


Introduction to Dynamical Systems



Michael Brin
and Garrett Stuck

CAMBRIDGE

CAMBRIDGE

more information - www.cambridge.org/0521808413

This page intentionally left blank

Introduction to Dynamical Systems

This book provides a broad introduction to the subject of dynamical systems, suitable for a one- or two-semester graduate course. In the first chapter, the authors introduce over a dozen examples, and then use these examples throughout the book to motivate and clarify the development of the theory. Topics include topological dynamics, symbolic dynamics, ergodic theory, hyperbolic dynamics, one-dimensional dynamics, complex dynamics, and measure-theoretic entropy. The authors top off the presentation with some beautiful and remarkable applications of dynamical systems to such areas as number theory, data storage, and Internet search engines.

This book grew out of lecture notes from the graduate dynamical systems course at the University of Maryland, College Park, and reflects not only the tastes of the authors, but also to some extent the collective opinion of the Dynamics Group at the University of Maryland, which includes experts in virtually every major area of dynamical systems.

Michael Brin is a professor of mathematics at the University of Maryland. He is the author of over thirty papers, three of which appeared in the *Annals of Mathematics*. Professor Brin is also an editor of *Forum Mathematicum*.

Garrett Stuck is a former professor of mathematics at the University of Maryland and currently works in the structured finance industry. He has coauthored several textbooks, including *The Mathematica Primer* (Cambridge University Press, 1998). Dr. Stuck is also a founder of Chalkfree, Inc., a software company.

Introduction to Dynamical Systems

MICHAEL BRIN

University of Maryland

GARRETT STUCK

University of Maryland



CAMBRIDGE
UNIVERSITY PRESS

PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS

The Edinburgh Building, Cambridge CB2 2RU, UK
40 West 20th Street, New York, NY 10011-4211, USA
477 Williamstown Road, Port Melbourne, VIC 3207, Australia
Ruiz de Alarcón 13, 28014 Madrid, Spain
Dock House, The Waterfront, Cape Town 8001, South Africa

<http://www.cambridge.org>

© Michael Brin, Garrett Stuck 2004

First published in printed format 2002

ISBN 0-511-02937-3 eBook (Adobe Reader)

ISBN 0-521-80841-3 hardback

To Eugenia, Pamela, Sergey, Sam, Jonathan, and Catherine
for their patience and support.

Contents

Introduction	<i>page xi</i>
1 Examples and Basic Concepts	1
1.1 The Notion of a Dynamical System	1
1.2 Circle Rotations	3
1.3 Expanding Endomorphisms of the Circle	5
1.4 Shifts and Subshifts	7
1.5 Quadratic Maps	9
1.6 The Gauss Transformation	11
1.7 Hyperbolic Toral Automorphisms	13
1.8 The Horseshoe	15
1.9 The Solenoid	17
1.10 Flows and Differential Equations	19
1.11 Suspension and Cross-Section	21
1.12 Chaos and Lyapunov Exponents	23
1.13 Attractors	25
2 Topological Dynamics	28
2.1 Limit Sets and Recurrence	28
2.2 Topological Transitivity	31
2.3 Topological Mixing	33
2.4 Expansiveness	35
2.5 Topological Entropy	36
2.6 Topological Entropy for Some Examples	41
2.7 Equicontinuity, Distality, and Proximality	45
2.8 Applications of Topological Recurrence to Ramsey Theory	48

3	Symbolic Dynamics	54
3.1	Subshifts and Codes	55
3.2	Subshifts of Finite Type	56
3.3	The Perron–Frobenius Theorem	57
3.4	Topological Entropy and the Zeta Function of an SFT	60
3.5	Strong Shift Equivalence and Shift Equivalence	62
3.6	Substitutions	64
3.7	Sofic Shifts	66
3.8	Data Storage	67
4	Ergodic Theory	69
4.1	Measure-Theory Preliminaries	69
4.2	Recurrence	71
4.3	Ergodicity and Mixing	73
4.4	Examples	77
4.5	Ergodic Theorems	80
4.6	Invariant Measures for Continuous Maps	85
4.7	Unique Ergodicity and Weyl’s Theorem	87
4.8	The Gauss Transformation Revisited	90
4.9	Discrete Spectrum	94
4.10	Weak Mixing	97
4.11	Applications of Measure-Theoretic Recurrence to Number Theory	101
4.12	Internet Search	103
5	Hyperbolic Dynamics	106
5.1	Expanding Endomorphisms Revisited	107
5.2	Hyperbolic Sets	108
5.3	ϵ -Orbits	110
5.4	Invariant Cones	114
5.5	Stability of Hyperbolic Sets	117
5.6	Stable and Unstable Manifolds	118
5.7	Inclination Lemma	122
5.8	Horseshoes and Transverse Homoclinic Points	124
5.9	Local Product Structure and Locally Maximal Hyperbolic Sets	128
5.10	Anosov Diffeomorphisms	130
5.11	Axiom A and Structural Stability	133
5.12	Markov Partitions	134
5.13	Appendix: Differentiable Manifolds	137

6	Ergodicity of Anosov Diffeomorphisms	141
6.1	Hölder Continuity of the Stable and Unstable Distributions	141
6.2	Absolute Continuity of the Stable and Unstable Foliations	144
6.3	Proof of Ergodicity	151
7	Low-Dimensional Dynamics	153
7.1	Circle Homeomorphisms	153
7.2	Circle Diffeomorphisms	160
7.3	The Sharkovsky Theorem	162
7.4	Combinatorial Theory of Piecewise-Monotone Mappings	170
7.5	The Schwarzian Derivative	178
7.6	Real Quadratic Maps	181
7.7	Bifurcations of Periodic Points	183
7.8	The Feigenbaum Phenomenon	189
8	Complex Dynamics	191
8.1	Complex Analysis on the Riemann Sphere	191
8.2	Examples	194
8.3	Normal Families	197
8.4	Periodic Points	198
8.5	The Julia Set	200
8.6	The Mandelbrot Set	205
9	Measure-Theoretic Entropy	208
9.1	Entropy of a Partition	208
9.2	Conditional Entropy	211
9.3	Entropy of a Measure-Preserving Transformation	213
9.4	Examples of Entropy Calculation	218
9.5	Variational Principle	221
	<i>Bibliography</i>	225
	<i>Index</i>	231

Introduction

The purpose of this book is to provide a broad and general introduction to the subject of dynamical systems, suitable for a one- or two-semester graduate course. We introduce the principal themes of dynamical systems both through examples and by explaining and proving fundamental and accessible results. We make no attempt to be exhaustive in our treatment of any particular area.

This book grew out of lecture notes from the graduate dynamical systems course at the University of Maryland, College Park. The choice of topics reflects not only the tastes of the authors, but also to a large extent the collective opinion of the Dynamics Group at the University of Maryland, which includes experts in virtually every major area of dynamical systems.

Early versions of this book have been used by several instructors at Maryland, the University of Bonn, and Pennsylvania State University. Experience shows that with minor omissions the first five chapters of the book can be covered in a one-semester course. Instructors who wish to cover a different set of topics may safely omit some of the sections at the end of Chapter 1, §§2.7–§2.8, §§3.5–3.8, and §§4.8–4.12, and then choose from topics in later chapters. Examples from Chapter 1 are used throughout the book. Chapter 6 depends on Chapter 5, but the other chapters are essentially independent. Every section ends with exercises (starred exercises are the most difficult).

The exposition of most of the concepts and results in this book has been refined over the years by various authors. Since most of these ideas have appeared so often and in so many variants in the literature, we have not attempted to identify the original sources. In many cases, we followed the written exposition from specific sources listed in the bibliography. These sources cover particular topics in greater depth than we do here, and we recommend them for further reading. We also benefited from the advice and guidance of a number of specialists, including Joe Auslander, Werner Ballmann,

Ken Berg, Mike Boyle, Boris Hasselblatt, Michael Jakobson, Anatole Katok, Michal Misiurewicz, and Dan Rudolph. We thank them for their contributions. We are especially grateful to Vitaly Bergelson for his contributions to the treatment of applications of topological dynamics and ergodic theory to combinatorial number theory. We thank the students who used early versions of this book in our classes, and who found many typos, errors, and omissions.

CHAPTER ONE

Examples and Basic Concepts

Dynamical systems is the study of the long-term behavior of evolving systems. The modern theory of dynamical systems originated at the end of the 19th century with fundamental questions concerning the stability and evolution of the solar system. Attempts to answer those questions led to the development of a rich and powerful field with applications to physics, biology, meteorology, astronomy, economics, and other areas.

By analogy with celestial mechanics, the evolution of a particular state of a dynamical system is referred to as an *orbit*. A number of themes appear repeatedly in the study of dynamical systems: properties of individual orbits; periodic orbits; typical behavior of orbits; statistical properties of orbits; randomness vs. determinism; entropy; chaotic behavior; and stability under perturbation of individual orbits and patterns. We introduce some of these themes through the examples in this chapter.

We use the following notation throughout the book: \mathbb{N} is the set of positive integers; $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$; \mathbb{Z} is the set of integers; \mathbb{Q} is the set of rational numbers; \mathbb{R} is the set of real numbers; \mathbb{C} is the set of complex numbers; \mathbb{R}^+ is the set of positive real numbers; $\mathbb{R}_0^+ = \mathbb{R}^+ \cup \{0\}$.

1.1 The Notion of a Dynamical System

A *discrete-time dynamical system* consists of a non-empty set X and a map $f: X \rightarrow X$. For $n \in \mathbb{N}$, the n th iterate of f is the n -fold composition $f^n = f \circ \dots \circ f$; we define f^0 to be the identity map, denoted Id . If f is invertible, then $f^{-n} = f^{-1} \circ \dots \circ f^{-1}$ (n times). Since $f^{n+m} = f^n \circ f^m$, these iterates form a group if f is invertible, and a semigroup otherwise.

Although we have defined dynamical systems in a completely abstract setting, where X is simply a set, in practice X usually has additional structure

that is preserved by the map f . For example, (X, f) could be a measure space and a measure-preserving map; a topological space and a continuous map; a metric space and an isometry; or a smooth manifold and a differentiable map.

A *continuous-time dynamical system* consists of a space X and a one-parameter family of maps of $\{f^t: X \rightarrow X\}$, $t \in \mathbb{R}$ or $t \in \mathbb{R}_0^+$, that forms a one-parameter group or semigroup, i.e., $f^{t+s} = f^t \circ f^s$ and $f^0 = \text{Id}$. The dynamical system is called a *flow* if the time t ranges over \mathbb{R} , and a *semiflow* if t ranges over \mathbb{R}_0^+ . For a flow, the *time- t map* f^t is invertible, since $f^{-t} = (f^t)^{-1}$. Note that for a fixed t_0 , the iterates $(f^{t_0})^n = f^{t_0 n}$ form a discrete-time dynamical system.

We will use the term *dynamical system* to refer to either discrete-time or continuous-time dynamical systems. Most concepts and results in dynamical systems have both discrete-time and continuous-time versions. The continuous-time version can often be deduced from the discrete-time version. In this book, we focus mainly on discrete-time dynamical systems, where the results are usually easier to formulate and prove.

To avoid having to define basic terminology in four different cases, we write the elements of a dynamical system as f^t , where t ranges over \mathbb{Z} , \mathbb{N}_0 , \mathbb{R} , or \mathbb{R}_0^+ , as appropriate. For $x \in X$, we define the *positive semiorbit* $\mathcal{O}_f^+(x) = \bigcup_{t \geq 0} f^t(x)$. In the invertible case, we define the *negative semiorbit* $\mathcal{O}_f^-(x) = \bigcup_{t \leq 0} f^t(x)$, and the *orbit* $\mathcal{O}_f(x) = \mathcal{O}_f^+(x) \cup \mathcal{O}_f^-(x) = \bigcup_t f^t(x)$ (we omit the subscript “ f ” if the context is clear). A point $x \in X$ is a *periodic point* of *period* $T > 0$ if $f^T(x) = x$. The orbit of a periodic point is called a *periodic orbit*. If $f^t(x) = x$ for all t , then x is a *fixed point*. If x is periodic, but not fixed, then the smallest positive T , such that $f^T(x) = x$, is called the *minimal period* of x . If $f^s(x)$ is periodic for some $s > 0$, we say that x is *eventually periodic*. In invertible dynamical systems, eventually periodic points are periodic.

For a subset $A \subset X$ and $t > 0$, let $f^t(A)$ be the image of A under f^t , and let $f^{-t}(A)$ be the preimage under f^t , i.e., $f^{-t}(A) = (f^t)^{-1}(A) = \{x \in X: f^t(x) \in A\}$. Note that $f^{-t}(f^t(A))$ contains A , but, for a non-invertible dynamical system, is generally not equal to A . A subset $A \subset X$ is *f -invariant* if $f^t(A) \subset A$ for all t ; *forward f -invariant* if $f^t(A) \subset A$ for all $t \geq 0$; and *backward f -invariant* if $f^{-t}(A) \subset A$ for all $t \geq 0$.

In order to classify dynamical systems, we need a notion of equivalence. Let $f^t: X \rightarrow X$ and $g^t: Y \rightarrow Y$ be dynamical systems. A *semiconjugacy* from (Y, g) to (X, f) (or, briefly, from g to f) is a surjective map $\pi: Y \rightarrow X$ satisfying $f^t \circ \pi = \pi \circ g^t$, for all t . We express this formula schematically by

saying that the following diagram commutes:

$$\begin{array}{ccc} Y & \xrightarrow{g} & Y \\ \pi \downarrow & & \downarrow \pi \\ X & \xrightarrow{f} & X \end{array}$$

An invertible semiconjugacy is called a *conjugacy*. If there is a conjugacy from one dynamical system to another, the two systems are said to be *conjugate*; conjugacy is an equivalence relation. To study a particular dynamical system, we often look for a conjugacy or semiconjugacy with a better-understood model. To classify dynamical systems, we study equivalence classes determined by conjugacies preserving some specified structure. Note that for some classes of dynamical systems (e.g., measure-preserving transformations) the word *isomorphism* is used instead of “conjugacy.”

If there is a semiconjugacy π from g to f , then (X, f) is a *factor* of (Y, g) , and (Y, g) is an *extension* of (X, f) . The map $\pi: Y \rightarrow X$ is also called a *factor map* or *projection*. The simplest example of an extension is the *direct product*

$$(f_1 \times f_2)^t: X_1 \times X_2 \rightarrow X_1 \times X_2$$

of two dynamical systems $f_i^t: X_i \rightarrow X_i, i = 1, 2$, where $(f_1 \times f_2)^t(x_1, x_2) = (f_1^t(x_1), f_2^t(x_2))$. Projection of $X_1 \times X_2$ onto X_1 or X_2 is a semiconjugacy, so (X_1, f_1) and (X_2, f_2) are factors of $(X_1 \times X_2, f_1 \times f_2)$.

An extension (Y, g) of (X, f) with factor map $\pi: Y \rightarrow X$ is called a *skew product* over (X, f) if $Y = X \times F$, and π is the projection onto the first factor or, more generally, if Y is a fiber bundle over X with projection π .

Exercise 1.1.1. Show that the complement of a forward invariant set is backward invariant, and vice versa. Show that if f is bijective, then an invariant set A satisfies $f^t(A) = A$ for all t . Show that this is false, in general, if f is not bijective.

Exercise 1.1.2. Suppose (X, f) is a factor of (Y, g) by a semiconjugacy $\pi: Y \rightarrow X$. Show that if $y \in Y$ is a periodic point, then $\pi(y) \in X$ is periodic. Give an example to show that the preimage of a periodic point does not necessarily contain a periodic point.

1.2 Circle Rotations

Consider the unit circle $S^1 = [0, 1] / \sim$, where \sim indicates that 0 and 1 are identified. Addition mod 1 makes S^1 an abelian group. The natural distance

on $[0, 1]$ induces a distance on S^1 ; specifically,

$$d(x, y) = \min(|x - y|, 1 - |x - y|).$$

Lebesgue measure on $[0, 1]$ gives a natural measure λ on S^1 , also called Lebesgue measure λ .

We can also describe the circle as the set $S^1 = \{z \in \mathbb{C}: |z| = 1\}$, with complex multiplication as the group operation. The two notations are related by $z = e^{2\pi i x}$, which is an isometry if we divide arc length on the multiplicative circle by 2π . We will generally use the additive notation for the circle.

For $\alpha \in \mathbb{R}$, let R_α be the rotation of S^1 by angle $2\pi\alpha$, i.e.,

$$R_\alpha x = x + \alpha \pmod{1}.$$

The collection $\{R_\alpha: \alpha \in [0, 1)\}$ is a commutative group with composition as group operation, $R_\alpha \circ R_\beta = R_\gamma$, where $\gamma = \alpha + \beta \pmod{1}$. Note that R_α is an isometry: It preserves the distance d . It also preserves Lebesgue measure λ , i.e., the Lebesgue measure of a set is the same as the Lebesgue measure of its preimage.

If $\alpha = p/q$ is rational, then $R_\alpha^q = \text{Id}$, so every orbit is periodic. On the other hand, if α is irrational, then every positive semiorbit is dense in S^1 . Indeed, the pigeon-hole principle implies that, for any $\epsilon > 0$, there are $m, n < 1/\epsilon$ such that $m < n$ and $d(R_\alpha^m, R_\alpha^n) < \epsilon$. Thus R^{n-m} is rotation by an angle less than ϵ , so every positive semiorbit is ϵ -dense in S^1 (i.e., comes within distance ϵ of every point in S^1). Since ϵ is arbitrary, every positive semiorbit is dense.

For α irrational, density of every orbit of R_α implies that S^1 is the only R_α -invariant closed non-empty subset. A dynamical system with no proper closed non-empty invariant subsets is called *minimal*. In Chapter 4, we show that any measurable R_α -invariant subset of S^1 has either measure zero or full measure. A measurable dynamical system with this property is called *ergodic*.

Circle rotations are examples of an important class of dynamical systems arising as group translations. Given a group G and an element $h \in G$, define maps $L_h: G \rightarrow G$ and $R_h: G \rightarrow G$ by

$$L_h g = hg \quad \text{and} \quad R_h g = gh.$$

These maps are called *left* and *right translation* by h . If G is commutative, $L_h = R_h$.

A *topological group* is a topological space G with a group structure such that group multiplication $(g, h) \mapsto gh$, and the inverse $g \mapsto g^{-1}$ are

continuous maps. A continuous homomorphism of a topological group to itself is called an *endomorphism*; an invertible endomorphism is an *automorphism*. Many important examples of dynamical systems arise as translations or endomorphisms of topological groups.

Exercise 1.2.1. Show that for any $k \in \mathbb{Z}$, there is a continuous semiconjugacy from R_α to $R_{k\alpha}$.

Exercise 1.2.2. Prove that for any finite sequence of decimal digits there is an integer $n > 0$ such that the decimal representation of 2^n starts with that sequence of digits.

Exercise 1.2.3. Let G be a topological group. Prove that for each $g \in G$, the closure $H(g)$ of the set $\{g^n\}_{n=-\infty}^\infty$ is a commutative subgroup of G . Thus, if G has a minimal left translation, then G is abelian.

***Exercise 1.2.4.** Show that R_α and R_β are conjugate by a homeomorphism if and only if $\alpha = \pm\beta \pmod{1}$.

1.3 Expanding Endomorphisms of the Circle

For $m \in \mathbb{Z}$, $|m| > 1$, define the *times- m* map $E_m: S^1 \rightarrow S^1$ by

$$E_mx = mx \pmod{1}.$$

This map is a non-invertible group endomorphism of S^1 . Every point has m preimages. In contrast to a circle rotation, E_m expands arc length and distances between nearby points by a factor of m : If $d(x, y) \leq 1/(2m)$, then $d(E_mx, E_my) = md(x, y)$. A map (of a metric space) that expands distances between nearby points by a factor of at least $\mu > 1$ is called *expanding*.

The map E_m preserves Lebesgue measure λ on S^1 in the following sense: if $A \subset S^1$ is measurable, then $\lambda(E_m^{-1}(A)) = \lambda(A)$ (Exercise 1.3.1). Note, however, that for a sufficiently small interval I , $\lambda(E_m(I)) = m\lambda(I)$. We will show later that E_m is ergodic (Proposition 4.4.2).

Fix a positive integer $m > 1$. We will now construct a semiconjugacy from another natural dynamical system to E_m . Let $\Sigma = \{0, \dots, m-1\}^{\mathbb{N}}$ be the set of sequences of elements in $\{0, \dots, m-1\}$. The *shift* $\sigma: \Sigma \rightarrow \Sigma$ discards the first element of a sequence and shifts the remaining elements one place to the left:

$$\sigma((x_1, x_2, x_3, \dots)) = (x_2, x_3, x_4, \dots).$$

A base- m expansion of $x \in [0, 1]$ is a sequence $(x_i)_{i \in \mathbb{N}} \in \Sigma$ such that $x = \sum_{i=1}^{\infty} x_i/m^i$. In analogy with decimal notation, we write $x = 0.x_1x_2x_3\dots$

Base- m expansions are not always unique: A fraction whose denominator is a power of m is represented both by a sequence with trailing $m - 1$ s and a sequence with trailing zeros. For example, in base 5, we have $0.144\dots = 0.200\dots = 2/5$.

Define a map

$$\phi: \Sigma \rightarrow [0, 1], \quad \phi((x_i)_{i \in \mathbb{N}}) = \sum_{i=1}^{\infty} \frac{x_i}{m^i}.$$

We can consider ϕ as a map into S^1 by identifying 0 and 1. This map is surjective, and one-to-one except on the countable set of sequences with trailing zeros or $m - 1$'s. If $x = 0.x_1x_2x_3\dots \in [0, 1)$, then $E_mx = 0.x_2x_3\dots$. Thus, $\phi \circ \sigma = E_m \circ \phi$, so ϕ is a semiconjugacy from σ to E_m .

We can use the semiconjugacy of E_m with the shift σ to deduce properties of E_m . For example, a sequence $(x_i) \in \Sigma$ is a periodic point for σ with period k if and only if it is a periodic sequence with period k , i.e., $x_{k+i} = x_i$ for all i . It follows that the number of periodic points of σ of period k is m^k . More generally, (x_i) is eventually periodic for σ if and only if the sequence (x_i) is eventually periodic. A point $x \in S^1 = [0, 1] / \sim$ is periodic for E_m with period k if and only if x has a base- m expansion $x = 0.x_1x_2\dots$ that is periodic with period k . Therefore, the number of periodic points of E_m of period k is $m^k - 1$ (since 0 and 1 are identified).

Let $\mathcal{F}_m = \bigcup_{k=1}^{\infty} \{0, \dots, m-1\}^k$ be the set of all finite sequences of elements of the set $\{0, \dots, m-1\}$. A subset $A \subset [0, 1]$ is dense if and only if every finite sequence $w \in \mathcal{F}_m$ occurs at the beginning of the base- m expansion of some element of A . It follows that the set of periodic points is dense in S^1 . The orbit of a point $x = 0.x_1x_2\dots$ is dense in S^1 if and only if every finite sequence from \mathcal{F}_m appears in the sequence (x_i) . Since \mathcal{F}_m is countable, we can construct such a point by concatenating all elements of \mathcal{F}_m .

Although ϕ is not one-to-one, we can construct a right inverse to ϕ . Consider the partition of $S^1 = [0, 1] / \sim$ into intervals

$$P_k = [k/m, (k+1)/m), \quad 0 \leq k \leq m-1.$$

For $x \in [0, 1]$, define $\psi_i(x) = k$ if $E_m^i x \in P_k$. The map $\psi: S^1 \rightarrow \Sigma$, given by $x \mapsto (\psi_i(x))_{i=0}^{\infty}$, is a right inverse for ϕ , i.e., $\phi \circ \psi = \text{Id}: S^1 \rightarrow S^1$. In particular, $x \in S^1$ is uniquely determined by the sequence $(\psi_i(x))$.

The use of partitions to code points by sequences is the principal motivation for *symbolic dynamics*, the study of shifts on sequence spaces, which is the subject of the next section and Chapter 3.

Exercise 1.3.1. Prove that $\lambda(E_m^{-1}([a, b])) = \lambda([a, b])$ for any interval $[a, b] \subset [0, 1]$.

Exercise 1.3.2. Prove that $E_k \circ E_l = E_l \circ E_k = E_{kl}$. When is $E_k \circ R_\alpha = R_\alpha \circ E_k$?

Exercise 1.3.3. Show that the set of points with dense orbits is uncountable.

Exercise 1.3.4. Prove that the set

$$C = \{x \in [0, 1]: E_3^k x \notin (1/3, 2/3) \forall k \in \mathbb{N}_0\}$$

is the standard middle-thirds Cantor set.

***Exercise 1.3.5.** Show that the set of points with dense orbits under E_m has Lebesgue measure 1.

1.4 Shifts and Subshifts

In this section, we generalize the notion of shift space introduced in the previous section. For an integer $m > 1$ set $\mathcal{A}_m = \{1, \dots, m\}$. We refer to \mathcal{A}_m as an *alphabet* and its elements as *symbols*. A finite sequence of symbols is called a *word*. Let $\Sigma_m = \mathcal{A}_m^{\mathbb{Z}}$ be the set of infinite two-sided sequences of symbols in \mathcal{A}_m , and $\Sigma_m^+ = \mathcal{A}_m^{\mathbb{N}}$ be the set of infinite one-sided sequences. We say that a sequence $x = (x_i)$ contains the word $w = w_1 w_2 \dots w_k$ (or that w *occurs* in x) if there is some j such that $w_i = x_{j+i}$ for $i = 1, \dots, k$.

Given a one-sided or two-sided sequence $x = (x_i)$, let $\sigma(x) = (\sigma(x)_i)$ be the sequence obtained by shifting x one step to the left, i.e., $\sigma(x)_i = x_{i+1}$. This defines a self-map of both Σ_m and Σ_m^+ called the *shift*. The pair (Σ_m, σ) is called the *full two-sided shift*; (Σ_m^+, σ) is the *full one-sided shift*. The two-sided shift is invertible. For a one-sided sequence, the leftmost symbol disappears, so the one-sided shift is non-invertible, and every point has m preimages. Both shifts have m^n periodic points of period n .

The shift spaces Σ_m and Σ_m^+ are compact topological spaces in the product topology. This topology has a basis consisting of *cylinders*

$$C_{j_1, \dots, j_k}^{n_1, \dots, n_k} = \{x = (x_l): x_{n_i} = j_i, i = 1, \dots, k\},$$

where $n_1 < n_2 < \dots < n_k$ are indices in \mathbb{Z} or \mathbb{N} , and $j_i \in \mathcal{A}_m$. Since the preimage of a cylinder is a cylinder, σ is continuous on Σ_m^+ and is a homeomorphism of Σ_m . The metric

$$d(x, x') = 2^{-l}, \quad \text{where } l = \min\{|i|: x_i \neq x'_i\}$$




Figure 1.1. Examples of directed graphs with labeled vertices and the corresponding adjacency matrices.

generates the product topology on Σ_m and Σ_m^+ (Exercise 1.4.3). In Σ_m , the open ball $B(x, 2^{-l})$ is the symmetric cylinder $C_{x_{-l}, -l+1, \dots, l}^{-l, -l+1, \dots, l}$, and in Σ_m^+ , $B(x, 2^{-l}) = C_{x_1, \dots, x_l}^{1, \dots, l}$. The shift is expanding on Σ_m^+ ; if $d(x, x') < 1/2$, then $d(\sigma(x), \sigma(x')) = 2d(x, x')$.

In the product topology, periodic points are dense, and there are dense orbits (Exercise 1.4.5).

Now we describe a natural class of closed shift-invariant subsets of the full shift spaces. These *subshifts* can be described in terms of *adjacency matrices* or their associated *directed graphs*. An adjacency matrix $A = (a_{ij})$ is an $m \times m$ matrix whose entries are zeros and ones. Associated to A is a directed graph Γ_A with m vertices such that a_{ij} is the number of edges from the i th vertex to the j th vertex. Conversely, if Γ is a finite directed graph with vertices v_1, \dots, v_m , then Γ determines an adjacency matrix B , and $\Gamma = \Gamma_B$. Figure 1.1 shows two adjacency matrices and the associated graphs.

Given an $m \times m$ adjacency matrix $A = (a_{ij})$, we say that a word or infinite sequence x (in the alphabet \mathcal{A}_m) is *allowed* if $a_{x_i x_{i+1}} > 0$ for every i ; equivalently, if there is a directed edge from x_i to x_{i+1} for every i . A word or sequence that is not allowed is said to be *forbidden*. Let $\Sigma_A \subset \Sigma_m$ be the set of allowed two-sided sequences (x_i) , and $\Sigma_A^+ \subset \Sigma_m^+$ be the set of allowed one-sided sequences. We can view a sequence $(x_i) \in \Sigma_A$ (or Σ_A^+) as an infinite walk along directed edges in the graph Γ_A , where x_i is the index of the vertex visited at time i . The sets Σ_A and Σ_A^+ are closed shift-invariant subsets of Σ_m and Σ_m^+ , and inherit the subspace topology. The pairs (Σ_A, σ) and (Σ_A^+, σ) are called the two-sided and one-sided *vertex shifts* determined by A .

A point $(x_i) \in \Sigma_A$ (or Σ_A^+) is periodic of period n if and only if $x_{i+n} = x_i$ for every i . The number of periodic points of period n (in Σ_A or Σ_A^+) is equal to the trace of A^n (Exercise 1.4.2).

Exercise 1.4.1. Let A be a matrix of zeros and ones. A vertex v_i can be *reached* (in n steps) from a vertex v_j if there is a path (consisting of n edges) from v_i to v_j along directed edges of Γ_A . What properties of A correspond to the following properties of Γ_A ?

- (a) Any vertex can be reached from some other vertex.
- (b) There are no terminal vertices, i.e., there is at least one directed edge starting at each vertex.
- (c) Any vertex can be reached in one step from any other vertex .
- (d) Any vertex can be reached from any other vertex in exactly n steps.

Exercise 1.4.2. Let A be an $m \times m$ matrix of zeros and ones. Prove that:

- (a) the number of fixed points in Σ_A (or Σ_A^+) is the trace of A ;
- (b) the number of allowed words of length $n + 1$ beginning with the symbol i and ending with j is the i, j th entry of A^n ; and
- (c) the number of periodic points of period n in Σ_A (or Σ_A^+) is the trace of A^n .

Exercise 1.4.3. Verify that the metrics on Σ_m and Σ_m^+ generate the product topology.

Exercise 1.4.4. Show that the semiconjugacy $\phi: \Sigma \rightarrow [0, 1]$ of §1.3 is continuous with respect to the product topology on Σ .

Exercise 1.4.5. Assume that all entries of some power of A are positive. Show that in the product topology on Σ_A and Σ_A^+ , periodic points are dense, and there are dense orbits.

1.5 Quadratic Maps

The expanding maps of the circle introduced in §1.3 are *linear maps* in the sense that they come from linear maps of the real line. The simplest non-linear dynamical systems in dimension one are the quadratic maps

$$q_\mu(x) = \mu x(1 - x), \quad \mu > 0.$$

Figure 1.2 shows the graph of q_3 and successive images $x_i = q_3^i(x_0)$ of a point x_0 .

If $\mu > 1$ and $x \notin [0, 1]$, then $q_\mu^n(x) \rightarrow -\infty$ as $n \rightarrow \infty$. For this reason, we focus our attention on the interval $[0, 1]$. For $\mu \in [0, 4]$, the interval $[0, 1]$ is forward invariant under q_μ . For $\mu > 4$, the interval $(1/2 - \sqrt{1/4 - 1/\mu}, 1/2 + \sqrt{1/4 - 1/\mu})$ maps outside $[0, 1]$; we show in Chapter 7 that the set of points Λ_μ whose forward orbits stay in $[0, 1]$ is a Cantor set, and (Λ_μ, q_μ) is equivalent to the full one-sided shift on two symbols.

Let X be a locally compact metric space and $f: X \rightarrow X$ a continuous map. A fixed point p of f is *attracting* if it has a neighborhood U such that \bar{U} is compact, $f(\bar{U}) \subset U$, and $\bigcap_{n \geq 0} f^n(U) = \{p\}$. A fixed point p is *repelling*




Figure 1.2. Quadratic map of q_3 .

if it has a neighborhood U such that $\bar{U} \subset f(U)$, and $\bigcap_{n \geq 0} f^{-n}(U) = \{p\}$. Note that if f is invertible, then p is attracting for f if and only if it is repelling for f^{-1} , and vice versa. A fixed point p is called *isolated* if there is a neighborhood of p that contains no other fixed points.

If x is a periodic point of f of period n , then we say that f is an *attracting (repelling) periodic point* if x is an attracting (repelling) fixed point of f^n . We also say that the periodic orbit $\mathcal{O}(x)$ is attracting or repelling, respectively.

The fixed points of q_μ are 0 and $1 - 1/\mu$. Note that $q'_\mu(0) = \mu$ and that $q'_\mu(1 - 1/\mu) = 2 - \mu$. Thus, 0 is attracting for $\mu < 1$ and repelling for $\mu > 1$, and $1 - 1/\mu$ is attracting for $\mu \in (1, 3)$ and repelling for $\mu \notin [1, 3]$ (Exercise 1.5.4).

The maps q_μ , $\mu > 4$, have interesting and complicated dynamical behavior. In particular, periodic points abound. For example,

$$\begin{aligned} q_\mu([1/\mu, 1/2]) &\supset [1 - 1/\mu, 1], \\ q_\mu([1 - 1/\mu, 1]) &\supset [0, 1 - 1/\mu] \supset [1/\mu, 1/2]. \end{aligned}$$

Hence, $q_\mu^2([1/\mu, 1/2]) \supset [1/\mu, 1/2]$, so the Intermediate Value Theorem implies that q_μ^2 has a fixed point $p_2 \in [1/\mu, 1/2]$. Thus, p_2 and $q_\mu(p_2)$ are non-fixed periodic points of period 2. This approach to showing existence of periodic points applies to many one-dimensional maps. We exploit this technique in Chapter 7 to prove the Sharkovsky Theorem (Theorem 7.3.1), which asserts, for example, that for continuous self-maps of the interval the existence of an orbit of period three implies the existence of periodic orbits of all orders.

Exercise 1.5.1. Show that for any $x \notin [0, 1]$, $q_\mu^n(x) \rightarrow -\infty$ as $n \rightarrow \infty$.

Exercise 1.5.2. Show that a repelling fixed point is an isolated fixed point.

Exercise 1.5.3. Suppose p is an attracting fixed point for f . Show that there is a neighborhood U of p such that the forward orbit of every point in U converges to p .

Exercise 1.5.4. Let $f: \mathbb{R} \rightarrow \mathbb{R}$ be a C^1 map, and p be a fixed point. Show that if $|f'(p)| < 1$, then p is attracting, and if $|f'(p)| > 1$, then p is repelling.

Exercise 1.5.5. Are 0 and $1 - 1/\mu$ attracting or repelling for $\mu = 1$? for $\mu = 3$?

Exercise 1.5.6. Show the existence of a non-fixed periodic point of q_μ of period 3, for $\mu > 4$.

Exercise 1.5.7. Is the period-2 orbit $\{p_2, q_\mu(p_2)\}$ attracting or repelling for $\mu > 4$?

1.6 The Gauss Transformation

Let $[x]$ denote the greatest integer less than or equal to x , for $x \in \mathbb{R}$. The map $\varphi: [0, 1] \rightarrow [0, 1]$ defined by

$$\varphi(x) = \begin{cases} 1/x - [1/x] & \text{if } x \in (0, 1], \\ 0 & \text{if } x = 0 \end{cases}$$

was studied by C. Gauss, and is now called the *Gauss transformation*. Note that φ maps each interval $(1/(n+1), 1/n]$ continuously and monotonically onto $[0, 1)$; it is discontinuous at $1/n$ for all $n \in \mathbb{N}$. Figure 1.3 shows the graph of φ .




Figure 1.3. Gauss transformation.

Gauss discovered a natural invariant measure μ for φ . The Gauss measure of an interval $A = (a, b)$ is

$$\mu(A) = \frac{1}{\log 2} \int_a^b \frac{dx}{1+x} = (\log 2)^{-1} \log \frac{1+b}{1+a}.$$

This measure is φ -invariant in the sense that $\mu(\varphi^{-1}(A)) = \mu(A)$ for any interval $A = (a, b)$. To prove invariance, note that the preimage of (a, b) consists of infinitely many intervals: In the interval $(1/(n+1), 1/n)$, the preimage is $(1/(n+b), 1/(n+a))$. Thus,

$$\begin{aligned} \mu(\varphi^{-1}((a, b))) &= \mu\left(\bigcup_{n=1}^{\infty} \left(\frac{1}{n+b}, \frac{1}{n+a}\right)\right) \\ &= \frac{1}{\log 2} \sum_{n=1}^{\infty} \log\left(\frac{n+a+1}{n+a} \cdot \frac{n+b}{n+b+1}\right) = \mu((a, b)). \end{aligned}$$

Note that in general $\mu(\varphi(A)) \neq \mu(A)$.

The Gauss transformation is closely related to continued fractions. The expression

$$[a_1, a_2, \dots, a_n] = \frac{1}{a_1 + \frac{1}{a_2 + \dots + \frac{1}{a_n}}}, \quad a_1, \dots, a_n \in \mathbb{N},$$

is called a *finite continued fraction*. For $x \in (0, 1]$, we have $x = 1/([\frac{1}{x}] + \varphi(x))$. More generally, if $\varphi^{n-1}(x) \neq 0$, set $a_i = [1/\varphi^{i-1}(x)] \geq 1$ for $i \leq n$. Then,

$$x = \frac{1}{a_1 + \frac{1}{a_2 + \frac{1}{\dots + \frac{1}{a_n + \varphi^n(x)}}}}$$

Note that x is rational if and only if $\varphi^m(x) = 0$ for some $m \in \mathbb{N}$ (Exercise 1.6.2). Thus any rational number is uniquely represented by a finite continued fraction.

For an irrational number $x \in (0, 1)$, the sequence of finite continued fractions

$$[a_1, a_2, \dots, a_n] = \frac{1}{a_1 + \frac{1}{a_2 + \frac{1}{\dots + \frac{1}{a_n}}}}$$

converges to x (where $a_i = [1/\phi^{i-1}(x)]$) (Exercise 1.6.4). This is expressed concisely with the infinite continued fraction notation

$$x = [a_1, a_2, \dots] = \frac{1}{a_1 + \frac{1}{a_2 + \dots}}.$$

Conversely, given a sequence $(b_i)_{i \in \mathbb{N}}$, $b_i \in \mathbb{N}$, the sequence $[b_1, b_2, \dots, b_n]$ converges, as $n \rightarrow \infty$, to a number $y \in [0, 1]$, and the representation $y = [b_1, b_2, \dots]$ is unique (Exercise 1.6.4). Hence $\varphi(y) = [b_2, b_3, \dots]$, because $b_n = [1/\phi^{n-1}(y)]$.

We summarize this discussion by saying that the continued fraction representation conjugates the Gauss transformation and the shift on the space of finite or infinite integer-valued sequences $(b_i)_{i=1}^\omega$, $\omega \in \mathbb{N} \cup \{\infty\}$, $b_i \in \mathbb{N}$. (By convention, the shift of a finite sequence is obtained by deleting the first term; the empty sequence represents 0.) As an immediate consequence, we obtain a description of the eventually periodic points of φ (see Exercise 1.6.3).

Exercise 1.6.1. What are the fixed points of the Gauss transformation?

Exercise 1.6.2. Show that $x \in [0, 1]$ is rational if and only if $\varphi^m(x) = 0$ for some $m \in \mathbb{N}$.

Exercise 1.6.3. Show that:

- (a) a number with periodic continued fraction expansion satisfies a quadratic equation with integer coefficients; and
- (b) a number with eventually periodic continued fraction expansion satisfies a quadratic equation with integer coefficients.

The converse of the second statement is also true, but is more difficult to prove [Arc70], [HW79].

***Exercise 1.6.4.** Show that given any infinite sequence $b_k \in \mathbb{N}$, $k = 1, 2, \dots$, the sequence $[b_1, \dots, b_n]$ of finite continued fractions converges. Show that for any $x \in \mathbb{R}$, the continued fraction $[a_1, a_2, \dots]$, $a_i = [1/\phi^{i-1}(x)]$, converges to x , and that this continued fraction representation is unique.

1.7 Hyperbolic Toral Automorphisms

Consider the linear map of \mathbb{R}^2 given by the matrix

$$A = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}.$$

The eigenvalues are $\lambda = (3 + \sqrt{5})/2 > 1$ and $1/\lambda$. The map expands by a factor of λ in the direction of the eigenvector $v_\lambda = ((1 + \sqrt{5})/2, 1)$, and contracts




Figure 1.4. The image of the torus under A .

by $1/\lambda$ in the direction of $v_{1/\lambda} = ((1 - \sqrt{5})/2, 1)$. The eigenvectors are perpendicular because A is symmetric.

Since A has integer entries, it preserves the integer lattice $\mathbb{Z}^2 \subset \mathbb{R}^2$ and induces a map (which we also call A) of the torus $\mathbb{T}^2 = \mathbb{R}^2/\mathbb{Z}^2$. The torus can be viewed as the unit square $[0, 1] \times [0, 1]$ with opposite sides identified: $(x_1, 0) \sim (x_1, 1)$ and $(0, x_2) \sim (1, x_2)$, $x_1, x_2 \in [0, 1]$. The map A is given in coordinates by

$$A \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} (2x_1 + x_2) \bmod 1 \\ (x_1 + x_2) \bmod 1 \end{pmatrix}$$

(see Figure 1.4). Note that \mathbb{T}^2 is a commutative group and A is an automorphism, since A^{-1} is also an integer matrix.

The periodic points of $A: \mathbb{T}^2 \rightarrow \mathbb{T}^2$ are the points with rational coordinates (Exercise 1.7.1).

The lines in \mathbb{R}^2 parallel to the eigenvector v_λ project to a family W^u of parallel lines on \mathbb{T}^2 . For $x \in \mathbb{T}^2$, the line $W^u(x)$ through x is called the *unstable manifold* of x . The family W^u partitions \mathbb{T}^2 and is called the *unstable foliation* of A . This foliation is invariant in the sense that $A(W^u(x)) = W^u(Ax)$. Moreover, A expands each line in W^u by a factor of λ . Similarly, the *stable foliation* W^s is obtained by projecting the family of lines in \mathbb{R}^2 parallel to $v_{1/\lambda}$. This foliation is also invariant under A , and A contracts each *stable manifold* $W^s(x)$ by $1/\lambda$. Since the slopes of v_λ and $v_{1/\lambda}$ are irrational, each of the stable and unstable manifolds is dense in \mathbb{T}^2 (Exercise 1.11.1).

In a similar way, any $n \times n$ integer matrix B induces a group endomorphism of the n -torus $\mathbb{T}^n = \mathbb{R}^n/\mathbb{Z}^n = [0, 1]^n/\sim$. The map is invertible (an

automorphism) if and only if B^{-1} is an integer matrix, which happens if and only if $|\det B| = 1$ (Exercise 1.7.2). If B is invertible and the eigenvalues do not lie on the unit circle, then $B: \mathbb{T}^n \rightarrow \mathbb{T}^n$ has expanding and contracting subspaces of complementary dimensions and is called a *hyperbolic toral automorphism*. The stable and unstable manifolds of a hyperbolic toral automorphism are dense in \mathbb{T}^n (§5.10). This is easy to show in dimension two (Exercise 1.7.3 and Exercise 1.11.1).

Hyperbolic toral automorphisms are prototypes of the more general class of *hyperbolic dynamical systems*. These systems have uniform expansion and contraction in complementary directions at every point. We discuss them in detail in Chapter 5.

Exercise 1.7.1. Consider the automorphism of \mathbb{T}^2 corresponding to a non-singular 2×2 integer matrix whose eigenvalues are not roots of 1.

- Prove that every point with rational coordinates is eventually periodic.
- Prove that every eventually periodic point has rational coordinates.

Exercise 1.7.2. Prove that the inverse of an $n \times n$ integer matrix B is also an integer matrix if and only if $|\det B| = 1$.

Exercise 1.7.3. Show that the eigenvalues of a two-dimensional hyperbolic toral automorphism are irrational (so the stable and unstable manifolds are dense by Exercise 1.11.1).

Exercise 1.7.4. Show that the number of fixed points of a hyperbolic toral automorphism A is $|\det(A - I)|$ (hence the number of periodic points of period n is $|\det(A^n - I)|$).

1.8 The Horseshoe

Consider a region $D \subset \mathbb{R}^2$ consisting of two semicircular regions D_1 and D_5 together with a unit square $R = D_2 \cup D_3 \cup D_4$ (see Figure 1.5).

Let $f: D \rightarrow D$ be a differentiable map that stretches and bends D into a horseshoe as shown in Figure 1.5. Assume also that f stretches $D_2 \cup D_4$ uniformly in the horizontal direction by a factor of $\mu > 2$ and contracts




Figure 1.5. The horseshoe map.




Figure 1.6. Horizontal rectangles.

uniformly in the vertical direction by $\lambda < 1/2$. Since $f(D_5) \subset D_5$, the Brouwer fixed point theorem implies the existence of a fixed point $p \in D_5$.

Set $R_0 = f(D_2) \cap R$ and $R_1 = f(D_4) \cap R$. Note that $f(R) \cap R = R_0 \cup R_1$. The set $f^2(R) \cap f(R) \cap R = f^2(R) \cap R$ consists of four horizontal rectangles R_{ij} , $i, j \in \{0, 1\}$, of height λ^2 (see Figure 1.6). More generally, for any finite sequence $\omega_0, \dots, \omega_n$ of zeros and ones,

$$R_{\omega_0\omega_1\dots\omega_n} = R_{\omega_0} \cap f(R_{\omega_1}) \cap \dots \cap f^n(R_{\omega_n})$$

is a horizontal rectangle of height λ^n , and $f^n(R) \cap R$ is the union of 2^n such rectangles. For an infinite sequence $\omega = (\omega_i) \in \{0, 1\}^{\mathbb{N}_0}$, let $R_\omega = \bigcap_{i=0}^\infty f^i(R_{\omega_i})$. The set $H^+ = \bigcap_{n=0}^\infty f^n(R) = \bigcup_\omega R_\omega$ is the product of a horizontal interval of length 1 and a vertical Cantor set C^+ (a Cantor set is a compact, perfect, totally disconnected set). Note that $f(H^+) = H^+$.

We now construct, in a similar way, a set H^- using preimages. Observe that $f^{-1}(R_0) = f^{-1}(R) \cap D_2$, and $f^{-1}(R_1) = f^{-1}(R) \cap D_4$ are vertical rectangles of width μ^{-1} . For any sequence $\omega_{-m}, \omega_{-m+1}, \dots, \omega_{-1}$ of zeros and ones, $\bigcap_{i=1}^m f^{-i}(R_{\omega_i})$ is a vertical rectangle of width μ^{-m} , and $H^- = \bigcap_{i=1}^\infty f^{-i}(R)$ is the product of a vertical interval (of length 1) and a horizontal Cantor set C^- .

The *horseshoe set* $H = H^+ \cap H^- = \bigcap_{i=-\infty}^\infty f^i(R)$ is the product of the Cantor sets C^- and C^+ and is closed and f -invariant. It is *locally maximal*, i.e., there is an open set U containing H such that any f -invariant subset of U containing H coincides with H (Exercise 1.8.2). The map $\phi: \Sigma_2 = \{0, 1\}^{\mathbb{Z}} \rightarrow H$ that assigns to each infinite sequence $\omega = (\omega_i) \in \Sigma_2$ the unique point $\phi(\omega) = \bigcap_{-\infty}^\infty f^i(R_{\omega_i})$ is a bijection (Exercise 1.8.3). Note that

$$f(\phi(\omega)) = \bigcap_{-\infty}^\infty f^{i+1}(R_{\omega_i}) = \phi(\sigma_r(\omega)),$$

where σ_r is the right shift in Σ_2 , $\sigma_r(\omega)_{i+1} = \omega_i$. Thus, ϕ conjugates $f|_H$ and the full two-sided 2-shift.

The *horseshoe* was introduced by S. Smale in the 1960s as an example of a *hyperbolic set* that “survives” small perturbations. We discuss hyperbolic sets in Chapter 5.

Exercise 1.8.1. Draw a picture of $f^{-1}(R) \cap f(R)$ and $f^{-2}(R) \cap f^2(R)$.

Exercise 1.8.2. Prove that H is a locally maximal f -invariant set.

Exercise 1.8.3. Prove that ϕ is a bijection, and that both ϕ and ϕ^{-1} are continuous.

1.9 The Solenoid

Consider the solid torus $\mathcal{T} = S^1 \times D^2$, where $S^1 = [0, 1] \bmod 1$ and $D^2 = \{(x, y) \in \mathbb{R}^2: x^2 + y^2 \leq 1\}$. Fix $\lambda \in (0, 1/2)$, and define $F: \mathcal{T} \rightarrow \mathcal{T}$ by

$$F(\phi, x, y) = (2\phi, \lambda x + \frac{1}{2} \cos 2\pi \phi, \lambda y + \frac{1}{2} \sin 2\pi \phi).$$

The map F stretches by a factor of 2 in the S^1 -direction, contracts by a factor of λ in the D^2 -direction, and wraps the image twice inside \mathcal{T} (see Figure 1.7).

The image $F(\mathcal{T})$ is contained in the interior $\text{int}(\mathcal{T})$ of \mathcal{T} , and $F^{n+1}(\mathcal{T}) \subset \text{int}(F^n(\mathcal{T}))$. Note that F is one-to-one (Exercise 1.9.1). A slice $F(\mathcal{T}) \cap \{\phi = \text{const}\}$ consists of two disks of radius λ centered at diametrically opposite points at distance $1/2$ from the center of the slice. A slice $F^n(\mathcal{T}) \cap \{\phi = \text{const}\}$




Figure 1.7. The solid torus and its image $F(\mathcal{T})$.

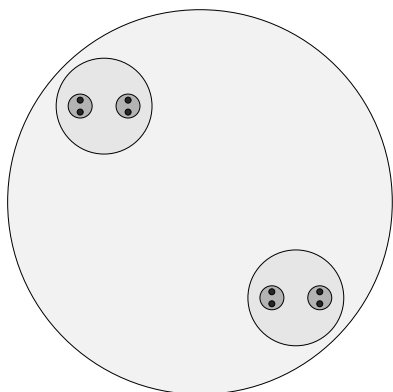


Figure 1.8. A cross-section of the solenoid.

consists of 2^n disks of radius λ^n : two disks inside each of the 2^{n-1} disks of $F^{n-1}(T) \cap \{\phi = \text{const}\}$. Slices of $F(T)$, $F^2(T)$, and $F^3(T)$ for $\phi = 1/8$ are shown in Figure 1.8.

The set $S = \bigcap_{n=0}^{\infty} F^n(T)$ is called a *solenoid*. It is a closed F -invariant subset of T on which F is bijective (Exercise 1.9.1). It can be shown that S is locally the product of an interval with a Cantor set in the two-dimensional disk.

The solenoid is an *attractor* for F . In fact, any neighborhood of S contains $F^n(T)$ for n sufficiently large, so the forward orbit of every point in T converges to S . Moreover, S is a hyperbolic set, and is therefore called a *hyperbolic attractor*. We give a precise definition of attractors in §1.13.

Let Φ denote the set of sequences $(\phi_i)_{i=0}^{\infty}$, where $\phi_i \in S^1$ and $\phi_i = 2\phi_{i+1} \pmod{1}$ for all i . The product topology on $(S^1)^{\mathbb{N}_0}$ induces the subspace topology on Φ . The space Φ is a commutative group under component-wise addition (mod 1). The map $(\phi, \psi) \mapsto \phi - \psi$ is continuous, so Φ is a topological group. The map $\alpha: \Phi \rightarrow \Phi$, $(\phi_0, \phi_1, \dots) \mapsto (2\phi_0, \phi_0, \phi_1, \dots)$ is a group automorphism and a homeomorphism (Exercise 1.9.3).

For $s \in S$, the first (angular) coordinates of the preimages $F^{-n}(s) = (\phi_n^-, x_n, y_n)$ form a sequence $h(s) = (\phi_0, \phi_1, \dots) \in \Phi$. This defines a map $h: S \rightarrow \Phi$. The inverse of h is the map $(\phi_0, \phi_1, \dots) \mapsto \bigcap_{n=0}^{\infty} F^n(\{\phi_n\} \times D^2)$, and h is a homeomorphism (Exercise 1.9.2). Note that $h: S \rightarrow \Phi$ conjugates F and α , i.e., $h \circ F = \alpha \circ h$. This conjugation allows one to study properties of (S, F) by studying properties of the algebraic system (Φ, α) .

Exercise 1.9.1. Prove that (a) $F: T \rightarrow T$ is injective, and (b) $F: S \rightarrow S$ is bijective.

Exercise 1.9.2. Prove that for every $(\phi_0, \phi_1, \dots) \in \Phi$ the intersection $\bigcap_{n=0}^{\infty} F^n(\{\phi_n\} \times D^2)$ consists of a single point s , and $h(s) = (\phi_0, \phi_1, \dots)$. Show that h is a homeomorphism.

Exercise 1.9.3. Show that Φ is a topological group, and α is an automorphism and homeomorphism.

Exercise 1.9.4. Find the fixed point of F and all periodic points of period 2.

1.10 Flows and Differential Equations

Flows arise naturally from first-order autonomous differential equations. Suppose $\dot{x} = F(x)$ is a differential equation in \mathbb{R}^n , where $F: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is continuously differentiable. For each point $x \in \mathbb{R}^n$, there is a unique solution $f^t(x)$ starting at x at time 0 and defined for all t in some neighborhood of 0. To simplify matters, we will assume that the solution is defined for all $t \in \mathbb{R}$; this will be the case, for example, if F is bounded, or is dominated in norm by a linear function. For fixed $t \in \mathbb{R}$, the time- t map $x \mapsto f^t(x)$ is a C^1 diffeomorphism of \mathbb{R}^n . Because the equation is autonomous, $f^{t+s}(x) = f^t(f^s(x))$, i.e., f^t is a flow.

Conversely, given a flow $f^t: \mathbb{R}^n \rightarrow \mathbb{R}^n$, if the map $(t, x) \mapsto f^t(x)$ is differentiable, then f^t is the time- t map of the differential equation

$$\dot{x} = \left. \frac{d}{dt} \right|_{t=0} f^t(x).$$

Here are some examples. Consider the linear autonomous differential equation $\dot{x} = Ax$ in \mathbb{R}^n , where A is a real $n \times n$ matrix. The flow of this differential equation is $f^t(x) = e^{At}x$, where e^{At} is the matrix exponential. If A is non-singular, the flow has exactly one fixed point at the origin. If all the eigenvalues of A have negative real part, then every orbit approaches the origin, and the origin is asymptotically stable. If some eigenvalue has positive real part, then the origin is unstable.

Most differential equations that arise in applications are non-linear. The differential equation governing an ideal frictionless pendulum is one of the most familiar:

$$\ddot{\theta} + \sin \theta = 0.$$

This equation cannot be solved in closed form, but it can be studied by qualitative methods. It is equivalent to the system

$$\begin{aligned} \dot{x} &= y, \\ \dot{y} &= -\sin x. \end{aligned}$$

The energy E of the system is the sum of the kinetic and potential energies, $E(x, y) = 1 - \cos x + y^2/2$. One can show (Exercise 1.10.2) by differentiating $E(x, y)$ with respect to t that E is constant along solutions of the differential equation. Equivalently, if f^t is the flow in \mathbb{R}^2 of this differential equation, then E is *invariant* by the flow, i.e., $E(f^t(x, y)) = E(x, y)$, for all $t \in \mathbb{R}$, $(x, y) \in \mathbb{R}^2$. A function that is constant on the orbits of a dynamical system is called a *first integral* of the system.

The fixed points in the phase plane for the undamped pendulum are $(k\pi, 0)$, $k \in \mathbb{Z}$. The points $(2k\pi, 0)$ are local minima of the energy. The points $(2(k+1)\pi, 0)$ are saddle points.

Now consider the damped pendulum $\ddot{\theta} + \gamma\dot{\theta} + \sin\theta = 0$, or the equivalent system

$$\begin{aligned}\dot{x} &= y, \\ \dot{y} &= -\sin x - \gamma y.\end{aligned}$$

A simple calculation shows that $\dot{E} < 0$ except at the fixed points $(k\pi, 0)$, $k \in \mathbb{Z}$, which are the local extrema of the energy. Thus the energy is strictly decreasing along every non-constant solution. In particular, every trajectory approaches a critical point of the energy, and almost every trajectory approaches a local minimum.

The energy of the pendulum is an example of a *Lyapunov function*, i.e., a continuous function that is non-increasing along the orbits of the flow. Any strict local minimum of a Lyapunov function is an asymptotically stable equilibrium point of the differential equation. Moreover, any bounded orbit must converge to the maximal invariant subset M of the set of points satisfying $\dot{E} = 0$. In the case of the damped pendulum, $M = \{(k\pi, 0) : k \in \mathbb{Z}\}$.

Here is another class of examples that appears frequently in applications, particularly optimization problems. Given a smooth function $f: \mathbb{R}^n \rightarrow \mathbb{R}$, the flow of the differential equation

$$\dot{x} = \text{grad } f(x)$$

is called the *gradient flow* of f . The function $-f$ is a Lyapunov function for the gradient flow. The trajectories are the projections to \mathbb{R}^n of paths of steepest ascent along the graph of f and are orthogonal to the level sets of f (Exercise 1.10.3).

A *Hamiltonian system* is a flow in \mathbb{R}^{2n} given by a system of differential equations of the form

$$\dot{q}_i = \frac{\partial H}{\partial p_i}, \quad \dot{p}_i = -\frac{\partial H}{\partial q_i}, \quad i = 1, \dots, n,$$

where the *Hamiltonian function* $H(p, q)$ is assumed to be smooth. Since the divergence of the right-hand side is 0, the flow preserves volume. The Hamiltonian function is a first integral, so that the level surfaces of H are invariant under the flow. If for some $C \in \mathbb{R}$ the level surface $H(p, q) = C$ is compact, the restriction of the flow to the level surface preserves a finite measure with smooth density. Hamiltonian flows have many applications in physics and mathematics. For example, the flow associated with the undamped pendulum is a Hamiltonian flow, where the Hamiltonian function is the total energy of the pendulum (Exercise 1.10.5).

Exercise 1.10.1. Show that the scalar differential equation $\dot{x} = x \log x$ induces the flow $f^t(x) = x^{\exp(t)}$ on the line.

Exercise 1.10.2. Show that the energy is constant along solutions of the undamped pendulum equation and strictly decreasing along non-constant solutions of the damped pendulum equation.

Exercise 1.10.3. Show that $-f$ is a Lyapunov function for the gradient flow of f , and that the trajectories are orthogonal to the level sets of f .

Exercise 1.10.4. Prove that any differentiable one-parameter group of linear maps of \mathbb{R} is the flow of a differential equation $\dot{x} = kx$.

Exercise 1.10.5. Show that the flow of the undamped pendulum is a Hamiltonian flow.

1.11 Suspension and Cross-Section

There are natural constructions for passing from a map to a flow, and vice versa. Given a map $f: X \rightarrow X$, and a function $c: X \rightarrow \mathbb{R}^+$ bounded away from 0, consider the quotient space

$$X_c = \{(x, t) \in X \times \mathbb{R}^+ : 0 \leq t \leq c(x)\} / \sim,$$

where \sim is the equivalence relation $(x, c(x)) \sim (f(x), 0)$. The *suspension* of f with *ceiling function* c is the semiflow $\phi^t: X_c \rightarrow X_c$ given by $\phi^t(x, s) = (f^n(x), s')$, where n and s' satisfy

$$\sum_{i=0}^{n-1} c(f^i(x)) + s' = t + s, \quad 0 \leq s' \leq c(f^n(x)).$$

In other words, flow along $\{x\} \times \mathbb{R}^+$ to $(x, c(x))$, then jump to $(f(x), 0)$ and continue along $\{f(x)\} \times \mathbb{R}^+$, and so on. See Figure 1.9. A suspension flow is also called a *flow under a function*.




Figure 1.9. Suspension and cross-section.

Conversely, a *cross-section* of a flow or semiflow $\psi^t: Y \rightarrow Y$ is a subset $A \subset Y$ with the following property: the set $T_y = \{t \in \mathbb{R}^+ : \psi^t(y) \in A\}$ is a non-empty discrete subset of \mathbb{R}^+ for every $y \in Y$. For $a \in A$, let $\tau(a) = \min T_a$ be the *return time* to A . Define the *first return map* $g: A \rightarrow A$ by $g(a) = \psi^{\tau(a)}(a)$, i.e., $g(a)$ is the first point after a in $\mathcal{O}_\psi^+(a) \cap A$ (see Figure 1.9). The first return map is often called the *Poincaré map*. Since the dimension of the cross-section is less by 1, in many cases maps in dimension n present the same level of difficulty as flows in dimension $n + 1$.

Suspension and cross-section are inverse constructions: the suspension of g with ceiling function τ is ψ^t , and $X \times \{0\}$ is a cross-section of ϕ with first return map f . If ϕ is a suspension of f , then the dynamical properties of f and ϕ are closely related, e.g., the periodic orbits of f correspond to the periodic orbits of ϕ . Both of these constructions can be tailored to specific settings (topological, measurable, smooth, etc.).

As an example, consider the 2-torus $\mathbb{T}^2 = \mathbb{R}^2/\mathbb{Z}^2 = S^1 \times S^1$, with topology and metric induced from the topology and metric on \mathbb{R}^2 . Fix $\alpha \in \mathbb{R}$, and define the linear flow $\phi_\alpha^t: \mathbb{T}^2 \rightarrow \mathbb{T}^2$ by

$$\phi_\alpha^t(x, y) = (x + \alpha t, y + t) \bmod 1.$$

Note that ϕ_α^t is the suspension of the circle rotation R_α with ceiling function 1, and $S^1 \times \{0\}$ is a cross-section for ϕ_α with constant return time $\tau(y) = 1$ and first return map R_α . The flow ϕ_α^t consists of left translations by the elements $g^t = (\alpha t, t) \bmod 1$, which form a one-parameter subgroup of \mathbb{T}^2 .

Exercise 1.11.1. Show that if α is irrational, then every orbit of ϕ_α is dense in T^2 , and if α is rational, then every orbit of ϕ_α is periodic.

Exercise 1.11.2. Let ϕ^t be a suspension of f . Show that a periodic orbit of ϕ^t corresponds to a periodic orbit of f , and that a dense orbit of ϕ^t corresponds to a dense orbit of f .

***Exercise 1.11.3.** Suppose $1, s,$ and αs are real numbers that are linearly independent over \mathbb{Q} . Show that every orbit of the time s map ϕ_α^s is dense in T^2 .

1.12 Chaos and Lyapunov Exponents

A dynamical system is *deterministic* in the sense that the evolution of the system is described by a specific map, so that the present (the initial state) completely determines the future (the forward orbit of the state). At the same time, dynamical systems often appear to be *chaotic* in that they have *sensitive dependence on initial conditions*, i.e., minor changes in the initial state lead to dramatically different long-term behavior. Specifically, a dynamical system (X, f) has sensitive dependence on initial conditions on a subset $X' \subset X$ if there is $\epsilon > 0$, such that for every $x \in X'$ and $\delta > 0$ there are $y \in X$ and $n \in \mathbb{N}$ for which $d(x, y) < \delta$ and $d(f^n(x), f^n(y)) > \epsilon$. Although there is no universal agreement on a definition of *chaos*, it is generally agreed that a chaotic dynamical system should exhibit sensitive dependence on initial conditions. Chaotic systems are usually assumed to have some additional properties, e.g., existence of a dense orbit.

The study of chaotic behavior has become one of the central issues in dynamical systems during the last two decades. In practice, the term *chaos* has been applied to a variety of systems that exhibit some type of random behavior. This random behavior is observed experimentally in some situations, and in others follows from specific properties of the system. Often a system is declared to be chaotic based on the observation that a typical orbit appears to be randomly distributed, and different orbits appear to be uncorrelated. The variety of views and approaches in this area precludes a universal definition of the word “chaos.”

The simplest example of a chaotic system is the circle endomorphism $(S^1, E_m), m > 1$ (§1.3). Distances between points x and y are expanded by a factor of m if $d(x, y) \leq 1/(2m)$, so any two points are moved at least $1/2m$ apart by some iterate of E_m , so E_m has sensitive dependence on initial conditions. A typical orbit is dense (§1.3) and is *uniformly distributed* on the circle (Proposition 4.4.2).

The simplest non-linear chaotic dynamical systems in dimension one are the quadratic maps $q_\mu(x) = \mu x(1 - x), \mu \geq 4$, restricted to the forward invariant set $\Lambda_\mu \subset [0, 1]$ (see §1.5 and Chapter 7).

Sensitive dependence on initial conditions is usually associated with positive *Lyapunov exponents*. Let f be a differentiable map of an open subset $U \subset \mathbb{R}^m$ into itself, and let $df(x)$ denote the derivative of f at x . For

$x \in U$ and a non-zero vector $v \in \mathbb{R}^m$ define the *Lyapunov exponent* $\chi(x, v)$ by

$$\chi(x, v) = \overline{\lim}_{n \rightarrow \infty} \frac{1}{n} \log \|df^n(x)v\|.$$

If f has uniformly bounded first derivatives, then χ is well defined for every $x \in U$ and every non-zero vector v .

The Lyapunov exponent measures the exponential growth rate of tangent vectors along orbits, and has the following properties:

$$\begin{aligned} \chi(x, \lambda v) &= \chi(x, v) \quad \text{for all real } \lambda \neq 0, \\ \chi(x, v + w) &\leq \max(\chi(x, v), \chi(x, w)), \\ \chi(f(x), df(x)v) &= \chi(x, v). \end{aligned} \tag{1.1}$$

See Exercise 1.12.1.

If $\chi(x, v) = \chi > 0$ for some vector v , then there is a sequence $n_j \rightarrow \infty$ such that for every $\eta > 0$

$$\|df^{n_j}(x)v\| \geq e^{(\chi - \eta)n_j} \|v\|.$$

This implies that, for a fixed j , there is a point $y \in U$ such that

$$d(f^{n_j}(x), f^{n_j}(y)) \geq \frac{1}{2} e^{(\chi - \eta)n_j} d(x, y).$$

In general, this does not imply sensitive dependence on initial conditions, since the distance between x and y cannot be controlled. However, most dynamical systems with positive Lyapunov exponents have sensitive dependence on initial conditions.

Conversely, if two close points are moved far apart by f^n , by the intermediate value theorem, there must exist points x and directions v for which $\|df^n(x)v\| > \|v\|$. Therefore, we expect f to have positive Lyapunov exponents if it has sensitive dependence on initial conditions, though this is not always the case.

The circle endomorphisms E_m , $m > 1$, have positive exponents at all points. A quadratic map q_μ , $\mu > 2 + \sqrt{5}$, has positive exponents at any point whose forward orbit does not contain 0.

Exercise 1.12.1. Prove (1.1).

Exercise 1.12.2. Compute the Lyapunov exponents for E_m .

Exercise 1.12.3. Compute the Lyapunov exponents for the solenoid, §1.9.

Exercise 1.12.4. Using a computer, calculate the first 100 points in the orbit of $\sqrt{2} - 1$ under the map E_2 . What proportion of these points is contained in each of the intervals $[0, \frac{1}{4})$, $[\frac{1}{4}, \frac{1}{2})$, $[\frac{1}{2}, \frac{3}{4})$, and $[\frac{3}{4}, 1)$?

1.13 Attractors

Let X be a compact topological space, and $f: X \rightarrow X$ be a continuous map. Generalizing the notion of an attracting fixed point, we say that a compact set $C \subset X$ is an *attractor* if there is an open set U containing C , such that $f(\bar{U}) \subset U$ and $C = \bigcap_{n \geq 0} f^n(U)$. It follows that $f(C) = C$, since $f(C) = \bigcap_{n \geq 1} f^n(U) \subset C$; on the other hand, $C = \bigcap_{n \geq 1} f^n(U) = f(C)$, since $f(U) \subset U$. Moreover, the forward orbit of any point $x \in U$ converges to C , i.e., for any open set V containing C , there is some $N > 0$ such that $f^n(x) \in V$ for all $n \geq N$. To see this, observe that X is covered by V together with the open sets $X \setminus f^n(\bar{U})$, $n \geq 0$. By compactness, there is a finite subcover, and since $f^n(U) \subset f^{n-1}(U)$, we conclude that there is some $N > 0$ such that $X = V \cup (X \setminus f^n(\bar{U}))$ for all $n \geq N$. Thus, $f^n(x) \in f^n(U) \subset V$ for $n \geq N$.

The *basin of attraction* of C is the set $\text{BA}(C) = \bigcup_{n \geq 0} f^{-n}(U)$. The basin $\text{BA}(C)$ is precisely the set of points whose forward orbits converge to C (Exercise 1.13.1).

An open set $U \subset X$ such that \bar{U} is compact and $f(\bar{U}) \subset U$ is called a *trapping region* for f . If U is a trapping region, then $\bigcap_{n \geq 0} f^n(U)$ is an attractor. For flows generated by differential equations, any region with the property that along the boundary the vector field points into the region is a trapping region for the flow. In practice, the existence of an attractor is proved by constructing a trapping region. An attractor can be studied experimentally by numerically approximating orbits that start in the trapping region.

The simplest examples of attractors are: the intersection of the images of the whole space (if the space is compact); attracting fixed points; and attracting periodic orbits. For flows, the examples include asymptotically stable fixed points and asymptotically stable periodic orbits.

Many dynamical systems have attractors of a more complicated nature. For example, recall that the solenoid S (§1.9) is a (hyperbolic) attractor for (\mathcal{T}, F) . Locally, S is the product of an interval with a Cantor set. The structure of hyperbolic attractors is relatively well understood. However, some nonlinear systems have attractors that are chaotic (with sensitive dependence on initial conditions) but not hyperbolic. These attractors are called *strange attractors*. The best-known examples of strange attractors are the Hénon attractor and the Lorenz attractor.

The study of strange attractors began with the publication by E. N. Lorenz in 1963 of the paper “Deterministic non-periodic flow” [Lor63]. In the process of investigating meteorological models, Lorenz studied the non-linear system of differential equations

$$\begin{aligned}\dot{x} &= \sigma(y - x), \\ \dot{y} &= Rx - y - xz, \\ \dot{z} &= -bz + xy,\end{aligned}\tag{1.2}$$

now called the *Lorenz system*. He observed that at parameter values $\sigma = 10$, $b = 8/3$, and $R = 28$, the solutions of (1.2) eventually start revolving alternately about two repelling equilibrium points at $(\pm\sqrt{72}, \pm\sqrt{72}, 27)$. The number of times the solution revolves about one equilibrium before switching to the other has no discernible pattern. There is a trapping region U that contains 0 but not the two repelling equilibrium points. The attractor contained in U is called the Lorenz attractor. It is an extremely complicated set consisting of uncountably many orbits (including a saddle fixed point at 0), and non-fixed periodic orbits that are known to be knotted [Wil84]. The attractor is not hyperbolic in the usual sense, though it has strong expansion and contraction and sensitive dependence on initial conditions. The attractor persists for small changes in the parameter values (see Figure 1.10).

The Hénon map $H = (f, g): \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is defined by

$$\begin{aligned}f(x, y) &= a - by - x^2, \\ g(x, y) &= x,\end{aligned}$$

where a and b are constants [Hén76]. The Jacobian of the derivative dH




Figure 1.10. Lorenz attractor.




Figure 1.11. Hénon attractor.

equals b . If $b \neq 0$, the Hénon map is invertible; the inverse is

$$(x, y) \mapsto (y, (a - x - y^2)/b).$$

The map changes area by a factor of $|b|$, and is orientation reversing if $b < 0$.

For the specific parameter values $a = 1.4$ and $b = -0.3$, Hénon showed that there is a trapping region U homeomorphic to a disk. His numerical experiments suggested that the resulting attractor has a dense orbit and sensitive dependence on initial conditions, though these properties have not been rigorously proved. Figure 1.11 shows a long segment of an orbit starting in the trapping region, which is believed to approximate the attractor. It is known that for a large set of parameter values $a \in [1, 2]$, $b \in [-1, 0]$, the attractor has a dense orbit and a positive Lyapunov exponent, but is not hyperbolic [BC91].

Exercise 1.13.1. Let A be an attractor. Show that $x \in B(A)$ if and only if the forward orbit of x converges to A .

Exercise 1.13.2. Find a trapping region for the flow generated by the Lorenz equations with parameter values $\sigma = 10$, $b = 8/3$, and $R = 28$.

Exercise 1.13.3. Find a trapping region for the Hénon map with parameter values $a = 1.4$, $b = -0.3$.

Exercise 1.13.4. Using a computer, plot the first 1000 points in an orbit of the Hénon map starting in a trapping region.

Topological Dynamics

A *topological dynamical system* is a topological space X and either a continuous map $f: X \rightarrow X$ or a continuous (semi)flow f^t on X , i.e., a (semi)flow f^t for which the map $(t, x) \mapsto f^t(x)$ is continuous. To simplify the exposition, we usually assume that X is locally compact, metrizable, and second countable, though many of the results in this chapter are true under weaker assumptions on X . As we noted earlier, we will focus our attention on discrete-time systems, though all general results in this chapter are valid for continuous-time systems as well.

Let X and Y be topological spaces. Recall that a continuous map $f: X \rightarrow Y$ is a *homeomorphism* if it is one-to-one and the inverse is continuous.

Let $f: X \rightarrow X$ and $g: Y \rightarrow Y$ be topological dynamical systems. A *topological semiconjugacy* from g to f is a surjective continuous map $h: Y \rightarrow X$ such that $f \circ h = h \circ g$. If h is a homeomorphism, it is called a *topological conjugacy*, and f and g are said to be *topologically conjugate* or *isomorphic*. Topologically conjugate dynamical systems have identical topological properties. Consequently, all the properties and invariants we introduce in this chapter, including minimality, topological transitivity, topological mixing, and topological entropy, are preserved by topological conjugacy.

Throughout this chapter, a metric space X with metric d is denoted (X, d) . If $x \in X$ and $r > 0$, then $B(x, r)$ denotes the open ball of radius r centered at x . If (X, d) and (Y, d') are metric spaces, then $f: X \rightarrow Y$ is an isometry if $d'(f(x_1), f(x_2)) = d(x_1, x_2)$ for all $x_1, x_2 \in X$.

2.1 Limit Sets and Recurrence

Let $f: X \rightarrow X$ be a topological dynamical system. Let x be a point in X . A point $y \in X$ is an *ω -limit point* of x if there is a sequence of natural numbers $n_k \rightarrow \infty$ (as $k \rightarrow \infty$) such that $f^{n_k}(x) \rightarrow y$. The *ω -limit set* of x is the

set $\omega(x) = \omega_f(x)$ of all ω -limit points of x . Equivalently,

$$\omega(x) = \bigcap_{n \in \mathbb{N}} \overline{\bigcup_{i \geq n} f^i(x)}.$$

If f is invertible, the α -limit set of x is $\alpha(x) = \alpha_f(x) = \bigcap_{n \in \mathbb{N}} \overline{\bigcup_{i \geq n} f^{-i}(x)}$. A point in $\alpha(x)$ is an α -limit point of x . Both the α - and ω -limit sets are closed and f -invariant (Exercise 2.1.1).

A point x is called (*positively*) *recurrent* if $x \in \omega(x)$; the set $\mathcal{R}(f)$ of recurrent points is f -invariant. Periodic points are recurrent.

A point x is *non-wandering* if for any neighborhood U of x there exists $n \in \mathbb{N}$ such that $f^n(U) \cap U \neq \emptyset$. The set $\text{NW}(f)$ of non-wandering points is closed, f -invariant, and contains $\omega(x)$ and $\alpha(x)$ for all $x \in X$ (Exercise 2.1.2). Every recurrent point is non-wandering, and in fact $\overline{\mathcal{R}(f)} \subset \text{NW}(f)$ (Exercise 2.1.3); in general, however, $\text{NW}(f) \not\subset \overline{\mathcal{R}(f)}$ (Exercise 2.1.11).

Recall the notation $\mathcal{O}(x) = \bigcup_{n \in \mathbb{Z}} f^n(x)$ for an invertible mapping f , and $\mathcal{O}^+(x) = \bigcup_{n \in \mathbb{N}_0} f^n(x)$.

PROPOSITION 2.1.1

1. Let f be a homeomorphism, $y \in \overline{\mathcal{O}(x)}$, and $z \in \overline{\mathcal{O}(y)}$. Then $z \in \overline{\mathcal{O}(x)}$.
2. Let f be a continuous map, $y \in \mathcal{O}^+(x)$, and $z \in \mathcal{O}^+(y)$. Then $z \in \mathcal{O}^+(x)$.

Proof. Exercise 2.1.7. □

Let X be compact. A closed, non-empty, forward f -invariant subset $Y \subset X$ is a *minimal set* for f if it contains no proper, closed, non-empty, forward f -invariant subset. A compact invariant set Y is minimal if and only if the forward orbit of every point in Y is dense in Y (Exercise 2.1.4). Note that a periodic orbit is a minimal set. If X itself is a minimal set, we say that f is minimal.

PROPOSITION 2.1.2. Let $f: X \rightarrow X$ be a topological dynamical system. If X is compact, then X contains a minimal set for f .

Proof. The proof is a straightforward application of Zorn's lemma. Let \mathcal{C} be the collection of non-empty, closed, f -invariant subsets of X , with the partial ordering given by inclusion. Then \mathcal{C} is not empty, since $X \in \mathcal{C}$. Suppose $\mathcal{K} \subset \mathcal{C}$ is a totally ordered subset. Then any finite intersection of elements of \mathcal{K} is non-empty, so by the finite intersection property for compact sets, $\bigcap_{K \in \mathcal{K}} K \neq \emptyset$. Thus, by Zorn's lemma, \mathcal{C} contains a minimal element, which is a minimal set for f . □

In a compact topological space, every point in a minimal set is recurrent (Exercise 2.1.4), so the existence of minimal sets implies the existence of recurrent points.

A subset $A \subset \mathbb{N}$ (or \mathbb{Z}) is *relatively dense* (or *syndetic*) if there is $k > 0$ such that $\{n, n + 1, \dots, n + k\} \cap A \neq \emptyset$ for any n . A point $x \in X$ is *almost periodic* if for any neighborhood U of x , the set $\{i \in \mathbb{N} : f^i(x) \in U\}$ is relatively dense in \mathbb{N} .

PROPOSITION 2.1.3. *If X is a compact Hausdorff space and $f: X \rightarrow X$ is continuous, then $\overline{\mathcal{O}^+(x)}$ is minimal for f if and only if x is almost periodic.*

Proof. Suppose x is almost periodic and $y \in \overline{\mathcal{O}^+(x)}$. We need to show that $x \in \overline{\mathcal{O}^+(y)}$. Let U be a neighborhood of x . There is an open set $U' \subset X$, $x \in U' \subset U$, and an open set $V \subset X \times X$ containing the diagonal, such that if $x_1 \in U'$ and $(x_1, x_2) \in V$, then $x_2 \in U$. Since x is almost periodic, there is $K \in \mathbb{N}$ such that for every $j \in \mathbb{N}$ we have that $f^{j+k}(x) \in U'$ for some $0 \leq k \leq K$. Let $V' = \bigcap_{i=0}^K f^{-i}(V)$. Note that V' is open and contains the diagonal of $X \times X$. There is a neighborhood $W \ni y$ such that $W \times W \subset V'$. Choose n such that $f^n(x) \in W$, and choose k such that $f^{n+k}(x) \in U'$ with $0 \leq k \leq K$. Then $(f^{n+k}(x), f^k(y)) \in V$, and hence $f^k(y) \in U$.

Conversely, suppose x is not almost periodic. Then there is a neighborhood U of x such that $A = \{i : f^i(x) \in U\}$ is not relatively dense. Thus, there are sequences $a_i \in \mathbb{N}$ and $k_i \in \mathbb{N}$, $k_i \rightarrow \infty$, such that $f^{a_i+j}(x) \notin U$ for $j = 0, \dots, k_i$. Let y be a limit point of $\{f^{a_i}(x)\}$. By passing to a subsequence, we may assume that $f^{a_i}(x) \rightarrow y$. Fix $j \in \mathbb{N}$. Note that $f^{a_i+j}(x) \rightarrow f^j(y)$, and $f^{a_i+j}(x) \notin U$ for i sufficiently large. Thus $f^j(y) \notin U$ for all $j \in \mathbb{N}$, so $x \notin \overline{\mathcal{O}^+(y)}$, which implies that $\overline{\mathcal{O}^+(x)}$ is not minimal. \square

Recall that an irrational circle rotation R_α is minimal (§1.2). Therefore every point is non-wandering, recurrent, and almost periodic. An expanding endomorphism E_m of the circle has dense orbits (§1.3), but is not minimal because it has periodic points. Every point is non-wandering, but not all points are recurrent (Exercise 2.1.5).

Exercise 2.1.1. Show that the α - and ω -limit sets of a point are closed invariant sets.

Exercise 2.1.2. Show that the set of non-wandering points is closed, is f -invariant, and contains $\omega(x)$ and $\alpha(x)$ for all $x \in X$.

Exercise 2.1.3. Show that $\overline{\mathcal{R}(f)} \subset \text{NW}(f)$.

Exercise 2.1.4. Let X be compact, $f: X \rightarrow X$ continuous.

- Show that $Y \subset X$ is minimal if and only if $\omega(y) = Y$ for every $y \in Y$.
- Show that Y is minimal if and only if the forward orbit of every point in Y is dense in Y .

Exercise 2.1.5. Show that there are points that are non-recurrent and not eventually periodic for an expanding circle endomorphism E_m .

Exercise 2.1.6. For a hyperbolic toral automorphism $A: \mathbb{T}^2 \rightarrow \mathbb{T}^2$, show that:

- (a) $\mathcal{R}(A)$ is dense, and hence $\text{NW}(A) = \mathbb{T}^2$, but
- (b) $\mathcal{R}(A) \neq \mathbb{T}^2$.

Exercise 2.1.7. Prove Proposition 2.1.1.

Exercise 2.1.8. Prove that a homeomorphism $f: X \rightarrow X$ is minimal if and only if for each non-empty open set $U \subset X$ there is $n \in \mathbb{N}$, such that $\bigcup_{k=-n}^n f^k(U) = X$.

Exercise 2.1.9. Prove that a homeomorphism f of a compact metric space X is minimal if and only if for every $\epsilon > 0$ there is $N = N(\epsilon) \in \mathbb{N}$, such that for every $x \in X$ the set $\{x, f(x), \dots, f^N(x)\}$ is ϵ -dense in X .

Exercise 2.1.10. Let $f: X \rightarrow X$ and $g: Y \rightarrow Y$ be continuous maps of compact metric spaces. Prove that $\overline{O_{f \times g}^+(x, y)} = \overline{O_f^+(x)} \times \overline{O_g^+(y)}$ if and only if $(x, g(y)) \in \overline{O_{f \times g}^+(x, y)}$.

Assume that f and g are minimal. Find necessary and sufficient conditions for $f \times g$ to be minimal.

***Exercise 2.1.11.** Give an example of a dynamical system where $\text{NW}(f) \not\subset \mathcal{R}(f)$.

2.2 Topological Transitivity

We assume throughout this section that X is second countable.

A topological dynamical system $f: X \rightarrow X$ is *topologically transitive* if there is a point $x \in X$ whose forward orbit is dense in X . If X has no isolated points, this condition is equivalent to the existence of a point whose ω -limit set is dense in X (Exercise 2.2.1).

PROPOSITION 2.2.1. *Let $f: X \rightarrow X$ be a continuous map of a locally compact Hausdorff space X . Suppose that for any two non-empty open sets U and V there is $n \in \mathbb{N}$ such that $f^n(U) \cap V \neq \emptyset$. Then f is topologically transitive.*

Proof. The hypothesis implies that given any open set $V \subset X$, the set $\bigcup_{n \in \mathbb{N}} f^{-n}(V)$ is dense in X , since it intersects every open set. Let $\{V_i\}$ be a countable basis for the topology of X . Then $Y = \bigcap_i \bigcup_{n \in \mathbb{N}} f^{-n}(V_i)$ is a countable intersection of open, dense sets and is therefore non-empty by

the Baire category theorem. The forward orbit of any point $y \in Y$ enters each V_i , hence is dense in X . \square

In most topological spaces, existence of a dense full orbit for a homeomorphism implies existence of a dense forward orbit, as we show in the next proposition. Note, however, that density of a particular full orbit $\mathcal{O}(x)$ does not imply density of the corresponding forward orbit $\mathcal{O}^+(x)$ (see Exercise 2.2.2).

PROPOSITION 2.2.2. *Let $f: X \rightarrow X$ be a homeomorphism of a compact metric space, and suppose that X has no isolated points. If there is a dense full orbit $\mathcal{O}(x)$, then there is a dense forward orbit $\mathcal{O}^+(y)$.*

Proof. Since $\overline{\mathcal{O}(x)} = X$, the orbit $\mathcal{O}(x)$ visits every non-empty open set U at least once, and therefore infinitely many times because X has no isolated points. Hence there is a sequence n_k , with $|n_k| \rightarrow \infty$, such that $f^{n_k}(x) \in B(x, 1/k)$ for $k \in \mathbb{N}$, i.e., $f^{n_k}(x) \rightarrow x$ as $k \rightarrow \infty$. Thus, $f^{n_k+l}(x) \rightarrow f^l(x)$ for any $l \in \mathbb{Z}$. There are either infinitely many positive or infinitely many negative indices n_k , and it follows that either $\mathcal{O}(x) \subset \overline{\mathcal{O}^+(x)}$ or $\mathcal{O}(x) \subset \overline{\mathcal{O}^-(x)}$. In the former case, $\overline{\mathcal{O}^+(x)} = X$, and we are done. In the latter case, let U, V be non-empty open sets. Since $\overline{\mathcal{O}^-(x)} = X$, there are integers $i < j < 0$ such that $f^i(x) \in U$ and $f^j(x) \in V$, so $f^{j-i}(U) \cap V \neq \emptyset$. Hence, by Proposition 2.2.1, f is topologically transitive. \square

Exercise 2.2.1. Show that if X has no isolated points and $\mathcal{O}^+(x)$ is dense, then $\omega(x)$ is dense. Give an example to show that this is not true if X has isolated points.

Exercise 2.2.2. Give an example of a dynamical system with a dense full orbit but no dense forward orbit.

Exercise 2.2.3. Is the product of two topologically transitive systems topologically transitive? Is a factor of a topologically transitive system topologically transitive?

Exercise 2.2.4. Let $f: X \rightarrow X$ be a homeomorphism. Show that if f has a non-constant first integral or Lyapunov function (§1.10), then it is not topologically transitive.

Exercise 2.2.5. Let $f: X \rightarrow X$ be a topological dynamical system with at least two orbits. Show that if f has an attracting periodic point, then it is not topologically transitive.

Exercise 2.2.6. Let α be irrational and $f: T^2 \rightarrow T^2$ be the homeomorphism of the 2-torus given by $f(x, y) = (x + \alpha, x + y)$.

- (a) Prove that every non-empty, open, f -invariant set is dense, i.e., f is topologically transitive.
- (b) Suppose the forward orbit of (x_0, y_0) is dense. Prove that for every $y \in S^1$ the forward orbit of (x_0, y) is dense. Moreover, if the set $\bigcup_{k=0}^n f^k(x_0, y_0)$ is ϵ -dense, then $\bigcup_{k=0}^n f^k(x_0, y)$ is ϵ -dense.
- (c) Prove that every forward orbit is dense, i.e., f is minimal.

2.3 Topological Mixing

A topological dynamical system $f: X \rightarrow X$ is *topologically mixing* if for any two non-empty open sets $U, V \subset X$, there is $N > 0$ such that $f^n(U) \cap V \neq \emptyset$ for $n \geq N$. Topological mixing implies topological transitivity by Proposition 2.2.1, but not vice versa. For example, an irrational circle rotation is minimal and therefore topologically transitive, but not topologically mixing (Exercise 2.3.1).

The following propositions establish topological mixing for some of the examples from Chapter 1.

PROPOSITION 2.3.1. *Any hyperbolic toral automorphism $A: \mathbb{T}^2 \rightarrow \mathbb{T}^2$ is topologically mixing.*

Proof. By Exercise 1.7.3, for each $x \in \mathbb{T}^2$, the unstable manifold $W^u(x)$ of A is dense in \mathbb{T}^2 . Thus for every $\epsilon > 0$, the collection of balls of radius ϵ centered at points of $W^u(x)$ covers \mathbb{T}^2 . By compactness, a finite subcollection of these balls also covers \mathbb{T}^2 . Hence, there is a bounded segment $S_0 \subset W^u(x)$ whose ϵ -neighborhood covers \mathbb{T}^2 . Since group translations of \mathbb{T}^2 are isometries, the ϵ -neighborhood of any translate $L_g S_0 = g + S_0 \subset W^u(g + x)$ covers \mathbb{T}^2 . To summarize: For every $\epsilon > 0$, there is $L(\epsilon) > 0$ such that every segment S of length $L(\epsilon)$ in an unstable manifold is ϵ -dense in \mathbb{T}^2 , i.e., $d(y, S) \leq \epsilon$ for every $y \in \mathbb{T}^2$.

Let U and V be non-empty open sets in \mathbb{T}^2 . Choose $y \in V$ and $\epsilon > 0$ such that $\bar{B}(y, \epsilon) \subset V$. The open set U contains a segment of length $\delta > 0$ in some unstable manifold $W^u(x)$. Let $\lambda, |\lambda| > 1$, be the expanding eigenvalue of A , and choose $N > 0$ such that $|\lambda|^N \delta \geq L(\epsilon)$. Then for any $n \geq N$, the image $A^n U$ contains a segment of length at least $L(\epsilon)$ in some unstable manifold, so $A^n U$ is ϵ -dense in \mathbb{T}^2 and therefore intersects V . \square

PROPOSITION 2.3.2. *The full two-sided shift (Σ_m, σ) and the full one-sided shift (Σ_m^+, σ) are topologically mixing.*

Proof. Recall from §1.4 that the topology on Σ_m has a basis consisting of open metric balls $B(\omega, 2^{-l}) = \{\omega' : \omega'_i = \omega_i, |i| \leq l\}$. Thus it suffices to

show that for any two balls $B(\omega, 2^{-l_1})$ and $B(\omega', 2^{-l_2})$, there is $N > 0$ such that $\sigma^n B(\omega, 2^{-l_1}) \cap B(\omega', 2^{-l_2}) \neq \emptyset$ for $n \geq N$. Elements of $\sigma^n B(\omega, 2^{-l_1})$ are sequences with specified values in the places $-n - l_1, \dots, -n + l_1$. Therefore the intersection is non-empty when $-n + l_1 < -l_2$, i.e., $n \geq N = l_1 + l_2 + 1$. This proves that (Σ_m, σ) is topologically mixing; the proof for (Σ_m^+, σ) is an exercise (Exercise 2.3.4). \square

COROLLARY 2.3.3. *The horseshoe (H, f) (§1.8) is topologically mixing.*

Proof. The horseshoe (H, f) is topologically conjugate to the full two-shift (Σ_2, σ) (see Exercise 1.8.3). \square

PROPOSITION 2.3.4. *The solenoid (S, F) is topologically mixing.*

Proof. Recall (Exercise 1.9.2) that (S, F) is topologically conjugate to (Φ, α) , where

$$\Phi = \{(\phi_i): \phi_i \in S^1, \phi_i = 2\phi_{i+1}, \forall i\} \subset \prod_{i=0}^{\infty} S^1 = \mathbb{T}^{\infty},$$

and $(\phi_0, \phi_1, \phi_2, \dots) \xrightarrow{\alpha} (2\phi_0, \phi_0, \phi_1, \dots)$. Thus, it suffices to show that (Φ, α) is topologically mixing. The topology in \mathbb{T}^{∞} has a basis consisting of open sets $\prod_{i=0}^{\infty} I_k$, where the I_j s are open in S^1 and all but finitely many are equal to S^1 . Let $U = (I_0 \times I_1 \times \dots \times I_k \times S^1 \times S^1 \times \dots) \cap \Phi$ and $V = (J_0 \times J_1 \times \dots \times J_l \times S^1 \times S^1 \times \dots) \cap \Phi$ be non-empty open sets from this basis. Choose $m > 0$ so that $2^m I_0 = S^1$. Then for $n > m + l$, the first $n - m$ components of

$$\alpha^n(U) = (2^n I_0 \times 2^{n-1} I_0 \times \dots \times I_0 \times I_1 \times \dots \times I_k \times S^1 \times S^1 \times \dots) \cap \Phi$$

are S^1 , so $\alpha^n(U) \cap V \neq \emptyset$. \square

Exercise 2.3.1. Show that a circle rotation is not topologically mixing. Show that an isometry is not topologically mixing if there is more than one point in the space.

Exercise 2.3.2. Show that expanding endomorphisms of S^1 are topologically mixing (see §1.3).

Exercise 2.3.3. Show that a factor of a topologically mixing system is also topologically mixing.

Exercise 2.3.4. Prove that (Σ_m^+, σ) is topologically mixing.

2.4 Expansiveness

A homeomorphism $f: X \rightarrow X$ is *expansive* if there is $\delta > 0$ such that for any two distinct points $x, y \in X$, there is some $n \in \mathbb{Z}$ such that $d(f^n(x), f^n(y)) \geq \delta$. A non-invertible continuous map $f: X \rightarrow X$ is *positively expansive* if there is $\delta > 0$ such that for any two distinct points $x, y \in X$, there is some $n \geq 0$ such that $d(f^n(x), f^n(y)) \geq \delta$. Any number $\delta > 0$ with this property is called an *expansiveness constant* for f .

Among the examples from Chapter 1, the following are expansive (or positively expansive): the circle endomorphisms E_m , $|m| \geq 2$; the full and one-sided shifts; the hyperbolic toral automorphisms; the horseshoe; and the solenoid (Exercise 2.4.2). For sufficiently large values of the parameter μ , the quadratic map q_μ is expansive on the invariant set Λ_μ . Circle rotations, group translations, and other equicontinuous homeomorphisms (see §2.7) are not expansive.

PROPOSITION 2.4.1. *Let f be a homeomorphism of an infinite compact metric space X . Then for every $\epsilon > 0$ there are distinct points $x_0, y_0 \in X$ such that $d(f^n(x_0), f^n(y_0)) \leq \epsilon$ for all $n \in \mathbb{N}_0$.*

Proof [Kin90]. Fix $\epsilon > 0$. Let E be the set of natural numbers m for which there is a pair $x, y \in X$ such that

$$d(x, y) \geq \epsilon \quad \text{and} \quad d(f^n(x), f^n(y)) \leq \epsilon \quad \text{for } n = 1, \dots, m. \quad (2.1)$$

Let $M = \sup E$ if $E \neq \emptyset$, and $M = 0$ if $E = \emptyset$.

If $M = \infty$, then for every $m \in \mathbb{N}$ there is a pair x_m, y_m satisfying (2.1). By compactness, there is a sequence $m_k \rightarrow \infty$ such that the limits

$$\lim_{k \rightarrow \infty} x_{m_k} = x', \quad \lim_{k \rightarrow \infty} y_{m_k} = y'$$

exist. By (2.1), $d(x', y') \geq \epsilon$ and, since f^j is continuous,

$$d(f^j(x'), f^j(y')) = \lim_{k \rightarrow \infty} d(f^j(x_{m_k}), f^j(y_{m_k})) \leq \epsilon$$

for every $j \in \mathbb{N}$. Thus, $x_0 = f(x')$, $y_0 = f(y')$ are the desired points.

Suppose now that M is finite. Since any finite collection of iterates of f is equicontinuous, there is $\delta > 0$ such that if $d(x, y) < \delta$, then $d(f^n(x), f^n(y)) < \epsilon$ for $0 \leq n \leq M$; the definition of M then implies that $d(f^{-1}(x), f^{-1}(y)) < \epsilon$. By induction, we conclude that $d(f^{-j}(x), f^{-j}(y)) < \epsilon$ for $j \in \mathbb{N}$ whenever $d(x, y) < \delta$. By compactness, there is a finite collection \mathcal{B} of open $\delta/2$ -balls that covers X . Let K be the cardinality of \mathcal{B} . Since X is infinite, we can choose a set $W \subset X$ consisting of $K + 1$ distinct points. By the pigeon-hole principle, for each $j \in \mathbb{Z}$, there are distinct points $a_j, b_j \in W$ such that $f^j(a_j)$

and $f^j(b_j)$ belong to the same ball $B_j \in \mathcal{B}$, so $d(f^j(a_j), f^j(b_j)) < \delta$. Thus, $d(f^n(a_j), f^n(b_j)) < \epsilon$ for $-\infty < n \leq j$. Since W is finite, there are distinct $x_0, y_0 \in W$ such that

$$a_j = x_0 \quad \text{and} \quad b_j = y_0$$

for infinitely many positive j and hence $d(f^n(x_0), f^n(y_0)) < \epsilon$ for all $n \geq 0$. \square

Proposition 2.4.1 is also true for non-invertible maps (Exercise 2.4.3).

COROLLARY 2.4.2. *Let f be an expansive homeomorphism of an infinite compact metric space X . Then there are $x_0, y_0 \in X$ such that $d(f^n(x_0), f^n(y_0)) \rightarrow 0$ as $n \rightarrow \infty$.*

Proof. Let $\delta > 0$ be an expansiveness constant for f . By Proposition 2.4.1, there are $x_0, y_0 \in X$ such that $d(f^n(x_0), f^n(y_0)) < \delta$ for all $n \in \mathbb{N}$. Suppose $d(f^n(x_0), f^n(y_0)) \not\rightarrow 0$. Then by compactness, there is a sequence $n_k \rightarrow \infty$ such that $f^{n_k}(x_0) \rightarrow x'$ and $f^{n_k}(y_0) \rightarrow y'$ with $x' \neq y'$. Then $f^{n_k+m}(x_0) \rightarrow f^m(x')$ and $f^{n_k+m}(y_0) \rightarrow f^m(y')$ for any $m \in \mathbb{Z}$. For k large, $n_k + m > 0$ and hence $d(f^m(x'), f^m(y')) \leq \delta$ for all $m \in \mathbb{Z}$, which contradicts expansiveness. \square

Exercise 2.4.1. Prove that every isometry of a compact metric space to itself is surjective and therefore is a homeomorphism.

Exercise 2.4.2. Show that the expanding circle endomorphisms E_m , $|m| \geq 2$, the full one- and two-sided shifts, the hyperbolic toral automorphisms, the horseshoe, and the solenoid are expansive, and compute expansiveness constants for each.

Exercise 2.4.3. Show that Proposition 2.4.1 is true for non-invertible continuous maps of infinite metric spaces.

2.5 Topological Entropy

Topological entropy is the exponential growth rate of the number of essentially different orbit segments of length n . It is a topological invariant that measures the complexity of the orbit structure of a dynamical system. Topological entropy is analogous to measure-theoretic entropy, which we introduce in Chapter 9.

Let (X, d) be a compact metric space, and $f: X \rightarrow X$ a continuous map. For each $n \in \mathbb{N}$, the function

$$d_n(x, y) = \max_{0 \leq k \leq n-1} d(f^k(x), f^k(y))$$

measures the maximum distance between the first n iterates of x and y . Each d_n is a metric on X , $d_n \geq d_{n-1}$, and $d_1 = d$. Moreover, the d_i are all equivalent metrics in the sense that they induce the same topology on X (Exercise 2.5.1).

Fix $\epsilon > 0$. A subset $A \subset X$ is (n, ϵ) -spanning if for every $x \in X$ there is $y \in A$ such that $d_n(x, y) < \epsilon$. By compactness, there are finite (n, ϵ) -spanning sets. Let $\text{span}(n, \epsilon, f)$ be the minimum cardinality of an (n, ϵ) -spanning set.

A subset $A \subset X$ is (n, ϵ) -separated if any two distinct points in A are at least ϵ apart in the metric d_n . Any (n, ϵ) -separated set is finite. Let $\text{sep}(n, \epsilon, f)$ be the maximum cardinality of an (n, ϵ) -separated set.

Let $\text{cov}(n, \epsilon, f)$ be the minimum cardinality of a covering of X by sets of d_n -diameter less than ϵ (the diameter of a set is the supremum of distances between pairs of points in the set). Again, by compactness, $\text{cov}(n, \epsilon, f)$ is finite.

The quantities $\text{span}(n, \epsilon, f)$, $\text{sep}(n, \epsilon, f)$, and $\text{cov}(n, \epsilon, f)$ count the number of orbit segments of length n that are distinguishable at scale ϵ . These quantities are related by the following lemma.

LEMMA 2.5.1. $\text{cov}(n, 2\epsilon, f) \leq \text{span}(n, \epsilon, f) \leq \text{sep}(n, \epsilon, f) \leq \text{cov}(n, \epsilon, f)$.

Proof. Suppose A is an (n, ϵ) -spanning set of minimum cardinality. Then the open balls of radius ϵ centered at the points of A cover X . By compactness, there exists $\epsilon_1 < \epsilon$ such that the balls of radius ϵ_1 centered at the points of A also cover X . Their diameter is $2\epsilon_1 < 2\epsilon$, so $\text{cov}(n, 2\epsilon, f) \leq \text{span}(n, \epsilon, f)$. The other inequalities are left as an exercise (Exercise 2.5.2). \square

Let

$$h_\epsilon(f) = \overline{\lim}_{n \rightarrow \infty} \frac{1}{n} \log(\text{cov}(n, \epsilon, f)). \quad (2.2)$$

The quantity $\text{cov}(n, \epsilon, f)$ increases monotonically as ϵ decreases, so $h_\epsilon(f)$ does as well. Thus the limit

$$h_{\text{top}} = h(f) = \lim_{\epsilon \rightarrow 0^+} h_\epsilon(f)$$

exists; it is called the *topological entropy* of f . The inequalities in Lemma 2.5.1 imply that equivalent definitions of $h(f)$ can be given using $\text{span}(n, \epsilon, f)$ or

$\text{sep}(n, \epsilon, f)$, i.e.,

$$h(f) = \lim_{\epsilon \rightarrow 0^+} \overline{\lim}_{n \rightarrow \infty} \frac{1}{n} \log(\text{span}(n, \epsilon, f)) \quad (2.3)$$

$$= \lim_{\epsilon \rightarrow 0^+} \overline{\lim}_{n \rightarrow \infty} \frac{1}{n} \log(\text{sep}(n, \epsilon, f)). \quad (2.4)$$

LEMMA 2.5.2. *The limit $\lim_{n \rightarrow \infty} \frac{1}{n} \log(\text{cov}(n, \epsilon, f)) = h_\epsilon(f)$ exists and is finite.*

Proof. Let U have d_m -diameter less than ϵ , and V have d_n -diameter less than ϵ . Then $U \cap f^{-m}(V)$ has d_{m+n} -diameter less than ϵ . Hence

$$\text{cov}(m+n, \epsilon, f) \leq \text{cov}(m, \epsilon, f) \cdot \text{cov}(n, \epsilon, f),$$

so the sequence $a_n = \log(\text{cov}(n, \epsilon, f)) \geq 0$ is subadditive. A standard lemma from calculus implies that a_n/n converges to a finite limit as $n \rightarrow \infty$ (Exercise 2.5.3). \square

It follows from Lemmas 2.5.1 and 2.5.2 that the lim sups in Formulas (2.2), (2.3), and (2.4) are finite. Moreover, the corresponding lim infs are finite, and

$$h(f) = \lim_{\epsilon \rightarrow 0^+} \underline{\lim}_{n \rightarrow \infty} \frac{1}{n} \log(\text{cov}(n, \epsilon, f)) \quad (2.5)$$

$$= \lim_{\epsilon \rightarrow 0^+} \underline{\lim}_{n \rightarrow \infty} \frac{1}{n} \log(\text{span}(n, \epsilon, f)) \quad (2.6)$$

$$= \lim_{\epsilon \rightarrow 0^+} \underline{\lim}_{n \rightarrow \infty} \frac{1}{n} \log(\text{sep}(n, \epsilon, f)). \quad (2.7)$$

The topological entropy is either $+\infty$ or a finite non-negative number. There are dramatic differences between dynamical systems with positive entropy and dynamical systems with zero entropy. Any isometry has zero topological entropy (Exercise 2.5.4). In the next section, we show that topological entropy is positive for several of the examples from Chapter 1.

PROPOSITION 2.5.3. *The topological entropy of a continuous map $f: X \rightarrow X$ does not depend on the choice of a particular metric generating the topology of X .*

Proof. Suppose d and d' are metrics generating the topology of X . For $\epsilon > 0$, let $\delta(\epsilon) = \sup\{d'(x, y) : d(x, y) \leq \epsilon\}$. By compactness, $\delta(\epsilon) \rightarrow 0$ as $\epsilon \rightarrow 0$. If U is a set of d_n -diameter less than ϵ , then U has d'_n -diameter at most $\delta(\epsilon)$. Thus $\text{cov}'(n, \delta(\epsilon), f) \leq \text{cov}(n, \epsilon, f)$, where cov and cov' correspond to the

metrics d and d' , respectively. Hence,

$$\lim_{\delta \rightarrow 0^+} \lim_{n \rightarrow \infty} \frac{1}{n} \log(\text{cov}(n, \delta, f)) \leq \lim_{\epsilon \rightarrow 0^+} \lim_{n \rightarrow \infty} \frac{1}{n} \log(\text{cov}(n, \epsilon, f)).$$

Interchanging d and d' gives the opposite inequality. \square

COROLLARY 2.5.4. *Topological entropy is an invariant of topological conjugacy.*

Proof. Suppose $f: X \rightarrow X$ and $g: Y \rightarrow Y$ are topologically conjugate dynamical systems, with conjugacy $\phi: Y \rightarrow X$. Let d be a metric on X . Then $d'(y_1, y_2) = d(\phi(y_1), \phi(y_2))$ is a metric on Y generating the topology of Y . Since ϕ is an isometry of (X, d) and (Y, d') , and the entropy is independent of the metric by Proposition 2.5.3, it follows that $h(f) = h(g)$. \square

PROPOSITION 2.5.5. *Let $f: X \rightarrow X$ be a continuous map of a compact metric space X .*

1. $h(f^m) = m \cdot h(f)$ for $m \in \mathbb{N}$.
2. If f is invertible, then $h(f^{-1}) = h(f)$. Thus $h(f^m) = |m| \cdot h(f)$ for all $m \in \mathbb{Z}$.
3. If $A_i, i = 1, \dots, k$ are closed (not necessarily disjoint) forward f -invariant subsets of X , whose union is X , then

$$h(f) = \max_{1 \leq i \leq k} h(f|A_i).$$

In particular, if A is a closed forward invariant subset of X , then $h(f|A) \leq h(f)$.

Proof. 1: Note that

$$\max_{0 \leq i < n} d(f^{mi}(x), f^{mi}(y)) \leq \max_{0 \leq j < mn} d(f^j(x), f^j(y)).$$

Thus, $\text{span}(n, \epsilon, f^m) \leq \text{span}(mn, \epsilon, f)$, so $h(f^m) \leq m \cdot h(f)$. Conversely, for $\epsilon > 0$, there is $\delta(\epsilon) > 0$ such that $d(x, y) < \delta(\epsilon)$ implies that $d(f^i(x), f^i(y)) < \epsilon$ for $i = 0, \dots, m$. Then $\text{span}(n, \delta(\epsilon), f^m) \geq \text{span}(mn, \epsilon, f)$, so $h(f^m) \geq m \cdot h(f)$.

2: The n th image of an (n, ϵ) -separated set for f is (n, ϵ) -separated for f^{-1} , and vice versa.

3: Any (n, ϵ) -separated set in A_i is (n, ϵ) -separated in X , so $h(f|A_i) \leq h(f)$. Conversely, the union of (n, ϵ) -spanning sets for the A_i s is an (n, ϵ) -spanning set for X . Thus if $\text{span}_i(n, \epsilon, f)$ is the minimum cardinality of an

(n, ϵ) -spanning subset of A_i , then

$$\text{span}(n, \epsilon, f) \leq \sum_{i=1}^k \text{span}_i(n, \epsilon, f) \leq k \cdot \max_{1 \leq i \leq k} \text{span}_i(n, \epsilon, f).$$

Therefore,

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n} \log(\text{span}(n, \epsilon, f)) &\leq \lim_{n \rightarrow \infty} \frac{1}{n} \log k + \lim_{n \rightarrow \infty} \frac{1}{n} \log \left(\max_{1 \leq i \leq k} \text{span}_i(n, \epsilon, f) \right) \\ &= 0 + \max_{1 \leq i \leq k} \lim_{n \rightarrow \infty} \frac{1}{n} \log(\text{span}_i(n, \epsilon, f)) \end{aligned}$$

The result follows by taking the limit as $\epsilon \rightarrow 0$. \square

PROPOSITION 2.5.6. *Let (X, d^X) and (Y, d^Y) be compact metric spaces, and $f: X \rightarrow X, g: Y \rightarrow Y$ continuous maps. Then:*

1. $h(f \times g) = h(f) + h(g)$; and
2. if g is a factor of f (or equivalently, f is an extension of g), then $h(f) \geq h(g)$.

Proof. To prove part 1, note that the metric $d((x, y), (x', y')) = \max\{d^X(x, x'), d^Y(y, y')\}$ generates the product topology on $X \times Y$, and

$$d_n((x, y), (x', y')) = \max\{d_n^X(x, x'), d_n^Y(y, y')\}.$$

If $U \subset X$ and $V \subset Y$ have diameters less than ϵ , then $U \times V$ has d -diameters less than ϵ . Hence

$$\text{cov}(n, \epsilon, f \times g) \leq \text{cov}(n, \epsilon, f) \cdot \text{cov}(n, \epsilon, g),$$

so $h(f \times g) \leq h(f) + h(g)$. On the other hand, if $A \subset X$ and $B \subset Y$ are (n, ϵ) -separated, then $A \times B$ is (n, ϵ) -separated for d . Hence

$$\text{sep}(n, \epsilon, f \times g) \geq \text{sep}(n, \epsilon, f) \cdot \text{sep}(n, \epsilon, g),$$

so, by (2.7), $h(f \times g) \geq h(f) + h(g)$.

The proof of part 2 is left as an exercise (Exercise 2.5.5). \square

PROPOSITION 2.5.7. *Let (X, d) be a compact metric space, and $f: X \rightarrow X$ an expansive homeomorphism with expansiveness constant δ . Then $h(f) = h_\epsilon(f)$ for any $\epsilon < \delta$.*

Proof. Fix γ and ϵ with $0 < \gamma < \epsilon < \delta$. We will show that $h_{2\gamma}(f) = h_\epsilon(f)$. By monotonicity, it suffices to show that $h_{2\gamma}(f) \leq h_\epsilon(f)$.

By expansiveness, for distinct points x and y , there is some $i \in \mathbb{Z}$ such that $d(f^i(x), f^i(y)) \geq \delta > \epsilon$. Since the set $\{(x, y) \in X \times X: d(x, y) \geq \gamma\}$ is compact, there is $k = k(\gamma, \epsilon) \in \mathbb{N}$ such that if $d(x, y) \geq \gamma$, then

$d(f^i(x), f^i(y)) > \epsilon$ for some $|i| \leq k$. Thus if A is an (n, γ) -separated set, then $f^{-k}(A)$ is $(n + 2k, \epsilon)$ -separated. Hence, by Lemma 2.5.1, $h_\epsilon(f) \geq h_{2\gamma}(f)$. \square

REMARK 2.5.8. *The topological entropy of a continuous (semi)flow can be defined as the entropy of the time-1 map. Alternatively, it can be defined using the analog d_T , $T > 0$, of the metrics d_n . The two definitions are equivalent because of the equicontinuity of the family of time- t maps, $t \in [0, 1]$.*

Exercise 2.5.1. Let (X, d) be a compact metric space. Show that the metrics d_i all induce the same topology on X .

Exercise 2.5.2. Prove the remaining inequalities in Lemma 2.5.1.

Exercise 2.5.3. Let $\{a_n\}$ be a subadditive sequence of non-negative real numbers, i.e., $0 \leq a_{m+n} \leq a_m + a_n$ for all $m, n \geq 0$. Show that $\lim_{n \rightarrow \infty} a_n/n = \inf_{n \geq 0} a_n/n$.

Exercise 2.5.4. Show that the topological entropy of an isometry is zero.

Exercise 2.5.5. Let $g: Y \rightarrow Y$ be a factor of $f: X \rightarrow X$. Prove that $h(f) \geq h(g)$.

Exercise 2.5.6. Let Y and Z be compact metric spaces, $X = Y \times Z$, and π be the projection to Y . Suppose $f: X \rightarrow X$ is an isometric extension of a continuous map $g: Y \rightarrow Y$, i.e., $\pi \circ f = g \circ \pi$ and $d(f(x_1), f(x_2)) = d((x_1), (x_2))$ for all $x_1, x_2 \in X$ with $\pi(x_1) = \pi(x_2)$. Prove that $h(f) = h(g)$.

Exercise 2.5.7. Prove that the topological entropy of a continuously differentiable map of a compact manifold is finite.

2.6 Topological Entropy for Some Examples

In this section, we compute the topological entropy for some of the examples from Chapter 1.

PROPOSITION 2.6.1. *Let \tilde{A} be a 2×2 integer matrix with determinant 1 and eigenvalues λ, λ^{-1} , with $|\lambda| > 1$; and let $A: \mathbb{T}^2 \rightarrow \mathbb{T}^2$ be the associated hyperbolic toral automorphism. Then $h(A) = \log |\lambda|$.*

Proof. The natural projection $\pi: \mathbb{R}^2 \rightarrow \mathbb{R}^2/\mathbb{Z}^2 = \mathbb{T}^2$ is a local homeomorphism, and $\pi \tilde{A} = A\pi$. Any metric \tilde{d} on \mathbb{R}^2 invariant under integer translations induces a metric d on \mathbb{T}^2 , where $d(x, y)$ is the \tilde{d} -distance between the sets $\pi^{-1}(x)$ and $\pi^{-1}(y)$. For these metrics, π is a local isometry.

Let v_1, v_2 be eigenvectors of A with (Euclidean) length 1 corresponding to the eigenvalues λ, λ^{-1} . For $x, y \in \mathbb{R}^2$, write $x - y = a_1 v_1 + a_2 v_2$ and

define $\tilde{d}(x, y) = \max(|a_1|, |a_2|)$. This is a translation-invariant metric on \mathbb{R}^2 . A \tilde{d} -ball of radius ϵ is a parallelogram whose sides are of (Euclidean) length 2ϵ and parallel to v_1 and v_2 . In the metric \tilde{d}_n (defined for \tilde{A}), a ball of radius ϵ is a parallelogram with side length $2\epsilon|\lambda|^{-n}$ in the v_1 -direction and 2ϵ in the v_2 -direction. In particular, the Euclidean area of a \tilde{d}_n -ball of radius ϵ is not greater than $4\epsilon^2|\lambda|^{-n}$. Since the induced metric d on \mathbb{T}^2 is locally isometric to \tilde{d} , we conclude that for sufficiently small ϵ , the Euclidean area of a d_n -ball of radius ϵ in \mathbb{T}^2 is at most $4\epsilon^2|\lambda|^{-n}$. It follows that the minimal number of balls of d_n -radius ϵ needed to cover \mathbb{T}^2 is at least

$$\text{area}(\mathbb{T}^2)/4\epsilon^2|\lambda|^{-n} = |\lambda|^n/4\epsilon^2.$$

Since a set of diameter ϵ is contained in an open ball of radius ϵ , we conclude that $\text{cov}(n, \epsilon, A) \geq |\lambda|^n/4\epsilon^2$. Thus, $h(A) \geq \log |\lambda|$.

Conversely, since the closed \tilde{d}_n -balls are parallelograms, there is a tiling of the plane by ϵ -balls whose interiors are disjoint. The Euclidean area of such a ball is $C\epsilon^2|\lambda|^{-n}$, where C depends on the angle between v_1 and v_2 . For small enough ϵ , any ϵ -ball that intersects the unit square $[0, 1] \times [0, 1]$ is entirely contained in the larger square $[-1, 2] \times [-1, 2]$. Therefore the number of the balls that intersect the unit square does not exceed the area of the larger square divided by the area of a \tilde{d}_n -ball of radius ϵ . Thus, the torus can be covered by $9\lambda^n/C\epsilon^2$ closed d_n -balls of radius ϵ . It follows that $\text{cov}(n, 2\epsilon, A) \leq 9\lambda^n/C\epsilon^2$, so $h(A) \leq \log |\lambda|$. \square

To establish the corresponding result in higher dimensions, we need some results from linear algebra. Let B be a $k \times k$ complex matrix. If λ is an eigenvalue of B , let

$$V_\lambda = \{v \in \mathbb{C}^k : (B - \lambda I)^i v = 0 \text{ for some } i \in \mathbb{N}\}.$$

If B is real and γ is a real eigenvalue, let

$$V_\gamma^{\mathbb{R}} = \mathbb{R}^k \cap V_\gamma = \{v \in \mathbb{R}^k : (B - \gamma I)^i v = 0 \text{ for some } i \in \mathbb{N}\}.$$

If B is real and $\lambda, \bar{\lambda}$ is a pair of complex eigenvalues, let

$$V_{\lambda, \bar{\lambda}}^{\mathbb{R}} = \mathbb{R}^k \cap (V_\lambda \oplus V_{\bar{\lambda}}).$$

These spaces are called *generalized eigenspaces*.

LEMMA 2.6.2. *Let B be a $k \times k$ complex matrix, and λ be an eigenvalue of B . Then for every $\delta > 0$ there is $C(\delta) > 0$ such that*

$$C(\delta)^{-1}(|\lambda| - \delta)^n \|v\| \leq \|B^n v\| \leq C(\delta)(|\lambda| + \delta)^n \|v\|$$

for every $n \in \mathbb{N}$ and every $v \in V_\lambda$.

Proof. It suffices to prove the lemma for a Jordan block. Thus without loss of generality, we assume that B has λ s on the diagonal, ones above and zeros elsewhere. In this setting, $V_\lambda = \mathbb{C}^k$ and in the standard basis e_1, \dots, e_k , we have $Be_1 = \lambda e_1$ and $Be_i = \lambda e_i + e_{i-1}$, for $i = 2, \dots, k$. For $\delta > 0$, consider the basis $e_1, \delta e_2, \delta^2 e_3, \dots, \delta^{k-1} e_k$. In this basis, the linear map B is represented by the matrix

$$B_\delta = \begin{pmatrix} \lambda & \delta & & & \\ & \lambda & \delta & & \\ & & \ddots & \ddots & \\ & & & \lambda & \delta \\ & & & & \lambda \end{pmatrix}.$$

Observe that $B_\delta = \lambda I + \delta A$ with $\|A\| \leq 1$, where $\|A\| = \sup_{v \neq 0} \|Av\|/\|v\|$. Therefore

$$(|\lambda| - \delta)^n \|v\| \leq \|B_\delta^n v\| \leq (|\lambda| + \delta)^n \|v\|.$$

Since B_δ is conjugate to B , there is a constant $C(\delta) > 0$ that bounds the distortion of the change of basis. \square

LEMMA 2.6.3. *Let B be a $k \times k$ real matrix and λ an eigenvalue of B . Then for every $\delta > 0$ there is $C(\delta) > 0$ such that*

$$C(\delta)^{-1} (|\lambda| - \delta)^n \|v\| \leq \|B^n v\| \leq C(\delta) (|\lambda| + \delta)^n \|v\|$$

for every $n \in \mathbb{N}$ and every $v \in V_\lambda$ (if $\lambda \in \mathbb{R}$) or every $v \in V_{\lambda, \bar{\lambda}}$ (if $\lambda \notin \mathbb{R}$).

Proof. If λ is real, then the result follows from Lemma 2.6.2. If λ is complex, then the estimates for V_λ and $V_{\bar{\lambda}}$ from Lemma 2.6.2 imply a similar estimate for $V_{\lambda, \bar{\lambda}}$, with a new constant $C(\delta)$ depending on the angle between V_λ and $V_{\bar{\lambda}}$ and the constants in the estimates for V_λ and $V_{\bar{\lambda}}$ (since $|\lambda| = |\bar{\lambda}|$). \square

PROPOSITION 2.6.4. *Let \tilde{A} be a $k \times k$ integer matrix with determinant 1 and with eigenvalues $\alpha_1, \dots, \alpha_k$, where*

$$|\alpha_1| \geq |\alpha_2| \geq \dots \geq |\alpha_m| > 1 > |\alpha_{m+1}| \geq \dots \geq |\alpha_k|.$$

Let $A: \mathbb{T}^k \rightarrow \mathbb{T}^k$ be the associated hyperbolic toral automorphism. Then

$$h(A) = \sum_{i=1}^m \log |\alpha_i|.$$

Proof. Let $\gamma_1, \dots, \gamma_j$ be the distinct real eigenvalues of \tilde{A} , and $\lambda_1, \bar{\lambda}_1, \dots, \lambda_m, \bar{\lambda}_m$ be the distinct complex eigenvalues of \tilde{A} . Then

$$\mathbb{R}^k = \bigoplus_{i=1}^j V_{\gamma_i} \oplus \bigoplus_{i=1}^m V_{\lambda_i, \bar{\lambda}_i},$$

any vector $v \in \mathbb{R}^k$ can be decomposed uniquely as $v = v_1 + \cdots + v_{j+m}$ with v_i in the corresponding generalized eigenspace. Given $x, y \in \mathbb{R}^k$, let $v = x - y$, and define $\bar{d}(x, y) = \max(|v_1|, \dots, |v_{j+m}|)$. This is a translation-invariant metric on \mathbb{R}^k , and therefore descends to a metric on \mathbb{T}^k . Now, using Lemma 2.6.3, the proposition follows by an argument similar to the one in the proof of Proposition 2.6.1 (Exercise 2.6.3). \square

The next example we consider is the solenoid from §1.9.

PROPOSITION 2.6.5. *The topological entropy of the solenoid map $F: S \rightarrow S$ is $\log 2$.*

Proof. Recall from §1.9 that F is topologically conjugate to the automorphism $\alpha: \Phi \rightarrow \Phi$, where

$$\Phi = \{(\phi_i)_{i=0}^\infty: \phi_i \in [0, 1), \phi_i = 2\phi_{i+1} \bmod 1\},$$

and α is coordinatewise multiplication by 2 (mod 1). Thus, $h(F) = h(\alpha)$. Let $|x - y|$ denote the distance on $S^1 = [0, 1] \bmod 1$. The distance function

$$d(\phi, \phi') = \sum_{n=0}^{\infty} \frac{1}{2^n} |\phi_n - \phi'_n|$$

generates the topology in Φ introduced in §1.9.

The map $\pi: \Phi \rightarrow S^1$, $(\phi_i)_{i=0}^\infty \mapsto \phi_0$, is a semiconjugacy from α to E_2 . Hence, $h(\alpha) \geq h(E_2) = \log 2$ (Exercise 2.6.1). We will establish the inequality $h(\alpha) \leq \log 2$ by constructing an (n, ϵ) -spanning set.

Fix $\epsilon > 0$ and choose $k \in \mathbb{N}$ such that $2^{-k} < \epsilon/2$. For $n \in \mathbb{N}$, let $A_n \subset \Phi$ consist of the 2^{n+2k} sequences $\psi^j = (\psi_i^j)$, where $\psi_i^j = j \cdot 2^{-(n+k+i)} \bmod 1$, $j = 0, \dots, 2^{n+2k} - 1$. We claim that A_n is (n, ϵ) -spanning. Let $\phi = (\phi_i)$ be a point in Φ . Choose $j \in \{0, \dots, 2^{n+2k} - 1\}$ so that $|\phi_k - j \cdot 2^{-(n+2k)}| \leq 2^{-(n+2k+1)}$. Then $|\phi_i - \psi_i^j| \leq 2^{k-i} 2^{-(n+2k+1)}$, for $0 \leq i \leq k$. It follows that for $0 \leq m \leq n$,

$$\begin{aligned} d(\alpha^m \phi, \alpha^m \psi^j) &= \sum_{i=0}^{\infty} \frac{|2^m \phi_i - 2^m \psi_i^j|}{2^i} < \sum_{i=0}^k \frac{2^m |\phi_i - \psi_i^j|}{2^i} + \frac{1}{2^k} \\ &< 2^m \sum_{i=0}^k \frac{2^{k-i} 2^{-(n+2k+1)}}{2^i} + \frac{1}{2^k} < \frac{1}{2^{k-1}} < \epsilon. \end{aligned}$$

Thus $d_n(\phi, \psi^j) < \epsilon$, so A_n is (n, ϵ) -spanning. Hence,

$$h(\alpha) \leq \lim_{n \rightarrow \infty} \frac{1}{n} \log \text{card } A_n = \log 2. \quad \square$$

Note that $\alpha: \Phi \rightarrow \Phi$ is expansive with expansiveness constant $1/3$ (Exercise 2.6.4), so by Proposition 2.5.7, $h_\epsilon(\alpha) = h(\alpha)$ for any $\epsilon < 1/3$.

Exercise 2.6.1. Compute the topological entropy of an expanding endomorphism $E_m: S^1 \rightarrow S^1$.

Exercise 2.6.2. Compute the topological entropy of the full one- and two-sided m -shifts.

Exercise 2.6.3. Finish the proof of Proposition 2.6.4.

Exercise 2.6.4. Prove that the solenoid map (§1.9) is expansive.

2.7 Equicontinuity, Distality, and Proximity¹

In this section, we describe a number of properties related to the asymptotic behavior of the distance between corresponding points on pairs of orbits.

Let $f: X \rightarrow X$ be a homeomorphism of a compact Hausdorff space. Points $x, y \in X$ are called *proximal* if the closure $\overline{\mathcal{O}((x, y))}$ of the orbit of (x, y) under $f \times f$ intersects the diagonal $\Delta = \{(z, z) \in X \times X: z \in X\}$. Every point is proximal to itself. If two points x and y are not proximal, i.e., if $\overline{\mathcal{O}((x, y))} \cap \Delta = \emptyset$, they are called *distal*. A homeomorphism $f: X \rightarrow X$ is *distal* if every pair of distinct points $x, y \in X$ is distal. If (X, d) is a compact metric space, then $x, y \in X$ are proximal if there is a sequence $n_k \in \mathbb{Z}$ such that $d(f^{n_k}(x), f^{n_k}(y)) \rightarrow 0$ as $k \rightarrow \infty$; points $x, y \in X$ are distal if there is $\epsilon > 0$ such that $d(f^n(x), f^n(y)) > \epsilon$ for all $n \in \mathbb{Z}$ (Exercise 2.7.2)

A homeomorphism f of a compact metric space (X, d) is said to be *equicontinuous* if the family of all iterates of f is an equicontinuous family, i.e., for any $\epsilon > 0$, there exists $\delta > 0$ such that $d(x, y) < \delta$ implies that $d(f^n(x), f^n(y)) < \epsilon$ for all $n \in \mathbb{Z}$. An isometry preserves distances and is therefore equicontinuous. Equicontinuous maps share many of the dynamical properties of isometries. The only examples from Chapter 1 that are equicontinuous are the group translations, including circle rotations.

We denote by $f \times f$ the induced action of f in $X \times X$, defined by $f \times f(x, y) = (f(x), f(y))$.

PROPOSITION 2.7.1. *An expansive homeomorphism of an infinite compact metric space is not distal.*

Proof. Exercise 2.7.1. □

¹ Several arguments in this section were conveyed to us by J. Auslander.

PROPOSITION 2.7.2. *Equicontinuous homeomorphisms are distal.*

Proof. Suppose the equicontinuous homeomorphism $f: X \rightarrow X$ is not distal. Then there is a pair of proximal points $x, y \in X$, so $d(f^{n_k}(x), f^{n_k}(y)) \rightarrow 0$ for some sequence $n_k \in \mathbb{Z}$. Let $x_k = f^{n_k}(x)$ and $y_k = f^{n_k}(y)$. Let $\epsilon = d(x, y)$. Then for any $\delta > 0$, there is some $k \in \mathbb{N}$ such that $d(x_k, y_k) < \delta$, but $d(f^{-n_k}(x_k), f^{-n_k}(y_k)) = \epsilon$, so f is not equicontinuous. \square

Distal homeomorphisms are not necessarily equicontinuous. Consider the map $F: \mathbb{T}^2 \rightarrow \mathbb{T}^2$ defined by

$$\begin{aligned} x &\mapsto x + \alpha \pmod{1}, \\ y &\mapsto x + y \pmod{1}. \end{aligned}$$

We view \mathbb{T}^2 as the unit square with opposite sides identified and use the metric inherited from the Euclidean metric. To see that this map is distal, let $(x, y), (x', y')$ be distinct points in \mathbb{T}^2 . If $x \neq x'$, then $d(F^n(x, y), F^n(x', y'))$ is at least $d((x, 0), (x', 0))$, which is constant. If $x = x'$, then $d(F^n(x, y), F^n(x', y')) = d((x, y), (x', y'))$. Therefore, the pair $(x, y), (x', y')$ is distal. To see that F is not equicontinuous, let $p = (0, 0)$ and $q = (\delta, 0)$. Then for all n , the difference between the first coordinates of $F^n(p)$ and $F^n(q)$ is δ . The difference between the second coordinates of $F^n(p)$ and $F^n(q)$ is $n\delta$ as long as $n\delta < 1/2$. Therefore there are points that are arbitrarily close together that are moved at least $1/4$ apart, so F is not equicontinuous.

The preceding map is an example of a *distal extension*. Suppose a homeomorphism $g: Y \rightarrow Y$ is an extension of a homeomorphism $f: X \rightarrow X$ with projection $\pi: Y \rightarrow X$. We say that the extension is *distal* if any pair of distinct points $y, y' \in Y$ with $\pi(y) = \pi(y')$ is distal. The map $F: \mathbb{T}^2 \rightarrow \mathbb{T}^2$ in the preceding paragraph is a distal extension of a circle rotation, with projection on the first factor as the factor map. A straightforward generalization of the argument in the previous paragraph shows that a distal extension of a distal homeomorphism is distal. Moreover, as we show later in this section, any factor of a distal map is distal. Thus, (X_1, f_1) and (X_2, f_2) are distal if and only if $(X_1 \times X_2, f_1 \times f_2)$ is distal.

Similarly, $\pi: Y \rightarrow X$ is an *isometric extension* if $d(g(y), g(y')) = d(y, y')$ whenever $\pi(y) = \pi(y')$. The extension $\pi: Y \rightarrow X$ is an *equicontinuous extension* if for any $\epsilon > 0$, there exists $\delta > 0$ such that if $\pi(y) = \pi(y')$ and $d(y, y') < \delta$, then $d(g^n(y), g^n(y')) < \epsilon$, for all n . An isometric extension is an equicontinuous extension; an equicontinuous extension is a distal extension.

To prove Theorem 2.7.4, we need the following notion: For a subset $A \subset X$ and a homeomorphism $f: X \rightarrow X$, denote by f_A the induced action of f in

the product space X^A (an element z of X^A is a function $z: A \rightarrow X$, and $f_A(z) = f \circ z$). We say that $A \subset X$ is *almost periodic* if every $z \in X^A$ with range A is an almost periodic point of (X^A, f_A) . That is, A is almost periodic if for every finite subset $a_1, \dots, a_n \in A$, and neighborhoods $U_1 \ni a_1, \dots, U_n \ni a_n$, the set $\{k \in \mathbb{Z}: f^k(a_i) \in U_i, 1 \leq i \leq n\}$ is syndetic in \mathbb{Z} . Every subset of an almost periodic set is an almost periodic set. Note that if x is an almost periodic point of f , then $\{x\}$ is an almost periodic set.

LEMMA 2.7.3. *Every almost periodic set is contained in a maximal almost periodic set.*

Proof. Let A be an almost periodic set, and \mathcal{C} be a collection, totally ordered by inclusion, of almost periodic sets containing A . The set $\bigcup_{C \in \mathcal{C}} C$ is an almost periodic set and a maximal element of \mathcal{C} . By Zorn's lemma there is a maximal almost periodic set containing A . \square

THEOREM 2.7.4. *Let f be a homeomorphism of a compact Hausdorff space X . Then every $x \in X$ is proximal to an almost periodic point.*

Proof. If x is an almost periodic point, then we are done, since x is proximal to itself. Suppose x is not almost periodic, and let A be a maximal almost periodic set. By definition, $x \notin A$. Let $z \in X^A$ have range A , and consider $(x, z) \in (X \times X^A)$. Let (x_0, z_0) be an almost periodic point (of $(f \times f_A)$) in $\mathcal{O}(x, z)$. Since z is almost periodic, $z \in \overline{\mathcal{O}(z_0)}$. Hence there is $x' \in X$ such that (x', z) is almost periodic and $(x', z) \in \overline{\mathcal{O}(x, z)}$ (Proposition 2.1.1). Therefore $\{x'\} \cup \text{range}(z) = \{x'\} \cup A$ is an almost periodic set. Since A is maximal, $x' \in A$, i.e., x' appears as one of the coordinates of z . It follows that $(x', x') \in \mathcal{O}(x, x')$, and x is proximal to x' . \square

A homeomorphism f of a compact Hausdorff space X is called *pointwise almost periodic* if every point is almost periodic. By Proposition 2.1.3, this happens if and only if X is a union of minimal sets.

PROPOSITION 2.7.5. *Let f be a distal homeomorphism of a compact Hausdorff space X . Then f is pointwise almost periodic.*

Proof. Let $x \in X$. Then, by Theorem 2.7.4, x is proximal to an almost periodic $y \in X$. Since f is distal, $x = y$ and x is almost periodic. \square

PROPOSITION 2.7.6. *A homeomorphism of a compact Hausdorff space is distal if and only if the product system $(X \times X, f \times f)$ is pointwise almost periodic.*

Proof. If f is distal, so is $f \times f$, and hence $f \times f$ is pointwise almost periodic. Conversely, assume that $f \times f$ is pointwise almost periodic, and let $x, y \in X$ be distinct points. If x and y are proximal, then there is z with

$(z, z) \in \overline{\mathcal{O}(x, y)}$. Recall that $\overline{\mathcal{O}(x, y)}$ is minimal (Proposition 2.1.3). Since $(x, y) \notin \overline{\mathcal{O}(z, z)}$, we obtain a contradiction. \square

COROLLARY 2.7.7. *A factor of a distal homeomorphism f of a compact Hausdorff space X is distal.*

Proof. Let $g: Y \rightarrow Y$ be a factor of f . Then $f \times f$ is pointwise almost periodic by Proposition 2.7.6. Since $(g \times g)$ is a factor of $f \times f$, it is pointwise almost periodic (Exercise 2.7.5), and hence is distal. \square

The class of distal dynamical systems is of special interest because it is closed under factors and isometric extensions. The class of minimal distal systems is the smallest such class of minimal systems: According to Furstenberg's structure theorem [Fur63], every minimal distal homeomorphism (or flow) can be obtained by a (possibly transfinite) sequence of isometric extensions starting with the one-point dynamical system.

Exercise 2.7.1. Prove Proposition 2.7.1.

Exercise 2.7.2. Prove the equivalence of the topological and metric definitions of distal and proximal points at the beginning of this section.

Exercise 2.7.3. Give an example of a homeomorphism f of a compact metric space (X, d) such that $d(f^n(x), f^n(y)) \rightarrow 0$ as $n \rightarrow \infty$ for every pair $x, y \in X$.

Exercise 2.7.4. Show that any infinite closed shift-invariant subset of Σ_m contains a proximal pair of points.

Exercise 2.7.5. Prove that a factor of a pointwise almost periodic system is pointwise almost periodic.

2.8 Applications of Topological Recurrence to Ramsey Theory²

In this section, we establish several Ramsey-type results to illustrate how topological dynamics is applied in combinatorial number theory. One of the main principles of the Ramsey theory is that a sufficiently rich structure is indestructible by finite partitioning (see [Ber96] for more information on Ramsey theory). An example of such a statement is van der Waerden's theorem, which we prove later in this section. We conclude this section by

² The exposition in this section follows, to a large extent, [Ber00].

proving a result in Ramsey theory about infinite-dimensional vector spaces over finite fields.

THEOREM 2.8.1 (van der Waerden). *For each finite partition $\mathbb{Z} = \bigcup_{k=1}^m S_k$, one of the sets S_k contains arbitrarily long (finite) arithmetic progressions.*

We will obtain van der Waerden’s theorem as a consequence of a general recurrence property in topological dynamics.

Recall from §1.4 that $\Sigma_m = \{1, 2, \dots, m\}^{\mathbb{Z}}$ with metric $d(\omega, \omega') = 2^{-k}$, where $k = \min\{|i|: \omega_i \neq \omega'_i\}$, is a compact metric space. The shift $\sigma: \Sigma_m \rightarrow \Sigma_m$, $(\sigma\omega)_i = \omega_{i+1}$, is a homeomorphism. A finite partition $\mathbb{Z} = \bigcup_{k=1}^m S_k$ can be viewed as a sequence $\xi \in \Sigma_m$ for which $\xi_i = k$ if $i \in S_k$. Let $X = \overline{\bigcup_{i=-\infty}^{\infty} \sigma^i \xi}$ be the orbit closure of ξ under σ , and let $A_k = \{\omega \in X: \omega_0 = k\}$. If $\omega \in A_k$, $\omega' \in X$, and $d(\omega', \omega) < 1$, then $\omega' \in A_k$. Hence if there are integers $p, q \in \mathbb{N}$ and $\omega \in X$ such that $d(\sigma^{ip}\omega, \omega) < 1$ for $0 \leq i \leq q - 1$, then there is $r \in \mathbb{Z}$ such that $\xi_j = \omega_0$ for $i = r, r + p, \dots, r + (q - 1)p$. Therefore, Theorem 2.8.1 follows from the following multiple recurrence property (Exercise 2.8.1).

PROPOSITION 2.8.2. *Let T be a homeomorphism of a compact metric space X . Then for every $\epsilon > 0$ and $q \in \mathbb{N}$ there are $p \in \mathbb{N}$ and $x \in X$ such that $d(T^{jp}(x), x) < \epsilon$ for $0 \leq j \leq q$.*

We will obtain Proposition 2.8.2 as a consequence of a more general statement (Theorem 2.8.3), which has other corollaries useful in combinatorial number theory.

Let \mathcal{F} be the collection of all finite non-empty subsets of \mathbb{N} . For $\alpha, \beta \in \mathcal{F}$, we write $\alpha < \beta$ if each element of α is less than each element of β . For a commutative group G , a map $T: \mathcal{F} \rightarrow G$, $\alpha \mapsto T_\alpha$, defines an *IP-system* in G if

$$T_{\{i_1, \dots, i_k\}} = T_{\{i_1\}} \cdot \dots \cdot T_{\{i_k\}}$$

for every $\{i_1, \dots, i_k\} \in \mathcal{F}$; in particular, if $\alpha, \beta \in \mathcal{F}$ and $\alpha \cap \beta = \emptyset$, then $T_{\alpha \cup \beta} = T_\alpha T_\beta$. Every IP-system T is generated by the elements $T_{\{n\}} \in G$, $n \in \mathbb{N}$.

Let G be a group of homeomorphisms of a topological space X . For $x \in X$, denote by Gx the orbit of x under G . We say that G *acts minimally on X* if for each $x \in X$, the orbit Gx is dense in X .

THEOREM 2.8.3 (Furstenberg–Weiss [FW78]). *Let G be a commutative group acting minimally on a compact topological space X . Then for every non-empty open set $U \subset X$, every $n \in \mathbb{N}$, every $\alpha \in \mathcal{F}$, and any IP-systems*

$T^{(1)}, \dots, T^{(n)}$ in G , there is $\beta \in \mathcal{F}$ such that $\alpha < \beta$ and

$$U \cap T_\beta^{(1)}(U) \cap \dots \cap T_\beta^{(n)}(U) \neq \emptyset.$$

Proof [Ber00]. Since G acts minimally, and X is compact, there are elements $g_1, \dots, g_m \in G$ such that $\bigcup_{i=1}^m g_i(U) = X$ (Exercise 2.8.2).

We argue by induction on n . For $n = 1$, let T be an IP-system and $U \subset X$ be open and not empty. Set $V_0 = U$. Define recursively $V_k = T_{\{k\}}(V_{k-1}) \cap g_{i_k}(U)$, where i_k is chosen so that $1 \leq i_k \leq m$ and $T_{\{k\}}(V_{k-1}) \cap g_{i_k}(U) \neq \emptyset$. By construction, $T_{\{k\}}^{-1}(V_k) \subset V_{k-1}$ and $V_k \subset g_{i_k}(U)$. In particular, by the pigeon-hole principle, there are $1 \leq i \leq m$ and arbitrarily large $p < q$ such that $V_p \cup V_q \subset g_i(U)$. Choose p so that $\beta = \{p+1, p+2, \dots, q\} > \alpha$. Then the set $W = g_i^{-1}(V_q) \subset U$ is not empty, and

$$T_\beta^{-1}(W) = g_i^{-1}(T_{\{p+1\}}^{-1} \dots T_{\{q\}}^{-1}(V_q)) \subset g_i^{-1}(T_{\{p+1\}}^{-1}(V_{p+1})) \subset g_i^{-1}(V_p) \subset U.$$

Therefore, $U \cap T_\beta(U) \supset W \neq \emptyset$.

Assume that the theorem holds true for any n IP-systems in G . Let $U \subset X$ be open and not empty. Let $T^{(1)}, \dots, T^{(n+1)}$ be IP-systems in G . We will construct a sequence of non-empty open subsets $V_k \subset X$ and an increasing sequence $\alpha_k \in \mathcal{F}$, $\alpha_k > \alpha$, such that $V_0 = U$, $\bigcup_{j=1}^{n+1} (T_{\alpha_k}^{(j)})^{-1}(V_k) \subset V_{k-1}$, and $V_k \subset g_{i_k}(U)$ for some $1 \leq i_k \leq m$.

By the inductive assumption applied to $V_0 = U$ and the n IP-systems $(T^{(n+1)})^{-1}T^{(j)}$, $j = 1, \dots, n$, there is $\alpha_1 > \alpha$ such that

$$V_0 \cap (T_{\alpha_1}^{(n+1)})^{-1}T_{\alpha_1}^{(1)}(V_0) \cap \dots \cap (T_{\alpha_1}^{(n+1)})^{-1}T_{\alpha_1}^{(n)}(V_0) \neq \emptyset.$$

Apply $T_{\alpha_1}^{(n+1)}$ and, for an appropriate $1 \leq i_1 \leq m$, set

$$V_1 = g_{i_1}(V_0) \cap T_{\alpha_1}^{(1)}(V_0) \cap T_{\alpha_1}^{(2)}(V_0) \cap \dots \cap T_{\alpha_1}^{(n+1)}(V_0) \neq \emptyset.$$

If V_{k-1} and α_{k-1} have been constructed, apply the inductive assumption to V_{k-1} and the IP-systems $(T^{(n+1)})^{-1}T^{(j)}$, $j = 1, \dots, n$, to get $\alpha_k > \alpha_{k-1}$ such that

$$V_{k-1} \cap (T_{\alpha_k}^{(n+1)})^{-1}T_{\alpha_k}^{(1)}(V_{k-1}) \cap \dots \cap (T_{\alpha_k}^{(n+1)})^{-1}T_{\alpha_k}^{(n)}(V_{k-1}) \neq \emptyset.$$

Apply $T_{\alpha_k}^{(n+1)}$ and, for an appropriate $1 \leq i_k \leq m$, set

$$V_k = g_{i_k}(V_0) \cap T_{\alpha_k}^{(1)}(V_{k-1}) \cap T_{\alpha_k}^{(2)}(V_{k-1}) \cap \dots \cap T_{\alpha_k}^{(n+1)}(V_{k-1}) \neq \emptyset.$$

By construction, the sequences α_k and V_k have the desired properties. Since $V_k \subset g_{i_k}(U)$, there is $1 \leq i \leq m$ such that $V_k \subset g_i(U)$ for infinitely many k 's. Hence there are arbitrarily large $p < q$ such that $V_p \cup V_q \subset g_i(U)$. Let

$W = g_i^{-1}(V_q) \subset U$ and $\beta = \alpha_{p+1} \cup \dots \cup \alpha_q$. Then $W \neq \emptyset$, and for each $1 \leq j \leq n+1$,

$$\begin{aligned} (T_\beta^{(j)})^{-1}(W) &= g_i^{-1}(T_{\alpha_{s+1}}^{(j)})^{-1}(V_q) \\ &\subset g_i^{-1}(T_{\alpha_{s+1}}^{(j)})^{-1}(V_{q-1}) \subset \dots \subset g_i^{-1}(V_p) \subset U. \end{aligned}$$

Therefore $\bigcup_{j=1}^{n+1} (T_\beta^{(j)})^{-1}W \subset U$, and hence $\bigcup_{j=1}^{n+1} (T_\beta^{(j)})^{-1}U \neq \emptyset$. \square

COROLLARY 2.8.4. *Let G be a commutative group of homeomorphisms of a compact metric space X and let $T^{(1)}, \dots, T^{(n)}$ be IP-systems in G . Then for every $\alpha \in \mathcal{F}$ and every $\epsilon > 0$ there are $x \in X$ and $\beta > \alpha$ such that $d(x, T_\beta^{(i)}(x)) < \epsilon$ for each $1 \leq i \leq n$.*

Proof. Similarly to Proposition 2.1.2, there is a non-empty closed G -invariant subset $X' \subset X$ on which G acts minimally (Exercise 2.8.3). Thus the corollary follows from Theorem 2.8.3. \square

Proof of Proposition 2.8.2. Let $G = \{T^k\}_{k \in \mathbb{Z}}$. For $\alpha \in \mathcal{F}$, denote by $|\alpha|$ the sum of the elements in α . Apply Corollary 2.8.4 to G , X , and the IP-systems $T_\alpha^{(j)} = T^{j|\alpha|}$, $1 \leq j \leq q-1$. \square

The following generalization of Theorem 2.8.1 also follows from Corollary 2.8.4.

THEOREM 2.8.5. *Let $d \in \mathbb{N}$, and let A be a finite subset of \mathbb{Z}^d . Then for each finite partition $\mathbb{Z}^d = \bigcup_{k=1}^m S_k$, there are $k \in \{1, \dots, m\}$, $z_0 \in \mathbb{Z}^d$, and $n \in \mathbb{N}$ such that $z_0 + na \in S_k$ for each $a \in A$, i.e., $z_0 + nA \subset S_k$.*

Proof. Exercise 2.8.5. \square

Let V_F be an infinite vector space over a finite field F . A subset $A \subset V_F$ is a d -dimensional affine subspace if there are $v \in V_F$ and linearly independent $x_1, \dots, x_d \in V_F$ such that $A = v + \text{Span}(x_1, \dots, x_d)$.

THEOREM 2.8.6 [GLR72], [GLR73]. *For each finite partition $V_F = \bigcup_{k=1}^m S_k$, one of the sets S_k contains affine subspaces of arbitrarily large (finite) dimension.*

Proof ([Ber00]; see Theorem 2.8.3). We say that a subset $L \subset V_F$ is *monochromatic of color j* if $L \subset S_j$.

Since V_F is infinite, it contains a countable subspace isomorphic to the abelian group

$$F_\infty = \{\mathbf{a} = (a_i)_{i=1}^\infty \in F^\mathbb{N} : a_i = 0 \text{ for all but finitely many } i \in \mathbb{N}\}.$$

Without loss of generality we assume that $V_F = F_\infty$. The set $\Omega = \{1, \dots, m\}^{F_\infty}$ of all functions $F_\infty \rightarrow \{1, \dots, m\}$ is naturally identified with the set of all partitions of F_∞ into m subsets. The discrete topology on $\{1, \dots, m\}$ and product topology on Ω make it a compact Hausdorff space.

Let $\xi \in \Omega$ correspond to a partition $F_\infty = \bigcup_{k=1}^m S_k$, i.e., $\xi: F_\infty \rightarrow \{1, \dots, m\}$, $\xi(\mathbf{a}) = k$ if and only if $\mathbf{a} \in S_k$. Each $\mathbf{b} \in F_\infty$ induces a homeomorphism $T_{\mathbf{b}}: \Omega \rightarrow \Omega$, $(T_{\mathbf{b}}\eta)(\mathbf{a}) = \eta(\mathbf{a} + \mathbf{b})$. Denote by $X \subset \Omega$ the orbit closure of ξ , $X = \overline{\{\bigcup_{\mathbf{b} \in F_\infty} T_{\mathbf{b}}\xi\}}$. Similarly to the argument in the proof of Proposition 2.1.2, Zorn's lemma implies that there is a non-empty closed subset $X' \subset X$ on which the group F_∞ acts minimally.

Let $g: \mathcal{F} \rightarrow F_\infty$ be an IP-system in F_∞ such that the elements $g_n, n \in \mathbb{N}$, are linearly independent. Define an IP-system T of homeomorphisms of X by setting $T_\alpha = T_{g_\alpha}$. For each $f \in F$, set $T_\alpha^{(f)} = T_{fg_\alpha}$ to get $|F| = \text{card } F$ IP-systems of commuting homeomorphisms of X . Let $\mathbf{0} = (0, 0, \dots)$ be the zero element of F_∞ and $A_i = \{\eta \in \Omega: \eta(\mathbf{0}) = i\}$. Then each A_i is open and $\bigcup_{i=1}^m A_i = \Omega$. Therefore, there is $j \in \{1, \dots, m\}$ such that $U = A_j \cap X' \neq \emptyset$. By Theorem 2.8.3, there is $\beta_1 \in \mathcal{F}$ such that $U_1 = \bigcap_{f \in F} T_{\beta_1}^{(f)}(U) \neq \emptyset$. If $\eta \in U_1$, then $\eta(fg_{\beta_1}) = j$ for each $f \in F$. In other words, η contains a monochromatic affine line of color j . Since the orbit of ξ is dense in X' , there is $\mathbf{b}_1 \in F_\infty$ such that $\xi(fg_{\beta_1} + \mathbf{b}_1) = \eta(fg_{\beta_1}) = j$. Thus, S_j contains an affine line.

To obtain a two-dimensional affine subspace in S_j apply Theorem 2.8.3 to U_1, β_1 and the same collection of IP-systems to get $\beta_2 > \beta_1$ such that $U_2 = \bigcap_{f \in F} T_{\beta_2}^{(f)}(U_1) \neq \emptyset$. Since g_{β_2} is linearly independent with every $g_\alpha, \alpha < \beta_2$, each $\eta \in U_2$ contains a monochromatic two-dimensional affine subspace of color j . Since η can be arbitrarily approximated by the shifts of ξ , the latter also contains a monochromatic two-dimensional affine subspace of color j .

Proceeding in this manner, we obtain a monochromatic subspace of arbitrarily large dimension. □

Exercise 2.8.1. Prove Theorem 2.8.1 using Proposition 2.8.2.

Exercise 2.8.2. Prove that a group G acts minimally on a compact topological space X if and only if for every non-empty open set $U \subset X$ there are elements $g_1, \dots, g_n \in G$ such that $\bigcup_{i=1}^n g_i(U) = X$.

Exercise 2.8.3. Prove the following generalization of Proposition 2.1.2. If a group G acts by homeomorphisms on a compact metric space X , then there is a non-empty closed G -invariant subset X' on which G acts minimally.

Exercise 2.8.4. Prove that van der Waerden's Theorem 2.8.1 is equivalent to the following finite version: For each $m, n \in \mathbb{N}$ there is $k \in \mathbb{N}$ such that if the set $\{1, 2, \dots, k\}$ is partitioned into m subsets, then one of them contains an arithmetic progression of length n .

***Exercise 2.8.5.** For $z \in \mathbb{Z}^d$, the translation by z in \mathbb{Z}^d induces a homeomorphism (shift) T_z in $\Sigma = \{1, \dots, m\}^{\mathbb{Z}^d}$. Prove Theorem 2.8.5 by considering the orbit closure under the group of shifts of the element $\xi \in \Sigma$ corresponding to the partition of \mathbb{Z}^d and the IP-systems in \mathbb{Z}^d generated by the translations $T_f, f \in A$.

Symbolic Dynamics

In §1.4, we introduced the symbolic dynamical systems (Σ_m, σ) and (Σ_m^+, σ) , and we showed by example throughout Chapter 1 how these shift spaces arise naturally in the study of other dynamical systems. In all of those examples, we encoded an orbit of the dynamical system by its *itinerary* through a finite collection of disjoint subsets. Specifically, following an idea that goes back to J. Hadamard, suppose $f: X \rightarrow X$ is a discrete dynamical system. Consider a partition $\mathcal{P} = \{P_1, P_2, \dots, P_m\}$ of X , i.e., $P_1 \cup P_2 \cup \dots \cup P_m = X$ and $P_i \cap P_j = \emptyset$ for $i \neq j$. For each $x \in X$, let $\psi_i(x)$ be the index of the element of \mathcal{P} containing $f^i(x)$. The sequence $(\psi_i(x))_{i \in \mathbb{N}_0}$ is called the *itinerary* of x . This defines a map

$$\psi: X \rightarrow \Sigma_m^+ = \{1, 2, \dots, m\}^{\mathbb{N}_0}, \quad x \mapsto \{\psi_i(x)\}_{i=0}^\infty,$$

which satisfies $\psi \circ f = \sigma \circ \psi$. The space Σ_m^+ is totally disconnected, and the map ψ usually is not continuous. If f is invertible, then positive and negative iterates of f define a similar map $X \rightarrow \Sigma_m = \{1, 2, \dots, m\}^{\mathbb{Z}}$. The image of ψ in Σ_m or Σ_m^+ is shift-invariant, and ψ semiconjugates f to the shift on the image of ψ . The indices $\psi_i(x)$ are *symbols* – hence the name *symbolic dynamics*. Any finite set can serve as the *symbol set*, or *alphabet*, of a symbolic dynamical system. Throughout this chapter, we identify every finite alphabet with $\{1, 2, \dots, m\}$.

Recall that the *cylinder* sets

$$C_{j_1, \dots, j_k}^{n_1, \dots, n_k} = \{\omega = (\omega_l) : \omega_{n_i} = j_i, i = 1, \dots, k\},$$

form a basis for the product topology of Σ_m and Σ_m^+ , and that the metric

$$d(\omega, \omega') = 2^{-l}, \quad \text{where } l = \min\{|i| : \omega_i \neq \omega'_i\}$$

generates the product topology.

3.1 Subshifts and Codes¹

In this section, we concentrate on two-sided shifts. The case of one-sided shifts is similar.

A *subshift* is a closed subset $X \subset \Sigma_m$ invariant under the shift σ and its inverse. We refer to Σ_m as the *full m -shift*.

Let $X_i \subset \Sigma_{m_i}$, $i = 1, 2$, be two subshifts. A continuous map $c: X_1 \rightarrow X_2$ is a *code* if it commutes with the shifts, i.e., $\sigma \circ c = c \circ \sigma$ (here and later, σ denotes the shift in any sequence space). Note that a surjective code is a factor map. An injective code is called an *embedding*; a bijective code gives a topological conjugacy of the subshifts and is called an *isomorphism* (since Σ_m is compact, a bijective code is a homeomorphism).

For a subshift $X \subset \Sigma_m$, denote by $W_n(X)$ the set of words of length n that occur in X , and by $|W_n(X)|$ its cardinality. Since different elements of X differ in at least one position, the restriction $\sigma|_X$ is expansive. Therefore, Proposition 2.5.7 allows us to compute the topological entropy of $\sigma|_X$ through the asymptotic growth rate of $|W_n(X)|$.

PROPOSITION 3.1.1. *Let $X \subset \Sigma_m$ be a subshift. Then*

$$h(\sigma|_X) = \lim_{n \rightarrow \infty} \frac{1}{n} \log |W_n(X)|.$$

Proof. Exercise 3.1.1. □

Let X be a subshift, $k, l \in \mathbb{N}_0$, $n = k + l + 1$, and let α be a map from $W_n(X)$ to an *alphabet* $\mathcal{A}_{m'}$. The (k, l) *block code* c_α from X to the full shift $\Sigma_{m'}$ assigns to a sequence $x = (x_i) \in X$ the sequence $c_\alpha(x)$ with $c_\alpha(x)_i = \alpha(x_{i-k}, \dots, x_i, \dots, x_{i+l})$. Any block code is a code, since it is continuous and commutes with the shift.

PROPOSITION 3.1.2 (Curtis–Lyndon–Hedlund). *Every code $c: X \rightarrow Y$ is a block code.*

Proof. Let \mathcal{A} be the symbol set of Y , and define $\tilde{\alpha}: X \rightarrow \mathcal{A}$ by $\tilde{\alpha}(x) = c(x)_0$. Since X is compact, $\tilde{\alpha}$ is uniformly continuous, so there is a $\delta > 0$ such that $\tilde{\alpha}(x) = \tilde{\alpha}(x')$ whenever $d(x, x') < \delta$. Choose $k \in \mathbb{N}$ so that $2^{-k} < \delta$. Then $\tilde{\alpha}(x)$ depends only on $x_{-k}, \dots, x_0, \dots, x_k$, and therefore defines a map $\alpha: W_{2k+1} \rightarrow \mathcal{A}$ satisfying $c(x)_0 = \alpha(x_{-k} \dots x_0 \dots x_k)$. Since c commutes with the shift, we conclude that $c = c_\alpha$. □

¹ The exposition of this section as well as §3.2, §3.4, and §3.5 follows in part the lectures of M. Boyle [Boy93].

There is a canonical class of block codes obtained by taking the alphabet of the target shift to be the set of words of length n in the original shift. Specifically, let $k, l \in \mathbb{N}, l < k$, and let X be a subshift. For $x \in X$ set

$$c(x)_i = x_{i-k+l+1} \dots x_i \dots x_{i+l}, \quad i \in \mathbb{Z}.$$

This defines a block code c from X to the full shift on the alphabet $W_k(X)$ which is an isomorphism onto its image (Exercise 3.1.2). Such a code (or sometimes its image) is called a *higher block presentation* of X .

Exercise 3.1.1. Prove Proposition 3.1.1.

Exercise 3.1.2. Prove that a higher block presentation of X is an isomorphism.

Exercise 3.1.3. Use a higher block presentation to prove that for any block code $c: X \rightarrow Y$, there is a subshift Z and an isomorphism $f: Z \rightarrow X$ such that $c \circ f: Z \rightarrow Y$ is a $(0, 0)$ block code.

Exercise 3.1.4. Show that the full shift has points whose full orbit is dense but whose forward orbit is nowhere dense.

3.2 Subshifts of Finite Type

The complement of a subshift $X \subset \Sigma_m$ is open and hence is a union of at most countably many cylinders. By shift invariance, if C is a cylinder and $C \subset \Sigma_m \setminus X$, then $\sigma^n(C) \subset \Sigma_m \setminus X$ for all $n \in \mathbb{Z}$, i.e., there is a countable list of forbidden words such that no sequence in X contains a forbidden word and each sequence in $\Sigma_m \setminus X$ contains at least one forbidden word. If there is a finite list of finite words such that X consists of precisely the sequences in Σ_m that do not contain any of these words, then X is called a *subshift of finite type* (SFT); X is a k -step SFT if it is defined by a set of words of length at most $k + 1$. A 1-step SFT is called a *topological Markov chain*.

In §1.4 we introduced a vertex shift Σ_A^v determined by an adjacency matrix A of zeros and ones. A *vertex shift* is an example of an SFT. The forbidden words have length 2 and are precisely those that are not allowed by A , i.e., a word uv is forbidden if there is no edge from u to v in the graph Γ_A determined by A . Since the list of forbidden words is finite, Σ_A^v is an SFT. A sequence in Σ_A^v can be viewed as an infinite path in the directed graph Γ_A , labeled by the vertices.

An infinite path in the graph Γ_A can also be specified by a sequence of edges (rather than vertices). This gives a subshift Σ_A^e whose alphabet is the set of edges in Γ_A . More generally, a finite directed graph Γ , possibly

with multiple directed edges connecting pairs of vertices, corresponds to an adjacency matrix B whose i, j th entry is a non-negative integer specifying the number of directed edges in $\Gamma = \Gamma_B$ from the i th vertex to the j th vertex. The set Σ_B^e of infinite directed paths in Γ_B , labeled by the edges, is closed and shift-invariant and is called the *edge shift* determined by B . Any edge shift is a subshift of finite type (Exercise 3.2.3).

For any matrix A of zeros and ones, the map $uv \mapsto e$, where e is the edge from u to v , defines a 2-block isomorphism from Σ_A^v to Σ_A^e . Conversely, any edge shift is naturally isomorphic to a vertex shift (Exercise 3.2.4).

PROPOSITION 3.2.1. *Every SFT is isomorphic to a vertex shift.*

Proof. Let X be a k -step SFT with $k > 0$. Let $W_k(X)$ be the set of words of length k that occur in X . Let Γ be the directed graph whose set of vertices is $W_k(X)$; a vertex $x_1 \dots x_k$ is connected to a vertex $x'_1 \dots x'_k$ by a directed edge if $x_1 \dots x_k x'_k = x_1 x'_1 \dots x'_k \in W_{k+1}(X)$. Let A be the adjacency matrix of Γ . The code $c(x)_i = x_i \dots x_{i+k-1}$ gives an isomorphism from X to Σ_A^v . \square

COROLLARY 3.2.2. *Every SFT is isomorphic to an edge shift.*

The last proposition implies that “the future is independent of the past” in an SFT; i.e., with appropriate one-step coding, if the sequences $\dots x_{-2}x_{-1}x_0$ and $x_0x_1x_2 \dots$ are allowed, then $\dots x_{-2}x_{-1}x_0x_1x_2 \dots$ is allowed.

Exercise 3.2.1. Show that the collection of all isomorphism classes of subshifts of finite type is countable.

***Exercise 3.2.2.** Show that the collection of all subshifts of Σ_2 is uncountable.

Exercise 3.2.3. Show that every edge shift is an SFT.

Exercise 3.2.4. Show that every edge shift is naturally isomorphic to a vertex shift. What are the vertices?

3.3 The Perron–Frobenius Theorem

The Perron–Frobenius Theorem guarantees the existence of special invariant measures, called *Markov measures*, for subshifts of finite type.

A vector or matrix all of whose coordinates are positive (non-negative) is called *positive (non-negative)*. Let A be a square non-negative matrix. If for any i, j there is $n \in \mathbb{N}$ such that $(A^n)_{ij} > 0$, then A is called *irreducible*; otherwise A is called *reducible*. If some power of A is positive, A is called *primitive*.

An integer non-negative square matrix A is primitive if and only if the directed graph Γ_A has the property that there is $n \in \mathbb{N}$ such that, for every pair of vertices u and v , there is a directed path from u to v of length n (see Exercise 1.4.2). An integer non-negative square matrix A is irreducible if and only if the directed graph Γ_A has the property that, for every pair of vertices u and v , there is a directed path from u to v (see Exercise 1.4.2).

A real non-negative $m \times m$ matrix is *stochastic* if the sum of the entries in each row is 1 or, equivalently, the column vector with all entries 1 is an eigenvector with eigenvalue 1.

THEOREM 3.3.1 (Perron). *Let A be a primitive $m \times m$ matrix. Then A has a positive eigenvalue λ with the following properties:*

1. λ is a simple root of the characteristic polynomial of A ,
2. λ has a positive eigenvector v ,
3. any other eigenvalue of A has modulus strictly less than λ ,
4. any non-negative eigenvector of A is a positive multiple of v .

Proof. Denote by $\text{int}(W)$ the interior of a set W . We will need the following lemma.

LEMMA 3.3.2. *Let $L: \mathbb{R}^k \rightarrow \mathbb{R}^k$ be a linear operator, and assume that there is a non-empty compact set P such that $0 \in \text{int}(P)$ and $L^i(P) \subset \text{int}(P)$ for some $i > 0$. Then the modulus of any eigenvalue of L is strictly less than 1.*

Proof. If the conclusion holds for L^i with some $i > 0$, then it holds for L . Hence we may assume that $L(P) \subset \text{int}(P)$. It follows that $L^n(P) \subset \text{int}(P)$ for all $n > 0$. The matrix L cannot have an eigenvalue of modulus greater than 1, since otherwise the iterates of L would move some vector in the open set $\text{int}(P)$ off to ∞ .

Suppose that σ is an eigenvalue of L and $|\sigma| = 1$. If $\sigma^j = 1$, then L^j has a fixed point on ∂P , a contradiction.

If σ is not a root of unity, there is a 2-dimensional subspace U on which L acts as an irrational rotation and any point $p \in \partial P \cap U$ is a limit point of $\bigcup_{n>0} L^n(P)$, a contradiction. \square

Since A is non-negative, it induces a continuous map f from the unit simplex $S = \{x \in \mathbb{R}^m : \sum x_j = 1, x_j \geq 0, j = 1, \dots, m\}$ into itself; $f(x)$ is the radial projection of Ax onto S . By the Brouwer fixed point theorem, there is a fixed point $v \in S$ of f , which is a non-negative eigenvector of A with eigenvalue $\lambda > 0$. Since some power of A is positive, all coordinates of v are positive.

Let V be the diagonal matrix that has the entries of v on the diagonal. The matrix $M = \lambda^{-1}V^{-1}AV$ is primitive, and the column vector $\mathbf{1}$ with all

entries 1 is an eigenvector of M with eigenvalue 1, i.e., M is a stochastic matrix. To prove parts 1 and 3, it suffices to show that 1 is a simple root of the characteristic polynomial of M and that all other eigenvalues of M have moduli strictly less than 1. Consider the action of M on row vectors. Since M is stochastic and non-negative, the row action preserves the unit simplex S . By the Brouwer fixed point theorem, there is a fixed row vector $w \in S$ all of whose coordinates are positive. Let $P = S - w$ be the translation of S by $-w$. Since for some $j > 0$ all entries of M^j are positive, $M^j(P) \subset \text{int}(P)$ and, by Lemma 3.3.2, the modulus of any eigenvalue of the row action of M in the $(m - 1)$ -dimensional invariant subspace spanned by P is strictly less than 1.

The last statement of the theorem follows from the fact that the codimension-one subspace spanned by P is M^t -invariant and its intersection with the cone of non-negative vectors in \mathbb{R}^n is $\{0\}$. \square

COROLLARY 3.3.3. *Let A be a primitive stochastic matrix. Then 1 is a simple root of the characteristic polynomial of A , both A and the transpose of A have positive eigenvectors with eigenvalue 1, and any other eigenvalue of A has modulus strictly less than 1.*

Frobenius extended Theorem 3.3.1 to irreducible matrices.

THEOREM 3.3.4 (Frobenius). *Let A be a non-negative irreducible square matrix. Then there exists an eigenvalue λ of A with the following properties: (i) $\lambda > 0$, (ii) λ is a simple root of the characteristic polynomial, (iii) λ has a positive eigenvector, (iv) if μ is any other eigenvalue of A , then $|\mu| \leq \lambda$, (v) if k is the number of eigenvalues of modulus $|\lambda|$, then the spectrum of A (with multiplicity) is invariant under the rotation of the complex plane by angle $2\pi/k$.*

A proof of Theorem 3.3.4 is outlined in Exercise 3.3.3. A complete argument can be found in [Gan59] or [BP94].

Exercise 3.3.1. Show that if A is a primitive integral matrix, then the edge shift Σ_A^e is topologically mixing.

Exercise 3.3.2. Show that if A is an irreducible integral matrix, then the edge shift Σ_A^e is topologically transitive.

Exercise 3.3.3. This exercise outlines the main steps in the proof of Theorem 3.3.4. Let A be a non-negative irreducible matrix, and let B be the matrix with entries $b_{ij} = 0$ if $a_{ij} = 0$ and $b_{ij} = 1$ if $a_{ij} > 0$. Let Γ be the graph whose adjacency matrix is B . For a vertex v in Γ , let $d = d(v)$ be the greatest common divisor of the lengths of closed paths in Γ starting from v . Let $V_k, k = 0, 1, \dots, d - 1$, be the set of vertices of Γ that can be connected

to v by a path whose length is congruent to $k \pmod d$.

- (a) Prove that d does not depend on v .
- (b) Prove that any path of length l starting in V_k ends in V_m with m congruent to $k + l \pmod d$.
- (c) Prove that there is a permutation of the vertices that conjugates B^d to a block-diagonal matrix with square blocks B_k , $k = 0, 1, \dots, d - 1$, along the diagonal and zeros elsewhere, each B_k being a primitive matrix whose size equals the cardinality of V_k .
- (d) What are the implications for the spectrum of A ?
- (e) Deduce Theorem 3.3.4.

3.4 Topological Entropy and the Zeta Function of an SFT

For an edge or vertex shift, dynamical invariants can be computed from the adjacency matrix. In this section, we compute the topological entropy of an edge shift and introduce the zeta function, an invariant that collects combinatorial information about the periodic points.

PROPOSITION 3.4.1. *Let A be a square non-negative integer matrix. Then the topological entropy of the edge shift Σ_A^e and the vertex shift Σ_A^v equals the logarithm of the largest eigenvalue of A .*

Proof. We consider only the edge shift. By Proposition 3.1.1, it suffices to compute the cardinality of $W_n(\Sigma_A)$ (the words of length n in Σ_A), which is the sum S_n of all entries of A^n (Exercise 1.4.2). The proposition now follows from Exercise 3.4.1. \square

For a discrete dynamical system f , denote by $\text{Fix}(f)$ the set of fixed points of f and by $|\text{Fix}(f)|$ its cardinality. If $|\text{Fix}(f^n)|$ is finite for every n , we define the *zeta function* $\zeta_f(z)$ of f to be the formal power series

$$\zeta_f(z) = \exp \sum_{n=1}^{\infty} \frac{1}{n} |\text{Fix}(f^n)| z^n.$$

The zeta function can also be expressed by the *product formula*:

$$\zeta_f(z) = \prod_{\gamma} (1 - z^{|\gamma|})^{-1},$$

where the product is taken over all periodic orbits γ of f , and $|\gamma|$ is the number of points in γ (Exercise 3.4.4). The *generating function* $g_f(z)$ is

another way to collect information about the periodic points of f :

$$g_f(z) = \sum_{n=1}^{\infty} |\text{Fix}(f^n)|z^n.$$

The generating function is related to the zeta function by $\zeta_f(z) = \exp(zg'_f(z))$.

The zeta function of the edge shift determined by an adjacency matrix A is denoted by ζ_A . *A priori*, the zeta function is merely a formal power series. The next proposition shows that the zeta function of an SFT is a rational function.

PROPOSITION 3.4.2. $\zeta_A(z) = (\det(I - zA))^{-1}$.

Proof. Observe that

$$\exp\left(\sum_{n=1}^{\infty} \frac{x^n}{n}\right) = \exp(-\log(1-x)) = \frac{1}{1-x},$$

and that $|\text{Fix}(\sigma^n | \Sigma_A)| = \text{tr}(A^n) = \sum_{\lambda} \lambda^n$, where the sum is over the eigenvalues of A , repeated with the proper multiplicity (see Exercise 1.4.2). Therefore, if A is $N \times N$,

$$\begin{aligned} \zeta_A(z) &= \exp\left(\sum_{n=1}^{\infty} \sum_{\lambda} \frac{(\lambda z)^n}{n}\right) = \prod_{\lambda} \exp\left(\sum_{n=1}^{\infty} \frac{(\lambda z)^n}{n}\right) = \prod_{\lambda} (1 - \lambda z)^{-1} \\ &= \frac{1}{z^N} \prod_{\lambda} \left(\frac{1}{z} - \lambda\right)^{-1} = \left(z^N \det\left(\frac{1}{z}I - A\right)\right)^{-1} = (\det(I - zA))^{-1}. \end{aligned}$$

□

The following theorem addresses the rationality of the zeta function for a general subshift.

THEOREM 3.4.3 (Bowen–Lanford [BL70]). *The zeta function of a subshift $X \subset \Sigma_m$ is rational if and only if there are matrices A and B such that $|\text{Fix}(\sigma^n | X)| = \text{tr} A^n - \text{tr} B^n$ for all $n \in \mathbb{N}_0$.*

Exercise 3.4.1. Let A be a non-negative, non-zero, square matrix, S_n the sum of entries of A^n , and λ the eigenvalue of A with largest modulus. Prove that $\lim_{n \rightarrow \infty} (\log S_n)/n = \log \lambda$.

Exercise 3.4.2. Calculate the zeta and generating functions of the full 2-shift.

Exercise 3.4.3. Let $A = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$. Calculate the zeta function of Σ_A^e .

Exercise 3.4.4. Prove the product formula for the zeta function.

Exercise 3.4.5. Calculate the generating function of an edge shift with adjacency matrix A .

Exercise 3.4.6. Calculate the zeta function of a hyperbolic toral automorphism (see Exercise 1.7.4).

Exercise 3.4.7. Prove that if the zeta function is rational, then so is the generating function.

3.5 Strong Shift Equivalence and Shift Equivalence

We saw in §3.2 that any subshift of finite type is isomorphic to an edge shift Σ_A^e for some adjacency matrix A . In this section, we give an algebraic condition on pairs of adjacency matrices that is equivalent to topological conjugacy of the corresponding edge shifts.

Square matrices A and B are *elementary strong shift equivalent* if there are (not necessarily square) non-negative integer matrices U and V such that $A = UV$ and $B = VU$. Matrices A and B are *strong shift equivalent* if there are (square) matrices A_1, \dots, A_n such that $A_1 = A$, $A_n = B$, and the matrices A_i and A_{i+1} are elementary strong shift equivalent. For example, the matrices

$$\begin{pmatrix} 1 & 1 & 0 \\ 1 & 1 & 1 \\ 2 & 2 & 1 \end{pmatrix} \text{ and } (3)$$

are strong shift equivalent but not elementary strong shift equivalent (Exercise 3.5.1).

THEOREM 3.5.1 (Williams [Wil73]). *The edge shifts Σ_A^e and Σ_B^e are topologically conjugate if and only if the matrices A and B are strong shift equivalent.*

Proof. We show here only that strong shift equivalence gives an isomorphism of the edge shifts. The other direction is much more difficult (see [LM95]).

It is sufficient to consider the case when A and B are elementary strong shift equivalent. Let $A = UV$, $B = VU$, and Γ_A, Γ_B be the (disjoint) directed graphs with adjacency matrices A and B . If A is $k \times k$ and B is $l \times l$, then U is $k \times l$ and V is $l \times k$. We interpret the entry U_{ij} as the number of (additional) edges from vertex i of Γ_A to vertex j of Γ_B , and similarly we interpret V_{ji} as the number of edges from vertex j of Γ_B to vertex i of Γ_A . Since




Figure 3.1. A graph constructed from an elementary strong shift equivalence.

$A_{pq} = \sum_{j=1}^l U_{pj} V_{jq}$, the number of edges in Γ_A from vertex p to vertex q is the same as the number of paths of length 2 from vertex p to vertex q through a vertex in Γ_B . Therefore we can choose a one-to-one correspondence ϕ between the edges a of Γ_A and pairs uv of edges determined by U and V , i.e., $\phi(a) = uv$, so that the starting vertex of u is the starting vertex of a , the terminal vertex of u is the starting vertex of v , and the terminal vertex of v is the terminal vertex of a . Similarly, there is a bijection ψ from the edges b of Γ_B to pairs vu of edges determined by V and U . For each sequence $\dots a_{-1}a_0a_1 \dots \in \Sigma_A^e$ apply ϕ to get

$$\dots \phi(a_{-1})\phi(a_0)\phi(a_1) \dots = \dots u_{-1}v_{-1}u_0v_0u_1v_1 \dots,$$

and then apply ψ^{-1} to get $\dots b_{-1}b_0b_1 \dots \in \Sigma_B^e$ with $b_i = \psi^{-1}(v_iu_{i+1})$ (see Figure 3.1). This gives an isomorphism from Σ_A^e to Σ_B^e . \square

Square matrices A and B are *shift equivalent* if there are (not necessarily square) non-negative integer matrices U, V , and a positive integer k (called the *lag*) such that

$$A^k = UV, \quad B^k = VU, \quad AU = UB, \quad BV = VA.$$

The notion of shift equivalence was introduced by R. Williams, who conjectured that if two primitive matrices are shift equivalent, then they are strong shift equivalent, or, in view of Theorem 3.5.1, that shift equivalence classifies subshifts of finite type. K. Kim and F. Roush [KR99] constructed a counterexample to this conjecture.

For other notions of equivalence for SFTs see [Boy93].

Exercise 3.5.1. Show that the matrices

$$A = \begin{pmatrix} 1 & 1 & 0 \\ 1 & 1 & 1 \\ 2 & 2 & 1 \end{pmatrix} \quad \text{and} \quad B = (3)$$

are strong shift equivalent but not elementary strong shift equivalent. Write down an explicit isomorphism from (Σ_A, σ) to (Σ_B, σ) .

Exercise 3.5.2. Show that strong shift equivalence and shift equivalence are equivalence relations and elementary strong shift equivalence is not.

3.6 Substitutions²

For an alphabet $\mathcal{A}_m = \{0, 1, \dots, m-1\}$, denote by \mathcal{A}_m^* the collection of all finite words in \mathcal{A}_m , and by $|w|$ the length of $w \in \mathcal{A}_m^*$. A *substitution* $s: \mathcal{A}_m \rightarrow \mathcal{A}_m^*$ assigns to every symbol $a \in \mathcal{A}_m$ a finite word $s(a) \in \mathcal{A}_m^*$. We assume throughout this section that $|s(a)| > 1$ for some $a \in \mathcal{A}_m$, and that $|s^n(b)| \rightarrow \infty$ for every $b \in \mathcal{A}_m$. Applying the substitution to each element of a sequence or a word gives maps $s: \mathcal{A}_m^* \rightarrow \mathcal{A}_m^*$ and $s: \Sigma_m^+ \rightarrow \Sigma_m^+$,

$$x_0 x_1 \dots \mapsto s(x_0) s(x_1) \dots$$

These maps are continuous but not surjective. If $s(a)$ has the same length for all $a \in \mathcal{A}_m$, then s is said to have *constant length*.

Consider the example $m = 2$, $s(0) = 01$, $s(1) = 10$. We have: $s^2(0) = 0110$, $s^3(0) = 01101001$, $s^4(0) = 0110100110010110, \dots$ If \bar{w} is the word obtained from w by interchanging 0 and 1, then $s^{n+1}(0) = s^n(0)\overline{s^n(0)}$. The sequence of finite words $s^n(0)$ stabilizes to an infinite sequence

$$\mathcal{M} = 01101001100101101001011001101001 \dots$$

called the *Morse sequence*. The sequences \mathcal{M} and $\bar{\mathcal{M}}$ are the only fixed points of s in Σ_m^+ .

PROPOSITION 3.6.1. *Every substitution s has a periodic point in Σ_m^+ .*

Proof. Consider the map $a \mapsto s(a)_0$. Since \mathcal{A}_m contains m elements, there are $n \in \{1, \dots, m\}$ and $a \in \mathcal{A}_m$ such that $s^n(a)_0 = a$. If $|s^n(a)| = 1$, then the sequence $aaa\dots$ is a fixed point of s^n . Otherwise, $|s^{ni}(a)| \rightarrow \infty$, and the sequence of finite words $s^{ni}(a)$ stabilizes to a fixed point of s^n in Σ_m^+ . \square

If a substitution s has a fixed point $x = x_0 x_1 \dots \in \Sigma_m^+$ and $|s(x_0)| > 1$, then $s(x_0)_0 = x_0$ and the sequence $s^n(x_0)$ stabilizes to x ; we write $x = s^\infty(x_0)$. If $|s(a)| > 1$ for every $a \in \mathcal{A}_m$, then s has at most m fixed points in Σ_m .

The closure $\Sigma_s(a)$ of the (forward) orbit of a fixed point $s^\infty(a)$ under the shift σ is a subshift.

We call a substitution $s: \mathcal{A}_m \rightarrow \mathcal{A}_m^*$ *irreducible* if for any $a, b \in \mathcal{A}_m$ there is $n(a, b) \in \mathbb{N}$ such that $s^{n(a,b)}(a)$ contains b ; s is *primitive* if there is $n \in \mathbb{N}$ such that $s^n(a)$ contains b for all $a, b \in \mathcal{A}_m$.

We assume from now on that $|s^n(b)| \rightarrow \infty$ for every $b \in \mathcal{A}_m$.

² Several arguments in this section follow in part those of [Que87].

PROPOSITION 3.6.2. *Let s be an irreducible substitution over \mathcal{A}_m . If $s(a)_0 = a$ for some $a \in \mathcal{A}_m$, then s is primitive and the subshift $(\Sigma_s(a), \sigma)$ is minimal.*

Proof. Observe that $s^n(a)_0 = a$ for all $n \in \mathbb{N}$. Since s is irreducible, for every $b \in \mathcal{A}_m$ there is $n(b)$ such that b appears in $s^{n(b)}(a)$, and therefore appears in $s^n(a)$ for all $n \geq n(b)$. Hence, $s^n(a)$ contains all symbols from \mathcal{A}_m if $n \geq N = \max n(b)$. Since s is irreducible, for every $b \in \mathcal{A}_m$ there is $k(b)$ such that a appears in $s^{k(b)}(b)$ and hence in $s^n(b)$ with $n \geq k(b)$. It follows that for every $c \in \mathcal{A}_m$, $s^n(c)$ contains all symbols from \mathcal{A}_m if $n \geq 2(N + \max k(b))$, so s is primitive.

Recall (Proposition 2.1.3) that $(\Sigma_s(a), \sigma)$ is minimal if and only if $s^\infty(a)$ is almost periodic, i.e., for every $n \in \mathbb{N}$ the word $s^n(a)$ occurs in $s^\infty(a)$ infinitely often, and the gaps between successive occurrences are bounded. This happens if and only if a recurs in $s^\infty(a)$ with bounded gaps, which holds true because s is primitive (Exercise 3.6.1). \square

For two words $u, v \in \mathcal{A}_m^*$ denote by $N_u(v)$ the number of times u occurs in v . The *composition matrix* $M = M(s)$ of a substitution s is the non-negative integer matrix with entries $M_{ij} = N_i(s(j))$. The matrix $M(s)$ is primitive (respectively, irreducible) if and only if the substitution s is primitive (respectively, irreducible). For a word $w \in \mathcal{A}_m^*$, the numbers $N_i(w)$, $i \in \mathcal{A}_m$, form a vector $N(w) \in \mathbb{R}^m$. Observe that $M(s^n) = (M(s))^n$ for all $n \in \mathbb{N}$ and $N(s(w)) = M(s)N(w)$. If s has constant length l , then the sum of every column of M is l and the transpose of $l^{-1}M$ is a stochastic matrix.

PROPOSITION 3.6.3. *Let $s: \mathcal{A}_m \rightarrow \mathcal{A}_m^*$ be a primitive substitution, and let λ be the largest in modulus eigenvalue of $M(s)$. Then for every $a \in \mathcal{A}_m$*

1. $\lim_{n \rightarrow \infty} \lambda^{-n} N(s^n(a))$ is an eigenvector of $M(s)$ with eigenvalue λ ,
2. $\lim_{n \rightarrow \infty} \frac{|s^{n+1}(a)|}{|s^n(a)|} = \lambda$,
3. $v = \lim_{n \rightarrow \infty} |s^n(a)|^{-1} N(s^n(a))$ is an eigenvector of $M(s)$ corresponding to λ , and $\sum_{i=0}^{m-1} v_i = 1$.

Proof. The proposition follows directly from Theorem 3.3.1 (Exercise 3.6.2). \square

PROPOSITION 3.6.4. *Let s be a primitive substitution, $s^\infty(a)$ be a fixed point of s , and l_n be the number of different words of length n occurring in $s^\infty(a)$. Then there is a constant C such that $l_n \leq C \cdot n$ for all $n \in \mathbb{N}$. Consequently, the topological entropy of $(\Sigma_s(a), \sigma)$ is 0.*

Proof. Let $\underline{v}_k = \min_{a \in \mathcal{A}_m} |s^k(a)|$ and $\bar{v}_k = \max_{a \in \mathcal{A}_m} |s^k(a)|$, and note that $\underline{v}_k, \bar{v}_k \rightarrow \infty$ monotonically in k . Hence for every $n \in \mathbb{N}$ there is $k = k(n) \in \mathbb{N}$ such that $\underline{v}_{k-1} \leq n \leq \underline{v}_k$. Therefore, every word of length n occurring in x

is contained in $s^k(ab)$ for a pair of consecutive symbols ab from x . Let λ be the maximal-modulus eigenvalue λ of the primitive composition matrix $M = M(s)$. Then for every non-zero vector v with non-negative components there are constants $C_1(v)$ and $C_2(v)$ such that for all $k \in \mathbb{N}$,

$$C_1(v)\lambda^k \leq \|M^k v\| \leq C_2(v)\lambda^k,$$

where $\|\cdot\|$ is the Euclidean norm. Hence, by Proposition 3.6.3(1), there are positive constants C_1 and C_2 such that for all $k \in \mathbb{N}$

$$C_1 \cdot \lambda^k \leq \underline{v}_k \leq \bar{v}_k \leq C_2 \cdot \lambda^k.$$

Since for every $a \in \mathcal{A}_m$ there are at most \bar{v}_k different words of length n in $s^k(ab)$ with initial symbol in $s^k(a)$, we have

$$l_n \leq m^2 \bar{v}_k \leq C_2 \lambda^k m^2 = \left(\frac{C_2}{C_1} m^2 \lambda\right) C_1 \lambda^{k-1} \leq \left(\frac{C_2}{C_1} m^2 \lambda\right) \underline{v}_{k-1} \leq \left(\frac{C_2}{C_1} m^2 \lambda\right) n.$$

□

Exercise 3.6.1. Prove that if s is primitive and $s(a)_0 = a$, then each symbol $b \in \mathcal{A}_m$ appears in $s^\infty(a)$ infinitely often and with bounded gaps.

Exercise 3.6.2. Prove Proposition 3.6.3.

3.7 Sofic Shifts

A subshift $X \subset \Sigma_m$ is called *sofic* if it is a factor of a subshift of finite type, i.e., there is an adjacency matrix A and a code $c: \Sigma_A^e \rightarrow X$ such that $c \circ \sigma = \sigma \circ c$. Sofic shifts have applications in finite-state automata and data transmission and storage [MRS95].

A simple example of a sofic shift is the following subshift of (Σ_2, σ) , called the *even system of Weiss* [Wei73]. Let A be the adjacency matrix of the graph Γ_A consisting of two vertices u and v , an edge from u to itself labeled 1, an edge from u to v labeled 0_1 , and an edge from v to u labeled 0_2 (see Figure 3.2). Let X be the set of sequences of 0s and 1s such that there is an even number of 0s between every two 1s. The surjective code $c: \Sigma_A \rightarrow X$ replaces both 0_1 and 0_2 by 0.

As Proposition 3.7.1 shows, every sofic shift can be obtained by the following construction. Let Γ be a finite directed labeled graph, i.e., the edges of Γ are labeled by an alphabet \mathcal{A}_m . Note that we do *not* assume that different edges of Γ are labeled differently. The subset $X_\Gamma \subset \Sigma_m$ consisting of all infinite directed paths in Γ is closed and shift invariant.




Figure 3.2. The directed graph used to construct the even system of Weiss.

If a subshift (X, σ) is isomorphic to (X_Γ, σ) for some directed labeled graph Γ , then we say that Γ is a *presentation* of X . For example, a presentation for the even system of Weiss is obtained by replacing the labels 0_1 and 0_2 with 0 in Figure 3.2.

PROPOSITION 3.7.1. *A subshift $X \subset \Sigma_m$ is sofic if and only if it admits a presentation by a finite directed labeled graph.*

Proof. Since X is sofic, there is a matrix A and a code $c: \Sigma_A^e \rightarrow X$ (see Corollary 3.2.2). By Proposition 3.1.2, c is a block code. By passing to a higher block presentation we may assume that c is a 1-block code. Hence, X admits a presentation by a finite directed labeled graph. The converse is Exercise 3.7.2. \square

Exercise 3.7.1. Prove that the even system of Weiss is not a subshift of finite type.

Exercise 3.7.2. Prove that for any directed labeled graph Γ , the set X_Γ is a sofic shift.

Exercise 3.7.3. Show that there are only countably many non-isomorphic sofic shifts. Conclude that there are subshifts that are not sofic.

3.8 Data Storage³

Most computer storage devices (floppy disk, hard drive, etc.) store data as a chain of magnetized segments on tracks. A magnetic head can either change or detect the polarity of a segment as it passes the head. Since it is technically much easier to detect a change of polarity than to measure the polarity, a common technique is to record a 1 as a change of polarity and a 0 as no change in polarity. The two major problems that restrict the effectiveness of this method are intersymbol interference and clock drift. Both of these

³ The presentation of this section follows in part [BP94].

problems can be ameliorated by applying a block code to the data before it is written to the storage device.

Intersymbol interference occurs when two polarity changes are adjacent to each other on the track; the magnetic fields from the adjacent positions partially cancel each other, and the magnetic head may not read the track correctly. This effect can be minimized by requiring that in the encoded sequence every two 1s are separated by at least one 0.

A sequence of n 0s with 1s on both ends is read off the track as two pulses separated by n non-pulses. The length n is obtained by measuring the time between the pulses. Every time a 1 is read, the clock is synchronized. However, for a long sequence of 0s, clock error accumulates, which may cause the data to read incorrectly. To counteract this effect the encoded sequence is required to have no long stretches of 0s.

A common coding scheme called *modified frequency modulation* (MFM) inserts a 0 between each two symbols unless they are both 0s, in which case it inserts a 1. For example, the sequence

$$10100110001$$

is encoded for storage as

$$100010010010100101001.$$

This requires twice the length of the track, but results in fewer read/write errors. The set of sequences produced by the MFM coding is a sofic system (Exercise 3.8.3).

There are other considerations for storage devices that impose additional conditions on the sequences used to encode data. For example, the total magnetic charge of the device should not be too large. This restriction leads to a subset of (Σ_2, σ) that is not of finite type and not sofic.

Recall that the topological entropy of the factor does not exceed the topological entropy of the extension (Exercise 2.5.5). Therefore in any one-to-one coding scheme, which increases the length of the sequence by a factor of $n > 1$, the topological entropy of the original subshift must be not more than n times the topological entropy of the target subshift.

Exercise 3.8.1. Prove that the sequences produced by MFM have at least one and at most three 0s between every two 1s.

Exercise 3.8.2. Describe an algorithm to reverse the MFM coding.

Exercise 3.8.3. Prove that the set of sequences produced by the MFM coding is a sofic system.

Ergodic Theory

Ergodic¹ theory is the study of statistical properties of dynamical systems relative to a measure on the underlying space of the dynamical system. The name comes from classical statistical mechanics, where the “ergodic hypothesis” asserts that, asymptotically, the time average of an observable is equal to the space average. Among the dynamical systems with natural invariant measures that we have encountered before are circle rotations (§1.2) and toral automorphisms (§1.7). Unlike topological dynamics, which studies the behavior of individual orbits (e.g., periodic orbits), ergodic theory is concerned with the behavior of the system on a set of full measure and with the induced action in spaces of measurable functions such as L^p (especially L^2).

The proper setting for ergodic theory is a dynamical system on a *measure space*. Most natural (non-atomic) measure spaces are measure-theoretically isomorphic to an interval $[0, a]$ with Lebesgue measure, and the results in this chapter are most important in that setting. The first section of this chapter recalls some notation, definitions, and facts from measure theory. It is not intended to serve as a complete exposition of measure theory (for a full introduction see, for example, [Hal50] or [Rud87]).

4.1 Measure-Theory Preliminaries

A non-empty collection \mathfrak{A} of subsets of a set X is called a σ -*algebra* if \mathfrak{A} is closed under complements and countable unions (and hence countable intersections). A *measure* μ on \mathfrak{A} is a non-negative (possibly infinite) function on \mathfrak{A} that is σ -*additive*, i.e., $\mu(\bigcup_i A_i) = \sum_i \mu(A_i)$ for any countable collection of disjoint sets $A_i \in \mathfrak{A}$. A set of measure 0 is called a *null set*. A set whose complement is a null set is said to have *full measure*. The σ -algebra

¹ From the Greek words $\acute{\epsilon}\rho\gamma\omicron\nu$, “work,” and $\delta\delta\omicron\varsigma$, “path.”

is *complete* (relative to μ) if it contains every subset of every null set. Given a σ -algebra \mathfrak{A} and a measure μ , the *completion* $\bar{\mathfrak{A}}$ is the smallest σ -algebra containing \mathfrak{A} and all subsets of null sets in \mathfrak{A} ; the σ -algebra $\bar{\mathfrak{A}}$ is complete.

A *measure space* is a triple (X, \mathfrak{A}, μ) , where X is a set, \mathfrak{A} is a σ -algebra of subsets of X , and μ is a σ -additive measure. We always assume that \mathfrak{A} is complete, and that μ is *σ -finite*, i.e., that X is a countable union of subsets of finite measure. The elements of \mathfrak{A} are called *measurable sets*.

If $\mu(X) = 1$, then (X, \mathfrak{A}, μ) is called a *probability space* and μ is a *probability measure*. If $\mu(X)$ is finite, then we can rescale μ by the factor $1/\mu(X)$ to obtain a probability measure.

Let (X, \mathfrak{A}, μ) and (Y, \mathfrak{B}, ν) be measure spaces. The product measure space is the triple $(X \times Y, \mathfrak{C}, \mu \times \nu)$, where \mathfrak{C} is the completion relative to $\mu \times \nu$ of the σ -algebra generated by $\mathfrak{A} \times \mathfrak{B}$.

Let (X, \mathfrak{A}, μ) and (Y, \mathfrak{B}, ν) be measure spaces. A map $T: X \rightarrow Y$ is called *measurable* if the preimage of any measurable set is measurable. A measurable map T is *non-singular* if the preimage of every set of measure 0 has measure 0, and is *measure-preserving* if $\mu(T^{-1}(B)) = \nu(B)$ for every $B \in \mathfrak{B}$. A non-singular map from a measure space into itself is called a *non-singular transformation* (or simply a *transformation*). If a transformation T preserves a measure μ , then μ is called *T -invariant*. If T is an invertible measurable transformation, and its inverse is measurable and non-singular, then the iterates T^n , $n \in \mathbb{Z}$, form a group of measurable transformations. Measure spaces (X, \mathfrak{A}, μ) and (Y, \mathfrak{B}, ν) are *isomorphic* if there is a subset X' of full measure in X , a subset Y' of full measure in Y , and an invertible bijection $T: X' \rightarrow Y'$ such that T and T^{-1} are measurable and measure-preserving with respect to (\mathfrak{A}, μ) and (\mathfrak{B}, ν) . An isomorphism from a measure space into itself is an *automorphism*.

Denote by λ the Lebesgue measure on \mathbb{R} . A flow T^t on a measure space (X, \mathfrak{A}, μ) is *measurable* if the map $T: X \times \mathbb{R} \rightarrow X$, $(x, t) \mapsto T^t(x)$, is measurable with respect to the product measure on $X \times \mathbb{R}$, and $T^t: X \rightarrow X$ is a non-singular measurable transformation for each $t \in \mathbb{R}$. A measurable flow T^t is a *measure-preserving flow* if each T^t is a measure-preserving transformation.

Let T be a measure-preserving transformation of a measure space (X, \mathfrak{A}, μ) , and S a measure-preserving transformation of a measure space (Y, \mathfrak{B}, ν) . We say that T is an *extension* of S if there are sets $X' \subset X$ and $Y' \subset Y$ of full measure and a measure-preserving map $\psi: X' \rightarrow Y'$ such that $\psi \circ T = S \circ \psi$. A similar definition holds for measure-preserving flows. If ψ is an isomorphism, then T and S are called *isomorphic*. The product $T \times S$ is a measure-preserving transformation of $(X \times Y, \mathfrak{C}, \mu \times \nu)$, where \mathfrak{C} is the completion of the σ -algebra generated by $\mathfrak{A} \times \mathfrak{B}$.

Let X be a topological space. The smallest σ -algebra containing all the open subsets of X is called the *Borel σ -algebra* of X . If \mathfrak{A} is the Borel σ -algebra, then a measure μ on \mathfrak{A} is a *Borel measure* if the measure of any compact set is finite. A Borel measure is *regular* in the sense that the measure of any set is the infimum of measures of open sets containing it, and the supremum of measures of compact sets contained in it.

A one-point subset with positive measure is called an *atom*. A finite measure space is a *Lebesgue space* if it is isomorphic to the union of an interval $[0, a]$ (with Lebesgue measure) and at most countably many atoms. Most natural measure spaces are Lebesgue spaces. For example, if X is a complete separable metric space, μ a finite Borel measure on X , and \mathfrak{A} the completion of the Borel σ -algebra with respect to μ , then (X, \mathfrak{A}, μ) is a Lebesgue space. In particular, the unit square $[0, 1] \times [0, 1]$ with Lebesgue measure is (measure-theoretically) isomorphic to the unit interval $[0, 1]$ with Lebesgue measure (Exercise 4.1.1).

A Lebesgue space without atoms is called *non-atomic*, and is isomorphic to an interval $[0, a]$ with Lebesgue measure.

A set has full measure if its complement has measure 0. We say a property holds *mod 0* in X , or holds for *μ -almost every (a.e.) x* , if it holds on a subset of full μ -measure in X . We also use the word *essentially* to indicate that a property holds *mod 0*.

Let (X, \mathfrak{A}, μ) be a measure space. Two measurable functions are equivalent if they coincide on a set of full measure. For $p \in (0, \infty)$, the space $L^p(X, \mu)$ consists of equivalence classes *mod 0* of measurable functions $f: X \rightarrow \mathbb{C}$ such that $\int |f|^p d\mu < \infty$. As a rule, if there is no ambiguity, we identify the function with its equivalence class. For $p \geq 1$, the L^p norm is defined by $\|f\|_p = (\int |f|^p d\mu)^{1/p}$. The space $L^2(X, \mu)$ is a Hilbert space with inner product $\langle f, g \rangle = \int f \cdot g d\mu$. The space $L^\infty(X, \mu)$ consists of equivalence classes of essentially bounded measurable functions. If μ is finite, then $L^\infty(X, \mu) \subset L^p(X, \mu)$ for all $p > 0$. If X is a topological space and μ is a Borel measure on X , then the space $C_0(X, \mathbb{C})$ of continuous, complex-valued, compactly supported functions on X is dense in $L^p(X, \mu)$ for all $p > 0$.

Exercise 4.1.1. Prove that the unit square $[0, 1] \times [0, 1]$ with Lebesgue measure is (measure-theoretically) isomorphic to the unit interval $[0, 1]$ with Lebesgue measure.

4.2 Recurrence

The following famous result of Poincaré implies that recurrence is a generic property of orbits of measure-preserving dynamical systems.

THEOREM 4.2.1 (Poincaré Recurrence Theorem). *Let T be a measure-preserving transformation of a probability space (X, \mathfrak{A}, μ) . If A is a measurable set, then for a.e. $x \in A$, there is some $n \in \mathbb{N}$ such that $T^n(x) \in A$. Consequently, for a.e. $x \in A$, there are infinitely many $k \in \mathbb{N}$ for which $T^k(x) \in A$.*

Proof. Let

$$B = \{x \in A : T^k(x) \notin A \text{ for all } k \in \mathbb{N}\} = A \setminus \bigcup_{k \in \mathbb{N}} T^{-k}(A).$$

Then $B \in \mathfrak{A}$, and all the preimages $T^{-k}(B)$ are disjoint, are measurable, and have the same measure as B . Since X has finite total measure, it follows that B has measure 0. Since every point in $A \setminus B$ returns to A , this proves the first assertion. The proof of the second assertion is Exercise 4.2.1. \square

For continuous maps of topological spaces, there is a connection between measure-theoretic recurrence and the topological recurrence introduced in Chapter 2. If X is a topological space, and μ is a Borel measure on X , then $\text{supp } \mu$ (the *support* of μ) is the complement of the union of all open sets with measure 0 or, equivalently, the intersection of all closed sets with full measure. Recall from §2.1 that the set of recurrent points of a continuous map $T: X \rightarrow X$ is $\mathcal{R}(T) = \{x \in X : x \in \omega(x)\}$.

PROPOSITION 4.2.2. *Let X be a separable metric space, μ a Borel probability measure on X , and $f: X \rightarrow X$ a continuous measure-preserving transformation. Then almost every point is recurrent, and hence $\text{supp } \mu \subset \overline{\mathcal{R}(f)}$.*

Proof. Since X is separable, there is a countable basis $\{U_i\}_{i \in \mathbb{Z}}$ for the topology of X . A point $x \in X$ is recurrent if it returns (in the future) to every basis element containing it. By the Poincaré recurrence theorem, for each i , there is a subset \tilde{U}_i of full measure in U_i such that every point of \tilde{U}_i returns to U_i . Then $X_i = \tilde{U}_i \cup (X \setminus U_i)$ has full measure in X , so $\tilde{X} = \bigcap_{i \in \mathbb{Z}} X_i = \mathcal{R}(T)$ has full measure in X . \square

We will discuss some applications of measure-theoretic recurrence in §4.11.

Given a measure-preserving transformation T in a finite measure space (X, \mathfrak{A}, μ) and a measurable subset $A \in \mathfrak{A}$ of positive measure, the *derivative transformation* $T_A: A \rightarrow A$ is defined by $T_A(x) = T^k(x)$, where $k \in \mathbb{N}$ is the smallest natural number for which $T^k(x) \in A$. The derivative transformation is often called the *first return map*, or the *Poincaré map*. By Theorem 4.2.1, T_A is defined on a subset of full measure in A .

Let T be a transformation on a measure space (X, \mathfrak{A}, μ) , and $f: X \rightarrow \mathbb{N}$ a measurable function. Let $X_f = \{(x, k) : x \in X, 1 \leq k \leq f(x)\} \subset X \times \mathbb{N}$. Let \mathfrak{A}_f be the σ -algebra generated by the sets $A \times \{k\}$, $A \in \mathfrak{A}$, $k \in \mathbb{N}$, and define

$\mu_f(A \times \{k\}) = \mu(A)$. Define the *primitive transformation* $T_f: X_f \rightarrow X_f$ by $T_f(x, k) = (x, k+1)$ if $k < f(x)$ and $T_f(x, f(x)) = (T(x), 1)$. If $\mu(X) < \infty$ and $f \in L^1(X, \mathfrak{A}, \mu)$, then $\mu_f(X_f) = \int_X f(x) d\mu$. Note that the derivative transformation of T_f on the set $X \times \{1\}$ is just the original transformation T .

Primitive and derivative transformations are both referred to as *induced transformations*; we will encounter them later.

Exercise 4.2.1. Prove the second assertion of Theorem 4.2.1.

Exercise 4.2.2. Suppose $T: X \rightarrow X$ is a continuous transformation of a topological space X , and μ is a finite T -invariant Borel measure on X with $\text{supp } \mu = X$. Show that every point is non-wandering and μ -a.e. point is recurrent.

Exercise 4.2.3. Prove that if T is a measure-preserving transformation, then so are the induced transformations.

4.3 Ergodicity and Mixing

A dynamical system induces an action on functions: T acts on a function f by $(T_* f)(x) = f(T(x))$. The ergodic properties of a dynamical system correspond to the degree of statistical independence between f and $T_*^n f$. The strongest possible dependence happens for an invariant function $f(T(x)) = f(x)$. The strongest possible independence happens when a non-zero L^2 function is orthogonal to its images.

Let T be a measure-preserving transformation (or flow) on a measure space (X, \mathfrak{A}, μ) . A measurable function $f: X \rightarrow \mathbb{R}$ is *essentially T -invariant* if $\mu(\{x \in X: f(T^t x) \neq f(x)\}) = 0$ for every t . A measurable set A is *essentially T -invariant* if its characteristic function $\mathbf{1}_A$ is essentially T -invariant; equivalently, if $\mu(T^{-1}(A) \Delta A) = 0$ (we denote by Δ the symmetric difference, $A \Delta B = (A \setminus B) \cup (B \setminus A)$).

A measure-preserving transformation (or flow) T is *ergodic* if any essentially T -invariant measurable set has either measure 0 or full measure. Equivalently (Exercise 4.3.1), T is ergodic if any essentially T -invariant measurable function is constant mod 0.

PROPOSITION 4.3.1. *Let T be a measure-preserving transformation or flow on a finite measure space (X, \mathfrak{A}, μ) , and let $p \in (0, \infty]$. Then T is ergodic if and only if every essentially invariant function $f \in L^p(X, \mu)$ is constant mod 0.*

Proof. If T is ergodic, then every essentially invariant function is constant mod 0.

To prove the converse, let f be an essentially invariant measurable function on X . Then for every $M > 0$, the function

$$f_M(x) = \begin{cases} f(x) & \text{if } f(x) \leq M, \\ 0 & \text{if } f(x) > M \end{cases}$$

is bounded, is essentially invariant, and belongs to $L^p(X, \mu)$. Therefore it is constant mod 0. It follows that f itself is constant mod 0. \square

As the following proposition shows, any essentially invariant set or function is equal mod 0 to a strictly invariant set or function.

PROPOSITION 4.3.2. *Let (X, \mathfrak{A}, μ) be a measure space, and suppose that $f: X \rightarrow \mathbb{R}$ is essentially invariant for a measurable transformation or flow T on X . Then there is a strictly invariant measurable function \tilde{f} such that $f(x) = \tilde{f}(x) \text{ mod } 0$.*

Proof. We prove the proposition for a measurable flow. The case of a measurable transformation follows by a similar but easier argument and is left as an exercise.

Consider the measurable map $\Phi: X \times \mathbb{R} \rightarrow \mathbb{R}$, $\Phi(x, t) = f(T^t x) - f(x)$, and the product measure $\nu = \mu \times \lambda$ in $X \times \mathbb{R}$, where λ is Lebesgue measure on \mathbb{R} . The set $A = \Phi^{-1}(0)$ is a measurable subset of $X \times \mathbb{R}$. Since f is essentially T -invariant, for each $t \in \mathbb{R}$ the set

$$A_t = \{(x, t) \in (X \times \mathbb{R}): f(T^t x) = f(x)\}$$

has full μ -measure in $X \times \{t\}$. By the Fubini theorem, the set

$$A_f = \{x \in X: f(T^t x) = f(x) \text{ for a.e. } t \in \mathbb{R}\}$$

has full μ -measure in X . Set

$$\tilde{f}(x) = \begin{cases} f(y) & \text{if } T^t x = y \in A_f \text{ for some } t \in \mathbb{R}, \\ 0 & \text{otherwise.} \end{cases}$$

If $T^t x = y \in A_f$ and $T^s x = z \in A_f$, then y and z lie on the same orbit, and the value of f along this orbit is equal λ -almost everywhere to $f(y)$ and to $f(z)$, so $f(y) = f(z)$. Therefore \tilde{f} is well defined and strictly T -invariant. \square

A measure-preserving transformation (or flow) T on a probability space (X, \mathfrak{A}, μ) is called (*strong*) *mixing* if

$$\lim_{t \rightarrow \infty} \mu(T^{-t}(A) \cap B) = \mu(A) \cdot \mu(B)$$

for any two measurable sets $A, B \in \mathfrak{A}$. Equivalently (Exercise 4.3.3), T is mixing if

$$\lim_{t \rightarrow \infty} \int_X f(T^t(x)) \cdot g(x) d\mu = \int_X f(x) d\mu \cdot \int_X g(x) d\mu$$

for any bounded measurable functions f, g .

A measure-preserving transformation T of a probability space (X, \mathfrak{A}, μ) is called *weak mixing* if for all $A, B \in \mathfrak{A}$,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} |\mu(T^{-i}(A) \cap B) - \mu(A) \cdot \mu(B)| = 0$$

or, equivalently (Exercise 4.3.3), if for all bounded measurable functions f, g ,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} \left| \int_X f(T^i(x))g(x) d\mu - \int_X f d\mu \cdot \int_X g d\mu \right| = 0.$$

A measure-preserving flow T^t on (X, \mathfrak{A}, μ) is *weak mixing* if for all $A, B \in \mathfrak{A}$

$$\lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t |\mu(T^{-s}(A) \cap B) - \mu(A) \cdot \mu(B)| ds = 0,$$

or, equivalently (Exercise 4.3.3), if for all bounded measurable functions f, g ,

$$\lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t \left| \int_X f(T^s(x))g(x) d\mu ds - \int_X f d\mu \cdot \int_X g d\mu \right| = 0.$$

In practice, the definitions of ergodicity and mixing in terms of L^2 functions are often easier to work with than the definitions in terms of measurable sets. For example, to establish a certain property for each L^2 function on a separable topological space with Borel measure it suffices to do it for a countable set of continuous functions that is dense in L^2 (Exercise 4.3.5). If the property is “linear”, it is enough to check it for a basis in L^2 , e.g., for the exponential functions $e^{2\pi i x}$ on the circle $[0, 1)$.

PROPOSITION 4.3.3. *Mixing implies weak mixing, and weak mixing implies ergodicity.*

Proof. Suppose T is a measure-preserving transformation of the probability space (X, \mathfrak{A}, μ) . Let A and B be measurable subsets of X . If T is mixing, then $|\mu(T^{-i}(A) \cap B) - \mu(A) \cdot \mu(B)|$ converges to 0, so the averages do as well; thus T is weak mixing.

Let A be an invariant measurable set. Then applying the definition of weak mixing with $B = A$, we conclude that $\mu(A) = \mu(A)^2$, so either $\mu(A) = 1$ or $\mu(A) = 0$. \square

For continuous maps, ergodicity and mixing have the following topological consequences.

PROPOSITION 4.3.4. *Let X be a compact metric space, $T: X \rightarrow X$ a continuous map, and μ a T -invariant Borel measure on X .*

1. *If T is ergodic, then the orbit of μ -almost every point is dense in $\text{supp } \mu$.*
2. *If T is mixing, then T is topologically mixing on $\text{supp } \mu$.*

Proof. Suppose T is ergodic. Let U be a non-empty open set in $\text{supp } \mu$. Then $\mu(U) > 0$. By ergodicity, the backward invariant set $\bigcup_{k \in \mathbb{N}} T^{-k}(U)$ has full measure. Thus the forward orbit of almost every point visits U . It follows that the set of points whose forward orbit visits every element of a countable open basis has full measure in X . This proves the first assertion.

The proof of the second assertion is Exercise 4.3.4. \square

Exercise 4.3.1. Show that a measurable transformation is ergodic if and only if every essentially invariant measurable function is constant mod 0 (see the remark after Corollary 4.5.7).

Exercise 4.3.2. Let T be an ergodic measure-preserving transformation in a finite measure space (X, \mathfrak{A}, μ) , $A \in \mathfrak{A}$, $\mu(A) > 0$, and $f \in L^1(X, \mathfrak{A}, \mu)$, $f: X \rightarrow \mathbb{N}$. Prove that the induced transformations T_A and T_f are ergodic.

Exercise 4.3.3. Show that the two definitions of strong and weak mixing given in terms of sets and bounded measurable functions are equivalent.

Exercise 4.3.4. Prove the second statement of Proposition 4.3.4.

Exercise 4.3.5. Let T be a measure-preserving transformation of (X, \mathfrak{A}, μ) , and let $f \in L^1(X, \mu)$ satisfy $f(T(x)) \leq f(x)$ for a.e. x . Prove that $f(T(x)) = f(x)$ for a.e. x .

Exercise 4.3.6. Let X be a compact topological space, μ a Borel measure, and $T: X \rightarrow X$ a transformation preserving μ . Suppose that for every continuous f and g with 0 integrals,

$$\int_X f(T^n(x)) \cdot g(x) d\mu \rightarrow 0 \quad \text{as } n \rightarrow \infty.$$

Prove that T is mixing.

Exercise 4.3.7. Show that if $T: X \rightarrow X$ is mixing, then $T \times T: X \times X \rightarrow X \times X$ is mixing.

4.4 Examples

We now prove ergodicity or mixing for some of the examples from Chapter 1.

PROPOSITION 4.4.1. *The circle rotation R_α is ergodic with respect to Lebesgue measure if and only if α is irrational.*

Proof. Suppose α is irrational. By Proposition 4.3.1, it is enough to prove that any bounded R_α -invariant function $f: S^1 \rightarrow \mathbb{R}$ is constant mod 0. Since $f \in L^2(S^1, \lambda)$, the Fourier series $\sum_{n=-\infty}^{\infty} a_n e^{2n\pi i x}$ of f converges to f in the L^2 norm. The series $\sum_{n=-\infty}^{\infty} a_n e^{2n\pi i(x+\alpha)}$ converges to $f \circ R_\alpha$. Since $f = f \circ R_\alpha$ mod 0, uniqueness of Fourier coefficients implies that $a_n = a_n e^{2n\pi i \alpha}$ for all $n \in \mathbb{Z}$. Since $e^{2n\pi i \alpha} \neq 1$ for $n \neq 0$, we conclude that $a_n = 0$ for $n \neq 0$, so f is constant mod 0.

The proof of the converse is left as an exercise. \square

PROPOSITION 4.4.2. *An expanding endomorphism $E_m: S^1 \rightarrow S^1$ is mixing with respect to Lebesgue measure.*

Proof. Since any measurable subset of S^1 can be approximated by a finite union of intervals, it is sufficient to consider two intervals $A = [p/m^i, (p+1)/m^i]$, $p \in \{0, \dots, m^i - 1\}$, and $B = [q/m^j, (q+1)/m^j]$, $q \in \{0, \dots, m^j - 1\}$. Recall that $E_m^{-1}(B)$ is the union of m uniformly spaced intervals of length $1/m^{j+1}$:

$$E_m^{-1}(B) = \bigcup_{k=0}^{m-1} [(km^j + q)/m^{j+1}, (km^j + q + 1)/m^{j+1}].$$

Similarly, $E_m^{-n}(B)$ is the union of m^n uniformly spaced intervals of length $1/m^{j+n}$. Thus for $n > i$, the intersection $A \cap E_m^{-n}(B)$ consists of m^{n-i} intervals of length $m^{-(n+j)}$. Thus

$$\mu(A \cap E_m^{-n}(B)) = m^{n-i} (1/m^{n+j}) = m^{-i-j} = \mu(A) \cdot \mu(B). \quad \square$$

PROPOSITION 4.4.3. *Any hyperbolic toral automorphism $A: \mathbb{T}^n \rightarrow \mathbb{T}^n$ is ergodic with respect to Lebesgue measure.*

Proof. We consider here only the case

$$A = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}: \mathbb{T}^2 \rightarrow \mathbb{T}^2;$$

the argument in the general case is similar. Let $f: \mathbb{T}^2 \rightarrow \mathbb{R}$ be a bounded A -invariant measurable function. The Fourier series $\sum_{m,n=-\infty}^{\infty} a_{mn} e^{2\pi i(mx+ny)}$ of f converges to f in L^2 . The series

$$\sum_{m,n=-\infty}^{\infty} a_{mn} e^{2\pi i(m(2x+y)+n(x+y))}$$

converges to $f \circ A$. Since f is invariant, uniqueness of Fourier coefficients implies that $a_{mn} = a_{(2m+n)(m+n)}$ for all m, n . Since A does not have eigenvalues on the unit circle, if $a_{mn} \neq 0$ for some $(m, n) \neq (0, 0)$, then $a_{ij} = a_{mn} \neq 0$ with arbitrarily large $|i| + |j|$, and the Fourier series diverges. \square

A toral automorphism of \mathbb{T}^n corresponding to an integer matrix A is ergodic if and only if no eigenvalue of A is a root of unity; for a proof see, for example, [Pet89]. A hyperbolic toral automorphism is mixing (Exercise 4.4.3).

Let A be an $m \times m$ stochastic matrix, i.e., A has non-negative entries, and the sum of every row is 1. Suppose A has a non-negative left eigenvector q with eigenvalue 1 and sum of entries equal to 1 (recall that if A is irreducible, then by Corollary 3.3.3, q exists and is unique). We define a Borel probability measure $P = P_{A,q}$ on Σ_m (and Σ_m^+) as follows: for a cylinder C_j^n of length 1, we define $P(C_j^n) = q_j$; for a cylinder $C_{j_0, j_1, \dots, j_k}^{n, n+1, \dots, n+k} \subset \Sigma_m$ (or Σ_m^+) with $k+1 > 1$ consecutive indices,

$$P(C_{j_0, j_1, \dots, j_k}^{n, n+1, \dots, n+k}) = q_{j_0} \prod_{i=0}^{k-1} A_{j_i j_{i+1}}.$$

In other words, we interpret q as an initial probability distribution on the set $\{1, \dots, m\}$, and A as the matrix of transition probabilities. The number $P(C_j^n)$ is the probability of observing symbol j in the n th place, and A_{ij} is the probability of passing from i to j . The fact that $qA = q$ means that the probability distribution q is invariant under transition probabilities A , i.e.,

$$q_j = P(C_j^{n+1}) = \sum_{i=0}^{m-1} P(C_i^n) A_{ij}.$$

The pair (A, q) is called a *Markov chain* on the set $\{1, \dots, m\}$.

It can be shown that P extends uniquely to a shift-invariant σ -additive measure defined on the completion \mathcal{C} of the Borel σ -algebra generated by the cylinders (Exercise 4.4.5); it is called the *Markov measure* corresponding to A and q . The measure space $(\Sigma_m, \mathcal{C}, P)$ is a non-atomic Lebesgue probability space. If A is irreducible, this measure is uniquely determined by A .

A very important particular case of this situation arises when the transition probabilities do not depend on the initial state. In this case each row of A is the left eigenvector q , the shift-invariant measure P is called a *Bernoulli measure*, and the shift is called a *Bernoulli automorphism*.

Let A' be the adjacency matrix defined by $A'_{ij} = 0$ if $A_{ij} = 0$ and $A'_{ij} = 1$ if $A_{ij} > 0$. Then the support of P is precisely $\Sigma_{A'}^v \subset \Sigma_m$ (Exercise 4.4.6).

PROPOSITION 4.4.4. *If A is a primitive stochastic $m \times m$ matrix, then the shift σ is mixing in Σ_m with respect to the Markov measure $P(A)$.*

Proof. Exercise 4.4.7. □

Markov chains can be generalized to the class of *stationary (discrete) stochastic processes*, dynamical systems with invariant measures on shift spaces with a *continuous* alphabet. Let $(\Omega, \mathfrak{A}, P)$ be a probability space. A *random variable* on Ω is a measurable real-valued function on Ω . A sequence $(f_i)_{i=-\infty}^{\infty}$ of *random variables* is *stationary* if, for any $i_1, \dots, i_k \in \mathbb{Z}$ and any Borel subsets $B_1, \dots, B_k \subset \mathbb{R}$,

$$P\{\omega \in \Omega: f_{i_j}(\omega) \in B_j, j = 1, \dots, k\} = P\{\omega \in \Omega: f_{i_j+n}(\omega) \in B_j, j = 1, \dots, k\}.$$

Define the map $\Phi: \Omega \rightarrow \mathbb{R}^{\mathbb{Z}}$ by

$$\Phi(\omega) = (\dots, f_{-1}(\omega), f_0(\omega), f_1(\omega), \dots),$$

and the measure μ on the Borel subsets of $\mathbb{R}^{\mathbb{Z}}$ by $\mu(A) = P(\Phi^{-1}(A))$. Since the sequence (f_i) is stationary, the shift $\sigma: \mathbb{R}^{\mathbb{Z}} \rightarrow \mathbb{R}^{\mathbb{Z}}$ defined by $(\sigma x)_n = x_{n+1}$ preserves μ (Exercise 4.4.8).

Exercise 4.4.1. Prove that the circle rotation R_α is not weak mixing.

Exercise 4.4.2. Let $\alpha \in \mathbb{R}$ be irrational, and let $F: \mathbb{T}^2 \rightarrow \mathbb{T}^2$ be the map $(x, y) \mapsto (x + \alpha, x + y) \bmod 1$ introduced in §2.4. Prove that F preserves the Lebesgue measure and is ergodic but not weak mixing.

Exercise 4.4.3. Prove that any hyperbolic automorphism of \mathbb{T}^n is mixing.

Exercise 4.4.4. Show that an isometry of a compact metric space is not mixing for any invariant Borel measure whose support is not a single point. In particular, circle rotations are not mixing.

Exercise 4.4.5. Prove that any Markov measure is shift invariant.

Exercise 4.4.6. Prove that $\text{supp } P_{A,q} = \Sigma_{A'}^v$.

Exercise 4.4.7. Prove Proposition 4.4.4.

Exercise 4.4.8. Prove that the measure μ on $\mathbb{R}^{\mathbb{Z}}$ constructed above for a stationary sequence (f_i) is invariant under the shift σ .

4.5 Ergodic Theorems²

The collection of all orbits represents a complete evolution of the dynamical system T . The values $f(T^n(x))$ of a (measurable) function f may represent observations such as position or velocity. Long-term averages $\frac{1}{n} \sum_{k=0}^{n-1} f(T^k(x))$ of these quantities are important in statistical physics and other areas. A central question in ergodic theory is whether these averages converge as $n \rightarrow \infty$ and, if so, whether the limit depends on x . In the context of statistical physics, the *ergodic hypothesis* states that the asymptotic time average $\lim_{n \rightarrow \infty} (1/n) \sum_{k=0}^{n-1} f(T^k(x))$ equals the space average $\int_X f d\mu$ for a.e. x . We show that this happens if T is ergodic.

Let (X, \mathfrak{A}, μ) be a measure space and $T: X \rightarrow X$ a measure-preserving transformation. For a measurable function $f: X \rightarrow \mathbb{C}$ set $(U_T f)(x) = f(T(x))$. The operator U_T is linear and multiplicative: $U_T(f \cdot g) = U_T f \cdot U_T g$. Since T is measure-preserving, U_T is an isometry of $L^p(X, \mathfrak{A}, \mu)$ for any $p \geq 1$, i.e., $\|U_T f\|_p = \|f\|_p$ for any $f \in L^p$ (Exercise 4.5.3). If T is an automorphism, then $U_T^{-1} = U_{T^{-1}}$ is also an isometry, and hence U_T is a unitary operator on $L^2(X, \mathfrak{A}, \mu)$. We denote the scalar product on $L^2(X, \mathfrak{A}, \mu)$ by $\langle f, g \rangle$, the norm by $\| \cdot \|$, and the adjoint operator of U by U^* .

LEMMA 4.5.1. *Let U be an isometry of a Hilbert space H . Then $Uf = f$ if and only if $U^*f = f$.*

Proof. For every $f, g \in H$ we have $\langle U^*Uf, g \rangle = \langle Uf, Ug \rangle = \langle f, g \rangle$ and hence $U^*Uf = f$. If $Uf = f$, then (multiplying both sides by U^*) $U^*Uf = f$. Conversely, if $U^*f = f$, then $\langle f, Uf \rangle = \langle U^*f, f \rangle = \|f\|^2$ and $\langle Uf, f \rangle = \langle f, U^*f \rangle = \|f\|^2$. Therefore $\langle Uf - f, Uf - f \rangle = \|Uf\|^2 - \langle f, Uf \rangle - \langle Uf, f \rangle + \|f\|^2 = 0$. \square

THEOREM 4.5.2 (von Neumann Ergodic Theorem). *Let U be an isometry of a separable Hilbert space H , and let P be orthogonal projection onto the subspace $I = \{f \in H: Uf = f\}$ of U -invariant vectors in H . Then for every $f \in H$*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} U^i f = Pf.$$

² Several proofs in this section are due to F. Riesz; see [Hal60].

Proof. Let $U_n = \frac{1}{n} \sum_{i=0}^{n-1} U^i$ and $L = \{g - Ug : g \in H\}$. Note that L and I are U -invariant, and I is closed. If $f = g - Ug \in L$, then $\sum_{i=0}^{n-1} U^i f = g - U^n g$ and hence $U_n f \rightarrow 0$ as $n \rightarrow \infty$. If $f \in I$, then $U_n f = f$ for all $n \in \mathbb{N}$. We will show that $L \perp I$ and $H = \bar{L} \oplus I$, where \bar{L} is the closure of L .

Let $\{f_k\}$ be a sequence in L , and suppose $f_k \rightarrow f \in \bar{L}$. Then $\|U_n f\| \leq \|U_n(f - f_k)\| + \|U_n f_k\| \leq \|U_n\| \cdot \|f - f_k\| + \|U_n f_k\|$, and hence $U_n f \rightarrow 0$ as $n \rightarrow \infty$.

Let \perp denote the orthogonal complement, and note that $\bar{L}^\perp = L^\perp$. If $h \in L^\perp$, then $0 = \langle h, g - Ug \rangle = \langle h - U^*h, g \rangle$ for all $g \in H$ so that $h = U^*h$, and hence $Uh = h$, by Lemma 4.5.1. Conversely (again using Lemma 4.5.1), if $h \in I$, then $\langle h, g - Ug \rangle = \langle h, g \rangle - \langle U^*h, g \rangle = 0$ for every $g \in H$, and hence $h \in L^\perp$.

Therefore, $H = \bar{L} \oplus I$, and $\lim_{n \rightarrow \infty} U_n$ is the identity on I and 0 on \bar{L} . □

The following theorem is an immediate corollary of the von Neumann ergodic theorem.

THEOREM 4.5.3. *Let T be a measure-preserving transformation of a finite measure space (X, \mathfrak{A}, μ) . For $f \in L^2(X, \mathfrak{A}, \mu)$, set*

$$f_N^+(x) = \frac{1}{N} \sum_{n=0}^{N-1} f(T^n(x)).$$

Then f_N^+ converges in $L^2(X, \mathfrak{A}, \mu)$ to a T -invariant function \bar{f} .

If T is invertible, then $f_N^-(x) = \frac{1}{N} \sum_{n=0}^{N-1} f(T^{-n}(x))$ also converges in $L^2(X, \mathfrak{A}, \mu)$ to \bar{f} .

Similarly, let T be a measure-preserving flow in a finite measure space (X, \mathfrak{A}, μ) . For a function $f \in L^2(X, \mathfrak{A}, \mu)$ set

$$f_\tau^+(x) = \frac{1}{\tau} \int_0^\tau f(T^t(x)) dt \quad \text{and} \quad f_\tau^-(x) = \frac{1}{\tau} \int_0^\tau f(T^{-t}(x)) dt.$$

Then f_τ^+ and f_τ^- converge in $L^2(X, \mathfrak{A}, \mu)$ to a T -invariant function \bar{f} . □

Our next objective is to prove a pointwise version of the preceding theorem. First, we need a combinatorial lemma. If a_1, \dots, a_m are real numbers and $1 \leq n \leq m$, we say that a_k is an n -leader if $a_k + \dots + a_{k+p-1} \geq 0$ for some p , $1 \leq p \leq n$.

LEMMA 4.5.4. *For every n , $1 \leq n \leq m$, the sum of all n -leaders is non-negative.*

Proof. If there are no n -leaders, the lemma is true. Otherwise, let a_k be the first n -leader, and $p \geq 1$ be the smallest integer for which $a_k + \cdots + a_{k+p-1} \geq 0$. If $k \leq j \leq k+p-1$, then $a_j + \cdots + a_{k+p-1} \geq 0$, by the choice of p , and hence a_j is an n -leader. The same argument can be applied to the sequence a_{k+p}, \dots, a_m , which proves the lemma. \square

THEOREM 4.5.5 (Birkhoff Ergodic Theorem). *Let T be a measure-preserving transformation in a finite measure space (X, \mathfrak{A}, μ) , and let $f \in L^1(X, \mathfrak{A}, \mu)$. Then the limit*

$$\bar{f}(x) = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} f(T^k(x))$$

exists for a.e. $x \in X$, is μ -integrable and T -invariant, and satisfies

$$\int_X \bar{f}(x) d\mu = \int_X f(x) d\mu.$$

If, in addition, $f \in L^2(X, \mathfrak{A}, \mu)$, then by Theorem 4.5.3, \bar{f} is the orthogonal projection of f to the subspace of T -invariant functions.

If T is invertible, then $\frac{1}{n} \sum_{k=0}^{n-1} f(T^{-k}(x))$ also converges almost everywhere to \bar{f} .

Similarly, let T be a measure-preserving flow in a finite measure space (X, \mathfrak{A}, μ) . Then

$$f_{\tau}^{+}(x) = \frac{1}{\tau} \int_0^{\tau} f(T^t(x)) dt \quad \text{and} \quad f_{\tau}^{-}(x) = \frac{1}{\tau} \int_0^{\tau} f(T^{-t}(x)) dt$$

converge almost everywhere to the same μ -integrable and T -invariant limit function \bar{f} , and $\int_X f(x) d\mu = \int_X \bar{f}(x) d\mu$.

Proof. We consider only the case of a transformation. We assume without loss of generality that f is real-valued. Let

$$A = \{x \in X: f(x) + f(T(x)) + \cdots + f(T^k(x)) \geq 0 \text{ for some } k \in \mathbb{N}_0\}.$$

LEMMA 4.5.6 (Maximal Ergodic Theorem). $\int_A f(x) d\mu \geq 0$.

Proof. Let $A_n = \{x \in X: \sum_{i=0}^k f(T^i(x)) \geq 0 \text{ for some } k, 0 \leq k \leq n\}$. Then $A_n \subset A_{n+1}$, $A = \bigcup_{n \in \mathbb{N}} A_n$ and, by the dominated convergence theorem, it suffices to show that $\int_{A_n} f(x) d\mu \geq 0$ for each n .

Fix an arbitrary $m \in \mathbb{N}$. Let $s_n(x)$ be the sum of the n -leaders in the sequence $f(x), f(T(x)), \dots, f(T^{m+n-1}(x))$. For $k \leq m-1$, let $B_k \subset X$ be the set of points for which $f(T^k(x))$ is an n -leader of this sequence. By

Lemma 4.5.4,

$$0 \leq \int_X s_n(x) d\mu = \sum_{k=0}^{m+n-1} \int_{B_k} f(T^k(x)) d\mu. \quad (4.1)$$

Note that $x \in B_k$ if and only if $T(x) \in B_{k-1}$. Therefore, $B_k = T^{-1}(B_{k-1})$ and $B_k = T^{-k}(B_0)$ for $1 \leq k \leq m-1$, and hence

$$\int_{B_k} f(T^k(x)) d\mu = \int_{T^{-k}(B_0)} f(T^k(x)) d\mu = \int_{B_0} f(x) d\mu.$$

Thus the first m terms in (4.1) are equal, and since $B_0 = A_n$,

$$m \int_{A_n} f(x) d\mu + n \int_X |f(x)| d\mu \geq 0.$$

Since m is arbitrary, the lemma follows. \square

Now we can finish the proof of the Birkhoff ergodic theorem. For any $a, b \in \mathbb{R}$, $a < b$, the set

$$X(a, b) = \left\{ x \in X: \underline{\lim}_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} f(T^i(x)) < a < b < \overline{\lim}_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} f(T^i(x)) \right\}$$

is measurable and T -invariant. We claim that $\mu(X(a, b)) = 0$. Apply Lemma 4.5.6 to $T|_{X(a,b)}$ and $f - b$ to obtain that $\int_{X(a,b)} (f(x) - b) d\mu \geq 0$. Similarly, $\int_{X(a,b)} (a - f(x)) d\mu \geq 0$, and hence $\int_{X(a,b)} (a - b) d\mu \geq 0$. Therefore $\mu(X(a, b)) = 0$. Since a and b are arbitrary, we conclude that the averages $\frac{1}{n} \sum_{i=0}^{n-1} f(T^i(x))$ converge for a.e. $x \in X$.

For $n \in \mathbb{N}$, let $f_n(x) = \frac{1}{n} \sum_{i=0}^{n-1} f(T^i(x))$. Define $\bar{f}: X \rightarrow \mathbb{R}$ by $\bar{f}(x) = \underline{\lim}_{n \rightarrow \infty} f_n(x)$. Then \bar{f} is measurable, and f_n converges a.e. to \bar{f} . By Fatou's lemma and invariance of μ ,

$$\begin{aligned} \int_X \underline{\lim}_{n \rightarrow \infty} |f_n(x)| q d\mu &\leq \underline{\lim}_{n \rightarrow \infty} \int_X |f_n(x)| d\mu \\ &\leq \underline{\lim}_{n \rightarrow \infty} \frac{1}{n} \sum_{j=0}^{n-1} \int_X |f(T^j(x))| d\mu = \int_X |f(x)| d\mu \end{aligned}$$

Thus $\int_X |\bar{f}(x)| d\mu = \int_X \underline{\lim}_{n \rightarrow \infty} |f_n(x)| d\mu$ is finite, so \bar{f} is integrable.

The proof that $\int_X f(x) d\mu = \int_X \bar{f}(x) d\mu$ is left as an exercise (Exercise 4.5.2). \square

The following facts are immediate corollaries of Theorem 4.5.5 (Exercise 4.5.4, Exercise 4.5.5).

COROLLARY 4.5.7. *A measure-preserving transformation T in a finite measure space (X, \mathfrak{A}, μ) is ergodic if and only if for each $f \in L^1(X, \mathfrak{A}, \mu)$*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} f(T^i(x)) = \frac{1}{\mu(X)} \int_X f(x) d\mu \quad \text{for a.e. } x, \quad (4.2)$$

i.e., if and only if the time average equals the space average for every L^1 function. \square

The preceding corollary implies that to check the ergodicity of a measure-preserving transformation, it suffices to verify (4.2) for a dense subset of $L^1(X, \mathfrak{A}, \mu)$, e.g., for all continuous functions if X is a compact topological space and μ is a Borel measure. Moreover, due to linearity it suffices to check the convergence for a countable collection of functions that form a basis.

COROLLARY 4.5.8. *A measure-preserving transformation T of a finite measure space (X, \mathfrak{A}, μ) is ergodic if and only if for every $A \in \mathfrak{A}$, for a.e. $x \in X$,*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} \chi_A(T^k(x)) = \frac{\mu(A)}{\mu(X)},$$

where χ_A is the characteristic function of A . \square

Exercise 4.5.1. Let T be a measure-preserving transformation of a finite measure space (X, \mathfrak{A}, μ) . Prove that T is ergodic if and only if

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} \mu(T^{-k}(A) \cap B) = \mu(A) \cdot \mu(B)$$

for any $A, B \in \mathfrak{A}$.

Exercise 4.5.2. Using the dominated convergence theorem, finish the proof of Theorem 4.5.5 by showing that the averages $\frac{1}{n} \sum_{j=0}^{n-1} f$ converge to \bar{f} in L^1 .

Exercise 4.5.3. Prove that if T is a measure-preserving transformation, then U_T is an isometry of $L^p(X, \mathfrak{A}, \mu)$ for any $p \geq 1$.

Exercise 4.5.4. Prove Corollary 4.5.7.

Exercise 4.5.5. Prove Corollary 4.5.8.

Exercise 4.5.6. A real number x is said to be *normal in base n* if for any $k \in \mathbb{N}$, every finite word of length k in the alphabet $\{0, \dots, n-1\}$ appears with asymptotic frequency n^{-k} in the base- n expansion of x . Prove that almost every real number is normal with respect to every base $n \in \mathbb{N}$.

4.6 Invariant Measures for Continuous Maps

In this section, we show that a continuous map T of a compact metric space X into itself has at least one invariant Borel probability measure. Every finite Borel measure μ on X defines a bounded linear functional $L_\mu(f) = \int_X f d\mu$ on the space $C(X)$ of continuous functions on X ; moreover, L_μ is positive in the sense that $L_\mu(f) \geq 0$ if $f \geq 0$. The Riesz representation theorem [Rud87] states that the converse is also true: for every positive bounded linear functional L on $C(X)$, there is a finite Borel measure μ on X such that $L = \int_X f d\mu$.

THEOREM 4.6.1 (Krylov–Bogolubov). *Let X be a compact metric space and $T: X \rightarrow X$ a continuous map. Then there is a T -invariant Borel probability measure μ on X .*

Proof. Fix $x \in X$. For a function $f: X \rightarrow \mathbb{R}$ set $S_f^n(x) = \frac{1}{n} \sum_{i=0}^{n-1} f(T^i(x))$. Let $\mathcal{F} \subset C(X)$ be a dense countable collection of continuous functions on X . For any $f \in \mathcal{F}$ the sequence $S_f^n(x)$ is bounded, and hence has a convergent subsequence. Since \mathcal{F} is countable, there is a sequence $n_j \rightarrow \infty$ such that the limit

$$S_f^\infty(x) = \lim_{j \rightarrow \infty} S_f^{n_j}(x)$$

exists for every $f \in \mathcal{F}$. For any $g \in C(X)$ and any $\epsilon > 0$ there is $f \in \mathcal{F}$ such that $\max_{y \in X} |g(y) - f(y)| < \epsilon$. Therefore, for a large enough j ,

$$|S_g^{n_j}(x) - S_f^\infty(x)| \leq S_{|g-f|}^{n_j}(x) + |S_f^{n_j}(x) - S_f^\infty(x)| \leq 2\epsilon,$$

so $S_g^{n_j}(x)$ is a Cauchy sequence. Thus, the limit $S_g^\infty(x)$ exists for every $g \in C(X)$ and defines a bounded positive linear functional L_x on $C(X)$. By the Riesz representation theorem, there is a Borel probability measure μ such that $L_x(g) = \int_X g d\mu$. Note that

$$|S_g^{n_j}(T(x)) - S_g^{n_j}(x)| = \frac{1}{n_j} |g(T^{n_j}(x)) - g(x)|.$$

Therefore, $S_g^\infty(T(x)) = S_g^\infty(x)$ and μ is T -invariant. \square

Let $\mathcal{M} = \mathcal{M}(X)$ denote the set of all Borel probability measures on X . A sequence of measures $\mu_n \in \mathcal{M}$ converges in the *weak* topology* to a measure $\mu \in \mathcal{M}$ if $\int_X f d\mu_n \rightarrow \int_X f d\mu$ for every $f \in C(X)$. If μ_n is any sequence in \mathcal{M} and $F \subset C(X)$ is a dense countable subset, then, by a diagonal process, there is a subsequence μ_{n_j} such that $\int_X f d\mu_{n_j}$ converges for every $f \in F$,

and hence the sequence $\int_X g d\mu_{n_j}$ converges for every $g \in C(X)$. Therefore, \mathcal{M} is compact in the weak* topology. It is also convex: $t\mu + (1-t)\nu \in \mathcal{M}$ for any $t \in [0, 1]$ and $\mu, \nu \in \mathcal{M}$. A point in a convex set is *extreme* if it cannot be represented as a non-trivial convex combination of two other points. The extreme points of \mathcal{M} are the probability measures supported on points; they are called *Dirac measures*.

Let $\mathcal{M}_T \subset \mathcal{M}$ denote the set of all T -invariant Borel probability measures on X . Then \mathcal{M}_T is closed, and therefore compact, in the weak* topology, and convex.

Recall that if μ and ν are finite measures on a space X with σ -algebra \mathfrak{A} , then ν is *absolutely continuous* with respect to μ if $\nu(A) = 0$ whenever $\mu(A) = 0$, for $A \in \mathfrak{A}$. If ν is absolutely continuous with respect to μ , then the Radon–Nikodym theorem asserts that there is an L^1 function $dv/d\mu$, called the *Radon–Nikodym derivative*, such that $\nu(A) = \int_A (dv/d\mu)(x) d\mu$ for every $A \in \mathfrak{A}$ [Roy88].

PROPOSITION 4.6.2. *Ergodic T -invariant measures are precisely the extreme points of \mathcal{M}_T .*

Proof. If μ is not ergodic, then there is a T -invariant measurable subset $A \subset X$ with $0 < \mu(A) < 1$. Let $\mu_A(B) = \mu(B \cap A)/\mu(A)$ and $\mu_{X \setminus A}(B) = \mu(B \cap (X \setminus A))/\mu(X \setminus A)$ for any measurable set B . Then μ_A and $\mu_{X \setminus A}$ are T -invariant and $\mu = \mu(A)\mu_A + \mu(X \setminus A)\mu_{X \setminus A}$, so μ is not an extreme point.

Conversely, assume that μ is ergodic and that $\mu = t\nu + (1-t)\kappa$ with $\nu, \kappa \in \mathcal{M}_T$ and $t \in (0, 1)$. Then ν is absolutely continuous with respect to μ and $\nu(A) = \int_A r d\mu$, where $r = dv/d\mu \in L^1(X, \mu)$ is the Radon–Nikodym derivative. Observe that $r \leq \frac{1}{t}$ almost everywhere. Therefore $r \in L^2(X, \mu)$. Let U be the isometry of $L^2(X, \mu)$ given by $Uf = f \circ T$. Invariance of ν implies that for every $f \in L^2(X, \mu)$

$$\langle Uf, r \rangle_\mu = \int (f \circ T)r d\mu = \int fr d\mu = \langle f, r \rangle_\mu.$$

It follows that $\langle f, U^*r \rangle_\mu = \langle Uf, r \rangle_\mu = \langle f, r \rangle_\mu$, and hence $U^*r = r$. By Lemma 4.5.1 $Ur = r$. Since μ is ergodic, the function r is essentially constant, so $\mu = \nu = \kappa$. \square

By the Krein–Milman theorem [Roy88], [Rud91], \mathcal{M}_T is the closed convex hull of its extreme points. Therefore, the set \mathcal{M}_T^e of all T -invariant, ergodic, Borel probability measures is not empty. However, \mathcal{M}_T^e may be rather complicated; for example, it may be dense in \mathcal{M}_T in the weak* topology (Exercise 4.6.5).

Exercise 4.6.1. Describe \mathcal{M}_T and \mathcal{M}_T^e for the homeomorphism of the circle $T(x) = x + a \sin 2\pi x \pmod{1}$, $0 < a \leq \frac{1}{2\pi}$.

Exercise 4.6.2. Describe \mathcal{M}_T and \mathcal{M}_T^e for the homeomorphism of the torus $T(x, y) = (x, x + y) \pmod{1}$.

Exercise 4.6.3

- (a) Give an example of a map of the circle that is discontinuous at exactly one point and does not have non-trivial finite invariant Borel measures.
- (b) Give an example of a continuous map of the real line that does not have non-trivial finite invariant Borel measures.

Exercise 4.6.4. Let X and Y be compact metric spaces and $T: X \rightarrow Y$ a continuous map. Show that T induces a natural map $\mathcal{M}(X) \rightarrow \mathcal{M}(Y)$, and that this map is continuous in the weak* topology.

***Exercise 4.6.5.** Prove that if σ is the two-sided 2-shift, then \mathcal{M}_σ^e is dense in \mathcal{M}_σ in the weak* topology.

4.7 Unique Ergodicity and Weyl's Theorem³

In this section T is a continuous map of a compact metric space X . By §4.6, there are T -invariant Borel probability measures. If there is only one such measure, then T is said to be *uniquely ergodic*. Note that this unique invariant measure is necessarily ergodic by Proposition 4.6.2.

An irrational circle rotation is uniquely ergodic (Exercise 4.7.1). Moreover, any topologically transitive translation on a compact abelian group is uniquely ergodic (Exercise 4.7.2). On the other hand, unique ergodicity does not imply topological transitivity (Exercise 4.7.3).

PROPOSITION 4.7.1. *Let X be a compact metric space. A continuous map $T: X \rightarrow X$ is uniquely ergodic if and only if $S_f^n = \frac{1}{n} \sum_{i=0}^{n-1} f \circ T^i$ converges uniformly to a constant function S_f^∞ for any continuous function $f \in C(X)$.*

Proof. Suppose first that T is uniquely ergodic and μ is the unique T -invariant Borel probability measure. We will show that

$$\lim_{n \rightarrow \infty} \max_{x \in X} \left| S^n f(x) - \int_X f d\mu \right| \rightarrow 0.$$

³ The arguments of this section follow in part those of [Fur81a] and [CFS82].

Assume, for a contradiction, that there are $f \in C(X)$ and sequences $x_k \in X$ and $n_k \rightarrow \infty$ such that $\lim_{k \rightarrow \infty} S_f^{n_k}(x_k) = c \neq \int_X f d\mu$. As in the proof of Proposition 4.6.1, there is a subsequence $n_{k_i} \rightarrow \infty$ such that the limit $L(g) = \lim_{i \rightarrow \infty} S_f^{n_{k_i}}(x_{k_i})$ exists for any $g \in C(X)$. As in Proposition 4.6.1, L defines a T -invariant, positive, bounded linear functional on $C(X)$. By the Riesz representation theorem, $L(g) = \int_X g d\nu$ for some $\nu \in \mathcal{M}_T$. Since $L(f) = c \neq \int_X f d\mu$, the measures μ and ν are different, which contradicts unique ergodicity.

The proof of the converse is left as an exercise (Exercise 4.7.4). \square

Uniform convergence of the time averages of continuous functions does not, by itself, imply unique ergodicity. For example, if (X, T) is uniquely ergodic and $I = [0, 1]$, then $(X \times I, T \times \text{Id})$ is not uniquely ergodic, but the time averages converge uniformly for all continuous functions.

PROPOSITION 4.7.2. *Let T be a topologically transitive continuous map of a compact metric space X . Suppose that the sequence of time averages S_f^n converges uniformly for every continuous function $f \in C(X)$. Then T is uniquely ergodic.*

Proof. Since the convergence is uniform, $S_f^\infty = \lim_{n \rightarrow \infty} S_f^n$ is a continuous function. As in the proof of Proposition 4.6.1, $S_f^\infty(T(x)) = S_f^\infty(x)$ for every x . Since T is topologically transitive, S_f^∞ is constant. As in previous arguments, the linear functional $f \mapsto S_f^\infty$ defines a measure $\mu \in \mathcal{M}_T$ with $\int_X f d\mu = S_f^\infty$. Let $\nu \in \mathcal{M}_T$. By the Birkhoff ergodic theorem (Theorem 4.5.5), $S_f^\infty(x) = \int_X f d\nu$ for every $f \in C(X)$ and ν a.e. $x \in X$. Therefore, $\nu = \mu$. \square

Let X be a compact metric space with a Borel probability measure μ . Let $T: X \rightarrow X$ be a homeomorphism preserving μ . A point $x \in X$ is called *generic* for (X, μ, T) if for every continuous function f

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} f(T^k(x)) = \int_X f d\mu.$$

If T is ergodic, then by Corollary 4.5.8, μ -a.e. x is generic.

For a compact topological group G , the Haar measure on G is the unique Borel probability measure invariant under all left and right translations. Let $T: X \rightarrow X$ be a homeomorphism of a compact metric space, G a compact group, and $\phi: X \rightarrow G$ a continuous function. The homeomorphism $S: X \times G \rightarrow X \times G$ given by $S(x, g) = (T(x), \phi(x)g)$ is a *group extension* (or *G-extension*) of T . Observe that S commutes with the right translations $R_g(x, h) = (x, hg)$. If μ is a T -invariant measure on X and m is the Haar measure on G , then the product measure $\mu \times m$ is S -invariant (Exercise 4.7.7).

PROPOSITION 4.7.3 (Furstenberg). *Let G be a compact group with Haar measure m , X a compact metric space with a Borel probability measure μ , $T: X \rightarrow X$ a homeomorphism preserving μ , $Y = X \times G$, $\nu = \mu \times m$, and $S: Y \rightarrow Y$ a G -extension of T . If T is uniquely ergodic and S is ergodic, then S is uniquely ergodic.*

Proof. Since ν is R_g -invariant for every $g \in G$, if (x, h) is generic for ν , then (x, hg) is generic for ν . Since S is ergodic, ν -a.e. (x, h) is ν -generic. Therefore for μ -a.e. $x \in X$, the point (x, h) is ν -generic for every h . If a measure $\nu' \neq \nu$ is S -invariant and ergodic, then ν' -a.e. (x, h) is ν' -generic. The points that are ν' -generic cannot be ν -generic. Hence there is a subset $N \subset X$ such that $\mu(N) = 0$ and the first coordinate x of every ν' -generic point (x, h) lies in N . However, the projection of ν' to X is T -invariant and therefore is μ . This is a contradiction. \square

PROPOSITION 4.7.4. *Let $\alpha \in (0, 1)$ be irrational, and let $T: \mathbb{T}^k \rightarrow \mathbb{T}^k$ be defined by*

$$T(x_1, \dots, x_k) = (x_1 + \alpha, x_2 + a_{21}x_1, \dots, x_k + a_{k1}x_1 + \dots + a_{kk-1}x_{k-1}),$$

where the coefficients a_{ij} are integers and $a_{ii-1} \neq 0$, $i = 2, \dots, k$. Then T is uniquely ergodic.

Proof. By Exercise 4.7.8, T is ergodic with respect to Lebesgue measure on \mathbb{T}^k . An inductive application of Proposition 4.7.3 yields the result. \square

Let X be a compact topological space with a Borel probability measure μ . A sequence $(x_i)_{i \in \mathbb{N}}$ in X is *uniformly distributed* if for any continuous function f on X ,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n f(x_k) = \int_X f d\mu.$$

THEOREM 4.7.5 (Weyl). *If $P(x) = b_k x^k + \dots + b_0$ is a real polynomial such that at least one of the coefficients b_i , $i > 0$, is irrational, then the sequence $(P(n) \bmod 1)_{n \in \mathbb{N}}$ is uniformly distributed in $[0, 1]$.*

Proof [Fur81a]. Assume first that $b_k = \alpha/k!$ with α irrational. Consider the map $T: \mathbb{T}^k \rightarrow \mathbb{T}^k$ given by

$$T(x_1, \dots, x_k) = (x_1 + \alpha, x_2 + x_1, \dots, x_k + x_{k-1}).$$

Let $\pi: \mathbb{R}^k \rightarrow \mathbb{T}^k$ be the projection. Let $P_k(x) = P(x)$ and $P_{i-1}(x) = P_i(x+1) - P_i(x)$, $i = k, \dots, 1$. Then $P_1(x) = \alpha x + \beta$. Observe that $T^n(\pi(P_1(0), \dots, P_k(0))) = \pi(P_1(n), \dots, P_k(n))$. Since T is uniquely ergodic by Proposition 4.7.4, this orbit (and any other orbit) is uniformly distributed

on \mathbb{T}^k . It follows that the last coordinate $P_k(n) = P(n)$ is uniformly distributed on S^1 .

Exercise 4.7.9 finishes the proof. \square

Exercise 4.7.1. Prove that an irrational circle rotation is uniquely ergodic.

Exercise 4.7.2. Prove that any topologically transitive translation on a compact abelian group is uniquely ergodic.

Exercise 4.7.3. Prove that the diffeomorphism $T: S^1 \rightarrow S^1$ defined by $T(x) = x + a \sin^2(\pi x)$, $a < 1/\pi$, is uniquely ergodic but not topologically transitive.

Exercise 4.7.4. Prove the remaining statement of Proposition 4.7.1.

Exercise 4.7.5. Prove that the subshift defined by a fixed point a of a primitive substitution s is uniquely ergodic.

Exercise 4.7.6. Let T be a uniquely ergodic continuous transformation of a compact metric space X , and μ the unique invariant Borel probability measure. Show that $\text{supp } \mu$ is a minimal set for T .

Exercise 4.7.7. Let $S: X \times G \rightarrow X \times G$ be a G -extension of $T: (X, \mu) \rightarrow (X, \mu)$, and let m be the Haar measure on G . Prove that the product measure $\mu \times m$ is S -invariant.

Exercise 4.7.8. Use Fourier series on \mathbb{T}^k to prove that T from Proposition 4.7.4 is ergodic with respect to Lebesgue measure.

Exercise 4.7.9. Reduce the general case of Theorem 4.7.5 to the case where the leading coefficient is irrational.

4.8 The Gauss Transformation Revisited⁴

Recall that the Gauss transformation (§1.6) is the map of the unit interval to itself defined by

$$\phi(x) = \frac{1}{x} - \left[\frac{1}{x} \right] \quad \text{for } x \in (0, 1], \quad \phi(0) = 0.$$

The Gauss measure μ defined by

$$\mu(A) = \frac{1}{\log 2} \int_A \frac{dx}{1+x} \quad (4.3)$$

is a ϕ -invariant probability measure on $[0, 1]$.

⁴ The arguments of this section follow in part those of [Bil65].

For an irrational $x \in (0, 1]$, the n th entry $a_n(x) = [1/\phi^{n-1}(x)]$ of the continued fraction representing x is called the n -th *quotient*, and we write $x = [a_1(x), a_2(x), \dots]$. The irreducible fraction $p_n(x)/q_n(x)$ that is equal to the truncated continued fraction $[a_1(x), \dots, a_n(x)]$ is called the n th *convergent* of x . The numerators and denominators of the convergents satisfy the following relations:

$$p_0(x) = 0, \quad p_1(x) = 1, \quad p_n(x) = a_n(x)p_{n-1}(x) + p_{n-2}(x), \quad (4.4)$$

$$q_0(x) = 1, \quad q_1(x) = a_1(x), \quad q_n(x) = a_n(x)q_{n-1}(x) + q_{n-2}(x) \quad (4.5)$$

for $n > 1$. We have

$$x = \frac{p_n(x) + (\phi^n(x))p_{n-1}(x)}{q_n(x) + (\phi^n(x))q_{n-1}(x)}.$$

By an inductive argument

$$p_n(x) \geq 2^{(n-2)/2} \quad \text{and} \quad q_n(x) \geq 2^{(n-1)/2} \quad \text{for } n \geq 2,$$

and

$$p_{n-1}(x)q_n(x) - p_n(x)q_{n-1}(x) = (-1)^n, \quad n \geq 1. \quad (4.6)$$

For positive integers $b_k, k = 1, \dots, n$, let

$$\Delta_{b_1, \dots, b_n} = \{x \in (0, 1]: a_k(x) = b_k, k = 1, \dots, n\}.$$

The interval Δ_{b_1, \dots, b_n} is the image of the interval $[0, 1)$ under the map ψ_{b_1, \dots, b_n} defined by

$$\psi_{b_1, \dots, b_n}(t) = [b_1, \dots, b_{n-1}, b_n + t].$$

If n is odd, ψ_{b_1, \dots, b_n} is decreasing; if n is even, it is increasing. For $x \in \Delta_{b_1, \dots, b_n}$

$$x = \psi_{b_1, \dots, b_n}(t) = \frac{p_n + tp_{n-1}}{q_n + tq_{n-1}}, \quad (4.7)$$

where p_n and q_n are given by the recursive relations (4.4) and (4.5) with $a_n(x)$ replaced by b_n . Therefore

$$\Delta_{b_1, \dots, b_n} = \left[\frac{p_n}{q_n}, \frac{p_n + p_{n-1}}{q_n + q_{n-1}} \right) \quad \text{if } n \text{ is even,}$$

and

$$\Delta_{b_1, \dots, b_n} = \left(\frac{p_n + p_{n-1}}{q_n + q_{n-1}}, \frac{p_n}{q_n} \right] \quad \text{if } n \text{ is odd.}$$

If λ is Lebesgue measure, then $\lambda(\Delta_{b_1, \dots, b_n}) = (q_n(q_n + q_{n-1}))^{-1}$.

PROPOSITION 4.8.1. *The Gauss transformation is ergodic for the Gauss measure μ .*

Proof. For a measure ν and measurable sets A and B with $\nu(B) \neq 0$, let $\nu(A|B) = \nu(A \cap B)/\nu(B)$ denote the conditional measure. Fix b_1, \dots, b_n , and let $\Delta_n = \Delta_{b_1, \dots, b_n}$, $\psi_n = \psi_{b_1, \dots, b_n}$. The length of Δ_n is $\pm(\psi_n(1) - \psi_n(0))$, and for $0 \leq x < y \leq 1$,

$$\lambda(\{z: x \leq \phi^n(z) < y\} \cap \Delta_n) = \pm(\psi_n(y) - \psi_n(x)),$$

where the sign depends on the parity of n . Therefore

$$\lambda(\phi^{-n}([x, y]) | \Delta_n) = \frac{\psi_n(y) - \psi_n(x)}{\psi_n(1) - \psi_n(0)},$$

and by (4.6) and (4.7),

$$\lambda(\phi^{-n}([x, y]) | \Delta_n) = (y - x) \cdot \frac{q_n(q_n + q_{n-1})}{(q_n + xq_{n-1})(q_n + yq_{n-1})}.$$

The second factor in the right-hand side is between $1/2$ and 2 . Hence

$$\frac{1}{2}\lambda([x, y]) \leq \lambda(\phi^{-n}([x, y]) | \Delta_n) \leq 2\lambda([x, y]).$$

Since the intervals $[x, y]$ generate the σ -algebra,

$$\frac{1}{2}\lambda(A) \leq \lambda(\phi^{-n}(A) | \Delta_n) \leq 2\lambda(A) \quad (4.8)$$

for any measurable set $A \subset [0, 1]$.

Because the density of the Gauss measure μ is between $1/(2 \log 2)$ and $1/\log 2$,

$$\frac{1}{2 \log 2} \lambda(A) \leq \mu(A) \leq \frac{1}{\log 2} \lambda(A).$$

By (4.8),

$$\frac{1}{4} \mu(A) \leq \mu(\phi^{-n}(A) | \Delta_n) \leq 4\mu(A)$$

for any measurable $A \subset [0, 1]$.

Let A be a measurable ϕ -invariant set with $\mu(A) > 0$. Then $\frac{1}{4}\mu(A) \leq \mu(A | \Delta_n)$, or, equivalently, $\frac{1}{4}\mu(\Delta_n) \leq \mu(\Delta_n | A)$. Since the intervals Δ_n generate the σ -algebra, $\frac{1}{4}\mu(B) \leq \mu(B | A)$ for any measurable set B . By choosing $B = [0, 1] \setminus A$ we obtain that $\mu(A) = 1$. \square

The ergodicity of the Gauss transformation has the following number-theoretic consequences.

PROPOSITION 4.8.2. *For almost every $x \in [0, 1]$ (with respect to μ measure or Lebesgue measure), we have the following:*

1. Each integer $k \in \mathbb{N}$ appears in the sequence $a_1(x), a_2(x), \dots$ with asymptotic frequency

$$\frac{1}{\log 2} \log \left(\frac{k+1}{k} \right).$$

2. $\lim_{n \rightarrow \infty} \frac{1}{n} (a_1(x) + \dots + a_n(x)) = \infty.$
3. $\lim_{n \rightarrow \infty} \sqrt[n]{a_1(x)a_2(x) \cdots a_n(x)} = \prod_{k=1}^{\infty} \left(1 + \frac{1}{k^2 + 2k} \right)^{\log k / \log 2}.$
4. $\lim_{n \rightarrow \infty} \frac{\log q_n(x)}{n} = \frac{\pi^2}{12 \log 2}$

Proof. 1: Let f be the characteristic function of the semiopen interval $[1/k, 1/(k+1))$. Then $a_n(x) = k$ if and only if $f(\phi^n(x)) = 1$. By the Birkhoff ergodic theorem, for almost every x ,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} f(\phi^i(x)) = \int_0^1 f \, d\mu = \mu \left(\left[\frac{1}{k}, \frac{1}{k+1} \right) \right) = \frac{1}{\log 2} \log \left(\frac{k+1}{k} \right),$$

which proves the first assertion.

2: Let $f(x) = [1/x]$, i.e., $f(x) = a_1(x)$. Note that $\int_0^1 f(x)/(1+x) \, dx = \infty$, since $f(x) > (1-x)/x$ and $\int_0^1 \frac{(1-x)}{x(1+x)} \, dx = \infty$. For $N > 0$, define

$$f_N(x) = \begin{cases} f(x) & \text{if } f(x) \leq N, \\ 0 & \text{otherwise.} \end{cases}$$

Then, for any $N > 0$, for almost every x ,

$$\begin{aligned} \liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} f(\phi^k(x)) &\geq \liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} f_N(\phi^k(x)) \\ &= \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} f_N(\phi^k(x)) \\ &= \frac{1}{\log 2} \int_0^1 \frac{f_N(x)}{1+x} \, dx. \end{aligned}$$

Since $\lim_{N \rightarrow \infty} \int_0^1 \frac{f_N(x)}{1+x} \, dx \rightarrow \infty$, the conclusion follows.

3: Let $f(x) = \log a_1(x) = \log\left(\frac{1}{1+x}\right)$. Then $f \in L^1([0, 1])$ with respect to the Gauss measure μ (Exercise 4.8.1). By the Birkhoff ergodic theorem,

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n \log a_k(x) &= \frac{1}{\log 2} \int_0^1 \frac{f(x)}{1+x} dx \\ &= \frac{1}{\log 2} \sum_{k=1}^{\infty} \int_{\frac{1}{k+1}}^{\frac{1}{k}} \frac{\log k}{1+x} dx \\ &= \sum_{k=1}^{\infty} \frac{\log k}{\log 2} \cdot \log \left(1 + \frac{1}{k^2 + 2k}\right). \end{aligned}$$

Exponentiating this expression gives part 3.

4: Note that $p_n(x) = q_{n-1}(\phi(x))$ (Exercise 4.8.2), so

$$\frac{1}{q_n(x)} = \frac{p_n(x)}{q_n(x)} \frac{p_{n-1}(\phi(x))}{q_{n-1}(\phi(x))} \cdots \frac{p_1(\phi^{n-1}(x))}{q_1(\phi^{n-1}(x))}.$$

Thus

$$\begin{aligned} -\frac{1}{n} \log q_n(x) &= \frac{1}{n} \sum_{k=0}^{n-1} \log \left(\frac{p_{n-k}(\phi^k(x))}{q_{n-k}(\phi^k(x))} \right) \\ &= \frac{1}{n} \sum_{k=0}^{n-1} \log(\phi^k(x)) + \frac{1}{n} \sum_{k=0}^{n-1} \left(\log \frac{p_{n-k}(\phi^k(x))}{q_{n-k}(\phi^k(x))} - \log(\phi^k(x)) \right). \quad (4.9) \end{aligned}$$

It follows from the Birkhoff Ergodic Theorem that the first term of (4.9) converges a.e. to $(1/\log 2) \int_0^1 \log x/(1+x) dx = -\pi^2/12$. The second term converges to 0 (Exercise 4.8.2). \square

Exercise 4.8.1. Show that $\log([1/x]) \in L^1([0, 1])$ with respect to the Gauss measure μ .

Exercise 4.8.2. Show that $p_n(x) = q_n(\phi(x))$ and that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} \left(\log(\phi^k(x)) - \log \frac{p_{n-k}(\phi^k(x))}{q_{n-k}(\phi^k(x))} \right) = 0.$$

4.9 Discrete Spectrum

Let T be an automorphism of a probability space (X, \mathfrak{A}, μ) . The operator $U_T: L^2(X, \mathfrak{A}, \mu) \rightarrow L^2(X, \mathfrak{A}, \mu)$ is unitary, and each of its eigenvalues is a complex number of absolute value 1. Denote by Σ_T the set of all eigenvalues of U_T . Since constant functions are T -invariant, 1 is an eigenvalue of U_T . Any T -invariant function is an eigenfunction of U_T with eigenvalue 1.

Therefore, T is ergodic if and only if 1 is a simple eigenvalue of U_T . If f, g are two eigenfunctions with different eigenvalues $\sigma \neq \kappa$, then $\langle f, g \rangle = 0$, since $\langle f, g \rangle = \langle U_T f, U_T g \rangle = \sigma \bar{\kappa} \langle f, g \rangle$. Note that U_T is a multiplicative operator, i.e., $U_T(f \cdot g) = U_T(f) \cdot U_T(g)$, which has important implications for its spectrum.

PROPOSITION 4.9.1. Σ_T is a subgroup of the unit circle $S^1 = \{z \in \mathbb{C} : |z| = 1\}$. If T is ergodic, then every eigenvalue of U_T is simple.

Proof. If $\sigma \in \Sigma_T$ and $f(T(x)) = \sigma f(x)$, then $\bar{f}(T(x)) = \bar{\sigma} \bar{f}(x)$, and hence $\bar{\sigma} = \sigma^{-1} \in \Sigma_T$. If $\sigma_1, \sigma_2 \in \Sigma_T$ and $f_1(T(x)) = \sigma_1 f_1(x)$, $f_2(T(x)) = \sigma_2 f_2(x)$, then $f = f_1 f_2$ has eigenvalue $\sigma_1 \sigma_2$, and hence $\sigma_1 \sigma_2 \in \Sigma_T$. Therefore, Σ_T is a subgroup of S^1 .

If T is ergodic, the absolute value of any eigenfunction f is essentially constant (and non-zero). Thus, if f and g are eigenfunctions with the same eigenvalue σ , then f/g is in L^2 and is an eigenfunction with eigenvalue 1, so it is essentially constant by ergodicity. Therefore every eigenvalue is simple. \square

An ergodic automorphism T has *discrete spectrum* if the eigenfunctions of U_T span $L^2(X, \mathfrak{A}, \mu)$. An automorphism T has *continuous spectrum* if 1 is a simple eigenvalue of U_T and U_T has no other eigenvalues.

Consider a circle rotation $R_\alpha(x) = x + \alpha \pmod{1}$, $x \in [0, 1)$. For each $n \in \mathbb{Z}$, the function $f_n(x) = \exp(2\pi i n x)$ is an eigenfunction of U_{R_α} with eigenvalue $2\pi n \alpha$. If α is irrational, the eigenfunctions f_n span L^2 , and hence R_α has discrete spectrum. On the other hand, every weak mixing transformation has continuous spectrum (Exercise 4.9.1).

Let G be an abelian topological group. A *character* is a continuous homomorphism $\chi: G \rightarrow S^1$. The set of characters of G with the compact-open topology forms a topological group \hat{G} called the *group of characters* (or the *dual group*). For every $g \in G$, the evaluation map $\chi \mapsto \chi(g)$ is a character $\iota_g \in \hat{G}$, the dual of \hat{G} , and the map $\iota: G \rightarrow \hat{\hat{G}}$ is a homomorphism. If $\iota_g(\chi) \equiv 1$, then $\chi(g) = 1$ for every $\chi \in \hat{G}$, and hence ι is injective. By the Pontryagin duality theorem [Hel95], ι is also surjective and $\hat{\hat{G}} \cong G$. Moreover, if G is discrete, \hat{G} is a compact abelian group, and conversely.

For example, each character $\chi \in \hat{\mathbb{Z}}$ is completely determined by the value $\chi(1) \in S^1$. Therefore $\hat{\mathbb{Z}} \cong S^1$. On the other hand, if $\lambda \in \hat{S^1}$, then $\lambda: S^1 \rightarrow S^1$ is a homomorphism, so $\lambda(z) = z^n$ for some $n \in \mathbb{Z}$. Therefore, $\hat{S^1} = \mathbb{Z}$.

On a compact abelian group G with Haar measure λ , every character is in L^∞ , and therefore in L^2 . The integral of any non-trivial character with respect to Haar measure is 0 (Exercise 4.9.3). If σ and σ' are characters of

G , then $\sigma\overline{\sigma'}$ is also a character. If σ and σ' are different, then

$$\langle \sigma, \sigma' \rangle = \int_G \sigma(g)\overline{\sigma'(g)} d\lambda(g) = \int_G (\sigma\overline{\sigma'})(g) d\lambda(g) = 0.$$

Thus the characters of G are pairwise orthogonal in $L^2(G, \lambda)$.

THEOREM 4.9.2. *For every countable subgroup $\Sigma \subset S^1$ there is an ergodic automorphism T with discrete spectrum such that $\Sigma_T = \Sigma$.*

Proof. The identity character $\text{Id}: \Sigma \rightarrow S^1$, $\text{Id}(\sigma) = \sigma$, is a character of Σ . Let $T: \hat{\Sigma} \rightarrow \hat{\Sigma}$ be the translation $\chi \mapsto \chi \cdot \text{Id}$. The normalized Haar measure λ on $\hat{\Sigma}$ is invariant under T . For $\sigma \in \Sigma$, let $f_\sigma \in \hat{\Sigma}$ be the character of $\hat{\Sigma}$ such that $f_\sigma(\chi) = \chi(\sigma)$. Since

$$U_T f_\sigma(\chi) = f_\sigma(\chi \text{Id}) = f_\sigma(\chi) f_\sigma(\text{Id}) = \sigma f_\sigma(\chi),$$

f_σ is an eigenfunction with eigenvalue σ .

We claim that the linear span \mathcal{A} of the set of characters $\{f_\sigma: \sigma \in \Sigma\}$, is dense in $L^2(\hat{\Sigma}, \lambda)$, which will complete the proof. The set of characters separates points of $\hat{\Sigma}$, is closed under complex conjugation, and contains the constant function 1. Since the set of characters is closed under multiplication, \mathcal{A} is closed under multiplication, and is therefore an algebra. By the Stone–Weierstrass theorem [Roy88], \mathcal{A} is dense in $C(\Sigma, \mathbb{C})$, and therefore in $L^2(\hat{\Sigma}, \lambda)$. \square

The following theorem (which we do not prove) is a converse to Theorem 4.9.2.

THEOREM 4.9.3 (Halmos–von Neumann). *Let T be an ergodic automorphism with discrete spectrum, and let $\Sigma \subset S^1$ be its spectrum. Then T is isomorphic to the translation on $\hat{\Sigma}$ by the identity character $\text{Id}: \Sigma \rightarrow S^1$.*

A measure-preserving transformation $T: (X, \mathfrak{A}, \mu) \rightarrow (X, \mathfrak{A}, \mu)$ is *aperiodic* if $\mu(\{x \in X: T^n(x) = x\}) = 0$ for every $n \in \mathbb{N}$.

Theorem 4.9.4 (which we do not prove) implies that every aperiodic transformation can be approximated by a periodic transformation with an arbitrary period n . Many of the examples and counterexamples in abstract ergodic theory are constructed using the method of *cutting and stacking* based on this theorem.

THEOREM 4.9.4 (Rokhlin–Halmos [Hal60]). *Let T be an aperiodic automorphism of a Lebesgue probability space (X, \mathfrak{A}, μ) . Then for every $n \in \mathbb{N}$ and $\epsilon > 0$ there is a measurable subset $A = A(n, \epsilon) \subset X$ such that the sets $T^i(A)$, $i = 0, \dots, n-1$, are pairwise disjoint and $\mu(X \setminus \bigcup_{i=0}^{n-1} T^i(A)) < \epsilon$.*

Exercise 4.9.1. Prove that every weak mixing measure-preserving transformation has continuous spectrum.

Exercise 4.9.2. Suppose that $\alpha, \beta \in (0, 1)$ are irrational and α/β is irrational. Let T be the translation of \mathbb{T}^2 given by $T(x, y) = (x + \alpha, y + \beta)$. Prove that T is topologically transitive and ergodic and has discrete spectrum.

Exercise 4.9.3. Show that on a compact topological group G , the integral of any non-trivial character with respect to the Haar measure is 0.

4.10 Weak Mixing⁵

The property of weak mixing is typical in the following sense. Since each non-atomic probability Lebesgue space is isomorphic to the unit interval with Lebesgue measure λ , every measure-preserving transformation can be viewed as a transformation of $[0, 1]$ preserving λ . The *weak topology* on the set of all measure-preserving transformations of $[0, 1]$ is given by $T_n \rightarrow T$ if $\lambda(T_n(A) \Delta T(A)) \rightarrow 0$ for each measurable $A \subset [0, 1]$. Halmos showed [Hal44] that a residual (in the weak topology) subset of transformations are weak mixing. V. Rokhlin showed [Roh48] that the set of strong mixing transformations is of first category (in the weak topology).

The weak mixing transformations, as Theorem 4.10.6 below shows, are precisely those that have continuous spectrum. To show this we first prove a splitting theorem for isometries in a Hilbert space.

We say that a sequence of complex numbers $a_n, n \in \mathbb{Z}$ is *non-negative definite* if for each $N \in \mathbb{N}$,

$$\sum_{k,m=-N}^N z_k \bar{z}_m a_{k-m} \geq 0$$

for each finite sequence of complex numbers $z_k, -N \leq k \leq N$.

For a (linear) isometry U in a separable Hilbert space H , denote by U^* the adjoint of U , and for $n \geq 0$ set $U_n = U^n$ and $U_{-n} = U^{*n}$.

LEMMA 4.10.1. *For every $v \in H$, the sequence $\langle U_n v, v \rangle$ is non-negative definite.*

Proof.

$$\sum_{k,m=-N}^N z_k \bar{z}_m \langle U_{k-m} v, v \rangle = \sum_{k,m=-N}^N z_k \bar{z}_m \langle U_k v, U_m v \rangle = \left\| \sum_{l=-N}^N z_l U_l v \right\|^2. \quad \square$$

⁵ The presentation of this section to a large extent follows §2.3 of [Kre85].

LEMMA 4.10.2 (Wiener). For a finite measure ν on $[0, 1)$ set $\hat{\nu}_k = \int_0^1 e^{2\pi i k x} \nu(dx)$. Then $\lim_{n \rightarrow \infty} n^{-1} \sum_{k=0}^{n-1} |\hat{\nu}_k| = 0$ if and only if ν has no atoms.

Proof. Observe that $n^{-1} \sum_{k=0}^{n-1} |\hat{\nu}_k| \rightarrow 0$ if and only if $n^{-1} \sum_{k=0}^{n-1} |\hat{\nu}_k|^2 \rightarrow 0$.
Now

$$\begin{aligned} \frac{1}{n} \sum_{k=0}^{n-1} |v_k|^2 &= \frac{1}{n} \sum_{k=0}^{n-1} \int_0^1 e^{2\pi i k x} \nu(dx) \int_0^1 e^{-2\pi i k y} \nu(dy) \\ &= \int_0^1 \int_0^1 \left[\frac{1}{n} \sum_{k=0}^{n-1} e^{2\pi i k(x-y)} \right] \nu(dx) \nu(dy). \end{aligned}$$

The functions $n^{-1} \sum_{k=0}^{n-1} \exp(2\pi i k(x-y))$ are bounded in absolute value by 1 and converge to 1 for $x = y$ and to 0 for $x \neq y$. Therefore the last integral tends to the product measure $\nu \times \nu$ of the diagonal of $[0, 1) \times [0, 1)$. It follows that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} |v_k|^2 = \sum_{0 \leq x < 1} (\nu(\{x\}))^2. \quad \square$$

For a (linear) isometry U of a separable Hilbert space H , set

$$H_w(U) = \left\{ v \in H : \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} |\langle U^k v, v' \rangle| = 0 \text{ for each } v' \in H \right\},$$

and denote by $H_e(U)$ the closure of the subspace spanned by the eigenvectors of U . Both $H_w(U)$ and $H_e(U)$ are closed and U -invariant.

PROPOSITION 4.10.3. Let U be a (linear) isometry of a separable Hilbert space H . Then

1. For each $v \in H$, there is a unique finite measure ν_v on the interval $[0, 1)$ (called the spectral measure) such that for every $n \in \mathbb{Z}$

$$\langle U_n v, v \rangle = \int_0^1 e^{2\pi i n x} \nu_v(dx).$$

2. If v is an eigenvector of U with eigenvalue $\exp(2\pi i \alpha)$, then ν_v consists of a single atom at α of measure 1.
3. If $v \perp H_e(U)$, then ν_v has no atoms and $v \in H_w(U)$.

Proof. The first statement follows immediately from Lemma 4.10.1 and the spectral theorem for isometries in a Hilbert space [Hel95], [Fol95]. The second statement follows from the first (Exercise 4.10.3).

To prove the last statement let $v \perp H_e$ and $W = e^{-2\pi i x} U$. Applying the von Neumann ergodic theorem 4.5.2, let $u = \lim_{n \rightarrow \infty} n^{-1} \sum_{k=0}^{n-1} W^k v$. Then

$Wu = u$. By Proposition 4.10.3,

$$\begin{aligned} \langle u, v \rangle &= \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} e^{-2\pi i x k} \langle U^k v, v \rangle \\ &= \lim_{n \rightarrow \infty} \int_0^1 \frac{1}{n} \sum_{k=0}^{n-1} e^{-2\pi i(x-y)k} v_v(dy) = v_v(x). \end{aligned}$$

If $v_v(x) > 0$, then u is a non-zero eigenvector of U with eigenvalue $e^{2\pi i x}$ and $v \not\perp u$, which is a contradiction. Therefore $v_v(x) = 0$ for each x , and Lemma 4.10.2 completes the proof. \square

For a finite subset $B \subset \mathbb{N}$ denote by $|B|$ the cardinality of B . For a subset $A \subset \mathbb{N}$, define the *upper density* $\bar{d}(A)$ by

$$\bar{d}(A) = \limsup_{n \rightarrow \infty} \frac{1}{n} |A \cap [1, n]|.$$

We say that a sequence b_n *converges in density* to b and write $\text{d-lim}_n b_n = b$ if there is a subset $A \subset \mathbb{N}$ such that $\bar{d}(A) = 0$ and $\lim_{n \rightarrow \infty, n \notin A} b_n = b$.

LEMMA 4.10.4. *If (b_n) is a bounded sequence, then $\text{d-lim}_n b_n = 0$ if and only if $\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} |b_n - b| = 0$.*

Proof. Exercise 4.10.1. \square

The following splitting theorem is an immediate consequence of Proposition 4.10.3.

THEOREM 4.10.5 (Koopman–von Neumann [KvN32]). *Let U be an isometry in a separable Hilbert space H . Then $H = H_e \oplus H_w$. A vector $v \in H$ lies in $H_w(U)$ if and only if $\text{d-lim}_n \langle U^n v, v \rangle = 0$, and if and only if $\text{d-lim}_n \langle U^n v, w \rangle = 0$ for each $w \in H$.*

Proof. The splitting follows from Proposition 4.10.3. To prove the remaining statement that $\text{d-lim} \langle U^n v, v \rangle = 0$ if and only if $\text{d-lim} \langle U^n v, w \rangle = 0$, observe that $\langle U^n v, w \rangle \equiv 0$ if $v \perp U^k v$ for all $k \in \mathbb{N}$. If $w = U^k v$, then $\langle U^n v, w \rangle = \langle U^n v, U^k v \rangle = \langle U^{n-k} v, v \rangle$. \square

Recall that if T and S are measure-preserving transformations in finite measure spaces (X, \mathfrak{A}, μ) and (Y, \mathfrak{B}, ν) , then $T \times S$ is a measure-preserving transformation in the product space $(X \times Y, \mathfrak{A} \times \mathfrak{B}, \mu \times \nu)$. As in §4.9, we denote by U_T the isometry $U_T f(x) = f(T(x))$ of $L^2(X, \mathfrak{A}, \mu)$.

THEOREM 4.10.6. *Let T be a measure-preserving transformation of a probability space (X, \mathfrak{A}, μ) . Then the following are equivalent:*

1. T is weak mixing.
2. T has continuous spectrum.
3. $\text{d-lim}_n \int_X f(T^n(x)) \overline{f(x)} d\mu = 0$ if $f \in L^2(X, \mathfrak{A}, \mu)$ and $\int_X f d\mu = 0$.
4. $\text{d-lim}_n \int_X f(T^n(x)) \overline{g(x)} d\mu = \int_X f d\mu \cdot \int_X g d\mu$ for all functions $f, g \in L^2(X, \mathfrak{A}, \mu)$.
5. $T \times T$ is ergodic.
6. $T \times S$ is weak mixing for each weak mixing S .
7. $T \times S$ is ergodic for each ergodic S .

Proof. The transformation T is weak mixing if and only if $H_w(U_T)$ is the orthogonal complement of the constants in $L^2(X, \mathfrak{A}, \mu)$. Therefore, by Proposition 4.10.3, $1 \Leftrightarrow 2$. By Lemma 4.10.4, $1 \Leftrightarrow 3$. Clearly $4 \Rightarrow 3$. Assume that 3 holds. It is enough to show 4 for f with $\int_X f d\mu = 0$. Observe that 4 holds for g satisfying $\int_X f(T^k(x)) \overline{g(x)} d\mu = 0$ for all $k \in \mathbb{N}$. Hence it suffices to consider $g(x) = \overline{f(T^k(x))}$. But $\int_X f(T^n(x)) \overline{f(T^k(x))} d\mu = \int_X f(T^{n-k}(x)) \overline{f(x)} d\mu \rightarrow 0$ as $n \rightarrow \infty$ by 3. Therefore $3 \Leftrightarrow 4$.

Assume 5. Observe that T is ergodic and if U_T has an eigenfunction f , then $|f|$ is T -invariant, and hence constant. Therefore $f(x)/f(y)$ is $T \times T$ -invariant and $5 \Rightarrow 2$. Clearly $6 \Rightarrow 2$ and $7 \Rightarrow 5$.

Assume 3. To prove 7 observe that $L^2(X \times Y, \mathfrak{A} \times \mathfrak{B}, \mu \times \nu)$ is spanned by functions of the form $f(x)g(y)$. Let $\int_X f d\mu = \int_Y g d\nu = 0$. Then

$$\begin{aligned} & \int_{X \times Y} f(T^n(x)) \overline{g(S^n(y))} f(x) \overline{g(y)} d\mu \times \nu \\ &= \int_X f(T^n(x)) \overline{f(x)} d\mu \cdot \int_Y g(S^n(y)) \overline{g(y)} d\nu. \end{aligned}$$

The first integral on the right-hand side converges in density to 0 by part 3, while the second one is bounded. Therefore the product converges in density to 0, and part 7 follows. The proof of $3 \Rightarrow 6$ is similar (Exercise 4.10.4). \square

Exercise 4.10.1. Let (b_n) be a bounded sequence. Prove that $\text{d-lim } b_n = b$ if and only if $\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} |b_n - b| = 0$.

Exercise 4.10.2. Prove that d-lim has the usual arithmetic properties of limits.

Exercise 4.10.3. Prove the second statement of Proposition 4.10.3.

Exercise 4.10.4. Prove that $3 \Rightarrow 6$ in Theorem 4.10.6.

Exercise 4.10.5. Let T be a weak mixing measure-preserving transformation, and let S be a measure-preserving transformation such that $S^k = T$ for some $k \in \mathbb{N}$ (S is called a k th root of T). Prove that S is weak mixing.

4.11 Applications of Measure-Theoretic Recurrence to Number Theory

In this section we give highlights of applications of measure-theoretic recurrence to number theory initiated by H. Furstenberg. As an illustration of this approach we prove Sárközy's Theorem (Theorem 4.11.5). Our exposition follows to a large extent [Fur77] and [Fur81a].

For a finite subset $F \subset \mathbb{Z}$, denote by $|F|$ the number of elements in F . A subset $D \subset \mathbb{Z}$ has *positive upper density* if there are $a_n, b_n \in \mathbb{Z}$ such that $b_n - a_n \rightarrow \infty$ and for some $\delta > 0$,

$$\frac{|D \cap [a_n, b_n]|}{b_n - a_n + 1} > \delta \quad \text{for all } n \in \mathbb{N}.$$

Let $D \subset \mathbb{Z}$ have positive upper density. Let $\omega_D \in \Sigma_2 = \{0, 1\}^{\mathbb{Z}}$ be the sequence for which $(\omega_D)_n = 1$ if $n \in D$ and $(\omega_D)_n = 0$ if $n \notin D$, and let X_D be the closure of its orbit under the shift σ in Σ_2 . Set $Y_D = \{\omega \in X_D : \omega_0 = 1\}$.

PROPOSITION 4.11.1 (Furstenberg). *Let $D \subset \mathbb{Z}$ have positive upper density. Then there exists a shift-invariant Borel probability measure μ on X_D such that $\mu(Y_D) > 0$.*

Proof. By §4.6, every σ -invariant Borel probability measure on X_D is a linear functional L on the space $C(X_D)$ of continuous functions on X_D such that $L(f) \geq 0$ if $f \geq 0$, $L(1) = 1$, and $L(f \circ \sigma) = L(f)$.

For a function $f \in C(X_D)$, set

$$L_n(f) = \frac{1}{b_n - a_n + 1} \sum_{i=a_n}^{b_n} f(\sigma^i(\omega_D)),$$

where a_n, b_n , and δ are associated with D as in the preceding paragraph. Observe that $L_n(f) \leq \max f$ for each n . Let $(f_j)_{j \in \mathbb{N}}$ be a countable dense subset in $C(X_D)$. By a diagonal process, one can find a sequence $n_k \rightarrow \infty$ such that $\lim_{k \rightarrow \infty} L_{n_k}(f_j)$ exists for each j . Since $(f_j)_{j \in \mathbb{N}}$ is dense in $C(X_D)$, we have that

$$L(f) = \lim_{k \rightarrow \infty} \frac{1}{b_{n_k} - a_{n_k} + 1} \sum_{i=a_{n_k}}^{b_{n_k}} f(\sigma^i(\omega_D))$$

exists for each $f \in C(X_D)$ and determines a σ -invariant Borel probability measure μ .

Let $\chi \in C(X_D)$ be the characteristic function of Y_D . Then

$$L(\chi) = \int \chi d\mu = \mu(Y_D) > 0. \quad \square$$

PROPOSITION 4.11.2. *Let $p(k)$ be a polynomial with integer coefficients and $p(0) = 0$. Let U be an isometry of a separable Hilbert space H , and $H_{\text{rat}} \subset H$ be the closure of the subspace spanned by the eigenvectors of U whose eigenvalues are roots of 1. Suppose $v \in H$ is such that $\langle U^{p(k)}v, v \rangle = 0$ for all $k \in \mathbb{N}$. Then $v \perp H_{\text{rat}}$.*

Proof. Let $v = v_{\text{rat}} + w$ with $v_{\text{rat}} \in H_{\text{rat}}$ and $w \perp H_{\text{rat}}$. We use the following lemma, whose proof is similar to the proof of Lemma 4.10.2 (Exercise 4.11.1). \square

LEMMA 4.11.3. $\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} U^{p(k)}w = 0$ for all $w \perp H_{\text{rat}}$.

Fix $\epsilon > 0$, and let $v'_{\text{rat}} \in H_{\text{rat}}$ and m be such that $\|v_{\text{rat}} - v'_{\text{rat}}\| < \epsilon$ and $U^m v'_{\text{rat}} = v'_{\text{rat}}$. Then $\|U^{mk}v_{\text{rat}} - v_{\text{rat}}\| < 2\epsilon$ for each k and, since $p(mk)$ is divisible by m ,

$$\left\| \frac{1}{n} \sum_{k=0}^{n-1} U^{p(mk)}v_{\text{rat}} - v_{\text{rat}} \right\| < 2\epsilon.$$

Since $(1/n) \sum_{k=0}^{n-1} U^{p(mk)}w \rightarrow 0$ by Lemma 4.11.3, for n large enough we have

$$\left\| \frac{1}{n} \sum_{k=0}^{n-1} U^{p(mk)}v - v_{\text{rat}} \right\| < 2\epsilon.$$

By assumption, $\langle U^{p(mk)}v, v \rangle = 0$. Hence $|\langle v_{\text{rat}}, v \rangle| < 2\epsilon\|v\|$, so $\langle v_{\text{rat}}, v \rangle = 0$. \square

As a corollary of the preceding proposition we obtain Furstenberg's polynomial recurrence theorem.⁶

THEOREM 4.11.4 (Furstenberg). *Let $p(t)$ be a polynomial with integer coefficients and $p(0) = 0$. Let T be a measure-preserving transformation of a finite measure space (X, \mathfrak{A}, μ) , and $A \in \mathfrak{A}$ be a set with positive measure. Then there is $n \in \mathbb{N}$ such that $\mu(A \cap T^{p(n)}A) > 0$.*

Proof. Let U be the isometry induced by T in $H = L^2(X, \mathfrak{A}, \mu)$, $(Uh)(x) = h(T^{-1}(x))$. If $\mu(A \cap T^{p(n)}A) = 0$ for each $n \in \mathbb{N}$, then the characteristic function χ_A of A satisfies $\langle U^{p(n)}\chi_A, \chi_A \rangle = 0$ for each n . By Proposition 4.11.2, χ_A is orthogonal to all eigenfunctions of U whose eigenvalues are roots of 1. However $\mathbf{1}(x) \equiv 1$ is an eigenfunction of U with eigenvalue 1 and $\langle \mathbf{1}, \chi_A \rangle = \mu(A) \neq 0$. \square

⁶ A slight modification of the arguments above yields Proposition 4.11.2 and Theorem 4.11.4 for polynomials with integer values at integer points (rather than integer coefficients).

Theorem 4.11.4 and Proposition 4.11.1 imply the following result in combinatorial number theory.

THEOREM 4.11.5 (Sárközy [Sár78]). *Let $D \subset \mathbb{Z}$ have positive upper density, and let p be a polynomial with integer coefficients and $p(0) = 0$. Then there are $x, y \in D$ and $n \in \mathbb{N}$ such that $x - y = p(n)$.*

The following extension of the Poincaré recurrence theorem (whose proof is beyond the scope of this book) was used by Furstenberg to give an ergodic-theoretic proof of the Szemerédi theorem on arithmetic progressions (Theorem 4.11.7).

THEOREM 4.11.6 (Furstenberg's Multiple Recurrence Theorem [Fur77]). *Let T be an automorphism of a probability space (X, \mathfrak{A}, μ) . Then for every $n \in \mathbb{N}$ and every $A \in \mathfrak{A}$ with $\mu(A) > 0$ there is $k \in \mathbb{N}$ such that*

$$\mu(A \cap T^{-k}(A) \cap T^{-2k}(A) \cap \dots \cap T^{-nk}(A)) > 0.$$

THEOREM 4.11.7 (Szemerédi [Sze69]). *Every subset $D \subset \mathbb{Z}$ of positive upper density contains arbitrarily long arithmetic progressions.*

Proof. Exercise 4.11.3. □

Exercise 4.11.1. Prove Lemma 4.11.3.

Exercise 4.11.2. Use Theorem 4.11.4 and Proposition 4.11.1 to prove Theorem 4.11.5.

Exercise 4.11.3. Use Proposition 4.11.1 and Theorem 4.11.6 to prove Theorem 4.11.7.

4.12 Internet Search⁷

In this section, we describe a surprising application of ergodic theory to the problem of searching the Internet. This approach is used by the Internet search engine GoogleTM (www.google.com).

The Internet offers enormous amounts of information. Looking for information on the Internet is analogous to looking for a book in a huge library without a catalog. The task of locating information on the web is performed by search engines. The first search engines appeared in the early 1990s. The most popular engines handle tens of millions of searches per day.

⁷ The exposition in this section follows to a certain extent that of [BP98].

The main tasks performed by a search engine are: gathering information from web pages; processing and storing this information in a database; and producing from this database a list of web pages relevant to a query consisting of one or more words. The gathering of information is performed by robot programs called *crawlers* that “crawl” the web by following links embedded in web pages. Raw information collected by the crawlers is parsed and coded by the *indexer*, which produces, for each web page, a set of word occurrences (including word position, font type, and capitalization) and records all links from this web page to other pages, thus creating the *forward index*. The *sorter* rearranges information by words (rather than web pages), thus creating the *inverted index*. The *searcher* uses the inverted index to answer the query, i.e., to compile a list of documents relevant to the keywords and phrases of the query.

The order of the documents on the list is extremely important. A typical list may contain tens of thousands of web pages, but at best only the first several dozen may be reviewed by the user. Google uses two characteristics of the web page to determine the order of the returned pages – the relevance of the document to the query and the *PageRank* of the web page. The relevance is based on the relative position, fontification, and frequency of the keyword(s) in the document. This factor by itself often does not produce good search results. For example, a query on the word “Internet” in one of the early search engines returned a list whose first entry was a web page in Chinese containing no English words other than “Internet.” Even now, many search engines return barely relevant results when searching on common terms.

Google uses Markov chains to rank web pages. The collection of all web pages and links between them is viewed as a directed graph G in which the web pages serve as vertices and the links as directed edges (from the web page on which they appear to the web page to which they point). At the moment there are about 1.5 billion web pages with about 10 times as many links. We number the vertices with positive integers $i = 1, 2, \dots, N$. Let \tilde{G} be the graph obtained from G by adding a vertex 0 with edges to and from all other vertices. Let $b_{ij} = 1$ if there is an edge from vertex i to vertex j in \tilde{G} , and let $O(i)$ be the number of outgoing edges adjacent to vertex i in \tilde{G} . Note that $O(i) > 0$ for all i . Fix a damping parameter $p \in (0, 1)$ (for example, $p = .75$). Set $B_{ii} = 0$ for $i \geq 0$. For $i, j > 0$ and $i \neq j$ set

$$B_{0i} = \frac{1}{N}, \quad B_{i0} = \begin{cases} 1 & \text{if } O(i) = 1, \\ 1 - p & \text{if } O(i) \neq 1, \end{cases} \quad B_{ij} = \begin{cases} 0 & \text{if } b_{ij} = 0, \\ \frac{p}{O(i)} & \text{if } b_{ij} = 1. \end{cases}$$

The matrix B is stochastic and primitive. Therefore, by Corollary 3.3.3, it has a unique positive left eigenvector q with eigenvalue 1 whose entries

add up to 1. The pair (B, q) is a Markov chain on the vertices of \tilde{G} . Google interprets q_i as the PageRank of web page i and uses it together with the relevance factor of the page to determine how high on the return list this page should be.

For any initial probability distribution q' on the vertices of \tilde{G} , the sequence $q' B^n$ converges exponentially to q . Thus one can find an approximation for q by computing $p B^n$, where q' is the uniform distribution. This approach to finding q is computationally much easier than trying to find an eigenvector for a matrix with 1.5 billion rows and columns.

Exercise 4.12.1. Let A be an $N \times N$ stochastic matrix, and let A_{ij}^n be the entries of A^n , i.e., A_{ij}^n is the probability of going from i to j in exactly n steps (§4.4). Suppose q is an invariant probability distribution, $qA = q$.

- (a) Suppose that for some j , we have $A_{ij} = 0$ for all $i \neq j$, and $A_{jk}^n > 0$ for some $k \neq j$ and some $n \in \mathbb{N}$. Show that $q_j = 0$.
- (b) Prove that if $A_{ij} > 0$ for some $j \neq i$ and $A_{ji}^n = 0$ for all $n \in \mathbb{N}$, then $q_i = 0$.

Hyperbolic Dynamics

In Chapter 1, we saw several examples of dynamical systems that were locally linear and had complementary expanding and/or contracting directions: expanding endomorphisms of S^1 , hyperbolic toral automorphisms, the horseshoe, and the solenoid. In this chapter, we develop the theory of *hyperbolic differentiable dynamical systems*, which include these examples. Locally, a differentiable dynamical system is well approximated by a linear map – its derivative. Hyperbolicity means that the derivative has complementary expanding and contracting directions.

The proper setting for a differentiable dynamical system is a *differentiable manifold* with a differentiable map, or flow. A detailed introduction to the theory of differentiable manifolds is beyond the scope of this book. For the convenience of the reader, we give a brief formal introduction to differentiable manifolds in §5.13, and an even briefer informal introduction here.

For the purposes of this book, and without loss of generality (see the embedding theorems in [Hir94]), it suffices to think of a differentiable manifold M^n as an n -dimensional differentiable surface, or *submanifold*, in \mathbb{R}^N , $N > n$. The implicit function theorem implies that each point in M has a *local coordinate system* that identifies a neighborhood of the point with a neighborhood of 0 in \mathbb{R}^n . For each point x on such a surface $M \subset \mathbb{R}^N$, the tangent space $T_x M \subset \mathbb{R}^N$ is the space of all vectors tangent to M at x . The standard inner product on \mathbb{R}^N induces an inner product $\langle \cdot, \cdot \rangle_x$ on each $T_x M$. The collection of inner products is called a *Riemannian metric*, and a manifold M together with a Riemannian metric is called a *Riemannian manifold*. The (intrinsic) distance d between two points in M is the infimum of the lengths of differentiable curves in M connecting the two points.

A one-to-one differentiable mapping with a differentiable inverse is called a *diffeomorphism*.

A discrete-time differentiable dynamical system on a differentiable manifold M is a differentiable map $f: M \rightarrow M$. The derivative df_x is a linear map from $T_x M$ to $T_{f(x)} M$. In local coordinates df_x is given by the matrix of partial derivatives of f . A continuous-time differentiable dynamical system on M is a *differentiable flow*, i.e., a one-parameter group $\{f^t\}$, $t \in \mathbb{R}$, of differentiable maps $f^t: M \rightarrow M$ that depend differentiably on t . Since $f^{-t} \circ f^t = \text{Id}$, each map f^t is a diffeomorphism. The derivative

$$v(\cdot) = \left. \frac{d}{dt} f^t(\cdot) \right|_{t=0}$$

is a *differentiable vector field* tangent to M , and the flow $\{f^t\}$ is the one-parameter group of time- t maps of the differential equation $\dot{x} = v(x)$.

Differentiability, and even subtle differences in the degree of differentiability, have important and sometimes surprising consequences. See, for example, Exercise 2.5.7 and §7.2.

5.1 Expanding Endomorphisms Revisited

To illustrate and motivate some of the main ideas of this chapter we consider again expanding endomorphisms of the circle $E_m x = mx \pmod{1}$, $x \in [0, 1)$, $m > 1$, introduced in §1.3.

Fix $\epsilon < 1/2$. A finite or infinite sequence of points (x_i) in the circle is called an ϵ -orbit of E_m if $d(x_{i+1}, E_m x_i) < \epsilon$ for all i . The point x_i has m preimages under E_m that are uniformly spread on the circle. Exactly one of them, y_i^{i-1} , is closer than ϵ/m to x_{i-1} . Similarly, y_i^{i-1} has m preimages under E_m ; exactly one of them, y_i^{i-2} , is closer than ϵ/m to x_{i-2} . Continuing in this manner, we obtain a point y_i^0 with the property that $d(E_m^j y_i^0, x_j) < \epsilon$ for $0 \leq j \leq i$. In other words, any finite ϵ -orbit of E_m can be approximated or *shadowed* by a real orbit. If $(x_i)_{i=0}^\infty$ is an infinite ϵ -orbit, then the limit $y = \lim_{i \rightarrow \infty} y_i^0$ exists (Exercise 5.1.1), and $d(E_m^i y, x_i) \leq \epsilon$ for $i \geq 0$. Since two different orbits of E_m diverge exponentially, there can be only one shadowing orbit for a given infinite ϵ -orbit. By construction, y depends continuously on (x_i) in the product topology (Exercise 5.1.2).

The above discussion of the ϵ -orbits of E_m is based solely on the uniform forward expansion of E_m . Similar arguments show that if f is C^1 -close to E_m , then each infinite ϵ -orbit (x_i) of f is shadowed by a unique real orbit of f that depends continuously on (x_i) (Exercise 5.1.3).

Consider now f that is C^1 -close enough to E_m . View each orbit $(f^i(x))$ as an ϵ -orbit of E_m . Let $y = \phi(x)$ be the unique point whose orbit $(E_m^i y)$ shadows $(f^i(x))$. By the above discussion, the map ϕ is a homeomorphism and

$E_m\phi(x) = \phi(f(x))$ for each x (Exercise 5.1.4). This means that any differentiable map that is C^1 -close enough to E_m is topologically conjugate to E_m . In other words, E_m is *structurally stable*; see §5.5 and §5.11.

Hyperbolicity is characterized by local expansion and contraction, in complementary directions. This property, which causes local instability of orbits, surprisingly leads to the global stability of the topological pattern of the collection of all orbits.

Exercise 5.1.1. Prove that $\lim_{i \rightarrow \infty} y_i^0$ exists.

Exercise 5.1.2. Prove that $\lim_{i \rightarrow \infty} y_i^0$ depends continuously on (x_i) in the product topology.

Exercise 5.1.3. Prove that if f is C^1 -close to E_m , then each infinite ϵ -orbit (x_i) of f is approximated by a unique real orbit of f that depends continuously on (x_i) .

Exercise 5.1.4. Prove that ϕ is a homeomorphism conjugating f and E_m .

5.2 Hyperbolic Sets

In this section, M is a C^1 Riemannian manifold, $U \subset M$ a non-empty open subset, and $f: U \rightarrow f(U) \subset M$ a C^1 diffeomorphism. A compact, f -invariant subset $\Lambda \subset U$ is called *hyperbolic* if there are $\lambda \in (0, 1)$, $C > 0$, and families of subspaces $E^s(x) \subset T_x M$ and $E^u(x) \subset T_x M$, $x \in \Lambda$, such that for every $x \in \Lambda$,

1. $T_x M = E^s(x) \oplus E^u(x)$,
2. $\|df_x^n v^s\| \leq C\lambda^n \|v^s\|$ for every $v^s \in E^s(x)$ and $n \geq 0$,
3. $\|df_x^{-n} v^u\| \leq C\lambda^n \|v^u\|$ for every $v^u \in E^u(x)$ and $n \geq 0$,
4. $df_x E^s(x) = E^s(f(x))$ and $df_x E^u(x) = E^u(f(x))$.

The subspace $E^s(x)$ (respectively, $E^u(x)$) is called the *stable (unstable) subspace* at x , and the family $\{E^s(x)\}_{x \in \Lambda}$ ($\{E^u(x)\}_{x \in \Lambda}$) is called the *stable (unstable) distribution* of $f|_\Lambda$. The definition allows the two extreme cases $E^s(x) = \{0\}$ or $E^u(x) = \{0\}$.

The horseshoe (§1.8) and the solenoid (§1.9) are examples of hyperbolic sets. If $\Lambda = M$, then f is called an *Anosov diffeomorphism*. Hyperbolic toral automorphisms (§1.7) are examples of Anosov diffeomorphisms. Any closed invariant subset of a hyperbolic set is a hyperbolic set.

PROPOSITION 5.2.1. *Let Λ be a hyperbolic set of f . Then the subspaces $E^s(x)$ and $E^u(x)$ depend continuously on $x \in \Lambda$.*

Proof. Let x_i be a sequence of points in Λ converging to $x_0 \in \Lambda$. Passing to a subsequence, we may assume that $\dim E^s(x_i)$ is constant. Let $w_{1,i}, \dots, w_{k,i}$ be an orthonormal basis in $E^s(x_i)$. The restriction of the unit tangent bundle T^1M to Λ is compact. Hence, by passing to a subsequence, $w_{j,i}$ converges to $w_{j,0} \in T_{x_0}^1M$ for each $j = 1, \dots, k$. Since condition 2 of the definition of a hyperbolic set is a closed condition, each vector from the orthonormal frame $w_{1,0}, \dots, w_{k,0}$ satisfies condition 2 and, by the invariance (condition 4), lies in $E^s(x_0)$. It follows that $\dim E^s(x_0) \geq k = \dim E^s(x_i)$. A similar argument shows that $\dim E^u(x_0) \geq \dim E^u(x_i)$. Hence, by (1), $\dim E^s(x_0) = \dim E^s(x_i)$ and $\dim E^u(x_0) = \dim E^u(x_i)$, and continuity follows. \square

Any two Riemannian metrics on M are equivalent on a compact set, in the sense that the ratios of the lengths of non-zero vectors are bounded above and away from zero. Thus the notion of a hyperbolic set does not depend on the choice of the Riemannian metric on M . The constant C depends on the metric, but λ does not (Exercise 5.2.2). However, as the next proposition shows, we can choose a particularly nice metric and $C = 1$ by using a slightly larger λ .

PROPOSITION 5.2.2. *If Λ is a hyperbolic set of f with constants C and λ , then for every $\epsilon > 0$ there is a C^1 Riemannian metric $\langle \cdot, \cdot \rangle'$ in a neighborhood of Λ , called the Lyapunov, or adapted, metric (to f), with respect to which f satisfies the conditions of hyperbolicity with constants $C' = 1$ and $\lambda' = \lambda + \epsilon$, and the subspaces $E^s(x)$, $E^u(x)$ are ϵ -orthogonal, i.e., $\langle v^s, v^u \rangle' < \epsilon$ for all unit vectors $v^s \in E^s(x)$, $v^u \in E^u(x)$, $x \in \Lambda$.*

Proof. For $x \in \Lambda$, $v^s \in E^s(x)$, and $v^u \in E^u(x)$, set

$$\|v^s\|' = \sum_{n=0}^{\infty} (\lambda + \epsilon)^{-n} \|df_x^n v^s\|, \quad \|v^u\|' = \sum_{n=0}^{\infty} (\lambda + \epsilon)^{-n} \|df_x^{-n} v^u\|. \quad (5.1)$$

Both series converge uniformly for $\|v^s\|, \|v^u\| \leq 1$ and $x \in \Lambda$. We have

$$\|df_x v^s\|' = \sum_{n=0}^{\infty} (\lambda + \epsilon)^{-n} \|df_x^{n+1} v^s\| = (\lambda + \epsilon)(\|v^s\|' - \|v^s\|) < (\lambda + \epsilon)\|v^s\|',$$

and similarly for $\|df_x^{-1} v^u\|'$. For $v = v^s + v^u \in T_x M$, $x \in \Lambda$, define $\|v\|' = \sqrt{(\|v^s\|')^2 + (\|v^u\|')^2}$. The metric is recovered from the norm:

$$\langle v, w \rangle' = \frac{1}{2}(\|v + w\|'^2 - \|v\|'^2 - \|w\|'^2).$$

With respect to this continuous metric, E^s and E^u are orthogonal and f satisfies the conditions of hyperbolicity with constant 1 and $\lambda + \epsilon$. Now, by

standard methods of differential topology [Hir94], $\langle \cdot, \cdot \rangle'$ can be uniformly approximated on Λ by a smooth metric defined in a neighborhood of Λ . \square

Observe that to construct an adapted metric it is enough to consider sufficiently long finite sums instead of infinite sums in (5.1).

A fixed point x of a differentiable map f is called *hyperbolic* if no eigenvalue of df_x lies on the unit circle. A periodic point x of f of period k is called *hyperbolic* if no eigenvalue of df_x^k lies on the unit circle.

Exercise 5.2.1. Construct a diffeomorphism of the circle that satisfies the first three conditions of hyperbolicity (with Λ being the whole circle) but not the fourth condition.

Exercise 5.2.2. Prove that if Λ is a hyperbolic set of $f: U \rightarrow M$ for some Riemannian metric on M , then Λ is a hyperbolic set of f for any other Riemannian metric on M with the same constant λ .

Exercise 5.2.3. Let x be a fixed point of a diffeomorphism f . Prove that $\{x\}$ is a hyperbolic set if and only if x is a hyperbolic fixed point. Identify the constants C and λ . Give an example when df_x has exactly two eigenvalues $\mu \in (0, 1)$ and μ^{-1} , but $\lambda \neq \mu$.

Exercise 5.2.4. Prove that the horseshoe (§1.8) is a hyperbolic set.

Exercise 5.2.5. Let Λ_i be a hyperbolic set of $f_i: U_i \rightarrow M_i, i = 1, 2$. Prove that $\Lambda_1 \times \Lambda_2$ is a hyperbolic set of $f_1 \times f_2: U_1 \times U_2 \rightarrow M_1 \times M_2$.

Exercise 5.2.6. Let M be a fiber bundle over N with projection π . Let U be an open set in M , and suppose that $\Lambda \subset U$ is a hyperbolic set of $f: U \rightarrow M$ and that $g: N \rightarrow N$ is a factor of f . Prove that $\pi(\Lambda)$ is a hyperbolic set of g .

Exercise 5.2.7. What are necessary and sufficient conditions for a periodic orbit to be a hyperbolic set?

5.3 ϵ -Orbits

An ϵ -orbit of $f: U \rightarrow M$ is a finite or infinite sequence $(x_n) \subset U$ such that $d(f(x_n), x_{n+1}) \leq \epsilon$ for all n . Sometimes ϵ -orbits are referred to as *pseudo-orbits*. For $r \in \{0, 1\}$, denote by dist_r the distance in the space of C^r -functions (see §5.13).

THEOREM 5.3.1. *Let Λ be a hyperbolic set of $f: U \rightarrow M$. Then there is an open set $O \subset U$ containing Λ and positive ϵ_0, δ_0 with the following property:*

for every $\epsilon > 0$ there is $\delta > 0$ such that for any $g: O \rightarrow M$ with $\text{dist}_1(g, f) < \epsilon_0$, any homeomorphism $h: X \rightarrow X$ of a topological space X , and any continuous map $\phi: X \rightarrow O$ satisfying $\text{dist}_0(\phi \circ h, g \circ \phi) < \delta$ there is a continuous map $\psi: X \rightarrow O$ with $\psi \circ h = g \circ \psi$ and $\text{dist}_0(\phi, \psi) < \epsilon$. Moreover, ψ is unique in the sense that if $\psi' \circ h = g \circ \psi'$ for some $\psi': X \rightarrow O$ with $\text{dist}_0(\phi, \psi') < \delta_0$, then $\psi' = \psi$.

Theorem 5.3.1 implies, in particular, that any collection of bi-infinite pseudo-orbits near a hyperbolic set is close to a unique collection of genuine orbits that *shadow* it (Corollary 5.3.2). Moreover, this property holds not only for f itself but for any diffeomorphism C^1 -close to f . In the simplest example, if X is a single point x (and h is the identity), Theorem 5.3.1 implies the existence of a fixed point near $h(x)$ for any diffeomorphism C^1 -close to f .

Proof.¹ By the Whitney embedding theorem [Hir94], we may assume that the manifold M is an m -dimensional submanifold in \mathbb{R}^N for some large N . For $y \in M$, let $D_\alpha(y)$ be the disk of radius α centered at y in the $(N - m)$ -plane $E^\perp(y) \subset \mathbb{R}^N$ that passes through y and is perpendicular to $T_y M$. Since Λ is compact, by the tubular neighborhood theorem [Hir94], for any relatively compact open neighborhood O of Λ in M there is $\alpha \in (0, 1)$ such that the α -neighborhood O_α of O in \mathbb{R}^N is foliated by the disks $D_\alpha(y)$. For each $z \in O_\alpha$ there is a unique point $\pi(z) \in M$ closest to y , and the map π is the projection to M along the disks $D_\alpha(y)$. Each map $g: O \rightarrow M$ can be extended to a map $\tilde{g}: O_\alpha \rightarrow M$ by

$$\tilde{g}(z) = g(\pi(z)).$$

Let $C(X, O_\alpha)$ be the set of continuous maps from X to O_α with distance dist_0 . Note that O_α is bounded and $\phi \in C(X, O_\alpha)$. Let Γ be the Banach space of bounded continuous vector fields $v: X \rightarrow \mathbb{R}^N$ with the norm $\|v\| = \sup_{x \in X} \|v(x)\|$. The map $\phi' \mapsto \phi' - \phi$ is an isometry from the ball of radius α centered at ϕ in $C(X, O_\alpha)$ onto the ball B_α of radius α centered at 0 in Γ . Define $\Phi: B_\alpha \rightarrow \Gamma$ by

$$(\Phi(v))(x) = \tilde{g}(\phi(h^{-1}(x)) + v(h^{-1}(x))) - \phi(x), \quad v \in B_\alpha, \quad x \in X.$$

If v is a fixed point of Φ and $\psi(x) = \phi(x) + v(x)$, then $\tilde{g}(\psi(h^{-1}(x))) = \psi(x)$. Observe that $\tilde{g}(y) \in M$ and hence $\psi(x) \in M$ for $x \in X$ and $g(\psi(h^{-1}(x))) = \psi(x)$. Thus to prove the theorem it suffices to show that Φ has a unique fixed point near ϕ , which depends continuously on g .

¹ The main idea of this proof was communicated to us by A. Katok.

The map Φ is differentiable as a map of Banach spaces, and the derivative

$$(d\Phi_v w)(x) = d\tilde{g}_{\phi(h^{-1}(x)+v(h^{-1}(x)))} w(h^{-1}(x))$$

is continuous in v . To establish the existence and uniqueness of a fixed point v and its continuous dependence on g , we study the derivative of Φ . By taking the maximum of appropriate derivatives we obtain that $\|(d\Phi_v w)(x)\| \leq L\|w\|$, where L depends on the first derivatives of g and on the embedding but does not depend on X, h , and ϕ . For $v = 0$,

$$(d\Phi_0 w)(x) = d\tilde{g}_{\phi(h^{-1}(x))} w(h^{-1}(x)).$$

Since Λ is a hyperbolic set, for some constants $\lambda \in (0, 1)$ and $C > 1$, we have for every $y \in \Lambda$ and $n \in \mathbb{N}$

$$\|df_y^n v\| \leq C\lambda^n \|v\| \quad \text{if } v \in E^s(y), \tag{5.2}$$

$$\|df_y^{-n} v\| \leq C\lambda^n \|v\| \quad \text{if } v \in E^u(y). \tag{5.3}$$

For $z \in O_\alpha$, let \tilde{T}_z denote the m -dimensional plane through z that is orthogonal to the disk $D_\alpha(\pi(z))$. The planes \tilde{T}_z form a differentiable distribution on O_α . Note that $\tilde{T}_z = T_z M$ for $z \in O$. Extend the splitting $T_y M = E^s(y) \oplus E^u(y)$ continuously from Λ to O_α (decreasing the neighborhood O and α if necessary) so that $E^s(z) \oplus E^u(z) = \tilde{T}_z$ and $T_z \mathbb{R}^N = E^s(z) \oplus E^u(z) \oplus E^\perp(\pi(z))$. Denote by P^s, P^u , and P^\perp the projections in each tangent space $T_z \mathbb{R}^N$ onto $E^s(z), E^u(z)$, and $E^\perp(\pi(z))$, respectively.

Fix $n \in \mathbb{N}$ so that $C\lambda^n < 1/2$. By (5.2)–(5.3) and continuity, for a small enough $\alpha > 0$ and small enough neighborhood $O \supset \Lambda$, there is $\epsilon_0 > 0$ such that for every g with $\text{dist}_1(f, g) < \epsilon_0$, every $z \in O_\alpha$, and every $v^s \in E^s(z), v^u \in E^u(z), v^\perp \in E^\perp(\pi(z))$ we have

$$\|P^s d\tilde{g}_z^n v^s\| \leq \frac{1}{2} \|v^s\|, \quad \|P^u d\tilde{g}_z^n v^s\| \leq \frac{1}{100} \|v^s\|, \tag{5.4}$$

$$\|P^u d\tilde{g}_z^n v^u\| \geq 2 \|v^u\|, \quad \|P^s d\tilde{g}_z^n v^u\| \leq \frac{1}{100} \|v^u\|, \tag{5.5}$$

$$d\tilde{g}_z^n v^\perp = 0. \tag{5.6}$$

Denote

$$\Gamma^v = \{v \in \Gamma : v(x) \in E^v(\phi(x)) \text{ for all } x \in X\}, \quad v = s, u, \perp.$$

The subspaces $\Gamma^s, \Gamma^u, \Gamma^\perp$ are closed and $\Gamma = \Gamma^s \oplus \Gamma^u \oplus \Gamma^\perp$. By construction,

$$d\Phi_0 = \begin{pmatrix} A^{ss} & A^{su} & 0 \\ A^{us} & A^{uu} & 0 \\ 0 & 0 & 0 \end{pmatrix},$$

where $A^{ij}: \Gamma^i \rightarrow \Gamma^j$, $i, j = s, u$. By (5.4)–(5.6), there are positive ϵ_0 and δ such that if $\text{dist}_1(f, g) < \epsilon_0$ and $\text{dist}_0(\phi \circ h, g \circ \phi) < \delta$, then the spectrum of $d\Phi_0$ is separated from the unit circle. Therefore the operator $d\Phi_0 - \text{Id}$ is invertible and

$$\|(d\Phi_0 - \text{Id})^{-1}\| < K,$$

where K depends only on f and ϕ .

As for maps of finite-dimensional linear spaces, $\Phi(v) = \Phi(0) + d\Phi_0 v + H(v)$, where $\|H(v_1) - H(v_2)\| \leq C_1 \max\{\|v_1\|, \|v_2\|\} \cdot \|v_1 - v_2\|$ for some $C_1 > 0$ and small $\|v_1\|, \|v_2\|$. A fixed point v of Φ satisfies

$$F(v) = -(d\Phi_0 - \text{Id})^{-1}(\Phi(0) + H(v)) = v.$$

If $\zeta > 0$ is small enough, then for any $v_1, v_2 \in \Gamma$ with $\|v_1\|, \|v_2\| < \zeta$,

$$\|F(v_1) - F(v_2)\| < \frac{1}{2}\|v_1 - v_2\|.$$

Thus for an appropriate choice of constants and neighborhoods in the construction, $F: \Gamma \rightarrow \Gamma$ is a contraction, and therefore has a unique fixed point, which depends continuously on g . \square

Theorem 5.3.1 implies that an ϵ -orbit lying in a small enough neighborhood of a hyperbolic set can be globally (i.e., for all times) approximated by a real f -orbit in the hyperbolic set. This property is called *shadowing* (the real orbit shadows the pseudo-orbit). A continuous map f of a topological space X has the *shadowing property* if for every $\epsilon > 0$ there is $\delta > 0$ such that every δ -orbit is ϵ -approximated by a real orbit.

For $\epsilon > 0$, denote by Λ_ϵ the open ϵ -neighborhood of Λ .

COROLLARY 5.3.2 (Anosov's Shadowing Theorem). *Let Λ be a hyperbolic set of $f: U \rightarrow M$. Then for every $\epsilon > 0$ there is $\delta > 0$ such that if (x_k) is a finite or infinite δ -orbit of f and $\text{dist}(x_k, \Lambda) < \delta$ for all k , then there is $x \in \Lambda_\epsilon$ with $\text{dist}(f^k(x), x_k) < \epsilon$.*

Proof. Choose a neighborhood O satisfying the conclusion of Theorem 5.3.1, and choose $\delta > 0$ such that $\Lambda_\delta \subset O$. If (x_k) is finite or semi-infinite, add to (x_k) the preimages of some $y_0 \in \Lambda$ whose distance to the first point of (x_k) is $< \delta$, and/or the images of some $y_m \in \Lambda$ whose distance to the

last point of (x_k) is $< \delta$, to obtain a doubly infinite δ -orbit lying in the δ -neighborhood of Λ . Let $X = (x_k)$ with discrete topology, $g = f, h$ be the shift $x_k \mapsto x_{k+1}$, and $\phi: X \rightarrow U$ be the inclusion, $\phi(x_k) = x_k$. Since (x_k) is a δ -orbit, $\text{dist}(\phi(h(x_k)), f(\phi(x_k))) < \delta$. Theorem 5.3.1 applies, and the corollary follows. \square

As in Chapter 2, denote by $\text{NW}(f)$ the set of non-wandering points, and by $\text{Per}(f)$ the set of periodic points of f . If Λ is f -invariant, denote by $\text{NW}(f|_\Lambda)$ the set of non-wandering points of f restricted to Λ . In general, $\text{NW}(f|_\Lambda) \neq \text{NW}(f) \cap \Lambda$.

PROPOSITION 5.3.3. *Let Λ be a hyperbolic set of $f: U \rightarrow M$. Then $\overline{\text{Per}(f|_\Lambda)} = \text{NW}(f|_\Lambda)$.*

Proof. Fix $\epsilon > 0$ and let $x \in \text{NW}(f|_\Lambda)$. Choose δ from Theorem 5.3.1, and let $V = B(x, \delta/2) \cap \Lambda$. Since $x \in \text{NW}(f|_\Lambda)$, there is $n \in \mathbb{N}$ such that $f^n(V) \cap V \neq \emptyset$. Let $z \in f^{-n}(f^n(V) \cap V) = V \cap f^{-n}(V)$. Then $\{z, f(z), \dots, f^{n-1}(z)\}$ is a δ -orbit, so by Theorem 5.3.1 there is a periodic point of period n within 2ϵ of z . \square

COROLLARY 5.3.4. *If $f: M \rightarrow M$ is Anosov, then $\overline{\text{Per}(f)} = \text{NW}(f)$.*

Exercise 5.3.1. Interpret Theorem 5.3.1 for $X = \mathbb{Z}_m$ and $h(z) = z + 1 \pmod{m}$.

Exercise 5.3.2. Let Λ be a hyperbolic set of $f: U \rightarrow M$. Prove that the restriction $f|_\Lambda$ is expansive.

Exercise 5.3.3. Let $T: [0, 1] \rightarrow [0, 1]$ be the *tent map*: $T(x) = 2x$ for $0 \leq x \leq 1/2$ and $T(x) = 2(1-x)$ for $1/2 \leq x \leq 1$. Does T have the shadowing property?

Exercise 5.3.4. Prove that a circle rotation does not have the shadowing property. Prove that no isometry of a manifold has the shadowing property.

Exercise 5.3.5. Show that every minimal hyperbolic set consists of exactly one periodic orbit.

5.4 Invariant Cones

Although hyperbolic sets are defined in terms of invariant families of linear spaces, it is often convenient, and in more general settings even necessary, to work with invariant families of linear cones instead of subspaces. In this section, we characterize hyperbolicity in terms of families of invariant cones.

Let Λ be a hyperbolic set of $f: U \rightarrow M$. Since the distributions E^s and E^u are continuous (Proposition 5.2.1), we extend them to continuous distributions \tilde{E}^s and \tilde{E}^u defined in a neighborhood $U(\Lambda) \supset \Lambda$. If $x \in U(\Lambda)$ and $v \in T_x M$, let $v = v^s + v^u$ with $v^s \in \tilde{E}^s(x)$ and $v^u \in \tilde{E}^u(x)$. Assume that the metric is adapted with constant λ . For $\alpha > 0$, define the *stable* and *unstable cones* of size α by

$$K_\alpha^s(x) = \{v \in T_x M: \|v^u\| \leq \alpha \|v^s\|\},$$

$$K_\alpha^u(x) = \{v \in T_x M: \|v^s\| \leq \alpha \|v^u\|\}.$$

For a cone K , let $\overset{\circ}{K} = \text{int}(K) \cup \{0\}$. Let $\Lambda_\epsilon = \{x \in U: \text{dist}(x, \Lambda) < \epsilon\}$.

PROPOSITION 5.4.1. *For every $\alpha > 0$ there is $\epsilon = \epsilon(\alpha) > 0$ such that $f^i(\Lambda_\epsilon) \subset U(\Lambda)$, $i = -1, 0, 1$, and for every $x \in \Lambda_\epsilon$*

$$df_x K_\alpha^u(x) \subset \overset{\circ}{K}_\alpha^u(f(x)) \quad \text{and} \quad df_{f(x)}^{-1} K_\alpha^s(f(x)) \subset \overset{\circ}{K}_\alpha^s(x).$$

Proof. The inclusions hold for $x \in \Lambda$. The statement follows by continuity. \square

PROPOSITION 5.4.2. *For every $\delta > 0$ there are $\alpha > 0$ and $\epsilon > 0$ such that $f^i(\Lambda_\epsilon) \subset U(\Lambda)$, $i = -1, 0, 1$, and for every $x \in \Lambda_\epsilon$*

$$\|df_x^{-1}v\| \leq (\lambda + \delta)\|v\| \quad \text{if} \quad v \in K_\alpha^u(x),$$

and

$$\|df_x v\| \leq (\lambda + \delta)\|v\| \quad \text{if} \quad v \in K_\alpha^s(x).$$

Proof. The statement follows by continuity for a small enough α and $\epsilon = \epsilon(\alpha)$ from Proposition 5.4.1. \square

The following proposition is the converse of Propositions 5.4.1 and 5.4.2.

PROPOSITION 5.4.3. *Let Λ be a compact invariant set of $f: U \rightarrow M$. Suppose that there is $\alpha > 0$ and for every $x \in \Lambda$ there are continuous subspaces $\tilde{E}^s(x)$ and $\tilde{E}^u(x)$ such that $\tilde{E}^s(x) \oplus \tilde{E}^u(x) = T_x M$, and the α -cones $K_\alpha^s(x)$ and $K_\alpha^u(x)$ determined by the subspaces satisfy*

1. $df_x K_\alpha^u(x) \subset K_\alpha^u(f(x))$ and $df_{f(x)}^{-1} K_\alpha^s(f(x)) \subset K_\alpha^s(x)$, and
2. $\|df_x v\| < \|v\|$ for non-zero $v \in K_\alpha^s(x)$, and $\|df_x^{-1}v\| < \|v\|$ for non-zero $v \in K_\alpha^u(x)$.

Then Λ is a hyperbolic set of f .

Proof. By compactness of Λ and of the unit tangent bundle of M , there is a constant $\lambda \in (0, 1)$ such that

$$\|df_x v\| \leq \lambda \|v\| \quad \text{for } v \in K_\alpha^s(x) \quad \text{and} \quad \|df_x^{-1} v\| \leq \lambda \|v\| \quad \text{for } v \in K_\alpha^u(x).$$

For $x \in \Lambda$, the subspaces

$$E^s(x) = \bigcap_{n \geq 0} df_{f^n(x)}^{-n} K^s(f^n(x)) \quad \text{and} \quad E^u(x) = \bigcap_{n \geq 0} df_{f^{-n}(x)}^n K^u(f^{-n}(x))$$

satisfy the definition of hyperbolicity with constants λ and $C = 1$. \square

Let

$$\begin{aligned} \Lambda_\epsilon^s &= \{x \in U: \text{dist}(f^n(x), \Lambda) < \epsilon \text{ for all } n \in \mathbb{N}_0\}, \\ \Lambda_\epsilon^u &= \{x \in U: \text{dist}(f^{-n}(x), \Lambda) < \epsilon \text{ for all } n \in \mathbb{N}_0\}. \end{aligned}$$

Note that both sets are contained in Λ_ϵ and that $f(\Lambda_\epsilon^s) \subset \Lambda_\epsilon^s$, $f^{-1}(\Lambda_\epsilon^u) \subset \Lambda_\epsilon^u$.

PROPOSITION 5.4.4. *Let Λ be a hyperbolic set of f with adapted metric. Then for every $\delta > 0$ there is $\epsilon > 0$ such that the distributions E^s and E^u can be extended to Λ_ϵ so that*

1. E^s is continuous on Λ_ϵ^s , and E^u is continuous on Λ_ϵ^u ,
2. if $x \in \Lambda_\epsilon \cap f(\Lambda_\epsilon)$ then $df_x E^s(x) = E^s(f(x))$ and $df_x E^u(x) = E^u(f(x))$,
3. $\|df_x v\| < (\lambda + \delta)\|v\|$ for every $x \in \Lambda_\epsilon$ and $v \in E^s(x)$,
4. $\|df_x^{-1} v\| < (\lambda + \delta)\|v\|$ for every $x \in \Lambda_\epsilon$ and $v \in E^u(x)$.

Proof. Choose $\epsilon > 0$ small enough that $\Lambda_\epsilon \subset U(\Lambda)$. For $x \in \Lambda_\epsilon^s$, let $E^s(x) = \lim_{n \rightarrow \infty} df_{f^n(x)}^{-n} (\tilde{E}^s(f^n(x)))$. By Proposition 5.4.2, the limit exists if δ , α , and ϵ are small enough. If $x \in \Lambda_\epsilon \setminus \Lambda_\epsilon^s$, let $n(x) \in \mathbb{N}$ be such that $f^n(x) \in \Lambda_\epsilon$ for $n = 0, 1, \dots, n(x)$ and $f^{n(x)+1}(x) \notin \Lambda_\epsilon$, and let $E^s(x) = df_{f^n(x)}^{-n(x)} (\tilde{E}^s(f^{n(x)}(x)))$. The continuity of E^s on Λ_ϵ^s and the required properties follow from Proposition 5.4.2. A similar construction with f replaced by f^{-1} gives an extension of E^u . \square

Exercise 5.4.1. Prove that the solenoid (§1.9) is a hyperbolic set.

Exercise 5.4.2. Let Λ be a hyperbolic set of f . Prove that there is an open set $O \supset \Lambda$ and $\epsilon > 0$ such that for every g with $\text{dist}_1(f, g) < \epsilon$, the invariant set $\Lambda_g = \bigcap_{n=-\infty}^{\infty} g^n(O)$ is a hyperbolic set of g .

Exercise 5.4.3. Prove that the topological entropy of an Anosov diffeomorphism is positive.

Exercise 5.4.4. Let Λ be a hyperbolic set of f . Prove that if $\dim E^u(x) > 0$ for each $x \in \Lambda$, then f has sensitive dependence on initial conditions on Λ (see §1.12).

5.5 Stability of Hyperbolic Sets

In this section, we use pseudo-orbits and invariant cones to obtain key properties of hyperbolic sets. The next two propositions imply that hyperbolicity is “persistent.”

PROPOSITION 5.5.1. *Let Λ be a hyperbolic set of $f: U \rightarrow M$. There is an open set $U(\Lambda) \supset \Lambda$ and $\epsilon_0 > 0$ such that if $K \subset U(\Lambda)$ is a compact invariant subset of a diffeomorphism $g: U \rightarrow M$ with $\text{dist}_1(g, f) < \epsilon_0$, then K is a hyperbolic set of g .*

Proof. Assume that the metric is adapted to f , and extend the distributions E_f^s and E_f^u to continuous distributions \tilde{E}_f^s and \tilde{E}_f^u defined in an open neighborhood $U(\Lambda)$ of Λ . For an appropriate choice of $U(\Lambda)$, ϵ_0 , and α , the stable and unstable α -cones determined by \tilde{E}_f^s and \tilde{E}_f^u satisfy the assumptions of Proposition 5.4.3 for the map g . \square

Denote by $\text{Diff}^1(M)$ the space of C^1 diffeomorphisms of M with the C^1 topology.

COROLLARY 5.5.2. *The set of Anosov diffeomorphisms of a given compact manifold is open in $\text{Diff}^1(M)$.*

PROPOSITION 5.5.3. *Let Λ be a hyperbolic set of $f: U \rightarrow M$. For every open set $V \subset U$ containing Λ and every $\epsilon > 0$, there is $\delta > 0$ such that for every $g: V \rightarrow M$ with $\text{dist}_1(g, f) < \delta$, there is a hyperbolic set $K \subset V$ of g and a homeomorphism $\chi: K \rightarrow \Lambda$ such that $\chi \circ g|_K = f|_\Lambda \circ \chi$ and $\text{dist}_0(\chi, \text{Id}) < \epsilon$.*

Proof. Let $X = \Lambda$, $h = f|_\Lambda$, and let $\phi: \Lambda \hookrightarrow U$ be the inclusion. By Theorem 5.3.1, there is a continuous map $\psi: \Lambda \rightarrow U$ such that $\psi \circ f|_\Lambda = g \circ \psi$. Set $K = \psi(\Lambda)$. Now apply Theorem 5.3.1 to $X = K$, $h = g|_K$, and the inclusion $\phi: K \hookrightarrow M$ to get $\psi': K \rightarrow U$ with $\psi' \circ g|_K = f|_\Lambda \circ \psi$. By uniqueness, $\psi^{-1} = \psi'$. For a small enough δ , the map $\chi = \psi'$ is close to the identity, and, by Proposition 5.5.1, K is hyperbolic. \square

A C^1 diffeomorphism f of a C^1 manifold M is called *structurally stable* if for every $\epsilon > 0$ there is $\delta > 0$ such that if $g \in \text{Diff}^1(M)$ and $\text{dist}_1(g, f) < \delta$, then there is a homeomorphism $h: M \rightarrow M$ for which $f \circ h = h \circ g$ and

$\text{dist}_0(h, \text{Id}) < \epsilon$. If one demands that the conjugacy h be C^1 , the definition becomes vacuous. For example, if f has a hyperbolic fixed point x , then any small enough perturbation g has a fixed point y nearby; if the conjugation is differentiable, then the matrices dg_y and df_x are similar. This condition restricts g to lie in a proper submanifold of $\text{Diff}^1(M)$.

COROLLARY 5.5.4. *Anosov diffeomorphisms are structurally stable.*

Exercise 5.5.1. Interpret Proposition 5.5.3 when Λ is a hyperbolic periodic point of f .

5.6 Stable and Unstable Manifolds

Hyperbolicity is defined in terms of infinitesimal objects: a family of linear subspaces invariant by the differential of a map. In this section, we construct the corresponding integral objects, the stable and unstable manifolds.

For $\delta > 0$, let $B_\delta = B(0, \delta) \subset \mathbb{R}^m$ be the ball of radius δ at 0.

PROPOSITION 5.6.1 (Hadamard–Perron). *Let $f = (f_n)_{n \in \mathbb{N}_0}$, $f_n: B_\delta \rightarrow \mathbb{R}^m$, be a sequence of C^1 diffeomorphisms onto their images such that $f_n(0) = 0$. Suppose that for each n there is a splitting $\mathbb{R}^m = E^s(n) \oplus E^u(n)$ and $\lambda \in (0, 1)$ such that*

1. $df_n(0)E^s(n) = E^s(n+1)$ and $df_n(0)E^u(n) = E^u(n+1)$,
2. $\|df_n(0)v^s\| < \lambda\|v^s\|$ for every $v^s \in E^s(n)$,
3. $\|df_n(0)v^u\| > \lambda^{-1}\|v^u\|$ for every $v^u \in E^u(n)$,
4. the angles between $E^s(n)$ and $E^u(n)$ are uniformly bounded away from 0,
5. $\{df_n(\cdot)\}_{n \in \mathbb{N}_0}$ is an equicontinuous family of functions from B_δ to $\text{GL}(m, \mathbb{R})$.

Then there are $\epsilon > 0$ and a sequence $\phi = (\phi_n)_{n \in \mathbb{N}_0}$ of uniformly Lipschitz continuous maps $\phi_n: B_\epsilon^s = \{v \in E^s(n) : \|v\| < \epsilon\} \rightarrow E^u(n)$ such that

1. $\text{graph}(\phi_n) \cap B_\epsilon = W_\epsilon^s(n) := \{x \in B_\epsilon : \|f_{n+k-1} \circ \dots \circ f_{n+1} \circ f_n(x)\| \rightarrow_{k \rightarrow \infty} 0\}$,
2. $f_n(\text{graph}(\phi_n)) \subset \text{graph}(\phi_{n+1})$,
3. if $x \in \text{graph}(\phi_n)$, then $\|f_n(x)\| \leq \lambda\|x\|$, so by (2), $f_n^k(x) \rightarrow 0$ exponentially as $k \rightarrow \infty$,
4. for $x \in B_\epsilon \setminus \text{graph}(\phi_n)$,

$$\|P_{n+1}^u f_n(x) - \phi_{n+1}(P_{n+1}^s f_n(x))\| > \lambda^{-1} \|P_n^u x - \phi_n(P_n^s x)\|,$$

where $P_n^s(P_n^u)$ denotes the projection onto $E^s(n)(E^u(n))$ parallel to $E^u(n)(E^s(n))$,

5. ϕ_n is differentiable at 0 and $d\phi_n(0) = 0$, i.e., the tangent plane to $\text{graph}(\phi_n)$ is $E^s(n)$.
6. ϕ depends continuously on f in the topologies induced by the following distance functions:

$$d_0(\phi, \psi) = \sup_{n \in \mathbb{N}_0, x \in B_\epsilon} 2^{-n} |\phi_n(x) - \psi_n(x)|,$$

$$d(f, g) = \sup_{n \in \mathbb{N}_0} 2^{-n} \text{dist}_1(f_n, g_n),$$

where dist_1 is the C^1 distance.

Proof. For positive constants L and ϵ , let $\Phi(L, \epsilon)$ be the space of sequences $\phi = (\phi_n)_{n \in \mathbb{N}_0}$, where $\phi_n: B_\epsilon^s \rightarrow E^u(n)$ is a Lipschitz-continuous map with Lipschitz constant L and $\phi_n(0) = 0$. Define distance on $\Phi(L, \epsilon)$ by $d(\phi, \psi) = \sup_{n \in \mathbb{N}_0, x \in B_\epsilon} |\phi_n(x) - \psi_n(x)|$. This metric is complete.

We now define an operator $F: \Phi(L, \epsilon) \rightarrow \Phi(L, \epsilon)$ called the *graph transform*. Suppose $\phi = (\phi_n) \in \Phi$. We prove in the next lemma that for a small enough ϵ , the projection of the set $f_n^{-1}(\text{graph}(\phi_{n+1}))$ onto $E^s(n)$ covers $E_\epsilon^s(n)$, and $f_n^{-1}(\text{graph}(\phi_{n+1}))$ contains the graph of a continuous function $\psi_n: B_\epsilon^s \rightarrow E_\epsilon^u(n)$ with Lipschitz constant L . We set $F(\phi)_n = \psi_n$.

Note that a map $h: \mathbb{R}^k \rightarrow \mathbb{R}^l$ is Lipschitz continuous at 0 with Lipschitz constant L if and only if the graph of h lies in the L -cone about \mathbb{R}^k , and is Lipschitz continuous at $x \in \mathbb{R}^k$ if and only if its graph lies in the L -cone about the translate of \mathbb{R}^k by $(x, h(x))$.

LEMMA 5.6.2. *For any $L > 0$, there exists $\epsilon > 0$ such that the graph transform F is a well-defined operator on $\Phi(L, \epsilon)$.*

Proof. For $L > 0$ and $x \in B_\epsilon$, let $K_L^s(n)$ denote the stable cone

$$K_L^s(n) = \{v \in \mathbb{R}^m: v = v^s + v^u, v^s \in E^s(n), v^u \in E^u(n), |v^u| \leq L|v^s|\}.$$

Note that $df_n^{-1}(0)K_L^s(n+1) \subset K_L^s(n)$ for any $L > 0$. Therefore, by the uniform continuity of df_n , for any $L > 0$ there is $\epsilon > 0$ such that $df_n^{-1}(x)K_L^s(n+1) \subset K_L^s(n)$ for any $n \in \mathbb{N}_0$ and $x \in B_\epsilon$. Hence the preimage under f_n of the graph of a Lipschitz-continuous function is the graph of a Lipschitz-continuous function. For $\phi \in \Phi(L, \epsilon)$, consider the following composition $\beta = P^s(n) \circ f_n^{-1} \circ \phi_n$, where $P^s(n)$ is the projection onto $E^s(n)$ parallel to $E^u(n)$. If ϵ is small enough, then β is an expanding map and its image covers $B_\epsilon^s(n)$ (Exercise 5.6.1). Hence $F(\phi) \in \Phi(L, \epsilon)$. \square

The next lemma shows that F is a contracting operator for an appropriate choice of ϵ and L .




Figure 5.1. Graph transform applied to ϕ and ψ .

LEMMA 5.6.3. *There are $\epsilon > 0$ and $L > 0$ such that F is a contracting operator on $\Phi(L, \epsilon)$.*

Proof. For $L \in (0, 0.1)$, let $K_L^u(n)$ denote the unstable cone

$$K_L^u(n) = \{v \in \mathbb{R}^m : v = v^s + v^u, v^s \in E^s(n), v^u \in E^u(n), |v^u| \geq L^{-1}|v^s|\},$$

and note that $df_n(0)K_L^u(n) \subset K_L^u(n+1)$. As in Lemma 5.6.2, by the uniform continuity of df_n , for any $L > 0$ there is $\epsilon > 0$ such that the inclusion $df_n K_L^u(n) \subset K_L^u(n+1)$ holds for every $n \in \mathbb{N}_0$ and $x \in B_\epsilon$.

Let $\phi, \psi \in \Phi(L, \epsilon)$, $\phi' = F(\phi)$, $\psi' = F(\psi)$ (see Figure 5.1). For any $\eta > 0$ there are $n \in \mathbb{N}_0$ and $y \in B_\epsilon^s$ such that $|\phi'_n(y) - \psi'_n(y)| > d(\phi', \psi') - \eta$. Let c^u be the straight line segment from $(y, \phi'_n(y))$ to $(y, \psi'_n(y))$. Since c^u is parallel to $E^u(n)$, we have that $\text{length}(f_n(c^u)) > \lambda^{-1} \text{length}(c^u)$. Let $f_n(y, \psi'_n(y)) = (z, \psi_{n+1}(z))$, and consider the curvilinear triangle formed by the straight line segment from $(z, \phi_{n+1}(z))$ to $(z, \psi_{n+1}(z))$, $f_n(c^u)$, and the shortest curve on the graph of ψ_{n+1} connecting the ends of these curves. For a small enough $\epsilon > 0$ the tangent vectors to the image $f_n(c^u)$ lie in $K_L^u(n+1)$ and the tangent vectors to the graph of ϕ_{n+1} lie in $K_L^s(n+1)$. Therefore,

$$\begin{aligned} |\phi_{n+1}(z) - \psi_{n+1}(z)| &\geq \frac{\text{length}(f_n(c^u))}{1 + 2L} - L(1 + L) \cdot \text{length}(f_n(c^u)) \\ &\geq (1 - 4L)\text{length}(f_n(c^u)), \end{aligned}$$

and

$$\begin{aligned} d(\phi, \psi) &\geq |\phi_{n+1}(z) - \psi_{n+1}(z)| \geq (1 - 4L)\text{length}(f_n(c^u)) \\ &> (1 - 4L)\lambda^{-1}\text{length}(c^u) = (1 - 4L)\lambda^{-1}(d(\phi', \psi') - \eta). \end{aligned}$$

Since η is arbitrary, F is contracting for small enough L and ϵ . □

Since F is contracting (Lemma 5.6.3) and depends continuously on f , it has a unique fixed point $\phi \in \Phi(L, \epsilon)$, which depends continuously on f (property 6) and automatically satisfies property 2. The invariance of the

stable and unstable cones (with a small enough ϵ) implies that ϕ satisfies properties 3 and 4. Property 1 follows immediately from 3 and 4. Since property 1 gives a geometric characterization of $\text{graph}(\phi_n)$, the fixed point of F for a smaller ϵ is a restriction of the fixed point of F for a larger ϵ to a smaller domain. As $\epsilon \rightarrow 0$ and $L \rightarrow \infty$, the stable cone $K_L^s(0, n)$ (which contains $\text{graph}(\phi_n)$) tends to $E^s(n)$. Therefore $E^s(n)$ is the tangent plane to $\text{graph}(\phi_n)$ at 0 (property 5). \square

The following theorem establishes the existence of *local stable manifolds* for points in a hyperbolic set Λ and in Λ_δ^s , and of *local unstable manifolds* for points in Λ and in Λ_δ^u (see §5.4); recall that $\Lambda_\delta^s \supset \Lambda$ and $\Lambda_\delta^u \supset \Lambda$.

THEOREM 5.6.4. *Let $f: M \rightarrow M$ be a C^1 diffeomorphism of a differentiable manifold, and let $\Lambda \subset M$ be a hyperbolic set of f with constant λ (the metric is adapted).*

Then there are $\epsilon, \delta > 0$ such that for every $x^s \in \Lambda_\delta^s$ and every $x^u \in \Lambda_\delta^u$ (see §5.4)

1. *the sets*

$$W_\epsilon^s(x^s) = \{y \in M: \text{dist}(f^n(x^s), f^n(y)) < \epsilon \text{ for all } n \in \mathbb{N}_0\},$$

$$W_\epsilon^u(x^u) = \{y \in M: \text{dist}(f^{-n}(x^u), f^{-n}(y)) < \epsilon \text{ for all } n \in \mathbb{N}_0\},$$

called the local stable manifold of x^s and the local unstable manifold of x^u , are C^1 embedded disks,

2. $T_{y^s} W_\epsilon^s(x^s) = E^s(y^s)$ for every $y^s \in W_\epsilon^s(x^s)$, and $T_{y^u} W_\epsilon^u(x^u) = E^u(y^u)$ for every $y^u \in W_\epsilon^u(x^u)$ (see Proposition 5.4.4),
3. $f(W_\epsilon^s(x^s)) \subset W_\epsilon^s(f(x^s))$ and $f^{-1}(W_\epsilon^u(f(x^u))) \subset W_\epsilon^u(x^u)$,
4. if $y^s, z^s \in W_\epsilon^s(x^s)$, then $d^s(f(y^s), f(z^s)) < \lambda d^s(y^s, z^s)$, where d^s is the distance along $W_\epsilon^s(x^s)$,
if $y^u, z^u \in W_\epsilon^u(x^u)$, then $d^u(f^{-1}(y^u), f^{-1}(z^u)) < \lambda d^u(y^u, z^u)$, where d^u is the distance along $W_\epsilon^u(x^u)$,
5. if $0 < \text{dist}(x^s, y) < \epsilon$ and $\exp_{x^s}^{-1}(y)$ lies in the δ -cone $K_\delta^u(x^s)$, then $\text{dist}(f(x^s), f(y)) > \lambda^{-1} \text{dist}(x^s, y)$,
if $0 < \text{dist}(x^u, y) < \epsilon$ and $\exp_{x^u}^{-1}(y)$ lies in the δ -cone $K_\delta^s(x^u)$, then $\text{dist}(f(x^u), f(y)) < \lambda \text{dist}(x^s, y)$,
6. if $y^s \in W_\epsilon^s(x^s)$, then $W_\alpha^s(y^s) \subset W_\epsilon^s(x^s)$ for some $\alpha > 0$,
if $y^u \in W_\epsilon^u(x^u)$, then $W_\beta^u(y^u) \subset W_\epsilon^u(x^u)$ for some $\beta > 0$,

Proof. Since $\Lambda_\delta^s \supset \Lambda$ is compact, for a small enough δ there is a collection \mathcal{U} of coordinate charts (U_x, ψ_x) , $x \in \Lambda_\delta^s$, such that U_x covers the δ -neighborhood of x and the changes of coordinates $\psi_x \circ \psi_y^{-1}$ between the charts have equicontinuous first derivatives. For any point $x^s \in \Lambda_\delta^s$, let

$f_n = \psi_{f^n(x^s)} \circ f \circ \psi_{f^{n-1}(x^s)}^{-1}$, $E^s(n) = d\psi_{f^n(x^s)}(x^s)E^s(f^n(x^s))$, and $E^u(n) = d\psi_{f^n(x)}(x)E^u(f^n(x))$, apply Proposition 5.6.1, and set $W_\epsilon^s(x) = W_0^s(\epsilon)$. Similarly, apply Proposition 5.6.1 to f^{-1} to construct the local unstable manifolds. Properties 1–6 follow immediately from Proposition 5.6.1. \square

Let Λ be a hyperbolic set of $f: U \rightarrow M$ and $x \in \Lambda$. The (global) stable and unstable manifolds of x are defined by

$$W^s(x) = \{y \in M: d(f^n(x), f^n(y)) \rightarrow 0 \text{ as } n \rightarrow \infty\},$$

$$W^u(x) = \{y \in M: d(f^{-n}(x), f^{-n}(y)) \rightarrow 0 \text{ as } n \rightarrow \infty\}.$$

PROPOSITION 5.6.5. *There is $\epsilon_0 > 0$ such that for every $\epsilon \in (0, \epsilon_0)$, for every $x \in \Lambda$,*

$$W^s(x) = \bigcup_{n=0}^{\infty} f^{-n}(W_\epsilon^s(f^n(x))), \quad W^u(x) = \bigcup_{n=0}^{\infty} f^n(W_\epsilon^u(f^{-n}(x))).$$

Proof. Exercise 5.6.2. \square

COROLLARY 5.6.6. *The global stable and unstable manifolds are embedded C^1 submanifolds of M homeomorphic to the unit balls in corresponding dimensions.*

Proof. Exercise 5.6.3. \square

Exercise 5.6.1. Suppose $f: \mathbb{R}^m \rightarrow \mathbb{R}^m$ is a continuous map such that $|f(x) - f(y)| \geq a|x - y|$ for some $a > 1$, for all $x, y \in \mathbb{R}^m$. If $f(0) = 0$, show that the image of a ball of radius $r > 0$ centered at 0 contains the ball of radius ar centered at 0.

Exercise 5.6.2. Prove Proposition 5.6.5.

Exercise 5.6.3. Prove Corollary 5.6.6.

5.7 Inclination Lemma

Let M be a differentiable manifold. Recall that two submanifolds $N_1, N_2 \subset M$ of complementary dimensions intersect transversely (or are transverse) at a point $p \in N_1 \cap N_2$ if $T_p M = T_p N_1 \oplus T_p N_2$. We write $N_1 \pitchfork N_2$ if every point of intersection of N_1 and N_2 is a point of transverse intersection.

Denote by B_ϵ^i the open ball of radius ϵ centered at 0 in \mathbb{R}^i . For $v \in \mathbb{R}^m = \mathbb{R}^k \times \mathbb{R}^l$, denote by $v^u \in \mathbb{R}^k$ and $v^s \in \mathbb{R}^l$ the components of $v = v^u + v^s$, and by $\pi^u: \mathbb{R}^m \rightarrow \mathbb{R}^k$ the projection to \mathbb{R}^k . For $\delta > 0$, let $K_\delta^u = \{v \in \mathbb{R}^m: \|v^s\| \leq \delta \|v^u\|\}$ and $K_\delta^s = \{v \in \mathbb{R}^m: \|v^u\| \leq \delta \|v^s\|\}$.




Figure 5.2. The image of the graph of ϕ under f^n .

LEMMA 5.7.1. *Let $\lambda \in (0, 1)$, $\epsilon, \delta \in (0, 0.1)$. Suppose $f: B_\epsilon^k \times B_\epsilon^l \rightarrow \mathbb{R}^m$ and $\phi: B_\epsilon^k \rightarrow B_\epsilon^l$ are C^1 maps such that*

1. 0 is a hyperbolic fixed point of f ,
2. $W_\epsilon^u(0) = B_\epsilon^k \times \{0\}$ and $W_\epsilon^s(0) = \{0\} \times B_\epsilon^l$,
3. $\|df_x(v)\| \geq \lambda^{-1}\|v\|$ for every $v \in K_\delta^u$ whenever both $x, f(x) \in B_\epsilon^k \times B_\epsilon^l$,
4. $\|df_x(v)\| \leq \lambda\|v\|$ for every $v \in K_\delta^s$ whenever both $x, f(x) \in B_\epsilon^k \times B_\epsilon^l$,
5. $df_x(K_\delta^u) \subset K_\delta^u$ whenever both $x, f(x) \in B_\epsilon^k \times B_\epsilon^l$,
6. $d(f^{-1})_x(K_\delta^s) \subset K_\delta^s$ whenever both $x, f^{-1}(x) \in B_\epsilon^k \times B_\epsilon^l$,
7. $T_{(y, \phi(y))}\text{graph}(\phi) \subset K_\delta^u$ for every $y \in B_\epsilon^k$,

Then for every $n \in \mathbb{N}$ there is a subset $D_n \subset B_\epsilon^k$ diffeomorphic to B^k and such that the image I_n under f^n of the graph of the restriction $\phi|_{D_n}$ has the following properties: $\pi^u(I_n) \supset B_{\epsilon/2}^k$ and $T_x I_n \subset K_{\delta\lambda^{2n}}^u$ for each $x \in I_n$.

Proof. The lemma follows from the invariance of the cones (Exercise 5.7.2). □

The meaning of the lemma is that the tangent planes to the image of the graph of ϕ under f^n are exponentially (in n) close to the “horizontal” space \mathbb{R}^k , and the image spreads over B_ϵ^k in the horizontal direction (see Figure 5.2).

The following theorem, which is also sometimes called the Lambda Lemma, implies that if f is C^r with $r \geq 1$, and D is any C^1 -disk that intersects transversely the stable manifold $W^s(x)$ of a hyperbolic fixed point x , then the forward images of D converge in the C^r topology to the unstable manifold $W^u(x)$ [PdM82]. We prove only C^1 convergence. Let B_R^u be the ball of radius R centered at x in $W^u(x)$ in the induced metric.

THEOREM 5.7.2 (Inclination Lemma). *Let x be a hyperbolic fixed point of a diffeomorphism $f: U \rightarrow M$, $\dim W^u(x) = k$, and $\dim W^s(x) = l$. Let $y \in W^s(x)$, and suppose that $D \ni y$ is a C^1 submanifold of dimension k intersecting $W^s(x)$ transversely at y .*

Then for every $R > 0$ and $\beta > 0$ there are $n_0 \in \mathbb{N}$ and, for each $n \geq n_0$, a subset $\tilde{D} = \tilde{D}(R, \beta, n)$, $y \in \tilde{D} \subset D$, diffeomorphic to an open k -disk and such that the C^1 distance between $f^n(\tilde{D})$ and B_R^u is less than β .

Proof. We will show that for some $n_1 \in \mathbb{N}$, an appropriate subset $D_1 \subset f^{n_1}(D)$ satisfies the assumptions of Lemma 5.7.1. Since $\{x\}$ is a hyperbolic set of f , for any $\delta > 0$ there is $\epsilon > 0$ such that $E^s(x)$ and $E^u(x)$ can be extended to invariant distributions \tilde{E}^s and \tilde{E}^u in the ϵ -neighborhood B_ϵ of x and the hyperbolicity constant is at most $\lambda + \delta$ (Proposition 5.4.4). Since $f^n(y) \rightarrow x$, there is $n_2 \in \mathbb{N}$ such that $z = f^{n_2}(y) \in B_\epsilon$. Since D intersects $W^s(x)$ transversely, so does $f^{n_2}(D)$. Therefore there is $\eta > 0$ such that if $v \in T_z f^{n_2}(D)$, $\|v\| = 1$, $v = v^s + v^u$, $v^s \in \tilde{E}^s(z)$, $v^u \in \tilde{E}^u(z)$, and $v^u \neq 0$, then $\|v^u\| \geq \eta \|v^s\|$. By Proposition 5.4.4, for a small enough $\delta > 0$, the norm $\|df^n v^s\|$ decays exponentially and $\|df^n v^u\|$ grows exponentially. Therefore, for an arbitrarily small cone size, there exists $n_2 \in \mathbb{N}$ such that $T_{f^{n_2}(y)} f^{n_2}(D)$ lies inside the unstable cone at $f^{n_2}(y)$. \square

Exercise 5.7.1. Prove that if x is a homoclinic point, then x is non-wandering but not recurrent.

Exercise 5.7.2. Prove Lemma 5.7.1.

Exercise 5.7.3. Let p be a hyperbolic fixed point of f . Suppose $W^s(p)$ and $W^u(p)$ intersect transversely at q . Prove that the union of p with the orbit of q is a hyperbolic set of Λ .

5.8 Horseshoes and Transverse Homoclinic Points

Let $\mathbb{R}^m = \mathbb{R}^k \times \mathbb{R}^l$. We will refer to \mathbb{R}^k and \mathbb{R}^l as the unstable and stable subspaces, respectively, and denote by π^u and π^s the projections to those subspaces. For $v \in \mathbb{R}^m$, denote $v^u = \pi^u(v) \in \mathbb{R}^k$ and $v^s = \pi^s(v) \in \mathbb{R}^l$. For $\alpha \in (0, 1)$, call the sets $K_\alpha^u = \{v \in \mathbb{R}^m: |v^s| \leq \alpha |v^u|\}$ and $K_\alpha^s = \{v \in \mathbb{R}^m: |v^u| \leq \alpha |v^s|\}$ the unstable and stable cones, respectively. Let $R^u = \{x \in \mathbb{R}^k: |x| \leq 1\}$, $R^s = \{x \in \mathbb{R}^l: |x| \leq 1\}$, and $R = R^u \times R^s$. For $z = (x, y)$, $x \in \mathbb{R}^k$, $y \in \mathbb{R}^l$, the sets $F^s(z) = \{x\} \times R^s$ and $F^u(z) = R^u \times \{y\}$ will be referred to as unstable and stable fibers, respectively. We say that a C^1 map $f: R \rightarrow \mathbb{R}^m$ has a *horseshoe* if there are $\lambda, \alpha \in (0, 1)$ such that

1. f is one-to-one on R ;
2. $f(R) \cap R$ has at least two components $\Delta_0, \dots, \Delta_{q-1}$;
3. if $z \in R$ and $f(z) \in \Delta_i$, $0 \leq i < q$, then the sets $G_i^u(z) = f(F^u(z)) \cap$




Figure 5.3. A non-linear horseshoe.

- Δ_i and $G_i^s(z) = f^{-1}(F^s(f(z)) \cap \Delta_i)$ are connected, and the restrictions of π^u to $G_i^u(z)$ and of π^s to $G_i^s(z)$ are onto and one-to-one;
4. if $z, f(z) \in R$, then the derivative df_z preserves the unstable cone K_α^u and $\lambda|df_z v| \geq |v|$ for every $v \in K_\alpha^u$, and the inverse $df_{f(z)}^{-1}$ preserves the stable cone K_α^s and $\lambda|df_{f(z)}^{-1} v| \geq |v|$ for every $v \in K_\alpha^s$.

The intersection $\Lambda = \bigcap_{n \in \mathbb{Z}} f^n(R)$ is called a *horseshoe*.

THEOREM 5.8.1. *The horseshoe $\Lambda = \bigcap_{n \in \mathbb{Z}} f^n(R)$ is a hyperbolic set of f . If $f(R) \cap R$ has q components, then the restriction of f to Λ is topologically conjugate to the full two-sided shift σ in the space Σ_q of bi-infinite sequences in the alphabet $\{0, 1, \dots, q - 1\}$.*

Proof. The hyperbolicity of Λ follows from the invariance of the cones and the stretching of vectors inside the cones. The topological conjugacy of $f|_\Lambda$ to the two-sided shift is left as an exercise (Exercise 5.8.2). □

COROLLARY 5.8.2. *If a diffeomorphism f has a horseshoe, then the topological entropy of f is positive.*

Let p be a hyperbolic periodic point of a diffeomorphism $f: U \rightarrow M$. A point $q \in U$ is called *homoclinic* (for p) if $q \neq p$ and $q \in W^s(p) \cap W^u(p)$; it is called *transverse homoclinic* (for p) if in addition $W^s(p)$ and $W^u(p)$ intersect transversely at q .

The next theorem shows that horseshoes, and hence hyperbolic sets in general, are rather common.

THEOREM 5.8.3. *Let p be a hyperbolic periodic point of a diffeomorphism $f: U \rightarrow M$, and let q be a transverse homoclinic point of p . Then for every $\epsilon > 0$ the union of the ϵ -neighborhoods of the orbits of p and q contains a horseshoe of f .*

Proof. We consider only the two-dimensional case; the argument for higher dimensions is a routine generalization of the proof below. We assume without loss of generality that $f(p) = p$ and f is orientation preserving. There is a C^1 coordinate system in a neighborhood $V = V^u \times V^s$ of p such that p is the origin and the stable and unstable manifolds of p coincide locally with the coordinate axes (Figure 5.4). For a point $x \in V$ and a vector $v \in \mathbb{R}^2$, we write $x = (x^u, x^s)$ and $v = (v^u, v^s)$, where s and u indicate the stable (vertical) and unstable (horizontal) components, respectively. We also assume that there is $\lambda \in (0, 1)$ such that $|df_p v^s| < \lambda |v^s|$ and $|df_p^{-1} v^u| < \lambda |v^u|$ for every $v \neq 0$. Fix $\delta > 0$, and let $K_{\delta/2}^s$ and $K_{\delta/2}^u$ be the stable and unstable $\delta/2$ -cones. Choose V small enough so that for every $x \in V$

$$\begin{aligned} df_x(K_{\delta/2}^u) &\subset K_{\delta/2}^u, & |df_x^{-1}v| &< \lambda|v| & \text{if } v \in K_{\delta/2}^u, \\ df_x^{-1}K_{\delta/2}^s &\subset K_{\delta/2}^s, & |df_x v| &< \lambda|v| & \text{if } v \in K_{\delta/2}^s. \end{aligned}$$

Since $q \in W^s(p) \cap W^u(p)$, we have that $f^n(q) \in V$ and $f^{-n}(q) \in V$ for all sufficiently large n . By invariance, $W^s(p)$ and $W^u(p)$ pass through all images $f^n(q)$. Since $W^u(p)$ intersects $W^s(p)$ transversely at q , by Theorem 5.7.2 there is n_u such that $f^n(q) \in V$ for $n \geq n_u$, and an appropriate neighborhood




Figure 5.4. A horseshoe at a homoclinic point.

D^u of $f^n(q)$ in $W^u(p)$ is a C^1 submanifold that “stretches across” V and whose tangent planes lie in $K_{\delta/2}^u$, i.e., D^u is the graph of a C^1 function $\phi^u: V^u \rightarrow V^s$ with $\|d\phi^u\| < \delta/2$. Similarly since $q \in W^u(p)$, there is $n_s \in \mathbb{N}$ such that $f^{-n}(q) \in U$ for $n \geq n_s$ and a small neighborhood D^s of $f^{-n}(q)$ in $W^s(p)$ is the graph of a C^1 function $\phi^s: V^s \rightarrow V^u$ with $\|d\phi^s\| < \delta/2$. Note that since f preserves orientation, the point $f^{n_u+1}(q)$ is not the next intersection of $W^u(p)$ with $W^s(p)$ after $f^{n_u}(q)$; in Figure 5.4 it is shown as the second intersection after $f^{n_u}(q)$ along $W^s(p)$.

Consider a narrow “box” R shown in Figure 5.4, and let $N = k + n_u + n_s + 1$. We will show that for an appropriate choice of the size and position of R and of $k \in \mathbb{N}$, the map $\tilde{f} = f^N$, the box R , and its image $\tilde{f}(R)$ satisfy the definition of a horseshoe. The smaller the width of the box and the closer it lies to $W^s(p)$, the larger is k for which $f^k(R)$ reaches the vicinity of $f^{-n_s}(q)$. The number $\bar{n} = n_u + n_s + 1$ is fixed. If v^u is a horizontal vector at $f^{-n_s}(q)$, its image $w = df_{f^{-n_s}(q)}^{\bar{n}} v^u$ is tangent to $W^u(p)$ at $f^{n_u+1}(q)$ and therefore lies in $K_{\delta/2}^u$. Moreover, $|w| \geq 2\beta|v^u|$ for some $\beta > 0$. For any sufficiently close vector v at a close enough base point, the image will lie in K_δ^u and $|df^{\bar{n}}v| \geq \beta|v|$. The same holds for “almost horizontal” vectors at points close to $f^{-n_s-1}(q)$.

On the other hand, $df_x(K_\alpha^u) \subset K_{\lambda\alpha}^u$ for every small enough $\alpha > 0$ and every $x \in V$. Therefore, if $x \in R$, $f(x), \dots, f^k(x) \in V$ and $v \in K_\delta^u$ is a tangent vector at x , then $df_x^k v \in K_{\delta\lambda^k}^u$ and $|df_x^k v| > \lambda^{-k}|v|$. Suppose now that $x \in R$ is such that $f^k(x)$ is close to either $f^{-n_s}(q)$ or $f^{-n_s-1}(q)$. Let k be large enough so that $\beta/\lambda^k > 10$. There is $\lambda' \in (0, 1)$ such that if $x \in R$ and $f^N(x)$ is close to either $f^{n_u}(q)$ or $f^{n_u+1}(q)$ (i.e., $f^k(x)$ is close to $f^{-n_s}(q)$ or $f^{-n_s-1}(q)$), then K_δ^u is invariant under df_x^N and $\lambda'|df_x^N v| \geq |v|$ for every $v \in K_\delta^u$. Similarly, for an appropriate choice of R and k , the stable δ -cones are invariant under df^{-N} and vectors from K_δ^s are stretched by df^{-N} by a factor at least $(\lambda')^{-1}$.

To guarantee the correct intersection of $f^N(R)$ with R we must choose R carefully. Choose the horizontal boundary segments of R to be straight line segments, and let R stretch vertically so that it crosses $W^u(p)$ near $f^{n_u}(q)$ and $f^{n_u+1}(q)$. By Theorem 5.7.2, the images of these horizontal segments under f^k are almost horizontal line segments. To construct the vertical boundary segments of R , take two vertical segments s_1 and s_2 to the left of $f^{-n_s-1}(q)$ and to the right of $f^{-n_s}(q)$, and truncate their preimages $f^{-k}(s_i)$ by the horizontal boundary segments. By Theorem 5.7.2, the preimages are almost vertical line segments. This choice of R satisfies the definition of a horseshoe. □

Exercise 5.8.1. Let $f: U \rightarrow M$ be a diffeomorphism, p a periodic point of f , and q a (non-transverse) homoclinic point (for p). Prove that every

arbitrarily small C^1 neighborhood of f contains a diffeomorphism g such that p is a periodic point of g and q is a transverse homoclinic point (for p).

Exercise 5.8.2. Prove that if $f(R) \cap R$ in Theorem 5.8.1 has q connected components, then the restriction of f to Λ is topologically conjugate to the full two-sided shift in the space Σ_q of bi-infinite sequences in the alphabet $\{1, \dots, q\}$.

Exercise 5.8.3. Let p_1, \dots, p_k be periodic points (of possibly different periods) of $f: U \rightarrow M$. Suppose $W^u(p_i)$ intersects $W^s(p_{i+1})$ transversely at $q_i, i = 1, \dots, k, p_{k+1} = p_1$ (in particular, $\dim W^s(p_i) = \dim W^s(p_1)$ and $\dim W^u(p_i) = \dim W^u(p_1), i = 2, \dots, k$). The points q_i are called *transverse heteroclinic points*. Prove the following generalization of Theorem 5.8.3: Any neighborhood of the union of the orbits of p_i s and q_i s contains a horseshoe.

5.9 Local Product Structure and Locally Maximal Hyperbolic Sets

A hyperbolic set Λ of $f: U \rightarrow M$ is called *locally maximal* if there is an open set V such that $\Lambda \subset V \subset U$ and $\Lambda = \bigcap_{n=-\infty}^{\infty} f^n(V)$. The horseshoe (§5.8) and the solenoid (§1.9) are examples of locally maximal hyperbolic sets (Exercise 5.9.1).

Since every closed invariant subset of a hyperbolic set is also a hyperbolic set, the geometric structure of a hyperbolic set may be very complicated and difficult to describe. However, due to their special properties, locally maximal hyperbolic sets allow a geometric characterization.

Since $E^s(x) \cap E^u(x) = \{0\}$, the local stable and unstable manifolds of x intersect at x transversely. By continuity, this transversality extends to a neighborhood of the diagonal in $\Lambda \times \Lambda$.

PROPOSITION 5.9.1. *Let Λ be a hyperbolic set of f . For every small enough $\epsilon > 0$ there is $\delta > 0$ such that if $x, y \in \Lambda$ and $d(x, y) < \delta$, then the intersection $W_\epsilon^s(x) \cap W_\epsilon^u(y)$ is transverse and consists of exactly one point $[x, y]$, which depends continuously on x and y . Furthermore, there is $C_p = C_p(\delta) > 0$ such that if $x, y \in \Lambda$ and $d(x, y) < \delta$, then $d^s(x, [x, y]) \leq C_p d(x, y)$ and $d^u(x, [x, y]) \leq C_p d(x, y)$, where d^s and d^u denote distances along the stable and unstable manifolds.*

Proof. The proposition follows immediately from the uniform transversality of E^s and E^u and Lemma 5.9.2. \square

Let $\epsilon > 0, k, l \in \mathbb{N}$, and let $B_\epsilon^k \subset \mathbb{R}^k, B_\epsilon^l \subset \mathbb{R}^l$ be the ϵ -balls centered at the origin.

LEMMA 5.9.2. *For every $\epsilon > 0$ there is $\delta > 0$ such that if $\phi: B_\epsilon^k \rightarrow \mathbb{R}^l$ and $\psi: B_\epsilon^l \rightarrow \mathbb{R}^k$ are differentiable maps and $|\phi(x)|, \|d\phi(x)\|, |\psi(y)|, \|d\psi(y)\| < \delta$ for all $x \in B_\epsilon^k$ and $y \in B_\epsilon^l$, then the intersection $\text{graph}(\phi) \cap \text{graph}(\psi) \subset \mathbb{R}^{k+l}$ is transverse and consists of exactly one point, which depends continuously on ϕ and ψ in the C^1 topology.*

Proof. Exercise 5.9.3. □

The following property of hyperbolic sets plays a major role in their geometric description and is equivalent to local maximality. A hyperbolic set Λ has a *local product structure* if there are (small enough) $\epsilon > 0$ and $\delta > 0$ such that (i) for all $x, y \in \Lambda$ the intersection $W_\epsilon^s(x) \cap W_\epsilon^u(y)$ consists of at most one point, which belongs to Λ , and (ii) for $x, y \in \Lambda$ with $d(x, y) < \delta$, the intersection consists of exactly one point of Λ , denoted $[x, y] = W_\epsilon^s(x) \cap W_\epsilon^u(y)$, and the intersection is transverse (Proposition 5.9.1). If a hyperbolic set Λ has a local product structure, then for every $x \in \Lambda$ there is a neighborhood $U(x)$ such that

$$U(x) \cap \Lambda = \{[y, z]: y \in U(x) \cap W_\epsilon^s(x), z \in U(x) \cap W_\epsilon^u(x)\}.$$

PROPOSITION 5.9.3. *A hyperbolic set Λ is locally maximal if and only if it has a local product structure.*

Proof. Suppose Λ is locally maximal. If $x, y \in \Lambda$ and $\text{dist}(x, y)$ is small enough, then by Proposition 5.9.1, $W_\epsilon^s(x) \cap W_\epsilon^u(y) = [x, y] =: z$ exists and, by Theorem 5.6.4(4), the forward and backward semiorbits of z stay close to Λ . Since Λ is locally maximal, $z \in \Lambda$.

Conversely, assume that Λ has a local product structure with constants ϵ, δ , and C_p from Proposition 5.9.1. We must show that if the whole orbit of a point q lies close to Λ , then the point lies in Λ . Fix $\alpha \in (0, \delta/3)$ such that $f(p) \in W_{\delta/3}^u(f(x))$ for each $x \in \Lambda$ and $p \in W_\alpha^u(x)$. Assume first that $q \in W_\alpha^u(x_0)$ for some $x_0 \in \Lambda$ and that there are $y_n \in \Lambda$ such that $d(f^n(q), y_n) < \alpha/C_p$ for all $n > 0$. Since $f(x_0), y_1 \in \Lambda$ and $d(f(x_0), y_1) < d(f(x_0), f(q)) + d(f(q), y_1) < \delta/3 + \alpha/C_p < \delta$, we have that $x_1 = [y_1, f(x_0)] \in \Lambda$ and, by Proposition 5.9.1, $f(q) \in W_\alpha^u(x_1)$. Similarly, $x_2 = [y_2, f(x_1)] \in \Lambda$ and $f^2(q) \in W_\alpha^u(x_2)$. By repeating this argument we construct points $x_n = [y_n, f^n(q)] \in \Lambda$ with $f^n(q) \in W_\alpha^u(x_n)$. Observe that $q_n = f^{-n}(x_n) \rightarrow q$ as $n \rightarrow \infty$. Since Λ is closed, $q \in \Lambda$. Similarly, if $q \in W_\alpha^s(x_0)$ for some $x_0 \in \Lambda$ and $f^n(q)$ stays close enough to Λ for all $n < 0$, then $q \in \Lambda$.

Assume now that $f^n(y)$ is close enough to $x_n \in \Lambda$ for all $n \in \mathbb{Z}$. Then $y \in \Lambda_\epsilon^s$ and $y \in \Lambda_\epsilon^u$. Hence, by Propositions 5.4.4 and 5.4.3, the union $\Lambda \cup \mathcal{O}_f(y)$ is a hyperbolic set (with close constants), and the local stable and unstable

manifolds of y are well defined. Observe that the forward semiorbit of $p = [y, x_0]$ and the backward semiorbit of $q = [x_0, y]$ stay close to Λ . Therefore, by the above argument, $p, q \in \Lambda$ and (by the local product structure) $y = [p, q] \in \Lambda$. \square

Exercise 5.9.1. Prove that horseshoes (§5.8) and the solenoid (§1.9) are locally maximal hyperbolic sets.

Exercise 5.9.2. Let p be a hyperbolic fixed point of f and $q \in W^s(p) \cap W^u(p)$ a transverse homoclinic point. By Exercise 5.7.3, the union of p with the orbit of q is a hyperbolic set of f . Is it locally maximal?

Exercise 5.9.3. Prove Lemma 5.9.2.

5.10 Anosov Diffeomorphisms

Recall that a C^1 diffeomorphism f of a connected differentiable manifold M is called *Anosov* if M is a hyperbolic set for f ; it follows directly from the definition that M is locally maximal and compact.

The simplest example of an Anosov diffeomorphism is the automorphism of \mathbb{T}^2 given by the matrix $\begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}$. More generally, any linear hyperbolic automorphism of the n -torus \mathbb{T}^n is Anosov. Such an automorphism is given by an $n \times n$ integer matrix with determinant ± 1 and with no eigenvalues of modulus 1.

Toral automorphisms can be generalized as follows. Let N be a simply connected nilpotent Lie group, and Γ a uniform discrete subgroup of N . The quotient $M = N/\Gamma$ is a compact *nilmanifold*. Let \bar{f} be an automorphism of N that preserves Γ and whose derivative at the identity is hyperbolic. The induced diffeomorphism f of M is Anosov. For specific examples of this type see [Sma67]. Up to finite coverings, all known Anosov diffeomorphisms are topologically conjugate to automorphisms of nilmanifolds.

The families of stable and unstable manifolds of an Anosov diffeomorphism form two foliations (see §5.13) called the *stable foliation* W^s and *unstable foliation* W^u (Exercise 5.10.1). These foliations are in general not C^1 , or even Lipschitz [Ano67], but they are Hölder continuous (Theorem 6.1.3). In spite of the lack of Lipschitz continuity, the stable and unstable foliations possess a uniqueness property similar to the uniqueness theorem for ordinary differential equations (Exercise 5.10.2).

Proposition 5.10.1 states basic properties of the stable and unstable distributions E^s and E^u , and the stable and unstable foliations W^s and W^u , of an Anosov diffeomorphism f . These properties follow immediately from the previous sections of this chapter. We assume that the metric is adapted

to f and denoted by d^s and d^u , the distances along the stable and unstable leaves.

PROPOSITION 5.10.1. *Let $f: M \rightarrow M$ be an Anosov diffeomorphism. Then there are $\lambda \in (0, 1)$, $C_p > 0$, $\epsilon > 0$, $\delta > 0$, and, for every $x \in M$, a splitting $T_x M = E^s(x) \oplus E^u(x)$ such that*

1. $df_x(E^s(x)) = E^s(f(x))$ and $df_x(E^u(x)) = E^u(f(x))$;
2. $\|df_x v^s\| \leq \lambda \|v^s\|$ and $\|df_x^{-1} v^u\| \leq \lambda \|v^u\|$ for all $v^s \in E^s(x)$, $v^u \in E^u(x)$;
3. $W^s(x) = \{y \in M: d(f^n(x), f^n(y)) \rightarrow 0 \text{ as } n \rightarrow \infty\}$
and $d^s(f(x), f(y)) \leq \lambda d^s(x, y)$ for every $y \in W^s(x)$;
4. $W^u(x) = \{y \in M: d(f^{-n}(x), f^{-n}(y)) \rightarrow 0 \text{ as } n \rightarrow \infty\}$
and $d^u(f^{-1}(x), f^{-1}(y)) \leq \lambda^n d^u(x, y)$ for every $y \in W^u(x)$;
5. $f(W^s(x)) = W^s(f(x))$ and $f(W^u(x)) = W^u(f(x))$;
6. $T_x W^s(x) = E^s(x)$ and $T_x W^u(x) = E^u(x)$;
7. if $d(x, y) < \delta$, then the intersection $W_\epsilon^s(x) \cap W_\epsilon^u(y)$ is exactly one point $[x, y]$, which depends continuously on x and y , and $d^s([x, y], x) \leq C_p d(x, y)$, $d^u([x, y], y) \leq C_p d(x, y)$.

For convenience we restate several properties of Anosov diffeomorphisms. Recall that a diffeomorphism $f: M \rightarrow M$ is *structurally stable* if for every $\epsilon > 0$ there is a neighborhood $\mathcal{U} \subset \text{Diff}^1(M)$ of f such that for every $g \in \mathcal{U}$ there is a homeomorphism $h: M \rightarrow M$ with $h \circ f = g \circ h$ and $\text{dist}_0(h, \text{Id}) < \epsilon$.

PROPOSITION 5.10.2

1. Anosov diffeomorphisms form an open (possibly empty) subset in the C^1 topology (Corollary 5.5.2).
2. Anosov diffeomorphisms are structurally stable (Corollary 5.5.4).
3. The set of periodic points of an Anosov diffeomorphism is dense in the set of non-wandering points (Corollary 5.3.4).

Here is a more direct proof of the density of periodic points. Let ϵ and δ satisfy Proposition 5.10.1. If $x \in M$ is non-wandering, then there is $n \in \mathbb{N}$ and $y \in M$ such that $\text{dist}(x, y)$, $\text{dist}(f^n(y), y) < \delta/(2C_p)$. Assume that $\lambda^n < 1/(2C_p)$. Then the map $z \mapsto [y, f^n(z)]$ is well defined for $z \in W_\delta^s(y)$. It maps $W_\delta^s(y)$ into itself and, by the Brouwer fixed point theorem, has a fixed point y_1 such that $d^s(y_1, y) < \delta$, $f^n(y_1) \in W^u(y_1)$ and $d^u(y_1, f^n(y_1)) < \delta$. The map f^{-n} sends $W_\delta^u(f^n(y_1))$ to itself and therefore has a fixed point.

THEOREM 5.10.3. *Let $f: M \rightarrow M$ be an Anosov diffeomorphism. Then the following are equivalent:*

1. $\text{NW}(f) = M$,
2. every unstable manifold is dense in M ,

3. every stable manifold is dense in M ,
4. f is topologically transitive,
5. f is topologically mixing.

Proof. We say that a set A is ϵ -dense in a metric space (X, d) if $d(x, A) < \epsilon$ for every $x \in X$.

1 \Rightarrow 2: We will show that every unstable manifold is ϵ -dense in M for an arbitrary $\epsilon > 0$. By Proposition 5.10.2(3), the periodic points are dense. Assume that $\epsilon > 0$ satisfies Proposition 5.10.1(7) and that periodic points $x_i, i = 1, \dots, N$, form an $\epsilon/4$ -net in M . Let P be the product of the periods of the x_i s, and set $g = f^P$. Note that the stable and unstable manifolds of g are the same as those of f .

LEMMA 5.10.4. *There is $q \in \mathbb{N}$ such that if $\text{dist}(W^u(y), x_i) < \epsilon/2$ and $\text{dist}(x_i, x_j) < \epsilon/2$ for some $y \in M, i, j$, then $\text{dist}(g^{nq}(W^u(y)), x_i) < \epsilon/2$ and $\text{dist}(g^{nq}(W^u(y)), x_j) < \epsilon/2$ for every $n \in \mathbb{N}$.*

Proof. By Proposition 5.10.2(3), there is $z \in W^u(y) \cap W_{C_p \epsilon_p}^s(x_i)$. Therefore $\text{dist}(g^t(z), x_i) < \epsilon/2$ for any $t \geq t_0$, where t_0 depends on ϵ but not on z . Since $\text{dist}(g^t(z), x_j) < \epsilon$, by Proposition 5.10.2(3) there exists a point $w \in W^u(g^t(z)) \cap W_{C_p \epsilon_p}^s(x_j)$. Hence $\text{dist}(g^\tau(w), x_j) < \epsilon/2$ for any $\tau \geq s_0$ which depends only on ϵ but not on w . The lemma follows with $q = s_0 + t_0$. \square

Since M is compact and connected, any x_i can be connected to any x_j by a chain of not more than N periodic points x_k with distance $< \epsilon/2$ between any two consecutive points. By Lemma 5.10.4, $g^{Nq}(W^u(y))$ is ϵ -dense in M for any $y \in M$. Hence, $W^u(x)$ is ϵ -dense for any $x = g^{-Nq}(y) \in M$. Therefore, $W^u(x)$ is dense for each x . Reversing the time gives 1 \Rightarrow 3.

LEMMA 5.10.5. *If every (un)stable manifold is dense in M , then for every $\epsilon > 0$ there is $R = R(\epsilon) > 0$ such that every ball of radius R in every (un)stable manifold is ϵ -dense in M .*

Proof. Let $x \in M$. Since $W^u(x) = \bigcup_R W_R^u(x)$ is dense, there is $R(x)$ such that $W_{R(x)}^u(x)$ is $\epsilon/2$ -dense. Since W^u is a continuous foliation, there is $\delta(x) > 0$ such that $W_{R(x)}^u(y)$ is ϵ -dense for any $y \in B(x, \delta(x))$. By the compactness of M , a finite collection \mathcal{B} of the $\delta(x)$ -balls covers M . The maximal $R(x)$ for the balls from \mathcal{B} satisfies the lemma. \square

2 \Rightarrow 5: Let $U, V \subset M$ be non-empty open sets. Let $x, y \in M$ and $\delta > 0$ be such that $W_\delta^u(x) \subset U$ and $B(y, \delta) \subset V$, and let $R = R(\delta)$ (see Lemma 5.10.5). Since f expands unstable manifolds exponentially and uniformly, there is $N \in \mathbb{N}$ such that $f^n(W_\delta^u(x)) \supset W_R^u(f^n(x))$ for $n \geq N$. By Lemma 5.10.5, $f^n(U) \cap V \neq \emptyset$ and hence f is topologically mixing. Similarly 3 \Rightarrow 5.

1 \Rightarrow 3 follows by reversing the time. Obviously 5 \Rightarrow 4 and 4 \Rightarrow 1. \square

Exercise 5.10.1. Prove that the stable and unstable manifolds of an Anosov diffeomorphism form foliations (see §5.13).

Exercise 5.10.2. Although the stable and unstable distributions of an Anosov diffeomorphism, in general, are not Lipschitz continuous, the following uniqueness property holds true. Let $\gamma(\cdot)$ be a differentiable curve such that $\dot{\gamma}(t) \in E^s(\gamma(t))$ for every t . Prove that γ lies in one stable manifold.

5.11 Axiom A and Structural Stability

Some of the results of §5.10 extend to a natural wider class of hyperbolic dynamical systems. Throughout this section we assume that f is a diffeomorphism of a compact manifold M . Recall that the set of non-wandering points $\text{NW}(f)$ is closed and f -invariant, and that $\overline{\text{Per}(f)} \subset \text{NW}(f)$.

A diffeomorphism f satisfies Smale's *Axiom A* if the set $\text{NW}(f)$ is hyperbolic and $\text{Per}(f) = \text{NW}(f)$. The second condition does not follow from the first. By Proposition 5.3.3, the set $\text{Per}(f)$ is dense in the set $\text{NW}(f|_{\text{NW}(f)})$ of non-wandering points of the restriction of f to $\text{NW}(f)$. However, in general $\text{NW}(f|_{\text{NW}(f)}) \neq \text{NW}(f)$ (Exercise 5.11.1, Exercise 5.11.2).

For a hyperbolic periodic point p of f , denote by $W^s(O(p))$ and $W^u(O(p))$ the unions of the stable and unstable manifolds of p and its images, respectively. If p and q are hyperbolic periodic points, we write $p \leq q$ when $W^s(O(p))$ and $W^u(O(q))$ have a point of transverse intersection. The relation \leq is reflexive. It follows from Theorem 5.7.2 that \leq is transitive (Exercise 5.11.3). If $p \leq q$ and $q \leq p$, we write $p \sim q$ and say that p and q are *heteroclinically related*. The relation \sim is an equivalence relation.

THEOREM 5.11.1 (Smale's Spectral Decomposition [Sma67]). *If f satisfies Axiom A, then there is a unique representation of $\text{NW}(f)$,*

$$\text{NW}(f) = \Lambda_1 \cup \Lambda_2 \cup \cdots \cup \Lambda_k,$$

as a disjoint union of closed f -invariant sets (called basic sets) such that

1. each Λ_i is a locally maximal hyperbolic set of f ;
2. f is topologically transitive on each Λ_i ; and
3. each Λ_i is a disjoint union of closed sets Λ_i^j , $1 \leq j \leq m_i$, the diffeomorphism f cyclically permutes the sets Λ_i^j , and f^{m_i} is topologically mixing on each Λ_i^j .

The basic sets are precisely the closures of the equivalence classes of \sim . For two basic sets, we write $\Lambda_i \leq \Lambda_j$ if there are periodic points $p \in \Lambda_i$ and $q \in \Lambda_j$ such that $p \leq q$.

Let f satisfy Axiom A. We say that f satisfies the *strong transversality condition* if $W^s(x)$ intersects $W^u(y)$ transversely (at all common points) for all $x, y \in \text{NW}(f)$.

THEOREM 5.11.2 (Structural Stability Theorem). *A C^1 diffeomorphism is structurally stable if and only if it satisfies Axiom A and the strong transversality condition.*

J. Robbin [Rob71] showed that a C^2 diffeomorphism satisfying Axiom A and the strong transversality condition is structurally stable. C. Robinson [Rob76] weakened C^2 to C^1 . R. Mañé [Mañ88] proved that a structurally stable C^1 diffeomorphism satisfies Axiom A and the strong transversality condition.

Exercise 5.11.1. Give an example of a diffeomorphism f such that $\text{NW}(f|_{\text{NW}(f)}) \neq \text{NW}(f)$.

Exercise 5.11.2. Give an example of a diffeomorphism f for which $\text{NW}(f)$ is hyperbolic and $\text{NW}(f|_{\text{NW}(f)}) \neq \text{NW}(f)$.

Exercise 5.11.3. Prove that \leq is a transitive relation.

Exercise 5.11.4. Suppose that f satisfies Axiom A. Prove that $\text{NW}(f)$ is a locally maximal hyperbolic set.

5.12 Markov Partitions

Recall (Chapter 1, Chapter 3) that a partition of the phase space of a dynamical system induces a coding of the orbits and hence a semiconjugacy with a subshift. For hyperbolic dynamical systems, there is a special class of partitions – *Markov partitions* – for which the target subshift is a subshift of finite type. A Markov partition \mathcal{P} for an invariant subset Λ of a diffeomorphism f of a compact manifold M is a collection of sets R_i called *rectangles* such that for all i, j, k

1. each R_i is the closure of its interior,
2. $\text{int } R_i \cap \text{int } R_j = \emptyset$ if $i \neq j$,
3. $\Lambda \subset \bigcup_i R_i$,
4. if $f^m(\text{int } R_i) \cap \text{int } R_j \cap \Lambda \neq \emptyset$ for some $m \in \mathbb{Z}$ and $f^n(\text{int } R_j) \cap \text{int } R_k \cap \Lambda \neq \emptyset$ for some $n \in \mathbb{Z}$, then $f^{m+n}(\text{int } R_i) \cap \text{int } R_k \cap \Lambda \neq \emptyset$.

The last condition guarantees the Markov property of the subshift corresponding to \mathcal{P} , i.e., the independence of the future from the past. For hyperbolic dynamical systems, each rectangle is closed under the local




Figure 5.5. Markov partition for the toral automorphism f_M .

product structure “commutator” $[x, y]$, i.e., if $x, y \in R_i$, then $[x, y] \in R_i$. For $x \in R_i$ let $W^s(x, R_i) = \bigcup_{y \in R_i} [x, y]$ and $W^u(x, R_i) = \bigcup_{y \in R_i} [y, x]$. The last condition means that if $x \in \text{int } R_i$ and $f(x) \in \text{int } R_j$, then $W^u(f(x), R_j) \subset f(W^u(x, R_i))$ and $W^s(x, R_i) \subset f^{-1}(W^s(f(x), R_j))$.

The partition of the unit interval $[0, 1]$ into m intervals $[k/m, (k + 1)/m]$ is a Markov partition for the expanding endomorphism E_m . The target subshift in this case is the full shift on m symbols.

We now describe a Markov partition for the hyperbolic toral automorphism $f = f_M$ given by the matrix

$$M = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix},$$

which was constructed by R. Adler and B. Weiss [AW67]. The eigenvalues are $(3 \pm \sqrt{5})/2$. We begin by partitioning the unit square representing the torus \mathbb{T}^2 in Figure 5.5 into two rectangles: A , consisting of three parts A_1, A_2, A_3 ; and B , consisting of two parts B_1, B_2 . The longer sides of the rectangles are parallel to the eigendirection of the larger eigenvalue $(3 + \sqrt{5})/2$, and the shorter sides are parallel to the eigendirection of the smaller eigenvalue $(3 - \sqrt{5})/2$. In Figure 5.5, the identified points and regions are marked by the same symbols. The images of A and B are shown in Figure 5.6. We subdivide




Figure 5.6. The image of the Markov partition under f_M .

A and B into five subrectangles $\Delta_1, \Delta_2, \Delta_3, \Delta_4, \Delta_5$ that are the connected components of the intersections of A and B with $f(A)$ and $f(B)$. The image of A consists of Δ_1, Δ_3' and Δ_4' ; the image of B consists of Δ_2' and Δ_5' . The part of the boundary of the Δ_i 's that is parallel to the eigendirection of the larger eigenvalue is called *stable*; the part that is parallel to the eigendirection of the smaller eigenvalue is called *unstable*. By construction, the partition Δ of \mathbb{T}^2 into five rectangles Δ_i has the property that the image of the stable boundary is contained in the stable boundary, and the preimage of the unstable boundary is contained in the unstable boundary (Exercise 5.12.1). In other words, for each i, j , the intersection $\Delta_{ij} = \Delta_i \cap f(\Delta_j)$ consists of one or two rectangles that stretch “all the way” through Δ_i , and the stable boundary of Δ_{ij} is contained in the stable boundary of Δ_i ; similarly, the intersection $\Delta_{ij}^{-1} = \Delta_i \cap f^{-1}(\Delta_j)$ consists of one or two rectangles that stretch “all the way” through Δ_i , and the unstable boundary of Δ_{ij}^{-1} is contained in the unstable boundary of Δ_i . Let $a_{ij} = 1$ if the interior of $f(\Delta_i) \cap \Delta_j$ is not empty, and $a_{ij} = 0$ otherwise, $i, j = 1, \dots, 5$. This defines the adjacency matrix

$$A = \begin{pmatrix} 1 & 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 1 & 0 \\ 0 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 \end{pmatrix}.$$

If $\omega = (\dots, \omega_{-1}, \omega_0, \omega_1, \dots)$ is an allowed infinite sequence for this adjacency matrix, then the intersection $\bigcap_{i=-\infty}^{\infty} f^{-i}(\Delta_{\omega_i})$ consists of exactly one point $\phi(\omega)$; it follows that there is a continuous semiconjugacy $\phi: \Sigma_A \rightarrow \mathbb{T}^2$, i.e., $f \circ \phi = \phi \circ \sigma$, where σ is the shift in Σ_A (Exercise 5.12.2). Conversely, let B_0 be the union of the boundaries of the Δ_i 's, and let $B = \bigcup_{i=-\infty}^{\infty} f^i(B_0)$. For $x \in \mathbb{T}^2 \setminus B$, set $\psi_i(x) = j$ if $f^i(x) \in \Delta_j$. The itinerary sequence $(\psi_i(x))_{i=-\infty}^{\infty}$ is an element of Σ_A , and $\phi \circ \psi = \text{Id}$ (Exercise 5.12.3).

In higher dimensions, this direct geometric construction does not work. Even for a hyperbolic toral automorphism, the boundary is nowhere differentiable. Nevertheless, as R. Bowen showed [Bow70], any locally maximal hyperbolic set Λ has a Markov partition [Bow70] which provides a semiconjugacy from a subshift of finite type to Λ .

Exercise 5.12.1. Prove that the stable boundary is forward invariant and the unstable boundary is backward invariant under f_M .

Exercise 5.12.2. Prove that for the toral automorphism f_M , the intersection of the preimages of rectangles Δ_i along an allowed infinite sequence ω consists of exactly one point. Prove that there is a semiconjugacy ϕ from $\sigma|_{\Sigma_A}$ to the toral automorphism f_M .

Exercise 5.12.3. Prove that the map ψ defined in the text above satisfies $\psi(x) \in \Sigma_A$ and that $\phi \circ \psi = \text{Id}$.

Exercise 5.12.4. Construct Markov partitions for the linear horseshoe (§1.8) and the solenoid (§1.9).

5.13 Appendix: Differentiable Manifolds

An m -dimensional C^k manifold M is a second-countable Hausdorff topological space together with a collection \mathcal{U} of open sets in M and for each $U \in \mathcal{U}$ a homeomorphism ϕ_U from U onto the unit ball $B^m \subset \mathbb{R}^m$ such that:

1. \mathcal{U} is a cover of M , and
2. for $U, V \in \mathcal{U}$, if $U \cap V \neq \emptyset$, the map $\phi_U \circ \phi_V^{-1}: \phi_V(U \cap V) \rightarrow \phi_U(U \cap V)$ is C^k .

We may take $k \in \mathbb{N} \cup \{\infty, \omega\}$, where C^ω denotes the class of real analytic functions.

We write M^m to indicate that M has dimension m . If $x \in M$ and $U \in \mathcal{U}$ contains x , then the pair (U, ϕ_U) , $U \in \mathcal{U}$, is called a *coordinate chart at x* , and the n component functions x_1, x_2, \dots, x_m of ϕ_U are called *coordinates* on

U . The collection of coordinate charts $\{(U, \phi_U)\}_{U \in \mathcal{U}}$ is called an *atlas* on M . Note that any open subset of \mathbb{R}^m is a C^k manifold, for any $k \in \mathbb{N} \cup \{\infty, \omega\}$.

If M^m and N^n are C^k manifolds, then a continuous map $f: M \rightarrow N$ is C^k if for any coordinate chart (U, ϕ_U) on M , and any coordinate chart (V, ψ_V) on N , the map $\psi_V \circ f \circ \phi_U^{-1}: \phi_U(U \cap f^{-1}(V)) \rightarrow \mathbb{R}^n$ is a C^k map. For $k \geq 0$, the set of C^k maps from M to N is denoted $C^k(M, N)$. We say that a sequence of functions $f_n \in C^k(M, N)$ converges if the functions and all their derivatives up to order k converge uniformly on compact sets. This defines a topology on $C^k(M, N)$ called the C^k topology.

We set $C^k(M) = C^k(M, \mathbb{R})$. The subset of $C^k(M, M)$ consisting of diffeomorphisms of M is denoted $\text{Diff}^k(M)$.

A C^k curve in M^m is a C^k map $\alpha: (-\epsilon, \epsilon) \rightarrow M$. The *tangent vector* to α at $\alpha(0) = p$ is the linear map $v: C^1(M) \rightarrow \mathbb{R}$ defined by

$$v(f) = \left. \frac{d}{dt} \right|_{t=0} f(\alpha(t))$$

for $f \in C^1(M)$. The *tangent space* at p is the linear space $T_p M$ of all tangent vectors at p .

Suppose (U, ϕ) is a coordinate chart, with coordinate functions x_1, \dots, x_m , and let $p \in U$. For $i = 1, \dots, m$, consider the curves

$$\alpha_i^p(t) = \phi^{-1}(x_1(p), \dots, x_{i-1}(p), x_i(p) + t, x_{i+1}(p), \dots, x_m(p)).$$

Define $(\partial/\partial x_i)_p$ to be the tangent vector to α_i^p at p , i.e., for $g \in C^1(M)$,

$$\left(\frac{\partial}{\partial x_i} \right)_p (g) = \left. \frac{d}{dt} \right|_{t=0} g(\alpha_i^p(t)) = \left(\frac{\partial}{\partial x_i} (g \circ \phi) \right)_{\phi(p)}.$$

The vectors $\partial/\partial x_i$, $i = 1, \dots, m$, are linearly independent at p , and span $T_p M$. In particular, $T_p M$ is a vector space of dimension m .

Let $f: M \rightarrow N$ be a C^k map, $k \geq 1$. For $p \in M$, the *tangent map* $df_p: T_p M \rightarrow T_{f(p)} N$ is defined by $df_p(v)(g) = v(g \circ f)$, for $g \in C^1(N)$. In terms of curves, if v is tangent to α at $p = \alpha(0)$, then $df_p(v)$ is tangent to $f \circ \alpha$ at $f(p)$.

The *tangent bundle* $TM = \bigcup_{x \in M} T_x M$ of M is a C^{k-1} manifold of twice the dimension of M with coordinate charts defined as follows. Let (U, ϕ_U) be a coordinate chart on M , $\phi_U = (x_1, \dots, x_m): U \rightarrow \mathbb{R}^m$. For each i , the derivative dx_i is a function from $TU = \bigcup_{p \in U} T_p M$ to \mathbb{R} , defined by $dx_i(v) = v(x_i)$, for $v \in TU$. The function $(x_1, \dots, x_m, dx_1, \dots, dx_m): TU \rightarrow \mathbb{R}^{2m}$ is a coordinate chart on TU , which we denote $d\phi_U$. Note that if $y, w \in \mathbb{R}^m$, then

$$d\phi_U \circ d\phi_V^{-1}(y, w) = (\phi_U \circ \phi_V^{-1}(y), d(\phi_U \circ \phi_V^{-1})_y(w)).$$

Let $\pi: TM \rightarrow M$ be the *projection map* that sends a vector $v \in T_p M$ to its base point p . A C^r *vector field* X on M is a C^r map $X: M \rightarrow TM$ such that $\pi \circ X$ is the identity on M . We write $X_p = X(p)$.

Let M^m and N^n be C^k manifolds. We say that M is a C^k *submanifold* of N if M is a subset of N and the inclusion map $i: M \rightarrow N$ is C^k and has rank m for each $x \in M$. If the topology of M coincides with the subspace topology, then M is an *embedded submanifold*. For each $x \in M$, the tangent space $T_x M$ is naturally identified with a subspace of $T_x N$. Two submanifolds $M_1, M_2 \subset N$ of complementary dimensions *intersect transversely* (or are *transverse*) at a point $p \in N_1 \cap N_2$ if $T_p N = T_p M_1 \oplus T_p M_2$.

A *distribution* E on a differentiable manifold M is a family of k -dimensional subspaces $E(x) \subset T_x M$, $x \in M$. The distribution is C^l , $l \geq 0$, if locally it is spanned by $k C^l$ vector fields.

Suppose W is a partition of a differentiable manifold M into C^1 submanifolds of dimension k . For $x \in M$, let $W(x)$ be the submanifold containing x . We say that W is a *k-dimensional continuous foliation with C^1 leaves* (or simply a *foliation*) if every $x \in M$ has a neighborhood U and a homeomorphism $h: B^k \times B^{m-k} \rightarrow U$ such that

1. for each $z \in B^{m-k}$, the set $h(B^k \times \{z\})$ is the connected component of $W(h(0, z)) \cap U$ containing $h(0, z)$, and
2. $h(\cdot, z)$ is C^1 and depends continuously on z in the C^1 topology.

The pair (U, h) is called a *foliation coordinate chart*. The sets $h(B^k \times \{z\})$ are called *local leaves* (or *plaques*), and the sets $h(\{y\} \times B^{m-k})$ are called *local transversals*. For $x \in U$, we denote by $W_U(x)$ the local leaf containing x . More generally, a differentiable submanifold $L^{m-k} \subset M$ is a *transversal* if L is transverse to the leaves of the foliation. Each submanifold $W(x)$ of the foliation is called a *leaf* of W .

A continuous foliation W is a C^k foliation, $k \geq 1$, if the maps h can be chosen to be C^k . For example, lines of constant slope on \mathbb{T}^2 form a C^∞ foliation.

A foliation W defines a distribution $E = TW$ consisting of the tangent spaces to the leaves. A distribution E is *integrable* if it is tangent to a foliation.

A C^k *Riemannian metric* on a C^{k+1} manifold M consists of a positive definite symmetric bilinear form $\langle \cdot, \cdot \rangle_p$ in each tangent space $T_p M$ such that for any C^k vector fields X and Y , the function $p \mapsto \langle X_p, Y_p \rangle_p$ is C^k . For each $v \in T_p M$, we write $\|v\| = (\langle v, v \rangle_p)^{1/2}$. If $\alpha: [a, b] \rightarrow M$ is a differentiable curve, we define the length of α to be $\int_a^b \|\dot{\alpha}(s)\| ds$. The (intrinsic) distance d between two points in M is defined to be the infimum of the lengths of differentiable curves in M connecting the two points.

A C^k Riemannian manifold is a C^{k+1} manifold with a C^k Riemannian metric. We denote by $T^1 M$ the set of tangent vectors of length 1 in a Riemannian manifold M .

A Riemannian manifold carries a natural measure called the *Riemannian volume*. Roughly speaking, the Riemannian metric allows one to compute the Jacobian of a differentiable map, and therefore allows one to define integration in a coordinate-free way.

If X is a topological space and (Y, d) is a metric space with metric, define a metric dist_0 on $C(X, Y)$ by

$$\text{dist}_0(f, g) = \min \left\{ 1, \sup_{x \in X} \max\{d(f(x), g(x))\} \right\}.$$

If X is compact, then this metric induces the topology of uniform convergence on compact sets. If X is not compact, this metric induces a finer topology. For example, the sequence of functions $f_n(x) = x^n$ in $C((0, 1), \mathbb{R})$ converges to 0 in the topology of uniform convergence on compact sets, but not in the metric dist_0 . The topology of uniform convergence on compact sets is metrizable even for non-compact sets, but we will not need this metric.

If M^m and N^n are C^1 Riemannian manifolds, we define a distance function dist_1 on $C^1(M, N)$ as follows: The Riemannian metric on N induces a metric (distance function) on the tangent bundle TN , making TN a metric space. For $f \in C^1(M, N)$, the differential of f gives a map $df: T^1 M \rightarrow TN$ on the unit tangent bundle of M . We set $\text{dist}_1(f, g) = \text{dist}_0(df, dg)$. If M is compact, the topology induced by this metric is the C^1 topology.

A differentiable manifold M is a (differentiable) *fiber bundle* over a differentiable manifold N with *fiber* F and (differentiable) *projection* $\pi: M \rightarrow N$ if for every $x \in N$ there is a neighborhood $V \ni x$ such that $\pi^{-1}(V)$ is diffeomorphic to $V \times F$ and $\pi^{-1}(y) \cong y \times F$. A diffeomorphism $f: M \rightarrow M$ is an *extension of* or a *skew product over* a diffeomorphism $g: N \rightarrow N$ if $\pi \circ f = g \circ \pi$; in this case g is called a *factor of* f .

Ergodicity of Anosov Diffeomorphisms

The purpose of this chapter is to establish the ergodicity of volume-preserving Anosov diffeomorphisms (Theorem 6.3.1). This result, which was first obtained by D. Anosov [Ano69] (see also [AS67]), shows that hyperbolicity has strong implications for the ergodic properties of a dynamical system. Moreover, since a small perturbation of an Anosov diffeomorphism is also Anosov (Proposition 5.10.2), this gives an open set of ergodic diffeomorphisms.

Our proof is an improvement of the arguments in [Ano69] and [AS67]. It is based on the classical approach called *Hopf's argument*. The first observation is that any f -invariant function is constant mod 0 on the stable and unstable manifolds (Lemma 6.3.2). Since these manifolds have complementary dimensions, one would expect the Fubini theorem to imply that the function is constant mod 0, and ergodicity would follow. The major difficulty is that, although the stable and unstable manifolds are differentiable, they need not depend differentiably on the point they pass through, even if f is real analytic. Thus the local product structure defined by the stable and unstable foliations does not yield a differentiable coordinate system, and we cannot apply the usual Fubini theorem. So we establish a property of the stable and unstable foliations called *absolute continuity* that implies the Fubini theorem.

The reason the stable and unstable manifolds do not vary differentiably is that they depend on the infinite future and past, respectively.

6.1 Hölder Continuity of the Stable and Unstable Distributions

For a subspace $A \subset \mathbb{R}^N$ and a vector $v \in \mathbb{R}^N$, set

$$\text{dist}(v, A) = \min_{w \in A} \|v - w\|.$$

For subspaces A, B in \mathbb{R}^N , define

$$\text{dist}(A, B) = \max \left(\max_{v \in A, \|v\|=1} \text{dist}(v, B), \max_{w \in B, \|w\|=1} \text{dist}(w, A) \right).$$

The following lemmas can be used to prove the Hölder continuity of invariant distributions for a variety of dynamical systems. Our objective is the Hölder continuity of the stable and unstable distributions of an Anosov diffeomorphism, which was first established by Anosov [Ano67]. We consider only the stable distribution; Hölder continuity of the unstable distribution follows by reversing the time.

LEMMA 6.1.1. *Let $L_n^i: \mathbb{R}^N \rightarrow \mathbb{R}^N, i = 1, 2, n \in \mathbb{N}$, be two sequences of linear maps. Assume that for some $b > 0$ and $\delta \in (0, 1)$,*

$$\|L_n^1 - L_n^2\| \leq \delta b^n$$

for each positive integer n .

Suppose that there are two subspaces $E^1, E^2 \subset \mathbb{R}^N$ and positive constants $C > 1$ and $\lambda < \mu$ with $\lambda < b$ such that

$$\begin{aligned} \|L_n^i v\| &\leq C\lambda^n \|v\| && \text{if } v \in E^i, \\ \|L_n^i w\| &\geq C^{-1}\mu^n \|w\| && \text{if } w \perp E^i. \end{aligned}$$

Then

$$\text{dist}(E^1, E^2) \leq 3C^2 \frac{\mu}{\lambda} \delta^{(\log \mu - \log \lambda) / (\log b - \log \lambda)}.$$

Proof. Set $K_n^1 = \{v \in \mathbb{R}^N: \|L_n^1 v\| \leq 2C\lambda^n \|v\|\}$. Let $v \in K_n^1$. Write $v = v^1 + v_\perp^1$, where $v^1 \in E^1$ and $v_\perp^1 \perp E^1$. Then

$$\|L_n^1 v\| = \|L_n^1(v^1 + v_\perp^1)\| \geq \|L_n^1 v_\perp^1\| - \|L_n^1 v^1\| \geq C^{-1}\mu^n \|v_\perp^1\| - C\lambda^n \|v^1\|,$$

and hence

$$\|v_\perp^1\| \leq C\mu^{-n} (\|L_n^1 v\| + C\lambda^n \|v^1\|) \leq 3C^2 \left(\frac{\lambda}{\mu}\right)^n \|v\|.$$

It follows that

$$\text{dist}(v, E^1) \leq 3C^2 \left(\frac{\lambda}{\mu}\right)^n \|v\|. \tag{6.1}$$

Set $\gamma = \lambda/b < 1$. There is a unique non-negative integer k such that $\gamma^{k+1} < \delta \leq \gamma^k$. Let $v^2 \in E^2$. Then

$$\begin{aligned} \|L_k^1 v^2\| &\leq \|L_k^2 v^2\| + \|L_k^1 - L_k^2\| \cdot \|v^2\| \\ &\leq C\lambda^k \|v^2\| + b^k \delta \|v^2\| \\ &\leq (C\lambda^k + (b\gamma)^k) \|v^2\| \leq 2C\lambda^k \|v^2\|. \end{aligned}$$

It follows that $v^2 \in K_k^1$ and hence $E^2 \subset K_k^1$. By symmetry, $E^1 \subset K_k^2$. By (6.1) and by the choice of k ,

$$\text{dist}(E^1, E^2) \leq 3C^2 \left(\frac{\lambda}{\mu}\right)^k \leq 3C^2 \frac{\mu}{\lambda} \delta^{(\log \mu - \log \lambda)/(\log b - \log \lambda)}. \quad \square$$

LEMMA 6.1.2. *Let f be a C^2 diffeomorphism of a compact C^2 submanifold $M \subset \mathbb{R}^N$. Then for each $n \in \mathbb{N}$ and all $x, y \in M$,*

$$\|df_x^n - df_y^n\| \leq b^n \cdot \|x - y\|,$$

where $b = \max_{z \in M} \|df_z\| (1 + \max_{z \in M} \|d_z^2 f\|)$.

Proof. Let $b_1 = \max_{z \in M} \|df_z\| \geq 1$ and $b_2 = \max_{z \in M} \|d_z^2 f\|$, so that $b = b_1(1 + b_2)$. Observe that $\|f^n(x) - f^n(y)\| \leq (b_1)^n \|x - y\|$ for all $x, y \in M$. The lemma obviously holds for $n = 1$. For the inductive step we have

$$\begin{aligned} \|df_x^{n+1} - df_y^{n+1}\| &\leq \|df_{f^n(x)}\| \cdot \|df_x^n - df_y^n\| + \|df_{f^n(x)} - df_{f^n(y)}\| \cdot \|df_y^n\| \\ &\leq b_1 b^n \|x - y\| + b_2 b_1^n \|x - y\| b_1 \leq b^{n+1} \|x - y\|. \quad \square \end{aligned}$$

Let M be a manifold embedded in \mathbb{R}^N , and suppose E is a distribution on M . We say that E is Hölder continuous with Hölder exponent $\alpha \in (0, 1]$ and Hölder constant L if

$$\text{dist}(E(x), E(y)) \leq L \cdot \|x - y\|^\alpha$$

for all $x, y \in M$ with $\|x - y\| \leq 1$.

One can define Hölder continuity for a distribution on an abstract Riemannian manifold by using parallel transport along geodesics to identify tangent spaces at nearby points. However, for a compact manifold M it suffices to consider Hölder continuity for some embedding of M in \mathbb{R}^N . This is so because on a compact manifold M , the ratio of any two Riemannian metrics is bounded above and below. So is the ratio between the intrinsic distance function on M and the extrinsic distance on M obtained by restricting the distance in \mathbb{R}^N to M . Thus the Hölder exponent is independent of both the Riemannian metric and the embedding, but the Hölder constant does change. So, without loss of generality, and to simplify the arguments in this section and the next one, we will deal only with manifolds embedded in \mathbb{R}^N .

THEOREM 6.1.3. *Let M be a compact C^2 manifold and $f: M \rightarrow M$ a C^2 Anosov diffeomorphism. Suppose that $0 < \lambda < 1 < \mu$ and $C > 0$ are such that $\|df_x^n v^s\| \leq C\lambda^n \|v^s\|$ and $\|df_x^n v^u\| \geq C\mu^n \|v^u\|$ for all $x \in M$, $v_s \in E^s(x)$, $v^u \in E^u(x)$, and $n \in \mathbb{N}$. Set $b = \max_{z \in M} \|df_z\| (1 + \max_{z \in M} \|d_z^2 f\|)$. Then the*

stable distribution E^s is Hölder continuous with exponent $\alpha = (\log \mu - \log \lambda) / (\log b - \log \lambda)$.

Proof. As indicated above, we may assume that M is embedded in \mathbb{R}^N . For $x \in M$, let $E^\perp(x)$ denote the orthogonal complement to the tangent plane $T_x M$ in \mathbb{R}^N . Since E^\perp is a smooth distribution, it is sufficient to prove the Hölder continuity of $E^s \oplus E^\perp$ on M .

Since M is compact, there is a constant $\bar{C} > 1$ such that for any $x \in M$, if $v \in T_x M$ is perpendicular to E^s , then $\|df_x^n v\| \geq \bar{C}^{-1} \mu^n \|v\|$.

For $x \in M$, extend df_x to a linear map $L(x): \mathbb{R}^N \rightarrow \mathbb{R}^N$ by setting $L(x)|_{E^\perp(x)} = 0$, and set $L_n(x) = L(f^{n-1}(x)) \circ \cdots \circ L(f(x)) \circ L(x)$. Note that $L_n(x)|_{T_x M} = df_x^n$.

Fix $x_1, x_2 \in M$ with $\|x_1 - x_2\| < 1$. By Lemma 6.1.2, the conditions of Lemma 6.1.1 are satisfied with $L_n^i = L_n(x_i)$ and $E^i = E^s(x_i)$, $i = 1, 2$, and the theorem follows. \square

Exercise 6.1.1. Let $\beta \in (0, 1]$, and M be a compact $C^{1+\beta}$ manifold, i.e., the first derivatives of the coordinate functions are Hölder continuous with exponent β . Let $f: M \rightarrow M$ be a $C^{1+\beta}$ Anosov diffeomorphism. Prove that the stable and unstable distributions of f are Hölder continuous.

6.2 Absolute Continuity of the Stable and Unstable Foliations

Let M be a smooth n -dimensional manifold. Recall (§5.13) that a continuous k -dimensional foliation W with C^1 leaves is a partition of M into C^1 submanifolds $W(x) \ni x$ which locally depend continuously in the C^1 topology on $x \in M$. Denote by m the Riemannian volume in M , and by m_N the induced Riemannian volume in a C^1 submanifold N . Note that every leaf $W(x)$ and every transversal carry an induced Riemannian volume.

Let (U, h) be a foliation coordinate chart on M (§5.13), and let $L = h(\{y\} \times B^{n-k})$ be a C^1 local transversal. The foliation W is called *absolutely continuous* if for any such L and U there is a measurable family of positive measurable functions $\delta_x: W_U(x) \rightarrow \mathbb{R}$ (called the *conditional densities*) such that for any measurable subset $A \subset U$

$$m(A) = \int_L \int_{W_U(x)} \mathbf{1}_A(x, y) \delta_x(y) dm_{W(x)}(y) dm_L(x).$$

Note that the conditional densities are automatically integrable.

PROPOSITION 6.2.1. *Let W be an absolutely continuous foliation of a Riemannian manifold M , and let $f: M \rightarrow M$ be a measurable function.*




Figure 6.1. Holonomy map p for a foliation W and transversals U_1 and U_2 .

Suppose there is a set $A \subset M$ of measure 0 such that f is constant on $W(x) \setminus A$ for every leaf $W(x)$.

Then f is essentially constant on almost every leaf, i.e., for any transversal L , the function f is $m_{W(x)}$ -essentially constant for m_L -almost every $x \in L$.

Proof. Absolute continuity implies that $m_{W(x)}(A \cap W(x)) = 0$ for m_L -almost every $x \in L$. □

Absolute continuity of the stable and unstable foliations is the property we need in order to prove the ergodicity of Anosov diffeomorphisms. However, we will prove a stronger property, called transverse absolute continuity; see Proposition 6.2.2.

Let W be a foliation of M , and (U, h) a foliation coordinate chart. Let $L_i = h(\{y_i\} \times B^{m-k})$ for $y_i \in B^k$, $i = 1, 2$. Define a homeomorphism $p: L_1 \rightarrow L_2$ by $p(h(y_1, z)) = h(y_2, z)$, for $z \in B^{m-k}$; p is called the *holonomy map* (see Figure 6.1). The foliation W is *transversely absolutely continuous* if the holonomy map p is absolutely continuous for any foliation coordinate chart and any transversals L_i as above, i.e., if there is a positive measurable function $q: L_1 \rightarrow \mathbb{R}$ (called the *Jacobian* of p) such that for any measurable subset $A \subset L_1$

$$m_{L_2}(p(A)) = \int_{L_1} \mathbf{1}_A q(z) dm_{L_1}(z).$$

If the Jacobian q is bounded on compact subsets of L_1 , then W is said to be *transversely absolutely continuous with bounded Jacobians*.

PROPOSITION 6.2.2. *If W is transversely absolutely continuous, then it is absolutely continuous.*




Figure 6.2. Holonomy maps for W and F .

Proof. Let L and U be as in the definition of an absolutely continuous foliation, $x \in L$ and let F be an $(n - k)$ -dimensional C^1 -foliation such that $F(x) \supset L$, $F_U(x) = L$, and $U = \bigcup_{y \in W_U(x)} F_U(y)$; see Figure 6.2. Obviously, F is absolutely continuous and transversely absolutely continuous. Let $\bar{\delta}_y(\cdot)$ denote the conditional densities for F . Since F is a C^1 foliation, $\bar{\delta}$ is continuous and hence measurable. For any measurable set $A \subset U$, by the Fubini theorem,

$$m(A) = \int_{W_U(x)} \int_{F_U(y)} \mathbf{1}_A(y, z) \bar{\delta}_y(z) dm_{F_U(y)}(z) dm_{W_U(x)}(y). \quad (6.2)$$

Let p_y denote the holonomy map along the leaves of W from $F_U(x) = L$ to $F_U(y)$, and let $q_y(\cdot)$ denote the Jacobian of p_y . We have

$$\int_{F_U(y)} \mathbf{1}_A(y, z) \bar{\delta}_y(z) dm_{F_U(y)}(z) = \int_L \mathbf{1}_A(p_y(s)) q_y(s) \bar{\delta}_y(p_y(s)) dm_L(s),$$

and by changing the order of integration in (6.2), which is an integral with respect to the product measure, we get

$$m(A) = \int_L \int_{W_U(x)} \mathbf{1}_A(p_y(s)) q_y(s) \bar{\delta}_y(p_y(s)) dm_{W_U(x)}(y) dm_L(s). \quad (6.3)$$

Similarly, let \bar{p}_s denote the holonomy map along the leaves of F from $W_U(x)$ to $W_U(s)$, $s \in L$, and let \bar{q}_s denote the Jacobian of \bar{p}_s . We transform the integral over $W_U(x)$ into an integral over $W_U(s)$ using the change of variables

$$r = p_y(s), y = \bar{p}_s^{-1}(r):$$

$$\begin{aligned} & \int_{W_U(x)} \mathbf{1}_A(p_y(s)) q_y(s) \bar{\delta}_y(p_y(s)) dm_{W(x)}(y) \\ &= \int_{W_U(s)} \mathbf{1}_A(r) q_y(s) \bar{\delta}_y(r) \bar{q}_s^{-1}(r) dm_{W(s)}(r). \end{aligned}$$

The last formula together with (6.3) gives the absolute continuity of W . \square

The converse of Proposition 6.2.2 is not true in general (Exercise 6.2.2).

LEMMA 6.2.3. *Let (X, \mathfrak{A}, μ) , (Y, \mathfrak{B}, ν) be two compact metric spaces with Borel σ -algebras and σ -additive Borel measures, and let $p_n: X \rightarrow Y$, $n = 1, 2, \dots$, and $p: X \rightarrow Y$ be continuous maps such that*

1. *each p_n and p are homeomorphisms onto their images,*
2. *p_n converges to p uniformly as $n \rightarrow \infty$,*
3. *there is a constant J such that $\nu(p_n(A)) \leq J\mu(A)$ for every $A \in \mathfrak{A}$.*

Then $\nu(p(A)) \leq J\mu(A)$ for every $A \in \mathfrak{A}$.

Proof. It is sufficient to prove the statement for an arbitrary open ball $B_r(x)$ in X . If $\delta < r$ then $p(B_{r-\delta}(x)) \subset p_n(B_r(x))$ for n large enough, and hence $\nu(p(B_{r-\delta}(x))) \leq \nu(p_n(B_r(x))) \leq J\mu(B_r(x))$. Observe now that $\nu(p(B_{r-\delta}(x))) \nearrow \nu(p(B_r(x)))$ as $\delta \searrow 0$. \square

For subspaces $A, B \subset \mathbb{R}^N$, set

$$\Theta(A, B) = \min\{\|v - w\| : v \in A, \|v\| = 1; w \in B, \|w\| = 1\}.$$

For $\theta \in [0, \sqrt{2}]$, we say that a subspace $A \subset \mathbb{R}^N$ is θ -transverse to a subspace $B \subset \mathbb{R}^N$ if $\Theta(A, B) \geq \theta$.

LEMMA 6.2.4. *Let \hat{E} be a smooth k -dimensional distribution on a compact subset of \mathbb{R}^N . Then for every $\xi > 0$ and $\epsilon > 0$ there is $\delta > 0$ with the following property. Suppose $Q_1, Q_2 \subset \mathbb{R}^N$ are $(N - k)$ -dimensional C^1 submanifolds with a smooth holonomy map $\hat{p}: Q_1 \rightarrow Q_2$ such that $\hat{p}(x) \in Q_2$, $\hat{p}(x) - x \in \hat{E}(x)$, $\Theta(T_x Q_1, \hat{E}(x)) \geq \xi$, $\Theta(T_{\hat{p}(x)} Q_2, \hat{E}(x)) \geq \xi$, $\text{dist}(T_x Q_1, T_{\hat{p}(x)} Q_2) \leq \delta$, and $\|\hat{p}(x) - x\| \leq \delta$ for each $x \in Q_1$. Then the Jacobian of \hat{p} does not exceed $1 + \epsilon$.*

Proof. Since only the first derivatives of Q_1 and Q_2 affect the Jacobian of \hat{p} at $x \in Q_1$, it equals the Jacobian at x of the holonomy map $\tilde{p}: T_x Q_1 \rightarrow T_{\hat{p}(x)} Q_2$ along \hat{E} . By applying an appropriate linear transformation L (whose determinant depends only on ξ), switching to new coordinates (u, v) in \mathbb{R}^N , and using the same notation for the images of all objects under L , we may

assume that (a) $x = (0, 0)$, (b) $T_{(0,0)}Q_1 = \{v = 0\}$, (c) $p(x) = (0, v_0)$, where $\|v_0\| = \|\hat{p}(x) - x\|$, (d) $T_{(0,v_0)}Q_2$ is given by the equation $v = v_0 + Bu$, where B is a $k \times (N - k)$ matrix whose norm depends only on δ , and (e) $\hat{E}(0, 0) = \{u = 0\}$, and $\hat{E}(w, 0)$ is given by the equation $u = w + A(w)v$, where $A(w)$ is an $(N - k) \times k$ matrix which is C^1 in w and $A(0) = 0$.

The image of $(w, 0)$ under \hat{p} is the intersection point of the planes $v = v_0 + Bu$ and $u = w + A(w)v$. Since the norm of B is bounded from above in terms of ξ , it suffices to estimate the determinant of the derivative $\partial u / \partial w$ at $w = 0$. We substitute the first equation into the second one,

$$u = w + A(w)v_0 + A(w)Bu;$$

differentiate with respect to w ,

$$\frac{\partial u}{\partial w} = I + \frac{\partial A(w)}{\partial w}v_0 + \frac{\partial A(w)}{\partial w}Bu + A(w)B\frac{\partial u}{\partial w};$$

and obtain for $w = 0$ (using $u(0) = 0$ and $A(0) = 0$)

$$\left. \frac{\partial u}{\partial w} \right|_{w=0} = I + \left. \frac{\partial A(w)}{\partial w} \right|_{w=0} v_0. \quad \square$$

THEOREM 6.2.5. *The stable and unstable foliations of a C^2 Anosov diffeomorphism are transversely absolutely continuous.*

Proof. Let $f: M \rightarrow M$ be a C^2 Anosov diffeomorphism with stable and unstable distributions E^s and E^u , and hyperbolicity constants C and $0 < \lambda < 1 < \mu$. We will prove the absolute continuity of the stable foliation W^s . Absolute continuity of the unstable foliation W^u follows by reversing the time. To prove the theorem, we are going to uniformly approximate the holonomy map by continuous maps with uniformly bounded Jacobians.

As in the proof of Theorem 6.1.3, we assume that M is a compact submanifold in \mathbb{R}^N [Hir94] and denote by $T_x M^\perp$ the orthogonal complement of $T_x M$ in \mathbb{R}^N . Let \hat{E}^s be a smooth distribution that approximates the continuous distribution $\tilde{E}^s(x) = E^s(x) \oplus T_x M^\perp$. \square

LEMMA 6.2.6. *For every $\theta > 0$ there is a constant $C_1 > 0$ such that for every $x \in M$, for every subspace $H \subset T_x M$ of the same dimension as $E^u(x)$ and θ -transverse to $E^s(x)$, and for every $k \in \mathbb{N}$,*

1. $\|df_x^k v\| \geq C_1 \mu^k \|v\|$ for every $v \in H$,
2. $\text{dist}(df_x^k H, df_x^k E^u(x)) \leq C_1 (\frac{\lambda}{\mu})^k \text{dist}(H, E^u(x))$.

Proof. Exercise 6.2.3. \square

By compactness of M , there is $\theta_0 > 0$ such that $\Theta(E^s(x), E^u(x)) \geq \theta_0$ for every $x \in M$. Also by compactness, there is a covering of M by finitely




Figure 6.3. Construction of approximating maps p_n .

many foliation coordinate charts (U_i, h_i) , $i = 1, \dots, l$, of the stable foliation W^s . It follows that there are positive constants ϵ and δ such that every $y \in M$ is contained in a coordinate chart U_j with the following property: If L is a compact connected submanifold of U_j such that

1. L intersects transversely every local stable leaf of U_j ,
2. $\Theta(T_z L, E^s) > \theta_0/3$ for all $z \in L$, and
3. $\text{dist}(y, L) < \delta$,

then for any subspace $E \subset \mathbb{R}^n$ with $\text{dist}(E, E^s(y) \oplus T_y M^\perp) < \epsilon$, the affine plane $y + E$ intersects L transversely in a unique point z_y , and $\|y - z_y\| < 6\delta/\theta_0$.

Let (U, h) be a foliation coordinate chart, and L_1, L_2 local transversals in U with holonomy map $p: L_1 \rightarrow L_2$. Define a map $\hat{p}: f^n(L_1) \rightarrow f^n(L_2)$ as follows: For $x \in L_1$, let $\hat{p}(f^n(x))$ be the unique intersection point of the affine plane $f^n(x) + \hat{E}(f^n(x))$ with $f^n(L_2)$ that is closest to $f^n(p(x))$ along $f^n(L_2)$ (note that there may be several such intersection points). The map \hat{p} is well defined by Lemma 6.2.6 and the remarks in the preceding paragraph.

For $x \in L_1$, set $p_n(x) = f^{-n}(\hat{p}(f^n(x)))$. Let $x_1 \in L_1, x_2 = p_n(x_1)$ and set $y_1 = f^n(x_1)$; see Figure 6.3. Observe that

$$\text{dist}(f^k(x_1), f^k(p(x_1))) \leq C\lambda^k \text{dist}(x_1, p(x_1)) \quad \text{for } k = 0, 1, 2, \dots \quad (6.4)$$

Assuming that \hat{E}^s is C^0 -close enough to \tilde{E}^s , it is, by Lemma 6.2.6, uniformly transverse to $f^n(L_1)$ and $f^n(L_2)$. Therefore, there is $C_2 > 0$ such that

$$\begin{aligned} \text{dist}(\hat{p}(f^n(x_1)), f^n(p(x_1))) &\leq C_2 \text{dist}(f^n(x_1), f^n(p(x_1))) \\ &\leq C_2 C \lambda^n \text{dist}(x_1, p(x_1)). \end{aligned}$$

Therefore, by (6.4) and Lemma 6.2.6,

$$\text{dist}(p_n(x_1), p(x_1)) \leq \frac{C_2 C}{C_1} \left(\frac{\lambda}{\mu}\right)^n \text{dist}(x_1, p(x_1)), \quad (6.5)$$

and hence p_n converges uniformly to p as $n \rightarrow \infty$.

Combining (6.4) and (6.5), we get

$$\begin{aligned} \text{dist}(f^k(x_1), f^k(x_2)) &\leq \text{dist}(f^k(x_1), f^k(p(x_1))) + \text{dist}(f^k(p(x_1)), f^k(x_2)) \\ &\leq C_3 \lambda^k. \end{aligned} \quad (6.6)$$

Let $J(f^k(x_i))$ be the Jacobian of f in the direction of the tangent plane $T_i^k(x_i) = T_{f^k(x_i)} L_i$, $i = 1, 2$, $k = 0, 1, 2, \dots$. Also, denote by Jac_{p_n} the Jacobian of p_n , and by $\text{Jac}_{\hat{p}}$ the Jacobian of $\hat{p}: f^n(L_1) \rightarrow f^n(L_2)$, which is uniformly bounded by Lemma 6.2.4. Then

$$\text{Jac}_{p_n}(x_1) = \prod_{k=0}^{n-1} (J(f^k(x_2)))^{-1} \cdot \text{Jac}_{\hat{p}}(f^n(x_1)) \cdot \prod_{k=0}^{n-1} J(f^k(x_1)).$$

To obtain a uniform bound on Jac_{p_n} we need to estimate the quantity $P = \prod_{k=0}^{n-1} (J(f^k(x_1))/J(f^k(x_2)))$ from above. By Theorem 6.1.3, Lemma 6.2.6, and (6.6), for some $C_4, C_5, C_6 > 0$ and $\bar{\alpha}$,

$$\begin{aligned} \text{dist}(T_1^k(x_1), T_2^k(x_2)) &\leq \text{dist}(T_1^k(x_1), \tilde{E}^u(f^k(x_1))) \\ &\quad + \text{dist}(\tilde{E}^u(f^k(x_1)), \tilde{E}^u(f^k(x_2))) \\ &\quad + \text{dist}(T_2^k(x_2), \tilde{E}^u(f^k(x_2))) \\ &\leq 2C_1 \left(\frac{\lambda}{\mu}\right)^k + C_4 (\text{dist}(f^k(x_1), f^k(x_2)))^\alpha \\ &\leq 2C_1 \left(\frac{\lambda}{\mu}\right)^k + C_5 \lambda^{\alpha k} \leq C_6 \lambda^{\alpha k}. \end{aligned} \quad (6.7)$$

Since f is a C^2 diffeomorphism, its derivative is Lipschitz continuous, and the Jacobians $J(f^k(x_1))$ and $J(f^k(x_2))$ are bounded away from 0 and ∞ . Therefore it follows from (6.7) that $|J(f^k(x_1)) - J(f^k(x_2))|/|J(f^k(x_2))| < C_7 \lambda^{\alpha k}$. Hence the product P converges and is bounded. \square

Exercise 6.2.1. Let W be a k -dimensional foliation of M , and let L be an $(n - k)$ -dimensional local transversal to W at $x \in M$, i.e., $T_x M = T_x W(x) \oplus T_x L$. Prove that there is a neighborhood $U \ni x$ and a C^1 coordinate chart $w: B^k \times B^{n-k} \rightarrow U$ such that the connected component of $L \cap U$ containing x is $w(0, B^{n-k})$ and there are C^1 functions $f_y: B^k \rightarrow B^{n-k}$, $y \in B^{n-k}$, with the following properties:

- (i) f_y depends continuously on y in the C^1 -topology;
- (ii) $w(\text{graph}(f_y)) = W_U(w(0, y))$.

Exercise 6.2.2. Give an example of an absolutely continuous foliation, which is not transversely absolutely continuous.

Exercise 6.2.3. Prove Lemma 6.2.6.

Exercise 6.2.4. Let $W_i, i = 1, 2$, be two transverse foliations of dimensions k_i on a smooth manifold M , i.e., $T_x W_1(x) \cap T_x W_2(x) = \{0\}$ for each $x \in M$. The foliations W_1 and W_2 are called *integrable* if there is a $(k_1 + k_2)$ -dimensional foliation W (called the *integral hull* of W_1 and W_2) such that $W(x) = \bigcup_{y \in W_1(x)} W_2(y) = \bigcup_{y \in W_2(x)} W_1(y)$ for every $x \in M$.

Let W_1 be a C^1 foliation and W_2 be an absolutely continuous foliation, and assume that W_1 and W_2 are integrable with integral hull W . Prove that W is absolutely continuous.

6.3 Proof of Ergodicity

The proof of Theorem 6.3.1 below follows the main ideas of E. Hopf's argument for the ergodicity of the geodesic flow on a compact surface of variable negative curvature.

We say that a measure μ on a differentiable Riemannian manifold M is *smooth* if it has a continuous density q with respect to the Riemannian volume m , i.e., $\mu(A) = \int_A q(x) dm(x)$ for each bounded Borel set $A \subset M$.

THEOREM 6.3.1. *A C^2 Anosov diffeomorphism preserving a smooth measure is ergodic.*

Proof. Let (X, \mathfrak{A}, μ) be a finite measure space such that X is a compact metric space with distance d , μ is a Borel measure, and \mathfrak{A} is the μ -completion of the Borel σ -algebra. Let $f: X \rightarrow X$ be a homeomorphism. For $x \in X$, define the *stable* set $V^s(x)$ and *unstable* set $V^u(x)$ by the formulas

$$V^s(x) = \{y \in X: d(f^n(x), f^n(y)) \rightarrow 0 \text{ as } n \rightarrow \infty\},$$

$$V^u(x) = \{y \in X: d(f^n(x), f^n(y)) \rightarrow 0 \text{ as } n \rightarrow -\infty\}.$$

LEMMA 6.3.2. *Let $\phi: X \rightarrow \mathbb{R}$ be an f -invariant measurable function. Then ϕ is constant mod 0 on stable and unstable sets, i.e., there is a null set N such that ϕ is constant on $V^s(x) \setminus N$ and on $V^u(x) \setminus N$ for every $x \in X \setminus N$.*

Proof. We will only deal with the stable sets. Without loss of generality assume that ϕ is non-negative. For a real C set $\phi_C(x) = \min(\phi(x), C)$. The function ϕ_C is f -invariant, and it suffices to prove the lemma for ϕ_C with arbitrary C . For $k \in \mathbb{N}$, let $\psi_k: X \rightarrow \mathbb{R}$ be a continuous function such that $\int_X |\phi_C - \psi_k| d\mu(x) < \frac{1}{k}$. By the Birkhoff ergodic theorem, the limit

$$\psi_k^+(x) = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} \psi_k(f^i(x))$$

exists for μ -a.e. x . By the invariance of μ and ϕ_C , for every $j \in \mathbb{Z}$,

$$\begin{aligned} \frac{1}{k} > \int_X |\phi_C(x) - \psi_k(x)| d\mu(x) &= \int_X |\phi_C(f^j(y)) - \psi_k(f^j(y))| d\mu(y) \\ &= \int_X |\phi_C(y) - \psi_k(f^j(y))| d\mu(y), \end{aligned}$$

and hence

$$\begin{aligned} \int_X \left| \phi_C(y) - \frac{1}{n} \sum_{i=0}^{n-1} \psi_k(f^i(y)) \right| d\mu(y) \\ \leq \frac{1}{n} \sum_{i=0}^{n-1} \int_X |\phi_C(y) - \psi_k(f^i(y))| d\mu(y) < \frac{1}{k}. \end{aligned}$$

Since ψ_k is uniformly continuous, $\psi_k^+(y) = \psi_k^+(x)$ whenever $y \in V^s(x)$ and $\psi_k^+(x)$ is defined. Therefore, there is a null set N_k such that ψ_k^+ exists and is constant on the stable sets in $X \setminus N_k$. It follows that $\phi_C^+(x) = \lim_{k \rightarrow \infty} \psi_k^+(x)$ is constant on the stable sets in $X \setminus \bigcup N_k$. Clearly $\phi_C(x) = \phi_C^+(x) \bmod 0$. \square

Let ϕ be a μ -measurable f -invariant function. By Lemma 6.3.2, there is a μ -null set N_s such that ϕ is constant on the leaves of W^s in $M \setminus N_s$ and another μ -null set N_u such that ϕ is constant on the leaves of W^u in $M \setminus N_u$. Let $x \in M$, and let $U \ni x$ be a small neighborhood, as in the definition of absolute continuity for W^s and W^u . Let $G_s \subset U$ be the set of points $z \in U$ for which $m_{W^s(z)}(N_s \cap W^s(z)) = 0$ and $z \notin N_s$. Let $G_u \subset U$ be the set of points $z \in U$ for which $m_{W^u(z)}(N_u \cap W^u(z)) = 0$ and $z \notin N_u$. By Proposition 6.2.1 and the absolute continuity of W^u and W^s (Theorem 6.2.5), both sets G_s and G_u have full μ -measure in U , and hence so does $G_s \cap G_u$. Again, by the absolute continuity of W^u , there is a full- μ -measure subset of points $z \in U$ such that $z \in G_s \cap G_u$ and $m_{W^u(z)}$ -a.e. point from $W^u(z)$ also lies in $G_s \cap G_u$. It follows that $\phi(x) = \phi(z)$ for μ -a.e. point $x \in U$. Since M is connected, ϕ is constant mod 0 on M . \square

Exercise 6.3.1. Prove that a C^2 Anosov diffeomorphism preserving a smooth measure is weak mixing.

Low-Dimensional Dynamics

As we have seen in the previous chapters, general dynamical systems exhibit a wide variety of behaviors and cannot be completely classified by their invariants. The situation is considerably better in low-dimensional dynamics and especially in one-dimensional dynamics. The two crucial tools for studying one-dimensional dynamical systems are the intermediate value theorem (for continuous maps) and conformality (for non-singular differentiable maps). A differentiable map f is *conformal* if the derivative at each point is a non-zero scalar multiple of an orthogonal transformation, i.e., if the derivative expands or contracts distances by the same amount in all directions. In dimension one, any non-singular differentiable map is conformal. The same is true for complex analytic maps, which we study in Chapter 8. But in higher dimensions, differentiable maps are rarely conformal.

7.1 Circle Homeomorphisms

The circle $S^1 = [0, 1] \bmod 1$ can be considered as the quotient space \mathbb{R}/\mathbb{Z} . The quotient map $\pi: \mathbb{R} \rightarrow S^1$ is a *covering map*, i.e., each $x \in S^1$ has a neighborhood U_x such that $\pi^{-1}(U_x)$ is a disjoint union of connected open sets, each of which is mapped homeomorphically onto U_x by π .

Let $f: S^1 \rightarrow S^1$ be a homeomorphism. We will assume throughout this section that f is orientation-preserving (see Exercise 7.1.3 for the orientation-reversing case). Since π is a covering map, we can lift f to an increasing homeomorphism $F: \mathbb{R} \rightarrow \mathbb{R}$ such that $\pi \circ F = f \circ \pi$. For each $x_0 \in \pi^{-1}(f(0))$ there is a unique lift F such that $F(0) = x_0$, and any two lifts differ by an integer translation. For any lift F and any $n \in \mathbb{Z}$, $F(x + n) = F(x) + n$ for any $x \in \mathbb{R}$.

THEOREM 7.1.1. *Let $f: S^1 \rightarrow S^1$ be an orientation-preserving homeomorphism, and $F: \mathbb{R} \rightarrow \mathbb{R}$ a lift of f . Then for every $x \in \mathbb{R}$, the limit*

$$\rho(F) = \lim_{n \rightarrow \infty} \frac{F^n(x) - x}{n}$$

exists, and is independent of the point x . The number $\rho(f) = \pi(\rho(F))$ is independent of the lift F , and is called the rotation number of f . If f has a periodic point, then $\rho(f)$ is rational.

Proof. Suppose for the moment that the limit exists for some $x \in [0, 1)$. Since F maps any interval of length 1 to an interval of length 1, it follows that $|F^n(x) - F^n(y)| \leq 1$ for any $y \in [0, 1)$. Thus

$$|(F^n(x) - x) - (F^n(y) - y)| \leq |F^n(x) - F^n(y)| + |x - y| \leq 2,$$

so

$$\lim_{n \rightarrow \infty} \frac{F^n(x) - x}{n} = \lim_{n \rightarrow \infty} \frac{F^n(y) - y}{n}.$$

Since $F^n(y + k) = F^n(y) + k$, the same holds for any $y \in \mathbb{R}$.

Suppose $F^q(x) = x + p$ for some $x \in [0, 1)$ and some $p, q \in \mathbb{N}$. This is equivalent to asserting that $\pi(x)$ is a periodic point for f with period q . For $n \in \mathbb{N}$, write $n = kq + r$, $0 \leq r < q$. Then $F^n(x) = F^r(F^{kq}x) = F^r(x + kp) = F^r(x) + kp$, and since $|F^r(x) - x|$ is bounded for $0 \leq r < q$,

$$\lim_{n \rightarrow \infty} \frac{F^n(x) - x}{n} = \frac{p}{q}.$$

Thus the rotation number exists and is rational whenever f has a periodic point.

Suppose now that $F^q(x) \neq x + p$ for all $x \in \mathbb{R}$ and $p, q \in \mathbb{N}$. By continuity, for each pair $p, q \in \mathbb{N}$, either $F^q(x) > x + p$ for all $x \in \mathbb{R}$, or $F^q(x) < x + p$ for all $x \in \mathbb{R}$. For $n \in \mathbb{N}$, choose $p_n \in \mathbb{N}$ so that $p_n - 1 < F^n(x) - x < p_n$ for all $x \in \mathbb{R}$. Then for any $m \in \mathbb{N}$,

$$m(p_n - 1) < F^{mn}(x) - x = \sum_{k=0}^{m-1} F^n(F^{kn}(x)) - F^{kn}(x) < mp_n,$$

which implies that

$$\frac{p_n}{n} - \frac{1}{n} < \frac{F^{mn}(x) - x}{mn} < \frac{p_n}{n}.$$

Interchanging the roles of m and n , we also have

$$\frac{p_m}{m} - \frac{1}{m} < \frac{F^{mn}(x) - x}{mn} < \frac{p_m}{m}.$$

Thus, $|p_m/m - p_n/n| < |1/m + 1/n|$, so $\{p_n/n\}$ is a Cauchy sequence. It follows that $(F^n(x) - x)/n$ converges as $n \rightarrow \infty$.

If $G = F + k$ is another lift of f , then $\rho(G) = \rho(F) + k$, so $\rho(f)$ is independent of the lift F . Moreover, there is a unique lift F such that $\rho(F) = \rho(f)$ (Exercise 7.1.1). \square

Since $S^1 = [0, 1] \bmod 1$, we will often abuse notation by writing $\rho(f) = x$ for some $x \in [0, 1]$.

PROPOSITION 7.1.2. *The rotation number depends continuously on the map in the C^0 topology.*

Proof. Let f be an orientation-preserving circle homeomorphism, and choose $p, q, p', q' \in \mathbb{N}$ such that $p/q < \rho(f) < p'/q'$. Let F be the lift of f such that $p < F^q(x) - x < p + q$. Then for all $x \in \mathbb{R}$, $p < F^q(x) - x < p + q$, since otherwise we would have $\rho(f) = p/q$. If g is another circle homeomorphism close to F , then there is a lift G close to F , and for g sufficiently close to f , the same inequality $p < G^q(x) - x < p + q$ holds for all $x \in \mathbb{R}$. Thus $p/q < \rho(g)$. A similar argument involving p' and q' completes the proof. \square

PROPOSITION 7.1.3. *Rotation number is an invariant of topological conjugacy.*

Proof. Let f and h be orientation-preserving homeomorphisms of S^1 , and let F and H be lifts of f and h . Then $H \circ F \circ H^{-1}$ is a lift of $h \circ f \circ h^{-1}$, and for $x \in \mathbb{R}$,

$$\begin{aligned} \frac{(HFH^{-1})^n(x) - x}{n} &= \frac{(HF^nH^{-1})(x) - x}{n} \\ &= \frac{H(F^nH^{-1}(x)) - F^nH^{-1}(x)}{n} + \frac{F^nH^{-1}(x) - H^{-1}(x)}{n} + \frac{H^{-1}(x) - x}{n}. \end{aligned}$$

Since the numerators in the first and third terms of the last expression are bounded independent of n , we conclude that

$$\rho(hfh^{-1}) = \lim_{n \rightarrow \infty} \frac{(HFH^{-1})^n(x) - x}{n} = \lim_{n \rightarrow \infty} \frac{F^n(x) - x}{n} = \rho(f). \quad \square$$

PROPOSITION 7.1.4. *If $f: S^1 \rightarrow S^1$ is a homeomorphism, then $\rho(f)$ is rational if and only if f has a periodic point. Moreover, if $\rho(f) = p/q$ where p and q are relatively prime non-negative integers, then every periodic point of f has minimal period q , and if $x \in \mathbb{R}$ projects to a periodic point of f , then $F^q(x) = x + p$ for the unique lift F with $\rho(F) = p/q$.*

Proof. The “if” part of the first assertion is contained in Theorem 7.1.1.

Suppose $\rho(f) = p/q$, where $p, q \in \mathbb{N}$. If F and $\tilde{F} = F + l$ are two lifts of f , then $\tilde{F}^q = F^q + lq$. Thus we may choose F to be the unique lift with $p \leq F^q(0) < p + q$. To show the existence of a periodic point of f , it suffices to show the existence of a point $x \in [0, 1]$ such that $F^q(x) = x + k$ for some $k \in \mathbb{N}$. We may assume that $x + p < F^q(x) < x + p + q$ for all $x \in [0, 1]$, since otherwise we have $F^q(x) = x + l$ for $k = p$ or $k = p + q$, and we are done. Choose $\epsilon > 0$ such that for any $x \in [0, 1]$, $x + p + \epsilon < F^q(x) < x + p + q - \epsilon$. The same inequality then holds for all $x \in \mathbb{R}$, since $F^q(x + k) = F^q(x) + k$ for all $k \in \mathbb{N}$. Thus

$$\frac{p + \epsilon}{q} = \frac{k(p + \epsilon)}{kq} < \frac{F^{kq}(x) - x}{kq} < \frac{k(p + q - \epsilon)}{kq} = \frac{p + 1 - \epsilon}{q}$$

for all $k \in \mathbb{N}$, contradicting $\rho(f) = p/q$. We conclude that $F^q(x) = x + p$ or $F^q(x) = x + p + q$ for some x , and x is periodic with period q .

Now assume $\rho(f) = p/q$, with p and q relatively prime, and suppose $x \in [0, 1]$ is a periodic point of f . Then there are integers $p', q' \in \mathbb{N}$ such that $F^{q'}(x) = x + p'$. By the proof of Theorem 7.1.1, $\rho(f) = p'/q'$, so if d is the greatest common divisor of p' and q' , then $q' = qd$ and $p' = pd$. We claim that $F^q(x) = x + p$. If not, then either $F^q(x) > x + p$ or $F^q(x) < x + p$. Suppose the former holds (the other case is similar). Then by monotonicity,

$$F^{dq}(x) > F^{(d-1)q}(x) + p > \dots > x + dp,$$

contradicting the fact that $F^{q'}(x) = x + p'$. Thus, x is periodic with period q . \square

Suppose f is a homeomorphism of S^1 . Given any subset $A \subset S^1$ and a distinguished point $x \in A$, we define an ordering on A by lifting A to the interval $[\tilde{x}, \tilde{x} + 1) \subset \mathbb{R}$, where $\tilde{x} \in \pi^{-1}(x)$, and using the natural ordering on \mathbb{R} . In particular, if $x \in S^1$, then the orbit $\{x, f(x), f^2(x), \dots\}$ has a natural order (using x as the distinguished point).

THEOREM 7.1.5. *Let $f: S^1 \rightarrow S^1$ be an orientation-preserving homeomorphism with rational rotation number $\rho = p/q$, where p and q are relatively prime. Then for any periodic point $x \in S^1$, the ordering of the orbit $\{x, f(x), f^2(x), \dots, f^{q-1}(x)\}$ is the same as the ordering of the set $\{0, p/q, 2p/q, \dots, (q-1)p/q\}$, which is the orbit of 0 under the rotation R_p .*

Proof. Let x be a periodic point of f , and let $i \in \{0, \dots, q-1\}$ be the unique number such that $f^i(x)$ is the first point to the right of x in the orbit of x . Then $f^{2i}(x)$ must be the first point to the right of $f^i(x)$, since if $f^l(x) \in (f^i(x), f^{2i}(x))$ then $l > i$ and $f^{l-i}(x) \in (x, f^i(x))$, contradicting the choice of i . Thus the points of the orbit are ordered as $x, f^i(x), f^{2i}(x), \dots, f^{(q-1)i}(x)$.

Let \tilde{x} be a lift of x . Since f^i carries each interval $[f^{ki}(x), f^{(k+1)i}(x)]$ to its successor, and there are q of these intervals, there is a lift \bar{F} of f^i such that $\bar{F}^q \tilde{x} = \tilde{x} + 1$. Let F be the lift of f with $F^q(x) = x + p$. Then F^i is a lift of f^i , so $F^i = \bar{F} + k$ for some k . We have

$$x + ip = F^{qi}(x) = (\bar{F} + k)^q(x) = \bar{F}^q(x) + qk = x + 1 + qk.$$

Thus $ip = 1 + qk$, so i is the unique number between 0 and q such that $ip \equiv 1 \pmod{q}$. Since the points of the set $\{0, p/q, 2p/q, \dots, (q-1)p/q\}$ are ordered as $0, ip/q, \dots, (q-1)ip/q$, the theorem follows. \square

Now we turn to the study of orientation-preserving homeomorphisms with irrational rotation number. If x and y are two points in S^1 , then we define the interval $[x, y] \subset S^1$ to be $\pi([\tilde{x}, \tilde{y}])$, where $\tilde{x} \in \pi^{-1}(x)$ and $\tilde{y} = \pi^{-1}(y) \cap [\tilde{x}, \tilde{x} + 1)$. Open and half-open intervals are defined in a similar way.

LEMMA 7.1.6. *Suppose $\rho(f)$ is irrational. Then for any $x \in S^1$ and any distinct integers $m > n$, every forward orbit of f intersects the interval $I = [f^m(x), f^n(x)]$.*

Proof. It suffices to show that $S^1 = \bigcup_{k=0}^{\infty} f^{-k}I$. Suppose not. Then

$$S^1 \not\subset \bigcup_{k=1}^{\infty} f^{-k(m-n)}I = \bigcup_{k=1}^{\infty} [f^{-(k-1)m+kn}(x), f^{-km+(k+1)n}(x)].$$

Since the intervals $f^{-k(m-n)}I$ abut at the endpoints, we conclude that $f^{-k(m-n)}f^n(x)$ converges monotonically to a point $z \in S^1$, which is a fixed point for f^{m-n} , contradicting the irrationality of $\rho(f)$. \square

PROPOSITION 7.1.7. *If $\rho(f)$ is irrational, then $\omega(x) = \omega(y)$ for any $x, y \in S^1$, and either $\omega(x) = S^1$ or $\omega(x)$ is perfect and nowhere dense.*

Proof. Fix $x, y \in S^1$. Suppose $f^{a_n}(x) \rightarrow x_0 \in \omega(x)$ for some sequence $a_n \nearrow \infty$. By Lemma 7.1.6, for each $n \in \mathbb{N}$, we can choose b_n such that $f^{b_n}(y) \in [f^{a_{n-1}}(x), f^{a_n}(x)]$. Then $f^{b_n}(y) \rightarrow x_0$, so $\omega(x) \subset \omega(y)$. By symmetry, $\omega(x) = \omega(y)$.

To show that $\omega(x)$ is perfect, we fix $z \in \omega(x)$. Since $\omega(x)$ is invariant, we have that $z \in \omega(z)$ is a limit point of $\{f^n(z)\} \subset \omega(x)$, so $\omega(x)$ is perfect.

To prove the last claim, we suppose that $\omega(x) \neq S^1$. Then $\partial\omega(x)$ is a non-empty closed invariant set. If $z \in \partial\omega(x)$, then $\omega(z) = \omega(x)$. Therefore, $\omega(x) \subset \partial\omega(x)$ and $\omega(x)$ is nowhere dense. \square

LEMMA 7.1.8. *Suppose $\rho(f)$ is irrational. Let F be a lift of f , and $\rho = \rho(F)$. Then for any $x \in \mathbb{R}$, $n_1\rho + m_1 < n_2\rho + m_2$ if and only if $F^{n_1}(x) + m_1 < F^{n_2}(x) + m_2$, for any $m_1, m_2, n_1, n_2 \in \mathbb{Z}$.*

Proof. Suppose $F^{n_1}(x) + m_1 < F^{n_2}(x) + m_2$ or, equivalently,

$$F^{(n_1-n_2)}(x) < x + m_2 - m_1.$$

This inequality holds for all x , since otherwise the rotation number would be rational. In particular, for $x = 0$ we have $F^{(n_1-n_2)}(0) < m_2 - m_1$. By an inductive argument, $F^{k(n_1-n_2)}(0) < k(m_2 - m_1)$. If $n_1 - n_2 > 0$, it follows that

$$\frac{F^{k(n_1-n_2)}(0) - 0}{k(n_1 - n_2)} < \frac{m_2 - m_1}{n_1 - n_2},$$

so $\rho = \lim_{k \rightarrow \infty} F^{k(n_1-n_2)}(0)/k(n_1 - n_2) \leq (m_2 - m_1)/(n_1 - n_2)$. Irrationality of ρ implies strict inequality, so $n_1\rho + m_1 < n_2\rho + m_2$. The same result holds in the case $n_1 - n_2 < 0$ by a similar argument. The converse follows by reversing the inequality. \square

THEOREM 7.1.9 (Poincaré Classification). *Let $f: S^1 \rightarrow S^1$ be an orientation-preserving homeomorphism with irrational rotation number ρ .*

1. *If f is topologically transitive, then f is topologically conjugate to the rotation R_ρ .*
2. *If f is not topologically transitive, then R_ρ is a factor of f , and the factor map $h: S^1 \rightarrow S^1$ can be chosen to be monotone.*

Proof. Let F be a lift of f , and fix $x \in \mathbb{R}$. Let $A = \{F^n(x) + m: n, m \in \mathbb{Z}\}$ and $B = \{n\rho + m: n, m \in \mathbb{Z}\}$. Then B is dense in \mathbb{R} (§1.2). Define $H: A \rightarrow B$ by $H(F^n(x) + m) = n\rho + m$. By the preceding lemma, H preserves order and is bijective. Extend H to a map $H: \mathbb{R} \rightarrow \mathbb{R}$ by defining

$$H(y) = \sup\{n\rho + m: F^n(x) + m < y\}.$$

Then $H(y) = \inf\{n\rho + m: F^n(x) + m > y\}$, since otherwise $\mathbb{R} \setminus B$ would contain an interval.

We claim that $H: \mathbb{R} \rightarrow \mathbb{R}$ is continuous. If $y \in \bar{A}$, then $H(y) = \sup\{H(z): z \in A, z < y\}$ and $H(y) = \inf\{H(z): z \in A, z > y\}$ implies that H is continuous

on \bar{A} . If I is an interval in $\mathbb{R} \setminus \bar{A}$, then H is constant on I and the constant agrees with the values at the endpoints. Thus $H: \mathbb{R} \rightarrow \mathbb{R}$ is a continuous extension of $H: A \rightarrow B$.

Note that H is surjective, non-decreasing, and that

$$\begin{aligned} H(y+1) &= \sup\{n\rho + m: F^n(x) + m < y+1\} \\ &= \sup\{n\rho + m: F^n(x) + (m-1) < y\} = H(y) + 1. \end{aligned}$$

Moreover,

$$\begin{aligned} H(F(y)) &= \sup\{n\rho + m: F^n(x) + m < F(y)\} \\ &= \sup\{n\rho + m: F^{n-1}(x) + m < y\} \\ &= \rho + H(y). \end{aligned}$$

We conclude that H descends to a map $h: S^1 \rightarrow S^1$ and $h \circ f = R_\rho \circ h$.

Finally, note that f is transitive if and only if $\{F^n(x) + m: n, m \in \mathbb{Z}\}$ is dense in \mathbb{R} . Since H is constant on any interval in $\mathbb{R} \setminus \bar{A}$, we conclude that h is injective if and only if f is transitive. (Note that by Proposition 7.1.7, either every orbit is dense or no orbit is dense.) \square

Exercise 7.1.1. Show that if F and $G = F + k$ are two lifts of f , then $\rho(F) = \rho(G) + k$, so $\rho(f)$ is independent of the choice of lift used in its definition. Show that there is a unique lift F of f such that $\rho(F) = \rho(f)$.

Exercise 7.1.2. Show that $\rho(f^m) = m\rho(f)$.

Exercise 7.1.3. Show that if f is an orientation-reversing homeomorphism of S^1 , then $\rho(f^2) = 0$.

Exercise 7.1.4. Suppose f has rational rotation number. Show that:

- if f has exactly one periodic orbit, then every non-periodic point is both forward and backward asymptotic to the periodic orbit; and
- if f has more than one periodic orbit, then every non-periodic orbit is forward asymptotic to some periodic orbit and backward asymptotic to a different periodic orbit.

Exercise 7.1.5. Show that Theorems 7.1.1 and 7.1.5 hold under the weaker hypothesis that $f: S^1 \rightarrow S^1$ is a continuous map such that any (and thus every) lift F of f is non-decreasing.

7.2 Circle Diffeomorphisms

The total variation of a function $f: S^1 \rightarrow \mathbb{R}$ is

$$\text{Var}(f) = \sup \sum_{k=1}^n |f(x_k) - f(x_{k+1})|,$$

where the supremum is taken over all partitions $0 \leq x_1 < \dots < x_n \leq 1$, for all $n \in \mathbb{N}$. We say that g has *bounded variation* if $\text{Var}(g)$ is finite. Note that any Lipschitz function has bounded variation. In particular, any C^1 function has bounded variation.

THEOREM 7.2.1 (Denjoy). *Let f be an orientation-preserving C^1 diffeomorphism of the circle with irrational rotation number $\rho = \rho(f)$. If f' has bounded variation, then f is topologically conjugate to the rigid rotation R_ρ .*

Proof. We know from Theorem 7.1.9 that if f is transitive, it is conjugate to R_ρ . Thus we assume that f is not transitive, and argue to obtain a contradiction. By Proposition 7.1.7, we may assume that $\omega(0)$ is a perfect, nowhere dense set. Then $S^1 \setminus \omega(0)$ is a disjoint union of open intervals. Let $I = (a, b)$ be one of these intervals. Then the intervals $\{f^n(I)\}_{n \in \mathbb{Z}}$ are pairwise disjoint, since otherwise f would have a periodic point. Thus $\sum_{n \in \mathbb{Z}} l(f^n(I)) \leq 1$, where $l(f^n(I)) = \int_a^b (f^n)'(t) dt$ is the length of $f^n(I)$.

LEMMA 7.2.2. *Let J be an interval in S^1 , and suppose the interiors of the intervals $J, f(J), \dots, f^{n-1}(J)$ are pairwise disjoint. Let $g = \log f'$, and fix $x, y \in J$. Then for any $n \in \mathbb{Z}$,*

$$\text{Var}(g) \geq |\log(f^n)'(x) - \log(f^n)'(y)|.$$

Proof. Using the fact that the intervals $J, f(J), \dots, f^n(J)$ are disjoint, we get

$$\begin{aligned} \text{Var}(g) &\geq \sum_{k=0}^{n-1} |g(f^k(y)) - g(f^k(x))| \geq \left| \sum_{k=0}^{n-1} g(f^k(y)) - g(f^k(x)) \right| \\ &= \left| \log \prod_{k=0}^{n-1} f'(f^k(y)) - \log \prod_{k=0}^{n-1} f'(f^k(x)) \right| \\ &= |\log(f^n)'(y) - \log(f^n)'(x)|. \quad \square \end{aligned}$$

Fix $x \in S^1$. We claim that there are infinitely many $n \in \mathbb{N}$ such that the intervals $(x, f^{-n}(x)), (f(x), f^{1-n}(x)), \dots, (f^n(x), x)$ are pairwise disjoint. It suffices to show that there are infinitely many n such that $f^k(x)$ is not in the interval $(x, f^n(x))$ for $0 \leq |k| \leq n$. Lemma 7.1.8 implies that the orbit of x is

ordered in the same way as the orbit of a point under the irrational rotation R_ρ . Since the orbit of a point under an irrational rotation is dense, the claim follows.

Choose n as in the preceding paragraph. Then by applying Lemma 7.2.2 with $y = f^{-n}(x)$, we obtain

$$\text{Var}(g) \geq \left| \log \frac{(f^n)'(x)}{(f^n)'(y)} \right| = |\log((f^n)'(x)(f^{-n})'(x))|.$$

Thus for infinitely many $n \in \mathbb{N}$, we have

$$\begin{aligned} l(f^n(I)) + l(f^{-n}(I)) &= \int_I (f^n)'(x) dx + \int_I (f^{-n})'(x) dx \\ &= \int_I [(f^n)'(x) + (f^{-n})'(x)] dx \\ &\geq \int_I \sqrt{(f^n)'(x)(f^{-n})'(x)} dx \\ &\geq \int_I \sqrt{\exp(-\text{Var}(g))} dx = \exp\left(-\frac{1}{2}\text{Var}(g)\right) l(I). \end{aligned}$$

This contradicts the fact that $\sum_{n \in \mathbb{Z}} l(f^n(I)) \leq \infty$, so we conclude that f is transitive, and therefore conjugate to R_ρ . \square

THEOREM 7.2.3 (Denjoy Example). *For any irrational number $\rho \in (0, 1)$, there is a non-transitive C^1 orientation-preserving diffeomorphism $f: S^1 \rightarrow S^1$ with rotation number ρ .*

Proof. We know from Lemma 7.1.8 that if $\rho(f) = \rho$, then for any $x \in S^1$, the orbit of x is ordered the same way as any orbit of R_ρ , i.e., $f^k(x) < f^l(x) < f^m(x)$ if and only if $R_\rho^k(x) < R_\rho^l(x) < R_\rho^m(x)$. Thus in constructing f , we have no choice about the order of the orbit of any point. We do, however, have a choice about the spacing between points in the orbit.

Let $\{l_n\}_{n \in \mathbb{Z}}$ be a sequence of positive real numbers such that $\sum_{n \in \mathbb{Z}} l_n = 1$ and l_n is decreasing as $n \rightarrow \pm\infty$ (we will impose additional constraints later). Fix $x_0 \in S^1$, and define

$$a_n = \sum_{\{k \in \mathbb{Z}: R_\rho^k(x_0) \in [x_0, R_\rho^n(x_0)]\}} l_k, \quad b_n = a_n + l_n.$$

The intervals $[a_n, b_n]$ are pairwise disjoint. Since $\sum_{n \in \mathbb{Z}} l_n = 1$, the union of these intervals covers a set of measure 1 in $[0, 1]$, and is therefore dense.

To define a C^1 homeomorphism $f: S^1 \rightarrow S^1$ it suffices to define a continuous, positive function g on S^1 with total integral 1. Then f will be defined

to be the integral of g . The function g should satisfy:

1. $\int_{a_n}^{b_n} g(t) dt = l_{n+1}$.

To construct such a g it suffices to define g on each interval $[a_n, b_n]$ so that it also satisfies:

2. $g(a_n) = g(b_n) = 1$.

3. For any sequence $\{x_k\} \subset \bigcup_{n \in \mathbb{Z}} [a_n, b_n]$, if $y = \lim x_k$, then $g(x_k) \rightarrow 1$.

We then define g to be 1 on $S^1 \setminus \bigcup_{n \in \mathbb{Z}} [a_n, b_n]$.

There are many such possibilities for $g|_{[a_n, b_n]}$. We use the quadratic polynomial

$$g(x) = 1 + \frac{6(l_{n+1} - l_n)}{l_n^3} (b_n - x)(x - a_n),$$

which clearly satisfies condition 1. For $n \geq 0$, we have $l_{n+1} - l_n < 0$, so

$$1 \geq g(x) \geq 1 - \frac{6(l_n - l_{n+1})}{l_n^3} \left(\frac{l_n}{2}\right)^2 = \frac{3l_{n+1} - l_n}{2l_n}.$$

For $n < 0$, we have $l_{n+1} - l_n > 0$, so

$$1 \leq g(x) \leq \frac{3l_{n+1} - l_n}{2l_n}.$$

Thus if we choose l_n such that $(3l_{n+1} - l_n)/2l_n \rightarrow 1$ as $n \rightarrow \pm\infty$, then condition 3 is satisfied. For example, we could choose $l_n = \alpha(|n| + 2)^{-1}(|n| + 3)^{-1}$, where $\alpha = 1/\sum_{n \in \mathbb{Z}} ((|n| + 2)^{-1}(|n| + 3)^{-1})$.

Now define $f(x) = a_1 + \int_0^x g(t) dt$. Using the results above, it follows that $f: S^1 \rightarrow S^1$ is a C^1 homeomorphism of S^1 with rotation number ρ (Exercise 7.2.1). Moreover, $f^n(0) = a_n$, and $\omega(0) = S^1 \setminus \bigcup_{n \in \mathbb{Z}} (a_n, b_n)$ is a closed, perfect, invariant set of measure zero. \square

Exercise 7.2.1. Verify the statements in the last paragraph of the proof of Theorem 7.2.3.

Exercise 7.2.2. Show directly that the example constructed in the proof of Theorem 7.2.3 is not C^2 .

7.3 The Sharkovsky Theorem

We consider the set $\mathbb{N}_{\text{Sh}} = \mathbb{N} \cup \{2^\infty\}$ obtained by adding the formal symbol 2^∞ to the set of natural numbers. The Sharkovsky ordering of this set is

$$\begin{aligned} 1 &< 2 < \dots < 2^n < \dots < 2^\infty < \dots \\ &< 2^m \cdot (2n+1) < \dots < 2^m \cdot 7 < 2^m \cdot 5 < 2^m \cdot 3 < \dots \\ &< 2(2n+1) < \dots < 14 < 10 < 6 < \dots \\ &< 2n+1 < \dots < 7 < 5 < 3. \end{aligned}$$

The symbol 2^∞ is added so that \mathbb{N}_{Sh} has the least-upper-bound property, i.e., every subset of \mathbb{N}_{Sh} has a supremum. The Sharkovsky ordering is preserved by multiplication by 2^k , for any $k \geq 0$ (where $2^k \cdot 2^\infty = 2^\infty$, by definition).

For $\alpha \in \mathbb{N}_{\text{Sh}}$, let $S(\alpha) = \{k \in \mathbb{N} : k \leq \alpha\}$ (note that $S(\alpha)$ is defined to be a subset of \mathbb{N} , not \mathbb{N}_{Sh}). For a map $f: [0, 1] \rightarrow [0, 1]$, we denote by $\text{MinPer}(f)$ the set of minimal periods of periodic points of f .

THEOREM 7.3.1 (Sharkovsky [Sha64]). *For every continuous map $f: [0, 1] \rightarrow [0, 1]$, there is $\alpha \in \mathbb{N}_{\text{Sh}}$ such that $\text{MinPer}(f) = S(\alpha)$. Conversely, for every $\alpha \in \mathbb{N}_{\text{Sh}}$, there is a continuous map $f: [0, 1] \rightarrow [0, 1]$ with $\text{MinPer}(f) = S(\alpha)$.*

The proof of the first assertion of the Sharkovsky theorem proceeds as follows: We assume that f has a periodic point x of minimal period $n > 1$, since otherwise there is nothing to show. The orbit of x partitions the interval $[0, 1]$ into a finite collection of subintervals whose endpoints are elements of the orbit. The endpoints of these intervals are permuted by f . By examining the combinatorial possibilities for the permutations of pairs of endpoints, and using the intermediate value theorem, one establishes the existence of periodic points of the desired periods.

The second assertion of the Sharkovsky theorem is proved as Lemma 7.3.9.

If I and J are intervals in $[0, 1]$ and $f(I) \supset J$, we say that I *f-covers* J , and we write $I \rightarrow J$. If $a, b \in [0, 1]$, then we will use $[a, b]$ to represent the closed interval between a and b , regardless of whether $a \geq b$ or $a \leq b$.

LEMMA 7.3.2

1. If $f(I) \supset I$, then the closure of I contains a fixed point of f .
2. Fix $m \in \mathbb{N} \cup \{\infty\}$, and suppose that $\{I_k\}_{1 \leq k < m}$ is a finite or infinite sequence of non-empty closed intervals in $[0, 1]$ such that $f(I_k) \supset I_{k+1}$ for $1 \leq k < m - 1$. Then there is a point $x \in I_1$ such that $f^k(x) \in I_{k+1}$ for $1 \leq k < m - 1$. Moreover, if $I_n = I_1$ for some $n > 0$, then I_1 contains a periodic point x of period n such that $f^k(x) \in I_{k+1}$ for $k = 1, \dots, n - 1$.

Proof. The proof of part 1 is a simple application of the intermediate value theorem.

To prove part 2, note that since $f(I_1) \supset I_2$, there are points $a_0, b_0 \in I_1$ that map to the endpoints of I_2 . Let J_1 be the subinterval of I_1 with endpoints a_0, b_0 . Then $f(J_1) = I_2$. Suppose we have defined subintervals $J_1 \supset J_2 \supset \dots \supset J_n$ in I_1 such that $f^k(J_k) = I_{k+1}$. Then $f^{n+1}(J_n) = f(I_{n+1}) \supset I_{n+2}$, so there is an interval $J_{n+1} \subset J_n$ such that $f^{n+1}(J_{n+1}) = I_{n+2}$. Thus we obtain a nested sequence $\{J_n\}$ of non-empty closed intervals. The intersection

$\bigcap_{i=1}^{m-1} J_i$ is non-empty, and for any x in the intersection, $f^k(x) \in I_{k+1}$ for $1 \leq k < m-1$.

The last assertion follows from the preceding paragraph together with part 1. \square

A *partition* of an interval I is a (finite or infinite) collection of closed subintervals $\{I_k\}$, with pairwise disjoint interiors, whose union is I . The *Markov graph* of f associated to the partition $\{I_k\}$ is the directed graph with vertices I_k , and a directed edge from I_i to I_j if and only if I_i f -covers I_j . By Lemma 7.3.2, any loop of length n in the Markov graph of f forces the existence of a periodic point of (not necessarily minimal) period n .

As a warmup to the proof of the full Sharkovsky theorem, we prove that the existence of a periodic point of minimal period three implies the existence of periodic points of all periods. This result was rediscovered in 1975 by T. Y. Li and J. Yorke, and popularized in their paper “Period three implies chaos” [LY75].

Let x be a point of period three. Replacing x with $f(x)$ or $f^2(x)$ if necessary, we may assume that $x < f(x)$ and $x < f^2(x)$. Then there are two cases: (1) $x < f(x) < f^2(x)$ or (2) $x < f^2(x) < f(x)$. In the first case, we let $I_1 = [x, f(x)]$ and $I_2 = [f(x), f^2(x)]$. The associated Markov graph is one of the two graphs shown in Figure 7.1.

For $k \geq 2$, the path $I_1 \rightarrow I_2 \rightarrow I_2 \rightarrow \cdots \rightarrow I_2 \rightarrow I_1$ of length k implies the existence of a periodic point y of period k with the itinerary $I_1, I_2, I_2, \dots, I_2, I_1$. If the minimal period of y is less than k , then $y \in I_1 \cap I_2 = \{f(x)\}$. But $f(x)$ does not have the specified itinerary for $k \neq 3$, so the minimal period of y is k . A similar argument applies to case (2), and this proves the Sharkovsky theorem for $n = 3$.

To prove the full Sharkovsky theorem it is convenient to use a subgraph of the Markov graph defined as follows. Let $P = \{x_1, x_2, \dots, x_n\}$ be a periodic orbit of (minimal) period $n > 1$, where $x_1 < x_2 < \cdots < x_n$. Let




Figure 7.1. The two possible Markov graphs for period three.

$I_j = [x_j, x_{j+1}]$. The P -graph of f is the directed graph with vertices I_j , and a directed edge from I_j to I_k if and only if $I_k \subset [f(x_j), f(x_{j+1})]$. Since $f(I_j) \supset [f(x_j), f(x_{j+1})]$, it follows that the P -graph is a subgraph of the Markov graph associated to the same partition. In particular, any loop in the P -graph is also a loop in the Markov graph. The P -graph has the virtue that it is completely determined by the ordering of the periodic orbit, and is independent of the behavior of the map on the intervals I_j . For example, in Figure 7.1, the top graph is the unique P -graph for a periodic orbit of period three with ordering $x < f(x) < f^2(x)$.

LEMMA 7.3.3. *The P -graph of f contains a trivial loop, i.e., there is a vertex I_j with a directed edge from I_j to itself.*

Proof. Let $j = \max\{i: f(x_i) > x_i\}$. Then $f(x_j) > x_j$ and $f(x_{j+1}) \leq x_{j+1}$, so $f(x_j) \geq x_{j+1}$ and $f(x_{j+1}) \leq x_j$. Thus $[f(x_j), f(x_{j+1})) \supset [x_j, x_{j+1}]$. \square

We will renumber the vertices of the P -graph (but not the points of P) so that $I_1 = [x_j, x_{j+1}]$, where $j = \max\{i: f(x_i) > x_i\}$. By the proof of the preceding lemma, I_1 is a vertex with a directed edge from itself to itself.

For any two points $x_i < x_k$ in P , define

$$\hat{f}([x_i, x_k]) = \bigcup_{l=i}^{k-1} [f(x_l), f(x_{l+1})].$$

In particular, $\hat{f}(I_k) = [f(x_k), f(x_{k+1})]$. If $\hat{f}(I_k) \supset I_l$, we say that $I_k \hat{f}$ -covers I_l . Since we will only be using P -graphs throughout the remainder of this section, we also redefine the notation $I_k \rightarrow I_l$ to mean that $I_k \hat{f}$ -covers I_l .

PROPOSITION 7.3.4. *Any vertex of the P -graph can be reached from I_1 .*

Proof. The nested sequence $I_1 \subset \hat{f}(I_1) \subset \hat{f}^2(I_1) \subset \dots$ must eventually stabilize, since $\hat{f}^k(I_1)$ is an interval whose endpoints are in the orbit of x . Then for k sufficiently large, $\mathcal{O}(x) \cap \hat{f}^k(I_1)$ is an invariant subset of $\mathcal{O}(x)$, and is therefore equal to $\mathcal{O}(x)$. It follows that $\hat{f}^k(I_1) = [x_1, x_n]$, so any vertex of the P -graph can be reached from I_1 . \square

LEMMA 7.3.5. *Suppose the P -graph has no directed edge from any interval I_k , $k \neq 1$, to I_1 . Then n is even, and f has a periodic point of period 2.*

Proof. Let $J_0 = [x_1, x_j]$ and $J_1 = [x_{j+1}, x_{n-1}]$, where $j = \max\{i: f(x_i) > x_i\}$ (the case $j = 1$ is not excluded *a priori*). Then $\hat{f}(J_0) \not\subset J_0$ (since $f(x_j) > x_j$) and $\hat{f}(J_0) \not\subset I_1$, so $\hat{f}(J_0) \subset J_1$, since $\hat{f}(J_0)$ is connected. Likewise, $\hat{f}(J_1) \subset J_0$. Now $\hat{f}(J_0) \cup \hat{f}(J_1) \supset \mathcal{O}(x)$, so $\hat{f}(J_0) = J_1$ and $\hat{f}(J_1) = J_0$. Thus $J_0 \hat{f}$ -covers




Figure 7.2. The P -graph for Lemmas 7.3.6 and 7.3.9.

J_1 and $J_1 f$ -covers J_0 , so f has a periodic point of minimal period 2, and $n = |\mathcal{O}(x)| = 2|\mathcal{O}(x) \cap J_0|$ is even. \square

LEMMA 7.3.6. *Suppose $n > 1$ is odd and f has no non-fixed periodic points of smaller odd period. Then there is a numbering of the vertices of the P -graph so that the graph contains the following edges, and no others (see Figure 7.2):*

1. $I_1 \rightarrow I_1$ and $I_{n-1} \rightarrow I_1$,
2. $I_i \rightarrow I_{i+1}$, for $i = 1, \dots, n - 2$,
3. $I_{n-1} \rightarrow I_{2i+1}$, for $0 \leq i < (n - 1)/2$.

Proof. By Lemma 7.3.5 and Lemma 7.3.4, there is a non-trivial loop in the P -graph starting from I_1 . By choosing a shortest such loop and renumbering the vertices of the graph, we may assume that we have a loop

$$I_1 \rightarrow I_2 \rightarrow \dots \rightarrow I_k \rightarrow I_1 \tag{7.1}$$

in the P -graph, $k \leq n - 1$. The existence of this loop implies that f has a periodic point of minimal period k . The path

$$I_1 \rightarrow I_1 \rightarrow I_2 \rightarrow \dots \rightarrow I_k \rightarrow I_1$$

implies the existence of a periodic point of minimal period $k + 1$. By the minimality of n , we conclude that $k = n - 1$, which proves statement 1.

Let $I_1 = [x_j, x_{j+1}]$. Note that $\hat{f}(I_1)$ contains I_1 and I_2 , but no other I_i , since otherwise we would have a shorter path than (7.1). Similarly, if $1 \leq i < n - 2$, then $\hat{f}(I_i)$ cannot contain I_k for $k > i + 1$. Thus $\hat{f}(I_1) = [x_j, x_{j+2}]$ or $\hat{f}(I_1) = [x_{j-1}, x_{j+1}]$. Suppose the latter holds (the other case is similar). Then $I_2 = [x_{j-1}, x_j]$, $f(x_{j+1}) = x_{j-1}$, and $f(x_j) = x_{j+1}$. If $2 < n - 1$, then $\hat{f}(I_2)$ can contain at most I_2 and I_3 , so $f(x_{j-1}) = x_{j+2}$. Continuing in this way (see Figure 7.3), we find that the intervals of the partition are ordered




Figure 7.3. The action of f from Lemma 7.3.6 on x_k is shown by arrows.

on the interval I as follows:

$$I_{n-1}, I_{n-3}, \dots, I_2, I_1, I_3, \dots, I_{n-2}.$$

Moreover, $f(x_2) = x_n$, $f(x_n) = x_1$, and $f(x_1) = x_j$, so $\hat{f}(I_{n-1}) = [x_j, x_{n-1}]$, and $\hat{f}(I_{n-1})$ contains all the odd-numbered intervals, which completes the proof of the lemma. \square

COROLLARY 7.3.7. *If n is odd, then f has a periodic point of minimal period q for any $q > n$ and for any even integer $q < n$.*

Proof. Let $m > 1$ be the minimal odd period of a non-fixed periodic point. By the preceding lemma, there are paths of the form

$$I_1 \rightarrow I_1 \rightarrow \dots \rightarrow I_1 \rightarrow I_2 \rightarrow \dots \rightarrow I_{m-1} \rightarrow I_1$$

of any length $q \geq m$. For $q = 2i < m$, the path

$$I_{m-1} \rightarrow I_{m-2i} \rightarrow I_{m-2i+1} \rightarrow \dots \rightarrow I_{m-1}$$

gives a periodic point of period q . The verification that these periodic points have minimal period q is left as an exercise (Exercise 7.3.3). \square

LEMMA 7.3.8. *If n is even, then f has a periodic point of minimal period 2.*

Proof. Let m be the smallest even period of a non-fixed periodic point, and let I_1 be an interval of the associated partition that \hat{f} -covers itself. If no other interval \hat{f} -covers I_1 , then Lemma 7.3.5 implies that $m = 2$.

Suppose then that some other interval \hat{f} -covers I_1 . In the proof of Lemma 7.3.6, we used the hypothesis that n is odd only to conclude the existence of such an interval. Thus the same argument as in the proof of that lemma implies that the P -graph contains the paths

$$I_1 \rightarrow I_2 \rightarrow \dots \rightarrow I_{n-1} \rightarrow I_1 \quad \text{and} \quad I_{n-1} \rightarrow I_{2i} \quad \text{for } 0 \leq i < n/2.$$

Then $I_{n-1} \rightarrow I_{n-2} \rightarrow I_{n-1}$ implies the existence of a periodic point of minimal period 2. \square

Conclusion of the proof of the first assertion of the Sharkovsky Theorem. There are two cases to consider:

1. $n = 2^k$, $k > 0$. If $q < n$, then $q = 2^l$ with $0 \leq l < k$. The case $l = 0$ is trivial. If $l > 0$, then $g = f^{q/2} = f^{2^{l-1}}$ has a periodic point of period 2^{k-l+1} , so by Lemma 7.3.8, g has a non-fixed periodic point of period 2. This point is a fixed point for f^q , i.e., it has period q for f . Since it is not fixed by g , its minimal period is q .

2. $n = p2^k$, p odd. The map f^{2^k} has a periodic point of minimal period p , so by Corollary 7.3.7, f^{2^k} has periodic points of minimal period m for all $m \geq p$ and all even $m < p$. Thus f has periodic points of minimal period $m2^k$ for all $m \geq p$ and all even $m < p$. In particular, f has a periodic point of minimal period 2^{k+1} , so by case 1, f has periodic points of minimal period 2^i for $i = 0, \dots, k$. □

The next lemma finishes the proof of the Sharkovsky theorem.

LEMMA 7.3.9. *For any $\alpha \in \mathbb{N}_{\text{Sh}}$, there is a continuous map $f: [0, 1] \rightarrow [0, 1]$ such that $\text{MinPer}(f) = S(\alpha)$.*

Proof. We distinguish three cases:

1. $\alpha \in \mathbb{N}$, α odd,
2. $\alpha \in \mathbb{N}$, α even, and
3. $\alpha = 2^\infty$.

Case 1. Suppose $n \in \mathbb{N}$ is odd, and $\alpha = n$. Choose points $x_0, \dots, x_{n-1} \in [0, 1]$ such that

$$0 = x_{n-1} < \dots < x_4 < x_2 < x_0 < x_1 < x_3 < \dots < x_{n-2} = 1,$$

and let $I_1 = [x_0, x_1]$, $I_2 = [x_2, x_0]$, $I_3 = [x_1, x_3]$, etc. Let $f: [0, 1] \rightarrow [0, 1]$ be the unique map defined by:

1. $f(x_i) = x_{i+1}$, $i = 0, \dots, n - 2$, and $f(x_{n-1}) = x_0$,
2. f is linear (or affine, to be precise) on each interval I_j , $j = 1, \dots, n - 1$.

Then x_0 is periodic of period n , and the associated P -graph is shown in Figure 7.2. Any path that avoids I_1 has even length. Loops of length less than n must be of the following form:

1. $I_i \rightarrow I_{i+1} \rightarrow \dots \rightarrow I_{n-1} \rightarrow I_{2j+1} \rightarrow I_{2j+2} \rightarrow \dots \rightarrow I_i$ for $i > 1$, or
2. $I_{n-1} \rightarrow I_{2i+1} \rightarrow \dots \rightarrow I_{n-1}$, or
3. $I_1 \rightarrow I_1 \rightarrow \dots \rightarrow I_1 \rightarrow I_1$

Paths of type 1 or 2 have even length, so no point in $\text{int}(I_j)$, $j = 2, \dots, n - 1$, can have odd period $k < n$. Since $f(I_1) = I_1 \cup I_2$, we have $|f'| > 1$ on I_1 , so every non-fixed point in $\text{int}(I_1)$ must move away from the (unique) fixed point in I_1 , and therefore eventually enters I_2 . Once a point enters I_2 , it must enter every I_j before it returns to $\text{int}(I_1)$. Thus there is no non-fixed periodic point in I_1 of period less than n . It follows that no point has odd period less than n . This finishes the proof of the theorem for n odd.




Figure 7.4. Graphs of $\mathcal{D}^k(f)$ for $f \equiv 1/2$.

Case 2. Suppose $n \in \mathbb{N}$ is even, and $\alpha = n$. For $f: [0, 1] \rightarrow [0, 1]$, define a new function $\mathcal{D}: [0, 1] \rightarrow [0, 1]$ by

$$\mathcal{D}(f)(x) = \begin{cases} \frac{2}{3} + \frac{1}{3} f(3x) & x \in [0, \frac{1}{3}], \\ (2 + f(1)) (\frac{2}{3} - x) & x \in [\frac{1}{3}, \frac{2}{3}], \\ x - \frac{2}{3} & x \in [\frac{2}{3}, 1]. \end{cases}$$

The operator $\mathcal{D}(f)$ is sometimes called the *doubling operator*, because $\text{MinPer}(\mathcal{D}(f)) = 2 \text{MinPer}(f) \cup \{1\}$, i.e., \mathcal{D} doubles the periods of a map. To see this, let $g = \mathcal{D}(f)$, and let $I_1 = [0, 1/3]$, $I_2 = [1/3, 2/3]$, and $I_3 = [2/3, 1]$. For $x \in I_1$, we have $g^2(x) = f(3x)/3$, so $g^{2k}(x) = f^k(3x)/3$. Thus $g^{2k}(x) = x$ if and only if $f^k(3x) = 3x$, so $\text{MinPer}(g) \supset 2 \text{MinPer}(f)$ (see Figure 7.4).

On the interval I_2 , $|g'| \geq 2$, so there is a unique repelling fixed point in $(1/3, 2/3)$, and every other point eventually leaves this interval and never returns, since $g(I_1 \cup I_3) \cap I_2 = \emptyset$. Thus no non-fixed point in I_2 is periodic.

Finally, any periodic point in I_3 enters I_1 , so its period is in $2\text{MinPer}(f)$, which verifies our claim that $\text{MinPer}(\mathcal{D}(f)) = 2\text{MinPer}(f) \cup \{1\}$.

Since n is even, we can write $n = p2^k$, where p is odd and $k > 0$. Let f be a map whose minimum odd period is p (see case 1). Then $\text{MinPer}(\mathcal{D}^k(f)) = 2^k \text{MinPer}(f) \cup \{2^{k-1}, 2^{k-2}, \dots, 1\}$, which settles case 2 of the lemma.

Case 3. Suppose $\alpha = 2^\infty$. Let $g_k = \mathcal{D}^k(\text{Id})$, where Id is the identity map. Then, by the induction and the remarks in the proof of case 2, $\text{MinPer}(g_k) = \{2^{k-1}, 2^{k-2}, \dots, 1\}$. The sequence $\{g_k\}_{k \in \mathbb{N}}$ converges uniformly to a continuous map $g_\infty: [0, 1] \rightarrow [0, 1]$, and $g_\infty = g_k$ on $[2/3^k, 1]$ (Exercise 7.3.4). It follows that $\text{MinPer}(g_\infty) \supset S(2^\infty)$.

Let x be a periodic point of g_∞ . If $0 \notin \mathcal{O}(x)$, then $\mathcal{O}(x) \subset [2/3^k, 1]$ for k sufficiently large, so x is a periodic point of g_k and has even period. Suppose then that 0 is periodic with period p . If $p > 2^\infty$, then there is $q \in \mathbb{N}$ such

that $p > q > 2^\infty$. By the first part of the Sharkovsky theorem, g_∞ has a periodic point y with minimal period q . Since $0 \in \mathcal{O}(y)$, we conclude by the preceding argument that q is even, which contradicts $q > 2^\infty$. Thus $\text{MinPer}(g_\infty) = S(2^\infty)$.

This concludes the proof of Lemma 7.3.9, and thus the proof of Theorem 7.3.1. \square

Exercise 7.3.1. Let σ be a permutation of $\{1, \dots, n-1\}$. Show that there is a continuous map $f: [0, 1] \rightarrow [0, 1]$ with a periodic point x of period n such that $x < f^{\sigma(1)} < \dots < f^{\sigma(n-1)}$.

Exercise 7.3.2. Show that there are maps $f, g: [0, 1] \rightarrow [0, 1]$, each with a periodic point of period n (for some n), such that the associated P -graphs are not isomorphic. (Note that for $n = 3$, all P -graphs are isomorphic.)

Exercise 7.3.3. Verify that the periodic points in the proof of Corollary 7.3.7 have *minimal* period q .

Exercise 7.3.4. Show that the sequence $\{g_k\}_{k \in \mathbb{N}}$ defined near the end of the proof of Lemma 7.3.9 converges uniformly, and the limit g_∞ satisfies $g_\infty = g_k$ on $[2/3^k, 1]$.

7.4 Combinatorial Theory of Piecewise-Monotone Mappings¹

Let $I = [a, b]$ be a compact interval. A continuous map $f: I \rightarrow I$ is *piecewise monotone* if there are points $a = c_0 < c_1 < \dots < c_l < c_{l+1} = b$ such that f is strictly monotone on each interval $I_i = [c_{i-1}, c_i]$, $i = 1, \dots, l+1$. We always assume that each interval $[c_{i-1}, c_i]$ is a maximal interval on which f is monotone, so the orientation of f reverses at the *turning points* c_1, \dots, c_l . The intervals I_i are called *laps* of f .

Note that any piecewise-monotone map $f: I \rightarrow I$ can be extended to a piecewise-monotone map of a larger interval J in such a way that $f(\partial J) \subset \partial J$. Thus we assume (without losing much generality) that $f(\partial I) \subset \partial I$. If f has l turning points and $f(\partial I) \subset \partial I$, then f is *l -modal*. If f has exactly one turning point, then f is *unimodal*.

The *address* of a point $x \in I$ is the symbol c_j if $x = c_j$ for some $j \in \{1, \dots, l\}$, or the symbol I_j if $x \in I_j$ and $x \notin \{c_1, \dots, c_l\}$. Note that c_0 and c_{l+1} are *not* included as addresses. The *itinerary* of x is the sequence $i(x) =$

¹ Our arguments in this section follow in part those of [CE80] and [MT88].

$(i_k(x))_{k \in \mathbb{N}_0}$, where $i_k(x)$ is the address of $f^k(x)$. Let

$$\Sigma = \{I_1, \dots, I_{l+1}, c_1, \dots, c_l\}^{\mathbb{N}_0}.$$

Then $i: I \rightarrow \Sigma$, and $i \circ f = \sigma \circ i$, where σ is the one-sided shift on Σ .

Example. Any quadratic map $q_\mu(x) = \mu x(1 - x)$, $0 < \mu \leq 4$, is a unimodal map of $I = [0, 1]$, with turning point $c_1 = 1/2$, $I_1 = [0, 1/2]$, $I_2 = [1/2, 1]$. If $0 < \mu < 2$, then $f(I) \subset [0, 1/2)$, so the only possible itineraries are (I_1, I_1, \dots) , (c_1, I_1, I_1, \dots) , and (I_2, I_1, I_1, \dots) . Note that the map $i: [0, 1] \rightarrow \Sigma$ is not continuous at c_1 .

If $\mu = 2$, then the possible itineraries are (I_1, I_1, \dots) , (c_1, c_1, \dots) , and (I_2, I_1, I_1, \dots) . If $2 < \mu < 3$, there is an attracting fixed point $(\mu - 1)/\mu \in (1/2, 2/3)$. Thus the possible itineraries are:

- (I_1, I_1, \dots) ,
- (c_1, I_2, I_2, \dots) ,
- (I_2, I_2, \dots) ,
- $(I_1, \dots, I_1, I_2, I_2, \dots)$,
- $(I_1, \dots, I_1, C_1, I_2, I_2, \dots)$,
- any of the above preceded by I_2 .

LEMMA 7.4.1. *The itinerary $i(x)$ is eventually periodic if and only if the iterates of x converge to a periodic orbit of f .*

Proof. If $i(x)$ is eventually periodic, then by replacing x by one of its forward iterates, we may assume that $i(x)$ is periodic, of period p . If $i_j(x) = c_j$ for some j , then c_j is periodic, and we are done. Thus we may assume that $f^k(x)$ is contained in the interior of a lap of f for each k . For $j = 0, \dots, p - 1$, let J_j be the smallest closed interval containing $\{f^k(x): k = j \text{ mod } p\}$. Since the itinerary is periodic of period p , each J_i is contained in a single lap, so $f: J_j \rightarrow J_{j+1}$ is strictly monotone. It follows that $f^p: J_0 \rightarrow J_0$ is strictly monotone.

Suppose $f^p: J_0 \rightarrow J_0$ is increasing. If $f^p(x) \geq x$, then by induction, $f^{kp}(x) \geq f^{(k-1)p}(x)$ for all $k > 0$, so $\{f^{kp}(x)\}$ converges to a point $y \in J_0$, which is fixed for f^p . A similar argument holds if $f^p(x) < x$.

If $f^p: J_0 \rightarrow J_0$ is decreasing, then $f^{2p}: J_0 \rightarrow J_0$ is increasing, and by the argument in the preceding paragraph, the sequence $\{f^{2kp}(x)\}$ converges to a fixed point of f^{2p} .

Conversely, suppose that $f^{kp}(x) \rightarrow y$ as $k \rightarrow \infty$, where $f^q(y) = y$. If the orbit of y does not contain any turning points, then eventually x has the same

itinerary as y . The case where $\mathcal{O}(y)$ does contain a turning point is left as an exercise (Exercise 7.4.1). \square

Let ϵ be a function defined on $\{I_1, \dots, I_l, c_1, \dots, c_l\}$ such that $\epsilon(I_1) = \pm 1$, $\epsilon(I_k) = (-1)^{k+1}\epsilon(I_1)$, and $\epsilon(c_k) = 1$ for $k = 0, \dots, l$. Associated to ϵ is a *signed lexicographic ordering* $<$ on Σ , defined as follows. For $s \in \Sigma$, define

$$\tau_n(s) = \prod_{0 \leq k < n} \epsilon(s_k).$$

We order the symbols $\{\pm I_j, \pm c_k\}$ by

$$-I_{l+1} < -c_l < -I_l < \dots < -I_1 < I_1 < c_1 < I_2 < \dots < c_l < I_{l+1}.$$

Given $s = (s_i), t = (t_i) \in \Sigma$, we say $s < t$ if and only if $s_0 < t_0$, or there is $n > 0$ such that $s_i = t_i$ for $i = 0, \dots, n - 1$, and $\tau_n(s)s_n < \tau_n(t)t_n$. The proof that $<$ is an ordering is left as an exercise.

Associated to an l -modal map f is a natural signed lexicographic ordering with $\epsilon(I_k) = 1$ if f is increasing on I_k and $\epsilon(I_k) = -1$ otherwise, and $\epsilon(c_k) = 1$, for $k = 1, \dots, l$. For $x \in I$, we define $\tau_n(x) = \tau_n(i(x))$. Note that if $\{x, f(x), \dots, f^{n-1}(x)\}$ contains no turning points, then $\tau_n(x)$ is the orientation of f^n at x : positive (i.e., increasing) if and only if $\tau_n(x) = 1$.

LEMMA 7.4.2. *For $x, y \in I$, if $x < y$, then $i(x) \leq i(y)$. Conversely, if $i(x) < i(y)$, then $x < y$.*

Proof. Suppose $i(x) \neq i(y)$, $i_k(x) = i_k(y)$ for $k = 0, \dots, n - 1$, and $i_n(x) \neq i_n(y)$. Then there is no turning point in the intervals $[x, y]$, $f([x, y])$, \dots , $f^{n-1}([x, y])$, so f^n is monotone on $[x, y]$, and is increasing if and only if $\tau_n(i(x)) = 1$. Thus $x < y$ if and only if $\tau_n(x)f^n(x) < \tau_n(y)f^n(y)$, and the latter holds if and only if $\tau_n(x)i_n(x) < \tau_n(y)i_n(y)$ since $i_n(x) \neq i_n(y)$. \square

LEMMA 7.4.3. *Let $I(x) = \{y: i(y) = i(x)\}$. Then:*

1. $I(x)$ is an interval (which may consist of a single point).
2. If $I(x) \neq \{x\}$, then $f^n(I(x))$ does not contain any turning points for $n \geq 0$. In particular, every power of f is strictly monotone on $I(x)$.
3. Either the intervals $I(x), f(I(x)), f^2(I(x)), \dots$ are pairwise disjoint, or the iterates of every point in $I(x)$ converge to a periodic orbit of f .

Proof. Lemma 7.4.2 implies immediately that $I(x)$ is an interval. To prove part 2, suppose there is $y \in I(x)$ such that $f^n(y)$ is a turning point. If $I(x)$ is not a single point, then there is some point $z \in I(x)$ such that $f^n(y) \neq f^n(z)$, since f^n is not constant on any interval. Thus $i_n(z) \neq i_n(y) = f(y)$, which contradicts the fact that $y, z \in I(x)$. Thus $I(x)$ must be a single point.

To prove part 3, suppose the intervals $I(x), f(I(x)), f^2(I(x)), \dots$ are not pairwise disjoint. Then there are integers $n \geq 0, p > 0$, such that $f^n(I(x)) \cap f^{n+p}(I(x)) \neq \emptyset$. Then $f^{n+kp}(I(x)) \cap f^{n+(k+1)p}(I(x)) \neq \emptyset$ for all $k \geq 1$. It follows that $L = \bigcup_{k \geq 1} f^{kp}(I(x))$ is a non-empty interval that contains no turning points and is invariant by f^p . Since f^p is strictly monotone on L , for any $y \in L$, the sequence $\{f^{2kp}(y)\}$ is monotone and converges to a fixed point of f^{2p} . \square

An interval $J \subset I$ is *wandering* if the intervals $J, f(J), f^2(J), \dots$ are pairwise disjoint, and $f^n(J)$ does not converge to a periodic orbit of f . Recall that if x is an attracting periodic point, then the *basin of attraction* $BA(x)$ of x is the set of all points whose ω -limit set is $\mathcal{O}(x)$.

COROLLARY 7.4.4. *Suppose f does not have wandering intervals, attracting periodic points, or intervals of periodic points. Then $i: I \rightarrow \Sigma$ is an injection, and therefore a bijective order-preserving map onto its image.*

Proof. To prove that i is injective we need only show that $I(x) = \{x\}$ for every $x \in I$. If not, then by the proof of Lemma 7.4.3, either $I(x)$ is wandering or there is an interval L with non-empty interior and $p > 0$ such that f^p is monotone on $L, f^p(L) \in L$, and the iterates of any point in L converge to a periodic orbit of f of period $2p$. The former case is excluded by hypothesis. In the latter case, by Exercise 7.4.2, either L contains an interval of periodic points, or some open interval in L converges to a single periodic point, contrary to the hypothesis. So $I(x) = \{x\}$. \square

Our next goal is to characterize the subset $i(I) \subset \Sigma$. As we indicated above, the map $i: I \rightarrow \Sigma$ is not continuous. Nevertheless, for any $x \in I$ and $k \in \mathbb{N}_0$, there is $\delta > 0$ such that $i_k(y)$ is constant on $(x, x + \delta)$ and on $(x - \delta, x)$ (but not necessarily the same on both intervals). Thus the limits $i(x^+) = \lim_{y \rightarrow x^+} i(y)$ and $i(x^-) = \lim_{y \rightarrow x^-} i(y)$ exist. Moreover, $i(x^+)$ and $i(x^-)$ are both contained in $\{I_1, \dots, I_l\}^{\mathbb{N}_0} \subset \Sigma$. For $j = 1, \dots, l$, we define the *j*th *kneading invariant* of f to be $v_j = i(c_j^+)$. For convenience we also define sequences $v_0 = i(c_0) = i(c_0^+)$ and $v_{l+1} = i(c_{l+1}) = i(c_{l+1}^-)$. Note that v_0 and v_{l+1} are eventually periodic of period 1 or 2, since by hypothesis the set $\{c_0, c_{l+1}\}$ is invariant. In fact, there are only four possibilities for the pair v_0, v_{l+1} , corresponding to the four possibilities for $f|_{\partial I}$.

LEMMA 7.4.5. *For any $x \in I, i(x)$ satisfies the following:*

1. $\sigma^n i(x) = i(c_k)$ if $f^n(x) = c_k$.
2. $\sigma v_k \leq \sigma^{n+1} i(x) \leq \sigma v_{k+1}$ if $f^n(x) \in I_{k+1}$ and f is increasing on I_{k+1} .
3. $\sigma v_k \geq \sigma^{n+1} i(x) \geq \sigma v_{k+1}$ if $f^n(x) \in I_{k+1}$ and f is decreasing on I_{k+1} .

Moreover, if f has no wandering intervals, attracting periodic points, or intervals of periodic points, then the inequalities in conditions 2 and 3 are strict.

Proof. The first assertion is obvious. To prove the second, suppose that $f^n(x) \in I_{k+1}$, and f is increasing on I_{k+1} . Then for $y \in (c_k, f^n(x))$, we have $f(c_k) < f(y) < f^{n+1}(x)$, so

$$i(f(c_k)) \leq i(f(y)) \leq i(f^{n+1}(x)) = \sigma^{n+1}i(x).$$

Since $v_k = \lim_{y \rightarrow c_k^+} i(y)$, we conclude that $\sigma v_k \leq \sigma^{n+1}i(x)$. The other inequalities are proved in a similar way.

If f has no wandering intervals, attracting periodic points, or intervals of periodic points, then Corollary 7.4.4 implies that i is injective, so \leq can be replaced by $<$ everywhere in the preceding paragraph. \square

The following immediate corollary of Lemma 7.4.5 gives an admissibility criterion for kneading invariants.

COROLLARY 7.4.6. *If $\sigma^n(v_j) = (I_{k+1}, \dots)$, then*

1. $\sigma v_k \leq \sigma^{n+1}v_j \leq \sigma v_{k+1}$ if f is increasing on I_{k+1} ,
2. $\sigma v_k \geq \sigma^{n+1}v_j \geq \sigma v_{k+1}$ if f is decreasing on I_{k+1} .

Let $f: I \rightarrow I$ be an l -modal map with kneading invariants v_1, \dots, v_l , and let v_0, v_{l+1} be the itineraries of the endpoints of I . Define Σ_f to be the set of all sequences $t = (t_n) \in \Sigma$ satisfying the following:

1. $\sigma^n t = i(c_k)$ if $t_n = c_k, k \in \{0, \dots, l\}$.
2. $\sigma v_k < \sigma^{n+1}t < \sigma v_{k+1}$ if $t_n = I_{k+1}$ and $\epsilon(I_{k+1}) = +1$.
3. $\sigma v_k > \sigma^{n+1}t > \sigma v_{k+1}$ if $t_n = I_{k+1}$ and $\epsilon(I_{k+1}) = -1$.

Similarly, we define $\hat{\Sigma}_f$ to be the set of sequences in Σ satisfying conditions 1–3 with $<$ replaced by \leq .

THEOREM 7.4.7. *Let $f: I \rightarrow I$ be an l -modal map with kneading invariants v_1, \dots, v_l , and let v_0, v_{l+1} be the itineraries of the endpoints. Then $i(I) \subset \hat{\Sigma}_f$. Moreover, if f has no wandering intervals, attracting periodic points, or intervals of periodic points, then $i(I) = \Sigma_f$, and $i: I \rightarrow \Sigma_f$ is an order-preserving bijection.*

Proof. Lemma 7.4.5 implies that $i(I) \subset \hat{\Sigma}_f$, and $i(I) \subset \Sigma_f$ if there are no wandering intervals, attracting periodic points, or intervals of periodic points.

Suppose f has no wandering intervals, attracting periodic points or intervals of periodic points. Let $t = (t_n) \in \Sigma_f$, and suppose $t \notin i(I)$. Then

$$L_t = \{x \in I: i(x) < t\}, \quad R_t = \{x \in I: i(x) > t\}$$

are disjoint intervals, and $I = L_t \cup R_t$.

We claim that L_t and R_t are non-empty. The proof of this claim breaks into four cases according to the four possibilities for $f|_{\partial I}$. We prove it in the case $f(c_0) = f(c_{l+1}) = c_0$. Then $v_0 = i(c_0^+) = (I_1, I_1, \dots)$, $v_{l+1} = i(c_{l+1}^-) = (I_{l+1}, I_1, I_1, \dots)$, $\epsilon(I_1) = 1$, and $\epsilon(I_{l+1}) = -1$. Note that $t \neq i(c_0) = v_0$ and $t \neq i(c_{l+1}) = v_{l+1}$, since $t \notin i(f)$. Thus $v_0 < t$, so $c_0 \in L_t$. If $t_0 < I_{l+1}$, then $t < v_{l+1}$, so $c_{l+1} \in R_t$, and we are done. So suppose $t_0 = I_{l+1}$. If $t_1 > I_1$, then $t < v_{l+1}$, and again we are done. If $t_1 = I_1$, then condition 2 implies that $\sigma v_0 < \sigma^2 t$, which implies in turn that $t < v_{l+1}$. Thus $v_{l+1} \in R_t$.

Let $a = \sup L_t$. We will show that $a \notin L_t$. Suppose for a contradiction that $a \in L_t$. Since $x \notin L_t$ for all $x > a$, we conclude that $i(a) < t \leq i(a^+)$. This implies that the orbit of a contains a turning point. Let $n \geq 0$ be the smallest integer such that $i_n(a) = c_k$ for some $k \in \{1, \dots, l\}$. Then $i_j(a) = t_j = i_j(a^+)$ for $j = 1, \dots, n-1$, and $i_n(a^+) = I_k$ or $i_n(a^+) = I_{k+1}$. Suppose the latter holds. Then f^n is increasing on a neighborhood of a . Since $i(a) < t \leq i(a^+)$ and $i_j(a) = t_j = i_j(a^+)$ for $j = 0, \dots, n-1$, it follows that

$$i(c_k) = \sigma^n(i(a)) < \sigma^n(t) < \sigma^n(i(a^+)) = v_k,$$

and $c_k \leq t_n \leq I_{k+1}$.

If $t_n = c_k$, then by condition 1, $\sigma^n(t) = i(c_k)$, so $t = i(a)$, contradicting the fact that $t \notin i(f)$. Thus we may assume that $t_n = I_{k+1}$. If f is increasing on I_{k+1} , then condition 2 implies that $\sigma^{n+1}(t) > \sigma v_k$. But $\sigma^n(t) \leq \sigma^n(i(a^+))$, $\tau_n(t) = +1$ and $t_n = i_n(a^+)$ imply that

$$\sigma^{n+1}(t) \leq \sigma^{n+1}(i(a^+)) = \sigma(v_k).$$

Similarly, if f is decreasing on I_{k+1} , then condition 3 implies that $\sigma^{n+1} < \sigma v_k$, which contradicts $\sigma^n(t) \leq \sigma^n(i(a^+))$, $\tau_{n+1}(t) = -1$, and $t_n = i_n(a^+)$.

We have shown that the case $i_n(a^+) = I_{k+1}$ leads to a contradiction. Similarly, the case $i_n(a^+) = I_k$ leads to a contradiction. Thus $a \notin L_t$. By similar arguments, $\inf R_t \notin R_t$, which contradicts the fact that I is the disjoint union of L_t and R_t . Thus $t \in i(I)$, so $i(I) = \Sigma_f$.

Lemma 7.4.2 now implies that $i: I \rightarrow \Sigma_f$ is an order-preserving bijection. □

COROLLARY 7.4.8. *Let f and g be l -modal maps of I with no wandering intervals, no attracting periodic points, and no intervals of periodic points. If f and g have the same kneading invariants and endpoint itineraries, then f and g are topologically conjugate.*

Proof. Let i_f and i_g be the itinerary maps of f and g , respectively. Then $i_f^{-1} \circ i_g: I \rightarrow \Sigma(v_0, v_1, \dots, v_{l+1}) \rightarrow I$ is an order-preserving bijection, and therefore a homeomorphism, which conjugates f and g . □

REMARK 7.4.9. *One can show that the following extension of Corollary 7.4.8 is also true: Let f and g be l -modal maps of I , and suppose f has no wandering intervals, no attracting periodic points, and no intervals of periodic points. If f and g have the same kneading invariants and endpoint itineraries, then f and g are topologically semiconjugate.*

Example. Consider the unimodal quadratic map $f: [-1, 1] \rightarrow [-1, 1]$, $f(x) = -2x^2 + 1$. This map is conjugate to the quadratic map $q_4: [0, 1] \rightarrow [0, 1]$, $q_4(x) = 4x(1 - x)$, via the homeomorphism $h: [-1, 1] \rightarrow [0, 1]$, $h(x) = \frac{1}{2}(x + 1)$. The orbit of the turning point $c = 0$ of f is $0, 1, -1, -1, \dots$, so the kneading invariant is $\nu = (I_2, I_2, I_1, I_1, \dots)$.

Now let $I = [-1, 1]$, and consider the *tent map* $T: I \rightarrow I$ defined by

$$T(x) = \begin{cases} 2x + 1, & x \leq 0, \\ -2x + 1, & x > 0. \end{cases}$$

The homeomorphism $\phi: I \rightarrow I$, $\phi(x) = (2/\pi) \sin^{-1}(x)$ conjugates f to T .

For any $n > 0$, the map f^{n+1} maps each of the intervals $[k/2^n, (k+1)/2^n]$, $k = -2^n, \dots, 2^n$, homeomorphically onto I . Thus the forward iterates of any open set cover I , or equivalently, the backward orbit of any point in I is dense in I . It follows from the next lemma that T has no wandering intervals, attracting periodic points, or intervals of periodic points, so any unimodal map with the same kneading invariants as T is semiconjugate to T . In particular, any unimodal map $g: [a, b] \rightarrow [a, b]$ with $g(a) = g(b) = a$ and $g(c) = b$ is semiconjugate to T .

LEMMA 7.4.10. *Let $I = [a, b]$ be an interval, and $f: I \rightarrow I$ a continuous map with $f(\partial I) \subset \partial I$. Suppose that every backward orbit is dense in I , and that f has a fixed point x_0 not in ∂I . Then f has no wandering intervals, no intervals of periodic points, and no attracting periodic points.*

Proof. Let $U \in I$ be an open interval. Fix $x \in U$. By density of $\bigcup f^{-n}(x)$, there is $n > 0$ such that $f^{-n}(x) \cap U \neq \emptyset$. Then $f^n(U) \cap U \neq \emptyset$, so U is not a wandering interval.

Suppose $z \in I$ is an attracting periodic point. Then the basin of attraction $\text{BA}(z)$ is a forward-invariant set with non-empty interior. Since backward orbits are dense, $\text{BA}(z)$ is a dense open subset of I and therefore intersects the backward orbit of x_0 . Thus $z = x_0$. On the other hand, the backward orbits of a and b are dense, and therefore intersect $\text{BA}(z)$, which is a contradiction. Thus there can be no attracting periodic point.

Any point in $\text{Per}(f)$ has finitely many preimages in $\text{Per}(f)$, so if $\text{Per}(f)$ had non-empty interior, the backward orbit of a point in $\text{Per}(f)$ would not be dense in $\text{Per}(f)$. Thus f has no intervals of periodic points. \square

The final result of this section is a realization theorem, which asserts that any “admissible” set of sequences in Σ is the set of kneading invariants of an l -modal map.

Note that for an l -modal map f , the endpoint itineraries are determined completely by the orientation of f on the first and last laps of f . Thus, given l and a function ϵ as in the definition of signed lexicographic orderings, we can define natural *endpoint itineraries* v_0 and v_{l+1} as sequences in the symbol space $\{I_1, I_{l+1}\}$.

THEOREM 7.4.11. *Let $v_1, \dots, v_l \in \{I_1, \dots, I_{l+1}\}^{\mathbb{N}_0}$, and $\epsilon(I_j) = \epsilon_0(-1)^j$, where $\epsilon_0 = \pm 1$. Let $<$ be the signed lexicographic ordering on $\Sigma = \{I_1, \dots, I_{l+1}, c_1, \dots, c_l\}^{\mathbb{N}_0}$ associated to ϵ . Let v_0, v_{l+1} be the endpoint itineraries determined uniquely by ϵ and l . If $\{v_0, \dots, v_{l+1}\}$ satisfies the admissibility criterion of Corollary 7.4.6, then there is a continuous l -modal map $f: [0, 1] \rightarrow [0, 1]$ with kneading invariants v_1, \dots, v_{l+1} .*

Proof. Define an equivalence relation \sim on Σ by the rule $t \sim s$ if and only if $t = s$, or $\sigma(t) = \sigma(s)$ and $t_0 = I_k, s_0 = I_{k\pm 1}$. To paraphrase: t and s are equivalent if and only if they differ at most in the first position, and then only if the first positions are adjacent intervals. (Thus, for example, $i(c_k^-) \sim i(c_k^+)$ for a turning point of an l -modal map.)

We will define a sequence of l -modal maps $f_N, N \in \mathbb{N}_0$, whose kneading invariants agree up to order N with v_1, \dots, v_l . The desired map f will be the limit in the C^0 topology of these maps.

Let $p_j^0 = c_j, j = 0, \dots, l + 1$. Choose points $p_j^1 \in [0, 1], j = 0, \dots, l + 1$, such that

1. if $\sigma^m(v_i) \sim \sigma^n(v_j)$ then $p_i^m = p_j^n$;
2. $p_i^m < p_j^n$ if and only if $\sigma^m(v_i) < \sigma^n(v_j)$ and $\sigma^m(v_i) \not\sim \sigma^n(v_j)$; and
3. the new points are equidistributed in each of the intervals $[p_j^0, p_{j+1}^0], j = 0, \dots, l + 1$.

Define $f_1: [0, 1] \rightarrow [0, 1]$ to be the piecewise-linear map specified by $f(p_j^0) = p_j^1$. Note that $p_j^1 < p_{j+1}^1$ if and only if $\sigma v_j < \sigma v_{j+1}$, which happens if and only if $\epsilon(I_{j+1}) = +1$. Thus f_1 is l -modal.

For $N > 0$ we define inductively points $p_j^N \in [0, 1], j = 0, \dots, l + 1$, satisfying conditions 1 and 2 for all $n, m \leq N$ and $j = 0, \dots, l + 1$, and so that in any subinterval defined by the points $\{p_j^n: 0 < n < N, 0 \leq j \leq l + 1\}$, the

new points $\{p_j^N\}$ in that interval are equidistributed. Then we define the map $f_N: I \rightarrow I$ to be the piecewise-linear map connecting the points (p_j^n, p_j^{n+1}) , $j = 0, \dots, l+1$, $n = 0, \dots, N-1$. It follows (Exercise 7.4.5) that:

1. f_N is l -modal for each $N > 0$;
2. $\{f_N\}$ converges in the C^0 topology to an l -modal map f with turning points c_1, \dots, c_l ; and
3. the kneading invariants of f are v_1, \dots, v_l . □

Exercise 7.4.1. Finish the proof of Lemma 7.4.1.

Exercise 7.4.2. Let L be an interval and $f: L \rightarrow L$ a strictly monotone map. Show that either L contains an interval of periodic points, or some open interval in L converges to a single periodic point.

Exercise 7.4.3. Work out the ordering on the set of itineraries of the quadratic map q_μ for $2 < \mu < 3$.

Exercise 7.4.4. Show that the tent map has exactly 2^n periodic points of period n , and the set of periodic points is dense in $[-1, 1]$.

Exercise 7.4.5. Verify the last three assertions in the proof of Theorem 7.4.11.

7.5 The Schwarzian Derivative

Let f be a C^3 function defined on an interval $I \subset \mathbb{R}$. If $f'(x) \neq 0$, we define the *Schwarzian derivative* of f at x to be

$$Sf(x) = \frac{f'''(x)}{f'(x)} - \frac{3}{2} \left(\frac{f''(x)}{f'(x)} \right)^2.$$

If x is an isolated critical point of f , we define $Sf(x) = \lim_{y \rightarrow x} Sf(y)$ if the limit exists.

For the quadratic map $q_\mu(x) = \mu x(1-x)$, we have that $Sq_\mu(x) = -6/(1-2x)^2$ for $x \neq 1/2$, and $Sf(1/2) = -\infty$. We also have $S \exp(x) = -1/2$ and $S \log(x) = 1/2x^2$.

LEMMA 7.5.1. *The Schwarzian derivative has the following properties:*

1. $S(f \circ g) = (Sf \circ g)(g')^2 + Sg$.
2. $S(f^n) = \sum_{i=0}^{n-1} Sf(f^i(x)) \cdot ((f^i)')^2$.
3. If $Sf < 0$, then $S(f^n) < 0$ for all $n > 0$.

The proof is left as an exercise (Exercise 7.5.3).

A function with negative Schwarzian derivative satisfies the following *minimum principle*.

LEMMA 7.5.2 (Minimum Principle). *Let I be an interval and $f: I \rightarrow I$ a C^3 map with $f'(x) \neq 0$ for all $x \in I$. If $Sf < 0$, then $|f'(x)|$ does not attain a local minimum in the interior of I .*

Proof. Let z be a critical point of f' . Then $f''(z) = 0$, which implies that $f'''(z)/f'(z) < 0$, since $Sf < 0$. Thus $f'''(z)$ and $f'(z)$ have opposite signs. If $f'(z) < 0$, then $f'''(z) > 0$ and z is a local minimum of f' , so z is a local maximum of $|f'|$. Similarly, if $f'(z) > 0$, then z is also a local maximum of $|f'|$. Since f' is never zero on I , this implies that $|f'|$ does not have a local minimum on I . \square

THEOREM 7.5.3 (Singer). *Let I be a closed interval (possibly unbounded), and $f: I \rightarrow I$ a C^3 map with negative Schwarzian derivative. If f has n critical points, then f has at most $n + 2$ attracting periodic orbits.*

Proof. Let z be an attracting periodic point of period m . Let $W(z)$ be the maximal interval about z such that $f^{mn}(y) \rightarrow z$ as $n \rightarrow \infty$ for all $y \in U$. Then $W(z)$ is open (in I), and $f^m(W(z)) \subset W(z)$.

Suppose that $W(z)$ is bounded and does not contain a point in ∂I , so $W(z) = (a, b)$ for some $a < b \in \mathbb{R}$. We claim that f^m has a critical point in $W(z)$. By maximality of $W(z)$, f^m must preserve the set of endpoints of $W(z)$. If $f^m(a) = f^m(b)$, then f^m must have a maximum or minimum in $W(z)$, and therefore a critical point in $W(z)$. If $f^m(a) \neq f^m(b)$, then f^m must permute a and b . Suppose $f^m(a) = a$ and $f^m(b) = b$. Then $(f^m)' \geq 1$ on ∂U , since otherwise a or b would be an attracting fixed point for f^m whose basin of attraction overlaps U . By the minimum principle, if f^m has no critical points in U , then $(f^m)' > 1$ on U , which contradicts $f^m(W(z)) = W(z)$, so f^m has a critical point in $W(z)$. If $f^m(a) = b$ and $f^m(b) = a$, then applying the preceding argument to f^{2m} , we conclude that f^{2m} has a critical point in $W(z)$. Since $f^m(W(z)) = W(z)$, it follows that f^m also has a critical point in $W(z)$.

By the chain rule, if $p \in W(z)$ is a critical point of f^m , then one of the points $p, f(p), \dots, f^{m-1}(p)$ is a critical point of f . Thus we have shown that either $W(z)$ is unbounded, or it meets ∂I , or there is a critical point of f whose orbit meets $W(z)$. Since there are only n critical points, and there are only two boundary points (or unbounded ends) of I , the theorem is proved. \square

COROLLARY 7.5.4. *For any $\mu > 4$, the quadratic map $q_\mu: \mathbb{R} \rightarrow \mathbb{R}$ has at most one (finite) attracting periodic orbit.*

Proof. The proof of Theorem 7.5.3 shows that if z is an attracting periodic point, then $W(z)$ either is unbounded or contains the critical point of q_μ . Since ∞ is an attracting periodic point, the basin of attraction of z must be bounded, and therefore must contain the critical point. \square

We now discuss a relation between the Schwarzian derivative and length distortion that is used in producing absolutely continuous invariant measures for maps of the interval with negative Schwarzian derivative.²

Let f be a piecewise-monotone real-valued function defined on a bounded interval I . Suppose $J \subset I$ is a subinterval such that $I \setminus J$ consists of disjoint non-empty intervals L and R . Denote by $|F|$ the length of an interval F . Define the *cross-ratios*

$$C(I, J) = \frac{|I| \cdot |J|}{|J \cup L| \cdot |J \cup R|}, \quad D(I, J) = \frac{|I| \cdot |J|}{|L| \cdot |R|}.$$

If f is monotone on I , set

$$A(I, J) = \frac{C(f(I), f(J))}{C(I, J)}, \quad B(I, J) = \frac{D(f(I), f(J))}{D(I, J)}.$$

The group \mathcal{M} of real Möbius transformations consists of maps of the extended real line $\mathbb{R} \cup \{\infty\}$ of the form $\phi(x) = (ax + b)/(cx + d)$, where $a, b, c, d \in \mathbb{R}$ and $ad - bc \neq 0$. Möbius transformations have Schwarzian derivative equal to 0 and preserve the cross-ratios C and D (Exercise 7.5.4). The group of Möbius transformations is *simply transitive* on triples of points in the extended real line, i.e., given any three distinct points $a, b, c \in \mathbb{R} \cup \{\infty\}$, there is a unique Möbius transformation $\phi \in \mathcal{M}$ such that $\phi(0) = a$, $\phi(1) = b$ and $\phi(\infty) = c$ (Exercise 7.5.5). Möbius transformations are also called *linear fractional transformations*.

PROPOSITION 7.5.5. *Let f be a C^3 real-valued function defined on a compact interval I such that f has negative Schwarzian derivative and $f'(x) \neq 0$, $x \in I$. Let $J \subset I$ be a closed subinterval that does not contain the endpoints of I . Then $A(I, J) > 1$ and $B(I, J) > 1$.*

Proof. Since every Möbius transformation has Schwarzian derivative 0 and preserves C and D , we may assume, by composing f on the left and on the right with appropriate Möbius transformations and using Lemma 7.5.1, that $I = [0, 1]$, $J = [a, b]$ with $0 < a < b < 1$, $f(0) = 0$, $f(a) = a$, and $f(1) = 1$. By Lemma 7.5.2, $|f'|$ does not have a local minimum in $[0, 1]$, and hence f cannot have fixed points except 0, a , and 1. Therefore $f(x) < x$ if $0 < x < a$

² Our exposition here follows to a large extent [vS88] and [dMvS93]

and $f(x) > x$ if $a < x < 1$; in particular, $f(b) > b$. We have

$$\begin{aligned} B(I, J) &= \frac{|f(1) - f(0)| \cdot |f(b) - f(a)|}{|f(a) - f(0)| \cdot |f(1) - f(b)|} \cdot \left(\frac{|1 - 0| \cdot |b - a|}{|a - 0| \cdot |1 - b|} \right)^{-1} \\ &= \frac{1 \cdot (f(b) - a) \cdot a \cdot (1 - b)}{a \cdot (1 - f(b)) \cdot 1 \cdot (b - a)} > 1. \end{aligned}$$

This proves the second inequality. The first one is left as an exercise (Exercise 7.5.6). \square

The following proposition, which we do not prove, describes bounded distortion properties of maps with negative Schwarzian derivative on intervals without critical points.

PROPOSITION 7.5.6 [vS88], [dMvS93]. *Let $f: [a, b] \rightarrow \mathbb{R}$ be a C^3 map. Assume that $Sf < 0$ and $f'(x) \neq 0$, for all $x \in [a, b]$. Then*

1. $|f'(a)| \cdot |f'(b)| \geq (|f(b) - f(a)| / (b - a))^2$;
2. $\frac{|f'(x)| \cdot |f(b) - f(a)|}{b - a} \geq \frac{|f(x) - f(a)|}{x - a} \cdot \frac{|f(b) - f(x)|}{b - x}$ for every $x \in (a, b)$.

Exercise 7.5.1. Prove that if $f: I \rightarrow \mathbb{R}$ is a C^3 diffeomorphism onto its image and $g(x) = \frac{d}{dx} \log |f'(x)|$, then

$$Sf(x) = g'(x) - \frac{1}{2}(g(x))^2 = -2\sqrt{|f'(x)|} \cdot \frac{d^2}{dx^2} \frac{1}{\sqrt{|f'(x)|}}.$$

Exercise 7.5.2. Show that any polynomial with distinct real roots has negative Schwarzian derivative.

Exercise 7.5.3. Prove Lemma 7.5.1.

Exercise 7.5.4. Prove that each Möbius transformation has Schwarzian derivative 0 and preserves the cross-ratios C and D .

Exercise 7.5.5. Prove that the action of the group of Möbius transformations on the extended real line is simply transitive on triples of points.

Exercise 7.5.6. Prove the remaining inequality of Proposition 7.5.5.

7.6 Real Quadratic Maps

In §1.5, we introduced the one-parameter family of real quadratic maps $q_\mu(x) = \mu x(1 - x)$, $\mu \in \mathbb{R}$. We showed that for $\mu > 1$, the orbit of any point




Figure 7.5. Quadratic map.

outside $I = [0, 1]$ converges monotonically to $-\infty$. Thus the interesting dynamics is concentrated on the set

$$\Lambda_\mu = \{x \in I \mid q_\mu^n(x) \in I \ \forall n \geq 0\}.$$

THEOREM 7.6.1. *Let $\mu > 4$. Then Λ_μ is a Cantor set, i.e., a perfect, nowhere dense subset of $[0, 1]$. The restriction $q_\mu|_{\Lambda_\mu}$ is topologically conjugate to the one-sided shift $\sigma: \Sigma_2^+ \rightarrow \Sigma_2^+$.*

Proof. Let $a = 1/2 - \sqrt{1/4 - 1/\mu}$ and $b = 1/2 + \sqrt{1/4 - 1/\mu}$ be the two solutions of $q_\mu(x) = 1$, and let $I_0 = [0, a]$, $I_1 = [b, 1]$. Then $q_\mu(I_0) = q_\mu(I_1) = I$, and $q_\mu((a, b)) \cap I = \emptyset$ (see Figure 7.5). Observe that the images $q_\mu^n(1/2)$ of the critical point $1/2$ lie outside I and tend to $-\infty$. Therefore the two inverse branches $f_0: I \rightarrow I_0$ and $f_1: I \rightarrow I_1$ and their compositions are well defined. For $k \in \mathbb{N}$, denote by W_k the set of all words of length k in the alphabet $\{0, 1\}$. For $w = \omega_1\omega_2 \dots \omega_k \in W_k$ and $j \in \{0, 1\}$, set $I_{wj} = f_j(I_w)$ and $g_w = f_{\omega_k} \circ \dots \circ f_{\omega_2} \circ f_{\omega_1}$, so that $I_w = g_w(I)$.

LEMMA 7.6.2. $\lim_{k \rightarrow \infty} \max_{w \in W_k} \max_{x \in I} |g'_w(x)| = 0$.

Proof. If $\mu > 2 + \sqrt{5}$, then $1 > |f'_j(1)| = \mu\sqrt{1 - 4/\mu} \geq |f'_j(x)|$ for every $x \in I$, $j = 0, 1$, and the lemma follows.

For $4 < \mu < 2 + \sqrt{5}$, the lemma follows from Theorem 8.5.10 (see also Theorem 8.5.11). \square

Lemma 7.6.2 implies that the length of the interval I_w tends to 0 as the length of w tends to infinity. Therefore, for each $\omega = \omega_1\omega_2 \dots \in \Sigma_2^+$

the intersection $\bigcap_{n \in \mathbb{N}} I_{\omega_1 \dots \omega_n}$ consists of exactly one point $h(\omega)$. The map $h: \Sigma_2^+ \rightarrow \Lambda_\mu$ is a homeomorphism conjugating the shift σ and $q_\mu|_{\Lambda_\mu}$ (Exercise 7.6.2). \square

Exercise 7.6.1. Prove that if $\mu > 4$ and $1/2 - \sqrt{1/4 - 1/\mu} < x < 1/2 + \sqrt{1/4 - 1/\mu}$, then $q_\mu^n(x) \rightarrow -\infty$ as $n \rightarrow \infty$.

Exercise 7.6.2. Prove that the map $h: \Sigma_2^+ \rightarrow \Lambda_\mu$ in the proof of Theorem 7.6.1 is a homeomorphism and that $q_\mu \circ h = h \circ \sigma$.

7.7 Bifurcations of Periodic Points³

The family of real quadratic maps $q_\mu(x) = \mu x(1 - x)$ (§1.5, §7.6) is an example of a (one-dimensional) parametrized family of dynamical systems. Although the specific quantitative behavior of a dynamical system depends on the parameter, it is often the case that the qualitative behavior remains unchanged for certain ranges of the parameter. A parameter value where the qualitative behavior changes is called a *bifurcation* value of the parameter. For example, in the family of quadratic maps, the parameter value $\mu = 3$ is a bifurcation value because the stability of the fixed point $1 - 1/\mu$ changes from repelling to attracting. The parameter value $\mu = 1$ is a bifurcation value because for $\mu < 1$, 0 is the only fixed point, and for $\mu > 1$, q_μ has two fixed points.

A bifurcation is called *generic* if the same bifurcation occurs for all nearby families of dynamical systems, where “nearby” is defined with respect to an appropriate topology (usually the C^2 or C^3 topology). For example, the bifurcation value $\mu = 3$ is generic for the family of quadratic maps. To see this, note that for μ close to the 3; the graph of q_μ crosses the diagonal transversely at the fixed point $x_\mu = 1 - 1/\mu$, and the magnitude of $q'_\mu(x_\mu)$ is less than 1 for $\mu < 3$ and greater than 1 for $\mu > 3$. If f_μ is another family of maps C^1 -close to q_μ , then the graph of $f_\mu(x)$ must also cross the diagonal at a point y_μ near x_μ , and the magnitude of $f'_\mu(y_\mu)$ must cross 1 at some parameter value close to 3. Thus f_μ has the same kind of bifurcation as q_μ . Similar reasoning shows that the bifurcation value $\mu = 1$ is also generic.

Generic bifurcations are the primary ones of interest. The notion of genericity depends on the dimension of the parameter space (e.g., a bifurcation may be generic for a one-parameter family, but not for a two-parameter

³ The exposition in this section follows to a certain extent that of [Rob95].

family). Bifurcations that are generic for one-parameter families of dynamical systems are called *codimension-one* bifurcations. In this section, we describe codimension-one bifurcations of fixed and periodic points for one-dimensional maps.

We begin with a non-bifurcation result. If the graph of a differentiable map f intersects the diagonal transversely at a point x_0 , then the fixed point x_0 persists under a small C^1 perturbation of f .

PROPOSITION 7.7.1. *Let $U \subset \mathbb{R}^m$ and $V \subset \mathbb{R}^n$ be open subsets, and let $f_\mu: U \rightarrow \mathbb{R}^m$, $\mu \in V$, be a family of C^1 maps such that*

1. *the map $(x, \mu) \mapsto f_\mu(x)$ is a C^1 map,*
2. *$f_{\mu_0}(x_0) = x_0$ for some $x_0 \in U$ and $\mu_0 \in V$,*
3. *1 is not an eigenvalue of $df_{\mu_0}(x_0)$.*

Then there are open sets $U' \subset U$, $V' \subset V$ with $x_0 \in U'$, $\mu_0 \in V'$ and a C^1 function $\xi: V' \rightarrow U'$ such that for each $\mu \in V'$, $\xi(\mu)$ is the only fixed point of f_μ in U' .

Proof. The proposition is an immediate consequence of the implicit function theorem applied to the map $(x, \mu) \mapsto f_\mu(x) - x$ (Exercise 7.7.1). \square

Proposition 7.7.1 shows that if 1 is not an eigenvalue of the derivative, then the fixed point does not bifurcate into multiple fixed points and does not disappear. The next proposition shows that periodic points cannot appear in a neighborhood of a hyperbolic fixed point.

PROPOSITION 7.7.2. *Under the assumption (and notation) of Proposition 7.7.1, suppose in addition that x_0 is a hyperbolic fixed point of f_{μ_0} , i.e., no eigenvalue of $df_{\mu_0}(x_0)$ has absolute value 1. Then for each $k \in \mathbb{N}$ there are neighborhoods $U_k \subset U'$ of x_0 and $V_k \subset V'$ of μ_0 such that $\xi(\mu)$ is the only fixed point of f_μ^k in U_k .*

If, in addition, x_0 is an attracting fixed point of f_{μ_0} , i.e., all eigenvalues of $df_{\mu_0}(x_0)$ are strictly less than 1 in absolute value, then the neighborhoods U_k and V_k can be chosen independent of k .

Proof. Since no eigenvalue of $df_{\mu_0}(x_0)$ has absolute value 1, it follows that 1 is not an eigenvalue of $df_{\mu_0}^k(x_0)$, so the first statement follows from Proposition 7.7.1.

The second statement is left as an exercise (Exercise 7.7.2). \square

Propositions 7.7.1 and 7.7.2 show that, for differentiable one-dimensional maps, bifurcations of fixed or periodic points can occur only if the absolute value of the derivative is 1. For one-dimensional maps there are only

two types of generic bifurcations: The *saddle-node bifurcation* (or the *fold bifurcation*) may occur if the derivative at a periodic point is 1, and the *period-doubling bifurcation* (or *flip bifurcation*) may occur if the derivative at a periodic point is -1 . We describe these bifurcations in the next two propositions. See [CH82] or [HK91] for a more extensive discussion of bifurcation theory, or [GG73] for a thorough exposition on the closely related topic of singularities of differentiable maps.

PROPOSITION 7.7.3 (Saddle-Node Bifurcation). *Let $I, J \subset \mathbb{R}$ be open intervals and $f: I \times J \rightarrow \mathbb{R}$ be a C^2 map such that*

1. $f(x_0, \mu_0) = x_0$ and $\frac{\partial f}{\partial x}(x_0, \mu_0) = 1$ for some $x_0 \in I$ and $\mu_0 \in J$,
2. $\frac{\partial^2 f}{\partial x^2}(x_0, \mu_0) < 0$ and $\frac{\partial f}{\partial \mu}(x_0, \mu_0) > 0$.

Then there are $\epsilon, \delta > 0$ and a C^2 function $\alpha: (x_0 - \epsilon, x_0 + \epsilon) \rightarrow (\mu_0 - \delta, \mu_0 + \delta)$ such that:

1. $\alpha(x_0) = \mu_0$, $\alpha'(x_0) = 0$, $\alpha''(x_0) = -\frac{\partial^2 f}{\partial x^2}(x_0, \mu_0) / \frac{\partial f}{\partial \mu}(x_0, \mu_0) > 0$.
2. Each $x \in (x_0 - \epsilon, x_0 + \epsilon)$ is a fixed point of $f(\cdot, \alpha(x))$, i.e., $f(x, \alpha(x)) = x$, and $\alpha^{-1}(\mu)$ is exactly the fixed point set of $f(\cdot, \mu)$ in $(x_0 - \epsilon, x_0 + \epsilon)$ for $\mu \in (\mu_0 - \delta, \mu_0 + \delta)$.
3. For each $\mu \in (\mu_0, \mu_0 + \delta)$, there are exactly two fixed points $x_1(\mu) < x_2(\mu)$ of $f(\cdot, \mu)$ in $(x_0 - \epsilon, x_0 + \epsilon)$ with

$$\frac{\partial f}{\partial x}(x_1(\mu), \mu) > 1 \quad \text{and} \quad 0 < \frac{\partial f}{\partial x}(x_2(\mu), \mu) < 1;$$

$$\alpha(x_i(\mu)) = \mu \text{ for } i = 1, 2.$$

4. $f(\cdot, \mu)$ does not have fixed points in $(x_0 - \epsilon, x_0 + \epsilon)$ for each $\mu \in (\mu_0 - \delta, \mu_0)$.

REMARK 7.7.4. *The inequalities in the second hypothesis of Proposition 7.7.3 correspond to one of the four possible generic cases when the two derivatives do not vanish. The other three cases are similar (Exercise 7.7.3).*

Proof. Consider the function $g(x, \mu) = f(x, \mu) - x$ (see Figure 7.6). Observe that

$$\frac{\partial g}{\partial \mu}(x_0, \mu_0) = \frac{\partial f}{\partial \mu}(x_0, \mu_0) > 0.$$

Therefore, by the implicit function theorem, there are $\epsilon, \delta > 0$ and a C^2 function $\alpha: (x_0 - \epsilon, x_0 + \epsilon) \rightarrow J$ such that $g(x, \alpha(x)) = 0$ for each $x \in (x_0 - \epsilon, x_0 + \epsilon)$ and there are no other zeros of g in $(x_0 - \epsilon, x_0 + \epsilon) \times (\mu_0 - \epsilon, \mu_0 + \epsilon)$. A direct calculation shows that α satisfies statement 1. Since




Figure 7.6. Saddle-node bifurcation.

$\alpha''(x_0) > 0$, statements 3 and 4 are satisfied for ϵ and δ sufficiently small (Exercise 7.7.4). □

PROPOSITION 7.7.5 (Period-Doubling Bifurcation). *Let $I, J \subset \mathbb{R}$ be open intervals, and $f: I \times J \rightarrow \mathbb{R}$ be a C^3 map such that:*

1. $f(x_0, \mu_0) = x_0$ and $\frac{\partial f}{\partial x}(x_0, \mu_0) = -1$ for some $x_0 \in I$ and $\mu_0 \in J$, so that by Proposition 7.7.1, there is a curve $\mu \mapsto \xi(\mu)$ of fixed points of $f(\cdot, \mu)$ for μ close to μ_0 .
2. $\eta = \frac{d}{d\mu} \Big|_{\mu=\mu_0} \frac{\partial f}{\partial x}(\xi(\mu), \mu) < 0$.
3. $\zeta = \frac{\partial^3 f(f(x_0, \mu_0), \mu_0)}{\partial x^3} = -2 \frac{\partial^3 f}{\partial x^3}(x_0, \mu_0) - 3 \left(\frac{\partial^2 f}{\partial x^2}(x_0, \mu_0) \right)^2 < 0$.

Then there are $\epsilon, \delta > 0$ and C^3 functions $\xi: (\mu_0 - \delta, \mu_0 + \delta) \rightarrow \mathbb{R}$ with $\xi(\mu_0) = x_0$ and $\alpha: (x_0 - \epsilon, x_0 + \epsilon) \rightarrow \mathbb{R}$ with $\alpha(x_0) = \mu_0, \alpha'(x_0) = 0$, and $\alpha''(x_0) = -2\eta/\zeta > 0$ such that:

1. $f(\xi(\mu), \mu) = \xi(\mu)$, and $\xi(\mu)$ is the only fixed point of $f(\cdot, \mu)$ in $(x_0 - \epsilon, x_0 + \epsilon)$ for $\mu \in (\mu_0 - \delta, \mu_0 + \delta)$.
2. $\xi(\mu)$ is an attracting fixed point of $f(\cdot, \mu)$ for $\mu_0 - \delta < \mu < \mu_0$ and is a repelling fixed point for $\mu_0 < \mu < \mu_0 + \delta$.
3. For each $\mu \in (\mu_0, \mu_0 + \delta)$, the map $f(\cdot, \mu)$ has, in addition to the fixed point $\xi(\mu)$, exactly two attracting period-2 points $x_1(\mu), x_2(\mu)$ in the interval $(x_0 - \epsilon, x_0 + \epsilon)$; moreover, $\alpha(x_i(\mu)) = \mu$ and $x_i(\mu) \rightarrow x_0$ as $\mu \searrow \mu_0$ for $i = 1, 2$.
4. For each $\mu \in (\mu_0 - \delta, \mu_0]$, the map $f(f(\cdot, \mu), \mu)$ has exactly one fixed point $\xi(\mu)$ in $(x_0 - \epsilon, x_0 + \epsilon)$.

REMARK 7.7.6. *The stability of the fixed point $\xi(\mu)$ and of the periodic points $x_1(\mu)$ and $x_2(\mu)$ depend on the signs of the derivatives in the third and fourth hypotheses of Proposition 7.7.5. Proposition 7.7.5 deals with only one of the four possible generic cases when the derivatives do not vanish. The other three cases are similar, and we do not consider them here (Exercise 7.7.5).*

Proof. Since

$$\frac{\partial f}{\partial x}(x_0, \mu_0) = -1 \neq 1,$$

we can apply the implicit function theorem to $f(x, \mu) - x = 0$ to obtain a differentiable function ξ such that $f(\xi(\mu), \mu) = \xi(\mu)$ for μ close to μ_0 and $\xi(\mu_0) = x_0$. This proves statement 1.

Differentiating $f(\xi(\mu), \mu) = \xi(\mu)$ with respect to μ gives

$$\frac{d}{d\mu} f(\xi(\mu), \mu) = \frac{\partial f}{\partial \mu}(\xi(\mu), \mu) + \frac{\partial f}{\partial x}(\xi(\mu), \mu) \cdot \xi'(\mu) = \xi'(\mu),$$

and hence

$$\xi'(\mu) = \frac{\frac{\partial f}{\partial \mu}(\xi(\mu), \mu)}{1 - \frac{\partial f}{\partial x}(\xi(\mu), \mu)}, \quad \xi'(\mu_0) = \frac{1}{2} \frac{\partial f}{\partial \mu}(x_0, \mu_0).$$

Therefore

$$\left. \frac{d}{d\mu} \right|_{\mu=\mu_0} \frac{\partial f}{\partial x}(\xi(\mu), \mu) = \frac{\partial^2 f}{\partial \mu \partial x}(x_0, \mu_0) + \frac{1}{2} \frac{\partial f}{\partial \mu}(x_0, \mu_0) \cdot \frac{\partial^2 f}{\partial x^2}(x_0, \mu_0) = \eta,$$

and assumption 2 yields statement 2.

To prove statements 3 and 4 consider the change of variables $y = x - \xi(\mu)$, $0 = x_0 - \xi(\mu_0)$ and the function $g(y, \mu) = f(f(y + \xi(\mu), \mu), \mu) - \xi(\mu)$. Observe that fixed points of $f(f(\cdot, \mu), \mu)$ correspond to solutions of $g(y, \mu) = y$. Moreover,

$$g(0, \mu) \equiv 0, \quad \frac{\partial g}{\partial y}(0, \mu_0) = 1, \quad \frac{\partial^2 g}{\partial y^2}(0, \mu_0) = 0,$$

i.e., the graph of the second iterate of $f(\cdot, \mu_0)$ is tangent to the diagonal at (x_0, μ_0) with second derivative 0. (See Figure 7.7.) A direct calculation shows that, by assumption 3, the third derivative does not vanish:

$$\frac{\partial^3 g}{\partial y^3}(0, \mu_0) = -2 \frac{\partial^3 f}{\partial x^3}(x_0, \mu_0) - 3 \left(\frac{\partial f^2}{\partial x^2}(x_0, \mu_0) \right)^2 = \zeta < 0.$$

Therefore

$$g(y, \mu_0) = y + \frac{1}{3!} \zeta y^3 + o(y^3).$$

Since $\xi(\mu)$ is a fixed point of $f(\cdot, \mu)$, we have that $g(0, \mu) \equiv 0$ in an interval




Figure 7.7. Period-doubling bifurcation: the graph of the second iterate.

about μ_0 . Therefore there is a differentiable function h such that $g(y, \mu) = y \cdot h(y, \mu)$, and to find the period-2 points of $f(\cdot, \mu)$ different from $\xi(\mu)$ we must solve the equation $h(y, \mu) = 1$. From (7.7.6) we obtain

$$h(y, \mu_0) = 1 + \frac{1}{3!}\zeta y^2 + o(y^2),$$

i.e.,

$$h(0, \mu_0) = 1, \quad \frac{\partial h}{\partial y}(0, \mu_0) = 0, \quad \text{and} \quad \frac{\partial^2 h}{\partial y^2}(0, \mu_0) = \frac{\zeta}{3}.$$

On the other hand,

$$\begin{aligned} \frac{\partial h}{\partial \mu}(0, \mu_0) &= \lim_{y \rightarrow 0} \frac{1}{y} \frac{\partial g}{\partial \mu}(y, \mu_0) = \frac{\partial^2 g}{\partial \mu \partial y}(0, \mu_0) \\ &= \frac{d}{d\mu} \left(\frac{\partial f}{\partial x}(\xi(\mu), \mu) \right)^2 \Big|_{\mu=\mu_0} = -2\eta > 0. \end{aligned}$$

By the implicit function theorem, there is $\epsilon > 0$ and a differentiable function $\beta: (-\epsilon, \epsilon) \rightarrow \mathbb{R}$ such that $h(y, \beta(y)) = 1$ for $|y| < \epsilon$ and $\beta(0) = \mu_0$. Differentiating $h(y, \beta(y)) = 1$ with respect to y , we obtain that $\beta'(0) = 0$. The second differentiation yields $\beta''(0) = \zeta/6\eta > 0$. Therefore $\beta(y) > 0$ for $y \neq 0$, and the new period-2 orbit appears only for $\mu > \mu_0$.

Note that since $g(\cdot, \mu)$ has three fixed points near x_0 for μ close to μ_0 , and the middle one, $\xi(\mu)$, is unstable, the other two must be stable. In fact, a direct calculation shows that

$$\begin{aligned} \frac{\partial g}{\partial y}(y, \beta(y)) &= \frac{\partial g}{\partial y}(0, \mu_0) + \frac{1}{2!} \frac{\partial^2 g}{\partial y^2}(0, \mu_0)y + \frac{1}{3!} \frac{\partial^3 g}{\partial y^3}(0, \mu_0)y^2 + o(y^2) \\ &= 1 + \frac{\zeta}{6}y^2 + o(y^2). \end{aligned}$$

Since $\zeta < 0$, the period-2 orbit is stable. □

Exercise 7.7.1. Prove Proposition 7.7.1.

Exercise 7.7.2. Prove the second statement of Proposition 7.7.2.

Exercise 7.7.3. State the analog of Proposition 7.7.3 for the remaining three generic cases when the derivatives from assumption 3 do not vanish.

Exercise 7.7.4. Prove statements 3 and 4 of Proposition 7.7.3.

Exercise 7.7.5. State the analog of Proposition 7.7.5 for the remaining three generic cases when the derivatives from assumptions 3 and 4 do not vanish.

Exercise 7.7.6. Prove that a period-doubling bifurcation occurs for the family $f_\mu(x) = 1 - \mu x^2$ at $\mu_0 = 3/4$, $x_0 = 2/3$.

7.8 The Feigenbaum Phenomenon

M. Feigenbaum [Fei79] studied the family

$$f_\mu(x) = 1 - \mu x^2, \quad 0 < \mu \leq 2,$$

of unimodal maps of the interval $[-1, 1]$. For $\mu < 3/4$, the unique attracting fixed point of f_μ is

$$x_\mu = \frac{\sqrt{1 + 4\mu} - 1}{2\mu}.$$

The derivative $f'_\mu(x_\mu) = 1 - \sqrt{1 + 4\mu}$ is greater than -1 for $\mu < 3/4$, equals -1 for $\mu = 3/4$, and is less than -1 for $\mu > 3/4$. A period-doubling bifurcation occurs at $\mu = 3/4$ (Exercise 7.7.6). For $\mu > 3/4$, the map f_μ has an attracting period-2 orbit. Numerical studies show that there is an increasing sequence of bifurcation values μ_n at which an attracting periodic orbit of period 2^n for f_μ loses stability and an attracting periodic orbit of period 2^{n+1} is born. The sequence μ_n converges, as $n \rightarrow \infty$, to a limit μ_∞ , and

$$\lim_{n \rightarrow \infty} \frac{\mu_\infty - \mu_{n-1}}{\mu_\infty - \mu_n} = \delta = 4.669201609 \dots \quad (7.2)$$

The constant δ is called the *Feigenbaum constant*. Numerical experiments show that the Feigenbaum constant appears for many other one-parameter families.

The Feigenbaum phenomenon can be explained as follows. Consider the infinite-dimensional space \mathcal{A} of real analytic maps $\psi: [-1, 1] \rightarrow [-1, 1]$ with $\psi(0) = 1$, and the map $\Phi: \mathcal{A} \rightarrow \mathcal{A}$ given by the formula

$$\Phi(\psi)(x) = \frac{1}{\lambda} \psi \circ \psi(\lambda x), \quad \lambda = \psi(1). \quad (7.3)$$

A fixed point g of Φ (which Feigenbaum estimated numerically) is an even function satisfying the Cvitanović–Feigenbaum equation

$$g \circ g(\lambda x) - \lambda g(x) = 0. \quad (7.4)$$




Figure 7.8. Fixed point and stable and unstable manifolds for the Feigenbaum map Φ .

The function g is a hyperbolic fixed point of Φ . The stable manifold $W^s(g)$ has codimension one, and the unstable manifold $W^u(g)$ has dimension one and corresponds to a simple eigenvalue $\delta = 4.669201609\dots$ of the derivative $d\Phi_g$. The codimension-one bifurcation set B_1 , of maps ψ for which an attracting fixed point loses stability and an attracting period two orbit is born, intersects $W^u(g)$ transversely. The preimage $B_n = \Phi^{1-n}(B_1)$ is the bifurcation set of maps for which an attracting orbit of period 2^{n-1} is replaced by an attracting orbit of period 2^n (Exercise 7.8.1). Figure 7.8 is a graphical depiction of the process underlying the Feigenbaum phenomenon.

By the infinite-dimensional version of the Inclination Lemma 5.7.2, the codimension-one bifurcation sets B_n accumulate to $W^s(g)$. Let f_μ be a one-parameter family of maps that intersects $W^s(g)$ transversely, and let μ_n be the sequence of period-doubling bifurcation parameters, $f_{\mu_n} \in B_n$. Using the inclination lemma, one can show that the sequence μ_n satisfies (7.2). O. E. Lanford established the correctness of this model through a computer-assisted proof [Lan84].

Exercise 7.8.1. Prove that if ψ has an attracting periodic orbit of period $2k$, then $\Phi(\psi)$ has an attracting periodic orbit of period k .

Complex Dynamics

In this chapter¹, we consider *rational maps* $R(z) = P(z)/Q(z)$ of the *Riemann sphere* $\bar{\mathbb{C}} = \mathbb{C} \cup \{\infty\}$, where P and Q are complex polynomials. These maps exhibit many interesting dynamical properties, and lend themselves to the computer-aided drawing of fractals and other fascinating pictures in the complex plane. For a more thorough exposition of the dynamics of rational maps see [Bea91] and [CG93].

8.1 Complex Analysis on the Riemann Sphere

We assume that the reader is familiar with the basic ideas of complex analysis (see, for example, [BG91] or [Con95]).

Recall that a function f from a domain $D \subset \mathbb{C}$ to $\bar{\mathbb{C}}$ is said to be *meromorphic* if it is analytic except at a discrete set of singularities, all of which are poles. In particular, rational functions are meromorphic.

The Riemann sphere is the one-point compactification of the complex plane, $\bar{\mathbb{C}} = \mathbb{C} \cup \{\infty\}$. The space $\bar{\mathbb{C}}$ has the structure of a complex manifold, given by the standard coordinate system on \mathbb{C} and the coordinate $z \mapsto z^{-1}$ on $\bar{\mathbb{C}} \setminus \{0\}$. If M and N are complex manifolds, then a map $f: M \rightarrow N$ is *analytic* if for every point $\zeta \in M$, there are complex coordinate neighborhoods U of ζ and V of $f(\zeta)$ such that $f: U \rightarrow V$ is analytic in the coordinates on U and V . An analytic map into $\bar{\mathbb{C}}$ is said to be *meromorphic*. This terminology is somewhat confusing, because in the modern sense (as maps of manifolds) meromorphic functions are analytic, while in the classical sense (as functions on \mathbb{C}), meromorphic functions are generally not analytic. Nevertheless, the terminology is so entrenched that it cannot be avoided.

¹ Many of the proofs in this chapter follow the corresponding arguments from [CG93].

It is easy to see that a map $f: \bar{\mathbb{C}} \rightarrow \bar{\mathbb{C}}$ is analytic (and meromorphic) if and only if both $f(z)$ and $f(1/z)$ are meromorphic (in the classical sense) on \mathbb{C} . It is known that every analytic map from the Riemann sphere to itself is a rational map. Note that the constant map $f(z) = \infty$ is considered to be analytic.

The group of Möbius transformations

$$\left\{ z \rightarrow \frac{az + b}{cz + d} : a, b, c, d \in \mathbb{C}; ad - bc = 1 \right\}$$

acts on the Riemann sphere and is *simply transitive* on triples of points, i.e., for any three distinct points $x, y, z \in \bar{\mathbb{C}}$, there is a unique Möbius transformation that carries x, y, z to $0, 1, \infty$, respectively (see §7.5).

Suppose $f: \bar{\mathbb{C}} \rightarrow \bar{\mathbb{C}}$ is a meromorphic map and ζ is a periodic point of minimal period k . If $\zeta \neq \infty$, the *multiplier* of ζ is the derivative $\lambda(\zeta) = (f^k)'(\zeta)$. If $\zeta = \infty$, the multiplier of ζ is $g'(0)$, where $g(z) = 1/f(1/z)$. The periodic point ζ is *attracting* if $0 < |\lambda(\zeta)| < 1$, *superattracting* if $\lambda(\zeta) = 0$, *repelling* if $|\lambda(\zeta)| > 1$, *rationally neutral* if $\lambda(\zeta)^m = 1$ for some $m \in \mathbb{N}$, and *irrationally neutral* if $|\lambda(\zeta)| = 1$ but $\lambda(\zeta)^m \neq 1$ for every $m \in \mathbb{N}$. One can prove that a periodic point is attracting or superattracting if and only if it is a topologically attracting periodic point in the sense of Chapter 1; similarly for repelling periodic points. The orbit of an attracting or superattracting periodic point is said to be an attracting or superattracting periodic orbit, respectively.

For an attracting or superattracting fixed point ζ of a meromorphic map f , we define the *basin of attraction* $\text{BA}(\zeta)$ as the set of points $z \in \bar{\mathbb{C}}$ for which $f^n(z) \rightarrow \zeta$ as $n \rightarrow \infty$. Since the multiplier of ζ is less than 1, there is a neighborhood U of ζ that is contained in $\text{BA}(\zeta)$, and $\text{BA}(\zeta) = \bigcup_{n \in \mathbb{N}} f^{-n}(U)$. The set $\text{BA}(\zeta)$ is open. The connected component of $\text{BA}(\zeta)$ containing ζ is called the *immediate basin of attraction*, and is denoted $\text{BA}^\circ(\zeta)$.

If ζ is an attracting or superattracting periodic point of period k , then the basin of attraction of the periodic orbit is the set of all points z for which $f^{nk}(z) \rightarrow f^j(\zeta)$ as $n \rightarrow \infty$ for some $j \in \{0, 1, \dots, k\}$ and is denoted $\text{BA}(\zeta)$. The union of the connected components of $\text{BA}(\zeta)$ containing a point in the orbit of ζ is called the *immediate basin of attraction* and is denoted $\text{BA}^\circ(\zeta)$.

A point ζ is a *critical point* (or *branch point*) of a meromorphic function f if f is not 1-to-1 on a neighborhood of ζ . A critical point ζ has *multiplicity* m if f is $(m+1)$ -to-1 on $U \setminus \{\zeta\}$ for a sufficiently small neighborhood U of ζ (this number is also called the *branch number* of f at ζ). Equivalently, ζ is a critical point of multiplicity m if ζ is a zero of f' (in local coordinates) of multiplicity m . If ζ is a critical point, then $f(\zeta)$ is called a *critical value*.

For a rational map $R = P/Q$, with P and Q relatively prime polynomials of degree p and q , respectively, the *degree* of R is $\deg(R) = \max(p, q)$. If R has degree d , then the map $R: \bar{\mathbb{C}} \rightarrow \bar{\mathbb{C}}$ is a *branched covering of degree d* , i.e., any $\xi \in \bar{\mathbb{C}}$ that is not a critical value has exactly d preimages; in fact, every point has exactly d preimages if critical points are counted with multiplicity. Since the number of preimages of a generic point is a topological invariant of R , the degree is invariant under conjugation by a Möbius transformation.

The rational maps of degree 1 are the Möbius transformations. A rational map is a polynomial if and only if the only preimage of ∞ is ∞ .

PROPOSITION 8.1.1. *Let R be a rational map of degree d . Then the number of critical points, counted with multiplicity, is $2d - 2$. If there are exactly two distinct critical points, then R is conjugate by a Möbius transformation to z^d or z^{-d} .*

Proof. By composing with a Möbius transformation we may assume that $R(\infty) = 0$ and that ∞ is neither a critical point nor a critical value. Then $R(\infty) = 0$ and the fact that ∞ is not a critical point imply that

$$R(z) = \frac{\alpha z^{d-1} + \dots}{\beta z^d + \dots},$$

where $\alpha \neq 0$ and $\beta \neq 0$. Hence

$$R'(z) = -\frac{\alpha\beta z^{2d-2} + \dots}{(\beta z^d + \dots)^2},$$

and the critical points of R are the zeros of the numerator (since ∞ is not a critical value).

The proof of the second assertion is left as an exercise (Exercise 8.1.5). \square

A family F of meromorphic functions in a domain $D \subset \bar{\mathbb{C}}$ is *normal* if every sequence from F contains a subsequence that converges uniformly on compact subsets of D in the standard spherical metric on $\bar{\mathbb{C}} \approx S^2$. A family F is *normal at a point* $z \in \bar{\mathbb{C}}$ if it is normal in a neighborhood of z .

The *Fatou set* $F(R) \subset \bar{\mathbb{C}}$ of a rational map $R: \bar{\mathbb{C}} \rightarrow \bar{\mathbb{C}}$ is the set of points $z \in \bar{\mathbb{C}}$ such that the family of forward iterates $\{R^n\}_{n \in \mathbb{N}}$ is normal at z . The *Julia set* $J(R)$ is the complement of the Fatou set. Both $F(R)$ and $J(R)$ are completely invariant under R (see Proposition 8.5.1). Points belonging to the same component of $F(R)$ have the same asymptotic behavior. As we will see later, the Fatou set contains all basins of attraction and the Julia set is the closure of the set of all repelling periodic points. The “interesting” dynamics is concentrated on the Julia set, which is often a fractal set. The case when $J(R)$ is a hyperbolic set is reasonably well understood (Theorem 8.5.10).

Exercise 8.1.1. Prove that any Möbius transformation is conjugate by another Möbius transformation to either $z \mapsto az$ or $z \mapsto z + a$.

Exercise 8.1.2. Prove that a non-constant rational map R is conjugate to a polynomial by a Möbius transformation if and only if $R^{-1}(z_0) = \{z_0\}$ for some $z_0 \in \bar{\mathbb{C}}$.

Exercise 8.1.3. Find all Möbius transformations that commute with $q_0(z) = z^2$.

Exercise 8.1.4. Let R be a rational map such that $R(\infty) = \infty$, and let f be a Möbius transformation such that $f(\infty)$ is finite. Define the multiplier $\lambda_R(\infty)$ of R at ∞ to be the multiplier of $f \circ R \circ f^{-1}$ at $f(\infty)$. Prove that $\lambda_R(\infty)$ does not depend on the choice of f .

Exercise 8.1.5. Prove the second assertion of Proposition 8.1.1.

Exercise 8.1.6. Let R be a non-constant rational map. Prove that

$$\deg(R) - 1 \leq \deg(R') \leq 2 \deg(R)$$

with equality on the left if and only if R is a polynomial and with equality on the right if and only if all poles of R are simple and finite.

8.2 Examples

The global dynamics of a rational map R depends heavily on the behavior of the critical points of R under its iterates. In most of the examples below, the Fatou set consists of finitely many components, each of which is a basin of attraction. Some of the assertions in the following examples will be proved in later sections of this chapter. Proofs of most of the assertions that are not proved here can be found in [CG93].

Let $q_a: \bar{\mathbb{C}} \rightarrow \bar{\mathbb{C}}$ be the quadratic map $q_a(z) = z^2 - a$, and denote by S^1 the unit circle $\{z \in \mathbb{C}: |z| = 1\}$. The critical points of q_a are 0 and ∞ , and the critical values are $-a$ and ∞ ; if $a \neq 0$, the only superattracting periodic (fixed) point is ∞ . In the examples below, we observe drastically different global dynamics depending on whether the critical point lies in the basin of a finite attracting periodic point, or in the basin of ∞ , or in the Julia set.

1. $q_0(z) = z^2$. There is a superattracting fixed point at 0, whose basin of attraction is the open disk $\Delta_1 = \{z \in \mathbb{C}: |z| < 1\}$, and a superattracting fixed point at ∞ , whose basin of attraction is the exterior of S^1 . There is also a repelling fixed point at 1, and for each $n \in \mathbb{N}$ there are 2^n repelling periodic points of period n on S^1 . The Julia set is S^1 ; the Fatou set is the complement of S^1 . The map q_0 acts on S^1 by




Figure 8.1. The Julia set for $a = 1$.

- $\phi \mapsto 2\phi \bmod 2\pi$ (where ϕ is the angular coordinate of a point $z \in S^1$). If U is a neighborhood of $\zeta \in S^1 = J(q_0)$, then $\bigcup_{n \in \mathbb{N}_0} q_0^n(U) = \mathbb{C} \setminus \{0\}$.
- $q_\epsilon(z) = z^2 - \epsilon$, $0 < \epsilon \ll 1$. There is an attracting fixed point near 0, a superattracting fixed point at ∞ , and, for each $n \in \mathbb{N}$, 2^n repelling periodic points near S^1 . The Julia set $J(q_\epsilon)$ is a closed continuous q_ϵ -invariant curve that is C^0 close to S^1 and is not differentiable at a dense set of points; it has a Hausdorff dimension greater than 1. The basins of attraction of the fixed points near 0 and at ∞ are, respectively, the interior and exterior of $J(q_\epsilon)$. The critical point and critical value lie in the immediate basin of attraction of the attracting fixed point near 0. The same properties hold true for maps of the form $f(z) = z^2 + \epsilon P(z)$, where P is a polynomial and ϵ is small enough.
 - $q_1(z) = z^2 - 1$. Note that $q_1(0) = -1$, $q_1(-1) = 0$. Therefore, 0 and -1 are superattracting periodic points of period 2. On the real line the repelling fixed point $(1 - \sqrt{5})/2$ separates the basins of attraction of 0 and -1 . The Julia set $J(q_1)$ contains two simple closed curves σ_0 and σ_{-1} that surround 0 and -1 and bound their immediate basins of attraction. The only preimage of -1 is 0; hence the only preimage of σ_{-1} is σ_0 . However, 0 has two preimages, $+1$ and -1 . Therefore σ_0 has two preimages, σ_{-1} and a closed curve σ_1 surrounding 1. Continuing in this manner and using the complete invariance of the Julia set (Proposition 8.5.1), we conclude that $J(q_1)$ contains infinitely many closed curves. Their interiors are components of the Fatou set. Figure 8.1 shows the Julia set for q_1 .²

² The pictures in this chapter were produced with *mandelspawm*; see <http://www.araneus.fi/gson/mandelspawm/>.




Figure 8.2. The Julia set for $a = -i$.

4. $q_{-i} = z^2 + i$. The critical point 0 is eventually periodic: $q_{-i}^2(0) = i - 1$, and $i - 1$ is a repelling periodic point of period 2. The only attracting periodic fixed point is ∞ . The Fatou set consists of one component and coincides with $\text{BA}(\infty)$. The Julia set is a *dendrite*, i.e., a compact, path-connected, locally connected, nowhere dense subset of \mathbb{C} that does not separate \mathbb{C} . Figure 8.2 shows the Julia set for q_{-i} .
5. $q_2(z) = z^2 - 2$. The change of variables $z = \zeta + \zeta^{-1}$ conjugates q_2 on $\bar{\mathbb{C}} \setminus [-2, 2]$ with $\zeta \mapsto \zeta^2$ on the exterior of S^1 . Hence $J(q_2) = [-2, 2]$, and $F(q_2) = \bar{\mathbb{C}} \setminus [-2, 2]$ is the basin of attraction of ∞ . The image of the critical point 0 is $-2 \in J(q_2)$. The change of variables $y = (2 - x)/4$ conjugates the action of q_2 on the real axis to $y \mapsto 4y(1 - y)$. The only attracting periodic point is ∞ .
6. $q_4(z) = z^2 - 4$. The only attracting periodic point is ∞ ; the critical value -4 lies in the (immediate) basin of attraction of ∞ ; and $J(q_4)$ is a Cantor set on the real axis; $\text{BA}(\infty)$ is the complement of $J(q_4)$.
7. This example illustrates the connection between the dynamics of rational maps and issues of convergence for the *Newton method*. Let $Q(z) = (z - a)(z - b)$ with $a \neq b$. To find the roots a and b using the Newton method one iterates the map

$$f(z) = z - \frac{Q(z)}{Q'(z)} = z - \frac{1}{\frac{1}{z-a} + \frac{1}{z-b}}.$$

The change of variables $\zeta = (z - a)/(z - b)$ sends a to 0, b to ∞ , ∞ to 1, and the line $l = \{(a + b)/2 + ti(a - b) : t \in \mathbb{R}\}$ to the unit circle, and conjugates f with $\zeta \mapsto \zeta^2$. Therefore the Newton method for Q converges to a or b if the initial point lies in the half plane of l

containing a or b , respectively; the Newton method diverges if the initial point lies on l .

Exercise 8.2.1. Prove the properties of q_0 described above.

Exercise 8.2.2. Let U be a neighborhood of a point $z \in S^1$. Prove that $\bigcup_{n \in \mathbb{N}} q_0(U) = \mathbb{C} \setminus \{0\}$.

Exercise 8.2.3. Check the above conjugacies for q_2 .

Exercise 8.2.4. Prove that ∞ is the only attracting periodic point of q_4 .

Exercise 8.2.5. Let $|a| > 2$ and $|z| \geq |a|$. Prove that $q_a^n(z) \rightarrow \infty$ as $n \rightarrow \infty$.

Exercise 8.2.6. Prove the statements in example 7.

8.3 Normal Families

The theory of normal families of meromorphic functions is a keystone in the study of complex dynamics. The principal result, Theorem 8.3.2, is due to P. Montel [Mon27].

PROPOSITION 8.3.1. *Suppose F is a family of analytic functions in a domain D , and suppose that for every compact subset $K \subset D$ there is $C(K) > 0$ such that $|f(z)| < C(K)$ for all $z \in K$ and $f \in F$. Then F is a normal family.*

Proof. Let $\delta = \frac{1}{2} \min_{z \in K} \text{dist}(z, \partial D)$. By the Cauchy formula,

$$f'(z) = \frac{1}{2\pi} \int_{\gamma} \frac{f(\xi)}{(\xi - z)^2} d\xi$$

for any smooth closed curve γ in D that contains z in its interior. Let $K \subset D$ be compact, K_δ be the closure of the δ -neighborhood of K , and γ be the circle of radius δ centered at z . Then $|f'(z)| < C(K_\delta)/\delta$ for every $f \in F$ and $z \in K$. Thus the family F is equicontinuous on K , and therefore normal by the Arzela–Ascoli theorem. \square

We say that a family F of functions on a domain D *omits* a point a if $f(z) \neq a$ for every $f \in F$ and $z \in D$.

THEOREM 8.3.2 (Montel). *Suppose that a family F of meromorphic functions in a domain $D \subset \mathbb{C}$ omits three distinct points $a, b, c \in \mathbb{C}$. Then F is normal in D .*

Proof. Since D is covered by disks, we may assume without loss of generality that D is a disk. By applying a Möbius transformation, we may also assume that $a = 0$, $b = 1$, and $c = \infty$. Let Δ_1 be the unit disk. By the uniformization

theorem [Ahl73], there is an analytic covering map $\phi: \Delta_1 \rightarrow (\mathbb{C} \setminus \{0, 1\})$ (ϕ is called the *modular function*). For every function $f: D \rightarrow (\mathbb{C} \setminus \{0, 1\})$ there is a lift $\tilde{f}: D \rightarrow \Delta_1$ such that $\phi \circ \tilde{f} = f$. The family $\tilde{F} = \{\tilde{f}: f \in F\}$ is bounded and therefore, by Proposition 8.3.1, normal. The normality of F follows immediately. \square

Exercise 8.3.1. Let f be a meromorphic map defined on a domain $D \subset \bar{\mathbb{C}}$, and let $k > 1$. Show that the family $\{f^n\}_{n \in \mathbb{N}}$ is normal on D if and only if the family $\{f^{kn}\}_{n \in \mathbb{N}}$ is normal.

8.4 Periodic Points

THEOREM 8.4.1. Let ζ be an attracting fixed point of a meromorphic map $f: \bar{\mathbb{C}} \rightarrow \bar{\mathbb{C}}$. Then there is a neighborhood $U \ni \zeta$ and an analytic map $\phi: U \rightarrow \mathbb{C}$ that conjugates f and $z \mapsto \lambda(\zeta)z$, i.e., $\phi(f(z)) = \lambda(\zeta)\phi(z)$ for all $z \in U$.

Proof. We abbreviate $\lambda(\zeta) = \lambda$. Conjugating by a translation (or by $z \mapsto 1/z$ if $\zeta = \infty$), we replace ζ by 0. Then on any sufficiently small neighborhood of 0, say $\Delta_{1/2} = \{z: |z| < 1/2\}$, there is a $C > 0$ such that $|f(z) - \lambda z| \leq C|z|^2$. Hence for every $\epsilon > 0$ there is a neighborhood U of 0 such that $|f(z)| < (|\lambda| + \epsilon)|z|$, for all $z \in U$, and, assuming that $|\lambda| + \epsilon < 1$,

$$|f^n(z)| < (|\lambda| + \epsilon)^n |z|.$$

Set $\phi_n(z) = \lambda^{-n} f^n(z)$. Then, for $z \in U$,

$$|\phi_{n+1}(z) - \phi_n(z)| = \left| \frac{f(f^n(z)) - \lambda f^n(z)}{\lambda^{n+1}} \right| \leq \frac{C(|\lambda| + \epsilon)^{2n} |z|^2}{|\lambda|^{n+1}},$$

and hence the sequence ϕ_n converges uniformly in U if $(|\lambda| + \epsilon)^2 < |\lambda|$.

By construction, $\phi_n(f(z)) = \lambda \phi_{n+1}(z)$. Therefore the limit $\phi = \lim_{n \rightarrow \infty} \phi_n$ is the required conjugation. \square

COROLLARY 8.4.2. Let ζ be a repelling fixed point of a meromorphic map f . Then there is a neighborhood U of ζ and an analytic map $\phi: U \rightarrow \bar{\mathbb{C}}$ that conjugates f and $z \mapsto \lambda(\zeta)z$ in U , i.e., $\phi(f(z)) = \lambda(\zeta)\phi(z)$ for $z \in U$.

Proof. Apply Theorem 8.4.1 to the branch g of f^{-1} with $g(\zeta) = \zeta$. \square

PROPOSITION 8.4.3. Let ζ be a fixed point of a meromorphic map f . Assume that $\lambda = f'(\zeta)$ is not 0 and is not a root of 1, and suppose that an analytic map ϕ conjugates f and $z \mapsto \lambda z$. Then ϕ is unique up to multiplication by a constant.

Proof. Again, we assume $\zeta = 0$. If there are two conjugating maps ϕ and ψ , then $\eta = \phi^{-1} \circ \psi$ conjugates $z \mapsto \lambda z$ with itself, i.e., $\eta(\lambda z) = \lambda \eta(z)$. If $\eta = a_1 z + a_2 z^2 + \cdots$, then $a_n \lambda^n = \lambda a_n$ and $a_n = 0$ for $n > 1$. \square

LEMMA 8.4.4. *Any rational map R of degree > 1 has infinitely many periodic points.*

Proof. Observe that the number of solutions of $R^n(z) - z = 0$ (counted with multiplicity) tends to ∞ as $n \rightarrow \infty$. Therefore, if R has only finitely many periodic points, their multiplicities cannot be bounded in n .

On the other hand, if ζ is a multiple root of $R^n(z) - z = 0$, then $(R^n)'(\zeta) = 1$ and $R^n(z) = \zeta + (z - \zeta) + a(z - \zeta)^m + \cdots$ for some $a \neq 0$ and $m \geq 2$. By induction, $R^{nk}(z) = \zeta + (z - \zeta) + ka(z - \zeta)^m + \cdots$ for $k \in \mathbb{N}$. Therefore, ζ has the same multiplicity as a fixed point of R^n and as a fixed point of R^{nk} . \square

PROPOSITION 8.4.5. *Let f be a meromorphic map of $\bar{\mathbb{C}}$. If ζ is an attracting or superattracting periodic point of f , then the family $\{f^n\}_{n \geq 0}$ is normal in $\text{BA}(\zeta)$.*

If ζ is a repelling periodic point of f , then the family $\{f^n\}$ is not normal at ζ .

Proof. Exercise 8.4.1. \square

THEOREM 8.4.6. *Let ζ be an attracting periodic point of a rational map R . Then the immediate basin of attraction $\text{BA}^\circ(\zeta)$ contains a critical point of f .*

Proof. Consider first the case when ζ is a fixed point. Suppose that $\text{BA}^\circ(\zeta)$ does not contain a critical point. For a small enough $\epsilon > 0$, there is a branch g of R^{-1} that is defined in the open ϵ -disk D_ϵ about ζ and satisfies $g(\zeta) = \zeta$. The map $g: D_\epsilon \rightarrow \text{BA}^\circ(\zeta)$ is a diffeomorphism onto its image, and therefore $g(D_\epsilon)$ is simply connected and does not contain a critical point. Thus g extends uniquely to a map on $g(D_\epsilon)$. By induction, g extends uniquely to $g^n(D_\epsilon)$, which is a simply connected subset of $\text{BA}^\circ(\zeta)$. The sequence $\{g^n\}$ is normal on D_ϵ , since it omits infinitely many periodic points of R different from ζ (Lemma 8.4.4). (Note that if R is a polynomial, then $\{g^n\}$ omits a neighborhood of ∞ , and Lemma 8.4.4 is not needed.) On the other hand, $|g'(\zeta)| > 1$ and hence $(g^n)'(\zeta) \rightarrow \infty$ as $n \rightarrow \infty$, and therefore the family $\{g^n\}$ is not normal (Proposition 8.4.5); a contradiction.

If ζ is a periodic point of period n , then the preceding argument shows that the immediate basin of attraction of ζ for the map R^n contains a critical point of R^n . Since the components of $\text{BA}^\circ(\zeta)$ are permuted by R , it follows from the chain rule that one of the components contains a critical point of R . \square

COROLLARY 8.4.7. *A rational map has at most $2d - 2$ attracting and super-attracting periodic orbits.*

Proof. The corollary follows immediately from Theorem 8.4.6 and Proposition 8.1.1. \square

More delicate analysis that is beyond the scope of this book leads to the following theorem.

THEOREM 8.4.8 (Shishikura [Shi87]). *The total number of attracting, super-attracting, and neutral periodic orbits of a rational map of degree d is at most $2d - 2$.*

The upper bound $6d - 6$ was obtained by P. Fatou.

Exercise 8.4.1. Prove Proposition 8.4.5.

Exercise 8.4.2. Let $D \subset \bar{\mathbb{C}}$ be a domain whose complement contains at least three points, and let $f: D \rightarrow D$ be a meromorphic map with an attracting fixed point $z_0 \in D$. Prove that the sequence of iterates f^n converges in D to z_0 uniformly on compact sets.

Exercise 8.4.3. Prove that every rational map $R \neq \text{Id}$ of degree $d \geq 1$ has $d + 1$ fixed points in $\bar{\mathbb{C}}$ counted with multiplicity.

8.5 The Julia Set

Recall that the Fatou set $F(R)$ of a rational map R is the set of points $z \in \bar{\mathbb{C}}$ such that the family of forward iterates R^n , $n \in \mathbb{N}$, is normal at z . The Julia set $J(R)$ is the complement of $F(R)$. The Julia set of a rational map is closed by definition, and non-empty by Lemma 8.4.4, Proposition 8.4.5, and Theorem 8.4.8. If U is a connected component of $F(R)$, then $R(U)$ is also a connected component of $F(R)$ (Exercise 8.5.1).

Suppose $V \neq \text{BA}(\infty)$ is a component of $\text{BA}(\infty)$. Then $R^n(V) \subset \text{BA}^\circ(\infty)$ for some $n > 0$. Moreover, $R^n(V)$ is both open and closed in $\text{BA}^\circ(\infty)$, since $R^n(V) = R^n(V \cup J(R)) \setminus J(R)$. It follows that $R^n(V) = \text{BA}^\circ(\infty)$.

PROPOSITION 8.5.1. *Let $R: \bar{\mathbb{C}} \rightarrow \bar{\mathbb{C}}$ be a rational map. Then $F(R)$ and $J(R)$ are completely invariant, i.e., $R^{-1}(F(R)) = F(R)$ and $R(J(R)) = J(R)$, and similarly for $J(R)$.*

Proof. Let $\zeta = R(\xi)$. Then R^{nk} converges in a neighborhood of ζ if and only if R^{nk+1} converges in a neighborhood of ξ . \square

PROPOSITION 8.5.2. *Let $R: \bar{\mathbb{C}} \rightarrow \bar{\mathbb{C}}$ be a rational map. Then either $J(R) = \bar{\mathbb{C}}$ or $J(R)$ has no interior.*

Proof. Suppose $U \subset J(R)$ is non-empty and open in $\bar{\mathbb{C}}$. Then the family $\{R^n\}_{n \in \mathbb{N}}$ is not normal on U and, in particular, by Theorem 8.3.2, $\bigcup_n R^n(U)$ omits at most two points in $\bar{\mathbb{C}}$. Since $J(R)$ is invariant and closed, $J(R) = \bar{\mathbb{C}}$. \square

Let $R: \bar{\mathbb{C}} \rightarrow \bar{\mathbb{C}}$ be a rational map, and U an open set such that $U \cap J(R) \neq \emptyset$. The family of iterates $\{R^n\}_{n \in \mathbb{N}_0}$ is not normal in U , so it omits at most two points in $\bar{\mathbb{C}}$. The set E_U of omitted points is called the exceptional set of R on U . The exceptional set of R is the set $E = \bigcup E_U$, where the union is over all open sets U with $U \cap J(R) \neq \emptyset$. A point in E is called an *exceptional point* of R .

PROPOSITION 8.5.3. *Let R be a rational map of degree greater than 1. Then the exceptional set of R contains at most two points. If the exceptional set consists of a single point, then R is conjugate by a Möbius transformation to a polynomial. If it consists of two points, then R is conjugate by a Möbius transformation to z^m or $1/z^m$, for some $m > 1$. The exceptional set is disjoint from $J(R)$.*

Proof. If E_U is empty for every U with $U \cap J(R) \neq \emptyset$, there is nothing to show.

Suppose $\{R^n\}_{n \in \mathbb{N}_0}$ omits two points z_0, z_1 on U for some open set U with $U \cap J(R) \neq \emptyset$. Then after conjugating by the rational map $\phi(z) = (z - z_1)/(z - z_0)$, R becomes a rational map whose family of iterates omits only 0 and ∞ on the set $\phi(U)$. Thus there are no solutions of $R(z) = \infty$ except possibly 0 or ∞ . If $R(0) \neq \infty$, then R has no poles, so it is a polynomial, and is therefore equal to z^m , $m > 0$, since $R(z) = 0$ has no non-zero solutions. If $R(0) = \infty$, then R has a unique pole at 0; since there are no finite solutions of $R(z) = 0$, it follows that $R(z) = 1/z^m$. We have shown that R is conjugate to z^m , $|m| > 1$, if the exceptional set of some open set U has two points. In this case the exceptional set is $\{0, \infty\}$.

Suppose that $\{R^n\}_{n \in \mathbb{N}_0}$ omits at most a single point on U for every open set U with $U \cap J(R) \neq \emptyset$. Fix such a set U with $E_U \neq \emptyset$, and let z_0 be the omitted point. Replacing R with its conjugate by the rational map $\phi(z) = 1/(z - z_0)$, we may take $z_0 = \infty$. Since $\{\infty\}$ is omitted, R has no poles, and is therefore a polynomial. Thus R omits ∞ on every open subset $U \subset \mathbb{C}$, and (by hypothesis) omits only a single point on U if $U \cap J(R) \neq \emptyset$, so ∞ is the only exceptional point of R .

In either case, $J(R)$ does not contain any exceptional points. \square

The following proposition shows that the Julia set possesses self-similarity, a characteristic property of fractal sets.

PROPOSITION 8.5.4. *Let $R: \mathbb{C} \rightarrow \mathbb{C}$ be a rational map of degree >1 with exceptional set E , and let U be a neighborhood of a point $\zeta \in J(R)$. Then $\bigcup_{n \in \mathbb{N}} R^n(U) = \mathbb{C} \setminus E$, and $J(R) \subset R^n(U)$ for some $n \in \mathbb{N}$.*

Proof. If E contains two points, then by Proposition 8.5.3, R is conjugate to z^m , $|m| > 1$, and the proof is left as an exercise (Exercise 8.5.4).

Suppose E is empty or consists of a single point. If the latter, we may and do assume that the omitted point is ∞ and R is a polynomial. Since repelling periodic points are dense in $J(R)$, we may choose a neighborhood $V \subset U$ and $n > 0$ such that $R^n(V) \supset V$. The family $\{R^{nk}\}_{k \in \mathbb{N}}$ on V does not omit any points in \mathbb{C} , and ∞ is omitted if and only if R is a polynomial, in which case $\infty \notin J(R)$. Hence $J(R) \subset \bigcup_n R^n(V)$. Since $J(R)$ is compact and $R^{nk}(V) \supset R^{n(k-1)}(V)$, the proposition follows. \square

COROLLARY 8.5.5. *Let $R: \mathbb{C} \rightarrow \mathbb{C}$ be a rational map of degree >1 . For any point $\zeta \notin E$, $J(R)$ is contained in the closure of the set of backward iterates of ζ . In particular, $J(R)$ is the closure of the set of backward iterates of any point in $J(R)$.*

PROPOSITION 8.5.6. *The Julia set of a rational map of degree >1 is perfect, i.e., it does not have isolated points.*

Proof. Exercise 8.5.3. \square

PROPOSITION 8.5.7. *Let $R: \mathbb{C} \rightarrow \mathbb{C}$ be a rational map of degree >1 . Then $J(R)$ is the closure of the set of repelling periodic points.*

Proof. We will show that $J(R)$ is contained in the closure of the set $\text{Per}(R)$ of the periodic points of R . The result will follow, since $J(R)$ is perfect and there are only finitely many non-repelling periodic points.

Suppose $\zeta \in J(R)$ has a neighborhood U that contains no periodic points, no poles, and no critical values of R . Since the degree of R is >1 , there are distinct branches f and g of R^{-1} in U , and $f(z) \neq g(z)$, $f(z) \neq R^n(z)$, and $g(z) \neq R^n(z)$ for all $n \geq 0$ and all $z \in U$. Hence the family

$$h_n(z) = \frac{R^n(z) - f(z)}{R^n(z) - g(z)} \cdot \frac{z - g(z)}{z - f(z)}, \quad n \in \mathbb{N},$$

omits $0, 1$, and ∞ in U and therefore is normal by Theorem 8.3.2. Since R^n can be expressed in terms of h_n , the family $\{R^n\}$ is also normal in U , a contradiction. Therefore $J(R) \subset \overline{\text{Per}(R)}$. \square

Let $P: \bar{\mathbb{C}} \rightarrow \bar{\mathbb{C}}$ be a polynomial. Then $P(\infty) = \infty$, and locally near ∞ there are $\deg P$ branches of P^{-1} . The complete preimage of any connected domain containing ∞ is connected, since $\infty = P^{-1}(\infty)$ must belong to every connected component of the preimage. Therefore $\text{BA}(\infty)$ is connected, i.e., $\text{BA}(\infty) = \text{BA}^\circ(\infty)$.

LEMMA 8.5.8. *Let $f: \bar{\mathbb{C}} \rightarrow \bar{\mathbb{C}}$ be a meromorphic function, and suppose ζ is an attracting periodic point. Then every component of $\text{BA}^\circ(\zeta)$ is simply connected.*

Proof. Since f cyclically permutes the components of $\text{BA}^\circ(\zeta)$, we may replace f by f^n , where n is the minimal period of ζ , and assume that ζ is fixed. After conjugating by a Möbius transformation, we may assume that ζ is finite.

Let γ be a smooth simple closed curve in $\text{BA}^\circ(\zeta)$, and let D be the simply connected region (in \mathbb{C}) that it bounds. Suppose $D \not\subseteq \text{BA}^\circ(\zeta)$. Let δ be the distance from ζ to the boundary of $\text{BA}^\circ(\zeta)$, and let U be the disk of radius $\delta/2$ around ζ . Because ζ is attracting, and γ is a compact subset of $\text{BA}^\circ(\zeta)$, there is $n > 0$ such that $f^n(\gamma) \subset U$. Let $g(z) = f^n(z) - \zeta$. Then $|g(z)| < \delta/2$ on γ , but $|g(z)| > \delta$ for some $z \in D$, since $f^n(D) \not\subseteq \text{BA}^\circ(\zeta)$. This contradicts the maximum principle for analytic functions. Thus $D \subset \text{BA}^\circ(\zeta)$, and $\text{BA}^\circ(\zeta)$ is simply connected. \square

PROPOSITION 8.5.9. *Let $R: \bar{\mathbb{C}} \rightarrow \bar{\mathbb{C}}$ be a rational map of degree > 1 . If U is any completely invariant component of $F(R)$, then $J(R) = \bar{U} \setminus U$, and $J(R) = \partial U$ if $F(R)$ is not connected. Every other component of $F(R)$ is simply connected. There are at most two completely invariant components. If R is a polynomial, then $\text{BA}(\infty)$ is completely invariant.*

Proof. Suppose U is a completely invariant component of $F(R)$. Then, by Corollary 8.5.5, $J(R)$ is contained in the closure of U , and also of $F(R) \setminus U$ if the latter is non-empty. This proves the first assertion. Since $J(R) \cup U = \bar{U}$ is connected, every component of the complement in $\bar{\mathbb{C}}$ is simply connected (by a basic result of homotopy theory).

Suppose there is more than one completely invariant component of $F(R)$. Then, by the preceding paragraph, each must be simply connected. Let U be such a component. Then $R: U \rightarrow U$ is a branched covering of degree d , so there must be $d - 1$ critical points, counted with multiplicity. Since the total number of critical points is $2d - 2$ (Proposition 8.1.1), this implies that there are at most two completely invariant components.

If R is a polynomial, then $\text{BA}(\infty)$ is completely invariant (Exercise 8.5.1). \square

The *postcritical set* of a rational map R is the union of the forward orbits of all critical points of R , and is denoted $\text{CL}(R)$.

THEOREM 8.5.10 (Fatou). *Let R be a rational map of degree >1 . Suppose that all critical points of R tend to attracting periodic points of R under the forward iterates of R . Then $J(R)$ is a hyperbolic set for R , i.e., there are $a > 1$ and $n \in \mathbb{N}$ such that $|(R^n)'(z)| \geq a$ for every $z \in J(R)$.*

Proof. If R has exactly two critical points, then it is conjugate to z^d or z^{-d} (Proposition 8.1.1), and the theorem follows by a direct computation.

We assume then that there are at least three critical points. Let $U = \bar{\mathbb{C}} \setminus \text{CL}(R)$; then $R^{-1}(U) \subset U$. By the uniformization theorem [Ahl73], there is an analytic covering map $\phi: \Delta_1 \rightarrow U$. Let $g: \Delta_1 \rightarrow \Delta_1$ be the lift of a locally defined branch of R^{-1} , so $R \circ \phi \circ g = \phi$.

The family $\{\phi \circ g^n\}$ is normal, since it omits $\text{CL}(R)$. Let f be the uniform limit of a sequence $\phi \circ g^{n_k}$. Let $z_0 \in \phi^{-1}(J(R))$, and let $O \subset \Delta_1$ be a neighborhood of z_0 such that $\phi(O)$ does not contain any attracting periodic points of R . Since $J(R)$ is invariant (Proposition 8.5.1) and closed, $f(z_0) \in J(R)$. If $f'(z_0) \neq 0$, then $f(O)$ contains a neighborhood of $f(z_0)$, and hence (by Proposition 8.5.9) contains a point $z_1 \in \text{BA}(\xi)$, where ξ is an attracting periodic point. Since $\phi \circ g^{n_k} \rightarrow f$, the value z_1 is taken on by every $\phi \circ g^{n_k}$ with k large enough. This implies that $R^{n_k}(z_1) \in \phi(O)$ for k sufficiently large, which contradicts the fact that $z_1 \in \text{BA}(\xi)$ and $\xi \notin \phi(O)$. Therefore, $f'(z_0) = 0$, so f is constant on $\phi^{-1}(J(R))$. It follows that $(R^{n_k})' = 1/(g^{n_k})'$ goes to infinity uniformly on $J(P)$, which proves the theorem. \square

THEOREM 8.5.11 (Fatou). *Let $P: \bar{\mathbb{C}} \rightarrow \bar{\mathbb{C}}$ be a polynomial such that $P^n(c) \rightarrow \infty$ as $n \rightarrow \infty$ for every critical point c . Then the Julia set $J(P)$ is totally disconnected, i.e., $J(P)$ is a Cantor set.*

Proof. Let D be a disk centered at 0 that contains $J(P)$, and choose N large enough that P^N carries all critical points outside of \bar{D} . Then for $n \geq N$, branches of P^{-n} are globally defined on D . Fix $z_0 \in J(P)$, and let g_n be the branch of P^{-n} with $g_n(P^n(z_0)) = z_0$, for $n \geq N$. The family $F = \{g_n\}_{n \geq N}$ is uniformly bounded on \bar{D} , and is therefore normal on \bar{D} . Let f be the uniform limit of a sequence in F . Since P is hyperbolic on $J(P)$ (Theorem 8.5.10), f must be constant on $J(P)$, and therefore constant on \bar{D} , since f is analytic and $J(P)$ has no isolated points. If $y \neq z_0$ is any other point of $J(P)$, then $y \notin g_n(D)$ for n sufficiently large, since the diameter of $g_n(D)$ converges to zero. The set $g_n(D) \cap J(P)$ is both open and closed in $J(P)$, because ∂D does not intersect $J(P)$. Thus z_0 and y are in different components of $J(P)$, so $J(P)$ is totally disconnected. \square

PROPOSITION 8.5.12. *Let $P: \bar{\mathbb{C}} \rightarrow \bar{\mathbb{C}}$ be a polynomial such that no critical point lies in $\text{BA}(\infty)$. Then $J(P)$ is connected.*

Proof. $\text{BA}(\infty)$ is simply connected (Lemma 8.5.8) and completely invariant. If $F(P)$ has only one component, then $J(P)$ is the complement in $\bar{\mathbb{C}}$ of $\text{BA}(\infty)$, and is therefore connected by a fundamental result of algebraic topology.

We assume then that $F(P)$ has at least two components. We conjugate by a Möbius transformation that carries ∞ to O , and one of the other components of $F(P)$ to a neighborhood of ∞ . We obtain a rational map R such that 0 is a superattracting fixed point and $\text{BA}(0)$ is a bounded, simply connected, completely invariant component of $F(R)$ that contains no critical points. Let g_n be the branch of R^n on $\text{BA}(0)$ with $g_n(0) = 0$. Let γ be the unit circle. Then $g_n(\gamma)$ converges to $J(R)$, so $J(R)$ is connected. \square

There are many other results about the Fatou and Julia sets that are beyond the scope of this book. For example, results of Wolff–Denjoy [Wol26], [Den26] and of Douady–Hubbard [Dou83] show that if a component of the Fatou set is eventually mapped back to itself, then its closure contains either an attracting periodic point or a neutral periodic point. A result of Sullivan [Sul85] shows that the Fatou set has no wandering components, i.e., no orbit in the set of components is infinite.

Exercise 8.5.1. Show that if U is a connected component of $F(R)$, then $R(U)$ is also a connected component of $F(R)$. Show that if P is a polynomial, then $\text{BA}(\infty)$ is completely invariant.

Exercise 8.5.2. Show that, for $m > 1$, the Julia set of $z \mapsto z^m$ is the unit circle S^1 , $\text{BA}(\infty)$ is the exterior of S^1 , and the α -limit set of every $z \neq 0$ is S^1 .

Exercise 8.5.3. Prove Proposition 8.5.6.

Exercise 8.5.4. Prove Proposition 8.5.4 for $R(z) = z^m$, $|m| > 1$.

Exercise 8.5.5. Let P be a polynomial of degree at least 2. Prove that $P^n \rightarrow \infty$ on the component of $F(P)$ that contains ∞ .

Exercise 8.5.6. Show that if R is a rational map of degree > 1 , and $F(R)$ has only finitely many components, then it has either 0, 1, or 2 components.

8.6 The Mandelbrot Set

For a general quadratic function $q(z) = \alpha z^2 + \beta z + \gamma$ with $\alpha \neq 0$, the change of variables $\zeta = z + \beta/2$ maps the critical point to 0 and conjugates q with




Figure 8.3. The Mandelbrot set.

$q_a(z) = z^2 - a$. Since the conjugation is unique, the maps q_a , $a \in \mathbb{C}$, are in one-to-one correspondence with conjugacy classes of quadratic maps. If $q_a^n(0) \rightarrow \infty$, then $J(q_a)$ is totally disconnected (see Theorem 8.5.11). Otherwise, the orbit $\{q_a^n(0)\}_{n \in \mathbb{N}}$ is bounded and $J(q_a)$ is connected (Proposition 8.5.12).

The *Mandelbrot set* M is the set of parameter values a for which the orbit of 0 is bounded, or equivalently, $M = \{a \in \mathbb{C} : 0 \notin \text{BA}(\infty) \text{ for } q_a\}$. The Mandelbrot set is shown in Figure 8.3.

THEOREM 8.6.1 (Douady–Hubbard [DH82]). $M = \{a \in \mathbb{C} : |q_a^n(0)| \leq 2 \text{ for all } n \in \mathbb{N}\}$. M is closed and simply connected.

Proof. Let $|a| > 2$. We have $|q_a(0)| = |a| > 2$, $|q_a^2(0)| = |q_a(a)| \geq |a^2| - |a| = |a|(|a| - 1)$, and $|q_a^n(0)| \geq |a|(|a| - 1)^{n-1}$ for $n \in \mathbb{N}$ (Exercise 8.6.1). Therefore $a \notin M$. If $|a| \leq 2$ and $|q_a^n(0)| = 2 + \alpha$ for some $n \in \mathbb{N}$ and $\alpha > 0$, then $|q_a^{n+1}(0)| \geq (2 + \alpha)^2 - 2 > 2 + 4\alpha$ and $|q_a^{n+k}(0)| \geq 2 + 4^k\alpha \rightarrow \infty$ as $k \rightarrow \infty$. Therefore $a \notin M$. The first and second statements follow.

If D is a bounded component of $\mathbb{C} \setminus M$, then $\max_{a \in \bar{D}} |q_a^n(0)| > 2$ for some $n \in \mathbb{N}$, and, by the maximum principle, $|q_a^n(0)| > 2$ for some $a \in \partial D \subset M$. This contradicts the first assertion of the theorem. Thus $\mathbb{C} \setminus M$ has no bounded components, has only one unbounded component containing ∞ , and is therefore connected. Hence M is simply connected. \square

The fixed points of q_a are $z_a^\pm = (1 \pm \sqrt{1 + 4a})/2$ with multipliers $\lambda^\pm = 1 \pm \sqrt{1 + 4a}$. The set $\{a \in \mathbb{C} : |1 \pm \sqrt{1 + 4a}| < 1\}$ is a subset of M (Exercise 8.6.3) and is called the *main cardioid* of M .

PROPOSITION 8.6.2. *Every point in ∂M is an accumulation point of the set of values of a for which q_a has a superattracting cycle.*

Proof. Since 0 is the only critical point of q_a , a periodic orbit is superattracting if and only if it contains 0. Let D be a disk that intersects ∂M and does not contain 0, and suppose that 0 is not a periodic point of q_a for any $a \in D$. Then $(q_a^n(0))^2 \neq a$ for all $a \in D$ and $n \in \mathbb{N}$. Let \sqrt{a} be a branch of the inverse of $z \mapsto z^2$ defined on D , and define $f_n(a) = q_a^n(0)/\sqrt{a}$ for $n \in \mathbb{N}$ and $a \in D$. Then the family $\{f_n\}_{n \in \mathbb{N}}$ omits 0, 1, and ∞ on D , and is therefore normal in D . On the other hand, since D intersects ∂M , it contains both points a for which $f_n(a)$ is bounded and points a for which $f_n(a) \rightarrow \infty$, and hence the family $\{f_n\}$ is not normal on D . Thus 0 must be periodic for q_a for some $a \in D$. \square

Exercise 8.6.1. Prove by induction that if $|a| > 2$, then $|q_a^n(0)| \geq |a|(|a| - 1)^{n-1}$ for $n \in \mathbb{N}$.

Exercise 8.6.2. Prove that the intersection of M with the real axis is $[-2, 1/4]$.

Exercise 8.6.3. Prove that the main cardioid is contained in M .

Exercise 8.6.4. Prove that the set of values a in \mathbb{C} for which q_a has an attracting periodic point of period 2 is the disk of radius $1/4$ centered at -1 (it is tangent to the main cardioid). Prove that this set is contained in M .

Measure-Theoretic Entropy

In this chapter, we give a short introduction to *measure-theoretic entropy*, also called *metric entropy*, for measure-preserving transformations. This invariant was introduced by A. Kolmogorov [Kol58], [Kol59] to classify Bernoulli automorphisms and developed further by Ya. Sinai [Sin59] for general measure-preserving dynamical systems. The measure-theoretic entropy has deep roots in thermodynamics, statistical mechanics, and information theory. We explain the interpretation of entropy from the perspective of information theory at the end of the first section.

9.1 Entropy of a Partition

Throughout this chapter (X, \mathfrak{A}, μ) is a Lebesgue space with $\mu(X) = 1$. We use the notation of Chapter 4. A (finite) *partition* of X is a finite collection ζ of essentially disjoint measurable sets C_i (called *elements* or *atoms* of ζ) whose union covers $X \bmod 0$. We say that a partition ζ' is a *refinement* of ζ and write $\zeta \leq \zeta'$ (or $\zeta' \geq \zeta$) if every element of ζ' is contained $\bmod 0$ in an element of ζ . Partitions ζ and ζ' are *equivalent* if each is a refinement of the other. We will deal with equivalence classes of partitions. The *common refinement* $\zeta \vee \zeta'$ of partitions ζ and ζ' is the partition into intersections $C_\alpha \cap C'_\beta$, where $C_\alpha \in \zeta$ and $C'_\beta \in \zeta'$; it is the smallest partition which is $\geq \zeta$ and ζ' . The *intersection* $\zeta \wedge \zeta'$ is the largest measurable partition which is $\leq \zeta$ and ζ' . The trivial partition consisting of a single element X is denoted by ν .

Although many definitions and statements in this chapter hold for infinite partitions, we discuss only finite partitions.

For $A, B \subset X$, let $A \triangle B = (A \setminus B) \cup (B \setminus A)$. Let $\xi = \{C_i: 1 \leq i \leq m\}$ and $\eta = \{D_j: 1 \leq j \leq n\}$ be finite partitions. By adding null sets if necessary, we

may assume that $m = n$. Define

$$d(\xi, \eta) = \min_{\sigma \in \mathcal{S}_m} \sum_{i=1}^m \mu(C_i \Delta D_{\sigma(i)}),$$

where the minimum is taken over all permutations of m elements. The axioms of distance are satisfied by d (Exercise 9.1.1).

Partitions ζ and ζ' are *independent*, and we write $\zeta \perp \zeta'$, if $\mu(C \cap C') = \mu(C) \cdot \mu(C')$ for all $C \in \zeta$ and $C' \in \zeta'$.

For a transformation T and partition $\xi = \{C_1, \dots, C_m\}$, let $T^{-1}(\xi) = \{T^{-1}(C_1), \dots, T^{-1}(C_m)\}$.

To motivate the definition of entropy below, consider a Bernoulli automorphism of Σ_m with probabilities $q_i > 0$, $q_1 + \dots + q_m = 1$ (see §4.4). Let ξ be the partition of Σ_m into m sets $C_i = \{\omega \in \Sigma_m: \omega_0 = i\}$, $\mu(C_i) = q_i$. Set $\eta_n = \bigvee_{k=0}^{n-1} \sigma^{-k}(\xi)$, and let $\eta_n(\omega)$ denote the element of η_n containing ω . For $\omega \in \Sigma_m$, let $f_i^n(\omega)$ be the relative frequency of symbol i in the word $\omega_1 \dots \omega_n$. Since σ is ergodic with respect to μ , by the Birkhoff ergodic theorem 4.5.5, for every $\epsilon > 0$ there are $N \in \mathbb{N}$ and a subset $A_\epsilon \subset \Sigma_m$ with $\mu(A_\epsilon) > 1 - \epsilon$ such that $|f_i^n(\omega) - q_i| < \epsilon$ for each $\omega \in A_\epsilon$ and $n \geq N$. Therefore, if $\omega \in A_\epsilon$, then

$$\mu(\eta_n(\omega)) = \prod_{i=1}^m q_i^{(q_i + \epsilon_i)n} = 2^{n \sum_{i=1}^m (q_i + \epsilon_i) \log q_i},$$

where $|\epsilon_i| < \epsilon$, and from now on \log denotes logarithm base 2 with $0 \log 0 = 0$. It follows that for μ -a.e. $\omega \in \Sigma_m$

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \mu(\eta_n(\omega)) = \sum_{i=1}^m q_i \log q_i,$$

and hence the number of elements of η_n with approximately correct frequency of symbols $1, \dots, m$ grows exponentially as 2^{nh} , where $h = -\sum_{i=1}^m q_i \log q_i$.

For a partition $\zeta = \{C_1, \dots, C_n\}$ define the *entropy of ζ* by

$$H(\zeta) = - \sum_{i=1}^n \mu(C_i) \log \mu(C_i)$$

(recall that $0 \log 0 = 0$). Note that $-x \log x$ is a strictly concave continuous function on $[0, 1]$, i.e., if $x_i \geq 0$, $\lambda_i \geq 0$, $i = 1, \dots, n$, and $\sum_i \lambda_i = 1$, then

$$- \left(\sum_{i=1}^n \lambda_i x_i \right) \cdot \log \sum_{i=1}^n \lambda_i x_i \geq - \sum_{i=1}^n \lambda_i x_i \log x_i \tag{9.1}$$

with equality if and only if all x_i s are equal. For $x \in X$, let $m(x, \zeta)$ denote the measure of the element of ζ containing x . Then

$$H(\zeta) = - \int_X \log m(x, \zeta) d\mu.$$

PROPOSITION 9.1.1. *Let ξ and η be finite partitions. Then*

1. $H(\xi) \geq 0$, and $H(\xi) = 0$ if and only if $\xi = \nu$;
2. if $\xi \leq \eta$, then $H(\xi) \leq H(\eta)$, and equality holds if and only if $\xi = \eta$;
3. if ξ has n elements, then $H(\xi) \leq \log n$, and equality holds if and only if each element of ξ has measure $1/n$;
4. $H(\xi \vee \eta) \leq H(\xi) + H(\eta)$ with equality if and only if $\xi \perp \eta$.

Proof. We leave the first three statements as exercises (Exercise 9.1.2). To prove the last statement, let μ_i, ν_j , and κ_{ij} be the measures of the elements of ξ, η , and $\xi \vee \eta$, respectively, so that $\sum_j \kappa_{ij} = \mu_i$ and $\sum_i \kappa_{ij} = \nu_j$. It follows from (9.1) that

$$-\nu_j \log \nu_j \geq - \sum_i \mu_i \frac{\kappa_{ij}}{\mu_i} \cdot \log \frac{\kappa_{ij}}{\mu_i} = - \sum_i \kappa_{ij} \log \kappa_{ij} + \sum_i \kappa_{ij} \log \mu_i,$$

and summation over j finishes the proof of the inequality. The equality is achieved if and only if $x_i = \kappa_{ij}/\mu_i$ does not depend on i for each j , which is equivalent to the independence of ξ and η . \square

The entropy of a partition has a natural interpretation as the ‘‘average information of the elements of the partition.’’ For example, suppose X represents the set of all possible outcomes of an experiment, and μ is the probability distribution of the outcomes. To extract information from the experiment, we devise a measuring scheme that effectively partitions X into finitely many observable subsets, or *events*, C_1, C_2, \dots, C_n . We define the information of an event C to be $I(C) = -\log \mu(C)$. This is a natural choice given that the information should have the following properties:

1. The information is a non-negative and decreasing function of the probability of an event; the lower the probability of an event, the greater the informational content of observing that event.
2. The information of the trivial event X is 0.
3. For independent events C and D , the information is additive, i.e., $I(C \cap D) = I(C) + I(D)$.

Up to a constant, $-\log \mu(C)$ is the only such function.

With this definition of information, the entropy of a partition is simply the average information of the elements of the partition.

Exercise 9.1.1. Prove: (i) $d(\xi, \eta) \geq 0$ with equality if and only if $\xi = \eta \bmod 0$ and (ii) $d(\xi, \zeta) \leq d(\xi, \eta) + d(\eta, \zeta)$.

Exercise 9.1.2. Prove the first three statements of Proposition 9.1.1.

Exercise 9.1.3. For $n \in \mathbb{N}$, let \mathcal{P}_n be the the space of equivalence classes of finite partitions with n elements with metric d . Prove that the entropy is a continuous function on \mathcal{P}_n .

9.2 Conditional Entropy

For measurable subsets $C, D \subset X$ with $\mu(D) > 0$, set $\mu(C|D) = \mu(C \cap D)/\mu(D)$. Let $\xi = \{C_i: i \in I\}$ and $\eta = \{D_j: j \in J\}$ be partitions. The *conditional entropy of ξ with respect to η* is defined by the formula

$$H(\xi|\eta) = - \sum_{j \in J} \mu(D_j) \sum_{i \in I} \mu(C_i|D_j) \log \mu(C_i|D_j).$$

The quantity $H(\xi|\eta)$ is the average entropy of the partition induced by ξ on an element of η . If $C(x) \in \xi$ and $D(x) \in \eta$ are the elements containing x , then

$$H(\xi|\eta) = - \int_X \log \mu(C(x)|D(x)) d\mu.$$

The following proposition gives several simple properties of conditional entropy.

PROPOSITION 9.2.1. *Let ξ, η , and ζ be finite partitions. Then*

1. $H(\xi|\eta) \geq 0$ with equality if and only if $\xi \leq \eta$;
2. $H(\xi|\nu) = H(\xi)$;
3. if $\eta \leq \zeta$, then $H(\xi|\eta) \geq H(\xi|\zeta)$;
4. if $\eta \leq \zeta$, then $H(\xi \vee \eta|\zeta) = H(\xi|\zeta)$;
5. if $\xi \leq \eta$, then $H(\xi|\zeta) \leq H(\eta|\zeta)$ with equality if and only if $\xi \vee \zeta = \eta \vee \zeta$;
6. $H(\xi \vee \eta|\zeta) = H(\xi|\zeta) + H(\eta|\xi \vee \zeta)$ and $H(\xi \vee \eta) = H(\xi) + H(\eta|\xi)$;
7. $H(\xi|\eta \vee \zeta) \leq H(\xi|\zeta)$;
8. $H(\xi|\eta) \leq H(\xi)$ with equality if and only if $\xi \perp \eta$.

Proof. To prove part 6, let $\xi = \{A_i\}$, $\eta = \{B_j\}$, $\zeta = \{C_k\}$. Then

$$\begin{aligned} H(\xi \vee \eta | \zeta) &= - \sum_{i,j,k} \mu(A_i \cap B_j \cap C_k) \cdot \log \frac{\mu(A_i \cap B_j \cap C_k)}{\mu(C_k)} \\ &= - \sum_{i,j,k} \mu(A_i \cap B_j \cap C_k) \log \frac{\mu(A_i \cap C_k)}{\mu(C_k)} \\ &\quad - \sum_{i,j,k} \mu(A_i \cap B_j \cap C_k) \log \frac{\mu(A_i \cap B_j \cap C_k)}{\mu(A_i \cap C_k)} \\ &= H(\xi | \zeta) + H(\eta | \xi \vee \zeta), \end{aligned}$$

and the first equality follows. The second equality follows from the first one with $\zeta = \nu$.

The remaining statements of Proposition 9.2.1 are left as exercises (Exercise 9.2.1). \square

For finite partitions ξ and η , define

$$\rho(\xi, \eta) = H(\xi | \eta) + H(\eta | \xi).$$

The function ρ , which is called the *Rokhlin metric*, defines a metric on the space of equivalence classes of partitions (Exercise 9.2.2).

PROPOSITION 9.2.2. *For every $\epsilon > 0$ and $m \in \mathbb{N}$ there is $\delta > 0$ such that if ξ and η are finite partitions with at most m elements and $d(\xi, \eta) < \delta$, then $\rho(\xi, \eta) < \epsilon$.*

Proof ([KH95], Proposition 4.3.5). Let partitions, $\xi = \{C_i: 1 \leq i \leq m\}$, $\eta = \{D_i: 1 \leq i \leq m\}$ satisfy $d(\xi, \eta) = \sum_{i=1}^m \mu(C_i \Delta D_i) = \delta$. We will estimate $H(\eta | \xi)$ in terms of δ and m .

If $\mu(C_i) > 0$, set $\alpha_i = \mu(C_i \setminus D_i) / \mu(C_i)$. Then

$$-\mu(C_i \cap D_i) \log \frac{\mu(C_i \cap D_i)}{\mu(C_i)} \leq -\mu(C_i)(1 - \alpha_i) \log(1 - \alpha_i)$$

and, by Proposition 9.1.1(3) applied to the partition of $C_i \setminus D_i$ induced by η ,

$$-\sum_{j \neq i} \mu(C_i \cap D_j) \log \frac{\mu(C_i \cap D_j)}{\mu(C_i)} \leq -\mu(C_i)\alpha_i(\log \alpha_i - \log(m-1)).$$

Therefore, since $\log x$ is concave,

$$\begin{aligned}
 & - \sum_j \mu(C_i \cap D_j) \log \frac{\mu(C_i \cap D_j)}{\mu(C_i)} \\
 & \leq \mu(C_i) \left((1 - \alpha_i) \log \frac{1}{1 - \alpha_i} + \alpha_i \log \frac{m - 1}{\alpha_i} \right) \leq \mu(C_i) \log m.
 \end{aligned}$$

It follows that

$$\begin{aligned}
 H(\eta|\xi) & \leq \sum_{\mu(C_i) < \sqrt{\delta}} \mu(C_i) \log m \\
 & + \sum_{\mu(C_i) \geq \sqrt{\delta}} \mu(C_i) (-(1 - \alpha_i) \log(1 - \alpha_i) - \alpha_i \log \alpha_i + \alpha_i \log(m - 1)).
 \end{aligned}$$

The first term does not exceed $\sqrt{\delta} m \log m$. To estimate the second term, observe that $\alpha_i \mu(C_i) \leq \delta$. Hence, if $\mu(C_i) \geq \sqrt{\delta}$, then $\alpha_i \leq \sqrt{\delta}$. Since the function $f(x) = -x \log x - (1 - x) \log(1 - x)$ is increasing on $(0, 1/2)$, for small δ the second term does not exceed $f(\sqrt{\delta}) + \sqrt{\delta} \log(m - 1)$, and

$$H(\eta|\xi) \leq f(\sqrt{\delta}) + \sqrt{\delta}(m \log m + \log(m - 1)).$$

Since $f(x) \rightarrow 0$ as $x \rightarrow 0$, the proposition follows. □

Exercise 9.2.1. Prove the remaining statements of Proposition 9.2.1.

Exercise 9.2.2. Prove that (i) $\rho(\xi, \eta) \geq 0$ with equality if and only if $\xi = \eta \bmod 0$ and (ii) $\rho(\xi, \zeta) \leq \rho(\xi, \eta) + \rho(\eta, \zeta)$.

9.3 Entropy of a Measure-Preserving Transformation

Let T be a measure-preserving transformation of a measure space (X, \mathfrak{A}, μ) and $\zeta = \{C_\alpha : \alpha \in I\}$ be a partition of X with finite entropy. For $k, n \in \mathbb{N}$, set $T^{-k}(\zeta) = \{T^{-k}(C_\alpha) : \alpha \in I\}$ and

$$\zeta^n = \zeta \vee T^{-1}(\zeta) \vee \dots \vee T^{-n+1}(\zeta).$$

Since $H(T^{-k}(\zeta)) = H(\zeta)$ and $H(\xi \vee \eta) \leq H(\xi) + H(\eta)$, we have that $H(\zeta^{m+n}) \leq H(\zeta^m) + H(\zeta^n)$. By subadditivity (Exercise 2.5.3), the limit

$$h(T, \zeta) = \lim_{n \rightarrow \infty} \frac{1}{n} H(\zeta^n)$$

exists, and is called the *metric* (or *measure-theoretic*) *entropy* of T relative to ζ . Note that $h(T, \zeta) \leq H(\zeta)$.

PROPOSITION 9.3.1. $h(T, \zeta) = \lim_{n \rightarrow \infty} H(\zeta | T^{-1}(\zeta^n))$.

Proof. Since $H(\xi|\eta) \geq H(\xi|\zeta)$ for $\eta \leq \zeta$, the sequence $H(\zeta|T^{-1}(\zeta^n))$ is non-increasing in n . Since $H(T^{-1}\xi) = H(\xi)$ and $H(\xi \vee \eta) = H(\xi) + H(\eta|\xi)$, we get

$$\begin{aligned} H(\zeta^n) &= H(T^{-1}(\zeta^{n-1}) \vee \zeta) = H(\zeta^{n-1}) + H(\zeta|T^{-1}(\zeta^{n-1})) \\ &= H(\zeta^{n-2}) + H(\zeta|T^{-1}(\zeta^{n-2})) + H(\zeta|T^{-1}(\zeta^{n-1})) = \dots \\ &= H(\zeta) + \sum_{k=1}^{n-1} H(\zeta|T^{-1}(\zeta^k)). \end{aligned}$$

Dividing by n and passing to the limit as $n \rightarrow \infty$ finishes the proof. \square

Proposition 9.3.1 means that $h(T, \zeta)$ is the average information added by the present state on condition that all past states are known.

PROPOSITION 9.3.2. *Let ξ and η be finite partitions. Then*

1. $h(T, T^{-1}(\xi)) = h(T, \xi)$; if T is invertible, then $h(T, T(\xi)) = h(T, \xi)$;
2. $h(T, \xi) = h(T, \bigvee_{i=0}^n T^{-i}(\xi))$ for $n \in \mathbb{N}$; if T is invertible, then $h(T, \xi) = h(T, \bigvee_{i=-n}^n T^{-i}(\xi))$ for $n \in \mathbb{N}$;
3. $h(T, \xi) \leq h(T, \eta) + H(\xi|\eta)$; if $\xi \leq \eta$, then $h(T, \xi) \leq h(T, \eta)$;
4. $|h(T, \xi) - h(T, \eta)| \leq \rho(\xi, \eta) = H(\xi|\eta) + H(\eta|\xi)$ (the Rokhlin inequality);
5. $h(T, \xi \vee \eta) \leq h(T, \xi) + h(T, \eta)$;

Proof. To prove statement 3 observe that, by the second statement of Proposition 9.2.1(6), $H(\xi^n) \leq H(\xi^n \vee \eta^n) = H(\eta^n) + H(\xi^n|\eta^n)$. We apply Proposition 9.2.1(6) n times to get

$$\begin{aligned} H(\xi^n|\eta^n) &= H(\xi \vee T^{-1}(\xi^{n-1})|\eta^n) = H(\xi|\eta^n) + H(T^{-1}(\xi^{n-1})|\xi \vee \eta^n) \\ &\leq H(\xi|\eta) + H(T^{-1}(\xi^{n-1})|\eta^n) \\ &\leq H(\xi|\eta) + H(T^{-1}(\xi)|T^{-1}(\eta)) + H(T^{-2}(\xi^{n-2})|\eta^n) \\ &\vdots \\ &\leq nH(\xi|\eta). \end{aligned}$$

Therefore

$$\frac{1}{n}H(\xi^n) \leq \frac{1}{n}H(\eta^n) + H(\xi|\eta),$$

and statement 3 follows.

The remaining statements of Proposition 9.3.2 are left as exercises (Exercise 9.3.2). \square

The *metric* (or *measure-theoretic*) *entropy* is the supremum of the entropies $h(T, \zeta)$ over all finite measurable partitions ζ of X .

If two measure-preserving transformations are isomorphic (i.e., if there exists a measure-preserving conjugacy), then their measure-theoretic entropies are equal. If the entropies are different, the transformations are not isomorphic.

We will need the following lemma.

LEMMA 9.3.3. *Let η be a finite partition, and let ζ_n be a sequence of finite partitions such that $d(\zeta_n, \eta) \rightarrow 0$. Then there are finite partitions $\xi_n \leq \zeta_n$ such that $H(\eta|\xi_n) \rightarrow 0$.*

Proof. Let $\eta = \{D_j : 1 \leq j \leq m\}$. For each j choose a sequence $C_j^n \in \zeta_n$ such that $\mu(D_j \Delta C_j^n) \rightarrow 0$. Let ξ_n consist of $C_j^n, 1 \leq j \leq m$, and $C_{m+1}^n = X \setminus \bigcup_{j=1}^m C_j^n$. Then $\mu(C_j^n) \rightarrow \mu(D_j)$ and $\mu(C_{m+1}^n) \rightarrow 0$. We have

$$\begin{aligned} H(\eta|\xi_n) &= - \sum_{i=1}^m \mu(C_i^n \cap D_i) \cdot \log \frac{\mu(C_i^n \cap D_i)}{\mu(C_i^n)} \\ &\quad - \sum_{j=1}^m \mu(C_{m+1}^n \cap D_j) \cdot \log \frac{\mu(C_{m+1}^n \cap D_j)}{\mu(C_{m+1}^n)} \\ &\quad - \sum_{i=1}^m \sum_{j \neq i} \mu(C_i^n \cap D_j) \cdot \log \frac{\mu(C_i^n \cap D_j)}{\mu(C_i^n)}. \end{aligned}$$

The first sum tends to 0 because $\mu(C_i^n \cap D_i) \rightarrow \mu(C_i^n)$. The second and third sums tend to 0 because $\mu(C_i^n \cap D_j) \rightarrow 0$ for $j \neq i$. □

A sequence (ζ_n) of finite partitions is called *refining* if $\zeta_{n+1} \geq \zeta_n$ for $n \in \mathbb{N}$.

A sequence (ζ_n) of finite partitions is called *generating* if for every finite partition ξ and every $\delta > 0$ there is $n_0 \in \mathbb{N}$ such that for every $n \geq n_0$ there is a partition ξ_n with $\xi_n \leq \bigvee_{i=1}^n \zeta_i$ and $d(\xi_n, \xi) < \delta$, or equivalently if every measurable set can be approximated by a union of elements of $\bigvee_{i=1}^n \zeta_i$ for a large enough n .

Every Lebesgue space has a generating sequence of finite partitions (Exercise 9.3.3). If X is a compact metric space with a non-atomic Borel measure μ , then a sequence of finite partitions ζ_n is generating if the maximal diameter of elements of ζ_n tends to 0 as $n \rightarrow \infty$ (Exercise 9.3.4).

PROPOSITION 9.3.4. *If (ζ_n) is a refining and generating sequence of finite partitions, then $h(T) = \lim_{n \rightarrow \infty} h(T, \zeta_n)$.*

Proof. Let ξ be a partition of X with m elements. Fix $\epsilon > 0$. Since (ζ_n) is a refining and generating, for every $\delta > 0$ there is $n \in \mathbb{N}$ and a partition ξ_n

with m elements such that $\xi_n \leq \bigvee_{i=1}^n \zeta_i$ and $d(\xi_n, \xi) < \delta$. By Proposition 9.2.2,

$$\rho(\xi, \zeta_n) = H(\xi|\zeta_n) + H(\zeta_n|\xi) < \epsilon.$$

By the Rokhlin inequality (Proposition 9.3.2(4)), $h(T, \xi) < h(T, \zeta_n) + \epsilon$. □

A (one-sided) *generator* for a non-invertible measure-preserving transformation T is a finite partition ξ such that the sequence $\xi^n = \bigvee_{k=0}^n T^{-k}(\xi)$ is generating. For an invertible T , a (two-sided) *generator* is a finite partition ξ such that the sequence $\bigvee_{k=-n}^n T^k(\xi)$ is generating. Equivalently, ξ is a generator if for any finite partition η there are partitions $\zeta_n \leq \bigvee_{k=0}^n T^{-k}(\xi)$ (or $\zeta_n \leq \bigvee_{k=-n}^n T^k(\xi)$) such that $d(\zeta_n, \eta) \rightarrow 0$.

The following corollary of Proposition 9.3.4 allows one to calculate the entropy of many measure-preserving transformations.

THEOREM 9.3.5 (Kolmogorov–Sinai). *Let ξ be a generator for T . Then $h(T) = h(T, \xi)$.*

Proof. We consider only the non-invertible case. Let η be a finite partition. Since ξ is a generator, there are partitions $\zeta_n \leq \bigvee_{i=0}^n T^{-i}(\xi)$ such that $d(\zeta_n, \eta) \rightarrow 0$. By Lemma 9.3.3 for any $\delta > 0$ there is $n \in \mathbb{N}$ and a partition $\xi_n \leq \zeta_n \leq \bigvee_{i=0}^n T^{-i}(\xi)$ with $H(\xi_n|\eta) < \delta$. By statements 3, 5, and 2 of Proposition 9.3.2,

$$h(T, \eta) \leq h(T, \xi_n) + H(\eta|\xi_n) \leq h\left(T, \bigvee_{i=0}^n T^{-i}(\xi)\right) + \delta = h(T, \xi) + \delta. \quad \square$$

PROPOSITION 9.3.6. *Let T and S be measure-preserving transformations of measure spaces (X, \mathfrak{A}, μ) and (Y, \mathfrak{B}, ν) , respectively.*

1. $h(T^k) = kh(T)$ for every $k \in \mathbb{N}$; if T is invertible, then $h(T^{-1}) = h(T)$ and $h(T^k) = |k|h(T)$ for every $k \in \mathbb{Z}$.
2. If T is a factor of S , then $h_\mu(T) \leq h_\nu(S)$.
3. $h_{\mu \times \nu}(T \times S) = h_\mu(T) + h_\nu(S)$.

Proof. To prove statement 3, consider refining and generating sequences of partitions ξ_k and η_k in X and Y , respectively. Then

$$\zeta_k = (\xi_k \times \nu) \vee (\mu \times \eta_k)$$

is a refining and generating sequence in $X \times Y$. Since

$$\zeta_k^n = (\xi_k^n \times \nu) \vee (\mu \times \eta_k^n) \quad \text{and} \quad (\xi_k^n \times \nu) \perp (\mu \times \eta_k^n),$$

we obtain, by Proposition 9.1.1 and Proposition 9.3.4, that

$$h(T \times S) = \lim_{k \rightarrow \infty} \lim_{n \rightarrow \infty} \frac{1}{n} H(\zeta_k^n) \lim_{k \rightarrow \infty} \lim_{n \rightarrow \infty} \frac{1}{n} (H(\xi_k^n) + H(\eta_k^n)) = h(T) + h(S).$$

The first two statements are left as exercises (Exercise 9.3.6). \square

Let T be a measure-preserving transformation of a probability space (X, \mathfrak{A}, μ) , and ζ a finite partition. As before, let $m(x, \zeta^n)$ be the measure of the element of ζ^n containing $x \in X$. The amount of information conveyed by the fact that x lies in a particular element of ζ^n (or that the points $x, T(x), \dots, T^{n-1}(x)$ lie in particular elements of ζ) is $I_{\zeta^n}(x) = -\log \times m(x, \zeta^n)$. A proof of the following theorem can be found in [Pet89] or [Mañ88].

THEOREM 9.3.7 (Shannon–McMillan–Breiman). *Let T be an ergodic measure-preserving transformation of a probability space (X, \mathfrak{A}, μ) , and ζ a finite partition. Then*

$$\lim_{n \rightarrow \infty} \frac{1}{n} I_{\zeta^n}(x) = h(T, \zeta) \quad \text{for a.e. } x \in X \text{ and in } L^1(X, \mathfrak{A}, \mu).$$

Theorem 9.3.7 implies that, for a typical point $x \in X$, the information $I_{\zeta^n}(x)$ grows asymptotically as $n \cdot h(T, \zeta)$ and the measure $m(x, \zeta^n)$ decays exponentially as $e^{-nh(T, \zeta)}$. The proof of the following corollary is left as an exercise (Exercise 9.3.8).

COROLLARY 9.3.8. *Let T be an ergodic measure-preserving transformation of a probability space (X, \mathfrak{A}, μ) , and ζ a finite partition. Then for every $\epsilon > 0$ there is $n_0 \in \mathbb{N}$ and for every $n \geq n_0$ a subset S_n of the elements of ζ^n such that the total measure of the elements from S_n is $\geq 1 - \epsilon$ and for each element $C \in S_n$*

$$-n(h(T, \zeta) + \epsilon) < \log \mu(C) < -n(h(T, \zeta) - \epsilon).$$

Exercise 9.3.1. Let T be a measure-preserving transformation of a non-atomic measure space (X, \mathfrak{A}, μ) . For a finite partition ξ and $x \in X$, let $\xi_n(x)$ be the element of ξ^n containing x . Prove that $\mu(\xi^n(x)) \rightarrow 0$ as $n \rightarrow \infty$ for a.e. x and every non-trivial finite partition ξ if and only if all powers T^n , $n \in \mathbb{N}$, are ergodic.

Exercise 9.3.2. Prove the remaining statements of Proposition 9.3.2.

Exercise 9.3.3. Prove that every Lebesgue space has a generating sequence of partitions.

Exercise 9.3.4. If ζ is a partition of a finite metric space, then we define the *diameter* of ζ to be $\text{diam}(\zeta) = \sup_{C \in \zeta} \text{diam}(C)$. Prove that a sequence (ζ_n) of finite partitions of a compact metric space X with a non-atomic Borel measure μ is generating if the diameter of ζ_n tends to 0 as $n \rightarrow \infty$.

Exercise 9.3.5. Suppose a measure-preserving transformation T has a generator with k elements. Prove that $h(T) \leq \log k$.

Exercise 9.3.6. Prove the first two statements of Proposition 9.3.6.

Exercise 9.3.7. Show that if an invertible transformation T has a one-sided generator, then $h(T) = 0$.

Exercise 9.3.8. Prove Corollary 9.3.8.

9.4 Examples of Entropy Calculation

Let (X, d) be a compact metric space, and μ a non-atomic Borel measure on X . By Exercise 9.3.4, any sequence of finite partitions whose diameter tends to 0 is generating. We will use this fact repeatedly in computing the metric entropy of some topological maps.

Rotations of S^1 . Let λ be the Lebesgue measure on S^1 . If α is rational, then $R_\alpha^n = \text{Id}$ for some n , so $h_\lambda(R_\alpha) = (1/n)h_\lambda(R_\alpha^n) = (1/n)h_\lambda(\text{Id}) = 0$. If α is irrational, let ξ_N be a partition of S^1 into N intervals of equal length. Then ξ_N^n consists of nN intervals, so $H(\xi_N^n) \leq \log nN$. Thus $h(R_\alpha, \xi_N) \leq \lim_{n \rightarrow \infty} (\log nN)/n = 0$. The collection of partitions ξ_N , $N \in \mathbb{N}$, is clearly generating, so $h(R_\alpha) = 0$.

This result can also be deduced from Exercise 9.3.7 by noting that every forward semiorbit is dense, so any non-trivial partition is a one-sided generator for R_α .

Expanding Maps. The partition

$$\xi = \{[0, 1/k), [1/k, 2/k), \dots, [(k-1)/k, 1)\}$$

is a generator for the expanding map $E_k: S^1 \rightarrow S^1$, since the elements of ξ^n are of the form $[i/k^n, (i+1)/k^n)$. We have

$$H(\xi^n) = - \sum \frac{1}{|k|^n} \log \left(\frac{1}{|k|^n} \right) = n \log |k|,$$

so $h_\lambda(E_k) = \log |k|$.

Shifts. Let $\sigma: \Sigma_m \rightarrow \Sigma_m$ be the one or two-sided shift on m symbols, and let $p = (p_1, \dots, p_m)$ be a non-negative vector with $\sum_{i=1}^m p_i = 1$. The vector p defines a measure on the alphabet $\{1, 2, \dots, m\}$. The associated product measure μ_p on Σ_m is called a *Bernoulli measure*. For a cylinder set, we have

$$\mu_p(C_{j_1, \dots, j_k}^{n_1, \dots, n_k}) = \prod_{i=1}^k p_{j_i}.$$

Let $\xi = \{C_j^0: j = 1, \dots, m\}$. Then ξ is a (one- or two-sided) generator for σ , since $\text{diam}(\bigvee_{i=0}^m \sigma^i \xi) \rightarrow 0$ with respect to the metric $d(\omega, \omega') = 2^{-l}$, where $l = \min\{|i|: \omega_i \neq \omega'_i\}$. Thus

$$h_{\mu_p}(\sigma) = h_{\mu_p}(\sigma, \xi) = \lim_{n \rightarrow \infty} \frac{1}{n} H\left(\bigvee_{i=0}^m \sigma^{-i} \xi\right)$$

For $i \neq j$, $\sigma^i \xi$ and $\sigma^j \xi$ are independent, so

$$H\left(\bigvee_{i=0}^m \sigma^{-i} \xi\right) = nH(\xi).$$

Thus $h_{\mu_p}(\sigma) = H(\xi) = -\sum p_i \log p_i$.

Recall that the topological entropy of σ is $\log m$. Thus the metric entropy of σ with respect to any Bernoulli measure is less than or equal to the topological entropy, and equality holds if and only if $p = (1/n, \dots, 1/n)$.

We next calculate the metric entropy of σ with respect to the Markov measures defined in §4.4. Let A be an irreducible $m \times m$ stochastic matrix, and q the unique positive left eigenvector whose entries sum to 1. Recall that for the measure $P = P_{A,q}$, the measure of a cylinder set is

$$P(C_{j_0, j_1, \dots, j_k}^{n, n+1, \dots, n+k}) = q_{j_0} \prod_{i=0}^{k-1} A_{j_i j_{i+1}}.$$

By Proposition 9.3.1, we have $h_P(\sigma, \xi) = \lim_{n \rightarrow \infty} H(\xi | \sigma^{-1}(\xi^n))$. By definition,

$$H(\xi | \sigma^{-1}(\xi^n)) = -\sum_{C \in \xi, D \in \sigma^{-1}(\xi^n)} P(C \cap D) \log \frac{P(C \cap D)}{P(D)}.$$

For $C = C_{j_0}^0 \in \xi$ and $D = C_{j_1, \dots, j_n}^{1, \dots, n} \in \sigma^{-1}(\xi^n)$, we have

$$P(C \cap D) = q_{j_0} \prod_{i=0}^{n-1} A_{j_i j_{i+1}} \quad \text{and} \quad P(D) = q_{j_1} \prod_{i=1}^{n-1} A_{j_i j_{i+1}}.$$

Thus

$$\begin{aligned}
 H(\xi|\sigma^{-1}(\xi^n)) &= - \sum_{j_0, j_1, \dots, j_n=1}^m q_{j_0} \prod_{i=1}^{n-1} A_{j_i, j_{i+1}} \log \left(\frac{q_{j_0} A_{j_0, j_1}}{q_{j_1}} \right) \\
 &= - \sum_{j_0, j_1, \dots, j_n=1}^m q_{j_0} \prod_{i=1}^{n-1} A_{j_i, j_{i+1}} (\log A_{j_0, j_1} + \log q_{j_0} - \log q_{j_1}).
 \end{aligned}
 \tag{9.2}$$

Using the identities $\sum_{i=1}^n q_i A_{i,k} = q_k$ and $\sum_{k=1}^n A_{i,k} = 1$, we find that

$$\sum_{j_0, j_1, \dots, j_n} q_{j_0} \prod_{i=0}^{n-1} A_{j_i, j_{i+1}} \log A_{j_0, j_1} = \sum_{j_0, j_1} q_{j_0} A_{j_0, j_1} \log A_{j_0, j_1}, \tag{9.3}$$

$$\sum_{j_0, j_1, \dots, j_n} q_{j_0} \prod_{i=0}^{n-1} A_{j_i, j_{i+1}} \log q_{j_0} = \sum_{j_0} q_0 \log q_{j_0}, \tag{9.4}$$

$$\sum_{j_0, j_1, \dots, j_n=1}^m q_{j_0} \prod_{i=1}^{n-1} A_{j_i, j_{i+1}} \log q_{j_1} = \sum_{j_1} q_{j_1} \log q_{j_1}. \tag{9.5}$$

It follows from (9.2)–(9.5) that

$$h_P(\sigma) = - \sum_{j_0, j_1} q_{j_0} A_{j_0, j_1} \log A_{j_0, j_1}.$$

There are many Markov measures for a given subshift. We now construct a special Markov measure, called the *Shannon–Parry measure*, that maximizes the entropy. By the results of the next section, a Markov measure maximizes the entropy if and only if the metric entropy with respect to the measure is the same as the topological entropy of the underlying subshift.

Let B be a primitive matrix of zeros and ones. Let λ be the largest positive eigenvalue of B , and let q be a positive left eigenvector of B with eigenvalue λ . Let v be a positive right eigenvector of B with eigenvalue λ normalized so that $\langle q, v \rangle = 1$. Let V be the diagonal matrix whose diagonal entries are the coordinates of v , i.e., $V_{ij} = \delta_{ij} v_j$. Then $A = \lambda^{-1} V^{-1} B V$ is a stochastic matrix: all elements of A are positive, and the rows sum to 1. The elements of A are $A_{ij} = \lambda^{-1} v_i^{-1} B_{ij} v_j$. Let $p = qV = (q_1 v_1, \dots, q_n v_n)$. Then p is a positive left eigenvector of A with eigenvalue 1, and $\sum_{i=1}^n p_i = \langle q, v \rangle = 1$.

The Markov measure $P = P_{A,p}$ is called the Shannon–Parry measure for the subshift σ_A . Recall that while P is defined on the full shift space Σ , its support is the subspace Σ_A . Thus $h_P(\sigma_A) = h_P(\sigma)$. Using the properties

$qB = \lambda q$, $\langle q, v \rangle = 1$, and $B_{ij} \log B_{ij} = 0$, we have

$$\begin{aligned}
 h_P(\sigma) &= - \sum_{i,j} p_i A_{ij} \log A_{ij} \\
 &= - \sum_{i,j} q_i v_i \lambda^{-1} v_i^{-1} B_{ij} v_j \log (\lambda^{-1} v_i^{-1} B_{ij} v_j) \\
 &= - \sum_{i,j} \lambda^{-1} q_i v_j B_{ij} \log (\lambda^{-1} v_i^{-1} B_{ij} v_j) \\
 &= \sum_{i,j} \lambda^{-1} q_i v_j B_{ij} \log \lambda + \sum_{i,j} \lambda^{-1} q_i v_j B_{ij} (\log v_i - \log B_{ij} v_j) \\
 &= \log \lambda + \sum_j q_j v_j \log v_j - \sum_{i,j} \lambda^{-1} q_i v_j B_{ij} \log v_j \\
 &\quad - \sum_{i,j} \lambda^{-1} q_i v_j B_{ij} \log B_{ij} \\
 &= \log \lambda + \sum_j v_j q_j \log v_j - \sum_i v_i q_i \log v_i = \log \lambda.
 \end{aligned}$$

Thus $h_P(\sigma_A) = \log \lambda$, which is the topological entropy of σ_A (Proposition 3.4.1).

Toral Automorphisms. We consider only the two-dimensional case. Let $A: \mathbb{T}^2 \rightarrow \mathbb{T}^2$ be a hyperbolic toral automorphism. The Markov partition constructed in §5.12 gives a (measurable) semiconjugacy $\phi: \Sigma_A \rightarrow \mathbb{T}^2$ between a subshift of finite type and A . Since the image of the Lebesgue measure under ϕ^* is the Parry measure, the metric entropy of A (with respect to the Lebesgue measure) is the logarithm of the largest eigenvalue of A (Exercise 9.4.1).

Exercise 9.4.1. Let A be a hyperbolic toral automorphism. Prove that the image of the Lebesgue measure on \mathbb{T}^2 under the semiconjugacy ϕ is the Parry measure, and calculate the metric entropy of A .

9.5 Variational Principle¹

In this section, we establish the *variational principle for metric entropy* [Din71], [Goo69], which asserts that for a homeomorphism of a compact metric space, the topological entropy is the supremum of the metric entropies for all invariant probability measures.

¹ The proof of the variational principle below follows the argument of M. Misiurewicz [Mis76]; see also [KH95] and [Pet89].

Let f be a homeomorphism of a compact metric space X , and \mathcal{M} the space of Borel probability measures on X .

LEMMA 9.5.1. *Let $\mu, \nu \in \mathcal{M}$ and $t \in (0, 1)$. Then for any measurable partition of ξ of X ,*

$$tH_\mu(\xi) + (1 - t)H_\nu(\xi) \leq H_{t\mu + (1-t)\nu}(\xi).$$

Proof. The proof is a straightforward consequence of the convexity of $x \log x$ (Exercise 9.5.1). \square

For a partition $\xi = \{A_1, \dots, A_k\}$, define the *boundary* of ξ to be the set $\partial\xi = \bigcup_{i=1}^k \partial A_i$, where $\partial A = \overline{A} \cap \overline{X - A}$.

LEMMA 9.5.2. *Let $\mu \in \mathcal{M}$. Then:*

1. *for any $x \in X$ and $\delta > 0$, there is $\delta' \in (0, \delta)$ such that $\mu(\partial B(x, \delta')) = 0$;*
2. *for any $\delta > 0$, there is a finite measurable partition $\xi = \{C_1, \dots, C_k\}$ with $\text{diam}(C_i) < \delta$ for all i and $\mu(\partial\xi) = 0$;*
3. *if $\{\mu_n\} \subset \mathcal{M}$ is a sequence of probability measures that converges to μ in the weak* topology, and A is a measurable set with $\mu(\partial A) = 0$, then $\mu(A) = \lim_{n \rightarrow \infty} \mu_n(A)$.*

Proof. Let $S(x, \delta) = \{y \in X : d(x, y) = \delta\}$. Then $B(x, \delta) = \bigcup_{0 \leq \delta' < \delta} S(x, \delta')$. This is an uncountable union, so at least one of these must have measure 0. Since $\partial B(x, \delta) \subset S(x, \delta)$, statement 1 follows.

To prove statement 2, let $\{B_1, \dots, B_k\}$ be an open cover by balls of radius less than $\delta/2$ and $\mu(\partial B_i) = 0$. Let $C_1 = \overline{B_1}$, $C_2 = \overline{B_2} \setminus \overline{B_1}$, $C_i = \overline{B_i} \setminus \bigcup_{j=1}^{i-1} \overline{B_j}$. Then $\xi = \{C_1, \dots, C_k\}$ is a partition, and $\partial\xi = \bigcup \partial C_i \subset \bigcup_{i=1}^k \partial B_i$.

To prove statement 3, let A be a measurable set with $\mu(\partial A) = 0$. Since X is a normal topological space, there is a sequence $\{f_k\}$ of non-negative continuous functions on X such that $f_k \searrow \chi_A$. Then, for fixed k ,

$$\overline{\lim}_{n \rightarrow \infty} \mu_n(A) \leq \overline{\lim}_{n \rightarrow \infty} \mu_n(\overline{A}) \leq \lim_{n \rightarrow \infty} \mu_n(f_k) = \mu(f_k).$$

Taking the limit as $k \rightarrow \infty$, we obtain

$$\overline{\lim}_{n \rightarrow \infty} \mu_n(A) \leq \lim_{k \rightarrow \infty} \mu(f_k) = \mu(\overline{A}) = \mu(A).$$

Similarly,

$$\overline{\lim}_{n \rightarrow \infty} \mu_n(X \setminus A) \leq \mu(X \setminus A),$$

from which the result follows. \square

Let $|E|$ denote the cardinality of a finite set E .

LEMMA 9.5.3. *Let E_n be an (n, ϵ) -separated set, $\nu_n = (1/|E_n|) \sum_{x \in E_n} \delta_x$, and $\mu_n = \frac{1}{n} \sum_{i=0}^{n-1} f_*^i \nu_n$. If μ is any weak* accumulation point of $\{\mu_n\}_{n \in \mathbb{N}}$, then μ is f -invariant and*

$$\overline{\lim}_{n \rightarrow \infty} \frac{1}{n} \log |E_n| \leq h_\mu(f).$$

Proof. Let μ be an accumulation point of $\{\mu_n\}_{n \in \mathbb{N}}$. Then μ is clearly f -invariant.

Let ξ be a measurable partition with elements of diameter less than ϵ and $\mu(\partial\xi) = 0$. If $C \in \xi^n$, then $\nu_n(C) = 0$ or $1/|E_n|$, since C contains at most one element of E_n . Thus $H_{\nu_n}(\xi^n) = \log |E_n|$.

Fix $0 < q < n$, and $0 \leq k < q$. Let $a(k) = \lfloor \frac{n-k}{q} \rfloor$.

Let $S = \{k + rq + i : 0 \leq r < a(k), 0 \leq i < q\}$, and let T be the complement of S in $\{0, 1, \dots, n-1\}$. The cardinality of T is at most $k + q - 1 \leq 2q$. Since

$$\xi^n = \bigvee_{i=0}^{n-1} f^{-i} \xi = \left(\bigvee_{r=0}^{a(k)-1} f^{-rq-k} \xi^q \right) \vee \left(\bigvee_{i \in T} f^{-i} \xi \right),$$

it follows that

$$\begin{aligned} \log |E_n| = H_{\nu_n}(\xi^n) &\leq \sum_{r=0}^{a(k)-1} H_{\nu_n}(f^{-(rq+k)} \xi^q) + \sum_{i \in T} H_{\nu_n}(f^{-i} \xi) \\ &\leq \sum_{r=0}^{a(k)-1} H_{f_*^{rq+k} \nu_n}(\xi^q) + 2q \log |\xi|. \end{aligned}$$

Summing over k and using Lemma 9.5.1, we get

$$\begin{aligned} \frac{q}{n} \log |E_n| &= \frac{1}{n} \sum_{k=0}^{q-1} H_{\nu_n}(\xi^n) \leq \sum_{k=0}^{q-1} \left(\sum_{r=0}^{a(k)-1} \frac{1}{n} H_{f_*^{rq+k} \nu_n}(\xi^q) \right) + \frac{2q^2}{n} \log |\xi| \\ &\leq H_{\mu_n}(\xi^q) + \frac{2q^2}{n} \log |\xi| \end{aligned}$$

Thus, by Lemma 9.5.2(3), for fixed q ,

$$\overline{\lim}_{n \rightarrow \infty} \frac{1}{n} \log |E_n| \leq \lim_{n \rightarrow \infty} \frac{1}{q} H_{\mu_n}(\xi^q) = \frac{1}{q} H_\mu(\xi^q).$$

Letting $q \rightarrow \infty$, we get $\overline{\lim}_{n \rightarrow \infty} \frac{1}{n} \log |E_n| \leq h_\mu(f, \xi)$. □

THEOREM 9.5.4 (Variational Principle). *Let f be a homeomorphism of a compact metric space (X, d) . Then $h_{\text{top}}(f) = \sup\{h_\mu(f) \mid \mu \in \mathcal{M}_f\}$.*

Proof. Lemma 9.5.3 shows that $h_{\text{top}}(f) \leq \sup_{\mu \in \mathcal{M}_f} h_\mu(f)$, so we need only demonstrate the opposite inequality.

Let $\mu \in \mathcal{M}_f$ be an f -invariant Borel probability measure on X , and $\xi = \{C_1, \dots, C_k\}$ a measurable partition of X . By the regularity of μ and Lemma 9.3.3, we may choose compact sets $B_i \subset C_i$ so that the partition $\beta = \{B_0 = X \setminus \bigcup_{i=1}^k B_i, B_1, \dots, B_k\}$ satisfies $H(\xi|\beta) < 1$. Thus

$$h_\mu(f, \xi) \leq h_\mu(f, \beta) + H_\mu(\xi|\beta) \leq h_\mu(f, \beta) + 1.$$

The collection $\mathcal{B} = \{B_0 \cup B_1, \dots, B_0 \cup B_k\}$ is a covering of X by open sets. Moreover, $|\beta^n| \leq 2^n |\mathcal{B}^n|$, since each element of \mathcal{B}^n intersects at most two elements of β . Thus

$$H_\mu(\beta^n) \leq \log |\beta^n| \leq n \log 2 + \log |\mathcal{B}^n|$$

Let δ_0 be the Lebesgue number of \mathcal{B} , i.e., the supremum of all δ such that for all $x \in X$, $B(x, \delta)$ is contained in some $B_0 \cup B_i$. Then δ_0 is also the Lebesgue number of \mathcal{B}^n with respect to the metric d_n .

No subcollection of \mathcal{B} covers X , and the same is true of \mathcal{B}^n . Thus each element $C \in \mathcal{B}$ contains a point x_C that is not contained in any other element, so $B(x_C, \delta_0, n) \subset C$. It follows that the collection of all x_C is an (n, δ_0) -separated set. Thus $\text{sep}(n, \delta_0, f) \geq |\mathcal{B}^n|$, from which it follows that

$$\begin{aligned} h(f, \delta_0) &= \overline{\lim}_{n \rightarrow \infty} \frac{1}{n} \log(\text{sep}(n, \delta_0, f)) \geq \overline{\lim}_{n \rightarrow \infty} \frac{1}{n} \log |\mathcal{B}^n| \\ &\geq \overline{\lim}_{n \rightarrow \infty} \frac{1}{n} (\log |\beta^n| - n \log 2) \geq \overline{\lim}_{n \rightarrow \infty} \frac{1}{n} H_\mu(\beta^n) - \log 2 \\ &= h_\mu(f, \beta) - \log 2 \geq h_\mu(f, \xi) - \log 2 - 1 \end{aligned}$$

We conclude that $h_\mu(f) = h_\mu(f^n)/n \leq \frac{1}{n}(h_{\text{top}}(f^n) + \log 2 + 1)$ for all $n > 0$. Letting $n \rightarrow \infty$, we see that $h_\mu(f) \leq h_{\text{top}}(f)$ for all $\mu \in \mathcal{M}$, which proves the theorem. \square

Exercise 9.5.1. Prove Lemma 9.5.1.

Exercise 9.5.2. Let f be an expansive map of a compact metric space with expansiveness constant δ_0 . Show that f has a *measure of maximal entropy*, i.e., there is $\mu \in \mathcal{M}_f$ such that $h_\mu(f) = h_{\text{top}}(f)$. (Hint: Start with a measure supported on an (n, ϵ) -separated set, where $\epsilon \leq \delta_0$.)

Bibliography

- [AF91] Roy Adler and Leopold Flatto. Geodesic flows, interval maps, and symbolic dynamics. *Bull. Amer. Math. Soc. (N.S.)*, 25(2):229–334, 1991.
- [Ahl73] Lars V. Ahlfors. *Conformal invariants: topics in geometric function theory*. McGraw-Hill Series in Higher Mathematics. McGraw-Hill Book Co., New York, 1973.
- [Ano67] D. V. Anosov. Tangential fields of transversal foliations in y -systems. *Math. Notes*, 2:818–823, 1967.
- [Ano69] D. V. Anosov. *Geodesic flows on closed Riemann manifolds with negative curvature*. American Mathematical Society, Providence, RI, 1969.
- [Arc70] Ralph G. Archibald. *An introduction to the theory of numbers*. Charles E. Merrill Publishing Co., Columbus, OH, 1970.
- [AS67] D. V. Anosov and Ya. G. Sinai. Some smooth ergodic systems. *Russian Math. Surveys*, 22(5):103–168, 1967.
- [AW67] R. L. Adler and B. Weiss. Entropy, a complete metric invariant for automorphisms of the torus. *Proc. Nat. Acad. Sci. U.S.A.*, 57:1573–1576, 1967.
- [BC91] M. Benedicks and L. Carleson. The dynamics of the Hénon map. *Ann. of Math. (2)*, 133:73–169, 1991.
- [Bea91] Alan F. Beardon. *Iteration of rational functions*. Springer-Verlag, New York, 1991.
- [Ber96] Vitaly Bergelson. Ergodic Ramsey theory – an update. In *Ergodic theory of \mathbf{Z}^d actions (Warwick, 1993–1994)*, pages 1–61. Cambridge Univ. Press, Cambridge, 1996.
- [Ber00] Vitaly Bergelson. Ergodic theory and Diophantine problems. In *Topics in symbolic dynamics and applications (Temuco, 1997)*, pages 167–205. Cambridge Univ. Press, Cambridge, 2000.
- [BG91] Carlos A. Berenstein and Roger Gay. *Complex variables*. Springer-Verlag, New York, 1991.
- [Bil65] Patrick Billingsley. *Ergodic theory and information*. John Wiley & Sons Inc., New York, 1965.
- [BL70] R. Bowen and O. E. Lanford, III. Zeta functions of restrictions of the shift transformation. In *Global Analysis (Proc. Sympos. Pure Math., Vol. XIV, Berkeley, CA, 1968)*, pages 43–49. American Mathematical Society, Providence, RI, 1970.

- [Bow70] Rufus Bowen. Markov partitions for Axiom a diffeomorphisms. *Amer. J. Math.*, 92:725–747, 1970.
- [Boy93] Mike Boyle. Symbolic dynamics and matrices. In *Combinatorial and graph-theoretical problems in linear algebra (Minneapolis, MN, 1991)*, IMA Vol. Math. Appl., volume 50, pages 1–38. Springer-Verlag, New York, 1993.
- [BP94] Abraham Berman and Robert J. Plemmons. *Nonnegative matrices in the mathematical sciences*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1994.
- [BP98] Sergey Brin and Lawrence Page. The anatomy of a large-scale hypertextual web search engine. In *Seventh International World Wide Web Conference (Brisbane, Australia, 1998)*. <http://www7.scu.edu.au/programme/fullpapers/1921/com1921.htm>, 1998.
- [CE80] Pierre Collet and Jean-Pierre Eckmann. *Iterated maps on the interval as dynamical systems*. Birkhäuser, Boston, MA, 1980.
- [CFS82] I. P. Cornfeld, S. V. Fomin, and Ya. G. Sinai. *Ergodic theory*, volume 245 of *Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, New York, 1982.
- [CG93] Lennart Carleson and Theodore W. Gamelin. *Complex dynamics*. Universitext: Tracts in Mathematics. Springer-Verlag, New York, 1993.
- [CH82] Shui Nee Chow and Jack K. Hale. *Methods of bifurcation theory*. Springer-Verlag, New York, 1982.
- [Con95] John B. Conway. *Functions of one complex variable. II*. Springer-Verlag, New York, 1995.
- [Den26] Arnaud Denjoy. Sur l'itération des fonctions analytique. *C. R. Acad. Sci. Paris Sér. A-B*, 182:255–257, 1926.
- [Dev89] Robert L. Devaney. *An introduction to chaotic dynamical systems*. Addison-Wesley, Redwood City, Calif., 1989.
- [DH82] Adrien Douady and John Hamal Hubbard. Itération des polynômes quadratiques complexes. *C. R. Acad. Sci. Paris Sér. I Math.*, 294(3):123–126, 1982.
- [Din71] Efim I. Dinaburg. On the relation among various entropy characterizistics of dynamical systems. *Math. USSR, Izvestia*, 5:337–378, 1971.
- [dMvS93] Wellington de Melo and Sebastian van Strien. *One-dimensional dynamics*. Springer-Verlag, Berlin, 1993.
- [Dou83] Adrien Douady. Systèmes dynamiques holomorphes. In *Bourbaki seminar, Vol. 1982/83*, pages 39–63. Soc. Math. France, Paris, 1983.
- [DS88] Nelson Dunford and Jacob T. Schwartz. *Linear operators. Part II*. John Wiley & Sons Inc., New York, 1988.
- [Fei79] Mitchell J. Feigenbaum. The universal metric properties of non-linear transformations. *J. Statist. Phys.*, 21(6):669–706, 1979.
- [Fol95] Gerald B. Folland. *A course in abstract harmonic analysis*. CRC Press, Boca Raton, FL, 1995.
- [Fri70] Nathaniel A. Friedman. *Introduction to ergodic theory*. Van Nostrand Reinhold Mathematical Studies, No. 29. Van Nostrand Reinhold Co., New York, 1970.
- [Fur63] H. Furstenberg. The structure of distal flows. *Amer. J. Math.*, 85:477–515, 1963.
- [Fur77] Harry Furstenberg. Ergodic behavior of diagonal measures and a theorem of Szemerédi on arithmetic progressions. *J. Analyse Math.*, 31:204–256, 1977.

- [Fur81a] H. Furstenberg. *Recurrence in ergodic theory and combinatorial number theory*. M. B. Porter Lectures. Princeton University Press, Princeton, NJ, 1981.
- [Fur81b] Harry Furstenberg. Poincaré recurrence and number theory. *Bull. Amer. Math. Soc. (N.S.)*, 5(3):211–234, 1981.
- [FW78] H. Furstenberg and B. Weiss. Topological dynamics and combinatorial number theory. *J. Analyse Math.*, 34:61–85 (1979), 1978.
- [Gan59] F. R. Gantmacher. *The theory of matrices. Vols. 1, 2*. Chelsea Publishing Co., New York, 1959. Translated by K. A. Hirsch.
- [GG73] M. Golubitsky and V. Guillemin. *Stable mappings and their singularities*. Graduate Texts in Mathematics, Vol. 14. Springer-Verlag, New York, 1973.
- [GH55] W. Gottschalk and G. Hedlund. *Topological Dynamics*. A.M.S. Colloquium Publications, volume XXXVI. Amer. Math. Soc., Providence, RI, 1955.
- [GLR72] R. L. Graham, K. Leeb, and B. L. Rothschild. Ramsey's theorem for a class of categories. *Advances in Math.*, 8:417–433, 1972.
- [GLR73] R. L. Graham, K. Leeb, and B. L. Rothschild. Errata: Ramsey's theorem for a class of categories. *Advances in Math.*, 10:326–327, 1973.
- [Goo69] L. Wayne Goodwyn. Topological entropy bounds measure-theoretic entropy. *Proc. Amer. Math. Soc.*, 23:679–688, 1969.
- [Hal44] Paul R. Halmos. In general a measure preserving transformation is mixing. *Ann. of Math. (2)*, 45:786–792, 1944.
- [Hal50] Paul R. Halmos. *Measure theory*. D. Van Nostrand Co., Princeton, NJ, 1950.
- [Hal60] Paul R. Halmos. *Lectures on ergodic theory*. Chelsea Publishing Co., New York, 1960.
- [Hel95] Henry Helson. *Harmonic analysis*. Henry Helson, Berkeley, CA, second edition, 1995.
- [Hén76] M. Hénon. A two-dimensional mapping with a strange attractor. *Comm. Math. Phys.*, 50:69–77, 1976.
- [Hir94] Morris W. Hirsch. *Differential topology*. Springer-Verlag, New York, 1994. Corrected reprint of the 1976 original.
- [HK91] Jack K. Hale and Hüseyin Koçak. *Dynamics and bifurcations*. Springer-Verlag, New York, 1991.
- [HW79] G. H. Hardy and E. M. Wright. *An introduction to the theory of numbers*. The Clarendon Press, Oxford University Press, New York, fifth edition, 1979.
- [Kat72] A. B. Katok. Dynamical systems with hyperbolic structure, pages 125–211, 1972. Three papers on smooth dynamical systems. Translations of the AMS (series 2), volume 11b, AMS, Providence, RI, 1981.
- [KH95] Anatole Katok and Boris Hasselblatt. *Introduction to the modern theory of dynamical systems*. Encyclopedia of Mathematics and Its Applications, volume 54. Cambridge University Press, Cambridge, 1995. With a supplementary chapter by Katok and Leonardo Mendoza.
- [Kin90] J. L. King. A map with topological minimal self-joinings in the sense of Del Junco. *Ergodic Theory Dynamical Systems*, 10(4):745–761, 1990.
- [Kol58] A. N. Kolmogorov. A new metric invariant of transient dynamical systems and automorphisms in Lebesgue spaces. *Dokl. Akad. Nauk SSSR (N.S.)*, 119:861–864, 1958.
- [Kol59] A. N. Kolmogorov. Entropy per unit time as a metric invariant of automorphisms. *Dokl. Akad. Nauk SSSR*, 124:754–755, 1959.

- [KR99] K. H. Kim and F. W. Roush. The Williams conjecture is false for irreducible subshifts. *Ann. of Math. (2)*, 149(2):545–558, 1999.
- [Kre85] Ulrich Krengel. *Ergodic theorems*. Walter de Gruyter & Co., Berlin, 1985. With a supplement by Antoine Brunel.
- [KvN32] B. O. Koopman and J. von Neumann. Dynamical systems of continuous spectra. *Proc. Nat. Acad. Sci. U.S.A.*, 18:255–263, 1932.
- [Lan84] Oscar E. Lanford, III. A shorter proof of the existence of the Feigenbaum fixed point. *Comm. Math. Phys.*, 96(4):521–538, 1984.
- [LM95] Douglas Lind and Brian Marcus. *An introduction to symbolic dynamics and coding*. Cambridge University Press, Cambridge, 1995.
- [Lor63] E. N. Lorenz. Deterministic non-periodic flow. *J. Atmos. Sci.*, 20:130–141, 1963.
- [LY75] Tien Yien Li and James A. Yorke. Period three implies chaos. *Amer. Math. Monthly*, 82(10):985–992, 1975.
- [Mañ88] Ricardo Mañé. A proof of the C^1 stability conjecture. *Inst. Hautes Études Sci. Publ. Math.*, 66:161–210, 1988.
- [Mat68] John N. Mather. Characterization of Anosov diffeomorphisms. *Nederl. Akad. Wetensch. Proc. Ser. A 71: Indag. Math.*, 30:479–483, 1968.
- [Mil65] J. Milnor. *Topology from the differentiable viewpoint*. University Press of Virginia, Charlottesville, VA, 1965.
- [Mis76] Michal Misiurewicz. A short proof of the variational principle for a F_+^n action on a compact space. In *International Conference on Dynamical Systems in Mathematical Physics (Rennes, 1975)*, pages 147–157. Astérisque, No. 40. Soc. Math. France, Paris, 1976.
- [Mon27] P. Montel. *Leçons sur les familles normales de fonctions analytiques et leurs applications*. Gauthier-Villars, Paris, 1927.
- [Moo66] C. C. Moore. Ergodicity of flows on homogeneous spaces. *Amer. J. Math.*, 88:154–178, 1966.
- [MRS95] Brian Marcus, Ron M. Roth, and Paul H. Siegel. Modulation codes for digital data storage. In *Different aspects of coding theory (San Francisco, CA, 1995)*, pages 41–94. American Mathematical Society, Providence, RI, 1995.
- [MT88] John Milnor and William Thurston. On iterated maps of the interval. In *Dynamical systems (College Park, MD, 1986–87)*, Lecture Notes in Mathematics, volume 1342, pages 465–563. Springer-Verlag, Berlin, 1988.
- [PdM82] Jacob Palis, Jr., and Welington de Melo. *Geometric theory of dynamical systems*. Springer-Verlag, New York, 1982. An introduction.
- [Pes77] Ya. Pesin. Characteristic Lyapunov exponents and smooth ergodic theory. *Russian Math. Surveys*, 32:4:55–114, 1977.
- [Pet89] Karl Petersen. *Ergodic theory*. Cambridge University Press, Cambridge, 1989.
- [Pon66] L. S. Pontryagin. *Topological groups*. Gordon and Breach Science Publishers, Inc., New York, 1966. Translated from the second Russian edition by Arlen Brown.
- [Que87] Martine Queffélec. *Substitution dynamical systems – spectral analysis*. Lecture Notes in Mathematics, volume 1294. Springer-Verlag, Berlin, 1987.
- [Rob71] J. W. Robbin. A structural stability theorem. *Ann. of Math. (2)*, 94:447–493, 1971.
- [Rob76] Clark Robinson. Structural stability of C^1 diffeomorphisms. *J. Differential Equations*, 22(1):28–73, 1976.

- [Rob95] Clark Robinson. *Dynamical systems*. Studies in Advanced Mathematics. CRC Press, Boca Raton, Fla., 1995.
- [Roh48] V. A. Rohlin. A “general” measure-preserving transformation is not mixing. *Dokl. Akad. Nauk SSSR (N.S.)*, 60:349–351, 1948.
- [Rok67] V. Rokhlin. Lectures on the entropy theory of measure preserving transformations. *Russian Math. Surveys*, 22:1–52, 1967.
- [Roy88] H. L. Royden. *Real analysis*. Macmillan, third edition, 1988.
- [Rud87] Walter Rudin. *Real and complex analysis*. McGraw-Hill Book Co., New York, third edition, 1987.
- [Rud91] Walter Rudin. *Functional analysis*. McGraw-Hill Book Co. Inc., New York, second edition, 1991.
- [Rue89] David Ruelle. *Elements of differentiable dynamics and bifurcation theory*. Academic Press Inc., Boston, Mass., 1989.
- [Sár78] A. Sárközy. On difference sets of sequences of integers. III. *Acta Math. Acad. Sci. Hungar.*, 31(3–4):355–386, 1978.
- [Sha64] A. N. Sharkovskii. Co-existence of cycles of a continuous mapping of the line into itself. *Ukrainian Math. J.*, 16:61–71, 1964.
- [Shi87] Mitsuhiro Shishikura. On the quasiconformal surgery of rational functions. *Ann. Sci. École Norm. Sup. (4)*, 20(1):1–29, 1987.
- [Sin59] Ya. G. Sinai. On the concept of entropy of a dynamical system. *Dokl. Akad. Nauk SSSR (N.S.)*, 124:768–771, 1959.
- [Sma67] S. Smale. Differentiable dynamical systems. *Bull. Amer. Math. Soc.*, 73:747–817, 1967.
- [Sul85] Dennis Sullivan. Quasiconformal homeomorphisms and dynamics. I. Solution of the Fatou–Julia problem on wandering domains. *Ann. of Math. (2)*, 122(3):401–418, 1985.
- [SW49] Claude E. Shannon and Warren Weaver. *The Mathematical Theory of Communication*. The University of Illinois Press, Urbana, IL, 1949.
- [Sze69] E. Szemerédi. On sets of integers containing no four elements in arithmetic progression. *Acta Math. Acad. Sci. Hungar.*, 20:89–104, 1969.
- [vS88] Sebastian van Strien. Smooth dynamics on the interval (with an emphasis on quadratic-like maps). In *New directions in dynamical systems*, pages 57–119. Cambridge Univ. Press, Cambridge, 1988.
- [Wal75] Peter Walters. *Ergodic theory – Introductory Lectures*. Lecture Notes in Mathematics, volume 458. Springer-Verlag, Berlin, 1975.
- [Wei73] Benjamin Weiss. Subshifts of finite type and sofic systems. *Monatsh. Math.*, 77:462–474, 1973.
- [Wil73] R. F. Williams. Classification of subshifts of finite type. *Ann. of Math. (2)*, 98:120–153, 1973. Errata, *ibid.*, 99 (1974), 380–381.
- [Wil84] R. F. Williams. Lorenz knots are prime. *Ergodic Theory Dynam. Systems*, 4(1):147–163, 1984.
- [Wol26] Julius Wolff. Sur l’itération des fonctions bornées. *C. R. Acad. Sci. Paris Sér. A-B*, 182:200–201, 1926.

Index

- (X, \mathfrak{A}, μ) , 70
- $[x, y]$, 128
- 1-step SFT, 56
- Δ , 73
- Δ_1 , 194
- \perp , 209
- \prec , 162, 172
- \sim , 3
- \vee , 208
- \wedge , 208
- α -limit point, 29
- α -limit set, 29
- ϵ -dense, 4
- ϵ -orbit, 107, 110
- $\zeta_f(z)$, 60
- ζ^n , 213
- θ -transversality, 147
- $\Lambda_\epsilon^s, \Lambda_\epsilon^u$, 116
- ν , 209
- ν_j , 173
- $\rho(\xi, \eta)$, 212
- $\rho(f)$, 154
- Σ_A^e , 56
- Σ_A^v , 56
- Σ_m , 7
- Σ_m^+ , 7
- σ -algebra, 69
- σ , 5, 7
- ω -limit point, 28
- ω -limit set, 28
- $\omega_f(x)$, 29
- $BA(\zeta)$, 192
- $B(x, r)$, 28
- \mathbb{C} , 1
- $C(X)$, 85
- $CL(R)$, 204
- C^k curve, 138
- C^k topology, 138
- $C^k(M, N)$, 138
- $C_0(X, \mathbb{C})$, 71
- $\text{Diff}^1(M)$, 117
- $d(\xi, \eta)$, 209
- d_n , 37
- E^s, E^u , 108
- E_m , 5
- $F(R)$, 193
- f -covers, 163
- G -extension, 88
- $H(\xi|\eta)$, 211
- $H(\zeta)$, 209
- $H_\epsilon(U)$, 98
- $H_w(U)$, 98
- $h(T) = h_\mu(T)$, 215
- $h(T, \zeta)$, 213
- $h(f)$, 37
- $I(C)$, 210
- $I \rightarrow J$, 163
- $i(x)$, 170
- $J(R)$, 193
- k -step SFT, 56
- $L^p(X, \mu)$, 71
- l -modal map, 170
- \mathcal{M} , 85
- \mathcal{M}_T , 85
- \mathbb{N} , 1
- \mathbb{N}_0 , 1
- $NW(f)$, 29, 114
- n -leader, 81
- $\mathcal{O}_f(x)$, 2
- P -graph, 165
- \mathbb{Q} , 1
- q_μ , 9

- $\mathcal{R}(f)$, 29
- R_α , 4
- \mathbb{R} , 1
- S^1 , 4
- S_f^∞ , 87
- S_f^r , 85, 87
- $Sf(x)$, 178
- supp, 72, 90
- U_T , 80
- $\text{Var}(f)$, 160
- $W^s(x)$, $W^u(x)$, 118, 122
- $W_\epsilon^s(x)$, $W_\epsilon^u(x)$, 121
- \mathbb{Z} , 1

- a.e., 71
- abelian group, 87
- absolutely continuous
 - foliation, 144
 - measure, 86
- adapted metric, 109
- address of a point, 170
- adjacency matrix, 8, 56
- Adler, 135
- affine subspace, 51
- allowed word, 8
- almost every, 71
- almost periodic point, 30
- almost periodic set, 47
- alphabet, 7, 54, 55
- analytic function, 191
- Anosov diffeomorphism, 108, 130, 141, 142
 - topological entropy of, 116
- Anosov's shadowing theorem, 113
- aperiodic transformation, 96
- arithmetic progression, 103
- atlas, 138
- atom, 71
- attracting periodic point, 192
- attracting point, 9, 25, 173
- attractor, 18, 25
 - Hénon, 25
 - hyperbolic, 18
 - Lorenz, 25
 - strange, 25
- automorphism
 - Bernoulli, 79, 209
 - group, 5
 - of a measure space, 70
- average
 - space, 80, 84
 - time, 80, 84
- Axiom A, 133
- backward invariant, set, 2
- base- m expansion, 5
- basic set, 133
- basin of attraction, 25, 173, 192
 - immediate, 192
- Bernoulli
 - automorphism, 79, 209
 - measure, 79, 219
- bifurcation, 183
 - codimension-one, 184
 - flip, 185
 - fold, 185
 - generic, 183
 - period-doubling, 185
 - saddle-node, 185
- Birkhoff ergodic theorem, 82, 93
- block code, 55, 67
- Borel
 - σ -algebra, 71
 - measure, 71, 72
- boundary of a partition, 222
- bounded distortion, 181
- bounded Jacobian, 145
- bounded variation, 160
- Bowen, 61, 137
- branch number, 192
- branch point, 192
- Breiman, 217

- Cantor set, 7, 16, 18, 25, 182, 196, 204
- ceiling function, 21
- chain, Markov, *see* Markov, chain
- chaos, 23
- chaotic behavior, 23
- chaotic dynamical system, 23
- character, 95
- circle endomorphism, 23, 35, 36
- circle rotation, 4, 35, 45, 95
 - ergodicity of, 77
 - irrational, 4, 30, 77, 90
 - is uniquely ergodic, 87
- clock drift, 67
- code, 55
- color, 51
- combinatorial number theory, 48
- common refinement, 208
- completion of a σ -algebra, 70
- composition matrix, 65
- conditional density, 144
- conditional entropy, 211
- conditional measure, 92

- cone
 - stable, 115
 - unstable, 115
- conformal map, 153
- conjugacy, 3
 - measure-theoretic, 70
 - topological, 28, 39, 55
- conjugate, 3
- constant-length substitution, 64
- continued fraction, 12, 91
- continuous (semi)flow, 28
- continuous spectrum, 95, 97
- continuous-time dynamical system, 2
- convergence in density, 99
- convergent, 91
- coordinate chart, 137
- coordinate system, local, 106
- covering map, 153
- $\text{cov}(n, \epsilon, f)$, 37
- critical point, 192
- critical value, 192
- cross-ratio, 180
- cross-section, of a flow, 22
- Curtis, 55
- cutting and stacking, 96
- Cvitanović–Feigenbaum equation, 189
- cylinder, 54
- cylinder set, 7, 78

- d -lim, 99
- $\text{deg}(R)$, 193
- degree, of a rational map, 193
- dendrite, 196
- Denjoy, 205
 - example, 161
 - theorem, 160
- dense
 - forward orbit, 32
 - full orbit, 32
 - orbit, 23, 32
- dense, ϵ -, 4
- derivative transformation, 72
- deterministic dynamical system, 23
- diameter, 37
 - of a partition, 218
- diffeomorphism, 106
 - Anosov, 130, 141
 - Axiom A, 133
 - group, 138
 - of the circle, 160
- differentiable
 - flow, 107
 - manifold, 106, 137
 - map, 107
 - vector field, 107
- differential equation, 19
- $\text{Diff}^k(M)$, 138
- Dirac measure, 86
- direct product, 3
- directed graph, 8, 56, 67
- discrete spectrum, 94–97
- discrete-time dynamical system, 1
- dist, 140
- distal
 - extension, 46
 - homeomorphism, 45
 - points, 45
- distribution, 139
 - Hölder-continuous, 143
 - integrable, 139
 - stable, 108
 - unstable, 108
- Douady, 205
- Douady–Hubbard theorem, 206
- doubling operator, 169
- dual group, 95
- dynamical system, 2
 - chaotic, 23
 - continuous-time, 2
 - deterministic, 23
 - differentiable, 106
 - discrete-time, 1
 - ergodic, 4
 - hyperbolic, 22, 106
 - minimal, 4, 29
 - symbolic, 54
 - topological, 28
 - topologically mixing, 33
 - topologically transitive, 31
- dynamics
 - symbolic, 6, 54
 - topological, 28

- edge shift, 57
- elementary strong shift equivalence, 62
- embedded submanifold, 139
- embedding of a subshift, 55
- endomorphism
 - expanding, 5
 - group, 5

- entropy
 - conditional, 211
 - horseshoe, 125
 - measure-theoretic, 208, 213, 214
 - metric, 208, 213, 214
 - of a partition, 209
 - of a transformation, 214
 - topological, 36, 37, 116, 125
 - of an SFT, 60
- equicontinuous
 - extension, 46
 - homeomorphism, 45
- equivalent partitions, 208
- ergodic
 - dynamical system, 4
 - hypothesis, 69, 80
 - measure, 86
 - theorem, von Neumann, 80
 - theory, 69
 - transformation, 73
- ergodicity, 73, 75, 77
- essentially, 71
- essentially invariant
 - function, 73
 - set, 73
- even system of Weiss, 66
- eventually periodic point, 2, 171
- exceptional point, 201
- exceptional set, 201
- expanding endomorphism, 5, 30, 34, 107
 - is mixing, 77
- expanding map, 5
 - entropy of, 218
- expansive
 - homeomorphism, 35, 40
 - map, 35
 - positively, 35
- expansiveness constant, 35
- exponent
 - Hölder, 143
 - Lyapunov, 23
- extension, 3, 140
 - distal, 46
 - equicontinuous, 46
 - group, 88
 - isometric, 41, 46
 - measure-theoretic, 70
- extreme point, 86
- factor, 3, 55, 140
- factor map, 3
- Fatou, 200
- Fatou set, 193
- Fatou theorem, 204
- Feigenbaum constant, 189
- Feigenbaum phenomenon, 189
- fiber, 140
- fiber bundle, 140
- first integral, 20, 21
- first return map, 22, 72
- fixed point, 2
- $\text{Fix}(f)$, 60
- flip bifurcation, 185
- flow, 2, 19
 - continuous, 28
 - differentiable, 107
 - entropy of, 41
 - gradient, 20
 - measurable, 70, 74
 - measure preserving, 70
 - under a function, 21
 - weak mixing, 75
- fold bifurcation, 185
- foliation, 139
 - absolutely continuous, 144
 - coordinate chart, 139
 - leaf of, 139
 - stable, 14, 130
 - transversely absolutely continuous, 145
 - unstable, 14, 130
- forbidden word, 8
- forward invariant set, 2
- Fourier series, 78
- fractal, 191, 193, 202
- Frobenius theorem, 59
- full m -shift, 55
- full measure, 69
- full one-sided shift, 7, 33
- full shift, 35, 36
- full two-sided shift, 7, 33
- function
 - analytic, 191
 - ceiling, 21
 - essentially invariant, 73
 - generating, 60
 - Hamiltonian, 21
 - Lyapunov, 20
 - meromorphic, 191
 - strictly invariant, 74
 - zeta, 60
- Furstenberg, 49, 101–103

- Gauss, 11
 - measure, 12–13, 90–94
 - transformation, 11–13, 90–94
- generalized eigenspaces, 42
- generating function, 60
- generating sequence of partitions, 215
- generator
 - one-sided, 216
 - two-sided, 216
- generic bifurcation, 183
- generic point, 88
- global stable manifold, 122
- global unstable manifold, 122
- Google,™ 103
- gradient flow, 20
- graph transform, 119
- group
 - automorphism, 5
 - dual, 95
 - endomorphism, 5, 14
 - extension, 88
 - of characters, 95
 - topological, 4, 95
 - translation, 4, 35, 45, 87
- Haar measure, 88
- Hadamard, 54
- Hadamard–Perron theorem, 118
- Halmos, 97
- Halmos–von Neumann theorem, 96
- Hamiltonian
 - function, 21
 - system, 20
- Hedlund, 55
- Hénon attractor, 25
- heteroclinic point, 128
- heteroclinically related, 133
- higher block presentation, 56, 67
- Hölder
 - constant, 143
 - continuity, 142, 143
 - exponent, 143
- holonomy map, 145
- homeomorphism, 28
 - distal, 45
 - equicontinuous, 45
 - expansive, 35, 40
 - minimal, 29
 - of the circle, 153
 - pointwise almost periodic, 47
- homoclinic point, 125
- Hopf, 151
- Hopf's argument, 141
- horseshoe, 34–36, 108, 124, 125, 128
 - entropy of, 125
 - linear, 15, 16
 - non-linear, 124
- Hubbard, 205
- hyperbolic
 - attractor, 18
 - dynamical system, 15, 106
 - fixed point, 110
 - periodic point, 110
 - set, 108
- hyperbolic toral automorphism, 14, 22, 33, 35, 36, 41, 108, 130, 135
 - is mixing, 77
 - metric entropy of, 221
- Hyperbolicity, 106, 108
- immediate basin of attraction, 192
- Inclination lemma, 123
- independent partitions, 209
- induced transformation, 73
- information, 210
- integrable distribution, 139
- integrable foliations, 151
- integral hull, 151
- Internet, 103
- intersection of partitions, 208
- intersymbol interference, 67
- invariant
 - kneading, 173
 - measure, 70, 85
 - set, 2
- IP-system, 49
- irrational circle rotation, 4, 33, 77, 87, 90
- irrationally neutral periodic point, 192
- irreducible
 - matrix, 57
 - substitution, 64
- isolated fixed point, 10
- isometric extension, 41, 46
- isometry, 28, 34, 36, 114
 - entropy of, 38
- isomorphism, 3
 - measure-theoretic, 70
 - of subshifts, 55
 - of topological dynamical systems, 28
- itinerary, 54, 170

- Jacobian, 145
- Julia set, 193, 200

- Katok, 111
- Kim, 63
- kneading invariant, 173
- Kolmogorov, 208, 216
- Kolmogorov–Sinai theorem, 216
- Koopman–von Neumann theorem, 99
- Krein–Milman theorem, 86
- Krylov–Bogolubov theorem, 85

- lag, 63
- Lambda lemma, 123
- Lanford, 61
- lap, 170
- leader, 81
- leaf, 139
- leaf, local, 139
- Lebesgue number, 224
- Lebesgue space, 71
- left translation, 4
- lemma
 - inclination, 123
 - lambda, 123
 - Rokhlin–Halmos, 96
- length distortion, 180
- Li, 164
- linear fractional transformation, 180
- linear functional, 85
- Lipshitz function, 160
- local
 - coordinate system, 106
 - leaf, 139
 - product structure, 129
 - stable manifold, 121
 - transversal, 139
 - unstable manifold, 121
- locally maximal
 - hyperbolic set, 128
 - invariant set, 16
- log, 209
- Lorenz, 26
- Lorenz attractor, 25
- Lyapunov
 - exponent, 23
 - function, 20
 - metric, 109
- Lyndon, 55

- main cardioid, 206
- Mandelbrot set, 206

- Mañé, 134
- manifold
 - differentiable, 106, 137
 - Riemannian, 106, 140
 - stable, 14, 118, 122
 - unstable, 14, 118, 122
- map
 - covering, 153
 - expanding, 5
 - first return, 72
 - l -modal, 170
 - measurable, 70
 - measure preserving, 70
 - non-singular, 70
 - piecewise-monotone, 170
 - Poincaré, 22, 72
 - quadratic, 9, 35, 171, 176, 178, 179, 182, 205
 - rational, 191
 - tent, 176
 - unimodal, 170
- Markov
 - chain, 78, 104, 105
 - graph, 164, 165
 - measure, 57, 78, 104, 105, 219
 - partition, 134
- matrix
 - adjacency, 56
 - composition, 65
 - irreducible, 57
 - non-negative, 57
 - positive, 57
 - primitive, 57, 79
 - reducible, 57
 - shift-equivalent, 62
 - stochastic, *see* stochastic matrix
 - strong shift-equivalent, 62
- maximal almost periodic set, 47
- maximal ergodic theorem, 82
- McMillan, 217
- measurable
 - flow, 70, 74
 - map, 70
 - set, 70
- measure, 69
 - absolutely continuous, 86
 - Bernoulli, 79, 219
 - Borel, 72
 - conditional, 92
 - Dirac, 86
 - ergodic, 86
 - Gauss, *see also* Gauss measure
 - Haar, 88

- invariant, 70, 85
- Markov, 57, 78, *see also* Markov measure
- of maximal entropy, 224
- regular, 71
- Shannon–Parry, 220
- smooth, 151
- space, 69
- spectral, 98
- measure-preserving
 - flow, 70
 - transformation, 70
- measure-theoretic
 - entropy, 213, 215
 - extension, 70
 - isomorphism, 70
- meromorphic function, 191
- metric
 - adapted, 109
 - entropy, 213, 215
 - Lyapunov, 109
 - Riemannian, 106, 139
 - space, 28
- minimal
 - action of a group, 49
 - dynamical system, 4, 29
 - homeomorphism, 29
 - period, 2
 - set, 29, 90
- minimum principle, 179
- mixing, 74, 77
 - transformation, 74
 - weak, 75, 100, 152
- Möbius transformation, 180, 192
- mod 0, 71
- modified frequency modulation, 68
- modular function, 198
- monochromatic, 51
- Montel theorem, 197
- Morse sequence, 64
- multiple recurrence theorem, 103
- multiple recurrence, topological, 49
- multiplier, 192

- negative semiorbit, 2
- Newton method, 196
- nilmanifold, 130
- non-atomic Lebesgue space, 71
- non-linear horseshoe, 124
- non-negative definite sequence, 97
- non-singular map, 70
- non-singular transformation, 70
- non-wandering point, 29

- normal family, 193, 197
 - at a point, 193
- normal number, 84
- null set, 69
- number theory, 49, 92, 103

- omit a point, 197
- one-parameter group, 107
- one-sided shift, 35
- orbit, 2
 - dense, 23, 32
 - periodic, 2
 - uniformly distributed, 23, 89

- Parry, 220
- partition, 208
 - boundary of, 222
 - Markov, 134
 - of an interval, 164
- partitions, independent, 209
- pendulum, 19
- $\text{Per}(f)$, 114
- period, 2
- period, minimal, 2
- period-doubling bifurcation, 185
- periodic
 - orbit, 2
 - point, 2, 162, 198
 - attracting, 10, 192
 - irrationally neutral, 192
 - rationally neutral, 192
 - repelling, 10, 192
 - superattracting, 192
- Perron theorem, 58
- Perron–Frobenius theorem, 57
- piecewise-monotone map, 170
- plaque, 139
- Poincaré, 71
 - map, 22, 72
- Poincaré classification theorem, 158
- Poincaré recurrence theorem, 72, 103
- pointwise almost periodic homeomorphism, 47
- polynomial recurrence theorem, 102
- Pontryagin duality theorem, 95
- positive semiorbit, 2
- positive upper density, 101
- positively expansive, 35
- positively recurrent point, 29
- postcritical set, 204
- presentation, 67

- primitive
 - matrix, 57, 79, 104
 - substitution, 64, 90
 - transformation, 73
- probability
 - measure, 70
 - space, 70
- product
 - direct, 3
 - formula, 60
 - skew, 3
 - structure, local, 129
- projection, 3, 139, 140
- proximal points, 45
- pseudo-orbit, 110

- quadratic map, 9, 35, 171, 176, 178, 179, 182, 205
- quotient, 91

- Radon–Nikodym
 - derivative, 86
 - theorem, 86
- Ramsey, 48
- Ramsey theory, 48
- random behavior, 23
- random variable, 79
- rational map, 191
 - degree of, 193
- rationally neutral periodic point, 192
- rectangle, 134
- recurrence
 - measure-theoretic, 71
 - topological, 29, 72
- recurrent point, 29, 72
- reducible matrix, 57
- refinement, 208
- refining sequence of partitions, 215
- regular measure, 71
- relatively dense subset, 30
- repelling periodic point, 192
- repelling point, 9
- return time, 22
- Riemann sphere, 191
- Riemannian
 - manifold, 106, 140
 - metric, 106, 139
 - volume, 140
- Riesz, 80, 85
- Riesz representation theorem, 85
- right translation, 4

- Robbin, 134
- Robinson, 134
- Rokhlin, 97
- Rokhlin metric, 212
- Rokhlin–Halmos lemma, 96
- rotation number, 154
- rotation, entropy of, 218
- Roush, 63

- saddle–node bifurcation, 184
- Sárközy, 103
- Schwarzian derivative, 178
- search engine, 103
- self-similarity, 202
- semiconjugacy, 2
 - topological, 28
- semiflow, 2
- semiorbit
 - negative, 2
 - positive, 2
- sensitive dependence on initial conditions, 23, 117
- $\text{sep}(n, \epsilon, f)$, 37
- separated set, 37
- sequence of partitions
 - generating, 215
 - refining, 215
- SFT, 56
- shadowing, 107, 111, 113
- Shannon, 220
- Shannon–McMillan–Breiman theorem, 217
- Shannon–Parry measure, 220
- Sharkovsky ordering, 162
- Sharkovsky theorem, 163
- shift, 5, 7, 54, 60
 - edge, 57
 - entropy of, 219
 - equivalence, 62
 - full, 7, 35, 36
 - one-sided, 7
 - two-sided, 7
 - vertex, 8, 56
- shift equivalence, strong, 62
- Shishikura theorem, 200
- signed lexicographic ordering, 172
- simply transitive, 180, 192
- Sinai, 208, 216
- Singer theorem, 179
- skew product, 3, 140
- Smale, 133

- Smale's spectral decomposition theorem, 133
- smooth measure, 151
- sofic subshift, 66
- solenoid, 18, 35, 36, 44, 45, 108, 128
- space average, 84
- space, metric, 28
- $\text{span}(n, \epsilon, f)$, 37
- spanning set, 37
- spectral measure, 98
- spectrum
 - continuous, *see* continuous spectrum
 - discrete, *see* discrete spectrum
- stable
 - cone, 115
 - distribution, 108
 - foliation, 14, 130
 - manifold, 14, 118, 122
 - local, 121
 - set, 151
 - subspace, 108
- stationary sequence, 79
- stochastic matrix, 58, 78, 104
- stochastic process, 79
- strange attractor, 25
- strong mixing transformation, 74
- strong transversality condition, 134
- structural stability, 108, 117, 131
- Structural Stability Theorem, 134
- structurally stable diffeomorphism, 117
- submanifold, 106, 139
 - embedded, 139
- subshift, 8, 55, 134
 - entropy of, 219
 - generating function of, 60
 - of finite type, 56
 - one-sided, 55
 - sofic, 66
 - topological entropy of, 60
 - two-sided, 55
 - zeta function of, 61
- substitution, 64
 - constant length, 64
 - irreducible, 64
 - primitive, 64, 90
- Sullivan, 205
- superattracting periodic point, 192
- support, 72
- suspension, 21
- symbol, 7, 54
- symbolic dynamics, 6, 54
- symmetric difference, 73
- syndetic subset, 30
- Szemerédi, 103
- Szemerédi theorem, 103
- tangent
 - bundle, 138
 - map, 138
 - space, 138
 - vector, 138
- tent map, 114, 176
- terminal vertex, 9
- theorem
 - Anosov's shadowing, 113
 - Birkhoff ergodic, 82
 - Bowen–Lanford, 61
 - Curtis–Lyndon–Hedlund, 55
 - Denjoy, 160
 - Douady–Hubbard, 206
 - Fatou, 204
 - Frobenius, 59
 - Furstenberg–Weiss, 49
 - Hadarnard–Perron, 118
 - Halmos–von Neumann, 96
 - Kolmogorov–Sinai, 216
 - Koopman–von Neumann, 99
 - Krein–Milman, 86
 - Krylov–Bogolubov, 85
 - maximal ergodic, 82
 - Montel, 197
 - multiple recurrence, 103
 - Perron, 58
 - Perron–Frobenius, 57
 - Poincaré classification, 158
 - Poincaré recurrence, 72
 - polynomial recurrence, 102
 - Pontryagin duality, 95
 - Riesz representation, 85
 - Sárközy, 103
 - Shannon–McMillan–Breiman, 217
 - Sharkovsky, 163
 - Shishikura, 200
 - Singer, 179
 - Smale's spectral decomposition, 133
 - structural stability, 134
 - Szemerédi, 103
 - van der Waerden, 49
 - von Neumann ergodic, 80
 - Weyl, 89
 - Williams, 62
- time average, 84, 88

- time- t map, 2
- times- m map, 5
- topological
 - conjugacy, 28, 39
 - dynamical system, 28
 - dynamics, 28
 - entropy, 36, 37, 60, 116, 125
 - group, 4, 95
 - Markov chain, 56
 - mixing, 33, 76
 - properties, 28
 - recurrence, 29, 72
 - semiconjugacy, 28
 - transitivity, 31, 33, 76, 88
- topology, C^k , 138
- toral automorphism, *see* hyperbolic toral automorphism
- transformation, 70
 - aperiodic, 96
 - derivative, 72
 - entropy of, 215
 - ergodic, 73
 - Gauss, *see* Gauss transformation
 - induced, 73
 - linear fractional, 180
 - measure-preserving, 70
 - mixing, 74
 - Möbius, 180, 192
 - non-singular, 70
 - primitive, 73
 - strong mixing, 74
 - uniquely ergodic, 87
 - weak mixing, 75
- transition probability, 78
- translation, 90
 - group, 4
 - left, 4
 - right, 4
- translation-invariant metric, 44
- transversal, 139
 - local, 139
- transverse, 139
- transverse homoclinic point, 125
- transverse submanifolds, 122
- transversely absolutely continuous foliation, 145
- trapping region, 25
- turning point, 170
- uniform convergence, 88
- uniformly distributed, 89, 90
- unimodal map, 170
- uniquely ergodic, 87–90
- unitary operator, 80
- unstable
 - cone, 115
 - distribution, 108
 - foliation, 14, 130
 - manifold, 14, 118, 122
 - local, 121
 - set, 151
 - subspace, 108
- upper density, 99
- van der Waerden's theorem, 49
- variational principle for metric entropy, 221
- vector
 - non-negative, 57
 - positive, 57
- vector field, 107, 139
- vertex shift, 8, 56
 - one-sided, 8
 - two-sided, 8
- vertex, terminal, 9
- von Neumann, 80
- von Neumann ergodic theorem, 80
- wandering domain, 205
- wandering interval, 173
- weak mixing, 75, 97, 100, 152
 - flow, 75
 - transformation, 75
- weak topology, 97
- weak* topology, 85
- Weiss, 49, 135
 - even system of, 66
- Weyl theorem, 89
- Whitney embedding theorem, 111
- Wiener lemma, 98
- Williams, 63
- Wolff, 205
- word, 7
 - allowed, 8, 56
 - forbidden, 8, 56
- Yorke, 164
- zeta function, 60
 - rational, 61