

Ergodic Theory of Interval Exchange Maps

Marcelo VIANA

IMPA

Estrada Dona Castorina 110
22460 Rio de Janeiro — Brazil
viana@impa.br

Received: November 10, 2005

Accepted: January 16, 2006

ABSTRACT

A unified introduction to the dynamics of interval exchange maps and related topics, such as the geometry of translation surfaces, renormalization operators, and Teichmüller flows, starting from the basic definitions and culminating with the proof that almost every interval exchange map is uniquely ergodic. Great emphasis is put on examples and geometric interpretations of the main ideas.

Key words: interval exchange map, translation surface, Abelian differential, Teichmüller flow.

2000 Mathematics Subject Classification: 37D25, 37D40, 37D50.

Partially supported by Pronex, CNPq, and Faperj, Brazil.

Contents

1. Interval exchange maps	11
2. Rauzy-Veech induction	14
3. Keane condition	18
4. Minimality	20
5. Dynamics of the induction map	23
6. Rauzy classes	27
7. Rauzy-Veech renormalization	32
8. Zorich transformations	34
9. Continued fractions	36
10. Symplectic form	37
11. Translation surfaces	38
12. Suspending interval exchange maps	41
13. Some translation surfaces	43
14. Computing the suspension surface	46
15. Zippered rectangles	51
16. Genus and dimension	54
17. Hyperelliptic Rauzy classes	56
18. Invertible Rauzy-Veech induction	58

19. Induction for zippered rectangles	61
20. Teichmüller flow	62
21. Volume measure	65
22. Hyperelliptic pairs	71
23. Combinatorial statement	74
24. Finite volume	76
25. Recurrence and inducing	80
26. Projective metrics	82
27. Ergodicity	87
28. Space of invariant measures	88
29. Unique ergodicity	93
30. Zorich measure	95

Introduction

The study of interval exchange maps is a classical topic in Dynamics that has drawn a great deal of attention over the last decades. This is due to two main sorts of reasons. On the one hand, these dynamical systems have a particularly simple formulation and, yet, exhibit very rich dynamical behavior. On the other hand, they relate closely to several other objects, in Dynamics as well as in many other areas of Mathematics: measured foliations, translation flows, Abelian differentials, Teichmüller flows, continued fraction expansions, polygonal billiards, renormalization theory, to mention just a few.

Interest in this topic was renewed in very recent years, corresponding to the solution of certain long standing problems. Firstly, Avila, Forni [1] have shown that almost all interval exchange maps are weakly mixing. Topological weak mixing had been established some years ago by Nogueira, Rudolph [18]. Then, Avila and the present author [3, 4] proved the Zorich-Kontsevich [14, 25, 26] conjecture on the Lyapunov spectrum of the Teichmüller, following important partial results by Forni [8]. Then, even more recently, Avila, Gouezel, Yoccoz [2] showed that the Teichmüller flow is exponentially mixing. Partial progress in this direction was also obtained by Bufetov [7]

Results such as these rely on a substantial amount of information amassed since the late seventies, starting from the pioneer works of Rauzy [19], Keane [9, 10], Masur [17], Veech [21–23], Zorich [25, 26], and other authors. Amidst all this information, it is often not easy to find the most relevant lines of development, nor to unveil the geometric motivation underlying main ideas and arguments. In this article we aim to bridge that gap.

Indeed, we give a unified treatment of the main classical results, starting from the very definition of interval interchange map and culminating with the proof of the Keane conjecture that almost every interval exchange map admits a unique invariant probability measure. We put great emphasis on examples and geometric interpretations of the main ideas. Besides the original papers mentioned previously, we have also benefited from the presentations of Marmi, Moussa, Yoccoz [16, 24].

Our text may be divided into three main parts, each containing roughly ten sections.

In the first part, we define and analyze the class of interval exchange maps. One of the main tools is the Rauzy-Veech induction operator, that assigns to each interval

exchange map its first return to a convenient subinterval. The largest subset where this operator can be iterated indefinitely has full Lebesgue measure and is perfectly characterized by an explicit condition on the interval exchange map that was introduced by Keane. Moreover, interval exchange maps that satisfy the Keane condition are minimal, that is, all their orbits are dense.

We also introduce the Rauzy-Veech renormalization operator, defined by composing the induction operator with a rescaling of the domain. In addition, we consider “accelerated” induction and renormalization operators, introduced by Zorich. The Zorich renormalization operator may be seen as a high-dimensional version of the classical continued fraction expansion, as we shall also see.

In the second third of the paper we define and study translation surfaces and their geodesic flows. Translation surfaces provide a natural setting for defining the suspensions of interval exchange transformations, and introducing invertible versions of the induction and renormalization operators introduced previously in terms of interval exchanges. We describe the suspension construction and explain how the resulting translation surface may be computed from the combinatorial and metric data of the exchange transformation.

Another important dynamical system in the space of translation surfaces, or of zippered rectangles, is the Teichmüller flow. It is related to the Rauzy-Veech and Zorich renormalization operators in that the latter may be seen as Poincaré return maps of the Teichmüller flow to convenient cross-sections in the space of all interval exchange maps.

The third and last part of the paper is devoted to the proof of the Keane conjecture: almost every interval exchange map is uniquely ergodic. The original proof is due to Masur [17] and Veech [22], and alternative arguments were given by Rees [20], Kerckhoff [12], and Boshernitzan [6]. Our presentation is based on the original strategy, where the crucial step is to prove that the renormalization operator admits a natural absolutely continuous invariant measure which, in addition, is ergodic. The conjecture then follows by observing that interval exchange maps that are typical for this invariant measure are uniquely ergodic.

It is in the nature of things that the Masur-Veech invariant measure is infinite. However, the Zorich renormalization operator does admit an absolutely continuous invariant probability measure, which is also ergodic. Since this probability plays an important role in subsequent developments, we also review its construction in the last section of the paper.

1. Interval exchange maps

Let $I \subset \mathbb{R}$ be an interval (All intervals will be bounded, closed on the left and open on the right. For notational simplicity, we take the left endpoint of I to coincide with 0.) and $\{I_\alpha : \alpha \in \mathcal{A}\}$ be a partition of I into subintervals, indexed by some alphabet \mathcal{A} with $d \geq 2$ symbols. An *interval exchange map* is a bijective map from I to I which




Figure 1

is a translation on each subinterval I_α . Such a map f is determined by combinatorial and metric data as follows:

- (i) A pair $\pi = (\pi_0, \pi_1)$ of bijections $\pi_\varepsilon : \mathcal{A} \rightarrow \{1, \dots, d\}$ describing the ordering of the subintervals I_α before and after the map is iterated. This will be represented as

$$\pi = \begin{pmatrix} \alpha_1^0 & \alpha_2^0 & \dots & \alpha_d^0 \\ \alpha_1^1 & \alpha_2^1 & \dots & \alpha_d^1 \end{pmatrix}$$

where $\alpha_j^\varepsilon = \pi_\varepsilon^{-1}(j)$ for $\varepsilon \in \{0, 1\}$ and $j \in \{1, 2, \dots, d\}$.

- (ii) A vector $\lambda = (\lambda_\alpha)_{\alpha \in \mathcal{A}}$ with positive entries, where λ_α is the length of the subinterval I_α .

We call $p = \pi_1 \circ \pi_0^{-1} : \{1, \dots, d\} \rightarrow \{1, \dots, d\}$ the *monodromy invariant* of the pair $\pi = (\pi_0, \pi_1)$. Observe that our notation, that we borrow from Marmi, Moussa, Yoccoz [16], is somewhat redundant. Given any (π, λ) as above and any bijection $\phi : \mathcal{A}' \rightarrow \mathcal{A}$, we may define

$$\pi'_\varepsilon = \pi_\varepsilon \circ \phi, \quad \varepsilon \in \{0, 1\} \quad \text{and} \quad \lambda'_{\alpha'} = \lambda_{\phi(\alpha')}, \quad \alpha' \in \mathcal{A}'.$$

Then (π, λ) and (π', λ') have the same monodromy invariant and they define the same interval exchange transformation. This means one can always normalize the combinatorial data by choosing $\mathcal{A} = \{1, 2, \dots, d\}$ and $\pi_0 = \text{id}$, in which case π_1 coincides with the monodromy invariant p . However, this notation hides the symmetric roles of π_0 and π_1 , and is not invariant under the induction and renormalization algorithms that we are going to present. On the contrary, the present notation $\pi = (\pi_0, \pi_1)$ allows for a very elegant formulation of these algorithms, as we are going to see.

Example 1.1. The interval exchange transformation described by figure 1 corresponds to the pair $\pi = \begin{pmatrix} C & B & A & D \\ D & B & A & C \end{pmatrix}$. The monodromy invariant is equal to $p = (4, 2, 3, 1)$.

Example 1.2. For $d = 2$ there is essentially only one combinatorics, namely

$$\pi = \begin{pmatrix} A & B \\ B & A \end{pmatrix}.$$

The interval exchange transformation associated to (π, λ) is given by

$$f(x) = \begin{cases} x + \lambda_B & \text{if } x \in I_A, \\ x - \lambda_A & \text{if } x \in I_B. \end{cases}$$

Identifying I with the circle $\mathbb{R}/(\lambda_A + \lambda_B)\mathbb{Z}$, we get

$$f(x) = x + \lambda_B \pmod{(\lambda_A + \lambda_B)\mathbb{Z}}. \tag{1}$$

That is, the transformation corresponds to the rotation of angle $\lambda_B/(\lambda_A + \lambda_B)$.

Example 1.3. The data (π, λ) is *not* uniquely determined by f . Indeed, let

$$\pi = \begin{pmatrix} A & B & C \\ B & C & A \end{pmatrix}.$$

Given any λ , the interval exchange transformation f defined is

$$f(x) = \begin{cases} x + \lambda_B + \lambda_C & \text{for } x \in I_A, \\ x - \lambda_A & \text{for } x \in I_B \cup I_C. \end{cases}$$

This shows that f is also the interval exchange transformation defined by either of the following data:

- (π, λ') for any other λ' such that $\lambda'_A = \lambda_A$ and $\lambda'_B + \lambda'_C = \lambda_B + \lambda_C$
- $(\tilde{\pi}, \tilde{\lambda})$ with $\tilde{\pi} = \begin{pmatrix} A & D \\ D & A \end{pmatrix}$ and $\lambda''_A = \lambda_A$ and $\lambda''_D = \lambda_B + \lambda_C$.

Translation vectors. Given $\pi = (\pi_0, \pi_1)$, define $\Omega_\pi : \mathbb{R}^A \rightarrow \mathbb{R}^A$ by

$$\Omega_\pi(\lambda) = w \quad \text{with} \quad w_\alpha = \sum_{\pi_1(\beta) < \pi_1(\alpha)} \lambda_\beta - \sum_{\pi_0(\beta) < \pi_0(\alpha)} \lambda_\beta. \tag{2}$$

Then the corresponding interval exchange transformation f is given by

$$f(x) = x + w_\alpha, \quad \text{for } x \in I_\alpha.$$

We call w the *translation vector* of f . Notice that the matrix $(\Omega_{\alpha,\beta})_{\alpha,\beta \in A}$ of Ω_π is given by

$$\Omega_{\alpha,\beta} = \begin{cases} +1 & \text{if } \pi_1(\alpha) > \pi_1(\beta) \text{ and } \pi_0(\alpha) < \pi_0(\beta), \\ -1 & \text{if } \pi_1(\alpha) < \pi_1(\beta) \text{ and } \pi_0(\alpha) > \pi_0(\beta), \\ 0 & \text{in all other cases.} \end{cases} \tag{3}$$

(Except where otherwise stated, all matrices are with respect to the canonical basis of \mathbb{R}^A .)

Example 1.4. In the case of figure 1,

$$(w_A, w_B, w_C, w_D) = (\lambda_D - \lambda_C, \lambda_D - \lambda_C, \lambda_D + \lambda_B + \lambda_A, -\lambda_C - \lambda_B - \lambda_A).$$

The image of Ω_π is the 2-dimensional subspace

$$\{w \in \mathbb{R}^A : w_A = w_B = w_C + w_D\}.$$

On the other hand, for $\pi = \begin{pmatrix} A & B & C & D \\ D & C & B & A \end{pmatrix}$ we have

$$(w_A, w_B, w_C, w_D) = (\lambda_D + \lambda_C + \lambda_B, \lambda_D + \lambda_C - \lambda_A, \lambda_D - \lambda_B - \lambda_A, -\lambda_C - \lambda_B - \lambda_A)$$

and Ω_π is a bijection from \mathbb{R}^A to itself.

Lemma 1.5. *We have $\lambda \cdot w = 0$.*

Proof. This is an immediate consequence of the fact that Ω_π is anti-symmetric. A detailed calculation follows. By definition

$$\lambda \cdot w = \sum_{\alpha \in \mathcal{A}} \lambda_\alpha w_\alpha = \sum_{\alpha \in \mathcal{A}} \lambda_\alpha \left(\sum_{\pi_1(\beta) < \pi_1(\alpha)} \lambda_\beta - \sum_{\pi_0(\beta) < \pi_0(\alpha)} \lambda_\beta \right)$$

and this is equal to

$$\sum_{\pi_1(\beta) < \pi_1(\alpha)} \lambda_\alpha \lambda_\beta - \sum_{\pi_0(\beta) < \pi_0(\alpha)} \lambda_\alpha \lambda_\beta = \frac{1}{2} \sum_{\alpha \neq \beta} \lambda_\alpha \lambda_\beta - \frac{1}{2} \sum_{\alpha \neq \beta} \lambda_\alpha \lambda_\beta = 0.$$

This proves the statement. □

The *canonical involution* is the operation in the space of (π, λ) corresponding to interchanging the roles of π_0 and π_1 while leaving λ unchanged. Clearly, under this operation the monodromy invariant p and the transformation f are replaced by their inverses. Moreover, Ω_π is replaced by $-\Omega_\pi$, and so the translation vector is also replaced by its symmetric.

2. Rauzy-Veech induction

Let (π, λ) represent an interval exchange transformation. For each $\varepsilon \in \{0, 1\}$, denote by $\alpha(\varepsilon)$ the last symbol in the expression of π_ε , that is

$$\alpha(\varepsilon) = \pi_\varepsilon^{-1}(d) = \alpha_d^\varepsilon$$

Let us assume the intervals $I_{\alpha(0)}$ and $I_{\alpha(1)}$ have different lengths. Then we say that (π, λ) has *type 0* if $\lambda_{\alpha(0)} > \lambda_{\alpha(1)}$ and *type 1* if $\lambda_{\alpha(0)} < \lambda_{\alpha(1)}$. In either case, the largest of the two intervals is called the *winner* and the shortest one is called the *loser* of (π, λ) . Let J be the subinterval of I obtained by removing the loser, that is, the shortest of these two intervals:

$$J = \begin{cases} I \setminus f(I_{\alpha(1)}) & \text{if } (\pi, \lambda) \text{ has type 0} \\ I \setminus I_{\alpha(0)} & \text{if } (\pi, \lambda) \text{ has type 1.} \end{cases}$$

The *Rauzy-Veech induction* of f is the first return map $\hat{R}(f)$ to the subinterval J . This is again an interval exchange transformation, as we are going to explain.




Figure 2




Figure 3

If (π, λ) has type 0, take $J_\alpha = I_\alpha$ for $\alpha \neq \alpha(0)$ and $J_{\alpha(0)} = I_{\alpha(0)} \setminus f(I_{\alpha(1)})$. These intervals form a partition of J . Note that $f(J_\alpha) \subset J$ for every $\alpha \neq \alpha(1)$. This means that $\hat{R}(f) = f$ restricted these J_α . On the other hand,

$$f(J_{\alpha(1)}) = f(I_{\alpha(1)}) \subset I_{\alpha(0)}$$

and so,

$$f^2(J_{\alpha(1)}) \subset f(I_{\alpha(0)}) \subset J.$$

Consequently, $\hat{R}(f) = f^2$ restricted to $J_{\alpha(1)}$. See figure 2.

If (π, λ) has type 1, define $J_{\alpha(0)} = f^{-1}(I_{\alpha(0)})$ and $J_{\alpha(1)} = I_{\alpha(1)} \setminus J_{\alpha(0)}$, and $J_\alpha = I_\alpha$ for all other values of α . See figure 3. Then $f(J_\alpha) \subset J$ for every $\alpha \neq \alpha(0)$, and so $\hat{R}(f) = f$ restricted these J_α . On the other hand,

$$f^2(J_{\alpha(0)}) = f(I_{\alpha(0)}) \subset J,$$

and so $\hat{R}(f) = f^2$ restricted to $J_{\alpha(0)}$.

The induction map $\hat{R}(f)$ is not defined when the two rightmost intervals $I_{\alpha(0)}$ and $I_{\alpha(1)}$ have the same length. We shall return to this point in sections 3 and 5.

Remark 2.1. Suppose the n -th iterate $\hat{R}^n(f)$ is defined, for some $n \geq 1$, and let I^n be its domain. It follows from the definition of the induction algorithm that $\hat{R}^n(f)$ is the first return map of f to I^n . Similarly, $\hat{R}^n(f)^{-1} = \hat{R}^n(f^{-1})$ is the first return map of f^{-1} to I^n .

Let us express the map $f \mapsto \hat{R}(f)$ in terms of the coordinates (π, λ) in the space of interval exchange transformations. It follows from the previous description that if (π, λ) has type 0 then the transformation $\hat{R}(f)$ is described by (π', λ') , where

- $$\pi' = \begin{pmatrix} \pi'_0 \\ \pi'_1 \end{pmatrix} = \begin{pmatrix} \alpha_1^0 & \cdots & \alpha_{k-1}^0 & \alpha_k^0 & \alpha_{k+1}^0 & \cdots & \cdots & \alpha_{d-1}^0 \\ \alpha_1^1 & \cdots & \alpha_{k-1}^1 & \alpha(0) & \alpha(1) & \alpha_{k+1}^1 & \cdots & \alpha_{d-1}^1 \end{pmatrix},$$

or, in other words,

$$\alpha_j^{0'} = \alpha_j^0 \quad \text{and} \quad \alpha_j^{1'} = \begin{cases} \alpha_j^1 & \text{if } j \leq k \\ \alpha(1) & \text{if } j = k + 1 \\ \alpha_{j-1}^1 & \text{if } j > k + 1, \end{cases} \tag{4}$$

where $k \in \{1, \dots, d - 1\}$ is defined by $\alpha_k^1 = \alpha(0)$.

- $\lambda' = (\lambda'_\alpha)_{\alpha \in \mathcal{A}}$ where

$$\lambda'_\alpha = \lambda_\alpha \quad \text{for } \alpha \neq \alpha(0), \quad \text{and} \quad \lambda'_{\alpha(0)} = \lambda_{\alpha(0)} - \lambda_{\alpha(1)}. \tag{5}$$

Analogously, if (π, λ) has type 1 then $\hat{R}(f)$ is described by (π', λ') , where

- $$\pi' = \begin{pmatrix} \pi'_0 \\ \pi'_1 \end{pmatrix} = \begin{pmatrix} \alpha_1^0 & \cdots & \alpha_{k-1}^0 & \alpha(1) & \alpha(0) & \alpha_{k+1}^0 & \cdots & \alpha_{d-1}^0 \\ \alpha_1^1 & \cdots & \alpha_{k-1}^1 & \alpha_k^1 & \alpha_{k+1}^1 & \cdots & \cdots & \alpha_{d-1}^1 \end{pmatrix},$$

or, in other words,

$$\alpha_j^{0'} = \begin{cases} \alpha_j^0 & \text{if } j \leq k \\ \alpha(0) & \text{if } j = k + 1 \\ \alpha_{j-1}^0 & \text{if } j > k + 1 \end{cases} \quad \text{and} \quad \alpha_j^{1'} = \alpha_j^1, \tag{6}$$

where $k \in \{1, \dots, d - 1\}$ is defined by $\alpha_k^0 = \alpha(1)$.

- $\lambda' = (\lambda'_\alpha)_{\alpha \in \mathcal{A}}$ where

$$\lambda'_\alpha = \lambda_\alpha \quad \text{for } \alpha \neq \alpha(1), \quad \text{and} \quad \lambda'_{\alpha(1)} = \lambda_{\alpha(1)} - \lambda_{\alpha(0)}. \tag{7}$$

Example 2.2. If $\pi = \begin{pmatrix} B & C & A & E & D \\ A & E & B & D & C \end{pmatrix}$ and $\lambda_D < \lambda_C$ (type 1 case) then

$$\pi' = \begin{pmatrix} B & C & D & A & E \\ A & E & B & D & C \end{pmatrix}$$

and $\lambda' = (\lambda_A, \lambda_B, \lambda_C - \lambda_D, \lambda_D, \lambda_E)$.

Operator Θ . Let us also compare the translation vectors w and w' of f and $\hat{R}(f)$, respectively. From figure 2 we see that, if (π, λ) has type 0,

$$w'_\alpha = w_\alpha \text{ for } \alpha \neq \alpha(1), \quad \text{and} \quad w'_{\alpha(1)} = w_{\alpha(1)} + w_{\alpha(0)}.$$

Analogously, if (π, λ) has type 1,

$$w'_\alpha = w_\alpha \text{ for } \alpha \neq \alpha(0), \quad \text{and} \quad w'_{\alpha(0)} = w_{\alpha(0)} + w_{\alpha(1)}.$$

This may be expressed as

$$w' = \Theta(w) \tag{8}$$

where $\Theta = \Theta_{\pi, \lambda} : \mathbb{R}^A \rightarrow \mathbb{R}^A$ is the linear operator whose matrix $(\Theta_{\alpha, \beta})_{\alpha, \beta \in A}$ is given by

$$\Theta_{\alpha, \beta} = \begin{cases} 1 & \text{if } \alpha = \beta, \\ 1 & \text{if } \alpha = \alpha(1) \text{ and } \beta = \alpha(0), \\ 0 & \text{in all other cases} \end{cases} \tag{9}$$

if (π, λ) has type 0, and

$$\Theta_{\alpha, \beta} = \begin{cases} 1 & \text{if } \alpha = \beta, \\ 1 & \text{if } \alpha = \alpha(0) \text{ and } \beta = \alpha(1), \\ 0 & \text{in all other cases} \end{cases} \tag{10}$$

if (π, λ) has type 1. Notice that Θ depends only on π and the type ε .

Observe that Θ is invertible and its inverse is given by

$$\Theta_{\alpha, \beta}^{-1} = \begin{cases} 1 & \text{if } \alpha = \beta, \\ -1 & \text{if } \alpha = \alpha(1) \text{ and } \beta = \alpha(0), \\ 0 & \text{in all other cases} \end{cases}$$

when (π, λ) has type 0, and

$$\Theta_{\alpha, \beta}^{-1} = \begin{cases} 1 & \text{if } \alpha = \beta, \\ -1 & \text{if } \alpha = \alpha(0) \text{ and } \beta = \alpha(1), \\ 0 & \text{in all other cases} \end{cases}$$

when (π, λ) has type 1. So, the relations (5) and (7) may be rewritten as

$$\lambda' = \Theta^{-1*}(\lambda) \quad \text{or} \quad \lambda = \Theta^*(\lambda') \tag{11}$$

where Θ^* denotes the adjoint operator of Θ , that is, the operator whose matrix is transposed of that of Θ .

Remark 2.3. The canonical involution does not affect the operator Θ : if $\tilde{\pi}$ is obtained by interchanging the lines of π , then $\Theta_{\tilde{\pi}, \lambda} = \Theta_{\pi, \lambda}$. Notice that $(\tilde{\pi}, \lambda)$ and (π, λ) have opposite types.

3. Keane condition

Summarizing the previous section, the Rauzy-Veech induction is expressed by the transformation

$$\hat{R} : \hat{R}(\pi, \lambda) = (\pi', \lambda')$$

where π' is given by (4) and (6), and λ' is given by (5) and (7). Recall that \hat{R} is not defined when the two rightmost intervals have the same length, that is, when $\lambda_{\alpha(0)} = \lambda_{\alpha(1)}$. We want to consider \hat{R} as a dynamical system in the space of interval exchange transformations, but for this we must restrict the map to an invariant subset of (π, λ) such that the iterates $\hat{R}^n(\pi, \lambda)$ are defined for all $n \geq 1$.

Let us start with the following observation. We say that a pair $\pi = (\pi_0, \pi_1)$ is *reducible* if there exists $k \in \{1, \dots, d-1\}$ such that

$$\pi_1 \circ \pi_0^{-1}(\{1, \dots, k\}) = \{1, \dots, k\}.$$

Then, for any choice of λ , the subinterval

$$J = \bigcup_{\pi_0(\alpha) \leq k} I_\alpha = \bigcup_{\pi_1(\alpha) \leq k} I_\alpha$$

is invariant under the transformation f , and so is its complement. This means that f splits into two interval exchange transformations, with simpler combinatorics. Moreover, $(\pi', \lambda') = \hat{R}(\pi, \lambda)$ is also reducible, with the same invariant subintervals. In what follows, we always restrict ourselves to irreducible data.

A natural possibility is to restrict the induction algorithm to the subset of *rationaly independent* vectors $\lambda \in \mathbb{R}_+^A$, that is, such that

$$\sum_{\alpha \in A} n_\alpha \lambda_\alpha \neq 0 \quad \text{for all nonzero integer vectors } (n_\alpha)_{\alpha \in A} \in \mathbb{Z}^A. \tag{12}$$

It is clear that this condition is invariant under iteration of (5) and (7), and that it ensures that all iterates $\hat{R}^n(\pi, \lambda)$ are defined. Observe also that the set of rationally independent vectors has full Lebesgue measure in the cone \mathbb{R}_+^A .

However, it was observed by Keane [9,10] that rational independence is a bit too strong: depending on the combinatorial data, failure of (12) for certain integer vectors may not be an obstruction to further iteration of \hat{R} . Let ∂I_γ be the left endpoint of each subinterval I_γ . Recall that we take the left endpoint of I to coincide with the origin. Then

$$\partial I_\gamma = \sum_{\pi_0(\eta) < \pi_0(\gamma)} \lambda_\eta$$

represents the left endpoint of each subinterval I_γ . A pair (π, λ) satisfies the *Keane condition* if the orbits of these endpoints are as disjoint as they can possible be (It is clear that if $\pi_0(\beta) = 1$ then $f(\partial I_\alpha) = \partial I_\beta$ for $\alpha = \pi_1^{-1}(1)$):

$$f^m(\partial I_\alpha) \neq \partial I_\beta \quad \text{for all } m \geq 1 \text{ and } \alpha, \beta \in A \text{ with } \pi_0(\beta) \neq 1. \tag{13}$$

This ensures that π is irreducible and $(\pi', \lambda') = \hat{R}(\pi, \lambda)$ is well-defined. Moreover, property (13) is invariant under iteration of \hat{R} , because $\hat{R}(f)$ -orbits are contained in f -orbits. Thus, the Keane condition is sufficient for all iterates $(\pi^n, \lambda^n) = \hat{R}^n(\pi, \lambda)$, $n \geq 0$ to be defined. We shall see in Corollary 5.4 that it is also *necessary*.

Remark 3.1. The Keane condition (13) is not affected if one restricts to the case $\pi_1(\alpha) > 1$. Indeed, suppose one has $f^m(\partial I_\alpha) = \partial I_\beta > 0$ with $\pi_1(\alpha) = 1$ and $m > 1$. Then $f(\partial I_\alpha) = 0 = \partial I_\gamma$ for some $\gamma \in \mathcal{A}$. Then, $f^{m-1}(\partial I_\gamma) = \partial I_\beta$. Moreover, $\pi_1(\gamma) > 1$ because π is irreducible and $\pi_0(\gamma) = 1$.

The next result shows that, assuming irreducibility, the Keane condition is indeed more general than rational independence. In particular, it also corresponds to full Lebesgue measure.

Proposition 3.2. *If λ is rationally independent and π is irreducible then (π, λ) satisfies the Keane condition.*

Proof. Assume there exist $m \geq 1$ and $\alpha, \beta \in \mathcal{A}$ such that $f^m(\partial I_\alpha) = \partial I_\beta$ and $\pi_0(\beta) > 1$. Define β_j , $0 \leq j \leq m$, by

$$f^j(\partial I_\alpha) \in I_{\beta_j}.$$

Notice that $\beta_0 = \alpha$ and $\beta_m = \beta$. Then

$$\partial I_\beta - \partial I_\alpha = \sum_{0 \leq j < m} w_{\beta_j}$$

where $w = (w_\gamma)_{\gamma \in \mathcal{A}}$ is the translation vector defined in (2). Equivalently,

$$\sum_{\pi_0(\gamma) < \pi_0(\beta_m)} \lambda_\gamma - \sum_{\pi_0(\gamma) < \pi_0(\beta_0)} \lambda_\gamma = \sum_{0 \leq j < m} \left(\sum_{\pi_1(\gamma) < \pi_1(\beta_j)} \lambda_\gamma - \sum_{\pi_0(\gamma) < \pi_0(\beta_j)} \lambda_\gamma \right).$$

This may be rewritten as $\sum_{\gamma \in \mathcal{A}} n_\gamma \lambda_\gamma = 0$, where

$$n_\gamma = \#\{0 \leq j < m : \pi_1(\beta_j) > \pi_1(\gamma)\} - \#\{0 < j \leq m : \pi_0(\beta_j) > \pi_0(\gamma)\}.$$

Since we assume rational independence, we must have $n_\gamma = 0$ for all $\gamma \in \mathcal{A}$. Now let D be the maximum of $\pi_0(\beta_j)$ over all $0 < j \leq m$ and $\pi_1(\beta_j)$ over all $0 \leq j < m$. Note that $D \geq \pi_0(\beta) > 1$. So, since we assume that π is irreducible, there exists $\gamma \in \mathcal{A}$ such that $\pi_0(\gamma) < D \leq \pi_1(\gamma)$. The last inequality implies that $\pi_1(\beta_j) \leq \pi_1(\gamma)$ for all $0 \leq j < m$. Since $n_\gamma = 0$, this implies that $\pi_0(\beta_j) \leq \pi_0(\gamma) < D$ for all $0 < j \leq m$. A symmetric argument shows that $\pi_1(\beta_j) < D$ for all $0 \leq j < m$. This contradicts the definition of D . This contradiction proves that (π, λ) satisfies the Keane condition, as stated. \square

Example 3.3. Suppose $d = 2$. By (1), the interval exchange transformation is given by $f(x) = x + \lambda_B \pmod{(\lambda_A + \lambda_B)\mathbb{Z}}$. So, the Keane condition means that, given any $m \geq 1$ and $n \in \mathbb{Z}$, both

$$m\lambda_B \neq \lambda_A + n(\lambda_A + \lambda_B) \quad \text{and} \quad \lambda_A + m\lambda_B \neq \lambda_A + n(\lambda_A + \lambda_B).$$

It is clear that this holds if and only if (λ_A, λ_B) is rationally independent.

Example 3.4. Starting from $d = 3$, the Keane condition may be strictly weaker than rational independence. Consider, for instance, $\pi = \begin{pmatrix} A & B & C \\ C & A & B \end{pmatrix}$. Then $f(x) = x + \lambda_C \pmod{(\lambda_A + \lambda_B + \lambda_C)\mathbb{Z}}$ and the Keane condition means that

$$m\lambda_C \quad \text{and} \quad \lambda_A + m\lambda_C \quad \text{and} \quad \lambda_A + \lambda_B + m\lambda_C$$

are different from $\lambda_A + n(\lambda_A + \lambda_B + \lambda_C)$ and $\lambda_A + \lambda_B + n(\lambda_A + \lambda_B + \lambda_C)$, for all $m \geq 1$ and $n \in \mathbb{Z}$. This may be restated in a more compact form, as follows: given any $p \in \mathbb{Z}$ and $q \in \mathbb{Z}$,

$$p\lambda_C \neq q(\lambda_A + \lambda_B + \lambda_C) \quad \text{and} \quad p\lambda_C \neq \lambda_A + q(\lambda_A + \lambda_B + \lambda_C).$$

Clearly, this may hold even if (λ_A, λ_B) is rationally dependent.

4. Minimality

A transformation is called *minimal* if every orbit is dense in the whole domain of definition or, equivalently, the domain is the only nonempty closed invariant set.

Proposition 4.1. *If (π, λ) satisfies the Keane condition then f is minimal.*

For the proof, we begin by noting that the first return map of f to some interval $J \subset I_\alpha$ is again an interval exchange transformation:

Lemma 4.2. *Given any subinterval $J = [a, b]$ of some I_α , there exists a partition $\{J_j : 1 \leq j \leq k\}$ of J and integers $n_1, \dots, n_k \geq 1$, where $k \leq d + 2$, such that*

- (i) $f^i(J_j) \cap J = \emptyset$ for all $0 < i < n_j$ and $1 \leq j \leq k$;
- (ii) each $f^{n_j} \mid J_j$ is a translation from J_j to some subinterval of J ;
- (iii) those subintervals $f^{n_j}(J_j)$, $1 \leq j \leq k$ are pairwise disjoint.

Proof. Let A be the union of the boundary $\{a, b\}$ of J with the set of endpoints of all the intervals I_γ , $\gamma \in \mathcal{A}$, the endpoints of I excluded. Note that $\#A \leq d + 1$. Let $B \subset J$ be the set of points $z \in J$ for which there exists some $m \geq 1$ such that $f^i(z) \notin J$ for all $0 < i < m$ and $f^m(z) \in A$. The map $B \ni z \mapsto f^m(z) \in A$ is injective, because f is injective and there are no iterates in J prior to time m . Consequently, $\#B \leq \#A$. Consider the partition of J determined by the points of B . This partition

has at most $d + 2$ elements. By the Poincaré recurrence theorem, for each element $J_j = [a_j, b_j)$ there exists $n_j \geq 1$ such that $f^{n_j}(J_j)$ intersects J . Take n_j smallest. From the definition of B it follows that the restriction $f^{n_j} \mid J_j$ is a translation and its image is contained in J . Finally, the $f^{n_j}(J_j)$, $1 \leq j \leq k$ are pairwise disjoint because f is injective and the n_j are the first return times to J . \square

In fact, the statement is true for any interval $J \subset I$. See [21, § 3].

Corollary 4.3. *Under the assumptions of Lemma 4.2, the union \hat{J} of all forward iterates of J is a finite union of intervals and a fully invariant set: $f(\hat{J}) = \hat{J}$.*

Proof. The first claim follows directly from the first part of Lemma 4.2:

$$\hat{J} = \bigcup_{n=0}^{\infty} f^n(J) = \bigcup_{j=1}^k \bigcup_{i=0}^{n_j-1} f^i(J_j).$$

Moreover, parts (ii) and (iii) of Lemma 4.2, together with the observation

$$\sum_{j=1}^k |f^{n_j}(J_j)| = \sum_{j=1}^k |J_j| = |J|,$$

(We use $|\cdot|$ to represent length.) give that J coincides with $\cup_{j=1}^k f^{n_j}(J_j)$. This implies that \hat{J} is fully invariant. \square

Lemma 4.4. *If (π, λ) satisfies the Keane condition then f has no periodic points.*

Proof. Suppose there exists $m \geq 1$ and $x \in I$ such that $f^m(x) = x$. Define β_j , $0 \leq j \leq m$ by the condition $f^j(x) \in I_{\beta_j}$. Let J be the set of all points $y \in I$ such that $f^j(y) \in I_{\beta_j}$ for all $0 \leq j < m$. Then J is an interval and f^m restricted to it is a translation. Since $f^m(x) = x$, we actually have $f^m \mid J = \text{id}$. In particular, $f^m(\partial J) = \partial J$. The definition of J implies that there are $1 \leq k \leq m$ and $\beta \in \mathcal{A}$ such that $f^k(\partial J) = \partial I_\beta$. Then $f^m(\partial I_\beta) = \partial I_\beta$. If $\pi_0(\beta) > 1$, this contradicts the Keane condition. If $\pi_0(\beta) = 1$ then there exists $\alpha \in \mathcal{A}$ such that $f(\partial I_\alpha) = 0 = \partial I_\beta$. Note that $\alpha \neq \beta$, and so $\partial I_\alpha > 0$, because π is irreducible. Hence, $f^m(\partial I_\alpha) = \partial I_\alpha$ contradicts the Keane condition. These contradictions prove that there is no such periodic point x . \square

Proof of Proposition 4.1. Suppose there exists $x \in I$ such that $\{f^n(x) : n \geq 0\}$ is not dense in I . Then we may choose a subinterval $J = [a, b)$ of some I_α that avoids the closure of the orbit. Let \hat{J} be the union of all forward iterates of J . By Corollary 4.3, this is a finite union of intervals, fully invariant under f . We claim that \hat{J} can not be of the form $[0, \hat{b})$. The proof is by contradiction. Let \mathcal{B} be the subset of $\alpha \in \mathcal{A}$ such




Figure 4

that I_α is contained in \hat{J} . Then $\pi_0(\mathcal{B}) = \{1, \dots, k\}$ for some k . Since \hat{J} is invariant, we also have $\pi_1(\mathcal{B}) = \{1, \dots, k\}$. Hence,

$$\pi_0^{-1}(\{1, \dots, k\}) = \mathcal{B} = \pi_1^{-1}(\{1, \dots, k\}). \tag{14}$$

It is clear that $k < d$, because \hat{J} avoids the closure of the orbit of x , and so it can not be the whole I . If $k = 0$ then \hat{J} would be contained in I_α , where $\pi_0(\alpha) = 1$; by invariance, it would also be contained in $f(I_\alpha)$, implying that $\pi_1(\alpha) = 1$; this would contradict irreducibility (which is a consequence of the Keane condition). Thus, k must be positive. Then (14) contradicts irreducibility, and this contradiction proves our claim.

As a consequence, there exists some connected component $[\hat{a}, \hat{b})$ of \hat{J} with $\hat{a} > 0$. If $f^n(\hat{a}) \neq \partial I_\beta$ for every $n \geq 0$ and $\beta \in \mathcal{A}$, then (by continuity of f and invariance of \hat{J}) every $f^n(\hat{a})$, $n \geq 0$ would be on the boundary of some connected component of \hat{J} . As there are finitely many components, f would have a periodic point, which is forbidden by Lemma 4.4. Similarly, if $f^n(\hat{a}) \neq f(\partial I_\alpha)$ for every $n \leq 0$ and $\alpha \in \mathcal{A}$, then every $f^n(\hat{a})$, $n \leq 0$ would be on the boundary of some connected component of \hat{J} . Just as before, this would imply the existence of some periodic point, which is forbidden by Lemma 4.4. This proves that there are $n_1 \leq 0 \leq n_2$ and $\alpha, \beta \in \mathcal{A}$ such that

$$f^{n_1}(\hat{a}) = f(\partial I_\alpha) \quad \text{and} \quad f^{n_2}(\hat{a}) = \partial I_\beta. \tag{15}$$

If $\partial I_\beta > 0$, this contradicts the Keane condition (take $m = n_2 - n_1 + 1$). If $\partial I_\beta = 0$ then $n_2 > 0$, because we have taken $\hat{a} > 0$. Moreover, $\partial I_\beta = f(\partial I_\gamma)$, where $\pi_1(\gamma) = 1$. This means that (15) remains valid if one replaces β by γ and n_2 by $n_2 - 1$. As $\gamma \neq \beta$, by irreducibility, we have $\partial I_\gamma > 0$ and this leads to a contradiction just as in the previous case. \square

Remark 4.5. The Keane condition is not necessary for minimality. Consider the interval exchange transformation f illustrated in figure 4, where $\lambda_A = \lambda_C$, $\lambda_B = \lambda_D$, and $\lambda_A/\lambda_B = \lambda_C/\lambda_D$ is irrational. Then f does not satisfy the Keane condition, yet it is minimal.

Unique ergodicity. A transformation is called *uniquely ergodic* if it admits exactly one invariant probability (which is necessarily ergodic). See Mañé [15]. Then the transformation is minimal restricted to the support of this probability. Observe that interval exchange transformations always preserve the Lebesgue measure. Thus, in

this context, unique ergodicity means that every invariant measure is a multiple of the Lebesgue measure.

Keane [9] conjectured that every minimal interval exchange transformation is uniquely ergodic, and checked that this is true for $d = 2, 3$. However, Keynes, Newton [13] gave an example with $d = 5$ and two ergodic invariant probabilities. In turn, they conjectured that rational independence should suffice for unique ergodicity. Again, a counterexample was given by Keane [10], with $d = 4$ and two ergodic invariant probabilities. He then went on to make the following

Conjecture 4.6. *Almost every interval exchange transformation is uniquely ergodic.*

This statement was proved in the early eighties, independently, by Masur [17] and Veech [22]. That unique ergodicity holds for a (Baire) residual subset had been proved by Keane, Rauzy [11].

5. Dynamics of the induction map

This section contains a number of useful facts on the dynamics of the induction algorithm in the space of interval exchange transformations. The presentation follows section 4.3 of Yoccoz [24].

Let (π, λ) be such that the iterates $(\pi^n, \lambda^n) = \hat{R}^n(\pi, \lambda)$ are defined for all $n \geq 0$. For instance, this is the case if (π, λ) satisfies the Keane condition. For each $n \geq 0$, let $\varepsilon^n \in \{0, 1\}$ be the type and $\alpha^n, \beta^n \in \mathcal{A}$ be, respectively, the winner and the loser of (π^n, λ^n) . In other words, α^n and β^n are the two rightmost symbols in the two lines of π^n , with $\lambda_{\alpha^n} > \lambda_{\beta^n}$. In yet another equivalent formulation, $\pi_{\varepsilon^n}(\alpha^n) = d = \pi_{1-\varepsilon^n}(\beta^n)$.

It is clear that the sequence $(\varepsilon^n)_n$ takes both values 0 and 1 infinitely many times. Indeed, suppose the type ε^n was eventually constant. Then α^n would also be eventually constant, and so would $\lambda_{\alpha^n}^n$ for all $\alpha \neq \alpha^n$. On the other hand,

$$\lambda_{\alpha^{n+1}}^{n+1} = \lambda_{\alpha^n}^{n+1} = \lambda_{\alpha^n}^n - \lambda_{\beta^n}^n$$

for all large n . Since the $\lambda_{\beta^n}^n$ are bounded from zero, the $\lambda_{\alpha^n}^n$ would be eventually negative, which is a contradiction.

Proposition 5.1. *Both sequences $(\alpha^n)_n$ and $(\beta^n)_n$ take every value $\alpha \in \mathcal{A}$ infinitely many times.*

Proof. Given any symbol $\alpha \in \mathcal{A}$, consider any maximal time interval $[p, q)$ such that $\alpha^n = \alpha$ for every $n \in [p, q)$. At the end of this interval the type must change:

$$\varepsilon^q = 1 - \varepsilon^{q-1} \quad \text{and} \quad \pi_{1-\varepsilon^q}^q(\alpha) = d.$$

In other words, $\alpha = \beta^q$. This shows that we only have to prove the statement for the sequence $(\alpha^n)_n$.

Let \mathcal{B} be the subset of symbols $\beta \in \mathcal{A}$ that occur only finitely many times in the sequence $(\alpha^n)_n$. Up to replacing (π, λ) by some iterate, we may suppose that those symbols do not occur at all in $(\alpha^n)_n$. Then $\lambda_\beta^n = \lambda_\beta$ for all $\beta \in \mathcal{B}$ and $n \geq 0$. Since

$$\lambda_{\alpha^{n+1}}^{n+1} = \lambda_{\alpha^n}^n - \lambda_{\beta^n}^n,$$

this implies that every $\beta \in \mathcal{B}$ occurs only finitely many times in the sequence (β_n) . Once more, up to replacing the initial point by an iterate, we may suppose they do not occur at all in (β_n) . It follows that, for every $\beta \in \mathcal{B}$, the sequences

$$\pi_0^n(\beta) \quad \text{and} \quad \pi_1^n(\beta), \quad n \geq 0,$$

are non-decreasing. So, replacing (π, λ) by an iterate one more time, if necessary, we may suppose that these sequences are constant. We claim that

$$\pi_\varepsilon(\beta) < \pi_\varepsilon(\alpha) \quad \text{for every } \alpha \in \mathcal{A} \setminus \mathcal{B}, \beta \in \mathcal{B}, \text{ and } \varepsilon = 0, 1. \tag{16}$$

Indeed, suppose there were α, β , and ε such that $\pi_\varepsilon(\alpha) < \pi_\varepsilon(\beta)$. Then, since the sequence $\pi_\varepsilon^n(\beta)$ is non-decreasing, so must be the sequence $\pi_\varepsilon^n(\alpha)$. In particular, $\pi_\varepsilon^n(\alpha) < d$ for all $n \geq 0$. Now, since $\alpha \notin \mathcal{B}$, this implies that $\pi_{1-\varepsilon}^n(\alpha) = d$ and $\varepsilon^n = 1 - \varepsilon$, for some value of n .

$$\begin{pmatrix} \cdots & \alpha & \cdots & \beta & \cdots & \gamma \\ \cdots & \cdots & \cdots & \cdots & \cdots & \alpha \end{pmatrix} \xrightarrow{\hat{R}} \begin{pmatrix} \cdots & \alpha & \gamma & \cdots & \beta & \cdots \\ \cdots & \cdots & \cdots & \cdots & \cdots & \alpha \end{pmatrix}$$

Then $\pi_\varepsilon^{n+1}(\beta) = \pi_\varepsilon^n(\beta) + 1$, contradicting the previous conclusion that $\pi_\varepsilon^n(\beta)$ is constant. This contradiction proves our claim. Finally, (16) implies that

$$\pi_0(\mathcal{B}) = \{1, \dots, k\} = \pi_1(\mathcal{B})$$

for some $k < d$. Since π is assumed to be irreducible, we must have $k = 0$, that is, \mathcal{B} is the empty set. This proves the statement for the sequence $(\alpha^n)_n$ and, hence, completes the proof of the proposition. \square

Corollary 5.2. *The length of the domain I^n of the transformation $\hat{R}^n(f)$ goes to zero when n goes to ∞ .*

Proof. Since the sequences λ_α^n are non-increasing, for all $\alpha \in \mathcal{A}$, it suffices to show that they all converge to zero. Suppose there was $\beta \in \mathcal{A}$ and $c > 0$ such that $\lambda_\beta^n \geq c$ for every $n \geq 0$. For any value of n such that $\beta^n = \beta$, we have

$$\lambda_{\alpha^{n+1}}^{n+1} = \lambda_{\alpha^n}^n - \lambda_{\beta^n}^n \leq \lambda_{\alpha^n}^n - c.$$

By Proposition 5.1, this occurs infinitely many times. As the alphabet \mathcal{A} is finite, it follows that there exists some $\alpha \in \mathcal{A}$ such that

$$\lambda_\alpha^{n+1} \leq \lambda_\alpha^n - c.$$

for infinitely many values of n . This contradicts the fact that $\lambda_\alpha^n > 0$. \square

Corollary 5.3. *For each $m \geq 0$ there exists $n \geq 1$ such that*

$$\Theta_{\pi^m, \lambda^m}^{*n} > 0. \quad (\text{All the entries of the matrix are positive.})$$

Proof. Given $\alpha, \beta \in \mathcal{A}$, $m \geq 0$, $n \geq 1$, we represent by $\Theta^*(\alpha, \beta, m, n)$ the entry on row α and column β of the matrix of $\Theta_{\pi^m, \lambda^m}^{*n}$. By definition (9) and (10),

$$\Theta^*(\alpha, \beta, m, 1) = 1 \text{ if either } \alpha = \beta \text{ or } (\alpha, \beta) = (\alpha^m, \beta^m), \tag{17}$$

and $\Theta^*(\alpha, \beta, m, 1) = 0$ in all other cases. Observe also that every $\Theta^*(\alpha, \beta, m, n)$ is non-decreasing on n :

$$\begin{aligned} \Theta^*(\alpha, \beta, m, n + 1) &= \sum_{\gamma} \Theta^*(\alpha, \gamma, m, n) \Theta^*(\gamma, \beta, m + n, 1) \\ &\geq \Theta^*(\alpha, \beta, m, n) \Theta^*(\beta, \beta, m + n, 1) \geq \Theta^*(\alpha, \beta, m, n). \end{aligned} \tag{18}$$

Let α be fixed. We are going to construct an enumeration $\gamma_1, \gamma_2, \dots, \gamma_d$ of \mathcal{A} and integers n_1, n_2, \dots, n_d such that

$$\Theta^*(\alpha, \gamma_i, m, n) > 0 \quad \text{for every } n > n_i \text{ and } i = 1, 2, \dots, d. \tag{19}$$

It is clear that this implies the corollary, as β must be one of the γ_i .

For $i = 1$ just take $\gamma_1 = \alpha$ and $n_1 = 0$. The relations (17) and (18) immediately imply (19). Next, use Proposition 5.1 to find $m_2 > m$ such that the winner α^{m_2} coincides with γ_1 . Let $\gamma_2 = \beta^{m_2}$ be the loser. Note that $\gamma_2 \neq \gamma_1$, by irreducibility. Moreover, (17) gives that $\Theta^*(\gamma_1, \gamma_2, m_2, 1) = 1$, and this implies $\Theta^*(\gamma_1, \gamma_2, m, n) > 0$ for every $n > m_2 - m$. This gives (19) for $i = 2$, with $n_2 = m_2 - m$. If $d = 2$ then there is nothing left to prove, so assume $d > 2$. Using Proposition 5.1 twice, one finds $p_2 > m_2$ such that the winner α^{p_2} is neither γ_1 nor γ_2 , and $m_3 > p_2$ such that the winner $\alpha^{m_3} = \gamma_j$ for either $j = 1$ or $j = 2$. Consider the smallest such m_3 , and let $\gamma_3 = \beta^{m_3}$ be the loser. Notice that $\gamma_3 = \alpha^{m_3-1}$ and so it is neither γ_1 nor γ_2 . Moreover, (17) gives that $\Theta^*(\gamma_j, \gamma_3, m_3, 1) = 1$ and this implies

$$\Theta^*(\gamma_1, \gamma_3, m, n) \geq \Theta^*(\gamma_1, \gamma_j, m, m_3 - m) \Theta^*(\gamma_j, \gamma_3, m_3, n - m_3 + m) > 0$$

for $n > m_3 - m$. Notice that $m_3 - m > m_2 - m = n_2$. This proves (19) for $i = 3$ with $n_3 = m_3 - m$.

The general step of the enumeration is analogous. Assume we have constructed $\gamma_1, \dots, \gamma_k \in \mathcal{A}$, all distinct, and integers n_1, n_2, \dots, n_k such that (19) holds for $1 \leq i \leq k$. Assuming $k < d$, we may use Proposition 5.1 twice to find $p_k > m_k$ such that the winner α^{p_k} is not an element of $\{\gamma_1, \dots, \gamma_k\}$ and $m_{k+1} > p_k$ such that the winner $\alpha^{m_{k+1}} = \gamma_j$ for some $j \in \{1, \dots, k\}$. Choose the smallest such m_{k+1} and let $\gamma_{k+1} = \beta^{m_{k+1}}$ be the loser. Then $\gamma_{k+1} = \alpha^{m_{k+1}-1}$ and so it is not an element of $\{\gamma_1, \dots, \gamma_k\}$. The relation (17) gives $\Theta^*(\gamma_j, \gamma_{k+1}, m_{k+1}, 1) = 1$, and then

$$\Theta^*(\gamma_1, \gamma_{k+1}, m, n) \geq \Theta^*(\gamma_1, \gamma_j, m, m_{k+1} - m) \Theta^*(\gamma_j, \gamma_{k+1}, m_{k+1}, n - m_{k+1} + m)$$




Figure 5

is strictly positive for all $n > n_{k+1} = m_{k+1} - m$. This completes our recurrence construction and, thus, finishes the proof of the corollary. \square

At this point we can prove that (π, λ) can be iterated indefinitely (if and) only if it satisfies the Keane condition:

Corollary 5.4. *If $(\pi^n, \lambda^n) = \hat{R}^n(\pi, \lambda)$ is defined for all $n \geq 0$ then (π, λ) satisfies the Keane condition.*

Proof. Suppose that, for some $\alpha, \beta \in \mathcal{A}$, and $m \geq 1$,

$$f^{m-1}(\partial f(I_\alpha)) = \partial I_\beta. \tag{20}$$

Choose m minimum. In particular, by Remark 3.1, we have $\partial f(I_\alpha) > 0$. The definition of $f_n = \hat{R}^n(f)$ gives

$$\partial f(I_\alpha) = \partial f_n(I_\alpha^n), \quad \text{and} \quad \partial I_\beta = \partial I_\beta^n$$

for every n such that $\partial f(I_\alpha)$ and ∂I_β are in the domain I^n of f_n . Take n maximum such that both points are in I^n (Corollary 5.2). Since f_n is the first return map of f to I^n (Remark 2.1), the hypothesis (20) implies that

$$f_n^k(\partial f(I_\alpha)) = \partial I_\beta \quad \text{for some } k \leq m - 1. \tag{21}$$

Moreover, either I_β or $f_n(I_\alpha^n)$ (or both) is a rightmost partition interval for f_n .

If $\partial f(I_\alpha) = \partial I_\beta$ then $f_n(I_\alpha^n) = I_\beta^n$, that is, the two rightmost intervals of f_n have the same length. See figure 5. Hence, $f_{n+1} = \hat{R}^{n+1}(f)$ is not defined, which contradicts the hypothesis. This proves the statement in this case.

Now suppose f_n has type 0, that is, $\partial I_\beta < \partial f(I_\alpha)$. By definition,

$$f_{n+1}(\partial I_\alpha^{n+1}) = f_n^2(\partial I_\alpha^n) = f_n(\partial f(I_\alpha)) \quad \text{and} \quad \partial I_\beta^{n+1} = \partial I_\beta^n = \partial I_\beta.$$

See figure 6. Comparing with (21) we get

$$f_n^{k-1}(\partial f_{n+1}(I_\alpha^{n+1})) = f_n^k(\partial I_\alpha) = \partial I_\beta = \partial I_\beta^{n+1}.$$




Figure 6




Figure 7

Since both points are in I_α^{n+1} and f_{n+1} is the return map of f_n to I_α^{n+1} , this may be rewritten as

$$f_{n+1}^{l-1}(\partial f_{n+1}(I_\alpha^{n+1})) = \partial I_\beta^{n+1} \quad \text{for some } l \leq k < m. \tag{22}$$

Now suppose f_n has type 1, that is, $\partial I_\beta > \partial f(I_\alpha)$. By definition,

$$\partial f_{n+1}(I_\alpha^{n+1}) = \partial f_n(I_\alpha^n) = \partial f(I_\alpha) \quad \text{and} \quad \partial I_\beta^{n+1} = f_n^{-1}(\partial I_\beta^n) = f_n^{-1}(\partial I_\beta).$$

See figure 7. Comparing with (21) we get

$$f_n^{k-1}(\partial f_{n+1}(I_\alpha^{n+1})) = f_n^{k-1}(\partial f(I_\alpha)) = f_n^{-1}(\partial I_\beta) = \partial I_\beta^{n+1}.$$

Since f_{n+1} is the return map of f_n to I_α^{n+1} , this may be rewritten as

$$f_{n+1}^{l-1}(\partial I_\alpha^{n+1}) = \partial I_\beta^{n+1} \quad \text{for some } l \leq k < m. \tag{23}$$

In both subcases, we have shown that (20) implies a similar relation, either (22) or (23), where f is replaced by some induced map f_{n+1} , and $m \geq 2$ is replaced by a smaller l . Iterating this procedure, we must eventually reach the case $m = 1$, which was treated previously. \square

6. Rauzy classes

Given pairs π and π' , we say that π' is a *successor* of π if there exist $\lambda, \lambda' \in \mathbb{R}_+^A$ such that $\hat{R}(\pi, \lambda) = (\pi', \lambda')$. Any pair π has exactly two successors, corresponding to types 0 and 1. Similarly, each π' is the successor of exactly two pairs π , obtained by reversing the relations (4) and (6). Notice that π is irreducible if and only if π' is

irreducible. Thus, this relation defines a partial order in the set of irreducible pairs, which we may represent as a directed graph G . We call *Rauzy classes* the connected components of this graph.

Lemma 6.1. *If π and π' are in the same Rauzy class then there exists an oriented path in G starting at π and ending at π' .*

Proof. Let $A(\pi)$ be the set of all pairs π' that can be attained through an oriented path starting at π . As we have just seen, each vertex of the graph G has exactly two outgoing and two incoming edges. By definition, every edge starting from a vertex of $A(\pi)$ must end at some vertex of $A(\pi)$. By a counting argument, it follows that every edge ending at a vertex of $A(\pi)$ starts at some vertex of $A(\pi)$. This means that $A(\pi)$ is a connected component of G , and so it coincides with the whole Rauzy class $C(\pi)$. \square

A result of Kontsevich, Zorich [14] yields a complete classification of the Rauzy classes. Here, let us calculate all Rauzy classes for the first few values of d . The results are summarized in the table at the end of this section.

For $d = 2$ there are two possibilities for the monodromy invariant, but only one is irreducible: $(2, 1)$. The Rauzy graph reduces to

$$0 \circlearrowleft \begin{pmatrix} A & B \\ B & A \end{pmatrix} \circlearrowright 1$$

For $d = 3$ there are six possibilities for the monodromy invariant, but only three are irreducible: $(2, 3, 1)$, $(3, 1, 2)$, $(3, 2, 1)$. They are all represented in the Rauzy class


$$0 \circlearrowleft \begin{pmatrix} A & C & B \\ C & B & A \end{pmatrix} \xrightarrow{1} \begin{pmatrix} A & B & C \\ C & B & A \end{pmatrix} \xrightarrow{0} \begin{pmatrix} A & B & C \\ C & A & B \end{pmatrix} \circlearrowright 1$$

So, there exists a unique Rauzy class for $d = 3$.


For $d = 4$ there are 24 possibilities for the monodromy invariant, 13 of which are irreducible:

$$\begin{aligned} &(4, 3, 2, 1), \quad (4, 1, 3, 2), \quad (3, 1, 4, 2), \quad (4, 2, 1, 3), \quad (2, 4, 3, 1), \\ &(3, 2, 4, 1), \quad (2, 4, 1, 3), \quad (4, 2, 3, 1), \quad (4, 1, 2, 3), \quad (4, 3, 1, 2), \\ &(3, 4, 1, 2), \quad (2, 3, 4, 1), \quad (3, 4, 2, 1) \end{aligned}$$

The following Rauzy class accounts for the first seven values:




The other six values of the monodromy invariant occur in the Rauzy class:



So, there are exactly two Rauzy classes for $d = 4$.

All these graphs are symmetric with respect to the vertical axis: this symmetry corresponds to the canonical involution, that is, to interchanging the roles of π_0 and π_1 . The last graph has an additional central symmetry: pairs that are opposite relative to the center have the same monodromy invariant, and so they correspond to essentially the same interval exchange transformation. Identifying such pairs, one obtains the corresponding *reduced Rauzy class*:



The Rauzy classes for $d \leq 5$ are listed below:

d	representative	# vertices (full class)	# vertices (reduced)
2	(2,1)	1	1
3	(3,2,1)	3	3
4	(4,3,2,1)	7	7
4	(4,2,3,1)	12	6
5	(5,4,3,2,1)	15	15
5	(5,3,2,4,1)		11
5	(5,4,2,3,1)		35
5	(5,2,3,4,1)		10

Standard pairs. A pair $\pi = (\pi_0, \pi_1)$ is called *standard* if the last symbol in each line coincides with the first symbol in the other line. In other words, the monodromy invariant satisfies

$$\pi_1 \circ \pi_0^{-1}(1) = d \quad \text{and} \quad \pi_1 \circ \pi_0^{-1}(d) = 1.$$

Inspection of the examples of Rauzy classes in section 6 shows that they all contain some standard pair. This turns out to be a general fact:

Proposition 6.2. *Every Rauzy class contains some standard pair.*

Notice that the Rauzy-Veech operator leaves the first symbols $\alpha_1^\varepsilon = \pi_\varepsilon^{-1}(1)$, in both top and bottom lines unchanged throughout the entire Rauzy class $C(\pi)$. The proof of Proposition 6.2 is based on the auxiliary lemma that we state below. The lemma can be easily deduced from Proposition 5.1, but we also give a short direct proof.

Lemma 6.3. *Given any $\varepsilon \in \{0, 1\}$ and any $\beta \in \mathcal{A}$ such that $\pi_\varepsilon(\beta) \neq 1$, there exists some pair π' in the Rauzy class $C(\pi)$ such that $\pi'_\varepsilon(\beta) = d$, that is, β is the last symbol in the line ε of π' .*

Proof. For each $\varepsilon \in \{0, 1\}$ let \mathcal{A}_ε be the subset of all $\beta \in \mathcal{A}$ such that $\pi'_\varepsilon(\beta) < d$ for every π' in the Rauzy class. In view of the previous remarks, $\alpha_1^\varepsilon \in \mathcal{A}_\varepsilon$. Let $\kappa(\varepsilon)$ be the rightmost position ever attained by these symbols, that is, the maximum value of $\pi'_\varepsilon(\beta)$ over all π' in $C(\pi)$ and $\beta \in \mathcal{A}_\varepsilon$. By definition, $\kappa(\varepsilon) < d$. Our goal is to prove that $\kappa(\varepsilon) = 1$, and so $\mathcal{A}_\varepsilon = \{\alpha_1^\varepsilon\}$, for both $\varepsilon \in \{0, 1\}$.

Fix any $\beta_\varepsilon \in \mathcal{A}_\varepsilon$ for which the maximum is attained. Then $\pi'_\varepsilon(\beta_\varepsilon) = \kappa(\varepsilon)$ for every π' in $C(\pi)$. That is because symbols γ with $\pi_\varepsilon(\gamma) < d$ can only move to the right under the Rauzy-Veech iteration and, were that to happen, it would contradict the assumption that $\kappa(\varepsilon)$ is maximum. Recall also Lemma 6.1. The same argument shows that all the symbols to the left of β_ε are also constant on the Rauzy class:

$$(\pi'_\varepsilon)^{-1}(i) = \pi_\varepsilon^{-1}(i) \quad \text{for all } 1 \leq i \leq \kappa(\varepsilon). \tag{24}$$

In particular, no symbol to the left of β_ε on the line ε can ever reach the last position in the line $1 - \varepsilon$:

$$\pi_\varepsilon(\alpha) < \kappa(\varepsilon) \quad \Rightarrow \quad \pi'_{1-\varepsilon}(\alpha) < d \quad \Rightarrow \quad \pi'_{1-\varepsilon}(\alpha) \leq \kappa(1 - \varepsilon), \tag{25}$$

for any pair π' in $C(\pi)$. Let us write

$$\pi' = \begin{pmatrix} \alpha_1^0 & \cdots & \alpha_{\kappa(0)}^0 & \cdots & \cdots & \alpha_d^0 \\ \alpha_1^1 & \cdots & \cdots & \alpha_{\kappa(1)}^1 & \cdots & \alpha_d^1 \end{pmatrix}, \quad \alpha_i^\varepsilon = (\pi'_\varepsilon)^{-1}(i).$$

In view of (24), the relation (25) implies

$$\{\alpha_1^\varepsilon, \dots, \alpha_{\kappa(\varepsilon)-1}^\varepsilon\} \subset \{\alpha_1^{1-\varepsilon}, \dots, \alpha_{\kappa(1-\varepsilon)}^{1-\varepsilon}\} \quad \text{for } \varepsilon \in \{0, 1\}. \tag{26}$$

In particular, $\kappa(\varepsilon) - 1 \leq \kappa(1 - \varepsilon) \leq \kappa(\varepsilon) + 1$. There are four possibilities:

- (i) $\kappa(0) = \kappa(1) + 1$: then the case $\varepsilon = 0$ of (26) implies $\{\alpha_1^0, \dots, \alpha_{\kappa(1)}^0\} = \{\alpha_1^1, \dots, \alpha_{\kappa(1)}^1\}$, and this contradicts the assumption of irreducibility.
- (ii) $\kappa(0) = \kappa(1) - 1$: this is analogous to the first case, using the case $\varepsilon = 1$ in (26) instead.
- (iii) $\kappa(0) = \kappa(1)$ and $\{\alpha_1^0, \dots, \alpha_{\kappa(0)-1}^0\} = \{\alpha_1^1, \dots, \alpha_{\kappa(1)-1}^1\}$: this also contradicts irreducibility, unless $\kappa(0) = \kappa(1) = 1$.
- (iv) $\kappa(0) = \kappa(1)$ and there exists $1 \leq i < \kappa(0)$ such that $\alpha_i^0 = \alpha_{\kappa(1)}^1$: together with the case $\varepsilon = 1$ of (26), this gives

$$\{\alpha_1^1, \dots, \alpha_{\kappa(1)-1}^1, \alpha_{\kappa(1)}^1\} = \{\alpha_1^0, \dots, \alpha_{\kappa(0)}^0\}$$

and this implies that the two sets coincide (hence, there exists $1 \leq j < \kappa(1)$ such that $\alpha_j^1 = \alpha_{\kappa(0)}^0$). Once more, this contradicts irreducibility.

This completes the proof of the lemma. □

Now we can give the proof of Proposition 6.2:

Proof. As observed before, the first symbols α_1^{ε} in both lines remain unchanged under Rauzy-Veech iteration. By irreducibility, they are necessarily distinct. So, using Lemma 6.3, we may find a pair π' in $C(\pi)$ such that $\pi'_0(\alpha_1^1) = d$, that is, the last symbol in the top line coincides with the first one in the bottom line. Now, iterating π' under type 0 Rauzy-Veech map, we keep the top line unchanged, while rotating all the symbols in the bottom line to the right of α_1^1 . So, we eventually reach a pair π'' which satisfies $\pi''_1(\alpha_1^0) = d$, in addition to $\pi''_0(\alpha_1^1) = d$. Then π'' is standard. □

7. Rauzy-Veech renormalization

We are especially interested in a variation of the induction algorithm where one scales the domains of all interval exchange transformations to length 1.

Let π and π' be irreducible pairs such that π' is the type ε successor of π , for $\varepsilon \in \{0, 1\}$. For each $\lambda \in \mathbb{R}_+^A$ satisfying

$$\lambda_{\alpha(\varepsilon)} > \lambda_{\alpha(1-\varepsilon)} \tag{27}$$

we have

$$\hat{R}(\pi, \lambda) = (\pi', \lambda') \quad \text{with} \quad \lambda'_\alpha = \begin{cases} \lambda_\alpha & \text{if } \alpha \neq \alpha(\varepsilon), \\ \lambda_{\alpha(\varepsilon)} - \lambda_{\alpha(1-\varepsilon)} & \text{if } \alpha = \alpha(\varepsilon). \end{cases}$$

The map $\lambda \mapsto \lambda'$ thus defined is a bijection from the set of length vectors satisfying (27) to the whole \mathbb{R}_+^A : the inverse is given by

$$\lambda_\alpha = \begin{cases} \lambda'_\alpha & \text{if } \alpha \neq \alpha(\varepsilon), \\ \lambda'_{\alpha(\varepsilon)} + \lambda'_{\alpha(1-\varepsilon)} & \text{if } \alpha = \alpha(\varepsilon). \end{cases}$$

Take the interval I to have unit length, that is, $\sum_{\alpha \in A} \lambda_\alpha = 1$. The induction $\hat{R}(f)$ is defined on a shorter interval, with length $1 - \lambda_{\alpha(1-\varepsilon)}$, but after appropriate rescaling we may see it as a map $R(f)$ on a unit interval. This means we are now considering

$$R : (\pi, \lambda) \mapsto (\pi', \lambda''), \quad \text{where} \quad \lambda'' = \frac{\lambda'}{1 - \lambda_{\alpha(1-\varepsilon)}},$$

that we refer to as the *Rauzy-Veech renormalization map*. Let Λ_A be the set of all length vectors $\lambda \in \mathbb{R}_+^A$ with $\sum_{\alpha \in A} \lambda_\alpha = 1$, and let

$$\Lambda_{\pi, \varepsilon} = \{ \lambda \in \Lambda_A : \lambda_{\alpha(\varepsilon)} > \lambda_{\alpha(1-\varepsilon)} \} \quad \text{for } \varepsilon \in \{0, 1\}.$$




Figure 8




Figure 9

The previous observations mean that $(\pi, \lambda) \mapsto (\pi', \lambda'')$ maps $\{\pi\} \times \Lambda_{\pi, \varepsilon}$ bijectively onto $\{\pi'\} \times \Lambda_{\mathcal{A}}$. Figure 8 illustrates the case $d = 3$:

For each Rauzy class C we have a map $R : (\pi, \lambda) \mapsto (\pi', \lambda'')$ from $C \times \Lambda_{\mathcal{A}}$ to itself, (More precisely, this map is defined on the full Lebesgue measure subset of length vectors λ that satisfy the Keane condition.) with the following Markov property: R sends each $\{\pi\} \times \Lambda_{\pi, \varepsilon}$ bijectively onto $\{\pi'\} \times \Lambda_{\mathcal{A}}$, where π' is the type ε successor of π . Note that

$$\lambda'' = \frac{\Theta^{-1*}(\lambda)}{1 - \lambda_{\alpha(1-\varepsilon)}} \tag{28}$$

and the operator Θ depends only on π and the type ε , that is, it is constant on each $\{\pi\} \times \Lambda_{\pi, \varepsilon}$.

Example 7.1. For $d = 2$ there is only one pair, $\pi = \begin{pmatrix} A & B \\ B & A \end{pmatrix}$. We have

$$\Lambda_{\mathcal{A}} = \{(\lambda_A, \lambda_B) : \lambda_A > 0, \lambda_B > 0, \text{ and } \lambda_A + \lambda_B = 1\} \sim (0, 1),$$

where \sim refers to the bijective correspondence $(\lambda_A, \lambda_B) \mapsto x = \lambda_A$. Under this correspondence, $\Lambda_{\pi, 0} \sim (0, 1/2)$ and $\Lambda_{\pi, 1} \sim (1/2, 1)$, and the Rauzy-Veech renormalization $(\pi, \lambda) \mapsto (\pi, \lambda'')$ is given by (see figure 9)

$$r(x) = \begin{cases} x/(1-x) & \text{for } x \in (0, 1/2), \\ 2 - 1/x & \text{for } x \in (1/2, 1). \end{cases}$$




Figure 10

Observe that r has a tangency of order 1 with the identity at $x = 0$ and $x = 1$.

The following fundamental result was proved independently by Masur [17] and Veech [22]. A proof will be given in last part of this paper. Let $d\pi$ denote the counting measure in the set of pairs π , and Leb be the Lebesgue measure (of dimension $d - 1$) in the simplex $\Lambda_{\mathcal{A}}$.

Theorem 7.2. *For each Rauzy class C , the Rauzy-Veech renormalization map $R : C \times \Lambda_{\mathcal{A}} \rightarrow C \times \Lambda_{\mathcal{A}}$ admits an invariant measure ν which is absolutely continuous with respect to $d\pi \times \text{Leb}$. This measure ν is unique, up to product by a scalar, and ergodic. Moreover, its density with respect to Lebesgue measure is given by a homogeneous rational function of degree $-d$ and bounded away from zero.*

8. Zorich transformations

In general, the measures ν in Theorem 7.2 have infinite mass. For instance, it is well-known that for maps with neutral fixed points such as the one in Example 7.1, absolutely continuous invariant measures are necessarily infinite. Zorich [25] introduced an accelerated version of the Rauzy-Veech algorithm for which there exists a (unique) invariant *probability* absolutely continuous with respect to Lebesgue measure on each simplex $\Lambda_{\mathcal{A}}$. This is defined as follows.

Let C be a Rauzy class, $\pi = (\pi_0, \pi_1)$ be a vertex of C , and $\lambda \in \mathbb{R}_+^{\mathcal{A}}$ satisfy the Keane condition. Let $\varepsilon \in \{0, 1\}$ be the type of (π, λ) and, for each $j \geq 1$, let $\varepsilon^{(j)}$ be the type of the iterate $(\pi^{(j)}, \lambda^{(j)}) = \hat{R}^j(\pi, \lambda)$. Then define $n = n(\pi, \lambda) \geq 1$ to be smallest such that $\varepsilon^{(j)} \neq \varepsilon$. The *Zorich induction map* is defined by

$$\hat{Z}(\pi, \lambda) = (\pi^{(n)}, \lambda^{(n)}) = \hat{R}^n(\pi, \lambda).$$

We also consider the *Zorich renormalization map*

$$Z : C \times \Lambda_{\mathcal{A}} \rightarrow C \times \Lambda_{\mathcal{A}}, \quad Z(\pi, \lambda) = R^n(\pi, \lambda).$$




Figure 11

This map admits a Markov partition, into countably many domains. Indeed, for each π in the Rauzy class and $\varepsilon \in \{0, 1\}$, let

$$\Lambda_{\pi, \varepsilon, n}^* = \{ \lambda \in \Lambda_{\pi, \varepsilon} : \varepsilon^{(1)} = \dots = \varepsilon^{(n-1)} = \varepsilon \neq \varepsilon^{(n)} \}.$$

Then Z maps every $\{\pi\} \times \Lambda_{\pi, \varepsilon, n}^*$ bijectively onto $\{\pi^{(n)}\} \times \Lambda_{\pi, 1-\varepsilon}$. Moreover, by (28),

$$\lambda^{(n)} = c_n \Theta^{-n*}(\lambda)$$

where $c_n > 0$ and Θ^{-n*} depend only on π, ε, n , that is, they are constant on each $\{\pi\} \times \Lambda_{\pi, \varepsilon, n}^*$.

Example 8.1. For $d = 2$ (recall Example 7.1), the Zorich transformation Z is described by the map $z(x) = r^n(x)$ where $n = n(x) \geq 1$ is the smallest integer such that

$$r^n(x) \in (1/2, 1), \text{ if } x \in (0, 1/2) \quad \text{or} \quad r^n(x) \in (0, 1/2), \text{ if } x \in (1/2, 1).$$

See figure 11. This map is Markov and uniformly expanding (the latter is specific to $d = 2$). It is well-known that such maps admit absolutely continuous invariant probabilities.

We shall also prove the following result, where the main novelty is that the invariant measure μ is finite:

Theorem 8.2. *For each Rauzy class C , the Zorich renormalization map $Z : C \times \Lambda_{\mathcal{A}} \rightarrow C \times \Lambda_{\mathcal{A}}$ admits an invariant probability measure μ which is absolutely continuous with respect to $d\pi \times \text{Leb}$. This probability μ is unique and ergodic. Moreover, its density with respect to Lebesgue measure is given by a homogeneous rational function of degree $-d$ and bounded away from zero.*




Figure 12

9. Continued fractions

The classical continued fraction algorithm associates to each irrational number $x_0 \in (0, 1)$ the sequences

$$n_k = \left[\frac{1}{x_{k-1}} \right] \quad \text{and} \quad x_k = \frac{1}{x_{k-1}} - n_k,$$

where $[\cdot]$ denotes the integer part. Observe that

$$x_0 = \frac{1}{n_1 + x_1} = \frac{1}{n_1 + \frac{1}{n_2 + x_2}} = \frac{1}{n_1 + \frac{1}{n_2 + \frac{1}{n_3 + x_3}}} = \dots$$

The algorithm may also be written as

$$x_k = G^k(x_0) \quad \text{and} \quad n_k = \left[\frac{1}{x_{k-1}} \right]$$

where G is the Gauss map (see figure 12)

$$G : (0, 1) \rightarrow [0, 1], \quad G(x) = \frac{1}{x} - \left[\frac{1}{x} \right]$$

The Gauss map is very much equivalent to the Zorich transformation for $d = 2$ (and so the cases $d > 2$ of the Zorich transformation may be seen as higher dimensional generalizations of the classical continued fraction expansion). To see this, consider the bijection

$$\phi : (\lambda_A, \lambda_B) \mapsto y = \frac{\lambda_A}{\lambda_B}$$

from Λ_A to $(0, \infty)$. Moreover, let P be the bijection of Λ_A defined by $P : (\lambda_A, \lambda_B) \mapsto (\lambda_B, \lambda_A)$. Consider (λ_A, λ_B) in $\Lambda_{\pi,0}$, that is, such that $\lambda_A < \lambda_B$. Then $y = \phi(\lambda_A, \lambda_B) \in (0, 1)$. By definition,

$$\hat{Z} \circ P(\lambda_A, \lambda_B) = \hat{Z}(\lambda_B, \lambda_A) = (\lambda_B - n\lambda_A, \lambda_A)$$

where n is the integer part of λ_B/λ_A . In terms of the variable y , this corresponds to

$$y \mapsto \frac{1}{y} - n = G(y).$$

In other words, we have just shown that ϕ conjugates $Z \circ P$, restricted to $\Lambda_{\pi,0}$, to the Gauss map G . Consequently, ϕ conjugates $(Z \circ P)^n$, restricted to $\Lambda_{\pi,0}$, to G^n , for every $n \geq 1$. Observe that $P^2 = \text{id}$ and Z commutes with P . (In other words, P conjugates the restriction of Z to $\Lambda_{\pi,0}$ to the restriction of Z to $\Lambda_{\pi,1}$.) Hence, we have shown that $Z^{2k} \mid \Lambda_{\pi,0}$ is conjugate to G^{2k} , and $Z^{2k-1} \circ P \mid \Lambda_{\pi,0}$ is conjugate to G^{2k-1} , for every $k \geq 1$.

10. Symplectic form

It is clear from (3) that the operator $\Omega_\pi : \mathbb{R}^A \rightarrow \mathbb{R}^A$ is anti-symmetric, that is,

$$\Omega_\pi^* = -\Omega_\pi \tag{29}$$

where Ω_π^* is the adjoint operator, relative to the Euclidean metric \cdot on \mathbb{R}^A . Thus,

$$\tilde{\omega}_\pi : \mathbb{R}^A \times \mathbb{R}^A \rightarrow \mathbb{R}, \quad \tilde{\omega}_\pi(u, v) = u \cdot \Omega_\pi(v)$$

defines an alternate bilinear form on \mathbb{R}^d . In general, this bilinear form is degenerate: indeed, if $v \in \ker \Omega_\pi$ then $\tilde{\omega}_\pi(u, v) = 0$ for every $u \in \mathbb{R}^A$. On the other hand, there is always a naturally associated non-degenerate bilinear form ω_π on the subspace $H_\pi = \Omega_\pi(\mathbb{R}^A)$, defined by

$$\omega_\pi : H_\pi \times H_\pi \rightarrow \mathbb{R}, \quad \omega_\pi(\Omega_\pi(u), \Omega_\pi(v)) = u \cdot \Omega_\pi(v)$$

Lemma 10.1. *The previous relation defines a symplectic form, that is, a non-degenerate alternate bilinear form ω_π on H_π .*

Proof. The relation (29) implies that the orthogonal complement H_π^\perp coincides with $\ker \Omega_\pi$. Suppose $\Omega_\pi(u) = \Omega_\pi(u')$. Then $u - u' \in \ker \Omega_\pi$ and so

$$u \cdot \Omega_\pi(v) = u' \cdot \Omega_\pi(v)$$

for every $v \in \mathbb{R}^A$. This shows that ω_π is well-defined. It is clear that it is bilinear. The fact that ω_π is alternate is an immediate consequence of (29):

$$\omega_\pi(\Omega_\pi(v), \Omega_\pi(u)) = v \cdot \Omega_\pi(u) = u \cdot \Omega_\pi^*(v) = -\omega_\pi(\Omega_\pi(u), \Omega_\pi(v)).$$

Finally, it is also easy to see that ω_π is non-degenerate:

$$\omega_\pi(\Omega_\pi(u), \Omega_\pi(v)) = 0 \quad \forall v \iff u \cdot \Omega_\pi(v) = 0 \quad \forall v \iff u \in H_\pi^\perp$$

and, since we are taking $u \in H_\pi$, this can only happen if u vanishes. □

Lemma 10.2. *If $(\pi', \lambda') = \hat{R}(\pi, \lambda)$ then $\Theta \Omega_\pi \Theta^* = \Omega_{\pi'}$, where $\Theta = \Theta_{\pi, \lambda}$. In particular, the operator Θ induces a symplectic isomorphism from H_π onto $H_{\pi'}$, relative to the symplectic forms in the two spaces.*

Proof. Let $\lambda' \in \mathbb{R}^A$ be given by $\lambda = \Theta^*(\lambda')$; compare (11). Then define $w = \Omega_\pi(\lambda)$ and $w' = \Omega_{\pi'}(\lambda')$; compare (2). We have seen in (8) that $w' = \Theta(w)$. That is,

$$\Omega_{\pi'}(\lambda') = \Theta \Omega_\pi \Theta^*(\lambda')$$

for all $\lambda' \in \mathbb{R}^A$. This proves the first claim. Next, the relation $\Theta \Omega_\pi \Theta^* = \Omega_{\pi'}$ together with the fact that Θ and Θ^* are invertible imply that $u \in H_\pi$ if and only if $\Theta(u)$ is in $H_{\pi'}$. Moreover, the operator $\Theta : H_\pi \rightarrow H_{\pi'}$ is symplectic:

$$\begin{aligned} \omega_{\pi'}(\Theta \Omega_\pi(u), \Theta \Omega_\pi(v)) &= \omega_{\pi'}(\Omega_{\pi'} \Theta^{-1*}(u), \Theta \Omega_\pi(v)) \\ &= \Theta^{-1*}(u) \cdot \Theta \Omega_\pi(v) = u \cdot \Omega_\pi(v) = \omega_\pi(\Omega_\pi(u), \Omega_\pi(v)), \end{aligned}$$

for any vectors $u, v \in \mathbb{R}^A$. □

Remark 10.3. Lemma 10.1 implies that the space H_π has even dimension: we write $\dim H_\pi = 2g$. Since Θ is always an isomorphism from H_π to $H_{\pi'}$, it follows that the dimension is constant on the whole Rauzy class. We shall later interpret g as the genus of an orientable surface canonically associated to the Rauzy class.

11. Translation surfaces

An *Abelian differential* α is a holomorphic complex 1-form on a Riemann surface. We assume the Riemann surface is compact, and α is not identically zero. Then it has a finite number of zeroes, that we call singularities. In local coordinates, $\alpha_z = \varphi(z) dz$ for some $\varphi(z) \in \mathbb{C}$ that depends holomorphically on the point z . Near any non-singular point p , one can always find so-called *adapted coordinates* ζ relative to which the Abelian differential takes the form $\alpha_\zeta = d\zeta$: it suffices to take

$$\zeta = \int_p^z \varphi(\xi) d\xi$$

If p is a singularity, with multiplicity $m \geq 1$ say, then one considers instead

$$\zeta = (m + 1) \left(\int_p^z \varphi(\xi) d\xi \right)^{1/(m+1)} : \tag{30}$$




Figure 13

in these coordinates $\alpha_\zeta = \zeta^m d\zeta$. Notice that all changes of adapted coordinates near a regular point are given by translations: if ζ and ζ' are adapted coordinates then $d\zeta' = d\zeta$, and so $\zeta' = \zeta + \text{const}$. We say that the adapted coordinates form a *translation atlas*, and call the resulting structure a *translation surface*. Coordinate changes near singularities are slightly more subtle. If ζ' is a non-singular adapted coordinate and ζ is a singular one, then $d\zeta' = \zeta^m d\zeta$ or, in other words, $(m + 1)\zeta' = \zeta^{m+1} + \text{const}$. Figure 13 illustrates this relation between the two types of coordinates.

The translation atlas defines a flat (zero curvature) Riemannian metric on the surface minus the singularities, transported from the complex plane through the adapted charts. The form of (30) gives that the zeroes of the Abelian differential are conical singularities: in appropriate polar coordinates (ρ, θ) , $\rho > 0$, $\theta \in [0, 2\pi)$ centered at the singularity, the Riemannian metric is given by

$$ds^2 = d\rho^2 + (c\rho d\theta)^2, \quad \text{where } c = m + 1.$$

In addition, the translation atlas defines a parallel unit vector field on the complement of the singularities, namely, the pull-back of the vertical vector field under the local charts.

Conversely, a flat metric with finitely many singularities, of conical type, together with a parallel unit vector field X , completely determine a translation structure. Indeed, the neighborhood of any regular point p is isometric to an open subset of \mathbb{C} . Choose the isometry so that it sends the vector X_p to the vertical vector $(0, 1)$. Then the isometry is uniquely determined, and sends X to the constant vector field $(0, 1)$. In particular, these isometries coincide in the intersection of their domains, and so they define a Riemann surface atlas on the complement of the singularities. Moreover, they transport the canonical Abelian differential dz from \mathbb{C} to the surface.

Construction of translation surfaces. Let us describe a simple construction of translation surfaces. In fact, this construction is general: every translation surface can be obtained in the way we are going to describe.




Figure 14

Consider a polygon in \mathbb{R}^2 having an even number $2d \geq 4$ of sides

$$s_1, \dots, s_d, s'_1, \dots, s'_d$$

such that s_i and s'_i are parallel (non-adjacent) and have the same length, for every $i = 1, \dots, d$. See figure 14 for an example with $d = 4$.

Identifying s_i with s'_i by translation, for each $i = 1, \dots, d$, we obtain a translation surface M : the singularities correspond to the points obtained by identification of the vertices of the polygon; the Abelian differential and the flat metric are inherited from $\mathbb{R}^2 = \mathbb{C}$, and the vertical vector field $X = (0, 1)$ is parallel.

Let a_1, \dots, a_κ , $\kappa = \kappa(\pi)$ be the singularities. The *angle* of a singularity a_i is the topological index around zero

$$\text{angle}(a_i) = 2\pi \text{ind}(\beta, 0) = \frac{1}{i} \int_0^1 \frac{\dot{\beta}(t)}{\beta(t)} dt$$

of the curve $\beta(t) = \alpha_{\gamma(t)}(\dot{\gamma}(t))$, where $\gamma : [0, 1] \rightarrow M$ is any small simple closed curve around a_i . It is clear that

$$\text{angle}(a_i) = 2\pi(m_i + 1) \tag{31}$$

where m_i denotes the order of the zero of α at a_i . We call the singularity *removable* if the angle is exactly 2π , that is, if a_i is actually not a zero of α .

Let the translation surface be constructed from a planar polygon with $2d$ sides, as described above. Then the sum of all angles at the singularities coincides with the sum of the internal angles of the $2d$ -gon, that is

$$\sum_{i=1}^{\kappa} \text{angle}(a_i) = 2\pi(d - 1). \tag{32}$$

Using (31) we deduce that

$$\sum_{i=1}^{\kappa} m_i = d - \kappa - 1.$$




Figure 15

The angles are also related to the genus $g(M)$ and the Euler characteristic $\mathcal{X}(M) = 2 - 2g(M)$ of the surface M . To this end, consider a decomposition into $4d$ triangles as described in figure 15: a central point is linked to the vertices of the polygon and to the midpoint of every side.

Recall that the sides of the polygon are identified pairwise. So, this decomposition has $6d$ edges, $2d$ of them corresponding to segments inside the sides of the polygon. Moreover, there are $d + \kappa + 1$ vertices: the central one, plus d vertices coming from the midpoints of the polygon sides, and κ more sitting at the singularities. Therefore,

$$2 - 2g(M) = \mathcal{X}(M) = \kappa + 1 - d. \tag{33}$$

From (32) and (33), we obtain a kind of Gauss-Bonnet theorem for these flat surfaces:

$$\sum_{i=1}^{\kappa} [2\pi - \text{angle}(a_i)] = -2\pi \sum_{i=1}^{\kappa} m_i = 2\pi(\kappa + 1 - d) = 2\pi\mathcal{X}(M). \tag{34}$$

In fact this is the only restriction imposed on the orders of the singularities by the topology of the surface: given any $g \geq 1$ and integers $m_i \geq 0$, $i = 1, \dots, \kappa$ with $\sum_{i=1}^{\kappa} m_i = 2g - 2$, there exists some translation surface with κ singularities of orders m_1, \dots, m_{κ} .

12. Suspending interval exchange maps

Let π be an irreducible pair and $\lambda \in \mathbb{R}_+^{\mathcal{A}}$ be a length vector. We denote by T_{π}^+ the subset of vectors $\tau = (\tau_{\alpha})_{\alpha \in \mathcal{A}} \in \mathbb{R}^{\mathcal{A}}$ such that

$$\sum_{\pi_0(\alpha) \leq k} \tau_{\alpha} > 0 \quad \text{and} \quad \sum_{\pi_1(\alpha) \leq k} \tau_{\alpha} < 0 \tag{35}$$

for all $1 \leq k \leq d - 1$. Clearly, T_{π}^+ is a convex cone. We say that τ has *type 0* if the total sum $\sum_{\alpha \in \mathcal{A}} \tau_{\alpha}$ is positive and *type 1* if the total sum is negative. Define




Figure 16

$\zeta_\alpha = (\lambda_\alpha, \tau_\alpha) \in \mathbb{R}^2$ for each $\alpha \in \mathcal{A}$. Then consider the closed curve $\Gamma = \Gamma(\pi, \lambda, \tau)$ on \mathbb{R}^2 formed by concatenation of

$$\zeta_{\alpha_1^0}, \zeta_{\alpha_2^0}, \dots, \zeta_{\alpha_d^0}, -\zeta_{\alpha_d^1}, -\zeta_{\alpha_{d-1}^1}, \dots, -\zeta_{\alpha_1^1}$$

with starting point at the origin. Condition (35) means that the endpoints of all $\zeta_{\alpha_1^0} + \dots + \zeta_{\alpha_k^0}$ are on the upper half plane, and the endpoints of all $\zeta_{\alpha_1^1} + \dots + \zeta_{\alpha_k^1}$ are in the lower half plane, for every $1 \leq k \leq d - 1$. See figure 16.

Assume, for the time being, that this closed curve Γ is simple. Then it defines a planar polygon with $2d$ sides organized in pairs of parallel segments with the same length, as considered in the previous section. The *suspension surface* $M = M(\pi, \lambda, \tau)$ is the translation surface obtained by identification of the sides in each of the pairs. Let $I \subset M$ be the horizontal segment of length $\sum_{\alpha \in \mathcal{A}} \lambda_\alpha$ with the origin as left endpoint, that is,

$$I = \left[0, \sum_{\alpha \in \mathcal{A}} \lambda_\alpha \right) \times \{0\}. \tag{36}$$

The interval exchange transformation f defined by (π, λ) corresponds to the first return map to I of the vertical flow on M . To see this, for each $\alpha \in \mathcal{A}$, let

$$I_\alpha = \left[\sum_{\pi_0(\beta) < \pi_0(\alpha)} \lambda_\beta, \sum_{\pi_0(\beta) \leq \pi_0(\alpha)} \lambda_\beta \right) \times \{0\}$$

Consider the vertical segment starting from $(x, 0) \in I_\alpha$ and moving upwards. It hits the side represented by ζ_α at some point (x, z) . This is identified with the point (x', z') in the side represented by $-\zeta_\alpha$, given by

$$\begin{aligned} x' &= x - \sum_{\pi_0(\beta) < \pi_0(\alpha)} \lambda_\beta + \sum_{\pi_1(\beta) < \pi_1(\alpha)} \lambda_\beta = x + w_\alpha \\ z' &= z - \sum_{\pi_0(\beta) < \pi_0(\alpha)} \tau_\beta + \sum_{\pi_1(\beta) < \pi_1(\alpha)} \tau_\beta = z - h_\alpha. \end{aligned} \tag{37}$$




Figure 17

($h_\alpha > 0$ is defined by the last equality.) Continuing upwards from (x', z') we hit I back at the point $(x', 0)$. This shows that the return map does coincide with $f(x) = x + w_\alpha$ on each I_α .

In some fairly exceptional situations, such as in figure 17, the closed curve Γ may have self-intersections. It is easy to extend the definition of the suspension surface to this case: just consider the simple polygon obtained by removing the self-intersections in the way described in the figure, and then take the translation surface M obtained by identification of parallel sides of this polygon. The horizontal segment I may still be viewed as a cross-section to the vertical flow on M , and the corresponding first return map coincides with the interval exchange transformation f .

We are going to focus our presentation on the case when Γ is simple and, in general, let the reader to adapt the arguments to the case when there are self-intersections. In some sense, the non-simple case can be avoided altogether:

Remark 12.1. The curve $\Gamma(\pi, \lambda, \tau)$ can have self-intersections only if either

$$\sum_{\alpha \in \mathcal{A}} \tau_\alpha > 0 \quad \text{and} \quad \lambda_{\alpha(0)} < \lambda_{\alpha(1)}, \quad \text{i.e., } \tau \text{ has type 0 and } (\pi, \lambda) \text{ has type 1,}$$

as is the case in figure 17, or

$$\sum_{\alpha \in \mathcal{A}} \tau_\alpha < 0 \quad \text{and} \quad \lambda_{\alpha(0)} > \lambda_{\alpha(1)}, \quad \text{i.e., } \tau \text{ has type 1 and } (\pi, \lambda) \text{ has type 0.}$$

In other words, if (π, λ) and τ have the same type then the curve $\Gamma(\pi, \lambda, \tau)$ is necessarily simple. Using this observation, we shall see in Remark 18.3 that by Rauzy-Veech induction one eventually finds data $(\pi^n, \lambda^n, \tau^n)$ that represents the same translation surface and for which the curve $\Gamma(\pi^n, \lambda^n, \tau^n)$ has no self-intersections.

13. Some translation surfaces

We shall see later that the type (genus and singularities) of the translation surface $M = M(\pi, \lambda, \tau)$ depends only on the Rauzy class of π . Here we consider a representative of each Rauzy class with $d \leq 5$, and we exhibit the corresponding translation




Figure 18




Figure 19

surface for generic vectors λ and τ . The conclusions are summarized in the table near the end of this section.

For $d = 2$ and $\pi = \begin{pmatrix} A & B \\ B & A \end{pmatrix}$, corresponding to monodromy invariant $p = (2, 1)$, the four vertices are identified to a single point a , and $\text{angle}(a) = 2\pi$. Using (34) we conclude that M is the torus, (and the singularity is removable). See figure 18.

For $d = 3$ and $\pi = \begin{pmatrix} A & B & C \\ C & B & A \end{pmatrix}$, corresponding to $p = (3, 2, 1)$, the six vertices are identified to two different points, with $\text{angle}(a) = \text{angle}(b) = 2\pi$. Thus, M is the torus, and both singularities are removable. See figure 19.

For $d = 4$ and $\pi = \begin{pmatrix} A & B & C & D \\ D & C & B & A \end{pmatrix}$, corresponding to $p = (4, 3, 2, 1)$, the eight vertices are identified to a single point, with $\text{angle}(a) = 6\pi$. Thus, M has genus 2 (bitorus). See figure 20.

For $d = 4$ and $\pi = \begin{pmatrix} A & B & C & D \\ D & B & C & A \end{pmatrix}$, hence $p = (4, 2, 3, 1)$, the vertices are identified to three different points, with $\text{angle}(a) = \text{angle}(b) = \text{angle}(c) = 2\pi$. M is the torus, and all singularities are removable. See figure 21.

For $d = 5$ and $\pi = \begin{pmatrix} A & B & C & D & E \\ E & D & C & B & A \end{pmatrix}$, hence $p = (5, 4, 3, 2, 1)$, the ten vertices are identified to two different points, a and b , with $\text{angle}(a) = \text{angle}(b) = 4\pi$. Thus, M




Figure 20




Figure 21

is the bitorus ($g = 2$).

For $d = 5$ and $\pi = \begin{pmatrix} A & B & C & D & E \\ E & C & B & D & A \end{pmatrix}$, hence $p = (5, 3, 2, 4, 1)$, the vertices are identified to two different points, a and b , with angle $(a) = 2\pi$ and angle $(b) = 6\pi$. M is, again, the bitorus.

For $d = 5$ and $\pi = \begin{pmatrix} A & B & C & D & E \\ E & D & B & C & A \end{pmatrix}$, hence $p = (5, 4, 2, 3, 1)$, the vertices are identified to two different points, a and b , with angle $(a) = 6\pi$ and angle $(b) = 2\pi$. M is, once more, the bitorus.

For $d = 5$ and $\pi = \begin{pmatrix} A & B & C & D & E \\ E & B & C & D & A \end{pmatrix}$, hence $p = (5, 2, 3, 4, 1)$, the vertices are identified to four different points, with angle $(a) = \text{angle}(b) = \text{angle}(c) = \text{angle}(d) = 2\pi$. M is the torus and all singularities are removable.

Summarizing, we have:

d	representative	# vertices	angles	orders	genus	\mathcal{X}
2	(2,1)	1	2π	0	1	0
3	(3,2,1)	3	$2\pi, 2\pi$	0, 0	1	0
4	(4,3,2,1)	7	6π	2	2	-2
4	(4,2,3,1)	8	$2\pi, 2\pi, 2\pi$	0, 0, 0	1	0
5	(5,4,3,2,1)	15	$4\pi, 4\pi$	1, 1	2	-2
5	(5,3,2,4,1)	11	$6\pi, 2\pi$	2, 0	2	-2
5	(5,4,2,3,1)	35	$6\pi, 2\pi$	2, 0	2	-2
5	(5,2,3,4,1)	10	$2\pi, 2\pi, 2\pi, 2\pi$	0, 0, 0, 0	1	0

Remark 13.1. Starting from $d = 5$, different Rauzy classes may give rise to translation surfaces with the same number and orders of singularities.

14. Computing the suspension surface

Let us explain how the number κ and the orders m_i of the singularities may be computed from π , in general. Consider the set of all pairs (α, S) with $\alpha \in \mathcal{A}$ and $S \in \{L, R\}$. We think of (α, L) and (α, R) as representing, respectively, the origin (left endpoint) and the end (right endpoint) of the sides of the polygon labelled by α . Then, under the identifications that define the suspension surface, one must identify

$$(\alpha, R) \sim (\beta, L) \quad \text{if} \quad \pi_0(\alpha) + 1 = \pi_0(\beta) \tag{38}$$

$$(\alpha, R) \sim (\beta, L) \quad \text{if} \quad \pi_1(\alpha) + 1 = \pi_1(\beta) \tag{39}$$

and also

$$(\alpha, L) \sim (\beta, L) \quad \text{if} \quad \pi_0(\alpha) = 1 = \pi_1(\beta) \tag{40}$$

$$(\alpha, R) \sim (\beta, R) \quad \text{if} \quad \pi_0(\alpha) = d = \pi_1(\beta). \tag{41}$$

Extend \sim to an equivalence relation in the set of pairs (α, S) . Then the number κ of singularities is, precisely, the number of equivalence classes for this relation.

Figure 22 describes a specific case with $d = 7$:

$$\pi = \begin{pmatrix} A & B & C & D & E & F & G \\ G & F & E & D & C & B & A \end{pmatrix}.$$




Figure 22

There are two equivalence classes:

$$(A, L) \sim (B, R) \sim (C, L) \sim (D, R) \sim (E, L) \sim (F, R) \sim (G, L) \sim (A, L)$$

and

$$(A, R) \sim (B, L) \sim (C, R) \sim (D, L) \sim (E, R) \sim (F, L) \sim (G, R) \sim (A, R)$$

It is also easy to guess what the angles of these singularities are. For instance, consider the singularity a associated to the first equivalence class (the other one is analogous). The angle corresponds to the sum of the internal angles of the polygon at the 9 vertices that are identified to a . This sum is readily computed by noting that the arcs describing these internal angles cut the vertical direction exactly 6 times: one for each vertex, except for the exceptional $(A, L) = (G, L)$. See figure 22. Thus, $\text{angle}(a) = 6\pi$ and the singularity has order 2.

The general rule can be formulated as follows. Let us call *irregular pairs* to

$$(\pi_0^{-1}(1), L), \quad (\pi_1^{-1}(1), L), \quad (\pi_0^{-1}(d), R), \quad (\pi_1^{-1}(d), R).$$

All other pairs are called *regular*. Then there is an even number $2k$ of regular pairs in each equivalence class (one half above the horizontal axis and the other half below), and the angle of the corresponding singularity is equal to $2k\pi$.

This calculation remains valid when the curve $\Gamma(\pi, \lambda, \tau)$ has self-intersections. Let us explain this in the case when τ has type 0, the other one being symmetric. Then (π, λ) has type 1, according to Remark 12.1. Begin by writing

$$\pi = \begin{pmatrix} \cdots & \cdots & \cdots & A & \cdots & B & C_1 & \cdots & C_s \\ \cdots & B & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & A \end{pmatrix},$$

where $A = \alpha(1)$ and B is the leftmost symbol on the top row such that the side ζ_B contains some self-intersection. Recall that the suspension surface is defined from




Figure 23

the simple polygon obtained by removing self-intersections in the way described in figure 17. Combinatorially, this polygon corresponds to the permutation pair

$$\tilde{\pi} = \begin{pmatrix} \cdots & \cdots & \cdots & \cdots & A_1 & B_2 & C_1 & \cdots & C_s & \cdots & B_1 \\ \cdots & B_1 & B_2 & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & A_1 \end{pmatrix},$$

and so the number and orders of the singularities are determined by the equivalence classes of $\tilde{\pi}$, according to the calculation described above. Our claim is that the same is true for the original permutation pair π . This can be seen as follows. Going from π to $\tilde{\pi}$ one replaces A, B by the symbols A_1, B_1, B_2 . Consider the map ϕ defined by

$$\phi(A, L) = (A_1, L), \quad \phi(A, R) = (A_1, R), \quad \phi(B, L) = (B_1, L), \quad \phi(B, R) = (B_2, R),$$

and $\phi(\alpha, S) = (\alpha, S)$ for any other (α, S) . This projects down to a map ψ from the set of equivalence classes of π to the set of equivalence classes of $\tilde{\pi}$ (for the corresponding equivalence relations \sim). Moreover, ψ is injective and leaves invariant the number of regular pairs in each class. The map ψ is *not* surjective: the image avoids, exactly, the equivalence class

$$(B_1, R) \sim (A_1, R) \sim (B_2, L)$$

of $\tilde{\pi}$. However, this equivalence class contains exactly two regular pairs, and so it corresponds to a removable singularity. For consistency, we do remove this singularity from the structure of the suspension surface M . Thus, the number and order of the singularities of M can be obtained from the equivalence classes of π , as we claimed.

Permutation σ . For computations, it is useful to introduce the following alternative terminology. Let us label the pairs (α, S) by integer numbers in the range $\{0, 1, \dots, d\}$ as follows:

$$(\alpha, L) \leftrightarrow \pi_0(\alpha) - 1 \quad \text{and} \quad (\alpha, R) \leftrightarrow \pi_0(\alpha).$$

See figure 23. Notice that this labeling incorporates (38). The remaining identifica-

tions can be expressed in terms of the monodromy invariant p :

$$j \sim k \quad \text{if } p(j) + 1 = p(k + 1), \quad j \notin \{0, p^{-1}(d)\},$$

corresponding to (39), and $0 \sim p^{-1}(1) - 1$, corresponding to (40), and $p^{-1}(d) \sim d$, corresponding to (41). Moreover, these relations may be condensed into

$$j \sim \sigma(j) \quad \text{for every } 0 \leq j \leq d, \tag{42}$$

where $\sigma : \{0, 1, \dots, d\} \rightarrow \{0, 1, \dots, d\}$ is the transformation defined by

$$\sigma(j) = \begin{cases} p^{-1}(1) - 1 & \text{if } j = 0, \\ d & \text{if } p(j) = d, \\ p^{-1}(p(j) + 1) - 1 & \text{otherwise.} \end{cases} \tag{43}$$

It is clear from the construction that σ is a bijection of $\{0, 1, \dots, d\}$, but that can also be checked directly, as follows. Extend p to a bijection P of the set $\{0, 1, \dots, d, d + 1\}$, simply, by defining $P(0) = 0$ and $P(d + 1) = d + 1$. Then (43) becomes

$$\sigma(j) = P^{-1}(P(j) + 1) - 1 \quad \text{for all } 0 \leq j \leq d. \tag{44}$$

This implies that σ is injective, because P is, and it is also clear that σ takes values in $\{0, 1, \dots, d\}$. Thus, it is a bijection, as claimed.

In view of (42), the orbits of σ are in 1-to-1 correspondence to the equivalence classes of \sim . Therefore, the number κ of singularities coincides with the number of distinct orbits of σ . The rule for calculating the angles also translates easily to this terminology. Let us call $1, 2, \dots, d - 1$ *regular*, and 0 and d *irregular* vertices. Then the angle of each singularity a_i is given by

$$\text{angle}(a_i) = 2d_i\pi \tag{45}$$

where $2d_i$ is the number of regular vertices in the corresponding orbit of σ .

Remark 14.1. We have shown that κ and the a_i are determined by σ and, hence, by the monodromy invariant p . In particular, they are independent of λ and τ . This can be understood geometrically by noting that these integer invariants are locally constant on the parameters λ and τ and the domains \mathbb{R}_+^A and T_π^+ are connected, since they are convex cones.

Remark 14.2. Under the canonical involution $(\pi_0, \pi_1) \mapsto (\pi_1, \pi_0)$, the monodromy invariant is replaced by its inverse. Thus, the permutation σ is replaced by

$$\tilde{\sigma}(j) = P(P^{-1}(j) + 1) - 1 \quad \text{for all } 0 \leq j \leq d.$$

This is not quite the same as $\sigma^{-1}(j) = P^{-1}(P(j + 1) - 1)$, but the two transformations are conjugate:

$$\tilde{\sigma} \circ P(j) = P(j + 1) - 1 = P \circ \sigma^{-1}(j).$$

Thus, $\tilde{\sigma}$ and σ have the same number of orbits and, since the conjugacy preserves the set of regular vertices, corresponding orbits have the same number of regular vertices. This shows that the number and orders of the singularities are preserved by the canonical involution.

Proposition 14.3. *The number and the orders of the singularities are constant on each Rauzy class and, consequently, so is the genus.*

Proof. It suffices to prove that the number and the orders of the singularities corresponding to (π, λ) and $(\pi', \lambda') = \hat{R}(\pi, \lambda)$ always coincide. To this end, let p and p' be the monodromy invariants of (π, λ) and (π', λ') , respectively, and σ and σ' be the corresponding permutations of $\{0, 1, \dots, d\}$ given by (43) and (44). Suppose first that (π, λ) has type 0. Then

$$p'(j) = \begin{cases} p(j) & \text{if } p(j) \leq p(d), \\ p(j) + 1 & \text{if } p(d) < p(j) < d, \\ p(d) + 1 & \text{if } p(j) = d \end{cases}$$

or, equivalently,

$$(p')^{-1}(j) = \begin{cases} p^{-1}(j) & \text{if } j \leq p(d), \\ p^{-1}(d) & \text{if } j = p(d) + 1, \\ p^{-1}(j - 1) & \text{if } p(d) + 1 < j \leq d. \end{cases}$$

(We suppose $p(d) \neq d - 1$, for otherwise $p' = p$ and so $\sigma' = \sigma$.) This gives

$$\sigma'(j) = \begin{cases} p^{-1}(d) - 1 & \text{if } j = d, \\ d & \text{if } p(j) = d - 1, \\ \sigma(d) & \text{if } p(j) = d, \\ \sigma(j) & \text{in all other cases.} \end{cases}$$

This means that after Rauzy-Veech induction we have

$$p^{-1}(d - 1) \xrightarrow{\sigma'} d \xrightarrow{\sigma'} p^{-1}(d) - 1 \quad \text{and} \quad p^{-1}(d) \xrightarrow{\sigma'} \sigma(d)$$

whereas, beforehand,

$$p^{-1}(d - 1) \xrightarrow{\sigma} p^{-1}(d) - 1 \quad \text{and} \quad p^{-1}(d) \xrightarrow{\sigma} d \xrightarrow{\sigma} \sigma(d).$$

In other words, replacing σ by σ' means that d is displaced from the orbit of $p^{-1}(d)$ to the orbit of $p^{-1}(d - 1)$ and $p^{-1}(d) - 1$, but the orbit structure is otherwise unchanged. Consequently, the two permutations have the same number of orbits, and corresponding orbits have the same number of *regular* vertices. It follows that the number and




Figure 24

orders of the singularities remain the same. Now suppose (π, λ) has type 1. Let $\tilde{\pi}$ and $\tilde{\pi}'$ be obtained from π and π' by canonical involution. Then $(\tilde{\pi}, \lambda)$ has type zero, and $(\tilde{\pi}', \lambda') = \hat{R}(\tilde{\pi}, \lambda)$. So, by the previous paragraph, the number and orders of the singularities are the same for $(\tilde{\pi}, \lambda)$ and for $(\tilde{\pi}', \lambda')$. By Remark 14.2, the same is true about (π, λ) and $(\tilde{\pi}, \lambda)$, and about (π', λ') and $(\tilde{\pi}', \lambda')$. Thus, the number and orders of the singularities for (π, λ) and (π', λ') are also the same, as claimed. \square

Example 14.4. Figure 24 illustrates the orbit displacement in the proof of the proposition. One has $1 \rightarrow 3$ and $4 \rightarrow 7 \rightarrow 1$ before inducing, and $1 \rightarrow 7 \rightarrow 3$ and $4 \rightarrow 1$ afterwards. In this example all the points concerned belong to the same orbit.

In section 18, we shall extend the Rauzy-Veech induction $\hat{R}(\pi, \lambda) = (\pi', \lambda')$ to an operator $\hat{\mathcal{R}}(\pi, \lambda, \tau) = (\pi', \lambda', \tau')$ in the space of translation surfaces, in such a way that the data (π, λ, τ) and (π', λ', τ') always define the same translation surface. As the number of orders of the singularities depend only on the combinatorial data, by Remark 14.1, that will provide an alternative proof of Proposition 14.3.

15. Zippered rectangles

We are going to describe a useful alternative construction of the suspension of an interval exchange transformation, due to Veech [22]. Given an irreducible pair π and a vector $\tau \in \mathbb{R}^A$, define $h \in \mathbb{R}^A$ by

$$h_\alpha = - \sum_{\pi_1(\beta) < \pi_1(\alpha)} \tau_\beta + \sum_{\pi_0(\beta) < \pi_0(\alpha)} \tau_\beta = -\Omega_\pi(\tau)_\alpha. \tag{46}$$

Observe that if $\tau \in T_\pi^+$, that is, if it satisfies (35) then

$$\sum_{\pi_0(\beta) < \pi_0(\alpha)} \tau_\beta > 0 > \sum_{\pi_1(\beta) < \pi_1(\alpha)} \tau_\beta,$$




Figure 25

and so $h_\alpha > 0$, for all $\alpha \in \mathcal{A}$. We shall consider the convex cones inside the subspace $H_\pi = W_\pi = \Omega_\pi(\mathbb{R}^A)$ defined by

$$W_\pi^+ = \Omega_\pi(\mathbb{R}_+^A) \quad \text{and} \quad H_\pi^+ = -\Omega_\pi(T_\pi^+).$$

Suppose $\tau \in T_\pi^+$. For each $\alpha \in \mathcal{A}$, consider the rectangles of width λ_α and height h_α defined by (see figure 25)

$$R_\alpha^0 = \left(\sum_{\pi_0(\beta) < \pi_0(\alpha)} \lambda_\beta, \sum_{\pi_0(\beta) \leq \pi_0(\alpha)} \lambda_\beta \right) \times [0, h_\alpha]$$

$$R_\alpha^1 = \left(\sum_{\pi_1(\beta) < \pi_1(\alpha)} \lambda_\beta, \sum_{\pi_1(\beta) \leq \pi_1(\alpha)} \lambda_\beta \right) \times [-h_\alpha, 0]$$

and consider also the vertical segments

$$S_\alpha^0 = \left\{ \sum_{\pi_0(\beta) \leq \pi_0(\alpha)} \lambda_\beta \right\} \times \left[0, \sum_{\pi_0(\beta) \leq \pi_0(\alpha)} \tau_\beta \right]$$

$$S_\alpha^1 = \left\{ \sum_{\pi_1(\beta) \leq \pi_1(\alpha)} \lambda_\beta \right\} \times \left[\sum_{\pi_1(\beta) \leq \pi_1(\alpha)} \tau_\beta, 0 \right].$$

That is, S_α^ε joins the horizontal axis to the endpoint of the vector

$$\sum_{\pi_\varepsilon(\beta) \leq \pi_\varepsilon(\alpha)} \zeta_\beta = \sum_{\pi_\varepsilon(\beta) \leq \pi_\varepsilon(\alpha)} (\lambda_\beta, \tau_\beta).$$

Notice that

$$S_{\alpha(0)}^0 = S_{\alpha(1)}^1 = \left\{ \sum_{\beta \in \mathcal{A}} \lambda_\beta \right\} \times \left[0, \sum_{\beta \in \mathcal{A}} \tau_\beta \right].$$

Figure 25 describes two situations where this last segment is above and below the horizontal axis, respectively, depending on the type of τ .

The suspension surface $M = M(\pi, \lambda, \tau, h)$ is the quotient of the union

$$\bigcup_{\alpha \in \mathcal{A}} \bigcup_{\varepsilon=0,1} R_\alpha^\varepsilon \cup S_\alpha^\varepsilon$$

of these objects by certain identifications, that we are going to describe. First, we identify each R_α^0 to R_α^1 through the translation

$$(x, z) \mapsto (x + w_\alpha, z - h_\alpha),$$

that maps one to the other. Note that this is just the same map we used before to identify the two sides of the polygon corresponding to the vector $\zeta_\alpha = (\lambda_\alpha, \tau_\alpha)$: recall (37).

We may think of the segments S_α^ε as “zipping” adjacent rectangles together *up to a certain height*. Observe that, in most cases, S_α^ε is shorter than the heights of both adjacent rectangles (compare figure 25):

Lemma 15.1. *For any $\varepsilon \in \{0, 1\}$ and $\alpha \in \mathcal{A}$,*

- (i) $(-1)^\varepsilon \sum_{\pi_\varepsilon(\beta) \leq \pi_\varepsilon(\alpha)} \tau_\beta < h_\alpha$ *except, possibly, if $\pi_{1-\varepsilon}(\alpha) = d$.*
- (ii) $(-1)^\varepsilon \sum_{\pi_\varepsilon(\beta) \leq \pi_\varepsilon(\alpha)} \tau_\beta < h_\gamma$, *where $\gamma \in \mathcal{A}$ is defined by $\pi_\varepsilon(\gamma) = \pi_\varepsilon(\alpha) + 1$ and we suppose $\pi_\varepsilon(\alpha) < d$.*

Proof. For $\varepsilon = 0$ the relations (35) and (46) give

$$h_\alpha - \sum_{\pi_0(\beta) \leq \pi_0(\alpha)} \tau_\beta = - \sum_{\pi_1(\beta) \leq \pi_1(\alpha)} \tau_\beta > 0 \tag{47}$$

except, possibly, if $\pi_1(\alpha) = d$, that is, $\alpha = \alpha(1)$. This takes care of the rectangle to the left of S_α^0 . The one to the right (when it exists) is handled similarly: Let $\gamma \in \mathcal{A}$ be such that $\pi_0(\gamma) = \pi_0(\alpha) + 1$. Then

$$h_\gamma - \sum_{\pi_0(\beta) \leq \pi_0(\alpha)} \tau_\beta = - \sum_{\pi_1(\beta) < \pi_1(\gamma)} \tau_\beta > 0.$$

The case $\varepsilon = 1$ is analogous. □

On the other hand, the calculation in (47) also shows that for $\alpha = \alpha(1)$ the length of S_α^0 may exceed the height of R_α^0 : this happens if the sum of all τ_β is positive. In that case, let

$$\tilde{S} = \left\{ \sum_{\pi_0(\beta) \leq \pi_0(\alpha)} \lambda_\beta \right\} \times \left[h_\alpha, \sum_{\pi_0(\beta) \leq \pi_0(\alpha)} \tau_\beta \right],$$

that is, \tilde{S} is the subsegment of length $\sum_{\beta \in \mathcal{A}} \tau_\beta$ at the top of S_α^0 . Dually, if the sum of all τ_β is negative then, for $\alpha = \alpha(0)$, the length of S_α^1 exceeds the height of R_α^1 . In this case, define

$$\tilde{S} = \left\{ \sum_{\pi_0(\beta) \leq \pi_0(\alpha)} \lambda_\beta \right\} \times \left[\sum_{\pi_1(\beta) \leq \pi_1(\alpha)} \tau_\beta, -h_\alpha \right],$$

instead. That is, \tilde{S} is the subsegment of length $-\sum_{\beta \in \mathcal{A}} \tau_\beta$ at the bottom of S_α^1 . In either case, we identify \tilde{S} with the vertical segment $S_{\alpha(0)}^0 = S_{\alpha(1)}^1$, by translation. This completes the definition of the suspension surface.

This construction is equivalent to the one in section 12, in the sense that they give rise to suspension surfaces that are isometric, by an isometry that preserves the vertical direction. This is clear from the previous observations, at least when the closed curve $\Gamma(\pi, \lambda, \tau)$ is simple; we leave it to the reader to check that it remains true when there are self-intersections.

There is a natural notion of *area* of a *zippered rectangle* (π, λ, τ, h) , namely

$$\text{area}(\pi, \lambda, \tau, h) = \lambda \cdot h = \sum_{\alpha \in \mathcal{A}} \lambda_\alpha h_\alpha. \tag{48}$$

Sometimes we write $\text{area}(\pi, \lambda, \tau)$ to mean $\text{area}(\pi, \lambda, \tau, h)$ with $h = -\Omega_\pi(\tau)$.

16. Genus and dimension

We have seen in Remark 10.3 that the vector space $H_\pi = \Omega_\pi(\mathbb{R}^A)$ has even dimension. We can now interpret this dimension in terms of the genus of the suspension surface:

Proposition 16.1. *The dimension of H_π coincides with $2g(M)$, where $g(M)$ is the genus of the suspension surface M .*

Proof. Rename the intervals I_α so that the permutation pair π becomes normalized to $\mathcal{A} = \{1, \dots, d\}$, $\pi_0 = \text{id}$ and, thus, $\pi_1 = p = \text{monodromy invariant}$. Write the translation vector as $w = \Omega_\pi(\lambda)$, that is

$$w_j = \sum_{p(i) < p(j)} \lambda_i - \sum_{i < j} \lambda_i \quad \text{for each } 1 \leq j \leq d.$$

It is convenient to extend the definition to $j = 0$ and $j = d + 1$, simply, by replacing p by its extension P in (44). Since $P(0) = 0$ and $P(d + 1) = d + 1$, by definition, this just means we take $w_0 = w_{d+1} = 0$. Define $a_j = \sum_{i \leq j} \lambda_i$ for $1 \leq j \leq d$, and $a_0 = 0$.

Lemma 16.2. *We have $w_{\sigma(j)+1} - w_j = a_j - a_{\sigma(j)}$ for every $0 \leq j \leq d$.*

Proof. As we have seen in (44), $\sigma(j) = P^{-1}(P(j) + 1) - 1$, and so

$$w_{\sigma(j)+1} = \sum_{P(i) < P(\sigma(j)+1)} \lambda_i - \sum_{i < \sigma(j)+1} \lambda_i = \sum_{P(i) \leq P(j)} \lambda_i - \sum_{i \leq \sigma(j)} \lambda_i.$$

It follows that

$$w_{\sigma(j)+1} - w_j = \lambda_j - \sum_{i \leq \sigma(j)} \lambda_i + \sum_{i < j} \lambda_i = \sum_{i \leq j} \lambda_i - \sum_{i \leq \sigma(j)} \lambda_i = a_j - a_{\sigma(j)}$$

as claimed. □

Recall that the number of orbits of σ is equal to the number κ of singularities.

Lemma 16.3. *A vector λ is in $\ker \Omega_\pi$ if and only if the $(d + 1)$ -dimensional vector $(0, a_1, \dots, a_d)$ is constant on the orbits of σ . Hence, $\dim \ker \Omega_\pi = \kappa - 1$.*

Proof. The only if part is a direct consequence of Lemma 16.2: if $w = 0$ then $a_{\sigma(j)} - a_j = 0$ for every $0 \leq j \leq d$. To prove the converse, let λ be such that $(0, a_1, \dots, a_d)$ is constant on orbits of σ . Then, by Lemma 16.2,

$$w_{P^{-1}(P(j)+1)} = w_{\sigma(j)+1} = w_j \quad \text{for all } 0 \leq j \leq d.$$

Writing $P(j) = i$, this relation becomes

$$w_{P^{-1}(i+1)} = w_{P^{-1}(i)} \quad \text{for all } 0 \leq i \leq d.$$

It follows that $w_{P^{-1}(i)}$ is constant on $\{0, 1, \dots, d + 1\}$ and, since it vanishes for $i = 0$, it follows that it must vanish for every $1 \leq i \leq d$. Consequently, $w = (w_1, \dots, w_d)$ vanishes, and this means that $\lambda \in \ker \Omega_\pi$. This proves the first part of the lemma.

To prove the second one, consider the linear isomorphism

$$\psi : \mathbb{R}^d \rightarrow \mathbb{R}^d, \quad (\lambda_1, \dots, \lambda_d) \mapsto (a_1, \dots, a_d), \quad a_j = \sum_{i=1}^j \lambda_i.$$

Let K_π be the subspace of all $(a_1, \dots, a_d) \in \mathbb{R}^d$ such that $(0, a_1, \dots, a_d)$ is constant on the orbits of σ . The dimension of K_π is $\kappa - 1$, because the value of a_j on the orbit of 0 is predetermined by $a_0 = 0$. The previous paragraph shows that

$$\ker \Omega_\pi = \psi^{-1}(K_\pi).$$

Consequently, the dimension of the kernel is $\kappa - 1$, as claimed. □

Using Lemma 16.3 and the relation (33), we find

$$\dim \Omega_\pi(\mathbb{R}^A) = d - \dim \ker \Omega_\pi = d - \kappa + 1 = 2g(M).$$

This proves Proposition 16.1. □

It is possible to give an explicit description of $\ker \Omega_\pi$ and H_π , as follows. For each orbit \mathcal{O} of σ not containing zero, and for each $1 \leq j \leq d$, define

$$\lambda(\mathcal{O})_j = \mathcal{X}_\mathcal{O}(j) - \mathcal{X}_\mathcal{O}(j-1) = \begin{cases} 1 & \text{if } j \in \mathcal{O} \text{ but } j-1 \notin \mathcal{O}, \\ -1 & \text{if } j \notin \mathcal{O} \text{ but } j-1 \in \mathcal{O}, \\ 0 & \text{in all other cases.} \end{cases}$$

Lemma 16.4. *Define $a(\mathcal{O}) = \psi(\lambda(\mathcal{O}))$, that is, $a(\mathcal{O})_j = \sum_{i \leq j} \lambda(\mathcal{O})_i$. Then*

$$a(\mathcal{O})_j = \mathcal{X}_\mathcal{O}(j) = \begin{cases} 1 & \text{if } j \in \mathcal{O}, \\ 0 & \text{if } j \notin \mathcal{O}. \end{cases}$$

Proof. For $j = 1$ this follows from a simple calculation: $a(\mathcal{O})_1 = 1$ if $1 \in \mathcal{O}$ (and $0 \notin \mathcal{O}$) and $a(\mathcal{O})_1 = 0$ if $1 \notin \mathcal{O}$ (and $0 \notin \mathcal{O}$). The proof proceeds by induction: if $a(\mathcal{O})_{j-1} = \mathcal{X}_\mathcal{O}(j-1)$ then

$$a(\mathcal{O})_j = a(\mathcal{O})_{j-1} + \lambda(\mathcal{O})_j = \mathcal{X}_\mathcal{O}(j-1) + \mathcal{X}_\mathcal{O}(j) - \mathcal{X}_\mathcal{O}(j-1) = \mathcal{X}_\mathcal{O}(j).$$

The argument is complete. □

Clearly, the $a(\mathcal{O})$ form a basis of the subspace K_π of vectors (a_1, \dots, a_d) such that $(0, a_1, \dots, a_d)$ is constant on orbits of σ . It follows that

$$\{ \lambda(\mathcal{O}) : \mathcal{O} \text{ is an orbit of } \sigma \text{ not containing } 0 \}$$

is a basis of $\ker \Omega_\pi$. Moreover, since Ω_π is anti-symmetric, the range H_π is just the orthogonal complement of the kernel. In other words, $w \in H_\pi$ if and only if $w \cdot \lambda(\mathcal{O}) = 0$ for every orbit \mathcal{O} of σ not containing zero.

17. Hyperelliptic Rauzy classes

Let $d \geq 2$ be fixed. We call *hyperelliptic* the Rauzy class which contains the pair

$$\pi = \begin{pmatrix} A_1 & A_2 & \cdots & \cdots & A_d \\ A_d & \cdots & \cdots & A_2 & A_1 \end{pmatrix},$$

that is, which corresponds to the monodromy invariant $p = (d, d-1, \dots, 2, 1)$ defined by $p(i) = d+1-i$ for all i .

Lemma 17.1.

- (i) If d is even then the number of singularities $\kappa(\pi) = 1$, the singularity has order $d - 2$ and the surface M has genus $g(M) = d/2$. Moreover, the operator $\Omega_\pi : \mathbb{R}^A \rightarrow \mathbb{R}^A$ is an isomorphism.
- (ii) If d is odd then there are $\kappa(\pi) = 2$ singularities, and they both have order $(d - 3)/2$. The surface M has genus $g(M) = (d - 1)/2$. Moreover, the kernel of Ω_π has dimension 1.

Proof. Observe that $p^{-1}(i) = p(i) = d + 1 - i$ for all $1 \leq i \leq d$. From (43) we find that the permutation σ is given by

$$\sigma(j) = \begin{cases} d - 1 & \text{for } j = 0, \\ d & \text{for } j = 1, \\ j - 2 & \text{in all other cases.} \end{cases}$$

That is, σ is the right rotation by two units

$$\sigma = (d - 1, d, 0, 1, \dots, d - 2).$$

If d is even, then this rotation has a unique orbit in $\{0, 1, \dots, d\}$. It follows that $\kappa = 1$ and, by (32) the singularity has angle $(2d - 2)\pi$, that is, order $d - 2$. Moreover, (33) gives $g(M) = d/2$. If d is odd then the rotation has exactly two orbits:

$$0 \rightarrow d - 1 \rightarrow d - 3 \rightarrow \dots \rightarrow 2 \rightarrow 0 \quad \text{and} \quad d \rightarrow d - 2 \rightarrow d - 4 \rightarrow \dots \rightarrow 1 \rightarrow d.$$

Each one involves $(d - 1)/2$ regular elements (that is, different from 0 and d). Using (45) we get that they both have angle $(a_i) = (d - 1)\pi$, and so their order is $(d - 3)/2$. Moreover, (33) gives $g(M) = (d - 1)/2$.

The statement about Ω_π is now an immediate consequence of Proposition 16.1, but it may also be proved directly. To this end, let us normalize the permutation pair π (rename the intervals) so that $\mathcal{A} = \{1, \dots, d\}$, $\pi_0 = \text{id}$ and, thus, π_1 is the monodromy invariant p . Then $\Omega_\pi(\lambda) = w$ is given by

$$w_j = \sum_{\pi_1(i) < \pi_1(j)} \lambda_i - \sum_{\pi_0(i) < \pi_0(j)} \lambda_i = \sum_{i > j} \lambda_i - \sum_{i < j} \lambda_i.$$

This gives $w_j - w_{j+1} = \lambda_j + \lambda_{j+1}$ for $j = 1, \dots, d - 1$, and also $w_d + w_1 = \lambda_d - \lambda_1$. Suppose λ is in the kernel, that is, $w = 0$. Then the λ_j must be alternately symmetric, and the first and the last one must coincide: $\lambda_1 = \lambda_d$. If d is even this can only happen for $\lambda = 0$: thus, Ω_π is an isomorphism. If d is odd, it means that $\lambda = (x, -x, x, -x, \dots, -x)$ for some real number x . It is easy to check that vectors of this form are, indeed, in the kernel. This proves that the kernel of Ω_π has dimension 1 in this case. □

The relation (33) shows that d and κ always have opposite parities. So, the situation described in Lemma 17.1 corresponds to the smallest possible number of singularities.




Figure 26

18. Invertible Rauzy-Veech induction

We are going to define a counterpart $\hat{\mathcal{R}} : (\pi, \lambda, \tau) \mapsto (\pi', \lambda', \tau')$ of the Rauzy-Veech induction $\hat{R} : (\pi, \lambda) \mapsto (\pi', \lambda')$ at the level of suspension data (π, λ, τ) . Recall that \hat{R} corresponds to replacing the original interval exchange map by its first return to a conveniently chosen subinterval of the domain. Similarly, this map $\hat{\mathcal{R}}$ we are introducing corresponds to replacing the horizontal cross-section in (36) by a shorter one. The Poincaré return map of the vertical flow to this new cross-section is precisely the interval exchange map described by (π', λ') , and we want to rewrite the ambient surface as a suspension over this map: the coordinate τ' is chosen with this purpose in mind. Thus, the data (π, λ, τ) and (π', λ', τ') are really different presentations of the same translation surface. We shall check that the transformation $\hat{\mathcal{R}}$ is invertible almost everywhere and has a Markov property. Later, we shall see that it is a realization of the inverse limit (natural extension) of \hat{R} .

Let $\hat{\mathcal{H}} = \hat{\mathcal{H}}(C) = \{ (\pi, \lambda, \tau) : \pi \in C, \lambda \in \mathbb{R}_+^A, \tau \in T_\pi^+ \}$. The transformation $\hat{\mathcal{R}}$ is defined on $\hat{\mathcal{H}}$ by $\hat{\mathcal{R}}(\pi, \lambda, \tau) = (\pi', \lambda', \tau')$, where $(\pi', \lambda') = \hat{R}(\pi, \lambda)$ and

$$\tau'_\alpha = \begin{cases} \tau_\alpha, & \alpha \neq \alpha(\varepsilon), \\ \tau_{\alpha(\varepsilon)} - \tau_{\alpha(1-\varepsilon)}, & \alpha = \alpha(\varepsilon), \end{cases} \quad \varepsilon = \text{type of } (\pi, \lambda).$$

In other words (compare (11) for the definition of λ'),

$$\tau' = \Theta^{-1*}(\tau). \tag{49}$$

Figures 26 and 27 provide a geometric interpretation of this Rauzy-Veech induction, in terms of the polygon defining the suspension surface: one cuts from the polygon the triangle determined by the sides $\zeta_{\alpha(0)}$ and $-\zeta_{\alpha(1)}$ and pastes it back, adjacently to the other side labeled by $\alpha(\varepsilon)$, where ε is the type of (π, λ) . Observe that the surface itself remains unchanged or, rather, the translation surfaces determined by (π, λ, τ) and (π', λ', τ') are *equivalent*, in the sense that there exists an isometry between the two that preserves the vertical direction. We leave it to the




Figure 27




Figure 28

reader to check how this geometric interpretation extends to the case when the closed curve $\gamma(\pi, \lambda, \tau)$ has self-intersections. An equivalent formulation of the Rauzy-Veech induction in terms of zippered rectangles will be given in section 19.

Recall that we defined the type of τ to be 0 if the sum of τ_α over all $\alpha \in \mathcal{A}$ is positive and 1 if the sum is negative. Figures 26 and 27 immediately suggest that

$$(\pi, \lambda) \text{ has type } \varepsilon \quad \Rightarrow \quad \tau' \text{ has type } 1 - \varepsilon. \tag{50}$$

This observation is also contained in the next, more precise, lemma. See also figure 28, that describes the action of $\hat{\mathcal{R}}$ on both variables λ and τ .

Lemma 18.1. *The linear transformation Θ^{-1*} sends T_π^+ injectively inside $T_{\pi'}^+$ and, denoting $\varepsilon = \text{type of } (\pi, \lambda)$, the image coincides with the set of $\tau' \in T_{\pi'}^+$ whose type is $1 - \varepsilon$.*

Proof. Suppose $\varepsilon = 0$, as the other case is analogous. We begin by checking that the image of Θ^{-1*} is contained in $T_{\pi'}^+$, that is, τ' satisfies (35) if τ does. Firstly, $\pi'_0 = \pi_0$ and $\tau'_\alpha = \tau_\alpha$ for every $\alpha \neq \alpha(0)$ imply

$$\sum_{\pi'_0(\alpha) \leq k} \tau'_\alpha = \sum_{\pi_0(\alpha) \leq k} \tau_\alpha > 0 \tag{51}$$

for every $k < d$. Now let $l = \pi_1(\alpha(0))$ be the position of $\alpha(0)$ in the bottom line of π .

Recall that π'_1 and π_1 coincide to the left of l . So, just as before,

$$\sum_{\pi'_1(\alpha) \leq k} \tau'_\alpha = \sum_{\pi_1(\alpha) \leq k} \tau_\alpha < 0 \tag{52}$$

for every $k < l$. The case $k = l$ is more interesting: using $\tau'_{\alpha(0)} = \tau_{\alpha(0)} - \tau_{\alpha(1)}$

$$\sum_{\pi'_1(\alpha) \leq l} \tau'_\alpha = \sum_{\pi_1(\alpha) \leq l} \tau_\alpha - \tau_{\alpha(1)}.$$

To prove that this is less than zero, rewrite the right hand side as (recall the definition (46) of h)

$$-h_{\alpha(0)} + \sum_{\pi_0(\alpha) \leq l} \tau_\alpha - \tau_{\alpha(1)} = -h_{\alpha(0)} + \sum_{\alpha \in \mathcal{A}} \tau_\alpha - \tau_{\alpha(1)} = -h_{\alpha(0)} + \sum_{\pi_1(\alpha) < \pi_1(\alpha(1))} \tau_\alpha.$$

Both terms in the last expression are negative, because the entries of h are positive and τ satisfies (35). This deals with the case $k = l$. Next, for $k = l + 1$, we use the fact that $\pi'_1(\alpha(1)) = l + 1$ to obtain

$$\sum_{\pi'_1(\alpha) \leq l+1} \tau'_\alpha = \sum_{\pi_1(\alpha) \leq l} \tau_\alpha < 0. \tag{53}$$

More generally, for $l < k \leq d$ we have

$$\sum_{\pi'_1(\alpha) \leq k} \tau'_\alpha = \sum_{\pi_1(\alpha) \leq k-1} \tau_\alpha < 0. \tag{54}$$

This proves that the image of T_π^+ is indeed contained in $T_{\pi'}^+$. Moreover, the case $k = d$ gives that every τ' in the image has type 1,

$$\sum_{\alpha \in \mathcal{A}} \tau'_\alpha < 0, \tag{55}$$

as claimed. To complete the proof we only have to check that if $\tau' \in T_{\pi'}^+$ satisfies (55) then $\tau = \Theta^*(\tau')$ is in T_π^+ . This is easily seen from the relations (51)–(54). The hypothesis (55) is needed only when $k = d - 1$. \square

Recall that the Rauzy-Veech induction $\hat{R} : (\pi, \lambda) \mapsto (\pi', \lambda')$ for interval exchange transformations is 2-to-1 on its domain, the two pre-images corresponding to the two possible values of the type ε . For each $\varepsilon \in \{0, 1\}$, let us denote

$$\mathbb{R}_{\pi, \varepsilon}^{\mathcal{A}} = \{ \lambda \in \mathbb{R}_+^{\mathcal{A}} : (\pi, \lambda) \text{ has type } \varepsilon \} \quad \text{and} \quad T_\pi^\varepsilon = \{ \tau \in T_\pi^+ : \tau \text{ has type } \varepsilon \}.$$

From the previous lemma we obtain




Figure 29

Corollary 18.2. *The transformation $\hat{\mathcal{R}} : \hat{\mathcal{H}} \rightarrow \hat{\mathcal{H}}$ is an (almost everywhere) invertible Markov map, and it preserves the natural area:*

- (i) $\hat{\mathcal{R}}(\{\pi\} \times \mathbb{R}_{\pi, \varepsilon}^A \times T_{\pi}^+) = \{\pi'\} \times \mathbb{R}_{\pi'}^A \times T_{\pi'}^{1-\varepsilon}$ for every π and ε ;
- (ii) every (π', λ', τ') such that $\sum_{\alpha \in \mathcal{A}} \tau'_{\alpha} \neq 0$ has exactly one preimage for $\hat{\mathcal{R}}$;
- (iii) if $\hat{\mathcal{R}}(\pi, \lambda, \tau) = (\pi', \lambda', \tau')$ then $\text{area}(\pi, \lambda, \tau) = \text{area}(\pi', \lambda', \tau')$.

Proof. The first claim is contained in Lemma 18.1. The second one follows from the injectivity in that lemma, together with the observation that the sets $\{\pi'\} \times \mathbb{R}_{\pi'}^A \times T_{\pi'}^{1-\varepsilon}$ are pairwise disjoint. Finally, Lemma 10.2 and the relations (11) and (56) give

$$-\lambda' \cdot \Omega_{\pi'}(\tau') = \lambda' \cdot h' = \Theta^{-1*}(\lambda) \cdot \Theta(h) = \lambda \cdot h = -\lambda \cdot \Omega_{\pi}(\tau).$$

and this proves the third claim. □

Remark 18.3. Let ε be the type of (π, λ) . If τ also has type ε then the curve $\Gamma(\pi, \lambda, \tau)$ is simple, according to Remark 18.3. Otherwise, let $n \geq 1$ be minimum such that the type of (π^n, λ^n) is $1 - \varepsilon$. By (50), the type of τ^n is also $1 - \varepsilon$. It follows that the curve $\Gamma(\pi^n, \lambda^n, \tau^n)$ has no self-intersections. Recall that (π, λ, τ) and $(\pi^n, \lambda^n, \tau^n)$ represent the same translation surface, up to an isometry that preserves the vertical direction.

19. Induction for zippered rectangles

The definition of the induction operator $\hat{\mathcal{R}}$ is, perhaps, more intuitive in the language of zippered rectangles. Indeed, as explained previously, the idea behind the definition is to rewrite the translation surface as a suspension of the Poincaré return map of the vertical flow to a shorter cross-section. In terms of zippered rectangles this is achieved by an especially simple geometric procedure, described in figure 29: one removes a rightmost subrectangle from the rectangle corresponding to the symbol $\alpha(\varepsilon)$ and

pastes it back on top of the rectangle corresponding to the symbol $\alpha(1 - \varepsilon)$. The precise definition goes as follows.

Let $\hat{\mathcal{H}} = \hat{\mathcal{H}}(C)$ be the set of (π, λ, τ, h) such that $\pi \in C$, $\lambda \in \mathbb{R}_+^A$, $\tau \in T_\pi^+$, and $h = -\Omega_\pi(\tau) \in H_\pi^+$. Then define $\hat{\mathcal{R}}(\pi, \lambda, \tau, h) = (\pi', \lambda', \tau', h')$, where

$$h'_\alpha = \begin{cases} h_\alpha, & \alpha \neq \alpha(1 - \varepsilon), \\ h_{\alpha(1-\varepsilon)} + h_{\alpha(\varepsilon)}, & \alpha = \alpha(1 - \varepsilon). \end{cases}$$

Compare figure 29. Equivalently (recall (8)),

$$h' = \Theta(h). \tag{56}$$

Let us relate this to the definition $\hat{\mathcal{R}}(\pi, \lambda, \tau) = (\pi', \lambda', \tau')$ with $\tau' = \Theta^{-1*}(\tau)$ that was given in the previous section.

By Lemma 10.2, we have $\Theta \Omega_\pi \Theta^* = \Omega_{\pi'}$. Since Θ is an isomorphism, this gives that $\tau \in \ker \Omega_\pi$ if and only if $\tau' \in \ker \Omega_{\pi'}$. In other words,

$$\{ \ker \Omega_\pi : \pi \in C \}$$

defines an invariant subbundle for $\tau \mapsto \tau' = \Theta^{-1*}(\tau)$. As we have seen before, H_π is the orthogonal complement of the kernel, because Ω_π is anti-symmetric. Hence

$$\{ H_\pi : \pi \in C \}$$

is an invariant subbundle for the adjoint cocycle Θ . The map defined by (56) is just the restriction of the adjoint to this invariant subbundle.

The relation $\Theta \Omega_\pi \Theta^* = \Omega_{\pi'}$ also says that there is a conjugacy

$$\begin{array}{ccc} \mathbb{R}^A / \ker \Omega_\pi & \xrightarrow{\Theta^{-1*}} & \mathbb{R}^A / \ker \Omega_{\pi'} \\ \Omega_\pi \downarrow & & \downarrow \Omega_{\pi'} \\ H_\pi & \xrightarrow{\Theta} & H_{\pi'} \end{array}$$

20. Teichmüller flow

Let C be any Rauzy class. We defined $\hat{\mathcal{H}} = \hat{\mathcal{H}}(C)$ to be the set of all (π, λ, τ) such that $\pi \in C$, $\lambda \in \mathbb{R}_+^A$, and $\tau \in T_\pi^+$. The *Teichmüller flow* on $\hat{\mathcal{H}}$ is the natural action $\mathcal{T} = (\mathcal{T}^t)_{t \in \mathbb{R}}$ of the diagonal subgroup

$$\begin{pmatrix} e^t & 0 \\ 0 & e^{-t} \end{pmatrix}, \quad t \in \mathbb{R}$$

defined by

$$\mathcal{T}^t : \hat{\mathcal{H}} \rightarrow \hat{\mathcal{H}}, \quad (\pi, \lambda, \tau) \mapsto (\pi, e^t \lambda, e^{-t} \tau)$$

This is well defined because both \mathbb{R}_+^A and T_π^+ are invariant under product by positive scalars. It is clear that the Teichmüller flow commutes with the Rauzy-Veech induction map $\hat{\mathcal{R}}$ and preserves the natural area (48). For each $\lambda \in \mathbb{R}_+^A$, define the total length $|\lambda| = \sum_{\alpha \in A} \lambda_\alpha$. Given any $c > 0$, the affine subset

$$\mathcal{H}_c = \{(\pi, \lambda, \tau) \in \hat{\mathcal{H}} : |\lambda| = c\} \quad (\text{We denote } \mathcal{H} = \mathcal{H}_1.)$$

is a global cross-section for the Teichmüller flow \mathcal{T} : each trajectory intersects \mathcal{H}_c exactly once. In particular, the map

$$\Psi : \mathcal{H} \times \mathbb{R} \rightarrow \hat{\mathcal{H}}, \quad \Psi(\pi, \lambda, \tau, s) = \mathcal{T}^s(\pi, \lambda, \tau) = (\pi, e^s \lambda, e^s \tau)$$

is a diffeomorphism onto $\hat{\mathcal{H}}$. In these new coordinates, the Teichmüller flow is described, simply, by

$$\mathcal{T}^t : \mathcal{H} \times \mathbb{R} \rightarrow \mathcal{H} \times \mathbb{R}, \quad (\pi, \lambda, \tau, s) \mapsto (\pi, \lambda, \tau, s + t). \quad (57)$$

For each $(\pi, \lambda, \tau) \in \hat{\mathcal{H}}$, define the *Rauzy renormalization time* (The renormalization time depends only on π and $\lambda/|\lambda|$.)

$$t_R = t_R(\pi, \lambda) = -\log\left(1 - \frac{\lambda_{\alpha(1-\varepsilon)}}{|\lambda|}\right), \quad \varepsilon = \text{type of } (\pi, \lambda). \quad (58)$$

Notice that if $(\pi', \lambda') = \hat{\mathcal{R}}(\pi, \lambda)$ then $|\lambda'| = e^{-t_R} |\lambda|$. This means that the transformation

$$\mathcal{R} = \hat{\mathcal{R}} \circ \mathcal{T}^{t_R} : (\pi, \lambda, \tau) \mapsto \hat{\mathcal{R}}(\pi, e^{t_R} \lambda, e^{-t_R} \tau) \quad (59)$$

maps each cross-section \mathcal{H}_c back to itself. We call the restriction $\mathcal{R} : \mathcal{H} \rightarrow \mathcal{H}$ to $\mathcal{H} = \mathcal{H}_1$ the *invertible Rauzy-Veech renormalization map*. Observe that for any $(\pi, \lambda, \tau) \in \mathcal{H}$ we have $\mathcal{R}(\pi, \lambda, \tau) = (\pi', \lambda'', \tau'')$ where

$$(\pi', \lambda', \tau') = \hat{\mathcal{R}}(\pi, \lambda, \tau), \quad \lambda'' = \lambda' / (1 - \lambda_{\alpha(1-\varepsilon)}), \quad \tau'' = \tau' (1 - \lambda_{\alpha(1-\varepsilon)}).$$

In particular, \mathcal{R} is a lift of the map $R(\pi, \lambda) = (\pi', \lambda'')$ introduced in section 7. From Corollary 18.2 one obtains

Corollary 20.1. *The transformation $\mathcal{R} : \mathcal{H} \rightarrow \mathcal{H}$ is an (almost everywhere) invertible Markov map, and it preserves the natural area.*

We call *pre-stratum* $\hat{\mathcal{S}} = \hat{\mathcal{S}}(C)$ associated to C the quotient of the fundamental domain $\{(\pi, \lambda, \tau) \in \hat{\mathcal{H}} : 0 \leq \log|\lambda| \leq t_R(\pi, \lambda)\}$ by the equivalence relation

$$\mathcal{T}^{t_R(\pi, \lambda)}(\pi, \lambda, \tau) \sim \mathcal{R}(\pi, \lambda, \tau) \quad \text{for all } (\pi, \lambda, \tau) \in \mathcal{H}.$$

See figure 30. Equivalently, the pre-stratum may be seen as the quotient of the whole $\hat{\mathcal{H}}$ by the equivalence relation generated by $\mathcal{T}^{t_R}(\pi, \lambda, \tau) \sim \mathcal{R}(\pi, \lambda, \tau)$. We




Figure 30

denote by \mathcal{S} the (injective) image of \mathcal{H} under the quotient map. Observe that the dimension of the pre-stratum is given by

$$\dim \hat{\mathcal{S}}(C) = 2d = 4g + 2\kappa - 2.$$

Since \mathcal{R} commutes with \mathcal{T}^t , the latter induces a flow $\mathcal{T} = (\mathcal{T}^t)_{t \in \mathbb{R}}$ on the pre-stratum, that we also call *Teichmüller flow*. The invertible Rauzy-Veech renormalization is naturally identified with the Poincaré return map of this flow to the cross-section $\mathcal{S} \subset \hat{\mathcal{S}}$. Notice that the Teichmüller flow preserves the natural volume measure on $\hat{\mathcal{S}}$, inherited from $\hat{\mathcal{H}}$. We shall see that this volume is finite, if one restricts to $\{\text{area}(\pi, \lambda, \tau) \leq 1\}$.

Invertible Zorich maps. We also use accelerated versions of $\hat{\mathcal{R}}$ and \mathcal{R} , that we call *invertible Zorich induction* and *invertible Zorich renormalization*, respectively, defined by

$$\hat{\mathcal{Z}}(\pi, \lambda, \tau) = \hat{\mathcal{R}}^n(\pi, \lambda, \tau) \quad \text{and} \quad \mathcal{Z}(\pi, \lambda, \tau) = \mathcal{R}^n(\pi, \lambda, \tau),$$

where $n = n(\pi, \lambda) \geq 1$ is the first time the type of $(\pi^n, \lambda^n) = \hat{\mathcal{R}}^n(\pi, \lambda)$ differs from the type of (π, λ) . See section 8. The domain of $\hat{\mathcal{Z}}$ is a subset $\hat{\mathcal{Z}}_*$ of $\hat{\mathcal{H}}$ that we describe in the sequel. Begin by recalling (50):

- if (π, λ) has type 0, that is, $\lambda_{\alpha(0)} > \lambda_{\alpha(1)}$ then τ' has type 1, that is, $\sum_{\alpha \in \mathcal{A}} \tau'_\alpha < 0$;
- if (π, λ) has type 1, that is, $\lambda_{\alpha(0)} < \lambda_{\alpha(1)}$ then τ' has type 0, that is, $\sum_{\alpha \in \mathcal{A}} \tau'_\alpha > 0$.

Define $\hat{\mathcal{Z}}_* = \hat{\mathcal{Z}}_0 \cup \hat{\mathcal{Z}}_1$ where, for each $\varepsilon \in \{0, 1\}$,

$$\hat{\mathcal{Z}}_\varepsilon = \{(\pi, \lambda, \tau) \in \hat{\mathcal{H}} : (\pi, \lambda) \text{ has type } \varepsilon \text{ and } \tau \text{ has type } \varepsilon\}.$$

Then $n = n(\pi, \lambda)$ is just the first positive iterate for which $\hat{\mathcal{R}}^n(\pi, \lambda, \tau)$ hits $\hat{\mathcal{Z}}_*$. Thus, we consider $\hat{\mathcal{Z}}$ defined on the domain $\hat{\mathcal{Z}}_*$. The previous observations mean that




Figure 31

$\hat{\mathcal{Z}} : \hat{Z}_* \rightarrow \hat{Z}_*$ is the first return map of $\hat{\mathcal{R}}$ to the domain \hat{Z}_* . It follows that $\hat{\mathcal{Z}}$ is invertible: the inverse is the first return map to \hat{Z}_* of the map $\hat{\mathcal{R}}^{-1}$.

Analogously, we consider $\mathcal{Z} : Z_* \rightarrow Z_*$ where Z_* is the set of $(\pi, \lambda, \tau) \in \hat{Z}_*$ such that $|\lambda| = 1$. Then \mathcal{Z} is the first return map of \mathcal{R} to Z_* .

21. Volume measure

For translation surfaces. Let C be a Rauzy class. The space $\hat{\mathcal{H}} = \hat{\mathcal{H}}(C)$ has a natural volume measure $\hat{m} = d\pi d\lambda d\tau$, where $d\pi$ is the counting measure on C , and $d\lambda$ and $d\tau$ are the restrictions to \mathbb{R}_+^A and T_π^+ , respectively, of the Lebesgue measure on \mathbb{R}^A . Clearly, \hat{m} is invariant under the Teichmüller flow

$$\mathcal{T}^t : (\pi, \lambda, \tau) \mapsto (\pi, e^t \lambda, e^{-t} \tau).$$

Let us consider the coordinate change $\mathcal{H} \times \mathbb{R} \rightarrow \hat{\mathcal{H}}$, $(\pi, \lambda, \tau, s) \mapsto (\pi, e^s \lambda, e^s \tau)$ introduced in (57). Observe that

$$d\lambda = e^{s(d-1)} d_1 \lambda e^s ds = e^{sd} d_1 \lambda ds,$$

where $d_1 \lambda$ denotes the Lebesgue measure induced on $\Lambda_{\mathcal{A}} = \{\lambda \in \mathbb{R}_+^A : |\lambda| = 1\}$ by the Riemannian metric of \mathbb{R}^A . See figure 31. Thus, $\hat{m} = e^{sd} d\pi d_1 \lambda d\tau ds$. We denote $m = d\pi d_1 \lambda d\tau$, and view it as a measure on $\mathcal{H} = \mathcal{H}(C)$.

Lemma 21.1. *The measure \hat{m} is invariant under the Rauzy-Veech maps $\hat{\mathcal{R}}$ and \mathcal{R} . Moreover, m is invariant under the restriction $\mathcal{R} : \mathcal{H} \rightarrow \mathcal{H}$.*

Proof. Recall that $\hat{\mathcal{R}}(\pi, \lambda, \tau) = (\pi', \lambda', \tau')$ where $\lambda = \Theta_{\pi, \lambda}^*(\lambda')$ and $\tau = \Theta_{\pi, \lambda}^*(\tau')$. Since $\hat{\mathcal{R}}$ is injective and

$$\det \Theta_{\pi, \lambda}^* = \det \Theta_{\pi, \lambda} = 1,$$

it follows that $\hat{\mathcal{R}}$ preserves $\hat{m} = d\pi d\lambda d\tau$, as claimed. Now, in view of the definition (59), to prove that \hat{m} is preserved by \mathcal{R} we only have to show that it is preserved by

$$(\pi, \lambda, \tau) \mapsto \mathcal{T}^{t_R(\pi, \lambda)}(\pi, \lambda, \tau).$$

Using the coordinates (π, λ, τ, s) introduced previously, this corresponds to showing that the measure $e^{sd} d_1\lambda d\tau ds$ is invariant under the map

$$\Phi : (\lambda, \tau, s) \mapsto (\lambda, e^{-t_R(\pi, \lambda)}\tau, s + t_R(\pi, \lambda)).$$

The Jacobian matrix of Φ has the form

$$D\Phi = \begin{pmatrix} I_{d-1} & 0 & 0 \\ * & e^{-t_R}I_d & 0 \\ * & 0 & 1 \end{pmatrix}$$

(I_j denotes the j -dimensional identity matrix.) and so its determinant is $e^{-t_R d}$. Hence,

$$e^{(s+t_R)d} d_1\lambda d\tau ds |\det D\Phi| = e^{sd} d_1\lambda d\tau ds,$$

which means that Φ does preserve $e^{sd} d_1\lambda d\tau ds$.

Finally, \mathcal{R} preserves every $\mathcal{H}_c = \{(\pi, \lambda, \tau, s) \in \hat{\mathcal{H}} : e^s = c\}$ and the measure $\hat{m} = e^{sd} d\pi d_1\lambda d\tau ds$ disintegrates to conditional measures $c^d d\pi d_1\lambda d\tau$ on each \mathcal{H}_c . So, the previous conclusion that \mathcal{R} preserves \hat{m} means that it preserves these conditional measures for almost every c . From the definition (59) we get that $\lambda \mapsto c\lambda$ conjugates the restrictions of \mathcal{R} to \mathcal{H} and to \mathcal{H}_c , respectively. Consequently, $\mathcal{R} \upharpoonright \mathcal{H}_c$ preserves $c^d d\pi d_1\lambda d\tau$ if and only if the renormalization map $\mathcal{R} : \mathcal{H} \rightarrow \mathcal{H}$ preserves $m = d\pi d_1\lambda d\tau$. It follows that $\mathcal{R} : \mathcal{H} \rightarrow \mathcal{H}$ does preserve m , as claimed. \square

Given any $c > 0$, we denote by \hat{m}_c the restriction of \hat{m} to the region $\{(\pi, \lambda, \tau) \in \hat{\mathcal{H}} : \text{area}(\lambda, \tau) \leq c\}$. Since this region is invariant under $\hat{\mathcal{R}}$, \mathcal{R} , and \mathcal{T} , so are all these measures \hat{m}_c . Similarly, we denote by m_c the restriction of m to the region $\{(\pi, \lambda, \tau) \in \mathcal{H} : \text{area}(\lambda, \tau) \leq c\}$. Then every m_c is invariant under the restriction of Rauzy-Veech renormalization \mathcal{R} .

Recall that the pre-stratum $\hat{\mathcal{S}} = \hat{\mathcal{S}}(C)$ is the quotient of the space $\hat{\mathcal{H}}$ by the equivalence relation generated by

$$\mathcal{T}^{t_R(\pi, \lambda)}(\pi, \lambda, \tau) = (\pi, e^{t_R(\pi, \lambda)}\lambda, e^{-t_R(\pi, \lambda)}\tau) \sim \mathcal{R}(\pi, \lambda, \tau).$$

Since the Teichmüller flow commutes with \mathcal{R} , it projects down to a flow on $\hat{\mathcal{S}}$, that we also denote by \mathcal{T} . The (injective) image $\mathcal{S} \subset \hat{\mathcal{S}}$ of \mathcal{H} under the quotient map is a global cross-section to this flow. Moreover, the restriction of \hat{m} to the fundamental domain

$$\{(\pi, \lambda, \tau) \in \hat{\mathcal{H}} : 0 \leq \log|\lambda| \leq t_R(\pi, \lambda)\}$$

defines a volume measure on $\hat{\mathcal{S}}$, that we also denote by \hat{m} . It is easy to check that \hat{m} is invariant under the Teichmüller flow \mathcal{T}^t on the pre-stratum $\hat{\mathcal{S}}$. Finally, since area is invariant under the equivalence relation above, it is well defined in the pre-stratum. Sometimes, we denote by \hat{m}_c the restriction of \hat{m} to the subset of elements of the pre-stratum with $\text{area}(\pi, \lambda, \tau) \leq c$. All these measures are invariant under the Teichmüller flow on $\hat{\mathcal{S}}$.

For interval exchange maps. Let $P : \hat{\mathcal{H}} \rightarrow C \times \mathbb{R}_+^A$ be the canonical projection $P(\pi, \lambda, \tau) = (\pi, \lambda)$. Then let $\hat{\nu} = P_*(\hat{m}_1)$ be the measure obtained by projecting \hat{m}_1 down to $C \times \mathbb{R}_+^A$:

$$\hat{\nu}(E) = \hat{m}_1(P^{-1}(E)) = \hat{m}(\{(\pi, \lambda, \tau) : (\pi, \lambda) \in E \text{ and } \text{area}(\lambda, \tau) \leq 1\}).$$

Let \hat{R} and R be the Rauzy-Veech transformations at the level of interval exchange maps, introduced in sections 2 and 7. Likewise, let T^t be the projected Teichmüller flow $T^t(\pi, \lambda) = (\pi, e^t \lambda)$. Since

$$P \circ \mathcal{T}^t = T^t \circ P, \quad P \circ \hat{\mathcal{R}} = \hat{R} \circ P, \quad \text{and} \quad P \circ \mathcal{R} = R \circ P,$$

the measure $\hat{\nu}$ is invariant under \hat{R} , R , and T . Moreover, let $\nu = P_*(m_1)$ be the measure obtained by projecting m_1 down to $C \times \Lambda_A$:

$$\nu(E) = m_1(P^{-1}(E)) = m(\{(\pi, \lambda, \tau) : (\pi, \lambda) \in E \text{ and } \text{area}(\lambda, \tau) \leq 1\}).$$

Then ν is invariant under Rauzy-Veech renormalization $R : C \times \Lambda_A \rightarrow C \times \Lambda_A$.

Let \hat{S} be the quotient of $C \times \mathbb{R}_+^A$ by the equivalence relation generated on $C \times \mathbb{R}_+^A$ by $T^{t_R(\pi, \lambda)}(\pi, \lambda) \sim R(\pi, \lambda)$. We represent by S the (injective) image of $C \times \Lambda_A$ under this quotient map. The flow T^t induces a semi-flow $T^t : \hat{S} \rightarrow \hat{S}$, $t > 0$ which admits S as a global cross-section and whose first return map to this cross-section is the Rauzy-Veech renormalization $R : C \times \Lambda_A \rightarrow C \times \Lambda_A$.

The projection $P : \hat{\mathcal{H}} \rightarrow C \times \mathbb{R}_+^A$ induces a projection $P : \hat{S} \rightarrow \hat{S}$ such that $P \circ \mathcal{T}^t = T^t \circ P$. The absolutely continuous measure $\hat{\nu}$ restricted to the fundamental domain

$$\{(\pi, \lambda) \in C \times \mathbb{R}_+^A : 0 \leq \log|\lambda| \leq t_R(\pi, \lambda)\}$$

induces an absolutely continuous measure on \hat{S} , that we also denote as $\hat{\nu}$. It may also be obtained as $\hat{\nu} = P_*(\hat{m})$ where \hat{m} denotes the volume measure on \hat{S} introduced previously. It follows from $P \circ \mathcal{T}^t = T^t \circ P$ that $\hat{\nu}$ is invariant under the semi-flow T^t .

Example 21.2. For $d = 2$, the domain \mathbb{R}_+^A may be identified with $\mathbb{R} \times (0, 1)$, through

$$(\lambda_A, \lambda_B) \mapsto (\log|\lambda|, \lambda_A).$$

Note that the simplex Λ_A is identified with the interval $(0, 1)$, via $(\lambda_A, \lambda_B) \mapsto x = \lambda_A$. Then $d_1\lambda$ corresponds to the measure dx , and the Rauzy renormalization time is

$$t_R(x) = \begin{cases} -\log(1-x) & \text{if } x < 1/2, \\ -\log x & \text{if } x > 1/2. \end{cases} \tag{60}$$

\hat{S} is the quotient of the domain $\{(s, x) : 0 \leq s \leq t_R(x)\}$ by an identification of the boundary segment on the left with each of the two boundary curves on the right. See figure 32. The semi-flow T^t is horizontal, pointing to the right, and its return map to $\{0\} \times (0, 1)$ is the renormalization map R as presented in Example 7.1. The pre-stratum $\hat{S} = \hat{S} \times T_\pi^+$, where T_π^+ is the set of pairs (τ_A, τ_B) such that $\tau_A > 0 > \tau_B$.




Figure 32

Invariant densities. Since P is a submersion, the measure $\hat{\nu}$ is absolutely continuous with respect to $d\lambda$ (or, more precisely, $d\pi \times d\lambda$), with density

$$\frac{d\hat{\nu}}{d\lambda}(\pi, \lambda) = \text{vol}(\{\tau \in T_{\pi}^{+} : \text{area}(\lambda, \tau) \leq 1\}) \quad \text{for } (\pi, \lambda) \in C \times \mathbb{R}_{+}^{\mathcal{A}},$$

where $\text{vol}(\cdot)$ represents d -dimensional volume in T_{π}^{+} . Analogously, ν is absolutely continuous with respect to $d_1\lambda$ (or, more precisely, $d\pi \times d_1\lambda$), with density

$$\frac{d\nu}{d_1\lambda}(\pi, \lambda) = \text{vol}(\{\tau \in T_{\pi}^{+} : \text{area}(\lambda, \tau) \leq 1\}) \quad \text{for } (\pi, \lambda) \in C \times \Lambda_{\mathcal{A}}.$$

The right hand side in these expressions may be calculated as follows. (An explicit example will be worked out in section 22.)

The polyhedral cone T_{π}^{+} may be written, up to a codimension 1 subset, as a finite union of *simplicial cones* T^1, \dots, T^k , that is, subsets of $\mathbb{R}^{\mathcal{A}}$ of the form

$$T^i = \left\{ \sum_{\beta \in \mathcal{A}} c_{\beta} \tau^{i, \beta} : c_{\beta} > 0 \text{ for each } \beta \in \mathcal{A} \right\},$$

for some basis $(\tau^{i, \beta})_{\beta \in \mathcal{A}}$ of $\mathbb{R}^{\mathcal{A}}$. We always assume that this basis has been chosen with volume 1, that is, it is the image of some orthonormal basis by a linear operator with determinant 1. The volume of each domain

$$\{\tau \in T^i : \text{area}(\lambda, \tau) \leq 1\} = \{\tau \in T^i : -\lambda \cdot \Omega_{\pi}(\tau) \leq 1\}$$

may be calculated using the following elementary fact:

Lemma 21.3. *Let $T \subset \mathbb{R}^{\mathcal{A}}$ be a simplicial cone, $(\tau^{\beta})_{\beta \in \mathcal{A}}$ be a volume 1 basis of generators of T , and $L : \mathbb{R}^{\mathcal{A}} \rightarrow \mathbb{R}^{\mathcal{A}}$ be a linear operator. Then, for any λ satisfying $\lambda \cdot L(\tau^{\beta}) > 0$ for all $\beta \in \mathcal{A}$, we have*

$$\text{vol}(\{\tau \in T : \lambda \cdot L(\tau) \leq 1\}) = \frac{1}{d!} \prod_{\beta \in \mathcal{A}} \frac{1}{\lambda \cdot L(\tau^{\beta})}.$$

Proof. Let $M : \mathbb{R}^{\mathcal{A}} \rightarrow \mathbb{R}^{\mathcal{A}}$ be a linear operator mapping the canonical basis $(e^\beta)_{\beta \in \mathcal{A}}$ of $\mathbb{R}^{\mathcal{A}}$ to the basis $(\tau^\beta)_{\beta \in \mathcal{A}}$. Then let $\tilde{T} = M^{-1}(T)$ and $\tilde{L} = LM$. Then

$$\text{vol}(\{\tau \in T : \lambda \cdot L(\tau) \leq 1\}) = \text{vol}(\{v \in \tilde{T} : \lambda \cdot \tilde{L}(v) \leq 1\}). \tag{61}$$

Since T is a simplicial cone, \tilde{T} is the cone of vectors $v = \sum_{\beta \in \mathcal{A}} c_\beta e^\beta$ with entries $c_\beta > 0$ relative to the orthonormal basis. Then the set on the right hand side of (61) is the simplex with vertices at the origin and at each one of the points

$$\frac{e^\beta}{\lambda \cdot \tilde{L}(e^\beta)} = \frac{e^\beta}{\lambda \cdot L(\tau^\beta)}, \quad \beta \in \mathcal{A}.$$

Therefore,

$$\text{vol}(\{v \in \tilde{T} : \lambda \cdot \tilde{L}(v) \leq 1\}) = \prod_{\beta \in \mathcal{A}} \frac{1}{\lambda \cdot L(\tau^\beta)} \text{vol}(\Sigma_{\mathcal{A}})$$

where $\Sigma_{\mathcal{A}}$ is the canonical d -dimensional simplex, with vertices at the origin and at each of the points e^β , $\beta \in \mathcal{A}$. The latter has volume $1/d!$, and so the proof is complete. \square

Applying this lemma to each $T = T^i$ with $L = -\Omega_\pi$, we obtain

Proposition 21.4. *The density of $\hat{\nu}$ relative to Lebesgue measure is*

$$\frac{d\hat{\nu}}{d\lambda}(\pi, \lambda) = \text{vol}(\{\tau \in T_\pi^+ : \text{area}(\lambda, \tau) \leq 1\}) = \frac{1}{d!} \sum_{i=1}^k \prod_{\beta \in \mathcal{A}} \frac{1}{\lambda \cdot h^{i,\beta}}$$

where $h^{i,\beta} = -\Omega_\pi(\tau^{i,\beta})$. Moreover, the same formula holds for $d\nu/d_1\lambda$. In particular, all these densities are homogeneous rational functions with degree $-d$ and bounded away from zero.

Example 21.5. Let $d = 2$ and $\pi = \begin{pmatrix} A & B \\ B & A \end{pmatrix}$. The conditions (35) defining T_π^+ reduce to $\tau_A > 0 > \tau_B$. The operator Ω_π is given by

$$\Omega_\pi(\tau_A, \tau_B) = (\tau_B, -\tau_A)$$

and $\text{area}(\lambda, \tau) = \lambda_B \tau_A - \lambda_A \tau_B$. The operator Θ is given by

$$\Theta = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} \quad \text{if type} = 0 \quad \text{and} \quad \Theta = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \quad \text{if type} = 1$$

Figure 33 illustrates the action of the Rauzy transformation $\hat{\mathcal{R}}$ on the space of translation surfaces.




Figure 33

The measure $\hat{m} = d\lambda d\tau$ on $\hat{\mathcal{H}} = \{(\lambda, \tau) : \lambda_A > 0, \lambda_B > 0, \tau_A > 0 > \tau_B\}$ projects down to a measure $\hat{\nu}$ on \mathbb{R}_+^A which is absolutely continuous with respect to Lebesgue measure $d\lambda$, with density

$$\begin{aligned} \frac{d\hat{\nu}}{d\lambda}(\lambda) &= \text{vol}(\{\tau \in T_\pi^+ : \text{area}(\lambda, \tau) \leq 1\}) \\ &= \text{vol}(\{\tau \in \mathbb{R}^A : \tau_A > 0 > \tau_B \text{ and } \lambda_B \tau_A - \lambda_A \tau_B \leq 1\}), \end{aligned}$$

that is,

$$\frac{d\hat{\nu}}{d\lambda}(\lambda) = \frac{1}{2\lambda_A \lambda_B}.$$

The same expression holds for $d\nu/d_1\lambda$, restricted to Λ_A . Notice that the measure ν is infinite. Indeed, identifying Λ_A with $(0, 1)$ and $d_1\lambda$ with dx , through $(\lambda_A, \lambda_B) \mapsto x = \lambda_A$,

$$\nu(\Lambda_A) = \int_{\Lambda_A} \frac{1}{2\lambda_A \lambda_B} d_1\lambda = \int_0^1 \frac{1}{2x(1-x)} dx = \infty.$$

However, the measure ν is *finite* on \hat{S} . Indeed (recall Example 21.2)

$$\hat{\nu}(\hat{S}) = \int_0^1 \int_0^{t_R(x)} \frac{1}{2e^s x e^s (1-x)} e^{2s} dx ds = \int_0^1 t_R(x) \frac{1}{2x(1-x)} dx.$$

Using the expression (60), this becomes

$$\hat{\nu}(\hat{S}) = 2 \int_0^{1/2} -\log(1-x) \frac{1}{2x(1-x)} dx \leq 2 \int_0^{1/2} -\log(1-x) \frac{1}{x} dx < \infty.$$

This may be restated, equivalently, as $\hat{m}(\hat{S}_1) < \infty$. These conclusions are typical for all $d \geq 2$, as we shall see.

22. Hyperelliptic pairs

We are going to compute an explicit expression for the density in the case when

$$\pi_1 \circ \pi_0^{-1}(j) = d - j + 1 \quad \text{for } j = 1, \dots, d. \tag{62}$$

Denote

$$b_\alpha^\varepsilon = \sum_{\pi_\varepsilon(\beta) \leq \pi_\varepsilon(\alpha)} \tau_\beta \quad \text{for each } \alpha \in \mathcal{A} \text{ and } \varepsilon \in \{0, 1\}. \tag{63}$$

Note that $\sum_{\alpha \in \mathcal{A}} \tau_\alpha = b_{\alpha(\varepsilon)}^\varepsilon$ for $\varepsilon = 0, 1$. The cone T_π^+ is defined by

$$b_\alpha^0 > 0 \quad \text{for } \alpha \neq \alpha(0) \quad \text{and} \quad b_\alpha^1 < 0 \quad \text{for } \alpha \neq \alpha(1), \tag{64}$$

which is just a reformulation of (35). Let T_π^0 and T_π^1 be the subsets of T_π^+ defined by

$$\tau \in T_\pi^0 \Leftrightarrow \sum_{\alpha \in \mathcal{A}} \tau_\alpha > 0 \quad \text{and} \quad \tau \in T_\pi^1 \Leftrightarrow \sum_{\alpha \in \mathcal{A}} \tau_\alpha < 0. \tag{65}$$

Clearly, $T_\pi^+ = T_\pi^0 \cup T_\pi^1$, up to a codimension 1 subset.

Given $\alpha \in \mathcal{A}$ and $\varepsilon \in \{0, 1\}$, denote by α_ε^- the symbol to the left and by α_ε^+ the symbol to the right of α in line ε . That is,

$$\begin{aligned} \alpha_\varepsilon^- &= \pi_\varepsilon^{-1}(\pi_\varepsilon(\alpha) - 1) \quad \text{if } \pi_\varepsilon(\alpha) > 1, \\ \alpha_\varepsilon^+ &= \pi_\varepsilon^{-1}(\pi_\varepsilon(\alpha) + 1) \quad \text{if } \pi_\varepsilon(\alpha) < d. \end{aligned} \tag{66}$$

Lemma 22.1. T_π^ε is a simplicial cone for every $\varepsilon \in \{0, 1\}$.

Proof. We treat the case $\varepsilon = 0$, the other one being entirely analogous. For notational simplicity, let $b_\alpha = b_\alpha^0$ for every $\alpha \in \mathcal{A}$. Note that, because of (62),

$$b_\alpha^0 + b_\alpha^1 = \sum_{\beta \in \mathcal{A}} \tau_\beta + \tau_\alpha.$$

Equivalently,

$$b_{\alpha_0^-}^0 + b_\alpha^1 = \sum_{\beta \in \mathcal{A}} \tau_\beta = b_\alpha^0 + b_{\alpha_1^-}^1$$

(the first equality is for $\alpha \neq \alpha(1)$, the second one for $\alpha \neq \alpha(0)$). In particular,

$$b_\alpha^1 = \sum_{\beta \in \mathcal{A}} \tau_\beta - b_{\alpha_0^-}^0 = b_{\alpha(0)}^0 - b_{\alpha_0^-}^0 = b_{\alpha(0)} - b_{\alpha_0^-}.$$

Notice that when α varies in $\mathcal{A} \setminus \{\alpha(1)\}$ the symbol α_0^- varies in $\mathcal{A} \setminus \{\alpha(0)\}$:

$$\begin{pmatrix} \alpha(1) & \cdots & \alpha_0^- & \alpha & \cdots & \alpha(0) \\ \alpha(0) & \cdots & \cdots & \cdots & \cdots & \alpha(1) \end{pmatrix}.$$

Then (64) becomes

$$b_\alpha > 0 \text{ for } \alpha \neq \alpha(0) \text{ and } b_{\alpha(0)} - b_\beta < 0 \text{ for } \beta \neq \alpha(0),$$

and (65) gives that the cone T_π^0 is described by

$$b_\alpha > 0 \text{ for all } \alpha \in \mathcal{A} \text{ and } 0 < b_{\alpha(0)} < \min_{\beta \neq \alpha(0)} b_\beta. \tag{67}$$

Now it is easy to exhibit a basis of generators: take $b^\alpha = (b_\beta^\alpha)_{\beta \in \mathcal{A}}$ with

$$b_\beta^\alpha = \begin{cases} 1 & \text{if } \beta = \alpha, \\ 0 & \text{otherwise,} \end{cases} \text{ if } \alpha \neq \alpha(0), \tag{68}$$

$$b_\beta^\alpha = 1 \text{ for every } \beta \in \mathcal{A} \text{ if } \alpha = \alpha(0).$$

A vector $b = (b_\beta)_{\beta \in \mathcal{A}}$ satisfies (67) if and only if it can be written in the form $b = \sum_{\alpha \in \mathcal{A}} c_\alpha b^\alpha$ with $c_\alpha > 0$ for all $\alpha \in \mathcal{A}$. It follows that T_π^0 is a simplicial cone admitting the basis $\tau^\alpha = (\tau_\beta^\alpha)_{\beta \in \mathcal{A}}$ given by $\tau_\beta^\alpha = b_\beta^\alpha - b_{\beta_0^-}^\alpha$, that is,

$$\tau_\beta^\alpha = \begin{cases} 1 & \text{if } \beta = \alpha, \\ -1 & \text{if } \beta = \alpha_0^+, \\ 0 & \text{in all other cases,} \end{cases} \text{ if } \alpha \neq \alpha(0), \tag{69}$$

$$\tau_\beta^\alpha = \begin{cases} 1 & \text{if } \pi_0(\beta) = 1, \\ 0 & \text{otherwise,} \end{cases} \text{ if } \alpha = \alpha(0).$$

This completes the proof. □

Let $h^\alpha = -\Omega_\pi(\tau^\alpha)$, where $(\tau^\alpha)_{\alpha \in \mathcal{A}}$ is the basis of T_π^0 we found in (69), that is,

$$h_\beta^\alpha = \begin{cases} 1 & \text{if } \beta = \alpha \text{ or } \beta = \alpha_0^+, \\ 0 & \text{otherwise,} \end{cases} \text{ if } \alpha \neq \alpha(0),$$

$$h_\beta^\alpha = \begin{cases} 0 & \text{if } \pi_0(\beta) = 1 \text{ or } \beta = \alpha(1), \\ 1 & \text{otherwise,} \end{cases} \text{ if } \alpha = \alpha(0).$$

It is clear that the basis $(b^\alpha)_{\alpha \in \mathcal{A}}$ defined by (68) has volume 1. Since the map

$$b \mapsto \tau, \quad \tau_\beta = b_\beta - b_{\beta_0^-}$$

has determinant 1, it follows that $(\tau^\alpha)_{\alpha \in \mathcal{A}}$ also has volume 1. So, by Lemma 21.3, the contribution of the cone T_π^0 to the density is

$$\frac{1}{d!} \prod_{\alpha \in \mathcal{A}} \frac{1}{\lambda \cdot h^\alpha} = \frac{1}{d!} \prod_{\alpha \neq \alpha(0)} \left(\frac{1}{\lambda_\alpha + \lambda_{\alpha_0^+}} \right) \cdot \frac{1}{\sum_{\beta \neq \alpha(1)} \lambda_\beta}.$$

There is a completely symmetric calculation for T_π^1 . In this way, we get the following formula for the density in this case:

Proposition 22.2. *If π satisfies (62) then the invariant density is*

$$\frac{d\hat{\nu}}{d\lambda}(\pi, \lambda) = \sum_{\varepsilon=0,1} \frac{1}{d!} \prod_{\alpha \neq \alpha(\varepsilon)} \left(\frac{1}{\lambda_\alpha + \lambda_{\alpha_\varepsilon^+}} \right) \cdot \frac{1}{\sum_{\beta \neq \alpha(1-\varepsilon)} \lambda_\beta}$$

and $d\nu/d_1\lambda$ is given by the same expression, restricted to $C \times \Lambda_{\mathcal{A}}$.

Example 22.3. Let $d = 5$ and $\mathcal{A} = \{A, B, C, D, E\}$. Then

$$\pi = \begin{pmatrix} A & B & C & D & E \\ E & D & C & B & A \end{pmatrix}.$$

The cone T_π^0 is described by

$$\begin{aligned} b_A^0 > 0, & \quad b_B^0 > 0, \quad b_C^0 > 0, \quad b_D^0 > 0, \quad b_E^0 > 0, \\ b_E^1 < 0, & \quad b_D^1 < 0, \quad b_C^1 < 0, \quad b_B^1 < 0, \end{aligned}$$

that is,

$$\begin{aligned} b_A > 0, & \quad b_B > 0, \quad b_C > 0, \quad b_D > 0, \quad b_E > 0, \\ b_E - b_D < 0, & \quad b_E - b_C < 0, \quad b_E - b_B < 0, \quad b_E - b_A < 0, \end{aligned}$$

or, equivalently,

$$b_A > 0, \quad b_B > 0, \quad b_C > 0, \quad b_D > 0, \quad 0 < b_E < \min\{b_A, b_B, b_C, b_D\}.$$

As a basis take

$$\begin{aligned} b^A &= (1, 0, 0, 0, 0), & b^B &= (0, 1, 0, 0, 0), & b^C &= (0, 0, 1, 0, 0), \\ b^D &= (0, 0, 0, 1, 0), & b^E &= (1, 1, 1, 1, 1) \end{aligned}$$

or, equivalently,

$$\begin{aligned} \tau^A &= (1, -1, 0, 0, 0), & \tau^B &= (0, 1, -1, 0, 0), & \tau^C &= (0, 0, 1, -1, 0), \\ \tau^D &= (0, 0, 0, 1, -1), & \tau^E &= (1, 0, 0, 0, 0). \end{aligned}$$

We may write any $\tau = (\tau_A, \tau_B, \tau_C, \tau_D, \tau_E) \in T_\pi^0$ as

$$\tau = (b_A - b_E)\tau^A + (b_B - b_E)\tau^B + (b_C - b_E)\tau^C + (b_D - b_E)\tau^D + b_E\tau^E$$

where the coefficients are all positive. Moreover,

$$\begin{aligned} h^A &= (1, 1, 0, 0, 0), & h^B &= (0, 1, 1, 0, 0), & h^C &= (0, 0, 1, 1, 0), \\ h^D &= (0, 0, 0, 1, 1), & h^E &= (0, 1, 1, 1, 1). \end{aligned}$$

Hence, the contribution of T_π^0 to the density is

$$\frac{1}{5!} \cdot \frac{1}{\lambda_A + \lambda_B} \cdot \frac{1}{\lambda_B + \lambda_C} \cdot \frac{1}{\lambda_C + \lambda_D} \cdot \frac{1}{\lambda_D + \lambda_E} \cdot \frac{1}{\lambda_B + \lambda_C + \lambda_D + \lambda_E}.$$

The cone T_π^1 contributes

$$\frac{1}{5!} \cdot \frac{1}{\lambda_E + \lambda_D} \cdot \frac{1}{\lambda_D + \lambda_C} \cdot \frac{1}{\lambda_C + \lambda_B} \cdot \frac{1}{\lambda_B + \lambda_A} \cdot \frac{1}{\lambda_A + \lambda_B + \lambda_C + \lambda_D},$$

and so the total density is (recall that $|\lambda| = \sum_{\alpha \in \mathcal{A}} \lambda_\alpha$)

$$\frac{1}{5!} \cdot \frac{1}{(\lambda_A + \lambda_B)(\lambda_B + \lambda_C)(\lambda_C + \lambda_D)(\lambda_D + \lambda_E)} \left(\frac{1}{|\lambda| - \lambda_A} + \frac{1}{|\lambda| - \lambda_E} \right).$$

23. Combinatorial statement

We want to prove that the intersection of every pre-stratum with the set of (π, λ, τ) such that $\text{area}(\pi, \lambda, \tau) \leq 1$ has finite volume. The crucial step is

Proposition 23.1. *Let $(\tau^\beta)_{\beta \in \mathcal{A}}$ be a basis of $\mathbb{R}^{\mathcal{A}}$ contained in the closure of T_π^δ for some $\delta \in \{0, 1\}$, and let $h^\beta = -\Omega_\pi(\tau^\beta)$ for $\beta \in \mathcal{A}$. Given any non-empty proper subset \mathcal{B} of \mathcal{A} , we have*

$$\#\{\beta \in \mathcal{A} : h_\alpha^\beta = 0 \text{ for all } \alpha \in \mathcal{B}\} + \#\mathcal{B} \leq d,$$

and the inequality is strict unless \mathcal{B} contains $\alpha(1 - \delta)$ but not $\alpha(\delta)$.

Proof. We suppose $\delta = 0$, as the other case is analogous. Let $h = -\Omega_\pi(\tau)$ for some τ in the closure of T_π^0 . By (37) and (63),

$$h_\alpha = b_\alpha^0 - b_\alpha^1 = b_{\alpha_0^-}^0 - b_{\alpha_1^-}^1. \tag{70}$$

The symbol α_ε^- is not defined when $\pi_\varepsilon(\alpha) = 1$, but (70) remains valid in that case, as long as one interprets $b_{\alpha_\varepsilon^-}^\varepsilon$ to be zero. By the definition of $T_\pi^0 \subset T_\pi^+$ in (35) and (65), and the assumption that τ is in the closure of T_π^0 ,

$$b_\alpha^0 \geq 0 \text{ for all } \alpha \in \mathcal{A} \quad \text{and} \quad b_\alpha^1 \leq 0 \text{ for all } \alpha \in \mathcal{A} \setminus \{\alpha(1)\}.$$

Therefore, given any $\alpha \neq \alpha(1)$,

$$h_\alpha = 0 \quad \Rightarrow \quad b_\alpha^0 = b_\alpha^1 = 0 = b_{\alpha_0^-}^0 = b_{\alpha_1^-}^1. \tag{71}$$

A part of (71) remains valid even when $\alpha = \alpha(1)$:

$$h_\alpha = 0 \quad \Rightarrow \quad b_{\alpha_0^-}^0 = b_{\alpha_1^-}^1 = 0, \tag{72}$$

because $\alpha_1^- \neq \alpha(1)$. Finally, adding the relations

$$h_{\alpha(0)} = b_{\alpha(0)}^0 - b_{\alpha(0)}^1 \quad \text{and} \quad h_{\alpha(1)} = b_{\alpha(1)}^0 - b_{\alpha(1)}^1,$$

and recalling that $b_{\alpha(0)}^0 = \sum_{\beta \in \mathcal{A}} \tau_\alpha = b_{\alpha(1)}^1$, we get that

$$h_\alpha = 0 \text{ for both } \alpha \in \{\alpha(0), \alpha(1)\} \quad \Rightarrow \quad b_{\alpha(1)}^0 = b_{\alpha(0)}^1 = 0. \quad (73)$$

Now let \mathcal{B} be a non-empty proper subset of \mathcal{A} , and assume $h_\alpha = 0$ for all $\alpha \in \mathcal{B}$.

Case 1: \mathcal{B} does not contain $\alpha(1)$. Define

$$\mathcal{B}_\varepsilon = \mathcal{B} \cup \{\alpha_\varepsilon^- : \alpha \in \mathcal{B}\} \quad \text{for } \varepsilon \in \{0, 1\}. \quad (74)$$

Then (71) gives that

$$b_\beta^\varepsilon = 0 \text{ for all } \beta \in \mathcal{B}_\varepsilon \text{ and } \varepsilon \in \{0, 1\}. \quad (75)$$

We claim that there exists $\varepsilon \in \{0, 1\}$ such that

$$\#\mathcal{B}_\varepsilon > \#\mathcal{B}. \quad (76)$$

Indeed, it follows from the definition (74) that \mathcal{B} is contained in \mathcal{B}_ε . Moreover, the two sets coincide only if $\alpha_\varepsilon^- \in \mathcal{B}$ for every $\alpha \in \mathcal{B}$ or, in other words, if

$$\mathcal{B} = \pi_\varepsilon^{-1}(\{1, \dots, k\}) \quad \text{for some } 1 \leq k \leq d. \quad (77)$$

Note that $k < d$, because \mathcal{B} is a proper subset of \mathcal{A} . So, since π is irreducible, (77) can not hold simultaneously for both $\varepsilon = 0$ and $\varepsilon = 1$. Hence, there exists ε such that $\mathcal{B}_\varepsilon \neq \mathcal{B}$. This proves the claim. Now fix any such ε . Since the map $\tau \mapsto b^\varepsilon$ is injective, and the $(\tau^\beta)_{\beta \in \mathcal{A}}$ are linearly independent, (75) and (76) give

$$\#\{\beta \in \mathcal{A} : h_\alpha^\beta = 0 \text{ for all } \alpha \in \mathcal{B}\} \leq d - \#\mathcal{B}_\varepsilon < d - \#\mathcal{B}.$$

Case 2: \mathcal{B} contains $\alpha(1)$ but not $\alpha(0)$. Let $\mathcal{B}_1 = \mathcal{B} \setminus \{\alpha(1)\} \cup \{\alpha_1^- : \alpha \in \mathcal{B}\}$. The relations (71) and (72) imply that

$$b_\beta^1 = 0 \text{ for all } \beta \in \mathcal{B}_1.$$

Let $k \geq \pi_1(\alpha(0))$ be maximum such that $\bar{\beta} = \pi_1^{-1}(k)$ is not in \mathcal{B} . The assumption that \mathcal{B} contains $\alpha(1)$ but not $\alpha(0)$ ensures that k is well defined and less than d . Then $\bar{\beta} = \alpha_1^-$ for some $\alpha \in \mathcal{B}$, and so $\bar{\beta} \in \mathcal{B}_1$. This shows that

$$\mathcal{B}_1 \supset \mathcal{B} \setminus \{\alpha(1)\} \cup \{\bar{\beta}\},$$

and so $\#\mathcal{B}_1 \geq \#\mathcal{B}$. Hence, just as before,

$$\#\{\beta \in \mathcal{A} : h_\alpha^\beta = 0 \text{ for all } \alpha \in \mathcal{B}\} \leq d - \#\mathcal{B}_1 \leq d - \#\mathcal{B}.$$

Case 3: \mathcal{B} contains both $\alpha(0)$ and $\alpha(1)$. Define $\mathcal{B}_0 = \mathcal{B} \cup \{\alpha_0^- : \alpha \in \mathcal{B}\}$. By (71)–(73),

$$b_\beta^0 = 0 \quad \text{for all } \beta \in \mathcal{B}_0.$$

It is easy to check that \mathcal{B}_0 contains \mathcal{B} strictly. Indeed, the two sets can coincide only if $\alpha_0^- \in \mathcal{B}$ for every $\alpha \in \mathcal{B}$, that is, if $\mathcal{B} = \pi_0^{-1}(\{1, \dots, k\})$ for some k . Since \mathcal{B} contains $\alpha(0) = \pi_0^{-1}(d)$, this would imply $\mathcal{B} = \mathcal{A}$, contradicting the hypothesis. It follows, just as in the first case, that

$$\#\{\beta \in \mathcal{A} : h_\alpha^\beta = 0 \text{ for all } \alpha \in \mathcal{B}\} \leq d - \#\mathcal{B}_0 < d - \#\mathcal{B}.$$

The proof of Proposition 23.1 is complete. □

Remark 23.2. The inequality in Proposition 23.1 is not always strict. Indeed, let τ^A, \dots, τ^E be the generators of T_π^0 in Example 22.3, and let $\mathcal{B} = \{A\}$. Then \mathcal{B} contains $A = \alpha(1)$ but not $E = \alpha(0)$. Note also that

$$\{\beta : h_A^\beta = 0\} = \{B, C, D, E\}$$

has exactly $4 = d - \#\mathcal{B}$ elements. Thus, the equality holds in this case. In fact, if the inequality were strict in all cases, then arguments as in the next section would imply that the measure ν is finite. However, the latter is usually not true, as we have already seen in Example 21.5.

24. Finite volume

Let C be a Rauzy class and $\hat{\mathcal{S}} = \hat{\mathcal{S}}(C)$ be the corresponding pre-stratum. Define the normalized pre-stratum to be the subset $\hat{\mathcal{S}}_1 = \hat{\mathcal{S}}_1(C)$ of all $(\pi, \lambda, \tau) \in \hat{\mathcal{S}}$ such that $\text{area}(\lambda, \tau) \leq 1$.

Theorem 24.1. *For every Rauzy class C , the normalized pre-stratum $\hat{\mathcal{S}}_1$ has finite volume: $\hat{m}(\hat{\mathcal{S}}_1) < \infty$.*

Proof. Recall that $\hat{\mathcal{S}}_1$ is obtained from the subset of all $(\pi, \lambda, \tau) \in \hat{\mathcal{H}}$ such that $\text{area}(\lambda, \tau) \leq 1$ and

$$\sum_{\alpha \neq \alpha(1-\varepsilon)} \lambda_\alpha \leq 1 \leq \sum_{\alpha \in \mathcal{A}} \lambda_\alpha \tag{78}$$

by identifying (π, λ, τ) with $\hat{\mathcal{R}}(\pi, \lambda, \tau)$ when $\sum_{\alpha \neq \alpha(1-\varepsilon)} \lambda_\alpha = 1$. Thus,

$$\text{vol}(\hat{\mathcal{S}}_1) = \sum_{\pi \in C} \int \rho(\pi, \lambda) d\lambda, \tag{79}$$

where $\rho(\pi, \lambda)$ is the d -dimensional volume of $\{\tau \in T_\pi^+ : \text{area}(\lambda, \tau) \leq 1\}$, and the integral is over the set of $\lambda \in \mathbb{R}_+^{\mathcal{A}}$ satisfying (78). Let $T^i, i = 1, \dots, k$ be a decomposition of T_π^+ (up to a codimension 1 subset) into simplicial cones. Then, by Proposition 21.4,

$$\rho(\pi, \lambda) = \frac{1}{d!} \sum_{i=1}^k \prod_{\beta \in \mathcal{A}} \frac{1}{\lambda \cdot h^{i,\beta}} \tag{80}$$

where $h^{i,\beta} = -\Omega_\pi(\tau^{i,\beta})$ and $(\tau^{i,\beta})_{\beta \in \mathcal{A}}$ is a basis of generators of T^i . We may assume that each T^i is contained either in T_π^0 or in T_π^1 , and we do so in what follows. Let us consider (compare (57) also)

$$\Lambda_{\mathcal{A}} \times \mathbb{R} \ni (\lambda, s) \mapsto e^s \lambda \in \mathbb{R}_+^{\mathcal{A}}.$$

Recall that $d\lambda = e^{sd} d_1\lambda ds$, where $d_1\lambda$ denotes the $(d-1)$ -dimensional volume induced on the simplex $\Lambda_{\mathcal{A}}$ by the Riemannian metric of $\mathbb{R}^{\mathcal{A}}$. Notice that, given $(\lambda, s) \in \Lambda_{\mathcal{A}} \times \mathbb{R}$, the vector $e^s \lambda$ satisfies (78) if and only if $0 \leq s \leq t_R(\pi, \lambda)$, where t_R is the Rauzy renormalization time defined in (58). Recall also that $\lambda \mapsto \rho(\pi, \lambda)$ is homogeneous of degree $-d$. Thus, after change of variables, (79) becomes

$$\text{vol}(\hat{\mathcal{S}}_1) = \sum_{\pi \in C} \int_{\Lambda_{\mathcal{A}}} \int_0^{t_R(\pi, \lambda)} \rho(\pi, e^s \lambda) e^{sd} ds d_1\lambda = \sum_{\pi \in C} \int_{\Lambda_{\mathcal{A}}} \rho(\pi, \lambda) t_R(\pi, \lambda) d_1\lambda.$$

Using (80) and the definition of $t_R(\pi, \lambda)$, this gives

$$\text{vol}(\hat{\mathcal{S}}_1) = \frac{1}{d!} \sum_{\pi \in C} \sum_{i=1}^k \int_{\Lambda_{\mathcal{A}}} -\log(1 - \lambda_{\alpha(1-\varepsilon)}) \prod_{\beta \in \mathcal{A}} \frac{1}{\lambda \cdot h^{i,\beta}} d_1\lambda, \tag{81}$$

where ε is the type of (π, λ) . Therefore, to prove the theorem we only have to show that the integral is finite, for every fixed $\pi \in C$ and $i = 1, \dots, k$.

For simplicity, we write $h^\beta = h^{i,\beta}$ in what follows. Also, we assume T^i is contained in T_π^0 ; the other case is analogous. This implies the corresponding basis of generators $(\tau^{i,\beta})_{\beta \in \mathcal{A}}$ is contained in the closure of T_π^0 .

Let \mathcal{N} denote the set of integer vectors $n = (n_\alpha)_{\alpha \in \mathcal{A}}$ such that $n_\alpha \geq 0$ for all $\alpha \in \mathcal{A}$, and the n_α are not all zero. For each $n \in \mathcal{N}$, define

$$\Lambda(n) = \{ \lambda \in \Lambda_{\mathcal{A}} : 2^{-n_\alpha} \leq \lambda_\alpha d < 2^{-n_\alpha+1} \text{ for every } \alpha \in \mathcal{A} \}, \tag{82}$$

except that for $n_\alpha = 0$ the second inequality is omitted.

Lemma 24.2. *There exists $c_1 > 0$ depending only on the dimension d such that*

$$\text{vol}_{d-1} \Lambda(n) \leq c_1 2^{-\sum_{\mathcal{A}} n_\alpha}$$

for all $n \in \mathcal{N}$. Moreover, the family $\Lambda(n), n \in \mathcal{N}$ covers $\Lambda_{\mathcal{A}}$.

Proof. If $\sum_{\alpha \in \mathcal{A}} \lambda_\alpha = 1$ then $\lambda_\beta \geq 1/d$ for some $\beta \in \mathcal{A}$, and so λ belongs to some $\Lambda(n)$ with $n_\beta = 0$. This shows that these sets $\Lambda(n)$ do cover $\Lambda_{\mathcal{A}}$. To prove the volume estimate, fix n and $\beta \in \mathcal{A}$ such that $n_\beta = 0$. When λ varies in $\Lambda(n)$, the $(d - 1)$ -dimensional vector $(\lambda_\alpha)_{\alpha \neq \beta}$ varies in some subset $S(n)$ of the product space $\prod_{\alpha \neq \beta} [0, 2^{-n_\alpha + 1}]$. The $(d - 1)$ -dimensional volume of $S(n)$ is bounded above by $2^{d-1} 2^{-\sum_{\alpha \in \mathcal{A}} n_\alpha}$. Then, since $\Lambda(n)$ is a graph over $S(n)$,

$$\text{vol}_{d-1} \Lambda(n) \leq \sqrt{d} \text{vol}_{d-1} S(n) \leq c_1 2^{-\sum_{\alpha \in \mathcal{A}} n_\alpha},$$

where $c_1 = \sqrt{d} 2^{d-1}$. The proof is complete. □

It is clear that $\lambda_{\alpha(1-\varepsilon)} < 1/2$, and so

$$-\log(1 - \lambda_{\alpha(1-\varepsilon)}) \leq 2\lambda_{\alpha(1-\varepsilon)} = 2 \min\{\lambda_{\alpha(0)}, \lambda_{\alpha(1)}\}.$$

Therefore, for each fixed π and i , the integral in (81) is bounded above by

$$\sum_{n \in \mathcal{N}} \int_{\Lambda(n)} 2 \min\{\lambda_{\alpha(0)}, \lambda_{\alpha(1)}\} \prod_{\beta \in \mathcal{A}} \frac{1}{\lambda \cdot h^\beta} d_1 \lambda.$$

For each $\beta \in \mathcal{A}$, let $\mathcal{A}(\beta)$ be the subset of $\alpha \in \mathcal{A}$ such that $h_\alpha^\beta > 0$. Let $c_2 > 0$ be the minimum of the non-zero h_α^β , over all α and β . Then

$$\lambda \cdot h^\beta = \sum_{\alpha \in \mathcal{A}(\beta)} h_\alpha^\beta \lambda_\alpha \geq \sum_{\alpha \in \mathcal{A}(\beta)} c_2 d^{-1} 2^{-n_\alpha} \geq c_2 d^{-1} 2^{-\min_{\alpha \in \mathcal{A}(\beta)} n_\alpha}$$

for every $\lambda \in \Lambda(n)$ and $\beta \in \mathcal{A}$. Using Lemma 24.2 we deduce that

$$\int_{\Lambda(n)} 2 \min\{\lambda_{\alpha(0)}, \lambda_{\alpha(1)}\} \prod_{\beta \in \mathcal{A}} \frac{1}{\lambda \cdot h^\beta} d_1 \lambda \leq K 2^{-\max_\varepsilon n_{\alpha(\varepsilon)} + \sum_\beta \min_{\alpha \in \mathcal{A}(\beta)} n_\alpha - \sum_\alpha n_\alpha}, \quad (83)$$

where the constant $K = (2c_1)(d/c_2)^d$. Using Proposition 23.1, we obtain

Lemma 24.3.

$$\max_{\varepsilon \in \{0,1\}} n_{\alpha(\varepsilon)} - \sum_{\beta \in \mathcal{A}} \min_{\alpha \in \mathcal{A}(\beta)} n_\alpha + \sum_{\alpha \in \mathcal{A}} n_\alpha \geq \max_{\alpha \in \mathcal{A}} n_\alpha.$$

Proof. Let $0 = n^0 < n^1 < \dots$ be the different values taken by n_α , and \mathcal{B}^i , $i \geq 0$ be the set of values of $\alpha \in \mathcal{A}$ such that $n_\alpha \geq n^i$. On the one hand,

$$\sum_{\alpha \in \mathcal{A}} n_\alpha = \sum_{i \geq 1} n^i (\#\mathcal{B}^i - \#\mathcal{B}^{i+1}) = \sum_{i \geq 1} \#\mathcal{B}^i (n^i - n^{i-1}).$$

On the other hand, $\min_{\mathcal{A}(\beta)} n_\alpha \geq n^i$ if and only if $\mathcal{A}(\beta) \subset \mathcal{B}^i$. Consequently,

$$\begin{aligned} \sum_{\beta \in \mathcal{A}} \min_{\mathcal{A}(\beta)} n_\alpha &= \sum_{i \geq 1} n^i (\#\{\beta : \mathcal{A}(\beta) \subset \mathcal{B}^i\} - \#\{\beta : \mathcal{A}(\beta) \subset \mathcal{B}^{i+1}\}) \\ &= \sum_{i \geq 1} \#\{\beta : \mathcal{A}(\beta) \subset \mathcal{B}^i\} (n^i - n^{i-1}). \end{aligned}$$

Observe that $\mathcal{A}(\beta) \subset \mathcal{B}^i$ if and only if $h_\alpha^\beta = 0$ for all $\alpha \in \mathcal{A} \setminus \mathcal{B}^i$. So, by Proposition 23.1 (with $\mathcal{B} = \mathcal{A} \setminus \mathcal{B}^i$),

$$\#\{\beta : \mathcal{A}(\beta) \subset \mathcal{B}^i\} < \#\mathcal{B}^i \tag{84}$$

except, possibly, if \mathcal{B}^i contains $\alpha(0)$ but not $\alpha(1)$. On the one hand, if (84) does hold then

$$\#\mathcal{B}^i (n^i - n^{i-1}) - \#\{\beta : \mathcal{A}(\beta) \subset \mathcal{B}^i\} (n^i - n^{i-1}) \geq (n^i - n^{i-1}). \tag{85}$$

On the other hand, if \mathcal{B}^i contains $\alpha(0)$ but not $\alpha(1)$ then $n_{\alpha(1)} < n^i \leq n_{\alpha(0)}$. Let i_1 be the smallest and i_2 be the largest value of i for which this happens. Then

$$\#\mathcal{B}^i - \#\{\beta : \mathcal{A}(\beta) \subset \mathcal{B}^i\} \geq 0 \quad \text{for } i_1 \leq i \leq i_2 \tag{86}$$

and

$$\max\{n_{\alpha(0)}, n_{\alpha(1)}\} = n_{\alpha(0)} \geq n^{i_2} - n^{i_1-1} = \sum_{i=i_1}^{i_2} (n^i - n^{i-1}).$$

Putting (85) and (86) together, we find that

$$\max_{\varepsilon \in \{0,1\}} n_{\alpha(\varepsilon)} - \sum_{\beta \in \mathcal{A}} \min_{\mathcal{A}(\beta)} n_\alpha + \sum_{\alpha \in \mathcal{A}} n_\alpha \geq \sum_{i \geq 1} (n^i - n^{i-1}) = \max_{\alpha \in \mathcal{A}} n_\alpha.$$

This proves the lemma. □

Replacing the conclusion of the lemma in (83) we obtain, for every $n \in \mathcal{N}$,

$$\int_{\Lambda(n)} 2 \min\{\lambda_{\alpha(0)}, \lambda_{\alpha(1)}\} \prod_{\beta \in \mathcal{A}} \frac{1}{\lambda \cdot h^\beta} d_1 \lambda \leq K 2^{-\max_{\mathcal{A}} n_\alpha}. \tag{87}$$

For each $m \geq 0$ there are at most $(m+1)^d$ choices of $n \in \mathcal{N}$ with $\max_{\mathcal{A}} n_\alpha = m$. So, (87) implies that the integral in (81) is bounded above by

$$\sum_{m=0}^{\infty} K(m+1)^d 2^{-m} < \infty$$

for every $\pi \in C$ and every $1 \leq i \leq k$. The proof of Theorem 24.1 is complete. □

25. Recurrence and inducing

Given a measurable map $f : M \rightarrow M$ and a measure μ on M , we call (f, μ) *recurrent* if for any positive measure set $E \subset M$ and μ -almost every $x \in E$ there exists $n \geq 1$ such that $f^n(x) \in E$. The classical Poincaré recurrence theorem asserts that if μ is invariant and *finite* then (f, μ) is recurrent. Similar observations hold for flows as well.

Lemma 25.1. *The Teichmüller flow $T^t : \hat{S} \rightarrow \hat{S}$ and semi-flow $T^t : \hat{S} \rightarrow \hat{S}$ are recurrent, for the corresponding invariant measures \hat{m} and $\hat{\nu}$. The Rauzy-Veech renormalization maps $\mathcal{R} : \mathcal{H} \rightarrow \mathcal{H}$ and $R : C \times \Lambda_{\mathcal{A}} \rightarrow C \times \Lambda_{\mathcal{A}}$ are also recurrent, for the corresponding invariant measures m and ν .*

Proof. Since \hat{m} is a finite measure, by Theorem 24.1, the claim that (T^t, \hat{m}) is recurrent is a direct consequence of the Poincaré recurrence theorem. The claim for $(T^t, \hat{\nu})$ follows immediately, because $\hat{\nu} = P_*(\hat{m})$ and $T^t \circ P = P \circ T^t$: given any positive measure set $D \subset \hat{S}$, the fact that \hat{m} -almost every point of $P^{-1}(D)$ returns to $P^{-1}(D)$ under T^t implies that $\hat{\nu}$ -almost every point of D returns to D under T^t . Similarly, the statement for (\mathcal{R}, m) follows immediately from the fact that (T^t, \hat{m}) is recurrent, \mathcal{R} is the return map of T^t to the cross-section \mathcal{S} , and a subset of the cross-section as positive m -measure if and only the set of flow orbits has positive \hat{m} -measure. For the same reasons, the fact that $(T^t, \hat{\nu})$ is recurrent implies that (R, ν) is recurrent. \square

If (f, μ) is recurrent then, given any positive measure $D \subset M$ there is a first-return map $f_D : D \rightarrow D$ of f to D , defined by

$$f_D(x) = f^n(x), \quad n = \min\{k \geq 1 : f^k(x) \in D\}$$

at almost every point $x \in D$. We call f_D the map *induced* by f on D .

Lemma 25.2. *The induced map f_D preserves the restriction of μ to D .*

Proof. Suppose first that f is invertible. Then, given any measurable set $E \subset D$, the pre-image $f_D^{-1}(E)$ is the disjoint union of all $f^{-k}(E_k)$, $k \geq 1$ where E_k is the set of points $x \in E$ such that $f^{-k}(x) \in D$ but $f^{-j}(x) \notin D$ for $0 < j < k$. Since these E_k are pairwise disjoint, we get

$$\mu(f_D^{-1}(E)) = \sum_{k \geq 1} \mu(f^{-k}(E_k)) = \sum_{k \geq 1} \mu(E_k) = \mu(E).$$

To treat the general, possibly non-invertible, case, consider the natural extension $(\tilde{f}, \tilde{\mu})$ of the system (f, μ) . This is defined by

$$\tilde{f} : \tilde{M} \rightarrow \tilde{M}, \quad \tilde{f}(\dots, x_n, \dots, x_0) = (\dots, x_n, \dots, x_0, f(x_0))$$

where \tilde{M} is the space of all sequences $(x_n)_n$ on M such that $f(x_n) = x_{n-1}$ for all $n \geq 1$. Moreover, $\tilde{\mu}$ is the unique \tilde{f} -invariant measure such that $\pi_*(\tilde{\mu}) = \mu$, where $\pi : \tilde{M} \rightarrow M$ is the projection $(x_n)_n \mapsto x_0$. Clearly, $\pi \circ \tilde{f} = f \circ \pi$. Moreover,

$$\pi \circ \tilde{f}_{\tilde{D}} = f_D \circ \pi$$

where $\tilde{f}_{\tilde{D}}$ denotes the map induced by \tilde{f} on $\tilde{D} = \pi^{-1}(D)$. Then, using the previous paragraph,

$$\mu(f_D^{-1}(E)) = \tilde{\mu}(\pi^{-1}(f_D^{-1}(E))) = \tilde{\mu}(\tilde{f}_{\tilde{D}}^{-1}(\pi^{-1}(E))) = \tilde{\mu}(\pi^{-1}(E)) = \mu(E).$$

This completes the proof. □

Remark 25.3. It is clear that if f is ergodic for μ then f_D is ergodic for the restricted measure $\mu \upharpoonright D$. Indeed, given any $E \subset D$, let $F = \cup_{n=0}^{\infty} F_n$, where $F_0 = E$ and

$$F_n = \{x \in M : f^n(x) \in E \text{ but } f^k(x) \notin E \text{ for all } 0 \leq k < n\} \quad \text{for } n \geq 1.$$

If E is f_D -invariant then F is f -invariant. Suppose $\mu(E) > 0$. Then $\mu(F) > 0$ and so, by hypothesis, $\mu(F) = 1$. Consequently, $\mu(E) = \mu(F \cap D) = \mu(D)$. This shows that f_D is ergodic if f is. We are going to prove a partial converse to this fact.

We say that (f, μ) is a *Markov system* if the measure μ is f -invariant and there exists a countable partition $(M_j)_j$ of a full measure subset of M , such that each M_j is mapped bijectively to a full measure subset of M . Such systems always admit a Jacobian. Indeed, let μ_j be the measure defined on each M_j by $\mu_j(E) = \mu(f(E))$. Since μ is invariant, $\mu \leq \mu_j$ and, in particular, μ is absolutely continuous with respect to μ_j . The set where the Radon-Nikodym derivative vanishes has zero μ -measure:

$$\mu\left(\left\{x : \frac{d\mu}{d\mu_j}(x) = 0\right\}\right) = \int_{\{x: \frac{d\mu}{d\mu_j}(x)=0\}} \frac{d\mu}{d\mu_j} d\mu_j = 0.$$

Hence, $J_\mu f(x) = (d\mu/d\mu_j)^{-1}(x)$ is well-defined at μ -almost every point in each M_j , and it is a Jacobian of f relative to μ :

$$\int_E J_\mu f d\mu = \int_E \left(\frac{d\mu}{d\mu_j}\right)^{-1} d\mu = \int_E d\mu_j = \mu_j(E) = \mu(f(E))$$

for every measurable set $E \subset M_j$ and every $j \geq 1$.

Lemma 25.4. *Assume (f, μ) is a Markov system. If the map induced by f on some of the Markov domains M_j is ergodic for the restriction of μ to M_j , then (f, μ) itself is ergodic.*

Proof. Let $F \subset M$ be f -invariant. Then $E = F \cap M_j$ is f_{M_j} -invariant and so, either $\mu(E) = 0$ or $\mu(M_j \setminus E) = 0$. In the first case, the existence of a Jacobian implies that

$\mu(f(E)) = 0$. Notice that $f(E) = F$, up to a zero measure set, because $f : M_j \rightarrow M$ is essentially surjective and F is an invariant set. It follows that $\mu(F) = 0$. In the second case, a similar argument shows that $\mu(M \setminus F) = 0$. This proves that f is ergodic. \square

We are going to apply these observations to the Rauzy-Veech renormalization map R , and the R -invariant measure ν constructed in Section 21. Recall that R maps each $\{\pi\} \times \Lambda_{\pi,\varepsilon}$ bijectively to $\{\pi'\} \times \Lambda_{\mathcal{A}}$, where π' is the type ε successor of π and

$$\Lambda_{\pi,\varepsilon} = \{ \lambda \in \Lambda_{\mathcal{A}} : (\pi, \lambda) \text{ has type } \varepsilon \}.$$

For each $n \geq 1$ and $\underline{\varepsilon} = (\varepsilon_0, \dots, \varepsilon_{n-1}) \in \{0, 1\}^n$, define

$$\Lambda_{\pi,n,\underline{\varepsilon}} = \{ \lambda \in \Lambda_{\mathcal{A}} : R^k(\pi, \lambda) \text{ has type } \varepsilon_k \text{ for } k = 0, 1, \dots, n-1 \}. \quad (88)$$

Then R^n maps every $\Lambda_{\pi,n,\underline{\varepsilon}}$ bijectively to $\pi^n \times \Lambda_{\mathcal{A}}$. As a consequence of (11), $\Lambda_{\pi,n,\underline{\varepsilon}}$ is the image of $\Lambda_{\mathcal{A}}$ under the projectivization of Θ^{n*} , where $\Theta^{n*} = \Theta_{\pi,\lambda}^{n*}$ for any $(\pi, \lambda) \in \Lambda_{\pi,n,\underline{\varepsilon}}$. By Corollary 5.3, one may find $N \geq 1$ and $\underline{\varepsilon} = (\varepsilon_0, \dots, \varepsilon_{N-1})$ such that $\Lambda_* = \Lambda_{\pi,N,\underline{\varepsilon}}$ is relatively compact in Λ . Let N and Λ_* be fixed from now on, and denote by $R_* : \Lambda_* \rightarrow \Lambda_*$ the map induced by R^N on Λ_* . For x in a full measure subset of Λ_* , let $k \geq 1$ be the smallest positive integer such that $R^{kN}(x) \in \Lambda_*$. Then the set $\Lambda_{\pi,(k+1)N,\underline{\theta}}$ that contains x satisfies

$$R^{kN}(\Lambda_{\pi,(k+1)N,\underline{\theta}}) = \Lambda_*.$$

In particular, $R_* = R^{kN}$ on the set $\Lambda_{\pi,(k+1)N,\underline{\theta}}$. This proves that $(R_*, (\nu | \Lambda_*))$ is a Markov system.

Proposition 25.5. *The Markov system $(R_*, (\nu | \Lambda_*))$ is ergodic.*

The proof of this proposition appears in section 27. It uses the notion of projective metric, that we recall in section 26. This notion will be useful again later. Also in section 27, we deduce from the proposition that the renormalization maps R and \mathcal{R} , and the Teichmüller flow \mathcal{T}^t are ergodic, relative to their invariant measures ν , $\hat{\nu}$, and \hat{m} .

26. Projective metrics

Birkhoff [5] introduced the notion of projective metric associated to a general convex cone C in any vector space. Here we only need the case $C = \mathbb{R}_+^A$.

Given any $u, v \in C$, define

$$a(u, v) = \inf \left\{ \frac{v_\alpha}{u_\alpha} : \alpha \in \mathcal{A} \right\} \quad \text{and} \quad b(u, v) = \sup \left\{ \frac{v_\beta}{u_\beta} : \beta \in \mathcal{A} \right\}. \quad (89)$$

Notice that

$$v - tu \in C \Leftrightarrow t < a(u, v) \quad \text{and} \quad su - v \in C \Leftrightarrow s > b(u, v). \tag{90}$$

We call *projective metric* associated to $C = \mathbb{R}_+^A$ the function $d_p(\cdot, \cdot)$ defined by

$$d_p(u, v) = \log \frac{b(u, v)}{a(u, v)} = \log \sup \left\{ \frac{u_\alpha v_\beta}{v_\alpha u_\beta} : \alpha, \beta \in \mathcal{A} \right\} \tag{91}$$

for each $u, v \in C$. This terminology is justified by the next lemma, which says that $d_p(\cdot, \cdot)$ induces a distance in the projective quotient of C . The lemma is an easy consequence of the definition (91).

Lemma 26.1. *For all $u, v, w \in C$,*

- (i) $d_p(u, v) = d_p(v, u)$
- (ii) $d_p(u, v) + d_p(v, w) \geq d_p(u, w)$
- (iii) $d_p(u, v) \geq 0$
- (iv) $d_p(u, v) = 0$ if and only if there exists $t > 0$ such that $u = tv$.

Let $G : \mathbb{R}^A \rightarrow \mathbb{R}^A$ be a linear operator such that $G(C) \subset C$ or, equivalently, such that all the entries $G_{\alpha, \beta}$ of the matrix of G are non-negative. Then

$$t < a(u, v) \Leftrightarrow v - tu \in C \Rightarrow G(v) - tG(u) \in C \Leftrightarrow t < a(G(u), G(v)).$$

This means that $a(u, v) \leq a(G(u), G(v))$ and a similar argument proves that $b(u, v) \geq b(G(u), G(v))$. Therefore,

$$d_p(G(u), G(v)) \leq d_p(u, v) \quad \text{for all } u, v \in C. \tag{92}$$

It follows from Lemma 26.1 that, restricted to the simplex $\Lambda_{\mathcal{A}}$, the function d_p is a genuine metric. We call $g : \Lambda_{\mathcal{A}} \rightarrow \Lambda_{\mathcal{A}}$ a *projective map* if there exists a linear isomorphism $G : \mathbb{R}^A \rightarrow \mathbb{R}^A$ such that $G(\mathbb{R}_+^A) \subset \mathbb{R}_+^A$ and

$$g(\lambda) = \frac{G(\lambda)}{\sum_{\alpha \in \mathcal{A}} G(\lambda)_\alpha} = \frac{G(\lambda)}{\sum_{\alpha, \beta \in \mathcal{A}} G_{\alpha, \beta} \lambda_\beta}. \tag{93}$$

We say g is the projectivization of G . The relation (92) means that projective maps never expand the projective metric on the simplex.

A set $K \subset \Lambda_{\mathcal{A}}$ is relatively compact in $\Lambda_{\mathcal{A}}$ if and only if the coordinates of its points are all larger than some positive constant. So, it follows directly from the definition (91) that if K is relatively compact in $\Lambda_{\mathcal{A}}$ then it has finite diameter relative to the projective metric

$$\sup_{x, y \in K} d_p(x, y) < \infty.$$

We shall see in Proposition 26.3 that if the entries of G are strictly positive or, equivalently, if the image of g is relatively compact in $\Lambda_{\mathcal{A}}$, then the inequality in (92) is strict. Thus, in that case the maps G and g are uniform contractions relative to the projective metrics in $\mathbb{R}_+^{\mathcal{A}}$ and $\Lambda_{\mathcal{A}}$, respectively.

Lemma 26.2. *Let $g : \Lambda_{\mathcal{A}} \rightarrow \Lambda_{\mathcal{A}}$ be a projective map and Dg be its derivative. Then $\log|\det Dg|$ is $(d + 1)$ -Lipschitz continuous for the projective distance.*

Proof. We use the following observation: if a functional $h(\lambda) = \sum_{\beta} h_{\beta} \lambda_{\beta}$ has non-negative coefficients, $h_{\beta} \geq 0$, then $\log h(\lambda)$ is 1-Lipschitz relative to the projective distance. Indeed,

$$\log h(\sigma) - \log h(\lambda) = \log \frac{\sum_{\beta} h_{\beta} \sigma_{\beta}}{\sum_{\beta} h_{\beta} \lambda_{\beta}} \leq \log \sup_{\beta} \frac{\sigma_{\beta}}{\lambda_{\beta}} = \log b(\lambda, \sigma).$$

Recall the definition (89). Since $\sum_{\beta} \lambda_{\beta} = 1 = \sum_{\beta} \sigma_{\beta}$, we also have $a(\lambda, \sigma) \leq 1$. It follows that $\log b(\lambda, \sigma) \leq d_p(\lambda, \sigma)$. This justifies our observation.

Now let g be the projectivization of some linear isomorphism G . We begin by expressing Dg in terms of G . Let $\dot{\Lambda}_{\mathcal{A}}$ represent the hyperplane tangent to the simplex $\Lambda_{\mathcal{A}}$. From (93) we find

$$Dg(\lambda)\dot{\lambda} = \frac{G(\dot{\lambda})}{s(\lambda)} - \frac{G(\lambda)}{s(\lambda)} \cdot \frac{\sum_{\alpha} G(\dot{\lambda})_{\alpha}}{s(\lambda)}, \quad s(\lambda) = \sum_{\alpha, \beta} G_{\alpha, \beta} \lambda_{\beta}.$$

This may be rewritten as $Dg(\lambda) = P_{\lambda} \circ s(\lambda)^{-1} \circ G$, where $G : \dot{\Lambda}_{\mathcal{A}} \rightarrow G(\dot{\Lambda}_{\mathcal{A}})$, we use $s(\lambda)^{-1}$ to mean division by the scalar $s(\lambda)$ on the vector hyperplane $G(\dot{\Lambda}_{\mathcal{A}})$, and $P_{\lambda} : G(\dot{\Lambda}_{\mathcal{A}}) \rightarrow \dot{\Lambda}_{\mathcal{A}}$ is the projection along the direction of $G(\lambda)$. Consequently,

$$\log \det Dg(\lambda) = \log \det P_{\lambda} - (d - 1) \log s(\lambda) + \log \det G$$

We are going to show that each of the three terms on the right hand side is Lipschitz relative to the projective metric. Indeed, $\log \det G$ is constant. By the observation in the first paragraph, $\log s(\lambda)$ is 1-Lipschitz. Finally,

$$\log \det P_{\lambda} = \log(n_0 \cdot G(\lambda)) - \log(n_1 \cdot G(\lambda))$$

where n_0 and n_1 are unit vectors orthogonal to the hyperplanes $\dot{\Lambda}_{\mathcal{A}}$ and $G(\dot{\Lambda}_{\mathcal{A}})$, respectively. Both n_i have non-negative coefficients: on the one hand, n_0 is collinear to $(1, \dots, 1)$; on the other, n_1 is collinear to $G^*(1, \dots, 1)$, and the adjoint operator G^* has non-negative coefficients since G does. Using the observation in the first paragraph once more, it follows that each $\log(n_i \cdot G(\lambda))$ is a 1-Lipschitz function. Altogether, $\log \det Dg(\lambda)$ is $(d + 1)$ -Lipschitz relative to the projective metric, as claimed. \square

For proving Proposition 25.5, this is all we need to know about projective metrics. In the remainder of the present section we prove a few other properties that will be useful at latter occasions.

Proposition 26.3. *For any $\delta > 0$ there is $\theta < 1$ such that, if $\delta < G_{\alpha,\beta} < \delta^{-1}$ for all $\alpha, \beta \in \mathcal{A}$, then*

$$d_p(G(u), G(v)) \leq \theta d_p(u, v) \quad \text{for all } u, v \in C.$$

Proof. Given any $z, w \in C$, one can always find some $c = c(z, w) > 0$ such that

$$a(G(z), G(w)) > c\delta^2 \quad \text{and} \quad b(G(z), G(w)) < c\delta^{-2}. \tag{94}$$

Indeed, the hypothesis $\delta \leq G_{\alpha,\beta} \leq \delta^{-1}$ implies that

$$a(G(z), G(w)) = \inf_{\alpha} \frac{\sum_{\beta} G_{\alpha,\beta} w_{\beta}}{\sum_{\beta} G_{\alpha,\beta} z_{\beta}} > \delta^2 \frac{\sum_{\beta} w_{\beta}}{\sum_{\beta} z_{\beta}}.$$

Just take $c(z, w)$ to be the last factor on the right hand side, and observe that the same kind of argument also gives $b(G(z), G(w)) < \delta^{-2}c(z, w)$.

Now let $u, v \in C$ and, for each $n \geq 1$, consider arbitrary $0 < t_n < a(u, v)$ and $b(u, v) < s_n < \infty$. In other words,

$$v - t_n u \in C \quad \text{and} \quad s_n u - v \in C.$$

Taking $z = v - t_n u$ and $w = s_n u - v$ in (94), we find $c_n > 0$ such that

$$G(s_n u - v) - c_n \delta^2 G(v - t_n u) \in C \quad \text{and} \quad c_n \delta^{-2} G(v - t_n u) - G(s_n u - v) \in C.$$

Write $T_n = c_n \delta^2$ and $S_n = c_n \delta^{-2}$. The previous relations may be rewritten as

$$(s_n + t_n T_n)G(u) - (1 + T_n)G(v) \in C \quad \text{and} \quad (1 + S_n)G(v) - (s_n + t_n S_n)G(u) \in C,$$

and, by (90), this is the same as

$$b(G(u), G(v)) < \frac{s_n + t_n T_n}{1 + T_n} \quad \text{and} \quad a(G(u), G(v)) > \frac{s_n + t_n S_n}{1 + S_n}.$$

Combining these two inequalities we see that $d_p(G(u), G(v))$ can not exceed

$$\log\left(\frac{s_n + t_n T_n}{1 + T_n} \cdot \frac{1 + S_n}{s_n + t_n S_n}\right) = \log\left(\frac{s_n/t_n + T_n}{1 + T_n} \cdot \frac{1 + S_n}{s_n/t_n + S_n}\right).$$

The last term can be rewritten as

$$\begin{aligned} \log\left(\frac{s_n}{t_n} + T_n\right) - \log(1 + T_n) - \log\left(\frac{s_n}{t_n} + S_n\right) + \log(1 + S_n) &= \\ = \int_0^{\log(s_n/t_n)} \left(\frac{e^x dx}{e^x + T_n} - \frac{e^x dx}{e^x + S_n}\right), \end{aligned}$$

and this is not larger than

$$\sup_{x>0} \frac{e^x(S_n - T_n)}{(e^x + T_n)(e^x + S_n)} \log\left(\frac{s_n}{t_n}\right).$$

Now we use the following elementary fact:

$$\sup_{y>0} \frac{y(S_n - T_n)}{(y + T_n)(y + S_n)} = \frac{1 - (T_n/S_n)^{1/2}}{1 + (T_n/S_n)^{1/2}}.$$

(The supremum is attained at $y = (S_n T_n)^{1/2}$.) Noting that $T_n/S_n = \delta^4$, we conclude that

$$d_p(G(u), G(v)) \leq \frac{1 - \delta^2}{1 + \delta^2} \log\left(\frac{s_n}{t_n}\right).$$

Making $s_n \rightarrow a(u, v)$ and $t_n \rightarrow b(u, v)$, the last factor converges to $d_p(u, v)$, and we obtain the conclusion of the proposition with $\theta = (1 - \delta^2)/(1 + \delta^2)$. \square

Thus, if $g : \Lambda_{\mathcal{A}} \rightarrow \Lambda_{\mathcal{A}}$ is a projective map such that $g(\Lambda_{\mathcal{A}})$ is relatively compact in $\Lambda_{\mathcal{A}}$ or, in other words, such that it is the projectivization of a linear isomorphism G with strictly positive coefficients, then g is a uniform contraction relative to the projective metric. We also note that this metric is complete:

Proposition 26.4. *Any d_p -Cauchy sequence $(\lambda^n)_n$ is d_p -convergent. Moreover, the normalization $(\lambda^n/|\lambda^n|)_n$ is norm-convergent.*

Proof. Let $(\lambda^n)_n$ be a d_p -Cauchy sequence in C : given any $\varepsilon > 0$, there exists $N \geq 1$ such that $d_p(\lambda^m, \lambda^n) \leq \varepsilon$ for all $m, n \geq N$. Up to dropping a finite number of terms, we may suppose that $d_p(\lambda^m, \lambda^n) \leq 1$ for all $m, n \geq 1$. Then,

$$\frac{1}{e} \leq \frac{\lambda_\alpha^m \lambda_\beta^n}{\lambda_\alpha^n \lambda_\beta^m} \leq e \quad \text{for all } \alpha, \beta \in \mathcal{A} \text{ and } m, n \geq 1.$$

As a consequence, writing $R = e \sup\{\lambda_\alpha^1/\lambda_\beta^1 : \alpha, \beta \in \mathcal{A}\}$ we get

$$\frac{1}{R} \leq \frac{\lambda_\alpha^n}{\lambda_\beta^n} \leq R \quad \text{for all } \alpha, \beta \in \mathcal{A} \text{ and } n \geq 1. \tag{95}$$

It is no restriction to suppose that $|\lambda^n| = 1$ for all $n \geq 1$. Then

$$\inf_{\alpha \in \mathcal{A}} \lambda_\alpha^n \leq 1 \leq \sup_{\beta \in \mathcal{A}} \lambda_\beta^n \quad \text{and} \quad \inf_{\alpha \in \mathcal{A}} \frac{\lambda_\alpha^n}{\lambda_\alpha^m} \leq 1 \leq \sup_{\beta \in \mathcal{A}} \frac{\lambda_\beta^n}{\lambda_\beta^m} \tag{96}$$

for all $m, n \geq 1$. The first part of (96) together with (95) imply

$$\frac{1}{R} \leq \lambda_\alpha^n \leq R \quad \text{for all } \alpha \in \mathcal{A} \text{ and } n \geq 1. \tag{97}$$

The second part of (96) together with $d_p(\lambda^m, \lambda^n) \leq \varepsilon$ give

$$e^{-\varepsilon} \leq \inf_{\alpha \in \mathcal{A}} \frac{\lambda_\alpha^n}{\lambda_\alpha^m} \leq 1 \leq \sup_{\beta \in \mathcal{A}} \frac{\lambda_\beta^n}{\lambda_\beta^m} \leq e^\varepsilon \tag{98}$$




Figure 34

for all $m, n \geq N$. It follows that

$$\sup_{\alpha \in \mathcal{A}} |\lambda_\alpha^m - \lambda_\alpha^n| \leq \sup_{\alpha \in \mathcal{A}} \lambda_\alpha^m \cdot \sup_{\alpha \in \mathcal{A}} \left| \frac{\lambda_\alpha^n}{\lambda_\alpha^m} - 1 \right| \leq R(e^\varepsilon - 1).$$

This shows that $(\lambda^n)_n$ is a Cauchy sequence with respect to the usual norm in $\mathbb{R}^{\mathcal{A}}$. It follows that the sequence converges to some $\lambda \in \mathbb{R}^{\mathcal{A}}$. Passing to the limit in (97) we find that $R^{-1} \leq \lambda_\alpha \leq R$ for all $\alpha \in \mathcal{A}$ and, in particular, $\lambda \in C$. Passing to the limit in (98), we get

$$e^{-\varepsilon} \leq \inf_{\alpha \in \mathcal{A}} \frac{\lambda_\alpha^n}{\lambda_\alpha} \leq 1 \leq \sup_{\beta \in \mathcal{A}} \frac{\lambda_\beta^n}{\lambda_\beta} \leq e^\varepsilon$$

for all $n \geq N$. This means that $a(\lambda, \lambda^n) \geq e^{-\varepsilon}$ and $b(\lambda, \lambda^n) \leq e^\varepsilon$, and so $d_p(\lambda, \lambda^n) \leq 2\varepsilon$ for all $n \geq N$. Therefore, $(\lambda^n)_n$ is d_p -convergent to λ . \square

27. Ergodicity

Applying the conclusions in the first half of the previous section to the inverse branches of the map $R_* : \Lambda_* \rightarrow \Lambda_*$ introduced in section 25, we can give the

Proof of Proposition 25.5. The domain Λ_* has finite diameter $D_* > 0$ for the projective metric d_p , because it is relatively compact in $\Lambda_{\mathcal{A}}$. By Lemma 26.2, $\log|\det R_*^{-k}|$ is $(d + 1)$ -Lipschitz continuous relative to d_p , for every inverse branch $R_*^{-k} : \Lambda_* \rightarrow \Lambda_{\pi, (k+1)N, \underline{\theta}}$ of any iterate R_*^k of the map R_* . Consequently,

$$\log \frac{|\det R_*^{-k}|(x)}{|\det R_*^{-k}|(y)} \leq (d + 1)D_* \tag{99}$$

for any $x, y \in \Lambda_*$ and every inverse branch. Now let $E \subset \Lambda_{\mathcal{A}}$ be any R_* -invariant set with $\nu(E) > 0$. Then E has positive Lebesgue measure as well. Then, for any $\delta > 0$ there exists $k \geq 1$ such that

$$d_1\lambda(\Lambda_{\pi, (k+1)N, \underline{\theta}} \setminus E) < \delta d_1\lambda(\Lambda_{\pi, (k+1)N, \underline{\theta}}).$$

Taking the images under R_*^k and using (99), we find that

$$d_1\lambda(\Lambda_{\mathcal{A}} \setminus E) < \delta e^{(d+1)D_*} d_1\lambda(\Lambda_{\mathcal{A}}).$$

Since δ is arbitrary, we conclude that E has full Lebesgue measure in $\Lambda_{\mathcal{A}}$. It follows that it also has full ν measure in $\Lambda_{\mathcal{A}}$. This proves ergodicity. \square

Remark 27.1. Each inverse branch $R_*^{-k} : \Lambda_* \rightarrow \Lambda_{\pi, (k+1)N, \underline{\theta}}$ is the projectivization of a linear map Θ^{kN*} , and so it extends to a bijection from the whole simplex $\Lambda_{\mathcal{A}}$ to the set $\Lambda_{\pi, kN, \underline{\tau}}$ that contains $\Lambda_{\pi, (k+1)N, \underline{\theta}}$. Notice that $\Lambda_{\pi, kN, \underline{\tau}}$ is contained in Λ_* , which is relatively compact in $\Lambda_{\mathcal{A}}$. Using Proposition 26.3, we get that all these inverse branches contract the projective metric, with contraction rate uniformly bounded from 1. Thus, $R_* : \Lambda_* \rightarrow \Lambda_*$ is a uniformly expanding map. Although we do not use this fact, it could be combined with Lemma 26.2 to give an alternative proof that R_* and R admit invariant measures absolutely continuous with respect to Lebesgue measure.

Corollary 27.2. *The Rauzy-Veech renormalization map R is ergodic relative to the invariant measure ν . Moreover, every R -invariant measure absolutely continuous with respect to Lebesgue measure coincides with a multiple of ν .*

Proof. We have seen in Proposition 25.5 that the map R_* induced by R^N is ergodic relative to the restriction of ν to Λ_* . Using Lemma 25.4, we conclude that (R^N, ν) is ergodic. This implies that (R, ν) is ergodic. The uniqueness statement is a consequence of ergodicity and the fact that ν is actually equivalent to Lebesgue measure. \square

Together with Proposition 21.4, Corollary 27.2 completes the proof of Theorem 7.2. From the previous arguments we also get

Corollary 27.3. *The invertible Rauzy-renormalization map \mathcal{R} is ergodic, for the invariant measure m , and the Teichmüller flow \mathcal{T} is ergodic, for the invariant measure \hat{m} , restricted to the subset $\{(\pi, \lambda, \tau) : \text{area}(\lambda, \tau) = 1\}$.*

28. Space of invariant measures

Having finished the proof of Theorem 7.2, we are now going to use it to give a positive solution to Conjecture 4.6. This is done in the next theorem, which is due to Masur [17] and Veech [22]. The proof occupies sections 28 and 29.

Theorem 28.1. *Almost every interval exchange transformation is uniquely ergodic.*

Let $f : I \rightarrow I$ be an interval exchange transformation, defined by data (π, λ) . Throughout, we assume that (π, λ) satisfies the Keane condition (13). Then f is minimal and the renormalization $R^n(\pi, \lambda)$ is defined for all $n \geq 1$.

Let $\mathcal{M}(\pi, \lambda)$ denote the cone of f -invariant (positive) measures. Since f is minimal, every $\mu \in \mathcal{M}(\pi, f)$ is non-atomic and positive on open intervals. Define

$$\phi_\mu : I \rightarrow [0, \mu(I)), \quad \phi_\mu(x) = \mu([0, x]).$$

Then ϕ is continuous and injective, and so it is a homeomorphism. Define $\lambda(\mu) = (\lambda_\alpha(\mu))_{\alpha \in \mathcal{A}}$ by

$$\lambda_\alpha(\mu) = \mu(I_\alpha) \quad \text{for all } \alpha \in \mathcal{A},$$

where $(I_\alpha)_{\alpha \in \mathcal{A}}$ is the partition of I defined by (π, λ) . Notice that

$$|\lambda(\mu)| = \sum_{\alpha \in \mathcal{A}} \lambda_\alpha(\mu) = \sum_{\alpha \in \mathcal{A}} \mu(I_\alpha) = \mu(I).$$

Now define $f_\mu : [0, \mu(I)) \rightarrow [0, \mu(I))$ by $f_\mu = \phi_\mu \circ f \circ \phi_\mu^{-1}$.

Lemma 28.2. f_μ is the interval exchange transformation defined by $(\pi, \lambda(\mu))$.

Proof. For every $\alpha \in \mathcal{A}$, define

$$I_\alpha(\mu) = \left[\sum_{\pi_0(\beta) < \pi_0(\alpha)} \lambda_\beta(\mu), \sum_{\pi_0(\beta) \leq \pi_0(\alpha)} \lambda_\beta(\mu) \right).$$

If $s \in I_\alpha(\mu)$ then $\phi_\mu^{-1}(s) \in I_\alpha$ and, by the definition of ϕ_μ ,

$$\mu[\partial I_\alpha, \phi_\mu^{-1}(s)) = s - \partial I_\alpha(\mu) = s - \sum_{\pi_0(\beta) < \pi_0(\alpha)} \lambda_\beta(\mu).$$

Then $f(\phi_\mu^{-1}(s)) \in f(I_\alpha)$ and, since μ is f -invariant,

$$\mu[\partial f(I_\alpha), f(\phi_\mu^{-1}(s))] = s - \sum_{\pi_0(\beta) < \pi_0(\alpha)} \lambda_\beta(\mu).$$

Observe that $\mu[0, \partial f(I_\alpha)) = \sum_{\pi_1(\beta) < \pi_1(\alpha)} \mu(I_\beta) = \sum_{\pi_1(\beta) < \pi_1(\alpha)} \lambda_\beta(\mu)$. It follows that

$$f_\mu(s) = \phi_\mu(f(I_\alpha), f(\phi_\mu^{-1}(s))) = s - \sum_{\pi_0(\beta) < \pi_0(\alpha)} \lambda_\beta(\mu) + \sum_{\pi_1(\beta) < \pi_1(\alpha)} \lambda_\beta(\mu)$$

for all $s \in I_\alpha(\mu)$. This proves that f_μ is an interval exchange transformation, with translation vector $w(\mu) = (w_\alpha(\mu))_{\alpha \in \mathcal{A}}$ given by

$$w_\alpha(\mu) = \sum_{\pi_1(\beta) < \pi_1(\alpha)} \lambda_\beta(\mu) - \sum_{\pi_0(\beta) < \pi_0(\alpha)} \lambda_\beta(\mu),$$

and so the claim follows. □

Proposition 28.3. *The map $\mu \mapsto \lambda(\mu)$ is a linear isomorphism from $\mathcal{M}(\pi, \lambda)$ to the cone*

$$\bigcap_{n=1}^{\infty} \Theta_{\pi, \lambda}^{n*}(\mathbb{R}_+^A).$$

The proof of this proposition occupies the remainder of this section. The first step is

Lemma 28.4. *The map $\mu \mapsto \lambda(\mu)$ is linear and injective.*

Proof. Linearity is clear from the definition. To prove injectivity, observe that $\phi_\mu(0) = 0$ and so

$$f_\mu^n(0) = \phi_\mu(f^n(0)) = \mu[0, f^n(0)].$$

This relation shows that, for a dense subset of values of x , the value of $\mu([0, x])$ is determined by f_μ and, hence (Lemma 28.2), by $(\pi, \lambda(\mu))$. As the measure μ has no atoms, it follows that it is completely determined by $\lambda(\mu)$, which proves the claim. \square

Let us denote by $\mathcal{C}(\pi, \lambda)$ the image of $\mathcal{M}(\pi, \lambda)$ under the map $\mu \mapsto \lambda(\mu)$. Now, to prove Proposition 28.3 we only have to show that

$$\mathcal{C}(\pi, \lambda) = \bigcap_{n=1}^{\infty} \Theta_{\pi, \lambda}^{n*}(\mathbb{R}_+^A). \tag{100}$$

Lemma 28.5. *An interval exchange transformation $g : J \rightarrow J$ is topologically conjugate to f if and only if there exists $\mu \in \mathcal{M}(\pi, \lambda)$ such that $g = f_\mu$.*

Proof. The only if part is obvious: If $g = f_\mu$ then, by definition, it is conjugate to f by ϕ_μ . Conversely, suppose $g = \phi \circ f \circ \phi^{-1}$ for some homeomorphism $\phi : I \rightarrow J$. Let m be Lebesgue measure on J and $\mu = \phi_*^{-1}m$. Since m is invariant under g , the measure μ is invariant under f . Moreover,

$$\mu([0, x]) = m(\phi([0, x])) = m([0, \phi(x)]) = \phi(x)$$

for every $x \in I$. In other words, $\phi_\mu = \phi$ and so $g = f_\mu$. \square

Remark 28.6. Suppose g is defined by data $(\tilde{\pi}, \tilde{\lambda})$. The previous lemma means that it is also defined by $(\pi, \lambda(\mu))$. In general, the two pairs of data need not coincide. For instance, we have seen in Example 1.3 that

$$\pi = \begin{pmatrix} A & B & C \\ B & C & A \end{pmatrix} \quad \text{and} \quad \lambda = (\lambda_A, \lambda_B, \lambda_C)$$

define the same transformation as

$$\tilde{\pi} = \begin{pmatrix} A & D \\ D & A \end{pmatrix} \quad \text{and} \quad \tilde{\lambda} = (\lambda_A, \lambda_B + \lambda_C).$$

Another mechanism for non-uniqueness is that the linear map Ω_π is usually not injective, and the transformation depends only on the translation vector $w = \Omega_\pi(\lambda)$.

Lemma 28.7. *Let $(\pi', \lambda') = \hat{R}(\lambda, \pi)$ and I' be the domain of definition of the interval exchange transformation $f' = \hat{R}(f)$ defined by (π', λ') . The map $\rho : \mu \mapsto \mu' = \mu \upharpoonright I'$ is a linear isomorphism from $\mathcal{M}(\pi, \lambda)$ to $\mathcal{M}(\pi', \lambda')$.*

Proof. It is clear that the map ρ is linear. We start by checking that it takes values in $\mathcal{M}(\pi', \lambda')$, that is, that $\mu' = \mu \upharpoonright I'$ is an f' -invariant measure if μ is f -invariant. Indeed, we may write any measurable set $E \subset I'$ as a disjoint union $E_1 \cup E_2$, where E_1 is the intersection of E with $f^{-1}(I')$, and $E_2 = E \setminus E_1$. Then $f'(E) = f(E_1) \cup f^2(E_2)$, and this union is also disjoint. Consequently,

$$\mu'(f'(E)) = \mu(f(E_1)) + \mu(f^2(E_2)) = \mu(E_1) + \mu(E_2) = \mu'(E)$$

Observe also that if E is a measurable subset of $I \setminus I'$ then $f(E) \subset I'$ and then

$$\mu(E) = \mu(f(E)) = \mu'(f(E)). \tag{101}$$

This implies that μ' determines μ uniquely, and so the map ρ is injective. Finally, given any $\mu' \in \mathcal{M}(\pi', \lambda')$, we may use (101) to extend it to a measure μ on the whole I , and this measure is f -invariant. Thus, ρ is also surjective. \square

Corollary 28.8. *If $(\pi', \lambda') = \hat{R}(\pi, \lambda)$ then*

$$\mathcal{C}(\pi, \lambda) = \Theta_{\pi, \lambda}^*(\mathcal{C}(\pi', \lambda')).$$

Proof. Let (π, λ) have type $\varepsilon \in \{0, 1\}$. Recall that $I_\alpha = I'_\alpha$ for $\alpha \neq \alpha(\varepsilon)$, and $I_{\alpha(0)} = I'_{\alpha(0)} \cup f(I'_{\alpha(1)})$ when $\varepsilon = 0$, and $I_{\alpha(1)} = I'_{\alpha(1)} \cup f^{-1}(I'_{\alpha(0)})$ when $\varepsilon = 1$. Let $\mu \in \mathcal{M}(\pi, \lambda)$ and $\mu' \in \mathcal{M}(\pi', \lambda')$ be as in Lemma 28.7. Then, in both cases,

$$\mu(I_\alpha) = \mu'(I'_\alpha) \text{ for all } \alpha \neq \alpha(\varepsilon) \quad \text{and} \quad \mu(I_{\alpha(\varepsilon)}) = \mu'(I'_{\alpha(\varepsilon)}) + \mu'(I'_{\alpha(1-\varepsilon)}).$$

Equivalently,

$$\lambda_\alpha(\mu) = \lambda_\alpha(\mu') \text{ for all } \alpha \neq \alpha(\varepsilon) \quad \text{and} \quad \lambda_{\alpha(\varepsilon)}(\mu) = \lambda_{\alpha(\varepsilon)}(\mu') + \lambda_{\alpha(1-\varepsilon)}(\mu').$$

In other words, $\lambda(\mu) = \Theta_{\pi, \lambda}^*(\lambda(\mu'))$. This proves the statement. \square

Denote $(\pi^n, \lambda^n) = \hat{R}^n(\pi, \lambda)$ for each $n \geq 1$. Since every $\mathcal{C}(\pi^n, \lambda^n)$ is a subset of \mathbb{R}_+^A , Corollary 28.8 implies that

$$\mathcal{C}(\pi, \lambda) = \Theta_{\pi, \lambda}^{n*}(\mathcal{C}(\pi^n, \lambda^n)) \subset \Theta_{\pi, \lambda}^{n*}(\mathbb{R}_+^A) \quad \text{for all } n \geq 1.$$

This proves the direct inclusion in (100). The main point in proving the converse is the lemma that we state next.

Remark 28.9. Notice that $\lambda \in \mathcal{C}(\pi, \lambda)$, since it is the image of the Lebesgue measure under the map $\mu \mapsto \lambda(\mu)$. Thus, we always have $\lambda \in \Theta_{\pi, \lambda}^{n*}(\mathbb{R}_+^A)$ for all $n \geq 1$.

Lemma 28.10. *Every $\tilde{\lambda} \in \bigcap_{n=1}^{\infty} \Theta_{\pi,\lambda}^{n*}(\mathbb{R}_+^A)$ satisfies the Keane condition.*

Proof. Consider the line segment $[0, 1] \ni s \mapsto \lambda_s$ connecting $\lambda_0 = \lambda$ to $\lambda_1 = \tilde{\lambda}$ in \mathbb{R}^A . Since every $\Theta_{\pi,\lambda}^{n*}(\mathbb{R}_+^A)$ is a convex set, the whole segment is contained in the intersection:

$$\lambda_s \in \Theta_{\pi,\lambda}^{n*}(\mathbb{R}_+^A) \quad \text{for all } n \geq 1 \text{ and } s \in [0, 1].$$

For $n = 1$ this gives that $\lambda'_s = \Theta_{\pi,\lambda}^{-1*}(\lambda_s) \in \mathbb{R}_+^A$ for every $s \in [0, 1]$. By connectedness, this implies that (π, λ_s) has the same type as (π, λ) for every $s \in [0, 1]$. Consequently, $\Theta_{\pi,\lambda_s}^{-1*} = \Theta_{\pi,\lambda}^{-1*}$ and $\hat{R}(\pi, \lambda_s) = (\pi', \lambda'_s)$ for every $s \in [0, 1]$. Arguing by induction, we get that $\hat{R}^n(\pi, \lambda_s)$ has the same type as $(\pi^n, \lambda^n) = \hat{R}^n(\pi, \lambda)$ and

$$\hat{R}^n(\pi, \lambda_s) = (\pi^n, \lambda_s^n) \quad \text{with} \quad \lambda_s^n = \Theta_{\pi,\lambda_s}^{-n*}(\lambda_s) = \Theta_{\pi,\lambda}^{-n*}(\lambda_s).$$

In particular, $\hat{R}^n(\pi, \lambda_s)$ is defined for every $n \geq 1$. By Corollary 5.4, it follows that λ_s satisfies the Keane condition, for all $s \in [0, 1]$. \square

Lemma 28.11. *Let \tilde{f} be the interval exchange transformation defined by $(\pi, \tilde{\lambda})$, for some $\tilde{\lambda} \in \bigcap_{n=1}^{\infty} \Theta_{\pi,\lambda}^{n*}(\mathbb{R}_+^A)$. Then \tilde{f} is topologically conjugate to f .*

Proof. Consider the line segment $[0, 1] \ni s \mapsto \lambda_s$ connecting $\lambda_0 = \lambda$ to $\lambda_1 = \tilde{\lambda}$, and let f_s be the interval exchange transformation defined by each (π, λ_s) . By Proposition 4.1, the orbit of 0 under each transformation f_s is dense in the corresponding domain I_s . We claim that, for any $r, s \in [0, 1]$ and $m, n \geq 1$

$$f_r^m(0) < f_r^n(0) \quad \Leftrightarrow \quad f_s^m(0) < f_s^n(0). \tag{102}$$

Indeed, suppose there were m, n, r, s such that $f_r^m(0) < f_r^n(0)$ and yet $f_s^m(0) \geq f_s^n(0)$. Since the iterates vary continuously with the parameter, we may always assume that $f_s^m(0) = f_s^n(0)$. Then the f_s -orbit of zero would be finite. However, by Lemma 28.10 and Proposition 4.1, the map f_s is minimal, and so every orbit must be infinite. This contradiction proves (102).

Now define $\phi(f^n(0)) = \tilde{f}^n(0)$ for each $n \geq 1$. By (102), this map ϕ is monotone increasing. Since both orbits of 0, for f and \tilde{f} , are dense, it extends continuously to a homeomorphism $\phi : I \rightarrow \tilde{I}$. Notice that $\phi(f(x)) = \tilde{f}(\phi(x))$ for every x in the f -orbit of zero and, consequently, for every $x \in I$. This shows that ϕ is a topological conjugacy between f and \tilde{f} . \square

Lemma 28.12. *The conjugacy $\phi : I \rightarrow \tilde{I}$ in Lemma 28.11 satisfies $\phi(I_\alpha) = \tilde{I}_\alpha$ for all $\alpha \in \mathcal{A}$.*

Proof. Let $(I_{s,\alpha})_{\alpha \in \mathcal{A}}$ be the partition associated to (π, λ_s) , where $[0, 1] \ni s \mapsto \lambda_s$ is the line segment connecting $\lambda_0 = \lambda$ to $\lambda_1 = \tilde{\lambda}$. We also write $I_\alpha = I_{0,\alpha}$ and $\tilde{I}_\alpha = I_{1,\alpha}$. We begin by noting that, for each $n \geq 1$,

$$f^n(0) \in I_\alpha \quad \Leftrightarrow \quad \tilde{f}^n(0) \in \tilde{I}_\alpha. \tag{103}$$

Indeed, otherwise there would exist $s \in [0, 1]$ and $\beta \in \mathcal{A}$ with $\pi_0(\beta) > 1$ such that $f_s^n(0) = \partial I_\beta(s)$. That would imply that (π, λ_s) does not satisfy the Keane condition, contradicting Lemma 28.10.

By minimality, $\{f^n(0) : n \in \mathbb{N}\} \cap I_\alpha$ is a dense subset of I_α and $\{\tilde{f}^n(0) : n \in \mathbb{N}\} \cap \tilde{I}_\alpha$ is a dense subset of \tilde{I}_α , for each $\alpha \in \mathcal{A}$. The conclusion of the previous paragraph means that ϕ maps the former set to the latter. By continuity, it follows that $\phi(I_\alpha) = \tilde{I}_\alpha$, as claimed. \square

Finally, consider any $\tilde{\lambda} \in \bigcap_{n=1}^\infty \Theta_{\pi, \lambda}^{n*}(\mathbb{R}_+^A)$. By Lemma 28.11, the transformations f and \tilde{f} are conjugated by a homeomorphism $\phi : I \rightarrow \tilde{I}$. Let $\mu = \phi_*^{-1}m$ be the pull-back of Lebesgue measure under the map ϕ . Then μ is an f -invariant measure, that is, $\mu \in \mathcal{M}(\pi, \lambda)$. Lemma 28.12 implies that $\tilde{\lambda}_\alpha = \mu(I_\alpha) = \lambda_\alpha(\mu)$, for every $\alpha \in \mathcal{A}$. This shows that $\tilde{\lambda} = \lambda(\mu) \in \mathcal{C}(\pi, \lambda)$. The proof of Proposition 28.3 is complete.

29. Unique ergodicity

Lemma 29.1. $\mathcal{C}(\pi, \lambda)$ is a (closed) simplicial cone with dimension at most $d - 1$.

Proof. Since \mathbb{R}_+^A is a simplicial cone, the same is true for each $\Theta_{\pi, \lambda}^{n*}(\mathbb{R}_+^A)$. In order to show that their intersection is a closed simplicial cone, let $(\tau^{n, \beta})_{\beta \in \mathcal{A}}$ be a basis of generators of $\Theta_{\pi, \lambda}^{n*}(\mathbb{R}_+^A)$ such that every $\tau^{n, \beta}$ has norm 1. Taking a subsequence, we may suppose every $\tau^{n, \beta}$ converges to some τ^β when $n \rightarrow \infty$. We claim that $(\tau^\beta)_{\beta \in \mathcal{A}}$ generates $\mathcal{C}(\pi, \lambda)$. Indeed, from $\mathcal{C}(\pi, \lambda) \subset \Theta_{\pi, \lambda}^{n*}(\mathbb{R}_+^A)$ we have that every $v \in \mathcal{C}(\pi, \lambda)$ may be written as $v = \sum_\beta a_{n, \beta} \tau^{n, \beta}$ with $a_{n, \beta} > 0$ for every n, β . By Corollary 5.3, there exists $N \geq 1$ such that all the coefficients of $\Theta_{\pi, \lambda}^{n*}$ are positive. Using also that the sequence $\Theta_{\pi, \lambda}^{n*}(\mathbb{R}_+^A)$ is non-increasing, we conclude that there exists $\delta > 0$ such that

$$u_\alpha \geq \delta \|u\| \quad \text{for every } u \in \Theta_{\pi, \lambda}^{n*}(\mathbb{R}_+^A) \text{ and } n \geq N.$$

Then $\|v\| \geq v_\alpha = \sum_\beta a_{n, \beta} \tau_\alpha^{n, \beta} \geq a_{n, \beta} \delta$ for every α . This proves that the coefficients $a_{n, \beta}$ are uniformly bounded. Then, taking a subsequence, we may assume that every $a_{n, \beta}$ converges to some $a_\beta \geq 0$ when n goes to infinity. It follows that $v = \sum_\beta a_\beta \tau^\beta$, which proves the claim. Finally, if the dimension of $\mathcal{C}(\pi, \lambda)$ were d then the cone would have nonempty interior. Then it would contain rationally dependent vectors, and that would contradict Lemma 28.10. \square

Corollary 29.2. *An interval exchange transformation defined by an alphabet with d symbols has at most $d - 1$ invariant ergodic probabilities.*

Proof. This is a direct consequence of Lemma 29.1 and the fact that the ergodic measures are the extremal elements of the cone of invariant measures. \square

As an immediate consequence of Proposition 28.3, we get that an interval exchange transformation defined by (π, λ) is uniquely ergodic if and only if the cone




Figure 35

$\bigcap_{n=1}^{\infty} \Theta_{\pi, \lambda}^{n*}(\mathbb{R}_+^{\mathcal{A}})$ reduces to a half line. We are going to show that this is the case for almost all cases:

Proof of Theorem 28.1. Fix $\pi_* \in C$, $N \geq 1$, and $\varepsilon \in \{0, 1\}^N$ such that the set $\Lambda_* = \Lambda_{\pi_*, N, \varepsilon}$, as defined in (88), is relatively compact in $\Lambda_{\mathcal{A}}$. By Proposition 26.3, there exists $\theta < 1$ such that any projective map $g : \Lambda_{\mathcal{A}} \rightarrow \Lambda_*$ contracts the projective metric by θ , at least. Since (R, ν) is ergodic, by Corollary 27.2, for almost every, $(\pi, \lambda) \in C \times \Lambda_{\mathcal{A}}$ there are infinitely many times $0 < n_1 < \dots < n_j < \dots$ such that $(\pi_j, \lambda_j) = R^{n_j}(\pi, \lambda)$ is in $\{\pi_*\} \times \Lambda_*$ and $n_{j+1} - n_j \geq N$, for all j , and $n_1 \geq N$. We are going to show that the interval exchange transformation defined by any such (π, λ) is uniquely ergodic. Begin by noticing that, for any $n \geq 1$,

$$\Theta_{\pi, \lambda}^{n*} = \Theta_{\pi_s, \lambda_s}^{(n-n_s)*} \dots \Theta_{\pi_j, \lambda_j}^{(n_{j+1}-n_j)*} \dots \Theta_{\pi, \lambda}^{n_1*}$$

where $s \geq 1$ is largest such that $n_s < n$. Since $(\pi_j, \lambda_j) \in \{\pi_*\} \times \Lambda_*$ and $n_{j+1} - n_j \geq N$, the set $\Lambda_{\pi_j, n_{j+1}-n_j, \underline{\theta}_j}$ that contains λ_j is a subset of Λ_* . See figure 35. This means that $\Lambda_{\mathcal{A}}$ is sent inside Λ_* by the projectivization of

$$\Theta_{\pi_j, \lambda_j}^{(n_{j+1}-n_j)*} : \mathbb{R}_+^{\mathcal{A}} \rightarrow \mathbb{R}_+^{\mathcal{A}}.$$

Thus, the latter contracts the projective metric by $\theta < 1$, for all $1 \leq j \leq s - 1$. Since we also assume $n_1 \geq N$, the set $\Lambda_{\pi, n_1, \underline{\theta}}$ that contains λ is also a subset of Λ_* . Consequently, it has finite projective diameter. Using also that the map $\Theta_{\pi_s, \lambda_s}^{(n-n_s)*}$ does not expand the projective metric, by (92), we conclude that

$$\text{diam } \Theta_{\pi, \lambda}^{n*}(\mathbb{R}_+^{\mathcal{A}}) \leq \theta^{s-1} \text{diam } \Lambda_*$$

for all $n \geq 1$, where diam stands for projective diameter. As goes to infinity, the right hand side goes to zero. This means that the intersection of all $\Theta_{\pi, \lambda}^{n*}(\mathbb{R}_+^{\mathcal{A}})$ is reduced to a half line, as claimed. \square

Remark 29.3. Given (π, λ) and $n \geq 1$, the set $\Lambda_{\pi, n, \underline{\varepsilon}}$ that contains λ is the image of $\Lambda_{\mathcal{A}}$ under the projectivization of $\Theta_{\pi, \lambda}^{n*}$. Thus, the conclusion means that the intersection $\bigcap_{n=1}^{\infty} \Lambda_{\pi, n, \underline{\varepsilon}}$ reduces to a single point, for almost every (π, λ) . According to Remark 28.9, this point must coincide with λ .

30. Zorich measure

Here we prove Theorem 8.2. Recall that the invertible Zorich renormalization $\mathcal{Z} : Z_* \rightarrow Z_*$ was defined in section 20 as the first return map of the Rauzy-Veech renormalization $\mathcal{R} : \mathcal{H} \rightarrow \mathcal{H}$ to the domain $Z_* = Z_0 \cup Z_1$, where

$$Z_0 = \left\{ (\pi, \lambda, \tau) \in \mathcal{H} : \lambda_{\alpha(0)} > \lambda_{\alpha(1)} \quad \text{and} \quad \sum_{\alpha \in \mathcal{A}} \tau_\alpha > 0 \right\}$$

and

$$Z_1 = \left\{ (\pi, \lambda, \tau) \in \mathcal{H} : \lambda_{\alpha(0)} < \lambda_{\alpha(1)} \quad \text{and} \quad \sum_{\alpha \in \mathcal{A}} \tau_\alpha < 0 \right\}.$$

It follows from the definition (and Lemma 25.2) that \mathcal{Z} preserves the restriction of the measure m to Z_* . Moreover, \mathcal{Z} preserves the restriction of m to the domain $\{\text{area}(\lambda, \tau) \leq c\} \cap Z_*$, for any $c > 0$. In this regard, observe that \mathcal{Z} preserves the area (48), since \mathcal{R} does.

Also by construction, $P \circ \mathcal{Z} = \mathcal{Z} \circ P$, where $P : \mathcal{H} \rightarrow C \times \Lambda_{\mathcal{A}}$ denotes the canonical projection and \mathcal{Z} is the Zorich renormalization map introduced in section 8. Therefore, \mathcal{Z} preserves the measure

$$\mu = P_*(m | Z_* \cap \{\text{area}(\lambda, \tau) \leq 1\}).$$

Arguing in just the same way as in section 21, we see that the measure μ is absolutely continuous with respect to Lebesgue measure, with density

$$\frac{d\mu}{d_1\lambda}(\pi, \lambda) = \text{vol}(\{\tau \in T_\pi^\varepsilon : \text{area}(\lambda, \tau) \leq 1\}) = \frac{1}{d!} \sum_{i=1}^k \prod_{\beta \in \mathcal{A}} \frac{1}{\lambda \cdot h^{i,\beta}}, \tag{104}$$

where ε is the type of (π, λ) . Here the notations are as before:

$$T_\pi^0 = \left\{ \tau \in T_\pi^+ : \sum_{\alpha \in \mathcal{A}} \tau_\alpha > 0 \right\} \quad \text{and} \quad T_\pi^1 = \left\{ \tau \in T_\pi^+ : \sum_{\alpha \in \mathcal{A}} \tau_\alpha < 0 \right\},$$

T^1, \dots, T^k are pairwise disjoint simplicial cones covering the polyhedral cone T_π^ε up to a positive codimension subset, $(\tau^{i,\beta})_{\beta \in \mathcal{A}}$ is a basis of generators of each T^i , and $h^{i,\beta} = -\Omega_\pi(\tau^{i,\beta})$ for each i and β .

The relation (104) shows that the density of the absolutely continuous Z -invariant measure μ is given by a rational function with degree $-d$ and bounded from zero. The next step is to show that this measure μ is finite.

Example 30.1. Let us give an explicit expression for the density of μ when π is the pair defined in (62). We consider first the case when $(\pi, \lambda) \in C \times \Lambda_{\mathcal{A}}$ has type 0. We

have seen in section 22 that T_π^0 is a simplicial cone, and admits $\tau^\alpha = (\tau_\beta^\alpha)_{\beta \in \mathcal{A}}$ defined by

$$\tau_\beta^\alpha = \begin{cases} 1 & \text{if } \beta = \alpha, \\ -1 & \text{if } \beta = \alpha_0^+, & \text{if } \alpha \neq \alpha(0), \\ 0 & \text{in all other cases,} \end{cases}$$

$$\tau_\beta^\alpha = \begin{cases} 1 & \text{if } \pi_0(\beta) = 1, \\ 0 & \text{otherwise,} & \text{if } \alpha = \alpha(0). \end{cases}$$

as a volume 1 basis of generators. Then $h^\alpha = -\Omega_\pi(\tau^\alpha)$ is given by

$$h_\beta^\alpha = \begin{cases} 1 & \text{if } \beta = \alpha \text{ or } \beta = \alpha_0^+, \\ 0 & \text{otherwise,} & \text{if } \alpha \neq \alpha(0), \end{cases}$$

$$h_\beta^\alpha = \begin{cases} 0 & \text{if } \pi_0(\beta) = 1 \text{ or } \beta = \alpha(1), \\ 1 & \text{otherwise,} & \text{if } \alpha = \alpha(0). \end{cases}$$

It follows that

$$\frac{d\mu}{d_1\lambda}(\pi, \lambda) = \frac{1}{d!} \prod_{\alpha \neq \alpha(0)} \left(\frac{1}{\lambda_\alpha + \lambda_{\alpha_0^+}} \right) \cdot \frac{1}{\sum_{\beta \neq \alpha(1)} \lambda_\beta}.$$

The case when (π, λ) has type 1 is analogous, and one gets

$$\frac{d\mu}{d_1\lambda}(\pi, \lambda) = \frac{1}{d!} \prod_{\alpha \neq \alpha(1)} \left(\frac{1}{\lambda_\alpha + \lambda_{\alpha_1^+}} \right) \cdot \frac{1}{\sum_{\beta \neq \alpha(0)} \lambda_\beta}.$$

In particular, for $d = 2$ this gives

$$\frac{d\mu}{d_1\lambda}(\lambda) = \begin{cases} 1/(2\lambda_B) & \text{if } \lambda_A < \lambda_B, \\ 1/(2\lambda_A) & \text{if } \lambda_B < \lambda_A. \end{cases}$$

Notice that the density is bounded on $\Lambda_{\mathcal{A}}$, and so the measure μ is finite. While boundedness is specific to the case $d = 2$, finiteness holds in general, as we are going to see.

Proposition 30.2. *The measure $\mu(C \times \Lambda_{\mathcal{A}})$ is finite.*

Proof. Given $\pi \in C$, let Λ_ε be the subset of $\lambda \in \mathbb{R}_+^{\mathcal{A}}$ such that $\lambda_{\alpha(\varepsilon)} > \lambda_{\alpha(1-\varepsilon)}$. Then

$$\mu(C \times \Lambda_{\mathcal{A}}) = \sum_{\pi \in C} \sum_{\varepsilon=0,1} \int_{\Lambda_\varepsilon} \text{vol}(\{\tau \in T_\pi^\varepsilon : \text{area}(\lambda, \tau) \leq 1\}) d_1\lambda.$$

Using (104) we deduce that

$$\mu(C \times \Lambda_{\mathcal{A}}) = \sum_{\pi \in C} \sum_{\varepsilon=0,1} \int_{\Lambda_{\varepsilon}} \sum_{i=1}^k \prod_{\beta \in \mathcal{A}} \frac{1}{\lambda \cdot h^{i,\beta}} d_1 \lambda. \tag{105}$$

To prove the proposition we only have to show that the integral is finite, for every fixed $\pi \in C$, $\varepsilon \in \{0, 1\}$, and $i = 1, \dots, k$. Let us consider $\varepsilon = 0$; the case $\varepsilon = 1$ is analogous. Then the basis of generators $(\tau^{i,\beta})_{\beta \in \mathcal{A}}$ is contained in the closure of T_{π}^0 . For simplicity, we write $h^{\beta} = h^{i,\beta}$ in what follows.

Let \mathcal{N} denote the set of integer vectors $n = (n_{\alpha})_{\alpha \in \mathcal{A}}$ such that $n_{\alpha} \geq 0$ for all $\alpha \in \mathcal{A}$, and the n_{α} are not all zero. As in (82), define

$$\Lambda(n) = \{ \lambda \in \Lambda_{\mathcal{A}} : 2^{-n_{\alpha}} \leq \lambda_{\alpha} d < 2^{-n_{\alpha}+1} \text{ for every } \alpha \in \mathcal{A} \},$$

except that for $n_{\alpha} = 0$ the second inequality is omitted. By Lemma 24.2, the $\Lambda(n)$ cover $\Lambda_{\mathcal{A}}$ and satisfy

$$c_1 2^{-\sum_{\alpha} n_{\alpha}} \leq \text{vol}_{d-1} \Lambda(n) \leq c_1^{-1} 2^{-\sum_{\alpha} n_{\alpha}} \tag{106}$$

for some $c_1 > 0$. In what follows we consider $n_{\alpha(0)} \leq n_{\alpha(1)}$, for the corresponding $\Lambda(n)$ suffice to cover $\Lambda_0 = \{ \lambda_{\alpha(0)} > \lambda_{\alpha(1)} \}$. For each $\beta \in \mathcal{A}$, let $\mathcal{A}(\beta)$ be the subset of $\alpha \in \mathcal{A}$ such that $h_{\alpha}^{\beta} > 0$. Let $c_2 > 0$ be the minimum of the non-zero h_{α}^{β} , over all α and β . Then

$$\lambda \cdot h^{\beta} = \sum_{\alpha \in \mathcal{A}(\beta)} h_{\alpha}^{\beta} \lambda_{\alpha} \geq \sum_{\alpha \in \mathcal{A}(\beta)} c_2 d^{-1} 2^{-n_{\alpha}} \geq c_2 d^{-1} 2^{-\min_{\alpha \in \mathcal{A}(\beta)} n_{\alpha}}$$

for every $\beta \in \mathcal{A}$. Using (106) we deduce that

$$\int_{\Lambda(n)} \prod_{\beta \in \mathcal{A}} \frac{1}{\lambda \cdot h^{\beta}} d_1 \lambda \leq K 2^{\sum_{\beta \in \mathcal{A}} \min_{\alpha \in \mathcal{A}(\beta)} n_{\alpha} - \sum_{\alpha \in \mathcal{A}} n_{\alpha}}, \tag{107}$$

where the constant $K = (2/c_1)(d/c_2)^d$.

Lemma 30.3. *Assuming $n_{\alpha(0)} \leq n_{\alpha(1)}$, we have*

$$-\sum_{\beta \in \mathcal{A}} \min_{\alpha \in \mathcal{A}(\beta)} n_{\alpha} + \sum_{\alpha \in \mathcal{A}} n_{\alpha} \geq \max_{\alpha \in \mathcal{A}} n_{\alpha}.$$

Proof. Let $0 = n^0 < n^1 < \dots$ be the different values taken by n_{α} , and \mathcal{B}^i , $i \geq 0$ be the set of values of $\alpha \in \mathcal{A}$ such that $n_{\alpha} \geq n^i$. On the one hand,

$$\sum_{\alpha \in \mathcal{A}} n_{\alpha} = \sum_{i \geq 1} n_i (\#\mathcal{B}^i - \#\mathcal{B}^{i+1}) = \sum_{i \geq 1} \#\mathcal{B}^i (n^i - n^{i-1}). \tag{108}$$

On the other hand, $\min_{\mathcal{A}(\beta)} n_\alpha \geq n_i$ if and only if $\mathcal{A}(\beta) \subset \mathcal{B}^i$. Consequently,

$$\begin{aligned} \sum_{\beta \in \mathcal{A}} \min_{\mathcal{A}(\beta)} n_\alpha &= \sum_{i \geq 1} n_i (\#\{\beta : \mathcal{A}(\beta) \subset \mathcal{B}^i\} - \#\{\beta : \mathcal{A}(\beta) \subset \mathcal{B}^{i+1}\}) \\ &= \sum_{i \geq 1} \#\{\beta : \mathcal{A}(\beta) \subset \mathcal{B}^i\} (n^i - n^{i-1}). \end{aligned} \tag{109}$$

Observe that $\mathcal{A}(\beta) \subset \mathcal{B}^i$ if and only if $h_\alpha^\beta = 0$ for all $\alpha \in \mathcal{A} \setminus \mathcal{B}^i$. Observe also that the assumption $n_{\alpha(0)} \leq n_{\alpha(1)}$ means that if $\alpha(1) \in \mathcal{B}^i$ then $\alpha(0) \in \mathcal{B}^i$. Using Proposition 23.1 (with $\mathcal{B} = \mathcal{A} \setminus \mathcal{B}^i$), we obtain

$$\#\{\beta : \mathcal{A}(\beta) \subset \mathcal{B}^i\} < \#\mathcal{B}^i. \tag{110}$$

Putting (108)–(110) together, we find that

$$- \sum_{\beta \in \mathcal{A}} \min_{\mathcal{A}(\beta)} n_\alpha + \sum_{\alpha \in \mathcal{A}} n_\alpha \geq \sum_{i \geq 1}^k (n^i - n^{i-1}) = \max_{\alpha \in \mathcal{A}} n_\alpha.$$

This proves the lemma. □

Replacing the conclusion of the lemma in (107) we obtain, for every $n \in \mathcal{N}$,

$$\int_{\Lambda(n)} \prod_{\beta \in \mathcal{A}} \frac{1}{\lambda \cdot h^\beta} d_1 \lambda \leq K 2^{-\max_{\mathcal{A}} n_\alpha}. \tag{111}$$

For each $m \geq 0$ there are at most $(m + 1)^d$ vectors $n \in \mathcal{N}$ with $\max_{\mathcal{A}} n_\alpha = m$. So, (111) implies that the integral in (105) is bounded above by

$$\sum_{m=0}^{\infty} K(m + 1)^d 2^{-m} < \infty$$

for every $\pi \in C$ and $1 \leq i \leq k$. The proof of Proposition 30.2 is complete. □

To finish the proof of Theorem 8.2 we only have to observe that the system (Z, μ) is ergodic. This can be shown in the same way we proved, in Corollary 27.2, that (R, ν) is ergodic. We just outline the arguments. As noted before, (Z, μ) is a Markov system. Since μ is invariant and finite, (Z, μ) is a recurrent system. Consider any relatively compact subsimplex $\{\pi\} \times \Lambda_*$ which is mapped to a whole $\{\pi_0\} \times \Lambda_{\mathcal{A}}$ by some iterate Z^N . The map induced by Z on Λ_* has a bounded distortion property as in Lemma 26.2. For the same reason as in Proposition 25.5, that implies the induced map is ergodic relative to μ restricted to Λ_* . It follows, using Lemma 25.4, that Z itself is ergodic relative to μ . This proves the claim.

References

- [1] A. Avila and G. Forni, *Weak mixing for interval exchange transformations and translation flows*, Ann. of Math., to appear.
- [2] A. Avila, S. Gouëzel, and J.-C. Yoccoz, *The Teichmüller flow is exponentially mixing*, preprint.
- [3] A. Avila and M. Viana, *Symplecticity of Lyapunov spectra: Proof of the Zorich-Kontsevich conjecture* (2005), preprint, available at www.preprint.impa.br.
- [4] ———, *Dynamics in the moduli space of Abelian differentials*, Port. Math. (N.S.) **62** (2005), 531–547.
- [5] G. Birkhoff, *Lattice theory*, American Mathematical Society Colloquium Publications, vol. XXV, American Mathematical Society, Providence, R.I., 1967.
- [6] M. Boshernitzan, *A condition for minimal interval exchange maps to be uniquely ergodic*, Duke Math. J. **52** (1985), no. 3, 723–752.
- [7] A. I. Bufetov, *Decay of correlations for the Rauzy-Veech-Zorich induction map on the space of interval exchanges transformations and central limit theorem for the Teichmüller flow in the the space of Abelian differentials* (2005), preprint, available at www.arxiv.org.
- [8] G. Forni, *Deviation of ergodic averages for area-preserving flows on surfaces of higher genus*, Ann. of Math. (2) **155** (2002), no. 1, 1–103.
- [9] M. Keane, *Interval exchange transformations*, Math. Z. **141** (1975), 25–31.
- [10] ———, *Non-ergodic interval exchange transformations*, Israel J. Math. **26** (1977), no. 2, 188–196.
- [11] M. S. Keane and G. Rauzy, *Stricte ergodicité des échanges d’intervalles*, Math. Z. **174** (1980), no. 3, 203–212.
- [12] S. P. Kerckhoff, *Simplicial systems for interval exchange maps and measured foliations*, Ergodic Theory Dynam. Systems **5** (1985), no. 2, 257–271.
- [13] H. B. Keynes and D. Newton, *A “minimal”, non-uniquely ergodic interval exchange transformation*, Math. Z. **148** (1976), no. 2, 101–105.
- [14] M. Kontsevich and A. Zorich, *Connected components of the moduli spaces of Abelian differentials with prescribed singularities*, Invent. Math. **153** (2003), no. 3, 631–678.
- [15] R. Mañé, *Ergodic theory and differentiable dynamics*, Ergebnisse der Mathematik und ihrer Grenzgebiete (3), vol. 8, Springer-Verlag, Berlin, 1987.
- [16] S. Marmi, P. Moussa, and J.-C. Yoccoz, *The cohomological equation for Roth-type interval exchange maps*, J. Amer. Math. Soc. **18** (2005), no. 4, 823–872.
- [17] H. Masur, *Interval exchange transformations and measured foliations*, Ann. of Math. (2) **115** (1982), no. 1, 169–200.
- [18] A. Nogueira and D. Rudolph, *Topological weak-mixing of interval exchange maps*, Ergodic Theory Dynam. Systems **17** (1997), no. 5, 1183–1209.
- [19] G. Rauzy, *Échanges d’intervalles et transformations induites*, Acta Arith. **34** (1979), no. 4, 315–328.
- [20] M. Rees, *An alternative approach to the ergodic theory of measured foliations on surfaces*, Ergodic Theory Dynamical Systems **1** (1981), no. 4, 461–488 (1982).
- [21] W. A. Veech, *Interval exchange transformations*, J. Analyse Math. **33** (1978), 222–272.
- [22] ———, *Gauss measures for transformations on the space of interval exchange maps*, Ann. of Math. (2) **115** (1982), no. 1, 201–242.
- [23] ———, *The Teichmüller geodesic flow*, Ann. of Math. (2) **124** (1986), no. 3, 441–530.

- [24] J.-C. Yoccoz, *Continued fraction algorithms for interval exchange maps: an introduction*, École de Physique des Houches (2006), *Frontiers in Number Theory, Physics and Geometry*, volume 1: On random matrices, zeta functions and dynamical systems, 2006.
- [25] A. Zorich, *Finite Gauss measure on the space of interval exchange transformations. Lyapunov exponents*, *Ann. Inst. Fourier (Grenoble)* **46** (1996), no. 2, 325–370.
- [26] ———, *How do the leaves of a closed 1-form wind around a surface?*, *Pseudoperiodic topology*, 1999, pp. 135–178.