

Renormalizable two-parameter piecewise isometries

J. H. Lowenstein 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 exhibit two distinct renormalization scenarios for two-parameter piecewise isometries (PWI's), based on $2\pi/5$ rotations of a rhombus and parameter-dependent translations. Both scenarios rely on the recently established renormalizability of a one-parameter triangle map, which takes place if and only if the parameter belongs to the algebraic number field $\mathbb{K} = \mathbb{Q}(\sqrt{5})$ associated with the rotation matrix. With two parameters, features emerge which have no counterpart in the single-parameter model. In the first scenario, we show that renormalizability is no longer rigid: whereas one of the two parameters is restricted to \mathbb{K} , the second parameter can vary continuously over a real interval without destroying self-similarity. The mechanism involves neighbouring atoms which recombine after traversing distinct return paths. We show that this phenomenon also occurs in the simpler context of Rauzy-Veech renormalization of interval exchange transformations, here regarded as parametric piecewise isometries on a real interval. We explore this analogy in some detail.

In the second scenario, which involves two-parameter deformations of a three-parameter rhombus map, we exhibit a weak form of rigidity. The phase space splits into several (non-convex) invariant components, on each of which the renormalization still has a free parameter. However, the foliations of the different components are transversal in parameter space; as a result, simultaneous self-similarity of the component maps requires that both of the original parameters belong to the field \mathbb{K} .

June 13, 2016

Since the late 1970's, renormalization has provided a powerful framework for probing the asymptotic short-distance behavior of dynamical systems. Prominent among the pioneering works on the subject was G. Rauzy's systematic renormalization scheme for one-dimensional mappings known as interval-exchange transformations (IET's) [21]. An IET maps a real interval onto itself, partitioning it into sub-intervals, then permuting the sub-intervals. Rauzy's aim was not simply to study individual renormalizable IET's, but rather to construct large classes of such mappings. His strategy was to consider the sub-interval lengths as continuously variable parameters, on which the condition of renormalizability (i.e., dynamical self-similarity) could be imposed by solving algebraic equations.

It is only recently that a strategy similar to Rauzy's has been successfully applied to the renormalization of mappings in the plane which are close relatives of the one-dimensional IET's [10, 16, 23]. These are the piecewise isometries (PWI's), in which a polygon is partitioned into convex sub-polygons (atoms), which are mapped back into it via translations and rotations. In [16] we constructed two families of renormalizable one-parameter PWI's, including a particularly simple three-atom parametric PWI on a triangle. The present work is a first exploration of renormalization in a two-dimensional parameter space, and includes two quite different routes to renormalizability. One of these involves a degenerate form of self-similarity, in which the values of one of the two parameters are restricted to a field of algebraic numbers, while the other parameter can be varied freely within a finite interval. We show that this phenomenon, in simpler form, plays an important role in the Rauzy renormalization of IET's.

1 Introduction

Piecewise isometries (PWI) are maps of polygonal domains partitioned into convex sub-domains, called atoms, in such a way that the restriction of the map to each atom is an isometry. The first-return map to any convex sub-domain is a new PWI, called the induced PWI on that domain. If by repeating the induction we obtain a sub-system conjugate to the original one via a suitable group of isometries and homotheties, then we consider the original PWI to be renormalizable.

General results for PWI's are scarce [18,19]. All early work on renormalization concerned specific models of PWI's defined over quadratic fields (the field of a PWI is determined by the entries of the rotation matrices and the translation vectors defining the isometries) [1,2,12,14,22]. A more intricate form of renormalization was found in a handful of cubic cases [9,15].

It is only recently that one has been able to deal with the renormalizability, in a more general sense, of continuously deformable PWI's. This is in contrast to the situation for the one-dimensional analogues of PWI's, namely interval exchange transformations (IET's). An IET maps a real interval onto itself via piecewise translations (i.e., permutations) of sub-intervals. In this case, the parameters are a discrete permutation and a continuously deformable array of sub-interval lengths. Rauzy-Veech induction [21,24,25] provides a systematic scheme for parameter-dependent renormalization.

In two dimensions, the first results on parametric families concerned polygon-exchange transformations, for which the isometry is strictly translational. This was work of Hooper [10], on the measure of the periodic and aperiodic sets in a two-parameter family of rectangle-exchange transformations, and Schwartz [23], on the renormalization group of a one-parameter family of polygon-exchange transformations. Subsequently, the present authors [16] studied two one-parameter families of piecewise isometries. Each family has a fixed rotational component defined over a quadratic field ($\mathbb{Q}(\sqrt{5})$ and $\mathbb{Q}(\sqrt{2})$, respectively), and parameter-dependent translations. It was shown that self-similarity occurs if and only if the parameter belongs to the relevant field. In this circumstance, we say that renormalizability is *rigid*.

The mapping for the “pentagonal” model based on $\mathbb{Q}(\sqrt{5})$ is shown in figure 1. For the parameter s restricted to a suitable interval, there is an induced PWI on a triangular sub-domain (the so-called *base triangle*), which reproduces itself after scaling and the reparametrization $s \mapsto r(s)$. After an affine change of parameter, the function r was found to be of Lüroth type—a piecewise affine version of Gauss's map [4,8,17]. In the pentagonal model, the discontinuities of r accumulate at the origin; in the octagonal case ($\mathbb{K} = \sqrt{2}$) one has $r = f \circ f$, where f has two accumulation points of discontinuities. In both cases r is expanding and preserves the Lebesgue measure.

The present work is a first attempt at constructing renormalizable piecewise isometries with more than one deformation parameter. We have identified two quite different scenarios for achieving this goal, both of which depend, at some stage of the induction process, on the known behaviour of the single-parameter base triangle map of the pentagonal model.

In our first construction we start with the rhombus PWI shown in figure




Figure 1: One-parameter rhombus map of the pentagonal model of [16]. The boundary lines a , b , c , d , e , f , are respectively defined by the equations $x_0 = -s$, $x_1 = -s$, $x_0 = 1 - s$, $x_1 = 1 - s$, $x_2 = -s$, $x_2 = -1 - s$, where $x_m = \mathbf{u}_m \cdot \mathbf{x}$ with basis vectors \mathbf{u}_m given in (4), and s a real parameter which varies continuously over a finite interval.

2, whose caption reveals the dependence of the boundary lines on two independent parameters, s and t . Our choice of parameter domain Π , shown in figure 9, has a boundary where $t = s$, where the PWI collapses into the three-atom rhombus map shown in figure 1. It is easy to verify that varying s and t within Π retains the character of the map as a piecewise isometry. The question of renormalizability is more subtle. As we shall see in section 4, after three successive induction steps, one arrives at an induced PWI on an isosceles triangle which, after eliminating superfluous atom boundaries, reduces to precisely the base triangle map parametrized by s . The criterion of self-similarity thus leads to a foliation of the parameter domain Π : we have rigidity in the s direction (s must be an algebraic number in \mathbb{K}), but continuous variation in the t direction.

The appearance of a hidden reduction of the number of atoms, which reveals itself only after induction, is present in a much simpler form in Rauzy-Veech induction of interval exchange transformations. In section 5, we discuss this phenomenon in detail, showing that it leads to considerable simplification in the classification scheme for renormalizable IET's. In particular, there are always so-called *proper* self-similar IET's, without free parameters, as long as the number of intervals is greater than three. (At a deeper level, the occurrence of non-rigidity in renormalizable IET's can be understood in terms of the properties of the associated translation surfaces [11].)

In view of the one-dimensional analogue, we speculate that there may exist renormalizable two parameter PWI's which are proper, in the sense




Figure 2: Two-parameter rhombus map. The boundary lines a, b, c, d, e, f, g are, respectively, defined by the equations $x_0 = -t, x_1 = -s, x_0 = 1 - t, x_1 = 1 - s, x_2 = -s, x_2 = -1 - s, x_1 = -t$, where $x_m = \mathbf{u}_m \cdot \mathbf{x}$ and the basis vectors \mathbf{u}_m are given in (4). Here s and t are real parameters, with (s, t) varying continuously over the convex domain shown in figure 9.

that their renormalizability does not derive from that of a one-parameter induced map. With this in mind, we conducted an extensive search for a rigidly renormalizable two-parameter PWI, using two-dimensional sections of the three-dimensional parameter space of a four-atom rhombus shown in figure 13. We only found a weak form of rigidity, resulting from the co-existence, in disjoint (non-convex) regions of the phase space, of several one-parameter induced PWI's with transversally intersecting foliations of the parameter domain. This provides a second scenario for two-parameter renormalizability, which exhibits rigidity in both parameters without being genuinely proper. One such model is studied in detail in section 6. Whether proper renormalizable two-parameter PWI's exist at all remains an open question.

The remainder of this article is organized as follows. After introducing our basic definitions and methods in section 2, we review in section 3 the renormalization of the one-parameter base triangle. In section 4, we introduce our first scenario for two-parameter renormalization, showing how atom recombination leads to non-rigidity of the self-similarity. As we explain in section 5, the same mechanism plays an important role in Rauzy-Veech renormalization of interval exchange transformations in one dimension. In

the final section 6, we exhibit a four-atom rhombus map which exemplifies our second scenario, in which weak rigidity of both parameters is achieved, but at the expense of a partition of the phase space into several invariant components.

ACKNOWLEDGEMENTS: JHL and FV would like to thank, respectively, the School of Mathematical Sciences at Queen Mary, University of London, and the Department of Physics of New York University, for their hospitality.

2 Preliminaries

Throughout this paper, we let

$$\alpha = \sqrt{5}, \quad \omega = (\alpha + 1)/2, \quad \beta = \omega^{-1} = (\alpha - 1)/2. \quad (1)$$

The arithmetical environment is the quadratic field $\mathbb{Q}(\omega)$ with its ring of integers $\mathbb{Z}[\omega]$, given by

$$\mathbb{Q}(\omega) = \{x + y\alpha : x, y \in \mathbb{Q}\}, \quad \mathbb{Z}[\omega] = \{m + n\omega : m, n \in \mathbb{Z}\}. \quad (2)$$

The number ω , which is the fundamental unit in $\mathbb{Z}[\omega]$ (see [6, chapter 6]), will determine the scaling under renormalization. The number $\beta = \omega - 1$ is also a unit.

2.1 Planar objects

A *tile* X with n edges is a convex polygon defined by the half-plane conditions

$$\begin{aligned} \mathbf{u}_{m_i} \cdot \mathbf{x} &< b_i \quad (\text{excluded edge}) \\ &\text{or} \\ \mathbf{u}_{m_i} \cdot \mathbf{x} &\geq b_i \quad (\text{included edge}) \end{aligned} \quad i = 1, \dots, n, \quad (3)$$

where $\mathbf{x} = (x, y)$, $b_i \in \mathbb{R}$, and the \mathbf{u}_m are the vectors

$$\mathbf{u}_m = \left(\cos \frac{2\pi m}{5}, \sin \frac{2\pi m}{5} \right) \quad m \in \{0, \dots, 4\}. \quad (4)$$

For the i th edge, defined by $\mathbf{u}_{m_i} \cdot \mathbf{x} = b_i$, we introduce an index ϵ_i , where $\epsilon_i = -1$ if the edge is included in X , and $\epsilon_i = 1$ if it is excluded. We then represent X as a triple of n -vectors

$$X = [(m_1, \dots, m_n), (\epsilon_1, \dots, \epsilon_n), (b_1, \dots, b_n)]. \quad (5)$$

We shall assume that n is minimal, namely that X is not definable by fewer conditions.

A *tiling* \mathbf{X} is a set of disjoint tiles,

$$\mathbf{X} = \{X_1, \dots, X_N\}$$

and is associated with a domain X (union of tiles)

$$X = \bigcup_{k=1}^N X_k.$$

Note that a domain need not be convex, or even connected. Note further that thanks to (3), if a pair of tiles have disjoint interiors but share a common boundary segment, that segment belongs to one and only one tile of the pair. This allows the possibility of gluing together adjacent tiles without disturbing the inclusion relation of the respective edges.

2.2 Similarity group

The transformation properties of planar objects are provided by a group \mathfrak{G} which comprises the rotations and reflections of the symmetry group of the regular pentagon (the dihedral group D_5) together with translations in \mathbb{K}^2 and real scale transformations.

We adopt the following notation:

U_m : reflection about the line generated by \mathbf{u}_m .

R_m : rotation by the angle $2m\pi/5$.

$T_{\mathbf{d}}$: translation by $\mathbf{d} \in \mathbb{K}^2$.

S_η : scaling by $\eta \in \mathbb{R}_+$.

We write $\mathcal{X} \sim \mathcal{Y}$ to indicate that \mathcal{X} is similar to \mathcal{Y} , i.e., that $\mathcal{X} = \mathbf{G}(\mathcal{Y})$ for some $\mathbf{G} \in \mathfrak{G}$. As \mathfrak{G} is a group, this is an equivalence relation. Within \mathfrak{G} we distinguish two important subgroups: the *isometry group* \mathfrak{I} generated by rotations, reflections, and translations, and the *dynamical group* \mathfrak{I}_+ , generated by rotations and translations.

2.3 Dressed domains and sub-domains

A *dressed domain* is a triple

$$\mathcal{X} = (X, \mathbf{X}, \rho), \quad (6)$$

where $\mathbf{X} = \{X_1, \dots, X_N\}$ is a tiling of the domain X , and $\rho = \{\rho_1, \dots, \rho_N\}$, where $\rho_i \in \mathfrak{J}_+$ is an orientation-preserving isometry acting on the tile X_k . Under the action of $\mathbf{G} \in \mathfrak{G}$, a dressed domain \mathcal{X} transforms as

$$\mathbf{G}(\mathcal{X}) = \mathbf{G}(X, \mathbf{X}, \rho) = (\mathbf{G}(X), \{\mathbf{G}(X_1), \dots, \mathbf{G}(X_k)\}, \mathbf{G} \circ \rho \circ \mathbf{G}^{-1})$$

where the conjugacy acts component-wise. To emphasize the association of a mapping ρ with a particular dressed domain \mathcal{X} , we use the notation $\rho_{\mathcal{X}}$.

Let $\mathcal{X} = (X, \mathbf{X}, \rho_{\mathcal{X}})$ be a dressed domain, and let Y be a sub-domain of X . We denote by ρ_Y the first-return map on Y induced by $\rho_{\mathcal{X}}$. We call the resulting dressed domain $\mathcal{Y} = (Y, \mathbf{Y}, \rho_Y)$ a *dressed sub-domain* of \mathcal{X} , and write

$$\mathcal{Y} \triangleleft \mathcal{X}. \quad (7)$$

The dressed sub-domain relation (7) is *scale invariant*, namely invariant under an homothety. Indeed, if \mathbf{S}_η denotes scaling by a factor η , then in the data (5) specifying a tile, the orientations m_k remain unchanged, while the pentagonal coordinates b_k scale by η . Moreover, the identity

$$\mathbf{S}_\eta \mathbf{T}_d \mathbf{R}_n = \mathbf{T}_{\eta d} \mathbf{R}_n \mathbf{S}_\eta.$$

shows that the piecewise isometries ρ scale in the same way. We conclude that the sub-domain relation (7) is preserved if the dressed domain parameters are scaled by the same factor for both members.

2.4 Parametric dressed domains

In this article we consider continuously deformable dressed domains $\mathcal{X} = \mathcal{X}(s)$ called *parametric dressed domains*, depending on a real parameter vector $s = (s_1, \dots, s_p)$.

These are domains whose tiles X_k and image tiles $\rho_k(X_k)$ depend on s only via the coefficients b_i , while the parameters n , m_i and ϵ_i remain fixed [see (5)]. We shall require that the b_i 's be affine functions of s_1, \dots, s_p , with coefficients in $\mathbb{Q}(\omega)$. Algebraically, this is expressed as

$$b_i \in \mathbb{S} \quad \mathbb{S} = \mathbb{Q}(\omega) + \mathbb{Q}(\omega)s_1 + \dots + \mathbb{Q}(\omega)s_p, \quad (8)$$

where s_1, \dots, s_p are regarded as indeterminates. The set \mathbb{S} is a $(p + 1)$ -dimensional vector space over $\mathbb{Q}(\omega)$ (a $\mathbb{Q}(\omega)$ -module).

The condition (8) gives us affine functions $b_i : \mathbb{R}^p \rightarrow \mathbb{R}$

$$b_i(s_1, \dots, s_p) = b_{i,0} + b_{i,1}s_1 + \dots + b_{i,p}s_p \quad b_{i,j} \in \mathbb{Q}(\omega). \quad (9)$$

We define the *bifurcation-free set* $\mathbf{\Pi}(\mathcal{X})$ to be the maximal open set such that all of the edges of all $X_k(s)$ have non-zero lengths. Note that other types of bifurcations may occur if \mathcal{X} is embedded within a larger domain (see section 2.6.)

2.5 Renormalizable dressed domains

A parametric dressed domain $\mathcal{X}(s)$ is said to be *renormalizable* over an open domain $\Pi \subset \mathbb{R}^p$ if there exists a piecewise smooth map $r : \Pi \rightarrow \bar{\Pi}$ such that for every choice of $s \in r^{-1}(\Pi)$ the dressed domain $\mathcal{X}(s)$ has a dressed sub-domain \mathcal{Y} similar to $\mathcal{X}(r(s))$ which satisfies the recursive tiling property. The function r depends only on s , a requirement of scale invariance. In general, we have $\mathcal{Y} = \mathcal{Y}_{i(s)}(s)$, where i is a discrete index. The set $r^{-1}(\Pi)$ need not be connected (even if i is constant), each connected component being a bifurcation-free domain of \mathcal{Y} . (To extend the renormalization function r to the closure of Π , one must include bifurcation parameter values, as in [16].)

If $s = s_0$ is eventually periodic under r , then we say that $\mathcal{X}(s_0)$ is *self-similar*. A self-similar system has an induced sub-system which reproduces itself on a smaller scale under induction.

Let a parametric dressed domain $\mathcal{X}(s)$ have induced $\mathcal{X}_j(s)$, such that, for $j = 1, \dots, N$ we have: (i) \mathcal{X}_j is renormalizable over a domain Π_j ; (ii) the \mathcal{X}_j recursively tile \mathcal{X} ; (iii) the Π_j have non-empty intersection Π . Then we still consider \mathcal{X} renormalizable over Π .

The definition of renormalizability given above is tailored to our model; it is not the most general possible, and it is local in parameter space. We allow \mathcal{Y} to depend on a discrete index (as in Rauzy induction for interval exchange transformations — see section 5) to obtain a simpler renormalization function r (section 3). We only require \mathcal{X} to be eventually renormalizable, and we allow \mathcal{X} to have sub-domains with independent renormalization schemes (which is a common phenomenon, see section 6).

2.6 Computations

For computations, we use the Mathematica[®] procedures listed in the supplemental material [7]. All computations reported in this work are exact,

employing integer and polynomial arithmetic, and the symbolic representation of algebraic numbers.

The geometrical objects defined in section 2.1 require arithmetic in a bi-quadratic field, since only the first component of the vectors \mathbf{u}_m is in $\mathbb{Q}(\omega)$. To circumvent this difficulty, we conjugate our PWI to a map of a square where the clockwise rotation $2\pi/5$ is represented by the following matrix over $\mathbb{Z}[\omega]$

$$\begin{pmatrix} 0 & 1 \\ -1 & \beta \end{pmatrix}$$

where β was defined in (1). (This is still a PWI with respect to a non-Euclidean metric.) In the new co-ordinates, the vectors \mathbf{u}_m become

$$\{(1, 0), (0, 1), (-1, \beta), (-\beta, -\beta), (\beta, -1)\} \quad (10)$$

which belong to $\mathbb{Z}[\omega]^2$. With this representation, all of our calculations can be performed within the module \mathbb{S} defined in (8). As an example, figure 3 displays the rhombus map of figure 2 using the basis (10). We shall still




Figure 3: The two-parameter map of figure 2 represented using the basis vectors of (10). The equations for the boundary lines are the same as those listed under figure 2. The mapping is still piecewise isometric, but with respect to a non-Euclidean metric.

display our figures in the original coordinates, where geometric relations (especially reflection symmetries) are more apparent to the eye.

In constructing a return map orbit of a domain $\mathcal{X}(s)$ by direct iteration, one determines inclusion and disjointness relations among domains, which

require evaluations of inequalities (3). Since the latter are expressed by affine functions of the parameter s in some polytope Π , it suffices to check the inequalities on the boundary of Π . All these boundary values belong to the field $\mathbb{Q}(\omega)$, and the inequalities can be reduced to integer inequalities.

Typically, \mathcal{X} will be immersed in a larger domain \mathcal{Y} (an atom, say). Therefore, in addition to the intrinsic bifurcation-free polytope $\Pi(\mathcal{X})$ defined in section 2.4, one must also consider the polytope $\Pi(\mathcal{X}, \mathcal{Y})$ defined by the inclusion $\mathcal{X}(s) \subset \mathcal{Y}(s)$, as well as intersection of these polytopes.

The recursive tiling property defined in section 2.5 is established by adding up the areas of the tiles of all the orbits, and comparing it with the total area of the parent domain.

With these techniques, we are able to establish rigorously a variety of statements valid over convex sets in parameter space.

3 Base triangle

The *base triangle* is the simplest one-parameter renormalizable piecewise isometry associated with rotations by $2\pi/5$; it is self-similar precisely for parameter in the quadratic field $\mathbb{Q}(\sqrt{5})$. It was instrumental to the proof of renormalizability of a one-parameter rhombus map in [16], and it will appear again in the many-parameter versions presented here.

We develop a variant of the model presented in [16], which includes boundary segments in the tiles and an improved renormalization scheme. The base triangle \mathcal{B} prototype is the following dressed domain (see figure 4):




Figure 4: Base triangle prototype.

$$\mathcal{B} = (B, (B_1, B_2, B_3), (\rho_1, \rho_2, \rho_3))$$

where

$$\begin{aligned} B &= [(1, 0, 2), (-1, 1, 1), (\tau - \omega^2, 0, 0)], \\ B_1 &= [(0, 2, 3), (1, 1, 1), (0, 0, \omega - \omega\tau)], \\ B_2 &= [(1, 4, 0, 3, 2), (-1, 1, 1, -1, 1), (\tau - \omega^2, \omega^2 - \omega\tau, 0, \omega - \omega\tau, 0)], \\ B_3 &= [(1, 0, 4), (-1, 1, -1), (\tau - \omega^2, 0, \omega^2 - \omega\tau)]. \end{aligned} \tag{11}$$

The dynamics is given by a local reflection of each atom about its own symmetry axis, followed by a global reflection about the symmetry axis of B , which can be written as:

$$\begin{aligned} \rho_1 &= \mathbb{T}_{(\omega\tau - \omega, -\omega^2 + \omega^2\tau)} \mathbb{R}_2 \\ \rho_2 &= \mathbb{T}_{(0, \omega^2\tau - \omega^3)} \mathbb{R}_3 \\ \rho_3 &= \mathbb{T}_{(\omega\tau - 2\omega, \omega^2\tau - 2\omega^2)} \mathbb{R}_2. \end{aligned} \tag{12}$$

Here we have chosen a coordinate system such that the peak of the isosceles triangle is at the origin and the altitude of the atom B_3 is the parameter τ , which varies over the interval $(0, 1)$ without the occurrence of a bifurcation. This parameter (together with time-reversal invariance) determines the scale-invariant properties of the dressed domain, since it is related to the ratio η of altitudes of B_3 and B by the formula

$$\eta = \frac{\tau}{\omega^2 - \tau}.$$

As τ varies from 0 to 1, η increases from 0 to β .

The edges of the domain \mathcal{B} are included or excluded as stipulated in section 2; a vertex joining two included edges is included, but is excluded otherwise. The renormalizability analysis will also require a second base triangle $\tilde{\mathcal{B}}$, differing from \mathcal{B} by a change of sign of all edge coordinates and translation vectors, as well as of the respective ϵ_i . The dressed domains \mathcal{B} and $\tilde{\mathcal{B}}$ are \mathfrak{G} -inequivalent: not only do they have different boundary conditions, but their interiors differ by a rotation by π , not an element of the similarity group.

The renormalizability analysis for the base triangle is summarized in the following lemma:

Lemma 1 *Let \mathcal{B} be as above. The following holds:*

- (i) *For $0 < \tau < \beta^2$, \mathcal{B} has a dressed sub-domain $\mathcal{B}_1 \sim \mathcal{B}$ which is scaled by a factor $(1 - \tau)/(\omega^2 - \tau)$ and has shape parameter $r(\tau) = \omega^2\tau$.*

- (ii) For $\beta^2 < \tau < \beta$, \mathcal{B} has a dressed sub-domain $\mathcal{B}_2 \sim \tilde{\mathcal{B}}$ which is scaled by a factor $\tau/(\omega^2 - \tau)$ and has shape parameter $r(\tau) = \omega^3(\beta - \tau)$.
- (iii) For $\beta < \tau < 1$, \mathcal{B} has a dressed sub-domain $\mathcal{B}_3 \sim \tilde{\mathcal{B}}$ which is scaled by a factor $\tau/(\omega^2 - \tau)$ and has shape parameter $r(\tau) = \omega^2(1 - \tau)$.




Figure 5: Renormalization function $r(\tau)$ for base triangles.

The renormalization function r has three branches (see figure 5). In cases (ii) and (iii) one induces on the atom B_3 , over two disjoint bifurcation-free parameter ranges. Since the size of B_3 vanishes as τ approaches 0, in the range (i) we induce on the triangle $[(1, 0, 2), (-1, 1, 1), (\beta^2\tau - 1, 0, 0)]$, which is not an atom. This device prevents the occurrence of infinitely many singularities in the renormalization function found in [16].

In each case, the return orbits of the atoms of \mathcal{B}_i , together with a finite number of periodic tiles, completely tile the triangle B . For $\tilde{\mathcal{B}}$, the prescriptions (i)-(iii) hold with the roles of \mathcal{B} and $\tilde{\mathcal{B}}$ exchanged. The induction relations are represented as the graph in figure 6.

As in [16], the proof of Lemma 1 is by direct iteration, as discussed in section 2.6. The main difference in the two computational algorithms lies in the procedures used to verify inclusion and disjointness of tiles (see [7]). Specifically, checking the sub-polygon relation $X \subset Y$ requires verifying that no included vertex of X has landed on an excluded edge of Y . Similarly, to




Figure 6: Renormalization graph for base triangles. A directed link from X to Y indicates that X has an induced dressed sub-domain equivalent to Y , subject to the the indicated parameter constraint.

decide that X and Y are disjoint, one must check that no included vertex of either tile lies on an included edge of the other.

With reference to lemma 1, we remark that the base triangle \mathcal{B} , with its definition extended to the two-atom limiting cases $\tau = 0, 1$, is in fact renormalizable also at the parameter values $0, \beta^2, \beta, 1$, with $r(\tau) = 0$ in all these cases (for the calculations, see supplemental material [7]. These additional parameter values are needed to make \mathcal{B} renormalizable over the whole interval $[0, 1] \cap \mathbb{Q}(\omega)$. We shall use this property in sections 4 and 6.

4 Non-rigid self-similarity

We now turn to the two-parameter rhombus map introduced in section 1 —see figures 2 and 3. In suitable coordinates, the dressed domain is given by

$$\mathcal{R} = (R, (R_1, \dots, R_5), (\rho_{R_1}, \dots, \rho_{R_5})),$$

with (see figure 2)

$$\begin{aligned} R &= [(0, 1, 0, 1), (-1, -1, 1, 1), (-t, -s, 1-t, 1-s)], \\ R_1 &= [(0, 2, 1), (-1, -1, 1), (-t, -s, 1-s)], \\ R_2 &= [(0, 1, 2, 0, 1, 2), (-1, -1, -1, 1, 1, 1), (-t, -t, -1-s, 1-t, 1-s, -s)], \\ R_3 &= [(0, 2, 1), (1, 1, -1), (1-t, -1-s, -t)], \\ R_4 &= [(0, 1, 2, 1), (-1, -1, -1, 1), (-t, -s, -1-s, -t)], \\ R_5 &= [(1, 0, 1, 2), (-1, 1, 1, 1), (-s, 1-t, -t, -1-s)] \end{aligned} \quad (13)$$

$$\rho_{R_1} = \mathbb{T}_{(0,0)} R_4, \quad \rho_{R_2} = \mathbb{T}_{(0,1)} R_4, \quad \rho_{R_3} = \mathbb{T}_{(0,2)} R_4, \quad (14)$$

$$\rho_{R_4} = \mathbb{T}_{(1,1)} R_4. \quad \rho_{R_5} = \mathbb{T}_{(1,2)} R_4.$$

The corresponding bifurcation-free parametric domain $\Pi(\mathcal{R})$, defined in section 2.4, is found to be the triangle with vertices at $(0, 0)$, $(-1/\alpha, -1/\alpha)$,

and $(\beta/\alpha, -\beta/\alpha)$. On the boundary of $\Pi(\mathcal{R})$ given by with $s = t$, the dressed domain \mathcal{R} collapses into the one-parameter pentagonal model of [16], and hence is self-similar for all $s \in \mathbb{Q}(\omega)$ within a suitable interval.




Figure 7: The first two steps of the triple induction $\mathcal{R} \triangleright \mathcal{R}_3 \triangleright \mathcal{R}_{31} \triangleright \mathcal{A}$. The third step produces the dressed domain \mathcal{A} shown in figure 8.

Our goal is to determine a parametric dressed domain \mathcal{A} with bifurcation-free sub-domain $\Pi(\mathcal{A}) \subset \Pi(\mathcal{R})$ over which \mathcal{R} is renormalizable. To this end, we choose a specific parameter pair close to the $s = t$ boundary: $(s_0, t_0) = (-19/200, -1/10)$, and we establish that at this value the renormalization is amenable to a complete analysis (with computer assistance). Specifically, we consider a three-step induction, first on the triangular atom R_3 , followed by two inductions on sub-triangles, as shown in figure 7. The last induction produces the dressed domain \mathcal{A} , shown in figure 8, which is given by:

$$\mathcal{A} = (A, (A_1, \dots, A_8), (\rho_{A_1}, \dots, \rho_{A_8})), \quad (15)$$

with

$$\begin{aligned} A &= [(2, 1, 0), (1, -1, 1), (-1 - s, \beta^4 - s, 1 - s)] \\ A_1 &= [(1, 0, 4), (-1, 1, -1), (\beta^4 - t, 1 - s, \alpha\beta^4 - t)], \\ A_2 &= [(1, 4, 0, 4), (-1, 1, 1, -1), (\beta^4 - t, \alpha\beta^4 - t, 1 - s, \alpha\beta^4 - s)], \\ A_3 &= [(0, 1, 4, 1), (1, 1, -1, -1), (1 - s, \beta^4 - t, \alpha\beta^4 - t, \beta^4 - s)], \\ A_4 &= [(4, 1, 4, 1), (1, 1, -1, -1), (\alpha\beta^4 - t, \beta^4 - t, \alpha\beta^4 - s, \beta^4 - s)], \\ A_5 &= [(2, 1, 4, 0, 3), (1, -1, 1, 1, -1) \\ &\quad (-1 - s, \beta^4 - t, \alpha\beta^4 - s, 1 - s, -4\beta^3 - t)], \end{aligned} \quad (16)$$

$$\begin{aligned}
A_6 &= [(2, 3, 0, 3), (1, 1, 1, -1), (-1 - s, -4\beta^3 - t, 1 - s, -4\beta^3 - s)], \\
A_7 &= [(2, 1, 4, 1), (1, -1, 1, 1), (-1 - s, \beta^4 - s, \alpha\beta^4 - s, \beta^4 - t)], \\
A_8 &= [(2, 3, 0), (1, 1, 1), (-1 - s, -4\beta^3 - s, 1 - s)],
\end{aligned}$$

$$\begin{aligned}
\rho_{A_1} = \rho_{A_2} = \rho_{A_3} = \rho_{A_4} &= \mathbf{T}_{(8\beta^3, -2\beta^5)} \mathbf{R}_2, \\
\rho_{A_5} = \rho_{A_6} = \rho_{A_7} &= \mathbf{T}_{(2, 1+\beta^5)} \mathbf{R}_3, \\
\rho_{A_8} &= \mathbf{T}_{(2-\beta^6, -\beta^5)} \mathbf{R}_2.
\end{aligned} \tag{17}$$




Figure 8: The dressed domain \mathcal{A} , with its 8 atoms numbered as in (15). The boundaries of the composite atoms C_1, C_2, C_3 are coloured red, green, and blue, respectively.

We find that $\Pi(\mathcal{A})$ is the triangle with vertices

$$\left(-1 + \frac{2}{\alpha}, -1 + \frac{2}{\alpha}\right), \quad \left(\frac{1}{2}(11 - 5\alpha), \frac{1}{4}(-25 + 11\alpha)\right), \quad \left(\frac{1}{2}(11 - 5\alpha), \frac{1}{2}(11 - 5\alpha)\right),$$

shown in figure 9. One verifies that $\Pi(\mathcal{A})$ is adjacent to the line $s = t$ and that (s_0, t_0) lies in its interior.

Using direct iteration, we verify that for all $(s, t) \in \Pi(\mathcal{A})$ the return orbits of the eight atoms of \mathcal{A} , together with those of 13 periodic tiles, completely tile the rhombus R (see figure 10).

A decisive simplification of the analysis results from the observation that the four atoms A_1, \dots, A_4 of \mathcal{A} are mapped by the same isometry, and hence, with regard to the first-return map to \mathcal{A} , can be merged into a single triangular tile, A_{1234} . Similarly, atoms A_5, A_6, A_7 can be merged into a single




Figure 9: The parametric domain $\Pi(\mathcal{A})$.

reflection-symmetric pentagon, A_{567} . The mergers have been suggested in the shading of the tiles in figure 8. The dressed domain thus simplifies into

$$\begin{aligned} \mathcal{C} &= (C, (C_1, C_2, C_3), (\rho_{C_1}, \rho_{C_2}, \rho_{C_3})) \\ &\stackrel{\text{def}}{=} (A, (A_8, A_{567}, A_{1234}), (\rho_{A_8}, \rho_{A_5}, \rho_{A_1})) \end{aligned} \quad (18)$$

Moreover, one verifies that over $\Pi(\mathcal{C}) = \Pi(\mathcal{A})$, we have $\mathcal{C} \sim \mathcal{B}$. The intrinsic shape parameter of \mathcal{C} can be calculated from the ratio $\eta_{\mathcal{C}}$ of the altitude of C_3 to that of C :

$$\tau_{\mathcal{C}} = \frac{\omega^2 \eta_{\mathcal{C}}}{1 + \eta_{\mathcal{C}}} = \omega^7 (\alpha s + \beta^3). \quad (19)$$

As we transverse $\Pi(\mathcal{C})$ from left to right, s increases from $2/\alpha - 1$ to $(11 - 5\alpha)/2$, with $\tau_{\mathcal{C}}$ increasing from 0 to 1.

The issue of recursive tiling is now rather subtle. The rhombus is certainly tiled by the return orbits of C_1, C_2, C_3 , and the 13 periodic tiles which arose in the induction on \mathcal{A} . (see figure 10). However, the return paths are not the same for all tiles. In (say) C_3 , the tiles A_1, A_2, A_3, A_4 have four distinct 78-step return paths, which go their separate ways, but recombine eventually to form an atom of the dressed domain with a unique isometry. The coincidence of the return times is not necessary to the recombination, as these times could differ by any integer multiples of 5. As a result, the partition of C_3 into A_1, A_2, A_3, A_4 is relevant to the recursive tiling of the original rhombus, but not to dynamical self-similarity. (We shall encounter again the same phenomenon —recombination with different return times— in section 6.)




Figure 10: Tiling of \mathcal{R}_3 by return orbits of the 8 atoms of \mathcal{A} (coloured) and 7 periodic tiles (grey). Note that A_1, A_2, A_3, A_4 , which comprise C_1 , have distinct return orbits.

The parameter pairs $(s, t) \in \Pi(\mathcal{C})$ corresponding to self-similarity for the rhombus map \mathcal{R} are now determined by the self-similarity of the induced dressed sub-domain \mathcal{C} . In turn, the latter are the values of s for which the base triangle is self-similar, namely $\tau_{\mathcal{C}}(s) \in (0, 1) \cap \mathbb{Q}(\omega)$, while t is unconstrained. Thus, by (19), \mathcal{R} is renormalizable in $\Pi(\mathcal{C})$ if and only if

$$(s, t) \in \Pi(\mathcal{C}) \cap (\mathbb{Q}(\omega) \times \mathbb{R}).$$

This is the main result of this section.

5 Non-rigid self-similarity in Rauzy induction

The merging and recombination of atoms, resulting in the appearance of a hidden free parameter in the renormalization process, is not peculiar to planar PWI's. In fact, this phenomenon is analogous to that encountered in the simpler, one-dimensional setting of Rauzy-Veech renormalization of interval exchange transformations (IET's) [21, 24, 25]. For completeness, we first review the basic constructs.

We fix a half-open interval $\Omega = [0, l)$ and a partition of Ω into n half-open sub-intervals Ω_i . An IET is a piecewise isometry of Ω which is a translation on each Ω_i . We represent it as a pair (π, Λ) , where $\Lambda = (\lambda_1, \dots, \lambda_n)$ is the vector of the lengths of the sub-intervals and π is the permutation of $\{1, \dots, n\}$ such that the transformed interval length vector is $(\lambda_{\pi^{-1}(1)}, \dots, \lambda_{\pi^{-1}(n)})$ (i.e., the interval of length λ_k is translated from position k to position $\pi(k)$ in the array).

We assume that π is *irreducible* in the sense that $\{1, \dots, k\}$ is mapped into itself only if $k = n$. If we fix n , then (π, Λ) is a parametric PWI, with discrete and continuous parameters π and Λ , respectively. IET's which differ only by an overall translation or scale transformation are considered equivalent.

Rauzy-Veech induction on (π, Λ) consists of inducing on the larger of the two intervals $\Omega^{(0)} = [0, l - \lambda_n)$ and $\Omega^{(1)} = [0, l - \lambda_{\pi^{-1}(n)})$, denoted by type 0 and type 1 induction, respectively (the case $|\Omega^{(0)}| = |\Omega^{(1)}|$ is excluded from consideration, as in this case the map is not minimal). Induction corresponds to a map $(\pi, \Lambda) \mapsto (\pi', \Lambda') = (a_i(\pi), A_i(\pi)^{-1}\Lambda)$, $i = 0, 1$, where $A_i(\pi)^{-1}$ is an $n \times n$ integral matrix (see [20]), and $a_i(\pi)$ is defined as follows:

$$\begin{aligned} a_0((k_1, \dots, k_{j-1}, n, k_{j+1}, \dots, k_{n-1}, m)) \\ = (k'_1, \dots, k'_{j-1}, m+1, k'_{j+1}, \dots, k'_{n-1}, m), \end{aligned} \tag{20}$$

where $k'_i = k_i$ if $k_i < m$ and $k'_i = k_i + 1$ if $k_i > m$, and

$$\begin{aligned} a_1((k_1, \dots, k_{j-1}, n, k_{j+1}, \dots, k_{n-1}, m)) \\ = (k_1, \dots, k_{j-1}, n, m, k_{j+1}, \dots, k_{n-1}), \end{aligned} \tag{21}$$

with the obvious interpretations if either of the sequences k_1, \dots, k_{j-1} or k_{j+1}, \dots, k_{n-1} is empty or consists of a single member.

The permutations π of n symbols are then represented as the vertices of a *Rauzy graph*. Each vertex has two outgoing and two incoming edges, associated with a_i and a_i^{-1} , respectively, for $i = 0, 1$. The *Rauzy classes*

are the connected components of the graph. These IET's are linked by a sequence of Rauzy inductions, and a self-similar IET corresponds to a path on a Rauzy graph which terminates in a closed circuit γ of permutations $\pi_1, \pi_2, \dots, \pi_p$. Transversing such a circuit produces a rescaled version of the initial IET if its length vector is an eigenvector of the product A_γ^{-1} of the matrices $A_{i_k}(\pi_k)^{-1}$, $k = 1, \dots, p$, with a scale factor given by the corresponding eigenvalue. In the language of section 2.5, the irreducible IET's of length n form a parametric dressed domain whose tiles are the sub-intervals and whose parameters are π and $\Lambda = (\lambda_1, \dots, \lambda_n)$. The renormalization function corresponding to a particular permutation π and Rauzy-graph circuit originating at π is just

$$r_\gamma(\Lambda) = \frac{A_\gamma^{-1} \cdot \Lambda}{|A_\gamma^{-1} \cdot \Lambda|}.$$

As usual, self-similarity corresponds to a fixed-point of r_γ , i.e., an eigenvalue condition for the matrix A_γ .

In figures 11 and 12 we display the Rauzy graphs for $n = 3$ and $n = 4$ [25, section 6]. Two of the three irreducible permutations of the Rauzy graph for $n = 3$ (figure 11) and all of the permutations for the first class for $n = 4$ (figure 12) contain a consecutive pair of the form $(\dots, j, j + 1, \dots)$. Such an IET will be called *degenerate*. For those IET's, the consecutive intervals Ω_j and Ω_{j+1} have the same translation vector, and hence the interval $\Omega_j \cup \Omega_{j+1}$ may be merged into a single interval of length $\lambda'_j = \lambda_j + \lambda_{j+1}$. The n -interval IET is thus identical, as a mapping, to an $n - 1$ -interval IET.

In figures 11 and 12 we see examples of simplification via atom merger, even though this is not necessarily true recombination. For the latter, we require that the degeneracy appear only after at least one step of induction, to ensure that the members of a consecutive pair have distinct return paths. The simplest example of a recombination is the one with permutation (321).




Figure 11: Rauzy graph for $n = 3$. All permutations are degenerate, and their translation surface is a torus. Any renormalizable IET with three intervals will have a free parameter.




Figure 12: Rauzy graph for $n = 4$. All permutations of the upper component are degenerate. No permutation of the lower component has this property.

This IET is clearly not degenerate, but each of its Rauzy-induced children is. In both cases the two members of the consecutive pair have distinct return paths, and even distinct return times. Now the renormalization analysis reduces naturally to that a single IET with $\pi = (12)$. Once again, one of the parameters of the original IET can be varied continuously without disturbing strict renormalizability.

For general n , we define as *recombinant* any IET which, under Rauzy induction, has the property that each of its children has a permutation with a consecutive pair of intervals having distinct return paths. Furthermore, we define as *improper* an IET which is degenerate or of the form

$$(\dots, n, k + 1, \dots, k)$$

for some $k < n - 1$. An IET which is not improper is defined to be *proper*. By a straightforward application of the definitions (20) and (21), one obtains the following results:

- (i) A nondegenerate, improper IET is recombinant.
- (ii) Under Rauzy induction, the parents and children of an improper IET are all improper.

From (ii) we conclude that either all members of a Rauzy class are proper, or all are improper. In the latter case, each of the IET's is characterized by at least one parameter which can be varied continuously without disturbing the criteria of strict renormalizability. Removing these parameters, with merging of the relevant atoms, reduces the renormalization analysis to that of a Rauzy class of IET's with a smaller number of intervals. In other words, a complete inventory of renormalizable IET's can be obtained by restricting one's attention to the proper Rauzy classes. Doing so leads to a considerable improvement in efficiency: for example, for $n = 7$, there are 13 Rauzy classes containing a total of 3447 irreducible IET's; of these, there are 6 proper classes, with 1340 IET's. Self-similarity within the remaining 2107 improper IET's can be analysed by studying the proper IET's with $n < 7$. Note that $n = 4$ is special in having only degenerate IET's in its single improper Rauzy class. For larger n , nondegenerate, hence recombinant, examples are plentiful within the improper Rauzy classes, since it is easy to find permutations of the form $(\dots, n, k+1, \dots, k)$ with no consecutive pairs.

A deeper understanding of the above results can be obtained by representing an IET as Poincaré section of a flow on a translation surface [11, 24, 25]. The latter is a polygon with $2n$ sides (n is the number of intervals), labelled according to the ordering of the intervals before and after the permutation. The sides that correspond to the same interval have equal length and are parallel, and they are to be identified. A rectilinear flow on the plane will develop conical singularities on the surface, in correspondence to the vertices of the $2n$ -gon. While the translation surface is not unique, its genus and singularities (given by the total angle $2\pi(m+1)$ at the identified vertices) depend only on the Rauzy class. The removable singularities ($m = 0$) correspond to degenerate permutations, and they signal the appearance of free parameters in renormalizability. Since the translation surface does not change under induction, these structures depend only on the Rauzy class to which the permutation belongs.

6 Weakly-rigid self-similarity

The analogy with Rauzy induction suggests that there might exist two-parameter planar PWI's which do not admit self-similarity with free parameters, meaning that both parameters would be algebraically constrained. We shall exhibit a weak form of this property, resulting from the coexistence of two systems with semi-rigid self-similarity.




Figure 13: Three-parameter rhombus map with four atoms. The boundary lines a, b, c, d, e, f, g are respectively defined by the equations $x_0 = -t, x_1 = -s, x_0 = 1 - t, x_1 = 1 - s, x_2 = -s, x_2 = -1 - s - u, x_1 = -t$, where $x_m = \mathbf{u}_m \cdot \mathbf{x}$ and the basis vectors \mathbf{u}_m are given in (4), with s, t , and u real parameters which vary continuously over finite intervals.

Our starting point is the four-atom, three-parameter PWI of the $2\pi/5$ rhombus shown in figure 13. The interpretation of the parameters s, t, u is made clear in the conjugate system of figure 14, where the rhombus appears as a unit square. (A similar strategy can be pursued for the five-atom family of figures 2 and 3, but we found that the four-atom family is somewhat easier to work with.)

In convenient coordinates, the dressed domain is

$$\mathcal{R} = (R, (R_1, \dots, R_4), (\rho_1, \dots, \rho_4)),$$




Figure 14: Representation of the dressed domain \mathcal{R} using basis vectors (10). The dependence of the atoms's shapes on the parameters s, t, u is shown.

with

$$\begin{aligned}
R &= [(0, 1, 0, 1), (-1, -1, 1, 1), (-t, -s, 1-t, 1-s)], \\
R_1 &= [(0, 1, 2, 1), (-1, -1, -1, 1), (-t, -t, -s, 1-s)], \\
R_2 &= [(0, 1, 2, 1), (1, 1, 1, -1), (1-t, 1-s, -s, -t)], \\
R_3 &= [(0, 1, 2, 1), (-1, -1, -1, 1), (-t, -s, -1-s+u, -t)], \\
R_4 &= [(0, 1, 2, 1), (1, 1, 1, -1), (1-t, -t, -1-s+u, -s)]
\end{aligned} \tag{22}$$

$$\begin{aligned}
\rho_1 &= \mathbb{T}_{(0,0)} \mathbf{R}_4, & \rho_2 &= \mathbb{T}_{(0,1)} \mathbf{R}_4, \\
\rho_3 &= \mathbb{T}_{(1,1-u)} \mathbf{R}_4 & \rho_4 &= \mathbb{T}_{(1,2-u)} \mathbf{R}_4.
\end{aligned} \tag{23}$$

From figure 14, we see that the bifurcation-free domain $\Pi(\mathcal{R}) \subset \mathbb{R}^3$ is the polytope bounded by the planes $s-t=0$, $s-t=1$, $u=1$, $s+\beta^2 t=0$, $u-s-\beta^2 t=0$, $\beta^2-\beta^2 s-t=0$, and $1+\beta^2 s+t-u=0$.

This system has a simple one-parameter subsystem on the line L defined by $s-t=\beta^2, u=\beta$. We shall consider a two-parameter perturbation of this subsystem in the plane $u=\beta$ which intersects L . (We have also considered other planes, obtaining other manageable examples: see remarks at the end of this section.)

Setting $u=\beta$, the parameter polytope reduces to the hexagonal domain shown in figure 17 (left). As done in section 4, we choose a parameter pair

close to L lying within such a domain: $(s_0, t_0) = (2/5, 1/25)$. By inducing on




Figure 15: Induced dressed domain \mathcal{F} .

the trapezoidal atom R_1 , we obtain the parametric dressed domain \mathcal{F} shown in figure 15. One readily verifies that the return orbits of the eight atoms of \mathcal{F} completely tile R , so that the renormalizability of \mathcal{R} will follow from that of \mathcal{F} . We find that the complete tiling of F by renormalizable dressed sub-domains, given in figure 16, requires the return orbits of three dressed triangles $\mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3$, plus seven periodic tiles \mathcal{P}_i (five regular pentagons, one trapezoid, and one rhombus).

Letting

$$\Pi^*(\mathcal{F}) = \bigcap_{i=1}^3 \Pi(\mathcal{F}_i) \bigcap_{i=1}^7 \Pi(\mathcal{P}_i) \quad (24)$$

we find that $\Pi^*(\mathcal{F})$ is the quadrilateral with vertices

$$(\beta^2, 0), \quad (\beta^2 + \beta^6/\alpha, \beta^6/\alpha), \quad (\beta^2 + \beta^6/\alpha, \beta^4/\alpha), \quad (\beta^2, \beta^4/\alpha),$$

shown in figure 17 (right). Note that one of the bounding edges of the parameter domain coincides with the line L which was the starting point of our perturbative exploration.

The dressed domains \mathcal{F}_1 and \mathcal{F}_2 are base triangles equivalent to the prototype \mathcal{B} , with respective shape parameters

$$\tau_1 = \omega^6 \alpha (s - \beta^2), \quad \tau_2 = \omega^4 \alpha t.$$

Examination of \mathcal{F}_3 shows that its atoms with four and five sides share the same isometry (in spite of having different return paths, and even different




Figure 16: Tiling of F by return orbits of \mathcal{F}_1 (red), \mathcal{F}_2 (green), \mathcal{F}_3 (blue), and seven periodic tiles (grey).

return times on the rhombus), and hence can be merged for the purpose of testing renormalizability. After the merger, \mathcal{F}_3 is also equivalent to \mathcal{B} , with shape parameter

$$\tau_3 = \omega^6 \alpha (s - \beta^2).$$

Since each \mathcal{F}_i is renormalizable for $\tau_i \in \mathbb{Q}(\sqrt{5})$, we conclude that \mathcal{F} (hence \mathcal{R}) is renormalizable when all three shape parameters are in $\mathbb{Q}(\sqrt{5})$, i.e., when (s, t) is constrained to belong to $\Pi^*(\mathcal{F}) \cap \mathbb{Q}(\sqrt{5})^2$.

Each of the three renormalizable dressed domains \mathcal{F}_i provides a sequence of nested coverings of a distinct invariant component of the exceptional set complementary to all periodic points of the rhombus. The number of distinct ergodic components of the exceptional set is thus at least three. For the model of section 4, on the other hand, we believe that there is a single ergodic component.

In closing, we summarize briefly the results of our explorations of the three-parameter space of the four- or five-atom rhombus maps.

Manageable renormalizations are likely to be found in systems specified by parameters of small height. [The height $H(\zeta)$ of the algebraic number $\zeta = (m/n) + (m'/n')\omega$ is defined as $H(\zeta) = \max(|m|, |n|, |m'|, |n'|)$]. Such was the case for the domain \mathcal{R} , and the two-parameter restriction $u = \beta$ described above. We have considered other planes of small height: $s - t - \omega u + \beta = 0$, $s - t + u - 1 = 0$, etc. Within such planes, and for carefully chosen parameter




Figure 17: Left: The bifurcation-free domains $\Pi(\mathcal{R})$ and $\Pi(\mathcal{F})$ for the parameters s and t , with $u = \beta$. Right: Detailed view of $\Pi(\mathcal{F})$, showing the trapezoidal domain $\Pi^*(\mathcal{F})$ of equation (24). The latter is constructed as the intersection of $\Pi(\mathcal{F}_1)$ (semi-transparent green trapezoid), $\Pi(\mathcal{F}_2)$ (semi-transparent blue triangle), and $\Pi(\mathcal{F}_3)$ (semi-transparent red square). The seven periodic tiles \mathcal{P}_i do not contribute any additional constraints. The unperturbed one-parameter model corresponds to the (yellow) line $s - t = \beta^2$, which lies along the south-east boundary of $\Pi(\mathcal{F})$

patches, we have encountered a number of two-parameter renormalizable models. These are characterized by a decomposition of the rhombus into N disjoint dressed domains, each tiled by the return orbits of a single base triangle (provided we ignore the “decorations” produced by the common boundaries of merged atoms), plus a finite number of periodic tiles.

For the plane $u = 0$ of the five-atom map, the case $N = 1$ appears to be the norm, so that in those models the exceptional set is likely to be uniquely ergodic. Elsewhere, a proliferation of ergodic components is typical (albeit not universal). Notably, we have found no example of a rigidly self-similar, single component piecewise isometry with two or more parameters. Whether such a dynamical system exists at all remains an important open question.

References

- [1] R. Adler, B. Kitchens and C. Tresser, Dynamics of non-ergodic piecewise affine maps of the torus, *Ergod. Th. and Dynam. Sys.* **21** (2001) 959–999.
- [2] S. Akiyama and H. Brunotte and A. Pethó and W. Steiner, Periodicity of certain piecewise affine integer sequences, *Tsukuba J. Math.* **32** (2008) 197–251.
- [3] P. Arnoux P. and J. Yoccoz, Construction de difféomorphismes pseudo-Anosov, *C. R. Acad. Sci. Paris* **292** (1981) 75–78.
- [4] J. Barrionuevo, R. M. Burton, K. Dajani and C. Kraaikamp, Ergodic properties of generalised Lüroth series, *Acta Arithm.* **LXXIV** (4) (1996) 311–327.
- [5] M. D. Boshernitzan and C. R. Carroll, An extension of Lagrange’s theorem to interval exchange transformations over quadratic fields, *Journal d’Analyse Mathématique* **72** (1997) 21–44.
- [6] H. Cohn, *Advanced number theory*, Dover, New York (1980). (First published as *A second course in number theory* John Wiley and Sons, New York (1962).)
- [7] Supplemental material, comprising all Mathematica[®] procedures employed in this investigation.
<https://nyu.box.com/s/idygjzsy2f5f24tzezvfc7fi00rpi9>
- [8] J. Galambos, *Representations of Real Numbers by Infinite Series*, Lecture Notes in Math. 502, Springer, Berlin (1982).
- [9] A. Goetz and G. Poggiaspalla, Rotation by $\pi/7$, *Nonlinearity* **17** (2004) 1787–1802.
- [10] W. P. Hooper, Renormalization of polygon exchange maps arising from corner percolation, *Invent. Math.*, **191** (2013) 255–320.
- [11] M. Kontsevich and A. Zorich, Connected components of the moduli spaces of Abelian differentials with prescribed singularities, *Invent. Math.* **153** (2003) 631–678.
- [12] K. L. Kouptsov, J. H. Lowenstein and F. Vivaldi, Quadratic rational rotations of the torus and dual lattice maps, *Nonlinearity* **15** (2002) 1795–1842.

- [13] J. H. Lowenstein, *Pseudochaotic kicked oscillators*, Higher Education Press, Beijing and Springer-Verlag, Berlin (2012).
- [14] J. H. Lowenstein, S. Hatjispyros and F. Vivaldi, Quasi-periodicity, global stability and scaling in a model of Hamiltonian round-off, *Chaos* **7** (1997) 49–66.
- [15] J. H. Lowenstein, K. L. Kouptsov and F. Vivaldi, Recursive tiling and geometry of piecewise rotations by $\pi/7$, *Nonlinearity* **17** (2004) 371–395.
- [16] J. H. Lowenstein and F. Vivaldi, Renormalization of one-parameter families of piecewise isometries, *Dynamical Systems* (2016). [DOI:10.1080/14689367.2016.1153605]
- [17] J. Lüroth, Ueber eine eindeutige Entwicklung von Zahlen in eine unendliche Reihe, *Math. Ann.* **21** (1883) 411–423.
- [18] G. Poggiaspalla, Auto-similarités dans les systèmes isométriques par morceaux, PhD thesis, Université de la Méditerranée, Aix-Marseille II, (2003).
- [19] G. Poggiaspalla, Self-similarity in piecewise isometric systems, *Dynamical Systems*, **21** (2006) 147–189.
- [20] G. Poggiaspalla, J. H. Lowenstein and F. Vivaldi, Geometric representation of interval exchange maps over algebraic number fields, *Nonlinearity* **21** (2008) 149–177.
- [21] G. Rauzy, Échange d’intervalles et transformations induites, *Acta Arith.*, **34** (1979) 315–328.
- [22] R. E. Schwartz, *Outer billiards on kites*, vol. 171 of Annals of Math. studies, Princeton University Press, Princeton (2009).
- [23] R. E. Schwartz, *The Octagonal Pet*, Mathematical Surveys and Monographs, Volume 97, American Mathematical Society, (2014).
- [24] W. Veech, Gauss measures for transformations on the space of interval exchange maps, *Acta Math.*, **115** (1982) 201–242.
- [25] M. Viana, The ergodic theory of interval exchange maps, *Rev. Mat. Complut.* **19** (2006) 7–100.

- [26] J.-C. Yoccoz, *Continued fractions algorithms for interval exchange maps: an introduction*, *Frontiers in Number Theory, Physics and Geometry*, Vol 1, P. Cartier, B. Jula, P. Moussa, P. Vanhove (editors), Springer-Verlag, Berlin 4030437 (2006).