear. Experiments (figure 3) indicate that the orbits are segregated into non-overlapping sectors, with boundaries consisting of elementary steps connecting neighbouring points. This behaviour was described in theorem 3 in the introduction, which will be proved in the course of sections 4.1–4.4 as a collection of lemmas. We shall utilize not only the general properties established in the preceding section, but also the detailed (and presumably non-universal) lattice dynamics of the map ρ .




Figure 3: (a) Orbits passing through $(0,0)$, $(-2,1)$, and $(-1,-1)$ and their boundary curves in \mathcal{L}_0 . (b) Orbits passing through $(-2,-1)$, $(0,-1)$, and $(1,1)$ and their boundary curves in \mathcal{L}_η . The central square of side 15 has been highlighted in both cases.

4.1 Absence of crossing

A striking feature of the orbits in \mathcal{L}_ξ is the absence of orbit crossing. We now show how this property arises from the lattice dynamics.

Lemma 14 *For any $\mathbf{m} \in \mathcal{L}_\xi$, let $\mathbf{d} = \rho(\mathbf{m}) - \mathbf{m}$ be the translation vector at \mathbf{m} , and let \mathcal{O} be the orbit through \mathbf{m} . Then no orbit other than \mathcal{O} can land on the segment connecting \mathbf{m} and $\mathbf{m} + \mathbf{d}$. No orbit, including \mathcal{O} itself, can cross the line segment connecting \mathbf{m} and $\mathbf{m} + \mathbf{d}$.*

Proof. The only “land on” case would have $\mathbf{d} = \mathbf{d}_1$. But, consulting table 4, we see that the two interior points of the segment can be identified as $\rho^{-1}(\mathbf{m})$ and $\rho^2(\mathbf{m})$, both points of \mathcal{O} .

To study the possibilities for crossing, we classify the translation vectors \mathbf{d}_j according to directions of the compass: $\text{EW} = \{\mathbf{d}_1, \mathbf{d}_6\}$, $\text{SW-NE} = \{\mathbf{d}_0, \mathbf{d}_3, \mathbf{d}_4\}$, $\text{NW-SE} = \{\mathbf{d}_2, \mathbf{d}_5\}$. We note that it is impossible to cross an EW segment, since no translation vector has a vertical displacement greater than 1 in magnitude. For crossing to occur, an SW-NE step must cross an NW-SE step. We consider in turn each of the pairings, illustrating the hypothetical translation vector in figure 4.




Figure 4: Hypothetical orbit crossing scenarios.

- (a) \mathbf{d}_0 crosses \mathbf{d}_2 or \mathbf{d}_5 : this would require a succession of steps 0, 6, 5 or 0, 6, 2, which is forbidden according to table 4;
- (b) \mathbf{d}_3 crosses \mathbf{d}_2 or \mathbf{d}_5 : this is forbidden, since 6,3,6 cannot be preceded or followed by 2 or 5;
- (c) \mathbf{d}_4 crosses \mathbf{d}_2 : this has 2 steps originating at the same point (forbidden);
- (d) \mathbf{d}_4 crosses \mathbf{d}_5 : this has 2 steps ending at the same point (forbidden). □

4.2 Boundary curves

The absence of orbit crossing leads us to consider the planar regions occupied by the orbits. In this section we explore the structure of the boundaries of these regions. We begin by defining the relevant geometrical objects.

- (i) Two lattice points $\mathbf{m}, \mathbf{n} \in \mathbb{Z}^2$ are called *neighbours* (or *neighbouring points*) if $0 < \|\mathbf{m} - \mathbf{n}\| < 2$, where $\|\cdot\|$ is the euclidean norm. Such points are either nearest neighbours (NN), with $\|\mathbf{m} - \mathbf{n}\| = 1$, or next-nearest neighbours (NNN), with $\|\mathbf{m} - \mathbf{n}\| > 1$.
- (ii) The *phase* $\theta(\mathbf{n})$ of a neighbour \mathbf{n} of $\mathbf{m} \in \mathbb{Z}^2$, is an integer modulo 8 obtained assigning $\theta(\mathbf{n}) = 0$ to $\mathbf{n} = x + (0, -1)$ and incrementing by 1 as one proceeds around \mathbf{m} counterclockwise.
- (iii) An *elementary walk* on \mathbb{Z}^2 is a polygonal line (possibly semi-infinite or infinite) whose successive vertices are neighbouring lattice points.
- (iv) A *boundary point* for a ρ -orbit \mathcal{O} is a point \mathbf{m} which has a nearest neighbour not belonging to \mathcal{O} . A *boundary curve* associated with \mathcal{O} , is an infinite, non-self-intersecting

elementary walk whose vertices are boundary points of \mathcal{O} . A *shadowing curve* for a ρ -orbit \mathcal{O} is an infinite, non-self-intersecting elementary walk whose vertices are nearest neighbours of boundary points of \mathcal{O} but do not themselves belong to \mathcal{O} .

- (v) A *right* (resp. *left*) *boundary curve* for \mathcal{O} is a boundary curve such that each point on it has a nearest neighbour to the right (resp. left) of the curve which does not belong to \mathcal{O} . (The orientation agrees with the dynamics.)

That the boundary curves in \mathcal{L}_ξ do not intersect (theorem 3) is a direct consequence of the following lemma:

Lemma 15 *Let \mathcal{O} and \mathcal{E} be, respectively, a ρ -orbit and an elementary walk disjoint from \mathcal{O} , on \mathcal{L}_ξ . Then it is not possible for \mathcal{O} , or any boundary curve of \mathcal{O} , to cross \mathcal{E} .*




Figure 5: Hypothetical boundary crossing scenarios.

Proof. As in the proof of lemma 14, we consider all possibilities for the crossing segment in \mathcal{O} , omitting types \mathbf{d}_1 and \mathbf{d}_6 because the vertical separation between successive points on \mathcal{E} cannot exceed unity. The remaining scenarios are shown in figure 5. In each case, one of the lattice sites cannot be assigned consistently to either \mathcal{O} or \mathcal{E} . Hence we are left with no way for \mathcal{O} to cross \mathcal{E} . This also precludes crossing by a boundary curve of \mathcal{O} , since the latter would place points of \mathcal{O} on both sides of \mathcal{E} , and thus \mathcal{O} would have to cross \mathcal{E} somewhere. \square

4.3 Iterative construction of boundaries

Starting with a boundary point of an orbit in \mathcal{L}_ξ , one can systematically construct an arbitrarily long segment of an infinite boundary curve containing that point. The basis for the construction is the following lemma.

Steps	$\rho^{-2}(\mathbf{m})$	$\rho^{-1}(\mathbf{m})$	$\rho(\mathbf{m})$	$\rho^2(\mathbf{m})$	$\rho^3(\mathbf{m})$
0	—	2	3	4	—
561	4	2	—	—	—
361	7	2	—	—	—
2	—	2	1	—	—
3	—	2	—	3	—
24	4	5	7	—	—
64	0	2	7	—	—
25	4	5	1	0	—
45	—	3	1	0	—
06	0	7	6	—	—
162	—	—	6	0	—
163	—	—	6	3	4
360	—	—	6	4	5
361	—	—	6	—	2
46	—	3	6	4	—
56	—	5	6	—	—

Table 6: Phases of neighbouring orbit points for every possible assignment of the translation vector. In the first column, the translation vector \mathbf{d}_n corresponding to $\rho(\mathbf{m})$ is given by its subscript in boldface, with enough predecessors and successors to uniquely define the orbit segment. For example, if $\rho(\mathbf{m})$ corresponds to a lattice translation \mathbf{d}_1 , it must be preceded by either \mathbf{d}_5 and \mathbf{d}_6 or \mathbf{d}_3 and \mathbf{d}_6 .

Lemma 16 *On a ρ -orbit $\mathcal{O} \subset \mathcal{L}_\xi$, every point has at least two neighbours and one nearest neighbour on \mathcal{O} .*

Proof. For an arbitrary $\mathbf{m} \in \mathcal{O}$ we can calculate the translation vector $\rho(\mathbf{m}) - \mathbf{m} \in \{\mathbf{d}_0, \mathbf{d}_1, \dots, \mathbf{d}_6\}$, as well as that of its immediate predecessors and successors $\rho^n(\mathbf{m})$, $n = -2, -1, 1, 2, 3$. The search for a pair of neighbours succeeds in all cases; the details are displayed in table 6. As to the second claim, we note that every row of table 6 has an even entry. The even phases correspond to the nearest neighbours. This result also follows from the no-crossing lemma 15, with \mathcal{E} taken to be the 4-step walk connecting the nearest neighbours of the given point. \square

We now show that, given any boundary point \mathbf{m}_0 of a ρ -orbit \mathcal{O} on \mathcal{L}_ξ , there exists an iterative construction for an unbounded boundary curve of \mathcal{O} containing \mathbf{m}_0 . The procedure is the following.

Boundary construction algorithm. *Input: boundary point $\mathbf{m}_0 \in \mathcal{O}$, a nearest neighbour of \mathbf{m}_0 not belonging to \mathcal{O} , and positive integers M and N . Output: sequence $\mathbf{m}_{-N}, \dots, \mathbf{m}_0, \dots, \mathbf{m}_M$ of points on a boundary curve of \mathcal{O} , and a segment of a shadowing curve of \mathcal{O} whose vertices are nearest neighbours of the points \mathbf{m}_i .*

The points $\mathbf{m}_1, \dots, \mathbf{m}_M$, together with a shadowing curve segment, are obtained by iterative application of the following procedure. Given $\mathbf{m}_i \in \mathcal{O}$ and a nearest neighbour

$\mathbf{n}_i \notin \mathcal{O}$, let θ be the phase of \mathbf{n}_i relative to \mathbf{m}_i . Consider in turn each neighbour of \mathbf{m}_i , starting at phase $\theta + 1$ and incrementing the phase by 1 until we find another member of \mathcal{O} at some phase θ' . By lemma 16 such a point exists and we choose it to be \mathbf{m}_{i+1} . It is easy to see that the neighbour of \mathbf{m}_i at phase $\theta' - 1$ is a nearest neighbour of \mathbf{m}_{i+1} , not on \mathcal{O} , and we choose it to be \mathbf{n}_{i+1} . By definition, \mathbf{m}_{i+1} is a boundary point. Moreover, the nearest neighbours (if any) of \mathbf{m}_i encountered between phases $\theta + 1$ and $\theta' - 2$, inclusive, are, together with \mathbf{n}_i and \mathbf{n}_{i+1} , the vertices of a segment of a shadowing curve for \mathcal{O} .

Applying the iterative step $M - 1$ times produces a sequence of pairwise neighbouring boundary points which constitutes the forward part of the desired boundary segment, with its corresponding shadowing curve segment. The backward part is obtained by an obvious modification of the above iterative process, to obtain, in order, $\mathbf{m}_{-1}, \dots, \mathbf{m}_{-N}$, with its shadowing curve segment. The order of the points \mathbf{m}_i induces an orientation on the boundary curve connecting them. This orientation is a convention which may or may not coincide with the orientation inferred from the iterated application of the map ρ .

Crucial to the execution of the above procedure is an efficient sub-routine for determining whether a point \mathbf{m} belongs to a given orbit \mathcal{O} . Each orbit is associated with a periodic-code point, which is unique except in the special case $\xi = 0$, where a single orbit contains the points $(0, 0)$ and $(1, 0)$. A periodic-code point is a fixed point of γ^τ , where τ is the code period, a multiple of the order $\mu(\xi)$. Having found all the periodic-code points, we can determine which orbit a point \mathbf{x} belongs to by iterating $\gamma^{\mu(\xi)}$ on \mathbf{x} until a match with a periodic-code point is obtained. The number of iterations required increases, roughly, proportional to the level in the hierarchy, i.e., logarithmically with respect to the orbit segment length.

We note that the construction algorithm is reversible, in the sense that if $\mathbf{y} \in \mathcal{O}$ and $\mathbf{y}' \notin \mathcal{O}$ are nearest neighbours obtained via the algorithm starting from the pair $\mathbf{x} \in \mathcal{O}$ and $\mathbf{x}' \notin \mathcal{O}$, then the same statement holds with the pairs interchanged. From theorem 1, the lattice orbits of the 7 discontinuity points of the interval exchange map ρ are restricted to \mathcal{L}_0 and \mathcal{L}_η , $\eta = \frac{1}{2} + \frac{\lambda^2}{2}$. These two lattices provide instructive examples of the layout of orbits and their boundary curves, as depicted in figures 3 and 6.

4.4 Arrangement of lattice orbits

We now have all the ingredients to describe how, in general, the ρ -orbits share the lattice \mathcal{L}_ξ . First, we note that the layout is completely organized by what happens in the *central square* \mathcal{S} of side $a = 2[r_\xi] + 1$ centred at the origin. By proposition 12, we may choose $a = 15$, uniformly, so that the perimeter of \mathcal{S} is $4a - 4 = 56$. The boundary $\partial\mathcal{S}$ of \mathcal{S} is the concatenation of n segments $\partial\mathcal{S}_i$, $i = 0, 1, \dots, n - 1$ containing points of ρ -orbits \mathcal{O}_i , with $\mathcal{O}_i \neq \mathcal{O}_j$, $j \equiv i + 1 \pmod{n}$. Since the core region B_ξ is contained in \mathcal{S} , the \mathcal{O}_i are the only orbits present on the infinite lattice. Note that each \mathcal{O}_i is associated with a periodic-code point $\mathbf{p}_i \in \mathcal{S}$, but we cannot exclude the possibility $\mathbf{p}_i = \mathbf{p}_j$ for some i, j pair.

If $n = 1$, so that all points of $\partial\mathcal{S}$ belong to a single orbit \mathcal{O} , then clearly \mathcal{O} visits

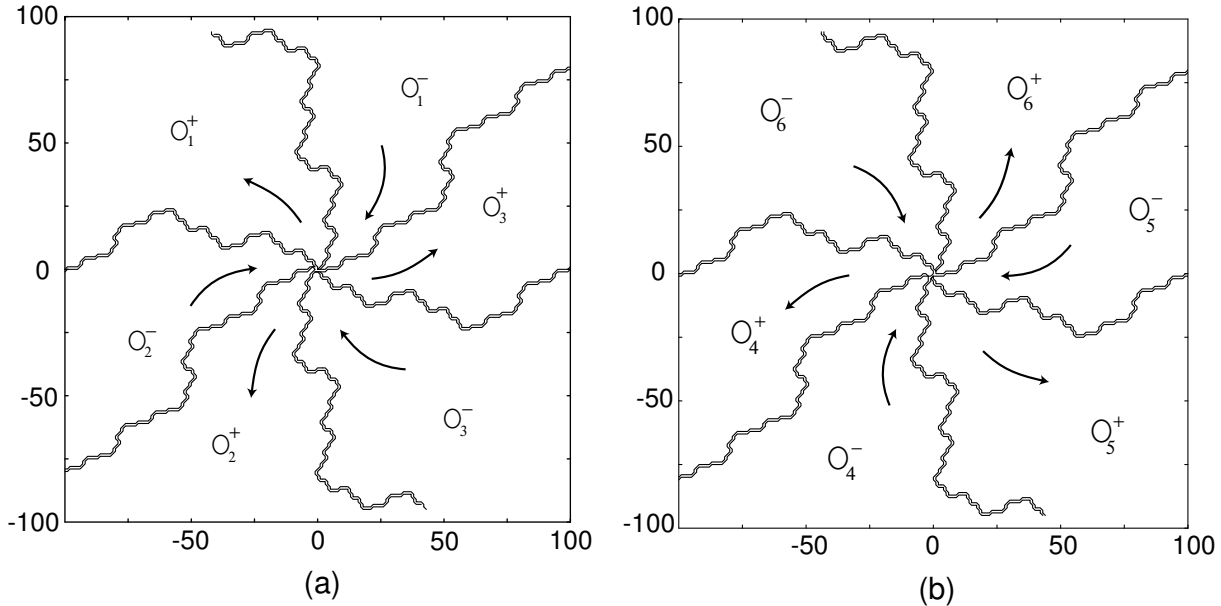


Figure 6: (a) Right and left boundary curves of the semi-orbits \mathcal{O}_i^\pm , $i = 1, 2, 3$ in \mathcal{L}_0 . (b) Right and left boundary curves of the semi-orbits \mathcal{O}_i^\pm , $i = 4, 5, 6$ in \mathcal{L}_η .

every point of \mathcal{L}_ξ , and there are no boundary curves. This establishes the first statement of theorem 3. Assume, then, that $n > 1$. Let \mathbf{x}_i^L and \mathbf{x}_i^R be the left- and right-hand (as viewed from inside $\partial\mathcal{S}$) endpoints of $\partial\mathcal{S}_i$, $i = 0, \dots, n-1$. Each of these points has a nearest neighbour on a different orbit, and so is a boundary point of \mathcal{O}_i . Using the boundary construction algorithm, one can construct an infinite boundary curve $\partial\mathcal{O}_i^L$ (resp. $\partial\mathcal{O}_i^R$) containing \mathbf{x}_i^L (resp. \mathbf{x}_i^R). These boundary curves are not necessarily distinct; however, every boundary point of every ρ -orbit in \mathcal{L}_ξ lies on boundary curve constructed via the algorithm from a $\mathbf{x}_i^R, \mathbf{x}_j^L$ pair on $\partial\mathcal{S}$. To prove this statement, which implies the rest of theorem 3, suppose \mathbf{x} is a boundary point of \mathcal{O} not lying on one of the constructed boundary curves. Choosing any nearest neighbour of \mathbf{x} not belonging to \mathcal{O} , we construct a new boundary curve $\partial\mathcal{O}$ containing \mathbf{x} , as well as its shadowing curve, $\partial\tilde{\mathcal{O}}$. Now $\partial\mathcal{O}$ must intersect $\partial\mathcal{S}$, since otherwise there would be an orbit which never enters the central square, which is forbidden. If $\partial\mathcal{O}$ intersects $\partial\mathcal{S}$ but $\partial\tilde{\mathcal{O}}$ does not, the same problem arises. Therefore both $\partial\mathcal{O}$ and $\partial\tilde{\mathcal{O}}$ intersect $\partial\mathcal{S}$, in points which can be identified with some $\mathbf{x}_i^R, \mathbf{x}_j^L$ pair. By the reversibility of the algorithm, the two boundary curves coincide and we get a contradiction.

Topologically, each distinct ρ -orbit has either a single left or right boundary, with a single opening to infinity, or a left-right pair of boundary curves, with two openings to infinity. In the single-opening case, the forward and backward parts of the orbit are side by side, and one can construct, using the algorithm, a pair of semi-bounded boundary curves, separating them. The half-boundary curves are, of course, not well defined within 3ρ -steps of their terminal points, since the two-neighbours lemma on which the algorithm depends breaks down there. In the double-opening case, the orbit comes in from infinity from one sector and, after crossing the central square, leaves from the other.

That every boundary curve $\partial\mathcal{O}$ is paired with an infinite, non-self-intersecting elementary walk $\partial\tilde{\mathcal{O}}$ satisfying the properties of theorem 3 follows from the explicit construction of section 4.3. This completes the proof of theorem 3.

Finally, we address, at least qualitatively, the apparent spiralling outward of the lattice orbits. For a given periodic-code point, \mathbf{p} , consider the sequence of points $\mathbf{x}_i = \omega^{i\tau}(\mathbf{p})$, $i = 0, 1, \dots$, where τ is the code period. The hierarchical codes of the \mathbf{x}_i have the same tail as \mathbf{p} , displaced i periods to the right with padding, and hence they lie on the same \mathcal{L}_ξ -orbit. After the first few iterations, the mapping $\gamma^\tau : \mathbf{x}_i \mapsto \mathbf{x}_{i+1}$ is well approximated by a rescaling by $\omega^{-\tau/2} \approx 2.5^\tau$ combined with a rotation by $\tau \times \theta$ with $\theta \approx 14^\circ$. In fact, all of the orbit points between \mathbf{x}_i and \mathbf{x}_{i+1} are mapped in this way by ω^τ , forming a skeleton of the orbit segment between \mathbf{x}_{i+1} and \mathbf{x}_{i+2} . In this way, an approximately self-similar lattice plume, spiralling outward to infinity, is generated. This behaviour is illustrated in figure 7 for the forward orbit of the period-1 point $(0,0)$ in \mathcal{L}_0 .

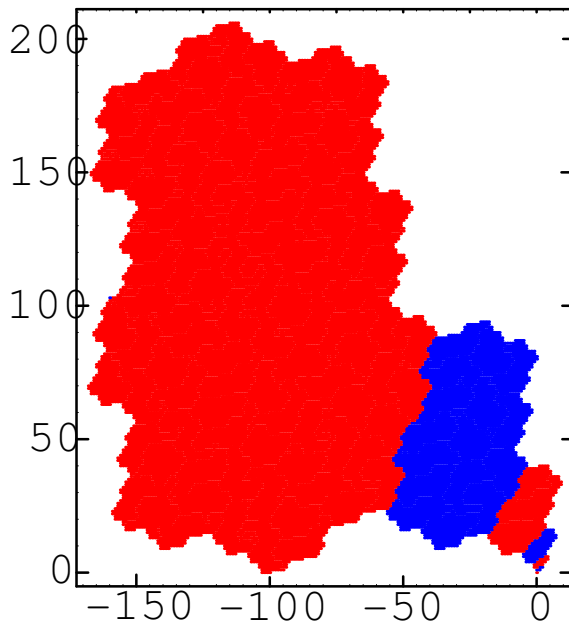


Figure 7: First 23433 points, constituting levels 0 through 5, of the forward orbit of $(0,0)$ in \mathcal{L}_0 . Even- and odd-level points are coloured blue and red, respectively.

4.5 Asymptotic boundary scaling

In section 4.3 we introduced an algorithm for prolonging boundary curves indefinitely. Using exact integer arithmetic, we have used the scheme to calculate the first 10^6 steps of one of these walks, namely the forward orbit $\partial\mathcal{O}_1^{R+}$ of $(0,0)$. The result, sampled at 1000-step intervals, is plotted in figure 8. Since the angular drift is very slow in comparison with the outward expansion, we have plotted the radial coordinate on a logarithmic scale. To quantify the apparent asymptotic fractal property of the boundary curve, we calculate a scaling dimension as follows. Our approach here is to construct heuristically a substitution

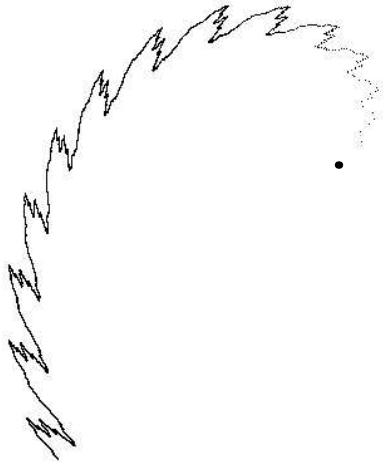


Figure 8: First 10^6 steps of boundary curve $\partial\mathcal{O}_1^{R+}$ are plotted using polar coordinates, with radial distance on a logarithmic scale to accentuate the spiralling outward.

rule, specified in table 7, which accounts for the first 10^6 steps of the boundary curve $\partial\mathcal{O}_1^{R+}$, conjecturing that the rule remains valid for the entire curve. The substitution has the alphabet $\{1, \dots, 42\}$, each symbol being mapped into to the (possibly empty) sequence of symbols specified in the third and sixth columns of the table. Geometrically, each symbol corresponds to a sequence of lattice displacement vectors, designated in the table by their phases ($(0, -1) \leftrightarrow 0$, $(1, -1) \leftrightarrow 1$, etc.). The infinite symbol sequence corresponding to $\partial\mathcal{O}_1^{R+}$ is the concatenation of finite pieces S_1, S_2, S_3, \dots , where $S_1 = \{1, 2\}$ and S_n , $n > 1$, is obtained from S_{n-1} by application of the substitution rule.

Roughly, going from S_n to S_{n+1} corresponds to application of the scaling map ω . The match is not exact because the boundary curve is not invariant with respect to ω . The substitution rule induces a map s on the points of the boundary, such that $\|s(\mathbf{m}) - \omega(\mathbf{m})\|$ is found empirically to be uniformly bounded. Recalling that the largest eigenvalue of the matrix W associated with the scaling map ω has magnitude $\lambda^{-\frac{3}{2}} = 2.29445$, we have that $\|s^n(\mathbf{m})\|$ is of order $\lambda^{-\frac{3}{2}n}$ for $n \rightarrow \infty$. Furthermore, the substitution rule corresponds to a 42×42 transition matrix, whose largest eigenvalue κ governs the rate of growth of the number N_n of boundary steps in S_n , i.e., $N_n = O(\kappa^n)$. Our conjectured scaling dimension D is the ratio of the logarithms of κ and $\lambda^{-\frac{3}{2}}$. For the substitution scheme of table 7, we find that $\kappa = 2.71667\dots$ is a real root of the polynomial $x^4 - 6x^3 + 12x^2 - 8x - 1$, hence

$$D = \frac{\ln \kappa}{\ln \lambda^{-\frac{3}{2}}} = 1.09336.$$

This number is precisely the Hausdorff dimension of the Rauzy fractal boundary calculated

by Ito and Kimura [7] and Messaoudi [11]. Considerable work remains to be done to prove the validity of the plume boundary substitution rule and to clarify its relation to the Rauzy fractal boundary.

Symbol	Translation	Substitution	Symbol	Translation	Substitution
1	3	39, 5	22	2, 2, 3, 2, 2, 2, 2	22, 3, 4, 23, 5, 7, 26
2	4	6	23	3	2, 24
3	3	3, 4, 27	24	3	2, 24
4	4	5, 8	25	2, 2, 3, 2, 2, 2, 1	22, 3, 4, 23, 5, 7, 25
5	3	3, 4, 29, 5	26	2, 2, 3, 2, 2, 2, 2	22, 3, 4, 23, 5, 7, 26
6	4, 5	9, 2, 36, 38, 11	27	3	2, 33
7	2, 2	7	28	5	34, 38, 40
8	4, 5	9, 2, 34, 38, 40	29	3	2, 35
9	4, 5	9, 2, 32, 38, 10	30	6	37, 38, 12
10	5	30, 38, 10	31	3	2, 31
11	5	30, 38, 11	32	6	2, 32, 38, 13
12	6	30, 38, 12	33	3	2, 33
13	6	30, 38, 13	34	6	2, 34, 38, 13
14	7	14, 15, 17, 18	35	3	2, 35
15	0	41, 20	36	6	2, 36, 38, 13
16	6	16	37	6	37, 38, 13
17	7	—	38	6	14, 16
18	7	14, 15, 17, 18	39	3	2, 31
19	2	21	40	5	30, 38, 40
20	1, 0, 7	19, 41, 20	41	0	41, 42
21	2, 2, 3, 2, 2, 2, 1	22, 3, 4, 23, 5, 7, 25	42	1, 0, 2	19, 41, 42

Table 7: List of translation vectors and substitution strings associated with the 42 symbols of our alphabet. The translation vectors are denoted by their phases (see definition in subsection 4.2).

4.6 Census of lattice orbits

In these final paragraphs, we exploit the results of Sec. 4.4 to investigate numerically how many distinct orbits exist for a typical restricted coset $\widehat{\xi + \mathcal{M}}$. Specifically, we count the number of elements of $\gamma^{40}(\mathcal{C})$, where \mathcal{C} is the set of 56 lattice points on the boundary of a square of side 15 centred at the origin. This gives us an upper bound. For the cases where the process has not converged to a unique orbit, we explicitly verify periodicity under further iteration of γ , thus guaranteeing that the number of orbits remains unchanged.

We considered all 23,526,584 restricted cosets in $\widehat{\frac{1}{n}\mathcal{M}}$ with $n \leq 100$. Of these, all but 276 have a unique periodic-code point, corresponding to a unique ρ -orbit in $\widehat{\xi + \mathcal{M}}$. Among the rare exceptions, 2 orbits share the lattice in each of 274 cases, and there are 3 orbits only for \mathcal{L}_0 and \mathcal{L}_η . The data are displayed in table 8. We have no explanation for the overwhelming prevalence of unique orbits.

n	order	orbits	period	cosets
1	1	3	1	1
2	1	3	1	1
4	4	2	4	8
11	10	2	20	20
14	2	2	2	12
22	10	2	20	60
26	4	2	4	24
43	7	2	14	14
58	14	2	14	28
86	7	2	14	14
94	46	2	46	92

Table 8: The cases for which the denominator n (column 1) does not exceed 100 and the number of orbits on \mathcal{L}_ξ (column 3) is larger than 1. Note that the hierarchical code period (column 4) and the number of relevant cosets $\xi + \mathcal{M}$ (column 5) are always multiples of the order $\mu(\xi)$ (column 2).

References

- [1] Akiyama, S., 1999, Self-affine tiling and Pisot numeration system. In *Number theory and its applications*, S. Kanemitsu and K. Györy (eds.) Kluwer Academic Publishers, The Netherlands p. 7–17.
- [2] Arnoux, P. and Rauzy, G., 1991, Représentation géométrique de suites de complexité $2n + 1$. (French) [Geometric representation of sequences of complexity $2n + 1$] *Bull. Soc. Math. France* **119** no. 2, 199–215.
- [3] Arnoux P. and Yoccoz, J., 1981, Construction de difféomorphismes pseudo-Anosov. *C. R. Acad. Sci. Paris* **292** 75–78.
- [4] Boshernitzan, M. D. and Carroll, C. R., 1997, An extension of Lagrange’s theorem. *J. d’Analyse Math.* **72** 21–44.
- [5] Canterini, V. and Siegel, A., 2001 Automate des préfixes-suffixes associé à une substitution primitive (in French. English summary) [Prefix-suffix automaton associated with a primitive substitution], *J. Théor. Nombres Bordeaux* **13**, no. 2, 353–369.
- [6] Holton, C. and Zamboni, L., 1998 Geometric realizations of substitutions, *Bull. Soc. Math. France* **126**, no. 2, 149–179.
- [7] Ito, S. and Kimura, M., 1991, On Rauzy fractal. *Japan J. Indust. Appl. Math.*, **8**, 461–486.
- [8] Kouptsov, K. L., Lowenstein J. H., and Vivaldi, F., 2002, Quadratic rational rotations of the torus and dual lattice maps. *Nonlinearity* **15** 1795–1482.
- [9] Lowenstein J. H., Kouptsov, K. L. and Vivaldi, F., 2004, Recursive tiling and geometry of piecewise rotations by $\pi/7$. *Nonlinearity* **17** 1–25.
- [10] Lowenstein, J. H., Poggiaspalla, G. and Vivaldi, F., 2005, Sticky orbits in a kicked-oscillator model. *Dynamical Systems* **20** 413–451.

- [11] Messaoudi, A., 2000, Frontière du fractal de Rauzy et système de numération complexe. *Acta Arithmetica*, **95**, 195–224.
- [12] Pytheas Fogg, N., 2002. Substitutions in dynamics, Arithmetics and Combinatorics. Springer-Verlag, Berlin.
- [13] Rauzy, G., 1982, Nombre Algébriques et Substitutions. *Bull. Soc. Math. France* **110** 147–178.
- [14] Veech, W., 1984, The metric theory of interval exchange transformations. III. The Sah-Arnoux-Fathi invariant, *Amer. J. Math.* **106**, no. 6, 1389–1422.
- [15] Vershik, A. M., 1985, A theorem on the Markov periodical approximation in ergodic theory. *J.Sov. Math.*, **28**, 667–674.