Contents

Chapter 1

Combinatorial and Number Theoretic Motivation

1.1. Additive combinatorics and topological dynamics

Ramsey Theory, named after the English mathematician Frank P. Ramsey, is the study of "large" structures that are preserved under partitions. If a space with an organized structure on it is subdivided into finitely many pieces, then at least one of these pieces still has an organized structure. There is no restriction how many pieces are used, other than that this number be finite, and there is no restriction on exactly how the subdivisions are created. Behind these patterns that persist under divisions is some notion of largeness. Any 'large' set contains the organized structure, and furthermore contains 'many' of these organized structures. However, the meanings of 'large' and 'many' change with each problem. Possible spaces include groups, semigroups, vector spaces and graphs, and the structures can be any sort of pattern in the space.

The simplest such result is the *Pigeonhole Principle*: if n pigeons are placed into m pigeonholes, then so long as $n > m$, at least one pigeonhole contains 2 pigeons. Placing this in a slightly more general context of subdivisions, if a set with $rn + 1$ elements is divided into r pieces, then some piece contains at least $n + 1$ elements. Ramsey Theory expands upon this simple observation.

The original Ramsey Theorem dealt with graphs and partitions of the edges of graphs. We restrict ourselves to the narrower context of additive combinatorics. The basic objects are sets of natural numbers and the basic questions are about these sets and their structures. Rather than algebraic properties of the natural numbers, we are concerned with properties that are more quantitative and combinatorial.

Partitioning $\mathbb{N}$ into finitely many subsets amounts to giving each $n \in \mathbb{N}$ a label from a prescribed set of labels. It turns out that *some* subset contains many different patterns. For example, there are two numbers a and b in the same subset, whose sum $a + b$ also lies in the same subset. Furthermore, there exist three numbers a, b and c in the same subset such that $a + b = 2c$. (The precise statements are given in Theorems 1.2 and 1.4.) The strength and beauty of such statements lie in the fact that we have made minimal assumptions, and yet we have strong conclusions: neither the number of subsets nor the mode of dividing up the numbers is limited, and yet at least one of the subsets contains specific structures. The downside is that the conclusion is only an existence statement: the desired pattern occurs in some subset, but we have no control over which subset contains the pattern.

Many of the original proofs in Ramsey Theory were combinatorial, with the Pigeonhole Principle and vast generalizations of it used in very intricate and subtle ways. We take a different approach to the problems, first translating a combinatorial problem into a problem in dynamical systems and then using techniques in dynamical systems to solve the reformulated problem. This beautiful connection between dynamics and Ramsey Theory was developed in the 1970's by Hillel Furstenberg and the connection has been elaborated upon by many others.

A dynamical system is a way of describing how one state evolves into another. The original motivation was Newtonian mechanics, the study of planetary motion around the sun. In Newtonian mechanics, a state is the positions and momentums of the planets. One can study the changes in these positions and momentums.

Dynamical systems can be defined with different properties on the underlying space and on the map defining the evolution of the states. We restrict ourselves to topological properties, which is the field of topological dynamics. For example, each planet returns to the same point after a certain amount of time, and this periodicity is a topological property.

The typical setting we consider is a compact metric space X with a continuous transformation $T\colon X \to X$. (The reader unfamiliar with metric spaces is advised to read the Appendices before proceeding on to Chapter 2.) Repeated applications, or iterations, of the transformation T represent the evolution of the dynamical system. These iterations correspond to the positive integers: applying the transformation n times to X corresponds in a sense (to be made precise later) to the integer n. To understand properties of the natural numbers under partition, we instead study properties of the iterations of T.

In the process of developing the tools in topological dynamics, we also gain the tools to understand results of a different nature in a different area of number theory, the field of Diophantine approximation. (The name Diophantine is after the Greek mathematician Diophantus from the third century.) The basic problem is to approximate real numbers by rational numbers. More precisely we want to closely approximate a real number by a rational number whose denominator has certain properties. Many of these number theoretic results are over a hundred years old, although the dynamical proofs are much more recent.

1.2. Coloring Theorems

We start with a definition that captures the idea of subdividing a space.

Definition 1.1. A *finite partition* of a set X is a (finite) collection of sets $C_1, C_2, \ldots, C_r$ such that

- $X = C_1 \cup C_2 \cup \ldots \cup C_r$
- $C_i \cap C_j = \varnothing$ for $i, j \in \{1, 2, \ldots, r\}$ with $i \neq j$.

In many examples we consider, X is either the natural numbers $\mathbb{N} = \{1, 2, 3, \ldots\}$ or the integers $\mathbb{Z} = \{\ldots, -2, -1, 0, 1, 2, \ldots\}$. One can partition $\mathbb{N}$ into two pieces by taking the evens $\{2, 4, 6, \ldots\}$ and the odds $\{1, 3, 5, \ldots\}$, or by taking the perfect squares $\{1, 4, 9, \ldots\}$ and the complement. Or the rule determining the partition could be probabilistic, for example, tossing a coin or rolling a die to determine in which piece of a partition a number should be placed.

One of the earliest examples of a pattern that is preserved under finite partitions is due to Issai Schur:

Theorem 1.2. *If $\mathbb{N} = C_1 \cup C_2 \cup \ldots \cup C_r$ is a finite partition, then for some $j \in \{1, 2, \ldots, r\}$, there exist $a, b, c \in C_j$ with $a + b = c$.*

The triple in the conclusion of the theorem is called a *Schur triple*.

Combinatorists often use more pictorial terminology for such theorems. Assign each piece of the finite partition a different color. Then we refer to a partition of the natural numbers $\mathbb{N}$ (or of a more general semigroup) into r pieces as a *coloring* of $\mathbb{N}$ (or of the semigroup) with r *colors*, and call this more succinctly an r*-coloring*. One can thus view an r-coloring of $\mathbb{N}$ as a function $f\colon \mathbb{N} \to \{1, 2, \ldots, r\}$. Extending the language of colorings, elements in the same partition piece are said to be *monochromatic*. Thus a set A is monochromatic if there is a fixed $j \in \{1, 2, \ldots, r\}$ such that $f(n) = j$ for all $n \in A$. Schur's Theorem can now be reformulated: for each $r \in \mathbb{N}$, any r-coloring of $\mathbb{N}$ contains a monochromatic Schur triple.

Schur's original proof used techniques from number theory. The proof we give uses topological dynamics, meaning that we translate Schur's statement about patterns persisting under partitions into a statement about points in a compact metric space returning close to themselves under iteration. Then we prove the new statement using topological dynamics. This is carried out in Section **??**.

Schur's Theorem can be reformulated in a finite version. So long as a subset of the natural numbers is sufficiently large, the conclusion of Schur's Theorem already holds:

Theorem 1.3. *Given $r \in \mathbb{N}$, there exists a positive integer $N = N(r)$ such that if $\{1, 2, \ldots, N\} = C_1 \cup C_2 \cup \ldots \cup C_r$, then for some $j \in \{1, 2, \ldots, r\}$, C_j contains a, b, c with $a + b = c$.*

It is easy to see that Theorem 1.3 implies Theorem 1.2: if there is a monochromatic Schur triple in any coloring of the first N natural numbers, then certainly there is a monochromatic Schur triple in any coloring of all natural numbers. More important is that the converse also holds. Suppose not. Then for each $N \in \mathbb{N}$, there is a coloring of $\{1, 2, \ldots, N\}$ with r colors that does not contain any monochromatic Schur triple. Use any of the r colors to color the integers $\{N+1, N+2, \ldots\}$, thus extending each coloring of $\{1, 2, \ldots, N\}$ to a coloring of $\mathbb{N}$ with r colors. Thus for each $N \in \mathbb{N}$, we have a corresponding coloring of $\mathbb{N}$. Since only finitely many colors were used, by the Pigeonhole Principle there are infinitely many $N \in \mathbb{N}$ such that the corresponding coloring uses the same color for the integer 1. Consider only this subsequence of colorings and note that we are still left with an infinite set of colorings. We can now repeat this process in this subsequence. By passing to a subsequence, if necessary, we can assume that all the colorings use the same color for the integer 2. Inductively, we can pass to a subsequence of the original list of colorings that fixes each $n \in \mathbb{N}$ with a single color. Taking the limit along this subsequence, we obtain a coloring of $\mathbb{N}$ that does not contain a monochromatic Schur triple, a contradiction.

This equivalence between statements about the natural numbers and statements about sufficiently large subsets is a general phenomenon: combinatorial statements about the natural numbers can be reformulated into statements about sufficiently large finite sets.

Schur conjectured that finite colorings of the integers had a finer combinatorial structure. Namely, he conjectured that any such coloring contains monochromatic arithmetic progressions of arbitrary length. An *arithmetic progression of length* $k \in \mathbb{N}$ is defined to be a sequence of the form $a, a+b, a+2b, \ldots, a+(k-1)b$ for some $a, b \in \mathbb{N}$. We call b the *common difference* of the arithmetic progression. For example, $9, 16, 23, 30, 37$ is an arithmetic progression of length 5 with common difference 7. Schur's conjecture was proved by Bartel Leendert van der Waerden:

Theorem 1.4. *If $\mathbb{N} = C_1 \cup C_2 \cup \ldots \cup C_r$ is a finite partition, then for some $j \in \{1, 2, \ldots, r\}$, C_j contains arbitrarily long arithmetic progressions.*

As was the case for Theorem 1.2, there is an equivalent finite formulation of van der Waerden's Theorem (Exercise 8.1). Given r colors, finding the bound $N(r, k)$ such that any r-coloring of $\{1, 2, \ldots, N\}$ with $N > N(r, k)$ contains a monochromatic arithmetic progression of length k is a difficult problem that has attracted the attention of many mathematicians. Unfortunately, methods of topological dynamics do not seem to give any insight into this problem.

Techniques from topological dynamics also can be used to prove coloring theorems on spaces other than the natural numbers or the integers. One such example is the geometric Ramsey Theorem. The object being partitioned is now an infinite dimensional vector space and the pattern that can always be found is an *affine subspace*, a translate of a vector subspace. (See Appendix B.2 for the terminology.)

Theorem 1.5. *Let F be a finite field, V an infinite dimensional vector space over F and $r \in \mathbb{N}$. If $V = C_1 \cup C_2 \cup \ldots \cup C_r$ is a finite partition, then for some $j \in \{1, 2, \ldots, r\}$, C_j contains affine subspaces of arbitrarily large finite dimension.*

To state the next coloring theorem, we need some notation. Let Λ denote a finite *alphabet*, meaning that Λ is a finite collection of symbols. We denote the number of elements in Λ by $|\Lambda|$. The elements of Λ are called the *letters* of the alphabet and a *word* of length N is an ordered set of N letters. For example, if Λ is the alphabet with the 5 letters $\{1, 2, 3, 4, 5\}$, then 452213 is a word of length 6 in Λ.

Let $W_N(\Lambda)$ denote the words of length N using the alphabet Λ. Assume that Λ does not contain the letter x. Let $W_N^*(\Lambda \cup \{x\})$ denote the words of length N in the expanded alphabet $\Lambda \cup \{x\}$ in which the letter x occurs in every word at least once. Thus

$$W_N^*(\Lambda \cup \{x\}) = W_N(\Lambda \cup \{x\}) \setminus W_N(\Lambda) \ .$$

In the example with $\Lambda = \{1, 2, 3, 4, 5\}$, 2153111 is a word in $W_7(\Lambda)$ and $2x53xx1$ is a word in $W_7^*(\Lambda \cup \{x\})$.

We refer to an element of $W_N^*(\Lambda \cup \{x\})$ as a function $f(x)$ mapping Λ into $W_N(\Lambda)$. Thus $f(x) = 2x53xx1$ is a function in $W_7^*(\Lambda \cup \{x\})$.

Evaluating f over the alphabet Λ, we have:

$$\begin{aligned}
f(1) &= 2153111 \\
f(2) &= 2253221 \\
f(3) &= 2353331 \\
f(4) &= 2453441 \\
f(5) &= 2553551 \ .
\end{aligned}$$

For a function $f(x)$ with coefficients in the alphabet Λ, we refer to $\{f(\lambda)\colon \lambda \in \Lambda\}$ as a *combinatorial line*. A combinatorial line is naturally viewed as a matrix, with the words in the line becoming the rows of the matrix. Thus the previous example becomes:

$$\begin{pmatrix}
2 & 1 & 5 & 3 & 1 & 1 & 1 \\
2 & 2 & 5 & 3 & 2 & 2 & 1 \\
2 & 3 & 5 & 3 & 3 & 3 & 1 \\
2 & 4 & 5 & 3 & 4 & 4 & 1 \\
2 & 5 & 5 & 3 & 5 & 5 & 1
\end{pmatrix}$$

One of the most fundamental result of Ramsey Theory is the following theorem, due to Alfred Hales and Robert Jewett:

Theorem 1.6. *Let Λ be a finite alphabet not containing the letter x and let $r \in \mathbb{N}$. There exists a positive integer $N(r, |\Lambda|)$ such that if $N \geq N(r, |\Lambda|)$, then for any finite partition $W_N(\Lambda) = C_1 \cup C_2 \cup \ldots \cup C_r$, there exists $f(x) \in W_N^*(\Lambda \cup \{x\})$ and $j \in \{1, 2, \ldots, r\}$ such that the combinatorial line $\{f(\lambda)\colon \lambda \in \Lambda\}$ belongs entirely to C_j.*

Theorem 1.6 is stronger than both Theorem 1.4 and Theorem 1.5. To see the first statement, view the alphabet Λ as a vector space over a finite field and for the second, use the letters of Λ to write numbers in a given base. The details are left to Exercise 12.1.

Thus far, all the theorems stated show the existence of some finite configuration in a sufficiently large set. However, one can generalize Theorem 1.2 in a different way and find certain infinite configurations. We start with a definition.

Definition 1.7. A subset $A \subset \mathbb{N}$ is an IP-*set* if there exists a sequence of positive integers $p_1, p_2, \ldots$ such that

$$A = \{p_{i_1} + p_{i_2} + \ldots + p_{i_k} : k \in \mathbb{N} \ \text{ and } \ i_1 < i_2 < \ldots < i_k\} \ .$$

The integers $p_1, p_2, \ldots$ are called the *generators* of the IP-set.

Thus starting with an infinite set of generators, the set of all possible finite sums of distinct generators form an IP-set. The name IP, amusingly enough, has two different interpretations. The generators can be interpreted geometrically, as the coordinates of parallelepipeds and so the IP set itself is an infinite parallelepiped. On the other hand, IP sets correspond to idempotents in a certain abstract system that can be built out of the integers (Section 14.1). Taking the I and P from idempotent also gives IP.

For example, the set of all multiples of 10 is an IP-set, by taking the generators $\{10, 10, 10, \ldots\}$. A more interesting example is the set of all numbers that can be written with only 0's and 1's, obtained using the generators $\{1, 10, 100, \ldots\}$.

Neil Hindman showed:

Theorem 1.8. *Let $r > 0$ be an integer. If $\mathbb{N} = C_1 \cup C_2 \cup \ldots \cup C_r$ is a finite partition, then for some $j \in \{1, 2, \ldots, r\}$, C_j contains an* IP*-set.*

Hindman's Theorem has an equivalent formulation in terms of subsets of the natural numbers (see Exercise 12.3):

Theorem 1.9. *If $\mathcal{F}$ is the collection of all finite subsets of $\mathbb{N}$ and $\mathcal{F} = C_1 \cup C_2 \cup \ldots \cup C_r$ is a finite partition, then there is a sequence of disjoint sets $\alpha_1, \alpha_2, \ldots$ such that all finite unions $\alpha_{i_1} \cup \alpha_{i_2} \cup \ldots \cup \alpha_{i_k}$ lie in some C_j for some $j \in \{1, 2, \ldots, r\}$.*

These examples illustrate a general principle of Ramsey Theory: some patterns are indestructible under finite partition. These are only a few of the examples of spaces and patterns on them that are preserved under arbitrary partitions.

1.3. Diophantine Approximation

Our second application of topological dynamics belongs to number theory. A basic question is how well one can use rational numbers $\mathbb{Q}$ to approximate real numbers $\mathbb{R}$. One of the simplest examples was proved by Leopold Kronecker:

Theorem 1.10. *Assume that $\alpha \notin \mathbb{Q}$ and $\gamma \in \mathbb{R}$. Then for all $\epsilon > 0$, there exist integers n and m such that $|n\alpha - \gamma - m| < \epsilon$.*

We reformulate this theorem in a more analytic manner:

Definition 1.11. A set A is *dense* in an interval $[x, y] \subseteq \mathbb{R}$ if for all $\epsilon > 0$, any subinterval of $[x, y]$ with length ϵ contains some element $a \in A$.

Kronecker's Theorem is equivalent to saying that the sequence

$$\{n\alpha \mod 1 \colon n \in \mathbb{Z}\}$$

is dense in $[0, 1]$ (Exercise 3.9).

More generally, Kronecker showed that one can simultaneously approximate a finite set of irrationals under a simple necessary hypothesis.

Definition 1.12. The numbers $\alpha_1, \alpha_2, \ldots, \alpha_k$ are *rationally independent* if the only solution to

$$a_1\alpha_1 + a_2\alpha_2 + \ldots + a_k\alpha_k = 0$$

with $a_1, a_2, \ldots, a_k \in \mathbb{Z}$ is $a_1 = a_2 = \ldots = a_k = 0$.

Kronecker showed:

Theorem 1.13. *Assume that $\alpha_1, \alpha_2, \ldots, \alpha_k$ are rationally independent, $\gamma_1, \gamma_2, \ldots, \gamma_k \in \mathbb{R}$ and $\epsilon > 0$. Then there exist integers n and $m_1, m_2, \ldots, m_k$ such that $|n\alpha_j - \gamma_j - m_j| < \epsilon$ for $j = 1, 2, \ldots, k$.*

Herman Weyl extended this, showing that the approximations can be made with polynomial expressions:

Theorem 1.14. *Assume that $p(x)$ is a polynomial with real coefficients and that at least one coefficient other than the constant term is irrational. Then for all $\epsilon > 0$, there exist integers n and m such that $|p(n) - m| < \epsilon$.*

To formulate our last Diophantine result, we need one more definition:

Definition 1.15. The positive integers $p, q \in \mathbb{N}$ are *multiplicatively independent* if they are not both powers of a single integer.

Thus the pairs $2, 3$ and $2, 6$ are multiplicatively independent, while the pairs $3, 9$ and $4, 8$ are not.

Furstenberg showed:

Theorem 1.16. *Assume that $p, q > 1$ are multiplicatively independent integers. Then for any irrational α,*

$$\{p^m q^n \alpha \mod 1 \colon m, n \in \mathbb{N}\}$$

is dense in $[0, 1]$.

Although at first glance this seems to differ from the other Diophantine results considered, it is of a similar nature. Namely, Theorem 1.16 means that any irrational can be approximated arbitrarily well by rationals whose denominator is of the form $p^m q^n$ for some $m, n \in \mathbb{N}$.

1.4. Complexity and periodicity

The field of topological dynamics has its origins in the work of George Birkhoff. Closely related is the field of symbolic dynamics, and this plays a major role in the translation of the combinatorial problems into dynamical ones. There are numerous other combinatorial properties of systems that arise, and we discuss a few of these (see Chapter 6).

If $\mathcal{A}$ is a finite set of symbols, called the *alphabet*, we can consider bi-infinite sequences in this alphabet, called the *words* in the alphabet. Words are elements of the space $\mathcal{A}^{\mathbb{Z}}$ and we can write a word $\eta \in \mathcal{A}^{\mathbb{Z}}$ as $\eta = (\eta_n)_{\{n \in \mathbb{Z}\}}$.

A word $\eta \in \mathcal{A}^{\mathbb{Z}}$ is said *periodic* if there exists $m \in \mathbb{N}$ such that $\eta_n = \eta_{n+m}$ for all $n \in \mathbb{Z}$.

For $n \in \mathbb{N}$, we define the *complexity function* $P_\eta(n)$ to be the number of distinct patterns of length n contained in η. Marston Morse and Gustav Hedlund showed that complexity and periodicity are closely related:

Theorem 1.17. *A word $\eta \in \mathcal{A}^{\mathbb{Z}}$ is periodic if and only if there exists $n_0 \in \mathbb{N}$ such that $P_\eta(n_0) \leq n_0$.*

While this completely characterizes sequences with sub-linear growth, there is a wide area of research on sequences whose complexity $P_\eta(n) \geq n+1$ for all $n \in \mathbb{N}$. Sequences that achieve this minimum, meaning that $P_\eta(n) = n+1$ for all $n \in \mathbb{N}$, are called *Sturmian sequences* and are well understood. Sequences with higher complexity are less well understood.

One can formulate analogous notions of periodicity and complexity in higher dimensions, and such results are further considered in Section 6.

Notes

The breakthrough paper of Furstenberg [**25**] establishing the connection between dynamics and additive combinatorics proved a generalization of van der Waerden's Theorem originally proven via combinatorial means by Szemerédi [**60**]. Szemerédi's Theorem states that if $A \subset \mathbb{N}$ has positive upper density, meaning that $\limsup_{N\to\infty} \frac{|[1,\ldots,N]\cap A|}{N} > 0$, then A contains arbitrarily long arithmetic progressions. (Note that $|S|$ represents the cardinality of the set S.) Szemerédi's Theorem was conjectured by Erdős and Turán [**20**], who sought to explain the phenomenon of van der Waerden's Theorem (Theorem 1.4). In any finite partition of $\mathbb{N}$, some piece trivially has positive upper density and so van der Waerden's Theorem follows immediately from Szemerédi's Theorem. Although Furstenberg's proof of Szemerédi's Theorem is dynamical in nature, we do not cover it in this book, as it requires the added strength of ergodic theory.

The first proofs using topological dynamics of many of the combinatorial statements we study, including Theorems 1.2, 1.4, and 1.9, were given by Furstenberg and Weiss [**30**]. An excellent, more advanced treatment on these connections is [**27**]. Theorem 1.2 was first proved by Schur [**57**], and the motivation was Fermat's Last Theorem. He showed that for all $n \in \mathbb{N}$, if p is a sufficiently large prime then the equation $x^n + y^n = z^n$ has a nontrivial solution $\mod p$. Theorem 1.4 was proved by van der Waerden in [**61**]. Van der Waerden first heard this conjectured by Baudet, and so it is sometimes referred to as Baudet's conjecture, but it seems that it was originally conjectured by Schur (see [**58**]). The problem of finding bounds for $N(r,k)$ in van der Waerden's Theorem has received a great deal of attention, including work of the logician Shelah [**59**] and the Fields medalist Gowers [**31**].

A generalization of van der Waerden's Theorem by Schur and Brauer was given in [**9**]. They showed that for any finite coloring of the integers, not only can one find a monochromatic arithmetic progression of arbitrary length, but one can also guarantee that the common difference of the arithmetic progression has the same color as the progression itself. The ultimate result generalizing these theorems was proved by Rado [**48**]. The precise statement of this theorem, as well as a multidimensional version of van der Waerden's Theorem due to Gallai is in Chapter 9. Although Gallai, also known in the literature by the name Grünwald, seems never to have published the result, Rado [**49**] attributes the result to him (but states that he gives a different proof than the original one).

The geometric Ramsey Theorem (Theorem 1.5) was first proven by Graham, Leeb, and Rothschild [**32**]. A dynamical proof was given by Furstenberg and Katznelson [**28**], which includes this as a corollary of a strengthening of a result of Carlson [**14**].

Theorem 1.6 was proven by Hales and Jewett [**34**]. A density version (again requiring the added strength of ergodic theory and so not addressed in this book) was proven by Furstenberg and Katznelson [**29**]. At this time, no combinatorial proof is known for this density version.

Hindman's proof of Theorems 1.9 and 1.8 was a major breakthrough, resolving a long standing conjecture. A dynamical proof appears in [**30**].

A good resource for Diophantine approximation is Hardy and Wright [**36**], which contains a proof of Theorems 1.10 and 1.13. Theorem 1.14 was proven by Weyl [**63**] and dynamical proofs are contained in [**27**].

Weyl proved a stronger statement on the equidistribution of values of a polynomial, showing that not only are the values of the fractional part of $p(n)$ dense in $[0,1]$, but that they are evenly distributed throughout the unit interval. We omit the precise statement, as its proof relies on the added strength of techniques from ergodic theory.

Theorem 1.16 was proven in the seminal paper [**24**].

Morse and Hedlund introduced symbolic systems and the complexity functions in [**43**], giving a formulation of Theorem 1.17. In [**44**], they gave characterizations of Sturmian sequences.

Chapter 2

Dynamical Systems

2.1. Basic examples and definitions

The setting we consider is a compact metric space[1] X endowed with a continuous transformation $T\colon X \to X$. We assume that the topology on X is compatible with the metric, meaning that the open sets of X are determined by the metric.

There may also be more than one mapping on the space X. If $T\colon X \to X$ and $S\colon X \to X$ are continuous transformations mapping X to itself, then the compositions $T \circ S\colon X \to X$ and $S \circ T\colon X \to X$ are also continuous transformations mapping X to itself. Thus we may also consider a semigroup (under composition) of transformations on the compact space. If each transformation in this semigroup is also invertible, we have a group of transformations mapping X to itself. We make these notions precise.

Definition 2.1. A *dynamical system* is a pair (X, G), where X is a compact metric space and G is a semigroup of continuous transformations mapping X to itself.

Often we assume the transformations in G are invertible, in wich case G is a group of continuous transformations mapping X to itself. The important point is that since T maps X to itself, we can repeatedly apply T

[1]More generally, one could consider a compact topological space endowed with a continuous transformation, and most of the results proven here hold in this more general setting. However, to minimize the technical considerations, we assume the extra structure that comes with the metric and use exercises to indicate how many of the proofs can be carried out in the more general setting. The reader unfamiliar with the notions of a topology and compactness should study the appendices before proceeding with this chapter.

to X. The long term behavior of these repeated applications is the main focus of the study of dynamics systems.

By T^0, we mean the identity map $T^0(x) = x$. For $n \geq 1$, we define

$$T^n(x) = T(T^{n-1}(x)).$$

Notation 2.2. While this notation is standard in dynamics, it differs from other standard uses of similar notation in other fields. Throughout this book, T^n does not denote the map T raised to the n^{th} power, nor does it denote the n^{th} derivative of the map T, but is always the map T applied n times.

When the transformations acting on X are generated by a single transformation T, we use the convention of denoting the dynamical system by (X, T) instead of listing all the maps generated by T that act on X. We use the common convention of omitting parentheses and writing Tx instead of $T(x)$.

We start with some simple examples of dynamical systems; the reader should check that each transformation given is continuous.

Example 2.3 (*Identity).* If X is any space, define the *identity transformation* $\mathrm{Id}\colon X \to X$ by $\mathrm{Id}(x) = x$ for all $x \in X$. Then (X, Id) is a dynamical system. Repeated applications of Id leave each $x \in X$ unchanged.

Example 2.4 (*Circle rotation).* Define an equivalence relation $\sim$ on the real numbers $\mathbb{R}$, by setting two points to be in the same equivalence class if they differ by an integer. Thus $x \sim y$ when $x - y \in \mathbb{Z}$. We write $[x]$ for the equivalence class of $x \in \mathbb{R}$:

$$[x] = \{y \in \mathbb{R}\colon x - y \in \mathbb{Z}\}.$$

This means that the equivalence class of 0 (or of any other integer) is $\mathbb{Z}$, and the equivalence class of any $x \in \mathbb{R}$ is $\{x + z\colon z \in \mathbb{Z}\}$. Since each equivalence class has a natural representative in $[0, 1)$, we can identify the set of all equivalence classes $X = \mathbb{R}/\mathbb{Z} = \{[x]\colon x \in \mathbb{R}\}$ with the compact interval $[0, 1]$, where the points 0 and 1 are identified. One can check that this space is compact (Exercise 2.1). It is traditional to denote this presentation of the circle as $\mathbb{T} = \mathbb{R}/\mathbb{Z}$, viewing it as a 1-torus.

The metric on $\mathbb{T}$ is induced from the natural (Euclidean) metric on $\mathbb{R}$: define the distance between $x \in \mathbb{T}$ and 0 to be $\min\{x, 1-x\}$, and denote this distance by $\|x\|$. More generally, for $x, y \in \mathbb{T}$, define the distance $\|x - y\|$ between x and y to be the distance between $x - y$ and 0. Equivalently, $\|x-y\| = \min\{|x-y|, 1-|x-y|\}$. This distance (see Exercise 2.1) is called the *distance to the nearest integer*. In terms of the representation of $\mathbb{T}$ as equivalence classes, the distance between the classes $[x]$ and $[y]$ is defined to be $\min\{|x' - y'|\colon x' \in [x], y' \in [y]\}$.

Because of the identification of the endpoints 0 and 1 in the interval $[0, 1]$, it is natural to view the interval $[0, 1]$ as a circle of radius 1 in $\mathbb{R}^2$, centered around the origin: a point $x \in [0, 1]$ corresponds to the point on the circle at an angle $2\pi x$.

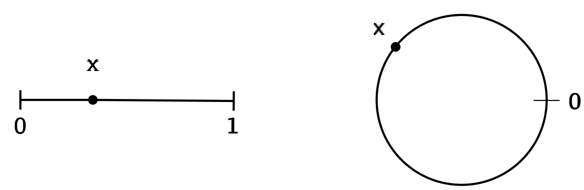

Figure 1. A point x in the interval $[0, 1]$ corresponds to a point x on the circle.

A natural group operation on $\mathbb{T}$ is addition modulo 1 and so we use additive notation. Let $\alpha \in [0, 1)$ and define $T\colon \mathbb{T} \to \mathbb{T}$ by

$$T(x) = (x + \alpha) \mod 1.$$

Because of the natural association of $\mathbb{T}$ with the circle of radius 1, the map T is called the *rotation* by α. Then $(\mathbb{T}, T)$ is a dynamical system. If $\alpha \notin \mathbb{Q}$, we call this an *irrational rotation* of the circle, and if $\alpha \in \mathbb{Q}$ we call this a *rational rotation* of the circle. Applying the transformation T n times to $\mathbb{T}$, we obtain

$$T^n(x) = (x + n\alpha) \mod 1$$

for any $n \in \mathbb{N}$. This also makes sense for negative n, as the inverse map T^{-1} is defined by

$$T^{-1}(x) = (x - \alpha) \mod 1$$

and more generally

$$T^{-n}(x) = (x - n\alpha) \mod 1.$$

This example can be generalized to a more abstract setting:

Example 2.5 (*Kronecker system*)**.** Assume that X is a compact metrizable group. (Thus X is a compact space with a group structure on it such that the group multiplication and taking of inverses are continuous; see Appendix B.1.) Let $a \in X$ be a fixed element. Define $T\colon X \to X$ by $Tx = ax$. Note that we have not assumed that X is abelian, and so unlike the particular case of the circle, we use multiplicative notation. Since we assumed that X is a group, the transformation T gives rise to a group of continuous transformations on X, given by $T^n(x) = a^n x$ for any $n \in \mathbb{Z}$. (Note that T^0 is the identity transformation.) Thus $(X, \{T^n : n \in \mathbb{Z}\})$ is a dynamical system. In analogy with the circle rotation, we call the transformation T a rotation. Since the whole system is generated by the single transformation T, we denote the system by (X, T). Such a system is called a *Kronecker system*.

Example 2.6 (*Isometry*)**.** If X is a compact metric space and $T\colon X \to X$ is an *isometry*, meaning that $d(Tx, Ty) = d(x, y)$ for all $x, y \in X$, then (X, T) is a dynamical system. (It follows immediately from the definition that T is continuous.) A circle rotation is an example of an isometry.

Given a dynamical system (X, G), we say that the transformations in G *act* on X. When we wish to emphasize that each transformation in G is continuous, we say that G *acts continuously* on X. In many applications, G is either $\mathbb{Z}$ or $\mathbb{N} \cup \{0\}$. For example, if T acts on the space X, then so do the transformations $T^0 = \mathrm{Id}, T^1 = T, T^2, T^3, \ldots$. Since the exponents of these transformations are $0, 1, 2, \ldots$ are exactly $\mathbb{N} \cup \{0\}$, it is natural to think of the semigroup $\mathbb{N} \cup \{0\}$ acting on X. This semigroup is generated by the single element 1 and the transformations acting on X are generated by the single transformation T.

If T is invertible, the group of transformations acting on X is

$$\{\ldots, T^{-2}, T^{-1}, T^0, T^1, T^2, \ldots\}.$$

Again, the transformation T^0 is the identity map and the transformation $T^1 = T$. This time, the group of transformations is in one to one correspondence with $\mathbb{Z}$, as the set of possible exponents for T are exactly the group $\mathbb{Z}$.

Dynamical systems with an invertible transformation are given a particular name:

Definition 2.7. A *homeomorphism* is a continuous, one-to-one transformation whose inverse is also continuous. A dynamical system (X, G) with each $T \in G$ a homeomorphism is called an *invertible dynamical system*.

Often we make further assumptions on the space X. For our goals it generally suffices to assume that X is metrizable, and if there is no ambiguity as to which metric is meant, the distance on X is denoted by d. If more than one space is relevant or we want to emphasize the space X, we denote the metric associated to the space X by d_X.

2.2. Distinguishing dynamical systems

In order to understand dynamical systems, we need a language to describe the systems, and we need ways in which to distinguish if two systems are the same. Primary tools for this are topological properties of a dynamical system, meaning those that are preserved under certain mappings from one system to another.

Definition 2.8. Assume that (X, T) and (Y, S) are dynamical systems. A *homomorphism* from (X, T) to (Y, S) is a continuous map $\pi \colon X \to Y$ satisfying

$$\pi(Tx) = S(\pi x) \text{ for all } x \in X.$$

We say that the homomorphism π *preserves* the dynamics. Pictorially, we can think of π as a map that makes the following diagram commute:

$$\begin{array}{ccc} X & \xrightarrow{T} & X \\ \downarrow \pi & & \downarrow \pi \\ Y & \xrightarrow{S} & Y \end{array}$$

Thus one passes from the system (X, T) to the system (Y, S) by a continuous change of coordinates.

This diagram is not symmetric, in the sense that some information about X can be lost in using π to pass from X to Y. We clarify the distinct roles that X and Y play:

Definition 2.9. A *topological semiconjugacy* is a surjective homomorphism $\pi \colon X \to Y$, and if such a map exists, we say that (X, T) and (Y, S) are *topologically semiconjugate*. If π is a homeomorphism, meaning that it is a surjective homomorphism that is also invertible and has a continuous inverse, it is called a *topological conjugacy*, and we say that (X, T) and (Y, S) are *topologically conjugate*. A topological conjugacy is also referred to as an *isomorphism*, and we can also say that the two systems are *isomorphic*.

While a topological semiconjugacy between the systems (X, T) and (Y, S) is not a symmetric relation, a topological conjugacy is. We usually omit the word topological and just write conjugacy or semiconjugacy.

Example 2.10 (Circle rotation, revisited)**.** We give another presentation of the circle rotation, isomorphic to that in Example 2.4. Consider the set

$$\mathcal{S}^1 \subset \mathbb{C} = \{z = x + iy \colon |z|^2 = x^2 + y^2 = 1\} \subset \mathbb{C}.$$

This is the complex presentation of the unit circle, where multiplication is the natural operation (rather than addition). Each point on this circle has modulus 1, and so each point can be uniquely written as $e^{2\pi i\theta}$ for some $\theta \in [0, 1)$. Assume that $T \colon \mathbb{T} \to \mathbb{T}$ is the rotation by $\alpha \in [0, 1)$ and define $S \colon \mathcal{S} \to \mathcal{S}$ by

$$S(e^{2\pi i\theta}) = e^{2\pi i(\theta+\alpha)}.$$

Then the transformations T and S are related by the homeomorphism $\rho\colon \mathbb{T} \to \mathcal{S}$ defined by $\rho(t) = e^{2\pi i t}$. If $z \in \mathbb{Z}$, then

$$\rho(t+z) = e^{2\pi i(t+z)} = e^{2\pi i t}e^{2\pi i z} = e^{2\pi i t} = \rho(t),$$

and so ρ does not depend on the equivalence class of the point t. The metric on $\mathbb{T}$ induces, under the map ρ, an equivalent metric on $\mathcal{S}$. Thus the systems $(\mathbb{T}, T)$ and $(\mathcal{S}, S)$ are conjugate and ρ is the conjugacy map.

Example 2.11 (*Squaring/Doubling map*)**.** Take $\mathcal{S}^1$ with the *squaring map* $S\colon \mathcal{S}^1 \to \mathcal{S}^1$ defined by $S(z) = z^2$. Thus S maps the point $z = e^{2\pi i\theta}$ to the point $e^{4\pi i\theta}$. This means that for the point $z = e^{i\theta}$, S doubles the angle θ and the doubling is taken modulo 2π.

The system $(\mathcal{S}^1, S)$ is conjugate to the *doubling map* $x \mapsto 2x \mod 1$ defined on $\mathbb{T}$. To define the conjugacy ρ, for $z = e^{2\pi i\theta} \in \mathcal{S}^1$, let $\rho(z) = \theta$. Then $\rho(z)$ is well-defined modulo 1. Identifying the points 0 and 1, we have an isomorphism between $\mathcal{S}^1$ and $\mathbb{T}$ that preserves the dynamics.

Not all semiconjugacies are conjugacies:

Example 2.12 (*Skew product*)**.** Let $X = \mathbb{T} \times \mathbb{T}$ and define

$$T(x,y) = (x + \alpha \mod 1, 2x + y + \alpha \mod 1),$$

where $\alpha \in [0,1)$. In general we omit the $\mod 1$ from the notation and write this as

$$T(x,y) = (x + \alpha, 2x + y + \alpha).$$

(It is clear from choice of the space that the entries must be taken $\mod 1$.) The projection map $\pi\colon X \to \mathbb{T}$ defined by $\pi(x,y) = x$ is a semiconjugacy. Thus the dynamics of the circle rotation, with rotation by α, are contained in this example. (See Exercise 2.5.)

Example 2.13 (*Logistic map*)**.** Let $X = [0,1]$ and define $Tx = kx(1-x)$, where $k \in [0,4]$. (The restriction on k is only to guarantee that T maps X to itself.) This dynamical system models population growth under some simplifying assumptions. Assume that there is a maximum population, so that if the population reaches this number then it suddenly dies out. For example, there may be a lack of some necessary resource such as space or food. Let P_n denote the fraction of the total maximum population that is alive at time n and thus $0 \le P_n \le 1$. The logistic model is

$$P_{n+1} = kP_n(1 - P_n),$$

where k is some constant determined by factors controlling the size of the population, such as space or food. Without the factor $(1 - P_n)$, the growth of the population is exponential and with this factor, the dynamical system has quadratic growth.

2.3. Orbits

Dynamics is the study of repeated applications of a transformation on the space and the behavior of points under these repeated applications. We have terminology to describe these objects.

Definition 2.14. In a system (X, T), the *forward orbit* $\mathcal{O}_T^+(x)$ of a point $x \in X$ is the set containing x and its forward images:

$$\mathcal{O}_T^+(x) = \{x, Tx, T^2x, \ldots\}.$$

If (X, T) is an invertible dynamical system, the *backwards orbit* $\mathcal{O}_T^-(x)$ of the point x is

$$\mathcal{O}_T^-(x) = \{x, T^{-1}x, T^{-2}x, \ldots\}.$$

The *full orbit* $\mathcal{O}_T(x)$ is the union of both the forward and backwards orbits:

$$\mathcal{O}_T(x) = \mathcal{O}_T^-(x) \ \cup \ \mathcal{O}_T^+(x).$$

The points in the orbit of x are called the *iterates* of x, and the powers of the transformation T are called *iterates* of the transformation.

The notion of a backwards orbit has only been defined in an invertible dynamical system. Although one may wish to study the pre-images of a point in an arbitrary system, this is a different notion than that of an orbit.

An orbit (forwards, backwards, or full) of a point may be finite or infinite. For example, under a rational rotation, the orbit of any point is finite. On the other hand, for an irrational rotation, the orbit of any point is infinite (see exercise 2.3). One can also easily construct examples where different points in the same system exhibit different types of behaviors, for example with some points having finite orbit and others infinite.

The definition of an orbit extends to the the setting of more than one transformation acting on the space:

Definition 2.15. If (X, G) is a dynamical system, the *orbit* $\mathcal{O}_G(x)$ of a point $x \in X$ is defined by

$$\mathcal{O}_G(x) = \{gx \colon g \in G\}.$$

In this context, it no longer makes sense to refer to the forwards or backwards orbit, as the group G does not necessarily have an order on it.

One of the simplest properties preserved under conjugacy is the existence of fixed points and of periodic points. For now, we restrict to dynamical systems endowed with a single transformation:

Definition 2.16. If (X, T) is a dynamical system, a point $x \in X$ is a *fixed point* if $Tx = x$. In particular, a fixed point x stays fixed under all higher iterations, meaning that $T^n x = x$ for all $n \in \mathbb{N}$.

If there exists some $n \in \mathbb{N}$ such that $x = T^n x$, then x is said to be a *periodic point* and the smallest such n is called the *period* of x. If no such n exists, then x is said to be *aperiodic*. A point $x \in X$ is said to be *preperiodic* if there exists some $m \in \mathbb{N}$ such that $T^m x$ is periodic.

If π is a conjugacy from the system (X, T) to the system (Y, S), then the image of the forward orbit of x is

$$\pi(x), \pi(Tx), \pi(T^2 x), \ldots$$

Thus a periodic point x is mapped onto a periodic point with the same period. Clearly, a fixed point is also mapped under a conjugacy (and more generally under a semiconjugacy) onto a fixed point.

A periodic point is a fixed point for some iterate T^n of T. However, the two sets are not the same, since the fixed points of T^n also include periodic points whose period is any divisor of n.

Example 2.17. All points are periodic (with the same period) for a rational rotation of the circle, but an irrational rotation of the circle has no periodic points (Exercise 2.3).

Example 2.18. The only fixed point of the doubling map is 0. Furthermore,

$$2^n x \equiv x \mod 1$$

has $2^n - 1$ distinct solutions $x = k/(2^n - 1)$, for $k = 0, 1, \ldots, 2^n - 2$. When n is prime, $k/(2^n - 1)$ for $k = 1, 2, \ldots, 2^n - 1$ is periodic with period n. This means that for prime n, the doubling map has $(2^n - 2)/n$ distinct periodic orbits of period n. By conjugacy, the same holds for the squaring map on the circle $\mathcal{S}^1$.

At times, we are interested not just in the behavior of individual points under a transformation but in the behavior of sets:

Definition 2.19. Let (X, T) be a dynamical system. If $A \subset X$, the *image* $T(A)$ of A under the map T is defined by:

$$T(A) = \{Tx \colon x \in A\}.$$

Similarly, the *pre-image* $T^{-1}(A)$ of a set A is defined by:

$$T^{-1}(A) = \{x \in X \colon Tx \in A\}.$$

Note that $T^{-1}(T(A))$ contains A, but if T is not invertible then these sets are not necessarily equal. One can easily construct a finite system such that these sets differ. A more interesting situation is described in the next section.

2.4. Shift systems

Our next example recurs throughout the book and is the main tool for translating combinatorial statements into dynamical ones. The idea is to use sequences of symbols from a certain alphabet to encode information about a dynamical system, the integers, or other dynamical objects.

Example 2.20 (*Two-sided shift on k symbols*)**.** Fix some $k \in \mathbb{N}$. Let $\Lambda = \{0, 1, \dots, k-1\}$, endowed with the discrete topology. Let

$$\Omega = \prod_{n=-\infty}^{\infty} \Lambda$$

be endowed with the product topology. We often write this using the shorthand $\Omega = \Lambda^{\mathbb{Z}}$ and call Λ the *alphabet*. (To simplify notation, we have taken all Λ to be equal; one could also define Λ_n for each $n \in \mathbb{Z}$, without an assumption that all Λ_n are identical, and take $\Omega = \prod_{n=-\infty}^{\infty} \Lambda_n$. This leads to more complicated dynamical systems, such as that of Exercise 2.15.) Then Ω is a compact, totally disconnected, and Hausdorff space; it can be made into a compact metric space by taking the metric

$$d(x, y) = \begin{cases} 0 & \text{if } x = y \\ 2^{-m} & \text{if } x \neq y \text{ and } m = \min\{|n| \colon x_n \neq y_n\} \end{cases}$$

(Exercise 2.6). Thus two points x and y are close if they agree for a large number of entries centered around 0. In particular, $d(x, y) = 1$ if $x_0 \neq y_0$ and $d(x, y) < 1$ otherwise. One can easily check that this metric generates the product topology on Ω. We write points $x \in X$ as $x = \{x_n\}_{n \in \mathbb{Z}}$. We define the *shift map* $T \colon \Omega \to \Omega$ by $Tx = y$, where $y_n = x_{n+1}$ for all $n \in \mathbb{Z}$.

$$\dots 001101011000101\overset{\downarrow}{0}001000001001011010 \dots$$

$$\dots 011010110001010\overset{\downarrow}{0}010000010010110101 \dots$$

Figure 2. The top is a bi-infinite sequence in the alphabet $\{0, 1\}$ with 0 marked by an arrow, and below is a shifted copy of the sequence.

The *cylinder sets* are defined to be sets $C \subset \Omega$ of the form

$$C = \{x \in \Omega \colon (x_{n_1}, x_{n_2}, \dots, x_{n_j}) \in A\} \tag{2.1}$$

with $A \subset \{0, 1, \dots, k-1\}^j$ a fixed set. These cylinder sets form a basis for the topology of Ω. Thus in order to check the continuity of T, it suffices to check it on the cylinder sets. Since for the cylinder set defined in (2.1),

$$T^{-1}(C) = \{x \in \Omega \colon (x_{n_1+1}, x_{n_2+1}, \dots, x_{n_j+1}) \in A\}$$

is itself a cylinder set, T is continuous. Since $T^{-1}(y) = x$, where $x_n = y_{n-1}$, the transformation T is invertible and the inverse is continuous, using the same argument that shows that T is continuous. The transformation T is called the *shift* on the k symbol alphabet Λ and the system (Ω, T) is called the *two-sided shift on k symbols*, or more simply is called the *k-shift*.

For example, if we take an alphabet with the two symbols $\{0, 1\}$, then Ω consists of all bi-infinite sequences of 0's and 1's. If x is one of the constant sequences, meaning that $x = \ldots 000 \ldots$ or $x = \ldots 111 \ldots$, then $Tx = x$. One can easily construct points with a given period n, by simply taking a pattern of length n that cannot be decomposed into disjoint copies of a smaller pattern, and copying this pattern in both directions. For a one-sided shift, one can use a similar idea to construct preperiodic points, where the preperiodic portion can have any length desired.

Often we build new dynamical systems out of existing ones. One of the simplest ways to do this is by finding the new system within the given system:

Definition 2.21. If (X, T) is a dynamical system and $Y \subseteq X$ is a closed set with $T(Y) \subseteq Y$, then the dynamical system (Y, T) is called a *subsystem* of the system (X, T). More generally, if $X \subseteq \Omega$ is a closed subset with $TX \subseteq X$, we call (X, T) a *symbolic dynamical system* or more succinctly, a *symbolic system*.

In order for these definitions to make sense, we must show that (Y, T) is in fact a dynamical system. This is left to Exercises 2.7 and 2.8. Symbolic systems are further discussed in Section 3.2.

By using the indexing set $\mathbb{N}$ instead of $\mathbb{Z}$ in Example 2.20, we obtain the *one-sided shift on k symbols*. In this case, when an element x is shifted, we lose the information that was contained in the first entry. For example, if (X, T) is the one-sided shift on the alphabet $\{0, 1\}$, then if A consists of the sequence of all 0's, $T(A) = A$, but $T^{-1}(A)$ contains the sequence of all 0's and the sequence with one 1 followed by all 0's. Using this idea, one can construct infinite sets A such that $T^{-1}(T(A))$ contains A, but the two sets are not equal.

Notes

Kronecker [**42**] proved results on Diophantine approximation (see Exercises 3.8 and 3.9), and these results can be proven using the dynamical system in Example 2.5. Furstenberg [**25**] dubbed this dynamical system a Kronecker system and it, and generalizations of it, play an important role in the proof of Szemerédi's Theorem.

The example of a symbolic dynamical system given in 2.20 has its origins in work of Hadamard [**33**] and the papers of Morse and Hedlund ([**43**] and [**44**]).

Exercises

Exercise 2.1. Show that the space $\mathbb{T}$ defined in Example 2.4 is compact, that the metric d defined is a metric, and that the transformation T defined is continuous.

Exercise 2.2. If p/q is rational, give all the periodic points and their periods under the rotation by p/q on the circle.

Exercise 2.3. Show that an irrational circle rotation has no periodic points. Show that every point has a dense orbit.

Exercise 2.4. If (X, T) is a dynamical system, show that $x \in X$ is periodic if and only if $\mathcal{O}_T^+(x)$ is compact.

Exercise 2.5. Compute the orbit of $(0, 0)$ in Example 2.12.

Exercise 2.6. Show that $\Omega = \{0, 1, \ldots, k-1\}^{\mathbb{Z}}$ is a metric space where

$$d(x, y) = \frac{1}{1 + \min\{|m| \colon x_m \neq y_m\}}$$

and that this metric generates the product topology. Show that it is a compact, totally disconnected, Hausdorff space.

Exercise 2.7. Let (Ω, T) denote the shift space. Assume that X is a closed subset of Ω with $TX \subseteq X$. Show that (X, T) is a dynamical system. This is generalized in Exercise 2.8.

Exercise 2.8. Let (X, T) be a dynamical system and assume that Y is a closed subset of X with $TY \subset Y$. Show that (Y, T) is a dynamical system.

Exercise 2.9. Let (X, T) be the 2-shift. Find all sequences in X with periods $2, 3, 4$ and 5. Determine which sequences lie in the same orbit.

Exercise 2.10. Show that the orbit of a point x in the shift space (Ω, T) on k symbols is dense if and only if every finite block of k symbols appears somewhere in x.

Exercise 2.11. Show that for the 2-shift, Ω is homeomorphic to the Cantor middle thirds set.

Exercise 2.12. Show that the preimage of a fixed (respectively, periodic) point under a semi-conjugacy need not be a fixed (respectively, periodic) point. More generally, show that every preimage under a semi-conjugacy of a fixed (respectively, periodic) point might not be a fixed (respectively, periodic) point.

Exercise 2.13. Show that if X is a compact metric space with metric d and $T\colon X \to X$ is an isometry, then T is invertible and $d(T^{-1}x, T^{-1}y) = d(x,y)$ for all $x, y \in X$.

Exercise 2.14. Show that the periodic points are dense in $[0,1]$ for the *tent map* $T\colon [0,1] \to [0,1]$ defined by

$$T(x) = \begin{cases} 2x & \text{if } 0 \le x < 1/2 \\ 2 - 2x & \text{if } 1/2 \le x \le 1. \end{cases}$$

Show that the nonperiodic points are also dense.

Exercise 2.15. Let $X_n = \{0, 1, \ldots, k_n\}$ and let $X = \prod_{n=1}^{\infty} X_n$. Define $T\colon X \to X$ by $Tx = y$ where

$$y = \begin{cases} (x_1 + 1, x_2, x_3, \ldots) & \text{if } x_1 < k_1 \\ (0, 0, \ldots, 0, x_m + 1, x_{m+1}, x_{m+2}, \ldots) & \text{if } x_1 = k_1, \ldots, \\ & \qquad x_{m-1} = k_{m-1}, x_m < k_m \\ (0, 0, \ldots, 0) & \text{if } x_1 = k_1, x_2 = k_2, \ldots \end{cases}$$

Show that (X, T) is a dynamical system. This is called the *adding machine.*

Exercise 2.16. Prove that a continuous, surjective transformation $T\colon [0,1] \to [0,1]$ has at least one fixed point. Show that T^2 has at least 2 fixed points. What about T^n for $n \in \mathbb{N}$?

Exercise 2.17. Characterized the periodic points for the map on $\mathbb{T}^2$ given by

$$T(x,y) = (x, y + x) \mod 1.$$

Chapter 3

Recurrence

3.1. Birkhoff recurrence

The idea of recurrence is central throughout dynamics and is the major tool that allows one to use dynamics to derive results in additive combinatorics. Under some natural assumptions on the system or perhaps on a subset in the system, the orbit of a point must return arbitrarily close to the original point. The most basic such result is the *Birkhoff Recurrence Theorem*, due to George Birkhoff:

Theorem 3.1. *Let (X, T) be a dynamical system. There exists $x \in X$ such that any neighborhood U of x contains some iterate $T^n x$ with $n \in \mathbb{N}$.*

Proof. Consider the family $\mathcal{F}$ of all nonempty closed subsets Y of X satisfying $TY \subseteq Y$. This family is nonempty, since it contains X. We can put a partially ordering $\prec$ on $\mathcal{F}$ by inclusion: we say that $Y_1 \prec Y_2$ if $Y_1 \subset Y_2$.

Let $\mathcal{G}$ be a totally ordered collection in the family $\mathcal{F}$. The intersection of any finite number of elements in $\mathcal{G}$ is nonempty (it is exactly the least element in this finite subset) and is again in $\mathcal{F}$. By the finite intersection property (Theorem E.14), the intersection of all elements in $\mathcal{G}$ is nonempty. Call this intersection G_0. Since G_0 is contained in each element of $\mathcal{G}$, we have that G_0 is the lower bound for $\mathcal{G}$. Therefore, by Zorn's Lemma (Axiom A.2), the family $\mathcal{F}$ has a minimal (again, with respect to the ordering $\prec$) element Z. For any $x \in Z$, let

$$Z_0 = \overline{\{T^n x \colon n \geq 1\}}.$$

Then Z_0 is a closed set and $TZ_0 \subseteq Z_0$. Since Z is also closed and satisfies $TZ \subseteq Z$, we have that $Z_0 \subseteq Z$. However, Z was chosen to be a minimal element and so we have that $Z_0 = Z$. In particular, $x \in Z_0$. Therefore any neighborhood of x contains some iterate $T^n x$ for some $n \geq 1$. $\square$

The Birkhoff Recurrence Theorem motivates several definitions. An important role in the proof is played by sets which get mapped into themselves under the transformation:

Definition 3.2. If (X, T) is a dynamical system, a set $Y \subseteq X$ is said to be *invariant* under T if $TY \subseteq Y$. We also call such a set T*-invariant.*

It follows immediately from the definition that if $Y \subseteq X$ is T-invariant, then $T^n Y \subseteq Y$ for all $n \in \mathbb{N}$.

For example, under the identity transformation Id on a space X, any subset $Y \subset X$ is T-invariant. If $X = [0, 1]$ with $Tx = x^2$, then the set $[0, 1/2]$ is T-invariant, while the set $[0, 1/4] \cup [1/2, 3/4]$ is not.

The type of point found in the conclusion of the theorem plays a major role in all our results:

Definition 3.3. If (X, T) is a dynamical system, a point $x \in X$ is said to be *recurrent* if for any neighborhood U of x, there exists an integer $n \geq 1$ such that $T^n x \in U$.

If X is a metric space and $x \in X$ is a recurrent point, then there is not just one iterate returning to a given neighborhood but infinitely many iterates return. That is, for all $\varepsilon > 0$, there exist infinitely many $n \in \mathbb{N}$ such that $d(T^n x, x) < \varepsilon$. If not, then we could always find some $\varepsilon' < \varepsilon$ such that the ε'-neighborhood of x contained no forward iterate of x.

Moreover, if we take a sequence of neighborhoods around a recurrent point x whose diameters shrink to 0, we can find an iterate of x lying in each neighborhood. This idea leads to an alternate characterization of a recurrent point: the point $x \in X$ is recurrence if there exists a sequence of integers $n_k \to \infty$ such that $T^{n_k} x \to x$.

It follows immediately from the definition that any fixed or periodic point is recurrent. However, more interesting behavior occurs when the iterates come close to, but do not equal, the original point. Every point in a Kronecker system exhibits such behavior:

Example 3.4. Let (X, T) be a Kronecker system and assume that $Tx = ax$ for some fixed element $a \in X$. By the Birkhoff Recurrence Theorem, we know that some point $x_0 \in X$ is recurrent. Since X is a group, given any $x \in X$, we can write $x = x_0 y$ for some $y \in X$. If U is any neighborhood of x, then Uy^{-1} is a neighborhood of x_0. Since x_0 is recurrent, there exists $n \in \mathbb{N}$ such that $a^n x_0 \in Uy^{-1}$. But this implies that $a^n(x_0 y) \in U$, which is equivalent to saying that $a^n x \in U$. Thus every point is recurrent.

The proof of the Birkhoff Recurrence Theorem (Theorem 3.1) gives a bit more than was stated. We showed that any dynamical system (X, T) contains a nonempty, closed T-invariant subset, which is the orbit closure of a recurrent point in X. This subset may be all of X, such as in the case of a minimal Kronecker system, may be a single point, such as in a system with a unique fixed point, or can be something in between.

More generally, the same proof can be extended to a dynamical system (X, G), where G is a semigroup of transformations acting continuously on X. We extend the definitions to this setting:

Definition 3.5. If (X, G) is a dynamical system, a subset $Y \subset X$ is *G-invariant* if $gY \subset Y$ for all $g \in G$.

A point $x \in X$ is *recurrent* if there exists a sequence $g_k \in G$ such that $g_k x \to x$.

The corresponding extension of the Birkhoff Recurrence Theorem is left to Exercise 3.6.

3.2. Symbolic systems and recurrence

We recall the definition of a symbolic system from Section 2.4. Let Λ be an alphabet and consider the one-sided shift $(\Lambda^{\mathbb{N}}, T)$. Assume that $X \subseteq \Lambda^{\mathbb{N}}$ is a closed T-invariant subspace. Recall (see Example 2.20) that the metric on (X, T) is defined by setting the distance between points x and y to be 0 if the points are equal, and otherwise the distance is 2^{-m}, where $m = \min\{|n| \colon x_n \neq y_n\}$.

A finite sequence of elements of Λ is called a *word* in the alphabet Λ and the *length* $|\Lambda|$ of this word is the number of letters in the sequence. If w and w' are words in Λ, then the *concatenation* ww' is also a word in Λ. An element of $\Lambda^{\mathbb{N}}$ can be viewed as a concatenation of infinitely many words. (Infinitely many words are needed, since by definition a word has finite length.)

A word $w = w_1 w_2 \dots w_n$ *occurs in* another word $w' = w'_1 w'_2 \dots w'_m$ if for some $j \in \{1, 2, \dots, m-n\}$, $w_1 = w'_{j+1}, w_2 = w'_{j+2}, \dots, w_n = w'_{j+n}$ and we say that w is a *subword* of w'. In this case, we can write $w' = uwv$ for some, possibly empty, words u and v.

If $x \in \Lambda^{\mathbb{N}}$ is periodic, then there exists some word w such that $x = www\dots$. We also write this as $x = \overline{w}$, where the bar indicates that w is repeated infinitely often. Similarly we can write a preperiodic point as $x = w'\overline{w}$, where w and w' are words in Λ. Thus in the one-sided shift on the symbols $\{0, 1\}$, the fixed point consisting of all 0's is written as $\overline{0}$ and the point alternating 0's and 1's is written as $\overline{01}$.

It follows immediately from the definitions that $x \in \Lambda^{\mathbb{N}}$ is recurrent if and only if every word in x occurs a second time. In particular, this means that every word in x occurs infinitely often. Furthermore, it suffices to consider words that start at the beginning of x, since all words occur in these words. Thus if a is the first symbol of a recurrent point $x \in \Lambda^{\mathbb{N}}$, then there is some word w_1 such that the initial portion of x is exactly aw_1a. Since the initial word aw_1a also must recur, we can find a word w_2 such that the initial portion of x is $aw_1aw_2aw_1a$. Continuing this process leads to:

Lemma 3.6. *If Λ is a finite alphabet, then $x \in \Lambda^{\mathbb{N}}$ is recurrent under the shift T if and only if*

$$x = aw_1aw_2aw_1aw_3aw_1aw_2aw_1aw_4aw_1aw_2aw_1aw_3aw_1aw_2aw_1a\ldots$$

for some symbol $a \in \Lambda$ and words $w_1, w_2, \ldots$ in the alphabet Λ.

It is easy to construct invariant sets in the shift system $(\Lambda^{\mathbb{N}}, T)$. Starting with any point $x \in \Lambda^{\mathbb{N}}$, let X be its orbit closure:

$$X = \overline{\{T^n x \colon n \in \mathbb{N}\}}.$$

Then X is a closed, T-invariant subset of $\Lambda^{\mathbb{N}}$. Depending on the choice of x, the subset X may be finite, infinite but not all of $\Lambda^{\mathbb{N}}$, or equal to $\Lambda^{\mathbb{N}}$ (Exercise 3.3).

3.3. Hilbert's Theorem: an early coloring theorem

In the late 1800's, Hilbert proved what is generally considered the first result in Ramsey Theory. We now have the tools to prove this result, first translating the problem on patterns in the natural numbers into a statement in topological dynamics and then using the Birkhoff Recurrence Theorem to prove the dynamical statement. This result and its proof are prototypes for the more complicated coloring theorems that follow.

For a subset $A = \{n_1, n_2, \ldots, n_k\}$ of integers, let $\mathcal{P}(A)$ denote the set of all finite (possibly empty) sums $n_{i_1} + n_{i_2} + \ldots + n_{i_m}$ with $i_1 < i_2 < \ldots < i_m$. Thus $\mathcal{P}(A)$ is the set of all finite sums of distinct elements of A. Hilbert showed:

Theorem 3.7. *If $\mathbb{N} = C_1 \cup C_2 \cup \ldots \cup C_r$ is a finite partition, then for all $\ell \in \mathbb{N}$, there exists a set $A = \{n_1, n_2, \ldots, n_\ell\}$ of natural numbers such that infinitely many translates of $\mathcal{P}(A)$ lie in C_j for some $j \in \{1, 2, \ldots, r\}$.*

Proof. We first define the particular dynamical system used for proving the coloring statement. Let $\Lambda = \{1, 2, \ldots, r\}$. Define

$$y_n = i \text{ if and only if } n \in C_i \text{ for } i = 1, 2, \ldots, r.$$

Thus $y = (y_n : n \in \mathbb{N}) \in \Lambda^{\mathbb{N}}$. We consider the system $(\overline{\mathcal{O}_T^+(y)}, T)$ generated by y under the shift T. By Theorem 3.1, this system contains a recurrent point x.

Let a denote the first coordinate x_1 of x. We show that C_a, the element of the partition containing all sequences whose first entry is a, contains the desired pattern. Since x is recurrent, by Lemma 3.6 x has the form

$$aw_1aw_2aw_1aw_3aw_1aw_2aw_1a\ldots$$

for some symbol $a \in \Lambda$ and words $w_1, w_2, \ldots$ in the alphabet Λ. Define

$$\begin{aligned} W_0 &= a \\ W_1 &= W_0w_1W_0 \\ W_2 &= W_1w_2W_1 \\ &\vdots \\ W_n &= W_{n-1}w_nW_{n-1}. \end{aligned}$$

By definition, W_n includes the sequence W_{n-1} as its initial portion and the point x is a limit of the finite sequences W_n. If some symbol occurs at the p-th entry of W_n, then the same symbol occurs both at p and $p + m_n$ in W_{n+1}, where $m_n = |W_nw_{n+1}|$. In particular, the symbol a occurs at each of the following entries:

$$1, 1 + m_1, 1 + m_2, 1 + m_1 + m_2, \ldots, 1 + m_1 + m_2 + \ldots + m_\ell.$$

This means that a occurs at each of the 2^ℓ sums in $\mathcal{P}(m_1, m_2, \ldots, m_\ell)$. By recurrence, again, every word occurs infinitely often in x and so C_a contains infinitely many translates of $\mathcal{P}(m_1, m_2, \ldots, m_\ell)$.

We now apply this in the system $(\overline{\mathcal{O}_T^+(y)}, T)$. This system contains a recurrent point x which has the above form. If $x = T^n y$ for some $n \in \mathbb{N}$, then the first symbol in x occurs in y on infinitely many translates of $\mathcal{P}(m_1, m_2, \ldots, m_\ell)$ and we are done. Otherwise, x is a limit point of $\overline{\mathcal{O}_T^+(y)}$ and so there exists a sequence $n_k \to \infty$ such that $T^{n_k} y \to x$ as $n_k \to \infty$. If y is periodic, then some translate of y is exactly x and so the symbol a appears in y on a translate of an arithmetic progression and again we are done. So assume that y is not periodic. If a is the first symbol in x, then a recurs at $1 + p$ for all $p \in \mathcal{P}(m_1, m_2, \ldots, m_\ell)$, as above. We can choose n_k sufficiently large such that $T^{n_k} y$ and x agree for $1 + m_1 + m_2 + \ldots + m_\ell$ entries. Then $y_{n_k+p} = a$ for all $p \in \mathcal{P}(m_1, m_2, \ldots, m_\ell)$. Taking $n_k \to \infty$, we have $1 + n_k + \mathcal{P}(m_1, m_2, \ldots, m_\ell) \subset C_a$ for an infinite sequence n_k. $\square$

Hilbert's Theorem is one of the earliest examples of what can be called a Ramsey type theorem, finding a monochromatic pattern that must occur in any finite coloring of the integers. The monochromatic subsets found in this theorem are parallelepipeds. More precisely, a subset A of a group is a

k-dimensional parallelepiped if there exist positive integers $p_1, p_2, \ldots, p_{k-1}$ such that $A = A_1 \cup A_2 \cup \ldots \cup A_k$, where $A_1 = \{a\}$, $A_2 = A_1 + p_1$, $A_3 = (A_1 \cup A_2) + p_2, \ldots, A_k = (A_1 \cup A_2 \cup \ldots \cup A_{k-1}) + p_{k-1}$. In this language, Hilbert's Theorem states that given a finite coloring of the natural numbers, there exists a k-dimensional parallelepiped such that infinitely many translates of it have the same color.

3.4. Limit sets

Recurrence and other dynamical properties are determined by the behavior of high iterates of points. One way to capture this long term behavior is by describing the set of points that are close to the orbit of a point. We begin with some definitions to make this notion precise.

Definition 3.8. If (X, T) is a metrizable dynamical system and $x \in X$ is a point, a point $y \in X$ is an *ω-limit point* of x if there exists a sequence of integers $n_k \to \infty$ such that $T^{n_k}x \to y$. The *ω-limit set* of x is the set of all ω-limit points of x and is denoted by $\omega(x)$.

In particular, y is an ω-limit point of x if $y \in \overline{\mathcal{O}_T^+(x)}$. We can define recurrence in this terminology: a point x is recurrent if and only if $x \in \omega(x)$. Symbolically, the ω-limit set is defined by

$$\omega(x) = \bigcap_{n=1}^{\infty} \overline{\bigcup_{j>n} T^j x}.$$

If (X, T) is the one-sided shift on the symbols $\{0, 1\}$ and $x = \overline{01}$, then $\omega(x) = \{\overline{01}, \overline{10}\}$. If $x = 010011000111\ldots$, then $\omega(x)$ consists of all sequences with any finite number (possibly 0) of 0's followed by infinitely many 1's and all sequences with any finite number of 1's (possibly 0) followed by infinitely many 0's. Under the squaring map $x \mapsto x^2$ on $[0, 1]$, we have that $\omega(x) = 0$ for any $0 \le x < 1$ and $\omega(1) = 1$.

In any dynamical system, the ω-limit set of a point is always nonempty, compact and T-invariant (Exercise 3.19). It follows that if X has no nonempty proper closed invariant subsets, then $\omega(x) = X$ for all $x \in X$.

If T is not invertible, then $T^{-1}\big(\omega(x)\big)$ can be larger than $\omega(x)$. For example, for the one-sided shift on k symbols, if we take $x = \overline{0}$, then $\omega(x) = x$, but $T^{-1}x$ contains k points.

In a dynamical system (X, T) with an invertible transformation T, it also makes sense to consider the long term behavior of orbits under backwards iteration:

Definition 3.9. If (X, T) is a dynamical system with invertible T and $x \in X$, the *α-limit set* $\alpha(x)$ of the point x is defined by

$$\alpha(x) = \bigcap_{n=1}^{\infty} \overline{\bigcup_{j>n} T^{-j}x}.$$

As for the ω-limit set, $\alpha(x)$ is nonempty, compact, and T-invariant.

The behavior of neighborhoods of a point can also give insight into the long term dynamics. We define:

Definition 3.10. In a dynamical system (X, T), a point $x \in X$ is *wandering* if there exists some neighborhood U of x such that $T^{-1}U, T^{-2}U, \ldots$ are all disjoint from U. Otherwise, the point x is said to be *nonwandering.* The set of nonwandering points in (X, T) is called the *nonwandering set* and is denoted by $\Omega = \Omega(T)$.

Thus for a wandering point x and for some neighborhood U of x, we have that for distinct integers $n, m \geq 1$, $T^{-n}U \cap T^{-m}U = \varnothing$. For a nonwandering point x and any neighborhood U of x, there exists $n \in \mathbb{N}$ such that $T^{-n}(U) \cap U \neq \varnothing$. Clearly, any recurrent point is nonwandering.

An equivalent formulation of the nonwandering set of T is:

$$\Omega(T) = \{x \in X \colon \text{for all neighborhoods } U \text{ of } x, \text{ there exists } n \geq 1 \text{ such that } T^{-n}U \cap U \neq \varnothing\}.$$

If T is a homeomorphism, then $T^{-n}U \cap U = T^{-n}(U \cap T^{n}U)$ and so $\Omega(T^{-1}) = \Omega(T)$ and we can define

$$\Omega(T) = \{x \in X \colon \text{for all neighborhoods } U \text{ of } x, \text{ there exists } n \neq 0 \text{ such that } T^{-n}U \cap U \neq \varnothing\}.$$

We can show a bit more:

Theorem 3.11. *If (X, T) is a metrizable dynamical system, then*

$$\Omega(T) = \{x \in X \colon \textit{for all neighborhoods } U \textit{ of } x \textit{ and all integers } N \geq 1, \textit{there exists } n \geq N \textit{ such that } T^{-n}U \cap U \neq \varnothing\}.$$

Proof. The right hand side is clearly a subset of $\Omega(T)$ and so it suffices to prove the opposite inclusion. Let $x \in \Omega(T)$, U be a neighborhood of x and let $N \geq 1$ be an integer. If x is periodic, then $T^{-n}U \cap U \neq \varnothing$ for infinitely many n and so certainly for some $n \geq N$. So assume that x is not periodic.

Fix an integer $N \geq 1$. Choose $\varepsilon > 0$ such that the ball $B(x; \varepsilon)$ around x of radius ε satisfies $B(x; \varepsilon) \subset U$. We show that there exists $0 < \delta < \varepsilon$ such that $B(x; \delta) \cap T^{-j}B(x; \delta) = \varnothing$ for all j with $1 \leq j \leq N - 1$. If no such δ exists, then for every $n \in \mathbb{N}$ with $1/n < \varepsilon$, there exists $x_n \in B(x; 1/n) \cap T^{-j_n} B(x; 1/n)$ for some $1 \leq j_n \leq N-1$. Since N is finite, by the Pigeonhole Principle we can choose a subsequence of integers $\{n_i\}$ such that $j_{n_i} = k$ for all i. Then $x_{n_i} \in B(x; 1/n_i)$ and so $x_{n_i} \to x$ as $i \to \infty$ and $T^k x_{n_i} \in B(x; 1/n_i)$. Therefore $T^k x_{n_i} \to x$. But $x_{n_i} \to x$ also implies that $T^k x_{n_i} \to T^k x$ by the continuity of T^k. Since $T^k x_{n_i}$ approaches both $T^k x$ and x, we must have that $T^k x = x$, a contradiction of the assumption that x is not periodic. This proves the existence of such δ. Thus for $x \in \Omega(T)$, there exists $n \geq N$ with $B(x; \delta) \cap T^{-n} B(x; \delta) \neq \varnothing$. □

We can also define the notion of nonwandering for an open subset in a system (X, T).

Definition 3.12. Let (X, T) be a dynamical system. An open set $A \subset X$ is *wandering* if $A, T^{-1}A, T^{-2}A, \ldots$ are all disjoint and otherwise it is *nonwandering*. The system (X, T) is said to be *nonwandering* if no nonempty open set is wandering.

Under the assumption of nonwandering, most points in a compact metric space must be recurrent:

Theorem 3.13. *Assume that X is a compact metric space. If the dynamical system (X, T) is nonwandering, then the set of recurrent points of (X, T) is residual.*

(For the definition of a residual set, see Appendix **??**.)

Proof. Define

$$f(x) = \inf_{n \geq 1} d(x, T^n x).$$

Then for any $\varepsilon > 0$ and $x_0 \in X$, if $f(x_0) = y$, then $d(x_0, T^n x_0) < y + \varepsilon$ for some $n \in \mathbb{N}$. By continuity, we have that $d(x, T^n x) < y + \varepsilon$ for all x in some neighborhood of x_0. Thus for all x in this neighborhood of x_0, we have that $f(x) < f(x_0) + \varepsilon$ and so f is upper semicontinuous.

By Theorem F.20, an upper semicontinuous function on a complete metric space has a residual set of points of continuity. Let x_0 be a point of continuity of $f(x)$. If $f(x_0) = 0$, then x_0 is a recurrent point. If $f(x_0) > 0$, then there exists a neighborhood U of x_0 and $\varepsilon > 0$ such that $f(x) > \varepsilon$ for all $x \in U$. Without loss, we can assume that the diameter of U is less than ε. Then U is not a wandering set and so for some $n \in \mathbb{N}$, $T^{-n}U \cap U \neq \varnothing$. This means that for some $x \in U$, we have $T^n x \in U$ and so $f(x) < \varepsilon$, a contradiction. Therefore, at the residual set of points of continuity, $f(x) = 0$ and so each of these points is recurrent. □

Notes

The first recurrence result was stated and proved by Poincaré [**46**], in the context of a measure preserving transformation on a finite measure space. Poincaré recurrence is the base case for the complicated induction in Furstenberg's proof [**25**] of Szemerédi's Theorem. Birkhoff proved recurrence in the setting of a compact topological group in [**5**]; he actually proved a stronger statement showing that one has not only recurrence, but also (in modern terminology) almost periodic points and minimal sets. The definitions and statements are left until Chapter 4.

Although the proof of Theorem 3.1 given uses Zorn's Lemma, one can give a constructive proof and we do so in Section 4.2.

Hilbert proved Theorem 3.7, as part of his "irreducibility" theorem on rational functions in [**37**]. He does not seem to have studied the combinatorial implications of the result. The proof given here follows that of Furstenberg [**27**]. Hilbert's Theorem can be strengthened in several ways, and Schur [**57**] proved the first such result, showing that any finite coloring of the integers contains a monochromatic cube (a 2-dimensional parallelepiped). We prove Schur's Theorem in Section 12.1. This was extended by Folkman [**21**] to k-dimensional parallelepipeds. In Section 12 we prove Hindman's far reaching generalization, showing that in any finite coloring of the integers, there is always a monochromatic infinite dimensional parallelepiped.

A proof of Theorem 3.13 is given in Furstenberg [**27**].

Exercises

For all the exercises, assume that (X, T) is a dynamical system.

Exercise 3.1. Show that Birkhoff's Theorem fails if you remove either the assumption that X is compact or the assumption that T is continuous.

Exercise 3.2. If (X, T) is a dynamical system and $Y \subset X$ is T-invariant, show that $\overline{Y}$ is also T-invariant.

Exercise 3.3. In the shift system $(\Lambda^{\mathbb{N}}, T)$, find x whose orbit closure is finite, x whose orbit closure is infinite but is not all of $\Lambda^{\mathbb{N}}$, and x whose orbit closure is equal to $\Lambda^{\mathbb{N}}$.

Exercise 3.4. Construct a dynamical system with exactly k recurrent points.

Exercise 3.5. Construct a dynamical system with a dense set of recurrent points, but not all points are recurrent.

Exercise 3.6. Show that any dynamical system (X, G) contains a recurrent point.

Exercise 3.7. Show that for all $\alpha \in \mathbb{R}$ and all $\varepsilon > 0$, there exist $m, n \in \mathbb{Z}$ such that $|n\alpha - m| < \varepsilon$.

Exercise 3.8. Prove Kronecker's Theorem: show that that for all $\alpha \notin \mathbb{Q}$, $\gamma \in \mathbb{R}$ and $\varepsilon > 0$, there exist $m, n \in \mathbb{Z}$ such that $|n\alpha - \gamma - m| < \varepsilon$.

Exercise 3.9. Show that Kronecker's (one dimensional) Theorem is equivalent to showing that for any $\alpha \notin \mathbb{Q}$, $\{n\alpha \mod 1 \colon n \in \mathbb{N}\}$ is dense in $[0, 1]$.

Exercise 3.10. Show that the preimage of a recurrent point under a semiconjugacy need not be a recurrent point. More generally, show that every preimage of a recurrent point under a semiconjugacy might not be recurrent.

Exercise 3.11. Show that if x is recurrent under T^n for some $n \in \mathbb{N}$ with $n \geq 2$, then it is recurrent under T.

Exercise 3.12. Show that if x is recurrent under T, then x is recurrent under T^n for all $n \in \mathbb{N}$.

Exercise 3.13. Show that the metric on a symbolic system (X, T) is equivalent to the metric defined as follows: if $a, b \in \Lambda$, set $d'(a, b) = 1$ if $a \neq b$ and $d'(a, b) = 0$ otherwise, and then for $x, y \in X$, define

$$d(x, y) = \sum_{n=1}^{\infty} \frac{d'(x_i, y_i)}{2^n}.$$

Exercise 3.14. If (X, T) is a symbolic space, the *cylinder set* of the word w is defined to be $\{x \in X \colon x_1 \dots x_n = w\}$. Show that every cylinder set is both open and closed.

Exercise 3.15. Show that any closed and open set in a symbolic space is a finite union of cylinder sets.

Exercise 3.16. If Λ is a finite alphabet, show that the symbolic space $\Lambda^{\mathbb{N}}$ is a Cantor space. That is, show that it is compact, totally disconnected, and perfect.

Exercise 3.17. Show that any compact, totally disconnected metric space is homeomorphic to a symbolic space. (This ia a converse to Exercise 3.16.)

Exercise 3.18. Show that if $x \in \overline{\mathcal{O}_T^+(y)}$ and $y \in \overline{\mathcal{O}_T^+(z)}$, then $x \in \overline{\mathcal{O}_T^+(z)}$. If T is invertible, then show that if $x \in \overline{\mathcal{O}_T(y)}$ and $y \in \overline{\mathcal{O}_T(z)}$, then $x \in \overline{\mathcal{O}_T(z)}$.

Exercise 3.19. Show that the α-limit set and the ω-limit set of a point are nonempty, compact and T-invariant.

Exercise 3.20. Define the ω-*limit set* $\omega(Y)$ of a set $Y \subseteq X$ by

$$\omega(Y) = \bigcap_{n=1}^{\infty} \overline{\bigcup_{j>n} T^j(Y)}.$$

Show that for any $Y \subseteq X$, $\omega(Y)$ is a compact, T-invariant subset.

Exercise 3.21. Show that for any $Y \subseteq X$, $\omega(\omega(Y)) = \omega(Y)$.

Exercise 3.22. Show that for any $Y \subseteq X$, $\omega(Y) = \omega(\overline{Y})$.

Exercise 3.23. Show that the nonwandering set is closed, T-invariant, and contains $\omega(x)$ and $\alpha(x)$ for all $x \in X$. (In particular, by Exercise 3.19, it is nonempty.)

Exercise 3.24. Show that the closure of the recurrent points is contained in the nonwandering set.

Exercise 3.25. (Difficult) Give an example showing the closure of the recurrent points is not always equal to the nonwandering set.

Chapter 4

Minimality

4.1. Minimal systems

Dynamical systems which have no nontrivial subsystems are in a sense indecomposable. These are the minimal systems, and play a fundamental role in topological dynamics:

Definition 4.1. If (X, T) is a dynamical system, $Y \subseteq X$ is a *minimal set* if it is closed, T-invariant and contains no proper, closed, nonempty T-invariant subset. The system (X, T) is said to be *minimal* if contains no nontrivial minimal subsets.

Generalizing this terminiology, if $Y \subseteq X$ is a minimal set, then we refer to the subsystem (Y, T) as minimal.

It is immediate that if $x \in X$ is a fixed point, then $\{x\}$ is a minimal set. Similarly, the orbit of a periodic point is minimal. More generally, a minimal set is the orbit closure of any of its points:

Proposition 4.2. *If (X, T) is a dynamical system, then $Y \subseteq X$ is a minimal set if and only if $\mathcal{O}_T^+(y)$ is dense in Y for all $y \in Y$.*

Proof. Assume that Y is minimal. If the forward orbit of some $y \in Y$ is not dense in Y, then $\overline{\mathcal{O}_T^+(y)}$ is a proper, closed, nonempty subset of Y that is T-invariant, contradicting the minimality of Y. Conversely, if $Y \subseteq X$ is not minimal, then it contains some point whose forward orbit closure is contained in a proper, closed subset of Y. But then this orbit is not dense in Y. $\square$

This proposition gives another equivalent way of formulating minimality. In the system (X,T), $Y \subseteq X$ is minimal if and only if for all $x, y \in Y$, $x \in \overline{\mathcal{O}_T^+(y)}$. In particular, this immediately implies that any point in a minimal set is recurrent. The converse does not hold: a recurrent point need not lie in a minimal set (Exercise 4.1).

Applying Proposition 4.2 to X, we have that X itself is minimal if no proper, closed, nonempty subset of X is T-invariant. Thus the system (X,T) is minimal if and only if for each $x \in X$, $\{T^n x \colon n \in \mathbb{N}\}$ is dense in X.

When X is a metric space, it is often convenient to reformulate minimality in terms of the metric: $Y \subseteq X$ is minimal if and only if for all $x, y \in X$ and all $\varepsilon > 0$, there exists $n \in \mathbb{N}$ such that $d(T^n x, y) < \varepsilon$.

In our proof of the Birkhoff Recurrence Theorem (Theorem 3.1) we showed not only the existence of a recurrent point, but of a minimal set:

Proposition 4.3. *Any dynamical system (X,T) contains a minimal subsystem.*

Proof. Take the minimal set Z_0 found in the proof of Theorem 3.1. Restricting T to Z_0 gives a minimal subsystem (Z_0, T). □

Example 4.4. The system (X, Id) is minimal if and only if X consists of a single point.

Example 4.5. An automorphism A of a compact metrizable group G is minimal if and only if $G = \{e\}$, since $Ae = e$.

Example 4.6. The (one or two-sided) shift on k symbols is minimal if and only if $k = 1$. On the other hand, it is easy to construct minimal sets in a shift system. For example, when $k = 2$, $\{\overline{01}, \overline{10}\}$ is a minimal set.

Example 4.7. If $X = [0,1]$ and $Tx = x^2$, then X is not minimal since the orbit of any point other than 0 or 1 does not recur. The only minimal sets in X are $\{0\}$ and $\{1\}$.

A Kronecker system is minimal under a simple condition on the element defining the rotation:

Proposition 4.8. *If (X,T) is a Kronecker system and $Tx = ax$ for some fixed $a \in X$, then X is minimal if and only if $\{a^n \colon n \in \mathbb{N}\}$ is dense in X.*

Proof. Let e denote the identity in X. Since $\mathcal{O}_T^+(e) = \{a^n \colon n \in \mathbb{N}\}$, minimality implies that $\{a^n \colon n \in \mathbb{N}\}$ is dense. Conversely, assume that $\{a^n \colon n \in \mathbb{N}\}$ is dense. For any $x \in X$, we show that $\overline{\mathcal{O}_T^+(x)} = X$. Let $y \in X$ be arbitrary. Multiplication by the fixed element x is a continuous map. By assumption, we can approximate yx^{-1} arbitrarily well by iterates

a^{n_j} for appropriately chosen n_j. Thus we have that $\{a^{n_j}x\}$ approximates y arbitrarily well. Therefore $\mathcal{O}_T^+(x)$ is dense in X. $\square$

This implies that an irrational circle rotation is minimal, while a rational circle rotation is not. However, the (finite) orbit of a point under a rational circle rotation is a minimal set.

It is useful to characterize minimality in terms of open sets:

Proposition 4.9. *For a dynamical system (X, T), the following are equivalent:*

(1) The system (X, T) is minimal.

(2) For all $x \in X$, $\overline{\mathcal{O}_T^+(x)} = X$.

(3) The only closed subsets $Y \subseteq X$ with $TY \subseteq Y$ are $\varnothing$ and X.

(4) For any nonempty open set $U \subseteq X$, $\bigcup_{n=0}^{\infty} T^{-n}U = X$.

(5) For any nonempty open set $U \subseteq X$, there exist finitely many positive integers $n_1, n_2, \ldots, n_k$ such that $\bigcup_{j=1}^{k} T^{-n_j}U = X$.

Proof. (1) $\Leftrightarrow$ (2) $\Leftrightarrow$ (3). These equivalences follow immediately from the definition and Proposition 4.2.

(3) $\Rightarrow$ (4). Assume that $U \subseteq X$ is a nonempty open set. It follows that $\bigcup_{n=0}^{\infty} T^{-n}U$ is also an open set and so

$$Y = X - \bigcup_{n=0}^{\infty} T^{-n}U$$

is closed. Furthermore, $TY \subseteq Y$. If $Y \neq X$, then by hypothesis $Y = \varnothing$, which means that $\bigcup_{n=0}^{\infty} T^{-n}U = X$.

(4) $\Rightarrow$ (5). If U is a nonempty open set, then $\bigcup_{n=0}^{\infty} T^{-n}U = X$. By compactness, finitely many of the iterates of U cover all of X.

(5) $\Rightarrow$ (1). Assume that X is not minimal. If Y is a closed, invariant subset of X, set $U = X - Y$. Then $\bigcup_{j=1}^{k} T^{-n_j}U$ is not all of X, a contradiction. $\square$

Minimality means that the space does not split into smaller pieces on which the transformation acts; the transformation is "indecomposable." An optimistic conjecture would be that system can be partitioned into minimal systems. Unfortunately, this does not hold, as is easily seen in the example $x \mapsto x^2$ on $[0, 1]$. On the other hand, since any system does contain minimal subsystems, for many theorems that we prove it will suffice to work in a minimal system. In this sense, they are the fundamental building blocks of dynamical systems.

We conclude this section with two examples of minimal systems, both constructed using symbolic systems.

Example 4.10. Let $\Omega = \{0,1\}^{\mathbb{N}}$ and let T be the shift on Ω. Given a finite word w in the alphabet $\{0,1\}$, define w' to be the word obtained by replacing every 0 in w by a 1 and every 1 in w by a 0. Inductively, define words w_n by setting $w_0 = 0$ and $w_{n+1} = w_n w_n'$:

$$\begin{aligned}
w_0 &= 0\\
w_1 &= 01\\
w_2 &= 0110\\
w_3 &= 01101001\\
w_4 &= 0110100110010110\\
w_5 &= 01101001100101101001011001101001\\
&\vdots
\end{aligned}$$

In this inductive procedure, at the n-th step the first $2n$ entries of the sequence have been defined and remain unchanged thereafter. Thus it makes sense to look at the limiting sequence of this procedure, and we define ω to be the limit of the words w_n, as $n \to \infty$. This sequence ω is known as the *Morse sequence*. Taking $X = \overline{\{\mathcal{O}_T^+(\omega)\}}$, the system (X,T) is known as the *Morse system*. It is minimal and infinite (Exercise 4.12).

Example 4.11. Let $\Omega = \{0,1\}^{\mathbb{N}}$ and let T be the shift on Ω. Define words w_n in $\{0,1\}$ inductively by setting $w_1 = 0010$ and $w_{n+1} = w_n w_n 1 w_n$. Once again, this inductive procedure has a limit and we set ω to be the limit of the words w_n, as $n \to \infty$:

$$\omega = 0010001010010001000101001010010001010010\ldots$$

This ω is called the *Chacon sequence* and the system (X,T), where $X = \overline{\{\mathcal{O}_T^+(\omega)\}}$ is called the *Chacon system*. It is minimal and infinite (Exercise 4.13).

4.2. Constructive proof of the existence of minimal sets

In Theorem 3.1, we used Zorn's Lemma to show the existence of minimal subsystems in any system. However, it is interesting to note that one can give a constructive proof of this fact.

Constructive proof of Proposition 4.3. Assume that (X,T) is a dynamical system. Let $U_1, U_2, \ldots$ be open sets forming a basis for the topology of X. (Note that since X is a compact metric space, it is second countable.) We construct the minimal set inductively. Consider the set

$\bigcup_{j=0}^{\infty} T^{-j}U_1$. If this set is equal to all of X, then set $X_1 = X$. Otherwise remove this set from X and set

$$X_1 = X - \bigcup_{j=0}^{\infty} T^{-j}U_1.$$

In both cases, $X_1 \subset X$ is a T-invariant set. If $U_1 \cap X_1 \neq \varnothing$, then the orbit of every point in X_1 intersects U_1.

We now repeat this argument for $U_2 \cap X_1$. Thus if $\bigcup_{j=0}^{\infty} T^{-j}U_2 \supseteq X_1$, then set $X_2 = X_1$ and otherwise set $X_2 = X_1 - \bigcup_{j=0}^{\infty} T^{-j}U_2$. We continue the process inductively. Once X_{n-1} has been defined, we define X_n:

$$\begin{cases} X_n = X_{n-1} & \text{if } \bigcup_{j=0}^{\infty} T^{-j}U_n \supseteq X_{n-1} \\ X_n = X_{n-1} - \bigcup_{j=0}^{\infty} T^{-j}U_n & \text{otherwise.} \end{cases}$$

Define

$$E = \bigcap_{n=1}^{\infty} X_n.$$

Since X is compact, $E \neq \varnothing$. For each U_k with $U_k \cap E \neq \varnothing$, we have that $\bigcup_{j=0}^{\infty} T^{-j}U_k \supseteq E$. Since $U_1 \cap E, U_2 \cap E, \ldots$ is a basis for the topology of E, every point in E has a dense orbit. Therefore E is a minimal set. $\square$

4.3. Uniform Recurrence

In a dynamical system (X, T), if $x \in X$ is recurrent and U is a neighborhood of x, then we know that $\{n \in \mathbb{N} \colon T^n x \in U\}$ is a nonempty set. This set may be small, in the sense that the gaps between consecutive entries can grow arbitrarily large. For example, in the 2-shift consider a sequence obtained by starting with a 1, followed by as many 0's as in the initial string, repeating the initial string, followed by as many 0's as in the new initial string, and then iterating this process:

$$1010001010000000000000000101000101000000000000000\ldots.$$

By construction, each initial segment recurs and so this point is recurrent. However, the number of iterates that need to be considered until the point recurs grows like 2^n, where n is the length of the initial segment.

Points in a minimal system can not have this sort of behavior and return to any neighborhood with greater frequency than arbitrary recurrent points. We introduce a definition to quantify this frequency:

Definition 4.12. A subset $S \subseteq \mathbb{N}$ is *syndetic* if S has bounded gaps.

A syndetic set is sometimes also referred to as *relatively dense*. If S is syndetic, there exists some $M \in \mathbb{N}$ such that S has nontrivial intersection with every interval of length M in $\mathbb{N}$. For example, the set

$\{1, 10, 100, 1000, \ldots\}$ is not syndetic, nor is the set of natural numbers with no 1's in their (base 10) expansion. On the other hand, the sets $\{9, 19, 29, 39, \ldots\}$ and $\{9, 18, 29, 38, 49, 58, \ldots\}$ are both syndetic.

The definition of syndetic can be extended in the obvious way for subsets of $\mathbb{Z}$: $S \subseteq \mathbb{Z}$ is syndetic if S has bounded gaps.

We now use the notion of syndeticity to strengthen the idea of recurrence, requiring that a point return to any neighborhood of itself with some frequency.

Definition 4.13. If (X, T) is a dynamical system, $x \in X$, and U is a neighborhood of x, then

$$\{n \in \mathbb{N} \colon T^n x \in U\}$$

is called the set of *return times* of x to U.

Definition 4.14. Let (X, T) be a dynamical system. A point $x \in X$ is *uniformly recurrent* if for every neighborhood U of x, the set of return times of x to U is syndetic.

A uniformly recurrent point is sometimes also referred to as an *almost periodic* point, emphasizing that in counting iterates which return to a neighborhood of the point, this growth rate is the same as the growth rate of a periodic sequence.

Clearly any fixed point, and more generally any periodic point, is uniformly recurrent. There is an intermediate, weaker than periodic but strong than uniform recurrence:

Definition 4.15. In a dynamical system (X, T), a point $x \in X$ is *quasiperiodic* if for any neighborhood U of x, there exists $p \in \mathbb{N}$, called the *quasiperiod*, such that for all $n \in \mathbb{N}$, we have $T^{np}x \in U$.

If a point is periodic, then it is quasiperiodic, if a point is quasiperiodic then it is uniformly recurrent, and if a point is uniformly recurrent then it is recurrent. None of these implications can be reversed.

Example 4.16. In the symbolic system $\Lambda^{\mathbb{N}}$, where Λ is a finite alphabet, a point x is uniformly recurrent if and only if every word that occurs in x occurs along a syndetic set (for example, by considering the first entry of the word). By Lemma 3.6, any recurrent point has the form

$$aw_1aw_2aw_1aw_3aw_1aw_2aw_1aw_4aw_1aw_2aw_1aw_3aw_1aw_2aw_1a\ldots$$

for a symbol $a \in \Lambda$ and words $w_1, w_2, \ldots$ in the alphabet Λ. For uniform recurrence, the length of the words w_n must be bounded. Thus one can easily construct points in the 2-shift that are recurrent but are not uniformly recurrent.

Example 4.17. In an irrational rotation, each point is uniformly recurrent, but not quasiperiodic. Another example is given in Exercise 4.9.

We now explain the close relation between minimality and uniform recurrence.

Theorem 4.18. *If the dynamical system (X, T) is minimal, then every $x \in X$ is uniformly recurrent.*

Proof. Let $x \in X$ and let U be a neighborhood of x. By Proposition 4.9, there exist finitely many iterates $n_1, n_2, \ldots, n_k$ such that

$$X = \bigcup_{j=1}^{k} T^{-n_j} U.$$

Therefore, for each $n \in \mathbb{N}$, there exists $n_j = n_j(n)$ with $j \in \{1, 2, \ldots, k\}$ such that $T^n x \in T^{-n_j} U$. This means that $T^{n+n_j} x \in U$ and so the gaps in the return times to U are bounded by $\max\{n_1, n_2, \ldots, n_k\}$. □

As was noted in Proposition 4.3, it follows from our proof of the Birkhoff Recurrence Theorem that any system (X, T) contains a minimal system. This leads to:

Corollary 4.19. *Any dynamical system (X, T) contains uniformly recurrent points.*

The converse of Theorem 4.18 does not hold. If each $x \in X$ is uniformly recurrent, it is not necessarily true that X is minimal, as is seen by taking the union of two minimal sets. However, we do have something close to a converse:

Theorem 4.20. *Let (X, T) be a dynamical system. If $x \in X$ is uniformly recurrent, then its forward orbit closure $\overline{\mathcal{O}_T^+(x)}$ is a minimal, T-invariant, closed subset of X.*

Proof. Assume that $x \in X$ is uniformly recurrent. It suffices to show that if $y \in \overline{\mathcal{O}_T^+(x)}$, then $x \in \overline{\mathcal{O}_T^+(y)}$.

We proceed by contradiction. Assume that $x \notin \overline{\mathcal{O}_T^+(y)}$. Let U be a neighborhood of x such that $\overline{U} \cap \overline{\mathcal{O}_T^+(y)} = \varnothing$. Since x is uniformly recurrent, there exist positive integers $n_1, n_2, \ldots, n_k$ such that for all $n \in \mathbb{N}$, $T^{n+n_j} x \in U$ for some $j \in \{1, 2, \ldots, k\}$. This means that for all $n \in \mathbb{N}$,

$$T^n x \in \bigcup_{j=1}^{k} T^{-n_j} U$$

and so

$$\mathcal{O}_T^+(x) \subseteq \bigcup_{j=1}^{k} T^{-n_j} U.$$

But $y \in \overline{\mathcal{O}_T^+(x)}$ and so

$$y \in \bigcup_{j=1}^{k} \overline{T^{-n_j} U} \subseteq \bigcup_{j=1}^{k} T^{-n_j} \overline{U}.$$

Therefore $T^{n_j} y \in \overline{U}$ for $j = 1, 2, \ldots, k$ and so $\mathcal{O}_T^+(y) \cap \overline{U} \neq \varnothing$, a contradiction. □

The same proof shows that if T is invertible, then the full orbit closure $\overline{\mathcal{O}_T(x)}$ of any uniformly recurrent point is a minimal, T-invariant, closed subset of X.

Combining the characterizations of minimal systems, for metrizable systems we obtain:

Corollary 4.21. *If (X, T) is a metrizable dynamical system, then $x \in X$ is uniformly recurrent if and only if for all $\varepsilon > 0$, there exists $N \in \mathbb{N}$ such that some iterate of x among any consecutive N iterates of x come within ε of any point in $\overline{\mathcal{O}_T^+(x)}$.*

Proof. Assume that x is uniformly recurrent. By Theorem 4.20, $\overline{\mathcal{O}_T^+(x)}$ is a minimal, T-invariant, closed subset of X. If $U \subseteq X$ is open and $U \cap \overline{\mathcal{O}_T^+(x)} \neq \varnothing$, then by Proposition 4.9 and the minimality of $\overline{\mathcal{O}_T^+(x)}$, we have

$$\overline{\mathcal{O}_T^+(x)} \subseteq \bigcup_{n=1}^{\infty} T^{-n} U.$$

But since $\overline{\mathcal{O}_T^+(x)}$ is compact, only finitely many of the iterates of U are needed to cover this set. Thus there exist $n_1, n_2, \ldots, n_k \in \mathbb{N}$ such that

$$\overline{\mathcal{O}_T^+(x)} \subseteq \bigcup_{j=1}^{k} T^{-n_j} U.$$

Therefore, for all $y \in \overline{\mathcal{O}_T^+(x)}$, there exists $n \in \mathbb{N}$ with $1 \leq n \leq \max\{n_1, n_2, \ldots n_k\}$ such that $T^n y \in U$. Taking U with diameter less than ε, we have the statement. The converse follows immediately from the definitions. □

4.4. Diophantine approximation

We use the existence of uniformly recurrent points to derive some number theoretic results, starting with a refinement of Kronecker's Diophantine result. First we prove the dynamical statement:

Proposition 4.22. *In a Kronecker system, every point is uniformly recurrent.*

Proof. Assume that (X, T) is Kronecker system and that $Tx = ax$ for some $a \in X$. By Corollary 4.19, some point $x_0 \in X$ is uniformly recurrent. Assume that U is an open set containing the identity of X. By the continuity of multiplication by x_0 on the group X, we have that $Ux_0 = \{ux_0 : u \in U\}$ is an open set containing the point x_0. Choose $n \in \mathbb{N}$ such that $T^n x_0 \in Ux_0$. Then for any $x \in X$, $T^n x = a^n x = a^n x_0 x_0^{-1} x \in Ux_0 x_0^{-1} x = Ux$. The set of return times of x_0 to any neighborhood of itself is syndetic a syndetic set. It follows that the set of return times of x to any neighborhood of itself is also syndetic. Thus x is uniformly recurrent. $\square$

We can apply this to an irrational circle rotation. It follows from the proposition that every point is uniformly recurrent and so Theorem 4.20, the forward orbit of any point is minimal. Thus an irrational circile rotation is a minimal system. We use this to derive a statement about the distributions of iterates of an irrational point in the unit interval:

Corollary 4.23. *Assume that $\alpha \notin \mathbb{Q}$ and that $0 \le a < b \le 1$. Then the set*

$$\{n \in \mathbb{N} : a < n\alpha \mod 1 < b\}$$

is syndetic.

Proof. Consider the Kronecker system $T(x) = x + \alpha \mod 1$ on the circle $\mathbb{T}$. (As usual for a rotation on the circle, we use additive notation.) By Proposition 4.22, every point is uniformly recurrent and so the forward orbit of 0 returns to any neighborhood of itself with bounded gaps. By Proposition 4.8, the forward orbit of 0 is dense and so this orbit enters any open set with bounded gaps. Since this orbit is exactly

$$\{n\alpha \mod 1 : n \in \mathbb{N}\},$$

the corollary follows by applying this to the open interval (a, b). $\square$

This immediately implies Kronecker's Theorem (Theorem 1.10). We could also prove Corollary 4.23 directly, by taking at least $[1/\varepsilon] + 1$ iterates in the orbit of α. Then using the Pigeonhole Principle, at least two points lie within ε of each other. By translating, we obtain the needed approximation.

The multidimensional version of Kronecker's Theorem is left to Exercise 4.15.

4.5. Piecewise syndetic sets

The existence of uniformly recurrent points has implications for coloring theorems. As for other coloring results, we translate the combinatorial statement into a dynamical one, and then prove the dynamical statement using topological dynamics. We start by defining the combinatorial objects:

Definition 4.24. A set $A \subset \mathbb{N}$ (or of $\mathbb{Z}$) is *thick* if A contains arbitrarily long intervals. The set A is *piecewise syndetic* if A is the intersection of a syndetic set and a thick set.

A thick set is piecewise syndetic, as is a syndetic set. It is easy to construct an example of a piecewise syndetic set that is neither thick nor syndetic. Given a syndetic set, any thick set eventually contains an interval whose length is greater than the gap size of the syndetic set. Thus any thick set and any syndetic set have nontrivial intersection.

A finite coloring of the integers contains a monochromatic piecewise syndetic set:

Theorem 4.25. *If $\mathbb{N} = C_1 \cup C_2 \cup \ldots \cup C_r$ is a finite partition, then for some $j \in \{1, 2, \ldots, r\}$, C_j is piecewise syndetic.*

Proof. Consider the alphabet $\Lambda = \{1, 2, \ldots, r\}$ and let T be the shift on $\Lambda^{\mathbb{N}}$. Define $x \in \Lambda^{\mathbb{N}}$ by $x_n = j$ if and only if $n \in C_j$ and let $X = \overline{\{T^n x \colon n \geq 0\}}$. Since (X, T) is a compact dynamical system, by Corollary 4.19, X contains some uniformly recurrent point y. If the first entry y_1 of y is the symbol j, then j recurs syndetically in y with some gap M. Since y lies in the orbit closure of x, there exist translates of x that come arbitrarily close to y. Therefore, given any $N \in \mathbb{N}$, we can pick $n \in \mathbb{N}$ (depending on N) such that $T^n x$ and y agree for the first N entries. This means that $x_{n+1} = y_1, x_{n+2} = y_2, \ldots, x_{n+N} = y_N$. In particular, within the interval $[n+1, n+2, \ldots, n+N]$ the letter j occurs in x with gaps bounded by M. Thus the letter j occurs in x with gaps bounded by M on arbitrarily long intervals and so the set C_j is piecewise syndetic. $\square$

4.6. Minimality in more general dynamical systems

The definitions given for a dynamical system (X, T) where the transformation is generated by a single transformation T extend to a general dynamical system (X, G), where G is a semigroup of transformations acting continuously on X.

Definition 4.26. If (X, G) is a dynamical system, then $Y \subseteq X$ is *minimal* if it is closed, G-invariant and no proper, closed nonempty subset of Y is g-invariant for any $g \in G$.

It follows immediately that this is equivalent to $\{gy\colon g \in G\}$ is dense in Y for any $y \in Y$. An analogous version of Proposition 4.9 can be proven in this context (Exercise 4.20).

The notions of syndetic and uniform recurrence can also be extended to this context:

Definition 4.27. A subset S of an abelian topological group (or semigroup) G is *syndetic* if there exists a compact set $K \subseteq G$ such that for all $g \in G$, there exists some $k \in K$ with $gk \in S$.

Definition 4.28. If (X, G) is a dynamical system, $x \in X$ is *uniformly recurrent* if for every neighborhood U of x, $\{g \in G\colon gx \in U\}$ is syndetic.

Again, the properties of uniformly recurrent points can be proven in this context and we leave this to the exercises.

4.7. Transitivity

For some of the applications thus far, we have not used the full strength of minimality, but rather a weaker condition known as transitivity. Instead of all points having a dense orbit, we have only needed the existence of *some* point whose orbit is dense.

Definition 4.29. The dynamical system (X, T) is *(topologically) transitive* if there exists some $x \in X$ such that $\mathcal{O}^+_T(x)$ is dense in X. More generally, the dynamical system (X, G) is *transitive* if there exists $x \in X$ such that $\overline{\{gx : g \in G\}} = X$. A point whose orbit is dense is said to be a *transitive point*.

In keeping with previous conventions, we say more succinctly that T is transitive.

In a system with noninvertible transformation T, transitivity is a one-sided condition: we are only concerned with the existence of a point whose forward orbit is dense. The existence of a point with a dense full orbit does not necessarily imply the existence of a point with dense forward orbit (Exercise 4.25).

If there are no isolated points in X and if $x \in X$ has a dense forward orbit, then for any $m \in \mathbb{N}$, $\{T^n x\colon n \geq m\}$ is dense, since the first m iterates are only finitely many points and so do not accumulate anywhere.

One of the differences between minimality and transitivity is that a transitive system can contain a dense set of periodic points without all points being periodic. In a minimal system, if some point is periodic, then all points are periodic, since the orbit closure of any point is the entire system. Similarly in a minimal system, if some point has a dense orbit, then all points have dense orbit. On the other hand:

Example 4.30. Consider the two-sided shift (Ω, T) and a point $x = \{x_n\}_{n=-\infty}^{\infty} \in \Omega$. Then $T^m x = x$ for some $m \in \mathbb{Z}$ if and only if $x_n = x_{n+m}$ for all $n \in \mathbb{Z}$. Thus the points fixed by T^m have the form

$$(\ldots, x_{m-1}, x_0, x_1, \ldots, x_{m-1}, x_0, x_1, \ldots, x_{m-1}, x_0, x_1, \ldots),$$

where $x_0, x_1, \ldots, x_{m-1}$ can be chosen freely. These periodic points are dense in Ω, but clearly not all points are of this form.

An isometry which is transitive is minimal (Exercise 4.3). We generalize this:

Proposition 4.31. *Assume that (X, T) is a transitive metrizable dynamical system with metric d. If there exists a metric on X equivalent to d with respect to which T is an isometry, then T is minimal.*

Note that we only assume that there is *some* metric on X with this property.

Proof. Assume that d' is a metric on X such that $d'(x, y) = d'(Tx, Ty)$ for all $x, y \in X$. Since T is an isometry with respect to d', T is invertible and $d'(T^{-1}x, T^{-1}y) = d'(x, y)$ for all $x, y \in X$ (Exercise 2.13). Furthermore, we cannot have any isolated points without all points being isolated, in which case the result is easy. Thus we can assume that there are no isolated points.

By transitivity, we can choose $x \in X$ such that $\overline{\mathcal{O}_T^+(x)} = X$. Let $y \in X$ and we show that $\overline{\mathcal{O}_T^+(y)} = X$.

Let $z \in X$ and let $\varepsilon > 0$. Since the forward orbit of x is dense, there exist $n, m \in \mathbb{N}$ with $n > m$ so that $d'(T^m x, y) < \varepsilon/2$ and $d'(T^n x, z) < \varepsilon/2$. Thus

$$\begin{aligned} d'(z, T^{n-m}y) &\le d'(z, T^n x) + d'(T^n x, T^{n-m} y) \\ &= d'(z, T^n x) + d'(x, T^{-m} y) \\ &= d'(z, T^n x) + d'(T^m x, y) \\ &< \varepsilon. \end{aligned}$$

Therefore, $\overline{\mathcal{O}_T^+(y)} = X$. $\square$

There are many equivalent formulations of transitivity:

Proposition 4.32. *Assume that (X, T) is a metrizable dynamical system and that T is a homeomorphism. The following are equivalent:*

(1) T is transitive.

(2) If $U \subset X$ is a nonempty open set and $TU \subset U$, then U is dense.

(3) If $U, V \subset X$ are nonempty open sets, then there exists $n \in \mathbb{N}$ so that $T^n U \cap V \neq \varnothing$.

(4) *The set of transitive points is residual.*

Because of condition (2), in analogy with ergodic theory, topological transitivity is sometimes called *(topological) ergodicity.*

Proof. (1) $\implies$ (2). Assume that $\overline{\mathcal{O}_T^+(x)} = X$ and let $U \subset X$ be open, invariant and nonempty. Choose $m \in \mathbb{N}$ so that $T^m x \in U$. Then by the invariance of U, for any $k \in \mathbb{N}$, we have $T^{m+k}x = T^k T^m x \in T^k U \subset U$. Therefore, the forward orbit of $T^m x$ lies entirely in U. But since T is a homeomorphism, it is onto. Thus the orbit of $T^m x$ is also dense in X and so U is dense in X.

(2) $\implies$ (3). Assume that $U, V \neq \varnothing$ are open sets in X. Then $\bigcup_{n=0}^{\infty} T^n U$ is also open and is T-invariant. By hypothesis this union is dense, which implies in particular that $T^n U \cap V \neq \varnothing$ for some $n \geq 1$.

(3) $\implies$ (4). Let $\{V_j\}$ be a countable basis for the topology of X. (Such a basis exists since a compact metric space is second countable.) Then for each j, $W = \bigcup_{n=1}^{\infty} T^{-n} V_j$ is open. By hypothesis, for any nonempty open set U in X, there exists $m \in \mathbb{N}$ such that $T^m U \cap W \neq \varnothing$ and so W dense in X. Thus $V = \bigcap_{j=1}^{\infty} \bigcup_{n=1}^{\infty} T^{-n}(V_j)$ is a countable intersection of dense, open sets and so by the Baire Category Theorem, it is residual. Take any $x \in V$. This means that for all j, we have $x \in \bigcup_{n=1}^{\infty} T^{-n} V_j$ and for for some $n \in \mathbb{N}$, $T^n x \in V_j$. Thus for any $x \in V$, $\mathcal{O}_T^+(x)$ enters each V_j and since $\{V_j\}$ is a basis, this orbit is dense in X. Thus we have a residual set of transitive points.

(4) $\implies$ (1). By the Baire Category Theorem, a residual set is dense and thus is nonempty. Therefore we have a point with dense orbit. $\square$

In the absence of isolated points, we have only used the fact that T is a homeomorphism in order to prove that condition (2) implies (3).

A two-sided shift has a dense orbit and so is transitive. Since it has fixed points, it is not minimal. However, under certain mild conditions, the existence of a dense two-sided orbit implies that of a dense one-sided orbit.

Proposition 4.33. *Assume that (X, T) is a metrizable dynamical system and that X has no isolated points. If there is some point with a dense full orbit, then there is a point with a dense forward orbit.*

Proof. Assume that $\overline{\mathcal{O}_T(x)} = X$. Since there are no isolated points, the orbit of x not only enters any nonempty open set, but does so infinitely often. Let B_j be a sequence of balls around x whose radii tend to 0. Then the orbit of x enters each of these balls and so there exists a sequence of integers n_j with $|n_j| \to \infty$ so that $T^{n_j} x \to x$. This convergence holds for infinitely many n_k and so either for infinitely many positive or infinitely many negative indices. For any fixed $m \in \mathbb{Z}$, by taking n_k large enough, we have that $n_k + m$ is always positive or always negative, and so we also have that $T^{n_k+m}x \to T^m x$. Therefore, we either have

$$\overline{\mathcal{O}_T(x)} \subset \overline{\mathcal{O}_T^+(x)}$$

or

$$\overline{\mathcal{O}_T(x)} \subset \overline{\mathcal{O}_T^-(x)}.$$

In the first case, since the closure of the full orbit is all of X, the closure of the forward orbit is also all of X and we are done. In the second case, assume that U and V are nonempty open subsets of X. Since $\overline{\mathcal{O}_T^-(x)} = X$, there exist $i < j < 0$ so that $T^i x \in U$ and $T^j x \in V$. Therefore $T^j x = T^{j-i}T^i x \in T^{j-i}U$ and so $T^{j-i}U \cap V \neq \varnothing$. By Proposition 4.32, T is transitive. $\square$

Notes

The constructive proof of the Birkhoff Recurrence Theorem given in Section 4.2 appears in Weiss [**62**].

Most of the material of Section 4.3 is based on Furstenberg [**27**]. What we are calling uniform recurrence is actually what Birkhoff called recurrence. In his original work [**5**], Birkhoff showed the existence not only of a recurrent point in an arbitrary dynamical system, but of a uniformly recurrent point.

Theorem 4.25 was proved by Tom Brown [**11**] using combinatorial methods. Furstenberg [**27**] proved this theorem using topological dynamics.

Exercises

Exercise 4.1. Contstruct a dynamical system (X, T) and a point $x \in X$ such that x is recurrent but x does not lie in any minimal set.

Exercise 4.2. Show that if the dynamical system (X, T) is minimal, then T is onto.

Exercise 4.3. Show (without using Proposition 4.31) that if T is an isometry of a compact metric space X and some $x \in X$ is transitive, then every $x \in X$ is transitive.

Exercise 4.4. Show that any minimal set in a dynamical system (X, T) is contained in the nonwandering set of T.

Exercise 4.5. Assume that T and S are rotations on the circle $\mathbb{T}$. Show that there exists a nonempty minimal closed set $A \subseteq \mathbb{T}$ that is invariant under both S and T.

Exercise 4.6. Show that there exist minimal systems (X, T) and (Y, S) such that $(X \times Y, T \times S)$ is not necessarily minimal. Find conditions on the systems (X, T) and (Y, S) such that $(X \times Y, T \times S)$ is minimal.

Exercise 4.7. Show that an irrational circle rotation is not quasiperiodic.

Exercise 4.8.
(1) Show that the preimage of a uniformly recurrent point under a semi-conjugacy need not be uniformly recurrent.
(2) Show that one can have a uniformly recurrent point none of whose preimages under a semi-conjugacy are uniformly recurrent.

Exercise 4.9. Show that the adding machine of Exercise 2.15 is minimal. Show that every point in this system is quasiperiodic, but none is periodic.

Exercise 4.10. The *Champernowne sequence* ω is defined by concatenating in lexicographic order all finite words on 2 symbols:

$$\omega = 0100011011000001010011100101111\ldots.$$

Show that this sequence is recurrent (under the shift map), but is not uniformly recurrent. Show that the orbit of ω is infinite.

Exercise 4.11. Consider the shift Ω on k symbols. Show that there exists a nonperiodic point in Ω that is uniformly recurrent. (Thus the orbit closure of this point is infinite and minimal.)

Exercise 4.12. Show that the Morse system (Example 4.10) contains infinitely many points and is minimal.

Exercise 4.13. Show that the Chacon system (Example 4.11) contains infinitely many points and is minimal.

Exercise 4.14. Show that there exists a uniformly recurrent real valued sequence $\omega(n)$ such that

$$\lim_{N\to\infty} \frac{1}{2N+1} \sum_{n=-N}^{N} \omega(n)$$

does not exist.

Exercise 4.15.
(1) Show that the dynamical system $(\mathbb{T}^k, T)$ is minimal, where $\alpha = (\alpha_1, \alpha_2, \ldots, \alpha_k) \in \mathbb{T}^k$, $\{1, \alpha_1, \alpha_2, \ldots, \alpha_k\}$ is a rationally independent set, and $T\colon X \to X$ is defined by $T(x) = x + \alpha$.
(2) Use this to derive a result in Diophantine approximation.

Exercise 4.16. Show that syndeticity is not partition regular, meaning that there is a finite partition of the integers such that no piece of the partition is syndetic.

Exercise 4.17. Assume that $A \subseteq \mathbb{N}$ is piecewise syndetic. If $A = A_1 \cup A_2 \cup \ldots \cup A_r$ is a finite partition of A, show that for some $j \in \{1, 2, \ldots, r\}$, A_j is piecewise syndetic.

Exercise 4.18. Show that the hypothesis that $\alpha_1, \alpha_2, \ldots, \alpha_k$ are rationally independent is necessary in Theorem 1.13.

Exercise 4.19. Using the existence of a uniformly recurrent point in any dynamical system, state and prove a strengthening of Hilbert's Theorem (Theorem 3.7).

Exercise 4.20. State and prove a version of Proposition 4.9 for a dynamical system (X, G), where G is a group of homeomorphisms acting continuously on X.

Exercise 4.21. Show that if the dynamical system (X, G) is minimal, then every point in X is uniformly recurrent. Conclude that any dynamical system (X, G) contains uniformly recurrent points.

Exercise 4.22. Show that in a dynamical system (X, G), if $x \in X$ is uniformly recurrent, then $\overline{\{gx\colon g \in G\}}$ is a closed minimal G-invariant subset of X.

Exercise 4.23. This exercise generalizes Exercise 4.17. If (X, G) is a dynamical system, a subset of G is *thick* if it has nonempty intersection with every syndetic set in G. A *piecewise syndetic set* is defined to be the intersection of a thick set and a syndetic set. Show that if $S \subseteq G$ is a piecewise syndetic set and $S = S_1 \cup S_2 \cup \ldots \cup S_r$ is a finite partition, then for some $j \in \{1, 2, \ldots, r\}$, S_j is piecewise syndetic.

Exercise 4.24. Show that the shift on k symbols is transitive.

Exercise 4.25. Construct a dynamical system with infinitely many points that have dense full orbit, but no point with dense forward orbit.

Exercise 4.26. Show that the product of two transitive systems need not be transitive.

Exercise 4.27. If the dynamical system (X, T) has no isolated points, show that (X, T) is transitive if and only if there exists some point in X whose ω-limit set is dense in X.

Exercise 4.28. State and prove a version of Proposition 4.32 for a dynamical system (X, G), where G is a group of homeomorphisms acting continuously on X.

Exercise 4.29. Show that if (X, G) is transitive, there exists a dense G_δ subset of points with dense orbit. (A G_δ is defined to be a countable intersection of open sets; see Appendix C.)

Chapter 5

Group extensions and factor maps

5.1. Lifting properties from factors

If one wishes to show that a dynamical system has a certain property, it is often easier to show that some subsystem has this property. Of course, this does not necessarily imply that the whole system has this property. Yet, one can still gain information if the subsystem is well chosen. Sometimes one can lift the property from the subsystem to the general system; sometimes in order to derive a combinatorial consequence it suffices to prove the property on a subsystem. We make these notions of subsystems and lifting precise, particularly defining some terms to clarify the unequal roles of the subsystem and the whole system.

Definition 5.1. If (X, T) and (Y, S) are dynamical systems, $\phi\colon X \to Y$ is a *semiconjugacy* if ϕ is a continuous, surjective mapping such that $\phi(Tx) = S(\phi x)$ for all $x \in X$. The system (Y, S) is a *factor* of (X, T) if there is a semiconjugacy $\phi\colon X \to Y$. The map ϕ is called the *factor map* and the system (X, T) is said to be an *extension* of (Y, S).

We often omit explicit mention of the transformations on the spaces and say that Y is a factor of X or that X an extension of Y.

In Example 2.11, we showed that the doubling map on $\mathbb{T}$ and the squaring map on $\mathcal{S}^1$ are conjugate. Thus each system is a factor of the other system. More interesting examples arise when this symmetry is broken:

Example 5.2. If (X, T) and (Y, S) are dynamical systems, then each is a factor of the product system $(X \times Y, T \times S)$. The projection map onto the first coordinate $\pi_1 \colon X \times Y \to X$, defined by $(x, y) \mapsto x$, is continuous and onto. Similarly, projection onto the second coordinate $\pi_2 \colon X \times Y \to Y$, defined by $(x, y) \mapsto y$, is continuous and onto. Therefore π_1 and π_2 are factor maps.

Example 5.3. A one-sided shift on k symbols is a factor of the two-sided shift on k symbols, where the factor map takes the bi-infinite sequence $x = (\dots, x_{-1}, x_0, x_1, \dots)$ to the one-sided sequence $(x_1, x_2, \dots)$.

Example 5.4. Let $\Omega = \{0,1\}^{\mathbb{Z}}$ and let T be the shift on Ω. Let (Y, S) be a dynamical system and assume that $f \colon Y \to \{0,1\}$ is a continuous map. Define $\phi \colon Y \to \Omega$ by

$$y \mapsto (\dots, f(S^{-1}y), f(y), f(Sy), f(S^2y), \dots).$$

Since $\phi(Y)$ is a closed, shift invariant subset of Ω, we have that $(\phi(Y), T)$ is a subsystem of (Ω, T). It is a factor of (Y, S) with factor map ϕ.

There is a natural identification between a system and a factor:

Definition 5.5. If (Y, S) is a factor of (X, T) with factor map $\phi \colon X \to Y$ and $y \in Y$, then the *fiber* above y is defined to be $\{\phi^{-1}(y) \colon y \in Y\}$.

If (Y, S) is a factor of (X, T), then we can identify the system (X, T) with the fibers $\{\phi^{-1}(y) \colon y \in Y\}$. Thus the dynamics inside a factor place constraints on the dynamics of the whole system.

It is clear that if Y is a factor of X, then the image of a fixed point is a fixed point under the factor map. The same holds for periodic and for recurrent points (Exercise 5.2). We are interested in the case that the converse holds, meaning that a fixed, periodic or recurrent point lifts to a fixed, periodic or recurrent point. In general, this is too much to expect: it is possible to construct extensions that do not preserve any of these properties (Exercises 2.12 and 3.10). However, under certain conditions on the factor map, these properties do lift:.

Definition 5.6. Assume that (Y, S) is a dynamical system, G is a compact metrizable group and $\phi \colon Y \to G$ is a continuous map. Let $X = Y \times G$ and define $T \colon X \to X$ by

$$(y, g) \mapsto (Sy, \phi(y)g).$$

The dynamical system (X, T) is a *group extension* of (Y, S) via the group G and factor map ϕ.

A group extension (X, T) of the system (Y, S) via the group G is sometimes also referred to as the *skew product* of Y and G. When Y is a metric space, X inherits the natural product measure from the metrics on Y and on G.

It follows immediately from the definition that X is an extension of Y, since the projection $(y, g) \mapsto y$ is a homomorphism of X onto Y.

Example 5.7. Let $T\colon \mathbb{T}^2 \to \mathbb{T}^2$ be given by $T(x,y) = (x+\alpha, x+y+\alpha)$. Then $(\mathbb{T}^2, T)$ is a group extension of the Kronecker system $(\mathbb{T}, S)$, where $S(x) = x + \alpha$ on $\mathbb{T}$. The compact group G defining this extension is $\mathbb{T}$ and the extension map ϕ is the identity.

Definition 5.8. Assume that (Y,S) is a dynamical system, G is a compact metrizable group, and let $(X = Y \times G, T)$ be a group extension of (Y,S). If $h \in G$, the *right translation* $R_h\colon X \to X$ is defined by $R_h(y,g) = (y, gh)$ for all $(y,g) \in X$. Similarly, the *left translation* $L_h\colon X \to X$ is defined by $L_h(y,g) = (y, hg)$ for all $(y,g) \in X$. An *automorphism* of (X,T) is a homeomorphism $P\colon X \to X$ that commutes with T.

An automorphism of X preserves recurrence properties of points. If $x \in X$ is recurrent under T, then $x \in \overline{\{T^n x\colon n \in \mathbb{N}\}}$. Therefore if P is an automorphism of (X,T), then $Px \in \overline{\{PT^n x\colon n \in \mathbb{N}\}} = \overline{\{T^n Px\colon n \in \mathbb{N}\}}$, and so Px is also recurrent under T.

Right (and left) translation act continuously on X, and are examples of automorphisms of X:

Lemma 5.9. *Assume that (X,T) is a group extension of (Y,S) via the group G and factor map ϕ. Then for all $h \in G$, right translation R_h and left translation L_h are automorphisms of X.*

Proof. It suffices to show this for the right translation R_h, as the proof for the left translation is analogous.

Since G is a group, R_h is a homeomorphism and we are left with checking that R_h and T commute for any $h \in G$. For all $(y,g) \in X$,

$$\begin{aligned} R_h\bigl(T(y,g)\bigr) =& R_h\bigl(Sy, \phi(y)g\bigr) = \bigl(Sy, \phi(y)gh\bigr) \\ =& T(y, gh) = T\bigl(R_h(y,g)\bigr). \end{aligned}$$

□

We use this to show that in a group extension, not only is some point in the preimage recurrent, but every point in the preimage is recurrent:

Theorem 5.10. *Assume that the dynamical system (X,T) is a group extension of (Y,S) via the group G. If $y \in Y$ is recurrent, then $(y,g) \in X$ is recurrent under T for all $g \in G$.*

Proof. Let e denote the identity in G. We first show that (y,e) is recurrent. Considering the orbit of (y,e) under T, recurrence of y implies that there exists $g_1 \in G$ such that

$$(y, g_1) \in \overline{\{T^n(y,e)\colon n \in \mathbb{N}\}}.$$

Applying the right translation R_{g_1}, we have that

$$(y, g_1^2) \in \overline{\{T^n(y, g_1) \colon n \in \mathbb{N}\}},$$

which in turn implies that

$$(y, g_1^2) \in \overline{\{T^n(y, e) \colon n \in \mathbb{N}\}}.$$

Iterating, we have that for each $m \in \mathbb{N}$,

$$(y, g_1^m) \in \overline{\{T^n(y, e) \colon n \in \mathbb{N}\}}.$$

But $\{g_1^m \colon m \in \mathbb{N}\}$ is the orbit of e in the compact group $G_1 = \overline{\{g_1^m \colon m \in \mathbb{N}\}}$, endowed with the transformation of rotation by g_1. In particular, it is a Kronecker system and so by Example 3.4 every point is recurrent. This means that

$$(y, e) \in \overline{\{T^n(y, e) \colon n \in \mathbb{N}\}}$$

and so (y, e) is recurrent. For arbitrary (y, g), note that $R_g(y, e) = (y, g)$. By Lemma 5.9, R_g is an automorphism and so (y, g) is also recurrent. $\square$

5.2. Diophantine approximation

We use the lifting of recurrent points to obtain classically results in Diophantine approximation:

Corollary 5.11. *For all $\alpha \in \mathbb{R}$ and all $\varepsilon > 0$, there exist $n, m \in \mathbb{Z}$ such that $|n^2\alpha - m| < \varepsilon$.*

Proof. Let $T \colon \mathbb{T}^2 \to \mathbb{T}^2$ be given by

$$T(x, y) = (x + \alpha, 2x + y + \alpha).$$

Then $(\mathbb{T}^2, T)$ is a group extension of the Kronecker system $x \mapsto x + \alpha$ on $\mathbb{T}$. Therefore every point in $\mathbb{T}^2$ is recurrent under T. In particular, the orbit of $(0, 0)$ under T returns to itself. Since this orbit is given by $T^n(0, 0) = (n\alpha, n^2\alpha)$, we obtain the statement. $\square$

More generally, via a series of group extensions we prove a Diophantine result for any polynomial:

Theorem 5.12. *Assume that $p(x)$ is a polynomial with real coefficients and $p(0) = 0$. Then for all $\varepsilon > 0$, there exist $n \in \mathbb{N}$ and $m \in \mathbb{Z}$ such that $|p(n) - m| < \varepsilon$.*

Proof. Assume that $p(x)$ has degree $d \geq 0$. Inductively, define

$$\begin{aligned}
p_d(x) &= p(x) \\
p_{d-1}(x) &= p_d(x+1) - p_d(x) \\
p_{d-2}(x) &= p_{d-1}(x+1) - p_{d-1}(x) \\
&\vdots \\
p_1(x) &= p_2(x+1) - p_2(x) \\
p_0(x) &= p_1(x+1) - p_1(x).
\end{aligned}$$

Then $p_j(x)$ has degree less than or equal to j for $j = 0, 1, \ldots, d$ and thus $p_0(x) = \alpha$ for some constant α. Define $T_1 \colon \mathbb{T}^1 \to \mathbb{T}^1$ by $T_1(x) = x + \alpha$. For $j \geq 2$, define $T_j \colon \mathbb{T}^j \to \mathbb{T}^j$ by

$$T_j(x_1, x_2, x_3, \ldots, x_j) = (x_1 + \alpha, x_2 + x_1, x_3 + x_2, \ldots, x_j + x_{j-1}).$$

(As usual, we omit the $\bmod 1$ from each term.) Then for $j = 1, \ldots, d$, each $(\mathbb{T}^j, T_j)$ is a group extension of the dynamical system $(\mathbb{T}^{j-1}, T_{j-1})$. By Theorem 5.10, each point is recurrent.

Consider the orbit of $\big(p_1(0), p_2(0), \ldots, p_d(0)\big)$ under $T = T_d$. Using the definition of the polynomials p_i, we have that the $(n+1)^{st}$ iterate under T is

$$\big(p_1(n), p_2(n), \ldots, p_d(n)\big).$$

Since this point is recurrent, it returns arbitrarily close to

$$(p_1(0), p_2(0), \ldots, p_d(0)).$$

But $p_d(n) = p(n)$, and so $p(n)$ returns arbitrarily close to $0 \mod 1$, which is the statement of the theorem. □

5.3. Lifting uniformly recurrent points

We have already seen that group extensions lift recurrent points to recurrent points. We now show that the same result holds for uniformly recurrent points:

Theorem 5.13. *Assume that (X, T) is a group extension of the dynamical system (Y, S) via the group G and factor map ϕ. If $y_0 \in Y$ is uniformly recurrent in (Y, S), then for any $g_0 \in G$, (y_0, g_0) is uniformly recurrent in $(Y \times G, T)$.*

Proof. Assume that $y_0 \in Y$ is uniformly recurrent. Then its orbit closure is a minimal set and so without loss of generality, we can assume that the system (Y, S) is minimal. Let $Z \subseteq Y \times G$ be a minimal set under the transformation $T(y, g) = (Sy, \phi(y)g)$, where $\phi \colon Y \to G$ is the factor map. Let $\pi \colon Y \times G \to Y$ denote projection onto the first coordinate, meaning that $\pi(y, g) = y$.

Since π is a homomorphism, $\pi(Z)$ is a minimal subset of Y (Exercise 5.8). As Y itself is minimal, $\pi(Z) = Y$. Consider right translation $R_{g'}\colon Y \times G \to Y \times G$, where $R_{g'}(y, g) = (y, gg')$. Since this is an automorphism of $(Y \times G, T)$, we also have that $R_{g'}(Z)$ is a minimal subset of $Y \times G$ and

$$\bigcup_{g' \in G} R_{g'}(Z) = \pi^{-1}\{\pi(Z)\} = \pi^{-1}(Y) = Y \times G.$$

Therefore each point of $Y \times G$ lies in a minimal subset and so each point is uniformly recurrent. □

We use this to improve the Diophantine results for polynomials, showing that not only can we solve the Diophantine inequality, but the set of solutions is syndetic. Hermann Weyl proved the multi-dimensional version:

Corollary 5.14. *Assume that $p_1(x), p_2(x), \ldots, p_k(x)$ are polynomials with real coefficients and $p_j(0) = 0$ for $j = 1, 2, \ldots, k$. Then for any $\varepsilon > 0$,*

$$\{n \in \mathbb{Z}\colon |p_j(n) - m_j| < \varepsilon \quad \text{for some } m_1, m_2, \ldots, m_k \in \mathbb{Z}\}$$

is syndetic.

Proof. Combining Theorem 5.13 and the group extensions defined in Theorem 5.12, we have that every point in these dynamical systems is uniformly recurrent. □

Notes

Most of the material in this chapter is based on Furstenberg [**27**].

Corollary 5.11 was proved by Hardy and Littlewood [**35**], using geometric properties of the numbers. The idea for another proof is contained in Exercise 8.7. Theorem 5.12 and Corollary 5.14 were originally proved by Weyl [**63**], using number theoretic methods. The proofs given here are based on Furstenberg [**27**].

Exercises

Exercise 5.1. If (X, T) and (Y, S) are minimal systems, S is invertible and $\phi\colon X \to Y$ is a semiconjugacy, show that for any nonempty open set U in X, $\phi(U)$ has nonempty interior.

Exercise 5.2. If (Y, S) is a factor of (X, T), show that a fixed point in X maps under the factor map to a fixed point in Y. Show that a recurrent point in X maps to a recurrent point in Y.

Exercise 5.3. Show that a factor of a transitive system is transitive.

Exercise 5.4. Show that a factor of a minimal system is minimal.

Exercise 5.5. Show that a factor map $\phi\colon X \to Y$ defines an equivalence relation $\sim$ on X, where $x \sim x'$ if and only if $\phi(x) = \phi(x')$. Conversely, show that an equivalence relation on X defines a factor of X.

Exercise 5.6. Assume that S is an automorphism of a minimal system (X, T) and assume that $Sx = x$ for some $x \in X$. Show that S is the identity map.

Exercise 5.7. Use a particular choice of a group extension to show that if $x \in X$ is a recurrent point under T, then x is recurrent under T^n for any $n \in \mathbb{N}$. (Group extensions are not needed for this and this problem has already appeared in Exercise 3.12.)

Exercise 5.8. Show that if (Y, S) is a factor map of (X, T) with factor map ϕ and $Z \subseteq X$ is minimal, then $\phi(Z)$ is minimal.

Exercise 5.9. Show that x is uniformly recurrent under T if and only if it is uniformly recurrent under T^n for any $n \in \mathbb{N}$.

Question 5.15. *Given a dynamical system (X, T), does there always exist $x \in X$ such that $(x, x, \ldots, x)$ is uniformly recurrent under $T \times T^2 \times \ldots \times T^k$?*

Chapter 6

Complexity

6.1. Alphabets and words

Definition 6.1. A finite set $\mathcal{A}$ of symbols is called the *alphabet*. A *word* $w = w_1 \dots w_N$ with $w_1, \dots, w_N \in \mathcal{A}$ is a finite string of elements from $\mathcal{A}$ and its length $|w|$ is the number of elements N in the string. The set of all words of length $N \geq 0$ is denote $\mathcal{W}_N(\mathcal{A})$, with the convention that the word of length 0 is the empty word, and the set of all words is denoted $\mathcal{W}(\mathcal{A})$.

Thus

$$\mathcal{W}(\mathcal{A}) = \bigcup_{N \geq 0} \mathcal{W}_N(\mathcal{A}).$$

Given two words $w, v \in \mathcal{W}$, the *concatenation* wv is the word of length $|w| + |v|$ formed by the string w followed by the string v. Concatenation is an associative operation and the empty word is the unit element of the operation.

Definition 6.2. A *one sided infinite word* in the alphabet $\mathcal{A}$ is an element $\eta = (\eta_n)_{n \in \mathbb{N}} \in \mathcal{A}^{\mathbb{N}}$. A *two sided infinite word* or *bi-infinite word* in the alphabet $\mathcal{A}$ is an element $\eta = (\eta_n)_{n \in \mathbb{Z}} \in \mathcal{A}^{\mathbb{Z}}$. A finite word $w = w_1 \dots w_k$ *occurs in* the (finite, one sided infinite, or bi-infinite) word $\eta = (\eta_n)$ if there exists some m such that $\eta_m = w_1, \eta_{m+1} = w_2, \dots, \eta_{m+k-1} = w_k$. We also say that the word η *contains* the word w.

Note that a word can refer either to a finite sequence of letters in the alphabet $\mathcal{A}$ or to an infinite sequence. It is usually clear from the context which is meant, and if confusion may arise, we write finite or infinite word, with the convention that infinite could be either one sided or two sided. We usually reserve Greek letters for infinite words and Roman letters for finite words.

The dynamical notions we have already defined extend to similar definitions for infinite words:

Definition 6.3. An infinite word $\eta \in \mathcal{A}^{\mathbb{N}}$ (respectively, a word $\eta \in \mathcal{A}^{\mathbb{Z}}$) is *periodic* if there exists $m \in \mathbb{N}$ such that $\eta_n = \eta_{n+m}$ for all $n \in \mathbb{N}$ (respectively, for all $n \in \mathbb{Z}$). For periodic η, the smallest such m is called the *period* of η.

An infinite word η (in $\mathcal{A}^{\mathbb{N}}$ or in $\mathcal{A}^{\mathbb{Z}}$) is *recurrent* if every finite word $w \in \eta$ that occurs in η occurs in η infinitely often.

An infinite word η (in $\mathcal{A}^{\mathbb{N}}$ or in $\mathcal{A}^{\mathbb{Z}}$) is *uniformly recurrent* if for every finite word $w \in \eta$ that occurs in η, there exists $M = M(w)$ such that any word of length M that occurs in η contains the word w.

It follows immediately from the definitions that a periodic word is recurrent and a uniformly recurrent word is recurrent, but the converse implications are false (Exercise 6.1).

6.2. Complexity of words

Definition 6.4. For an infinite word η, we define the *complexity function* $P_\eta : \mathbb{N} \to \mathbb{N}$ by setting $P_\eta(n)$ to be the number of distinct words of length n that occur in η.

By definition, $P_\eta(1)$ is the number of distinct letters that appear in η and in particular, for any non-constant η, we have that $P_\eta(1) \geq 2$. It also follows immediately from the definition that for all $n \in \mathbb{N}$,

$$P_\eta(n) \leq P_\eta(n+1). \tag{6.1}$$

Complexity and periodicity are closely related:

Theorem 6.5. *For recurrent $\eta \in \mathcal{A}^{\mathbb{Z}}$, the following are equivalent:*

(1) η is periodic.

(2) $P_\eta(n)$ is bounded for all $n \in \mathbb{N}$.

(3) There exists $n_0 \in \mathbb{N}$ such that $P_\eta(n_0) \leq n_0$.

Proof. The implication (1)⇒(2) is obvious by considering the period of η and the implication (2)⇒(3) is trivial. We are left with showing (3)⇒(1).

We can assume that $P_\eta(1) \geq 2$, or the statement trivially holds. Combining (6.1) and the assumption that there exists $n_0 \in \mathbb{N}$ such that $P_\eta(n_0) \leq n_0 - 1$, we have that there exists $k \in \mathbb{N}$ with $k \leq n_0$ such that

$P_\eta(k) = P_\eta(k+1)$. This means that for any word w of length k, there is a unique way to continue the word to the right and a unique way to continue the word to the left. Since each word of length k appears infinitely often, in particular a word of length k occurs twice. As each such occurrence is continued in a unique way, once we have the same window of length k twice, it continues identically in both directions and thus is periodic. □

Note that the assumption of recurrence was only used in the implication (3)⇒(1). Rephrasing this result, and using the fact that there are only finitely many words of any given length, we have the basic relation between complexity and periodicity that was proven by Morse and Hedlund:

Corollary 6.6. *A word $\eta \in \mathcal{A}^{\mathbb{Z}}$ is periodic if and only if there exists $n_0 \in \mathbb{N}$ such that $P_\eta(n_0) \leq n_0$.*

More generally, we have:

Corollary 6.7. *If $\eta \in \mathcal{A}^{\mathbb{Z}}$ is recurrent, then either η is periodic or $P_\eta(n) \geq n+1$ for all $n \in \mathbb{N}$.*

The class of non-periodic sequences achieving the minimal complexity of Corollary 6.7 are called *Sturmian sequences*, and are studied further in Section 6.4. For one sided infinite sequences, we have a similar result (see Exercise 6.2).

6.3. An associated dynamical system

As before, we construct a dynamical system associated to a sequence $\eta \in \mathcal{A}^{\mathbb{Z}}$, by building a compact metric space with a shift. Namely, endow $\mathcal{A}$ with the discrete topology and $\mathcal{A}^{\mathbb{Z}}$ with the product topology. Setting $X = \mathcal{A}^{\mathbb{Z}}$, we define the shift $T\colon X \to X$ by

$$T\eta(n) = \eta(n+1).$$

It is an easy exercise to check that the shift T is continuous, and more ore generally, for all $m \in \mathbb{Z}$, this defines the continuous iterates $T^m\colon X \to X$ by

$$T^m\eta(n) = \eta(m+n).$$

Defining

$$\mathcal{O}(\eta) = \{T^m\eta\colon m \in \mathbb{Z}\}$$

and setting $X_\eta = \overline{\mathcal{O}(\eta)}$, we obtain a topological dynamical system (X_η, T) (Exercise 6.3).

Periodicity of the sequence η is reflected in the associated dynamical system. Namely, a sequence η is periodic if and only if the orbit $\mathcal{O}(\eta)$ is finite, meaning that the associated dynamical system (X_η, T) is finite.

Definition 6.8. Given $\eta \in \mathcal{A}^{\mathbb{Z}}$ and a finite word $w = w_m \dots w_{m+k}$ that occurs in η, the *cylinder set* associated to w is the set

$$\{\rho \in X_\eta \colon \rho_m = w_m, \dots, \rho_{m+k} = w_{m+k}\}.$$

The cylinder sets are clopen sets and form a basis for the topology of X_η (Exercise 6.4).

For one sided $\eta \in \mathcal{A}^{\mathbb{N}}$, the same constructions of a dynamical system and cylinder sets apply.

6.4. Sturmian sequences

Definition 6.9. A recurrent sequence $\eta \in \mathcal{A}^{\mathbb{Z}}$ is said to be *Sturmian* if $P_\eta(n) = n+1$ for all $n \in \mathbb{N}$.

The restriction of the definition to recurrent sequences is to eliminate trivial examples, such as the sequence with a 1 at η_0 and 0 everywhere else.

The dynamical system associated to a Sturmian sequence is minimal (Exercise 6.8).

It turns out that Sturmian sequences can be produced from rotations. For $\alpha \in [0,1)$, recall that the rotation $R_\alpha \colon [0,1] \to [0,1]$ is defined by $R_\alpha(x) = x+\alpha \pmod 1$. Consider the partition $\mathcal{P}$ of $[0,1)$ into the intervals $[0, 1-\alpha)$ and $[1-\alpha, 1)$. Fixing $x \in [0,1)$, we code the orbit of x under the rotation $R = R_\alpha$ by which interval of this partition the iterates of x fall into. Namely, define a bi-infinite sequence $\eta \in \{0,1\}^{\mathbb{Z}}$ by

$$\eta_n = \begin{cases} 0 & \text{if } R^n x \in [0, 1-\alpha) \\ 1 & \text{if } R^n x \in [1-\alpha, 1). \end{cases} \tag{6.2}$$

This is a coding of the orbit of the rotation with respect to the partition $\{[0,1-\alpha), [1-\alpha,1)\}$ of $[0,1]$. The choice of how to deal with the endpoints of the intervals is arbitrary and alternately, we could have taken the coding corresponding to the partition $\{(0, 1-\alpha], (1-\alpha, 1]\}$. For irrational α, the codings obtained with respect to each of these partitions would agree after a certain point, and so we work with the first partition.

Set $I_0 = [0, 1-\alpha)$ and $I_1 = [1-\alpha, 1)$. Given a finite word $w = w_1 \dots w_n$ in the alphabet $\{0,1\}$, define the set $I(w)$ generated by it as

$$I(w) = I(w_1 \dots w_n) = \bigcap_{i=0}^{n-1} R_{-i}(I_{w_{i+1}}).$$

Then the sequence w is contained in η, as defined in (6.2), if and only if there exists $k \in \mathbb{Z}$ such that

$$x + k\alpha \pmod 1 \in I(w).$$

For irrational α, for any $x \in [0,1)$, we have that the sequence $(x + n\alpha \pmod 1)_{n\in\mathbb{Z}}$ is dense in $[0,1)$, meaning that w is contained in η if and only if $I(w) \neq \varnothing$. This implies that the set of words contained in η does not actually depend on the particular choice of $x \in [0,1)$. For a fixed irrational α, the collection of all such sequences η is a shift invariant subset of $\{0,1\}^{\mathbb{Z}}$. Taking the closure, we obtain a closed, shift invariant space $X(\alpha) \subset \{0,1\}^{\mathbb{Z}}$, called a *Sturmian shift*.

Moreover, for irrational α, we have that $X(\alpha)$ is minimal and the number of words of size n in $X(\alpha)$ is exactly $n+1$ (Exercises 6.9 and 6.10).

Combining this information, Hedlund and Morse showed that irrational rotations characterize all Sturmian sequences over an alphabet with two letters:

Theorem 6.10. *A sequence $\eta \in \{0,1\}^{\mathbb{Z}}$ is Sturmian if and only if there exists an irrational $\alpha \in (0,1)$ and $x \in [0,1)$ such that η is the coding of the orbit of x under rotation by α either with respect to the partition $\{[0,1-\alpha),[1-\alpha,1)\}$ or with respect to the partition $\{(0,1-\alpha],(1-\alpha,1]\}$.*

6.5. Higher dimensions

A natural generalization of studying sequences $\eta \in \mathcal{A}^{\mathbb{Z}}$ for some finite alphabet $\mathcal{A}$ is to consider higher dimensional configurations, meaning $\eta \in \mathcal{A}^{\mathbb{Z}^d}$ for some $d \geq 1$. Similarly, the notions of periodicity and complexity can be extended to higher dimensions, with some choices as to the way each notion is generalized.

For periodicity, one can have up to d directions of periodicity:

Definition 6.11. A configuration $\eta \in \mathcal{A}^{\mathbb{Z}^d}$ is *periodic* if there exists $\vec{m} \in \mathbb{Z}^d$, $\vec{m} \neq \vec{0}$, such that $\eta(\vec{m}+\vec{n}) = \eta(\vec{n})$ for all $\vec{n} \in \mathbb{Z}^d$. If there are d independent directions of periodicity, then η is said to be *fully periodic*.

Definition 6.12. The *rectangular complexity* $P_\eta(n_1,\ldots,n_d)$ of $\eta \in \mathcal{A}^{\mathbb{Z}^d}$ is defined to be the number of distinct patterns in η that fill an $n_1 \times \ldots \times n_d$ parallelepiped.

The *square complexity* $P_\eta(n,\ldots,n)$ of $\eta \in \mathcal{A}^{\mathbb{Z}^d}$ is defined to be the number of distinct patterns in an $n \times \ldots \times n$ cube. More generally, if S is any shape, we can define $P_\eta(S)$ to be the number of distinct patterns in η that fill a shape S.

Full periodicity does not give much information:

Proposition 6.13. *For $d \geq 1$, the configuration $\eta \in \mathcal{A}^{\mathbb{Z}^d}$ is fully periodic if and only if $P_\eta(n_1,\ldots,n_d)$ is bounded.*

Proof. If η is periodic, the result follows immediately by counting number of patterns in a period.

Conversely, assume that $P_\eta(n_1, \ldots, n_d) \leq N$ for all $n_1, \ldots, n_d \in \mathbb{N}$. Then each 1-dimensional complexity, determined by fixing all coordinates but the i^{th} one, satisfies

$$P_\eta(1, \ldots, 1, n_i, 1, \ldots, 1) \leq N.$$

By Corollary 6.6, each row is periodic with period $\leq N$. Thus $N!d$ is a period for η. $\square$

Instead of using rectangular complexity, one could prove the same result using square complexity. In particular, bounded complexity does not give much information.

Meurice Nivat conjectured that there is a relationship between periodicity and complexity in two dimensions:

Conjecture 6.14. *Let $\eta \in \mathcal{A}^{\mathbb{Z}^2}$. If there exist $n, k \in \mathbb{N}$ such that $P_\eta(n, k) \leq nk$, then η is periodic.*

This conjectured relation is the existence of a single direction of periodicity, and not necessarily double periodicity (Exercise 6.11). If the conjecture holds, it is optimal in some sense: for $d = 2$, there exist non-periodic configurations of complexity $nk + 1$, by setting $\eta_{0,0} = 1$ and $\eta_{n_1,n_2} = 0$ for $(n_1, n_2) \neq (0, 0)$.

It is significant that such behavior is a 2-dimensional phenomenon. In three dimensions, Robert Tijdeman gave the following example:

Example 6.15. Fix $k \neq 0$. Set

$$\begin{aligned} \eta_{i,0,0} &= 1 \text{ for all } i \in \mathbb{N} \\ \eta_{0,j,k} &= 1 \text{ for all } j, k \in \mathbb{N} \\ \eta_{i,j,k} &= 0 \text{ otherwise.} \end{aligned}$$

This example has 2 non-parallel, non-intersecting lines of 1's and 0's elsewhere, which is not possible unless there are at least 3 dimensions of freedom available.

Then one can check that for all $n \geq 3$, $P_\eta(n, n, \ldots, n) \leq n^3$, but η is not periodic.

More precisely,

$$P_\eta(n, n, \ldots, n) = \begin{cases} 2n^2 + 1 & \text{if } 2 \leq n \leq k \\ n^2(n - k) + 2nk & \text{if } n \geq k. \end{cases}$$

Unlike the Morse-Hedlund Theorem, the Nivat Conjecture is not an equivalence:

Example 6.16. There exist periodic configurations η in $d = 2$ with $P_\eta(n,k) > nk$ for all $n, k \in \mathbb{N}$, and even

$$P_\eta(n,k) = 2^{n+k-1} \text{ for all } n, k \geq 1.$$

One can construct such configurations by setting $\eta_{i,j} = u_{i+j}$, where u is the binary Champernowne word, meaning that one takes all words in lexicographic order and concatenates them. Then $P_k(u) = 2^n$ and η is periodic in a single direction.

There are some partial results towards the Nivat Conjecture, and the strongest result to date is:

Theorem 6.17. *If $\eta \in \mathcal{A}^{\mathbb{Z}^2}$ and there exist $n, k \in \mathbb{N}$ such that $P_\eta(n,k) \leq nk/2$, then η is periodic.*

Notes

Morse and Hedlund proved Corollaries 6.6 and 6.7 in [**43**]. The term Sturmian sequences was introduced in [**44**] and such sequences further studied by Coven and Hedlund [**16**].

Tijdeman [**53**] gave Example 6.15 in $d = 3$ and this was generalized to all dimensions by Sander and Tijdeman [**53**], who showed there exists aperiodic η with $P_\eta(n,\ldots,n) = 2n^{d-1}+1$. Berthé and Vuillon [**4**] and Cassaigne [**15**] constructed examples showing that Nivat is not an equivalence, such as the one in Example 6.16.

Nivat's Conjecture, under the stronger assumption that there exists $n \in \mathbb{N}$ with $P_\eta(2,n) \leq 2n$ (equivalently $P_\eta(n,2) \leq 2n$) was proven by Sander and Tijdeman [**55**]. If one assumes the existence of $n, k \in \mathbb{N}$ with $P_\eta(n,k) \leq nk/144$, periodicity of η was proven by Epifanio, Koskas and Mignosi [**18**] and periodicity was proven under the assumption of existence of $n, k \in \mathbb{N}$ with $P_\eta(n,k) \leq nk/16$ was proven by Quas and Zamboni [**47**]. Theorem 6.17 was proven in [**17**].

Exercises

Exercise 6.1. Give an example of an infinite word in a finite alphabet $\mathcal{A}$ that is recurrent but is not uniformly recurrent.

Exercise 6.2. The infinite word $\eta \in \mathcal{A}^{\mathbb{N}}$ is *ultimately periodic* if there exist integers $N, m \in \mathbb{N}$ such that for all $n \geq N$, $\eta(n+m) = \eta(n)$. If $\eta \in \mathcal{A}^{\mathbb{N}}$ and there exists $n_0 \in \mathbb{N}$ such that $P_\eta(n_0) \leq n_0$, show that η is ultimately periodic.

Exercise 6.3. Show that the system (X_η, T) defined in Section 6.3 is a topological dynamical system.

Exercise 6.4. Show that the cylinder sets of Definition 6.8 and that they form a basis for the open sets in the topology of X_η.

Exercise 6.5. Show that a word $\eta \in \mathcal{A}^{\mathbb{Z}}$ is recurrent if and only if the shift T in the associated dynamical system (X_η, T) is onto.

Exercise 6.6. Show that every Sturmian sequence is recurrent.

Exercise 6.7. For any $k \in \mathbb{N}$, show that there exists a sequence $\eta \in \mathcal{A}^{\mathbb{N}}$ such that $P_\eta(n) = n + k$ for all $n \in \mathbb{N}$.

Exercise 6.8. Show that the dynamical system associated to a Sturmian sequence $\eta \in \mathcal{A}^{\mathbb{Z}}$ is minimal.

Exercise 6.9. Show that for irrational α, the Sturmian system $X(\alpha)$ is minimal. (Hint: Show that the orbit closure $\overline{\mathcal{O}(\eta)}$ is the set of sequences obtained by coding every point on the unit circle.)

Exercise 6.10. Show that for irrational α, the sequence η defined in (6.2) is Sturmian. (Hint: check that for a word w of length n, the sets $I(w)$ are connected and are bounded by the end points $\{k(1-\alpha) \pmod 1 \colon k = 0, \dots, n-1\}$.)

Exercise 6.11. Show that there exists $\eta \in \mathcal{A}^{\mathbb{Z}^2}$ such that for some $n, k \in \mathbb{N}$, $P_\eta(n,k) \leq nk$ but η is not doubly periodic.

Chapter 7

Times p, Times q

7.1. Multiplicatively independent semigroups

We have already seen examples of subsets A of the integers such that for any irrational α, the subset $\{n\alpha \mod 1 \colon n \in A\}$ is dense in the interval $[0,1]$. Examples of such A include the set of all natural numbers (Corollary 4.23) and polynomial sequences (Theorem 5.12). In this chapter we produce other examples of such sequences.

Definition 7.1. Integers $p, q > 1$ are *multiplicatively independent* if they are not both rational powers of a single integer.

Thus relatively prime integers, such as 2 and 3, are multiplicatively independent, but so are some pairs of integers that are not relatively prime, such as 2 and 6. The integers 4 and 8 are not multiplicatively independent.

Given integers $p, q > 1$, by taking products of powers of these integers we obtain a semigroup S of natural numbers, meaning a subset of natural numbers that is closed under multiplication. We can write this subset as

$$S = \{p^m q^n \colon m, n \in \mathbb{Z}, m, n \geq 0\}.$$

(A semigroup differs from a group in that there is no requirement that it contains an identity or inverses; see Appendix B.1.)

Definition 7.2. A subset $S = \{s_1 < s_2 < \ldots\}$ of natural numbers is *nonlacunary* if

$$\lim_{j \to \infty} \frac{s_{j+1}}{s_j} = 1.$$

Lemma 7.3. *A semigroup in $\mathbb{N}$ that contains a pair of multiplicatively independent integers is nonlacunary.*

Proof. Assume that $p, q > 1$ are multiplicatively independent integers in the semigroup. Consider the semigroup

$$\{p^m q^n \colon m, n \in \mathbb{N} \cup \{0\}\}.$$

Enumerate the elements of this semigroup in increasing order and write them as $s_1 < s_2 < \dots$. We show that $\frac{s_{j+1}}{s_j} \to 1$ as $j \to \infty$.

Let $\varepsilon > 0$. Since $\log p / \log q$ is irrational (Exercise 7.1), there exist $m, n \in \mathbb{N}$ such that

$$0 < m - n\frac{\log p}{\log q} < \frac{\log(1+\varepsilon)}{\log q}.$$

Multiplying by $\log q$ and exponentiating, we have that

$$1 < \frac{q^m}{p^n} < 1 + \varepsilon.$$

Similarly, there exist $m', n' \in \mathbb{N}$ such that

$$1 < \frac{p^{m'}}{q^{n'}} < 1 + \varepsilon.$$

Let $N \in \mathbb{N}$ be chosen such that $s_N = p^n q^{n'}$. We claim that for any $j > N$,

$$1 < \frac{s_{j+1}}{s_j} < 1 + \varepsilon.$$

For $j > N$, writing $s_j = p^a q^b$, then either $a > n$ or $b > n'$. If $a > n$, then $s_j = p^n p^{a-n} q^b$. Setting $x = p^{a-n} q^{m+b}$, we have that

$$1 < \frac{x}{s_j} = \frac{q^m}{p^n} < 1 + \varepsilon.$$

Thus $x \geq s_{j+1}$ and

$$1 < \frac{s_{j+1}}{s_j} \leq \frac{x}{s_j} = \frac{q^m}{p^n} < 1 + \varepsilon.$$

Similarly, if $b > n'$, the same inequality holds. Since ε was arbitrary, the limit of the ratios is 1 and we have the statement of the lemma. $\square$

Furstenberg showed that the orbit of any irrational under a nonlacunary semigroup is dense in the circle:

Theorem 7.4. *If $S \subset \mathbb{N}$ is a nonlacunary semigroup and $\alpha \in \mathbb{R}$ is an irrational, then*

$$\{s\alpha \mod 1 \colon s \in S\}$$

is dense in $[0, 1]$.

This immediately gives a Diophantine result, as one can use rationals with denominators from the nonlacunary semigroup to approximate any point in the circle arbitrarily well. (An analogous equivalence is stated in Exercise 3.9.) However, this theorem differs from the other Diophantine results we have proven thus far, such as Corollary 4.23 or Theorem 5.12. The previous results can all be strengthened to prove a stronger result, showing that not only does one have density of the iterates, but they are spread out uniformly (called equidistributed) throughout the circle. This means that the probability of finding an iterate in a particular interval is proportional to the length of this interval. The dynamical methods needed to prove these theorems rely on measures and ergodic theory, which is beyond the scope of this book. On the other hand, there exist irrationals whose orbit is not equidistributed in the circle under the semigroup generated by a pair of multiplicatively independent integers and so Theorem 7.4 cannot be strengthened in this way.

7.2. Proof of Furstenberg's Diophantine Theorem

Given an integer $p > 1$, we say that $A \subset \mathbb{T}$ is *p-invariant* if it is invariant under multiplication by p. Thus,

$$pa \mod 1 \in A \text{ for all } a \in A.$$

In order to prove Theorem 7.4, it suffices to show:

Theorem 7.5. *Let $p, q > 1$ be multiplicatively independent integers. Assume that $A \subset \mathbb{T}$ is closed, infinite and p and q-invariant. Then $A = \mathbb{T}$.*

Theorem 7.4 follows, since the closure of an irrational point under a nonlacunary semigroup satisfies these hypotheses, and so by Theorem 7.5 must be all of $\mathbb{T}$. In fact, Theorem 7.4 is equivalent to Theorem 7.5 (Exercise 7.2).

For $x \in \mathbb{T}$, we let $\|x\|$ denotes the distance between x and 0 taken in the circle. Thus $\|x\| \leq 1/2$ for all $x \in \mathbb{T}$.

We start with by proving Theorem 7.5 under the additional hypothesis that 0 is a limit point of the subset:

Proposition 7.6. *Assume that $A \subset \mathbb{T}$ is closed, infinite, p and q-invariant and assume that 0 is a limit point of A. Then $A = \mathbb{T}$.*

Proof. Assume that $A \subset \mathbb{T}$ satisfies the hypotheses of the proposition. Let S be the semigroup generated by p and q and order the elements of S by size. Thus we can write $S = \{s_1 < s_2 < \ldots\}$.

Fix $\varepsilon > 0$. By Lemma 7.3, we can choose N so that for all $n \geq N$,

$$\frac{s_{n+1}}{s_n} < 1 + \varepsilon.$$

Since 0 is a limit point of A, we can pick $a \in A$ with $a \neq 0$ such that $\|a\| < \varepsilon / s_N$. Consider the (finite) set:

$$\{sa \mod 1 \colon s \in S, s_N \leq s \leq \frac{1}{\|a\|}\}.$$

This is a subset of A, since A is p and q-invariant. Furthermore, we claim that there is an element of this subset in every interval of size ε in $\mathbb{T}$: if $n > N$, then

$$\begin{aligned} \|(s_{n+1} - s_n)a\| &= \left\| \left(\frac{s_{n+1}}{s_n} - 1\right) a \right\| \\ &< \|\varepsilon s_n a\| < \varepsilon, \end{aligned}$$

since $s_n \leq 1/\|a\|$. This means that for $n > N$, the differences $(s_{n+1} - s_n)a$ are bounded by ε and for $s \geq 1/\|a\|$, these iterates wrap around the circle. It follows that A has nontrivial intersection with every interval of size ε. Since ε is arbitrary, $A = \mathbb{T}$. □

The next step is to generalize this proposition for rational limit points. We begin with a lemma.

Lemma 7.7. *If $p, q, r \in \mathbb{N}$ are pairwise relatively prime, then there exists $t \in \mathbb{N}$ such that $p^t \equiv q^t \equiv 1 \mod r$.*

Proof. Consider the powers $p, p^2, p^3, \ldots$, taken modulo r. (Note that none are $0 \mod r$, since p and r are relatively prime.) By the Pigeonhole Principle, for some $i \neq j$, we have that $p^i \equiv p^j \mod r$. Without loss, we can assume that $j > i$. Then $p^u \equiv 1 \mod r$, where $u = j - i$. Similarly, there exists $v \in \mathbb{N}$ such that $q^v \equiv 1 \mod r$. Set $t = uv$. Then

$$p^t = (p^u)^v \equiv 1 \mod r$$

and

$$q^t = (q^v)^u \equiv 1 \mod r.$$

□

We use this lemma to reduce the case of a rational limit point to that of zero being a limit point:

Proposition 7.8. *Assume that $A \subseteq \mathbb{T}$ is closed, infinite, p and q-invariant and assume that A has a rational limit point. Then $A = \mathbb{T}$.*

Proof. Assume that m/r is a rational limit point of A. Without loss, we can assume that m and r are relatively prime. By Lemma 7.7, there exists $t \in \mathbb{N}$ such that $p^t \equiv q^t \equiv 1 \mod r$. Then p^t and q^t still generate a nonlacunary semigroup. Consider the set

$$A' = A - m/r = \{a - m/r \colon a \in A\}.$$

This set is infinite, closed, and invariant under multiplication by p^t and q^t, since A is. Moreover, since m/r is a limit point of A, we have that 0 is a limit point of A'. Thus we can apply Proposition 7.6 to obtain that $A' = \mathbb{T}$. Therefore $A = \mathbb{T}$. $\square$

By Proposition 7.8, the proof of Theorem 7.5 has been reduced to the following: any closed, infinite subgroup of the circle that is invariant under multiplication by p and q contains a rational limit point. We now complete this argument:

Proof. (of Theorem 7.5) Let A' denote the set of limit points of A. Then A' is also nonempty, closed, and p and q-invariant. If A' contains a rational point, then we are finished. So assume that A' consists only of irrational points. Thus, we can assume that A' is a closed, p and q-invariant set of irrationals in $\mathbb{T}$. Clearly it is infinite, since if A' contains one irrational point and is p-invariant, then it contains infinitely many (irrational) points.

Consider the difference set

$$A' - A' = \{a_1 - a_2 \colon a_1, a_2 \in A'\}.$$

Then $0 \in A' - A'$ is a limit point of this difference, since A' is infinite and so accumulates somewhere in $\mathbb{T}$. Since A' is closed and p and q-invariant, so is $A' - A'$. Thus by Proposition 7.8, $A' - A' = \mathbb{T}$. In particular, A' is uncountable.

Define

$$P = \{x \in A' \colon \text{ every neighborhood of } x \text{ contains uncountably many points of } A'\}.$$

Since multiplication by p and q is continuous, P is p and q-invariant. Set $C = A' \setminus P$ (the set theoretic difference). Let $\mathcal{B}$ be a countable basis of open sets for the topology of the circle. If $x \in C$, letting U be a neighborhood of x that contains only countably many points of A', we have some basis element $B \in \mathcal{B}$ such that $x \in B \subseteq U$. Thus B contains countably many points of A'. Using these basis elements, we can cover C. Since there are countably many choices for the basis elements, it follows that C itself is countable. Furthermore C is p and q-invariant, since both A' and P are.

If $x \in \overline{P}$ and U is any neighborhood of x, then U contains a point of P. It follows that U contains uncountably many points of A'. Therefore $x \in P$ and so P is closed. If $x \in P$ and U is any neighborhood of x, then

U contains uncountably many points in A', of which only countably many are in C. Thus U contains at least 2 points of P, meaning that any point in P is a limit point of P. Thus P is perfect.

Define

$$P_k = \{(x_2 - x_1, x_3 - x_2, \dots, x_{k+1} - x_k) : x_j \in P \text{ for } j = 1, 2, \dots, k\}. \tag{7.1}$$

Claim: $P_k = \mathbb{T}^k$.

We proceed by induction. For $k = 1$, note that P is a closed, infinite p and q-invariant set. (More holds: since P is perfect, it must be uncountable.) Therefore $P - P$ is closed, p and q-invariant, and has 0 as a limit point. By Proposition 7.6, $P - P = \mathbb{T}$.

Assume that for some $k \geq 2$, $P_{k-1} = \mathbb{T}^{k-1}$ and consider a point $y = (y_1, y_2, \dots, y_k) \in P_k$. The first $k - 1$ coordinates of y can be written in the form given in (7.1), and by the inductive assumption we can take $y_1, y_2, \dots, y_{k-1}$ to be any points in the torus. In particular, choose these $k - 1$ coordinates to be rationals m/r, where $m, r \in \mathbb{N}$. By Lemma 7.7, for some $t \in \mathbb{N}$, we have that

$$p^t \equiv q^t \equiv 1 \mod r.$$

This determines some choice (not necessarily a unique choice) of the points $x_1, x_2, \dots, x_k \in P$, satisfying

$$y_1 = x_2 - x_1, y_2 = x_3 - x_2, \dots, y_{k-1} = x_k - x_{k-1}.$$

However, x_{k+1} is not yet determined and can still be chosen freely in P. Since P is perfect, we can choose $x_{k+1} \in P$ to approximate x_k arbitrarily well. This means that we can choose a sequence of points in P approaching x_k. For each choice of x_{k+1} in this sequence, the corresponding vector of differences $y = (y_1, y_2, \dots, y_k)$ still lies in P_k. Each of these points has the same first $k - 1$ coordinates. This means that the k-th coordinate of y has 0 as a limit point. Furthermore, P_k is p^t and q^t-invariant and so the choices for the last coordinate are also p^t and q^t-invariant. It follows from Proposition 7.8 that we can have any y_k in the last coordinate of y.

We can carry out this procedure for any choice of $k - 1$ rationals $m_1/r, m_2/r, \dots, m_{k-1}/r$ in the first $k - 1$ coordinates of the vector y. Since r can be taken arbitrarily large, rationals of this form are dense in $\mathbb{T}$. By construction, the first $k - 1$ coordinates of y are unchanged by multiplication by p^t and by q^t for some $t \in \mathbb{N}$. Thus P_k is closed and dense and so $P_k = \mathbb{T}^k$, proving the claim.

In particular, for each $k \in \mathbb{N}$, the point $(1/k, 1/k, \dots, 1/k) \in P_k$. Taking a choice of $x_1, x_2, \dots, x_{k+1}$ in (7.1) that gives rise to this point, the points $x_1, x_2, \dots, x_k$ are determined by $x_j = x_{j-1} + 1/k$ for $j = 2, \dots, k$, with each $x_j \in P$. Thus P has nontrivial intersection with any interval of length $1/k$ in $\mathbb{T}$. Since k is arbitrary, P is dense in $\mathbb{T}$. $\square$

7.3. Generalizations

Based on Furstenberg's Theorem, one can construct other dense subsets of the circle:

Theorem 7.9. *Let $p, q > 1$ be multiplicatively independent integers and let τ_m be any sequence of real numbers. Then for any irrational α*

$$\{p^n q^m \alpha + \tau_m : n, m \in \mathbb{N}\}$$

is dense in $\mathbb{T}$.

Before proving the theorem, we define a distance between subsets of the circle:

Definition 7.10. If $A, B \subseteq \mathbb{T}$, then the *Hausdorff distance* $d(A, B)$ between A and B is defined by

$$d(A, B) = \inf\{|a - b| : a \in A, b \in B\}.$$

The *Hausdroff metric* on $\mathbb{T}$ is the metric induced by the Hausdorff distance d.

Using this notion of distance between subsets of the circle and the associated notion of a limit, Theorem 7.9 is a corollary of the following lemma:

Lemma 7.11. *Let p, q be multiplicatively independent, integers, $\varepsilon > 0$, and assume that A is an infinite p-invariant subset of $\mathbb{T}$. There exists $n \in \mathbb{N}$ such that $q^n A$ has nontrivial intersection with any interval of size ϵ in $\mathbb{T}$.*

Proof. Without loss, we can assume that A is closed (by replacing A with its closure). Let

$$\mathcal{X} = \overline{\{q^n A : n \in \mathbb{N}\}}.$$

Since A is p-invariant, so is each $X \in \mathcal{X}$. Let $B = \bigcup_{X \in \mathcal{X}} X$. Then B is infinite, since it contains A; B is closed in $\mathbb{T}$, since $\mathcal{X}$ is closed in the Hausdorff topology; B is both p and q-invariant, since each $X \in \mathcal{X}$ is. By Theorem 7.5, $B = \mathbb{T}$. In particular, there exists $\alpha \in B$ such that the closure $\overline{\{p^n \alpha : n \in \mathbb{N}\}} = \mathbb{T}$. By definition, $\alpha \in X$ for some $X = \lim_j q^{n_j} A$, where the limit is taken along some sequence $n_j \to \infty$. Since X is p-invariant,

$$X \supseteq \overline{\{p^n \alpha : n \in \mathbb{N}\}}$$

and so $X = \mathbb{T}$. Thus along the sequence n_j, for sufficiently large n_j, $q^{n_j} A$ has nontrivial intersection with every interval of size ε in $\mathbb{T}$. $\square$

of Theorem ??. Fix $\varepsilon > 0$ and let

$$A = \{p^n\alpha : n \in \mathbb{N}\}.$$

By Lemma 7.11, there exists $m \in \mathbb{N}$ such that $q^m A$ has nontrivial intersection with every interval of size ε in $\mathbb{T}$. Since $q^m A + \tau_m$ is a translate of this set, it too has nontrivial intersection with every interval of size ε in $\mathbb{T}$. □

Theorem 7.12. *Let $k \in \mathbb{N}$. Let $p_i, q_i \in \mathbb{N}$ with $1 < p_i < q_i$ are pairs of multiplicatively independent integers for $i = 1, \ldots, k$ and assume that $p_1 \le p_2 \le \ldots \le p_k$. Then for distinct $\alpha_1, \ldots, \alpha_k \in \mathbb{T}$ with at least one $\alpha_i \notin \mathbb{Q}$,*

$$\Big\{\sum_{i=1}^{k} p_i^n q_i^m \alpha_i : n, m \in \mathbb{N}\Big\}$$

is dense in $\mathbb{T}$.

The proof of this theorem is left to a series of exercises.

There are simple to state (but not so simple to answer...) open questions related to the dynamics of nonlacunary subsets of integers. Assume that $p, q > 1$ are multiplicatively independent integers. It follows from Furtsenberg's Theorem that if A is a closed, infinite, p-invariant subset of the circle, then there exists a sequence $n_j \to \infty$ such that $q^{n_j} A \to \mathbb{T}$ in the Hausdorff metric. It is unknown if one needs to pass to a subsequence:

Question 7.13. *Assume that $p, q > 1$ are multiplicatively independent integers and that $A \subset \mathbb{T}$ is closed, infinite, and p-invariant. Let $\varepsilon > 0$. Does there exist $N \in \mathbb{N}$ such that for all $n > N$, $q^n A$ is ε dense?*

One can also try and generalize the results for nonlacunary subgroups for higher dimensions:

Question 7.14. *If $p, q > 1$ are relatively prime integers and α, β are irrationals that are independent over $\mathbb{Q}$, then is*

$$\{(p^n q^m \alpha, p^n q^m \beta) : m, n \in \mathbb{N}\}$$

dense in $\mathbb{T} \times \mathbb{T}$?

This question is phrased a bit differently from Theorem 7.4, requiring that p and q be relatively prime and not merely multiplicatively independent. It is possible to construct examples of relatively independent integers such that the corresponding conclusion of density does not hold.

Notes

Furstenberg's original paper [**24**] includes a proof of Theorem 7.4 as a corollary of the development of an important new notion in topolgical dynamics (and in ergodic theory) called disjointness. Since then, there have been several elementary proofs, such as one by Michael Boshernitzan [**6**]. The argument given here is a variant of an argument of Daniel Rudolph (unpublished). For the statements and proofs of the stronger equidistribution results alluded to after the statement of Theorem 7.4, see [**27**].

Theorems 7.9 and 7.12 appear in Kra [**41**]. The examples alluded to after Question 7.14 were constructed by Daniel Berend (unpublished). A hint on how to construct such examples is given in Exercise 7.9.

Exercises

Exercise 7.1. Show that $p, q > 1$ are multiplicatively independent integers if and only if $\log p / \log q \notin \mathbb{Q}$.

Exercise 7.2. Show that Theorem 7.4 implies Theorem 7.5.

Exercise 7.3. Let $p \in \mathbb{N}$ with $p > 1$. Show that exists an irrational $\alpha \in [0,1]$ such that $\{p^n\alpha \mod 1 \colon n \in \mathbb{N}\}$ is not dense. Conclude that not all irrational numbers can be approximated arbitrarily well by rationals with denominator of the form p^n for some $n \in \mathbb{N}$.

For Exercises 7.4–7.8, assume that $p_i, q_i > 1$ are pairs of multiplicatively independent integers and let $M_1 = \begin{pmatrix} p_1 & 0 \\ 0 & p_2 \end{pmatrix}$ and $M_2 = \begin{pmatrix} q_1 & 0 \\ 0 & q_2 \end{pmatrix}$.

Exercise 7.4. Assume that $A \subset \mathbb{T}^2$ is nonempty, closed, and invariant under multiplication by M_1 and M_2. Show that if all points of $A \cap \mathbb{Q}^2$ are isolated in A, then A is finite. Hint: the possibilities for the fibers $A_x = \{t \in \mathbb{T} \colon (t,x) \in A\}$.

Exercise 7.5. Assume that A is a closed subset of $\mathbb{T}^2$ that is invariant under M_1 and M_2 and assume that $(r,s) \in A \cap \mathbb{Q}^2$. Show that there exist $n, m \in \mathbb{N}$ such that $A - (r,s)$ is invariant under M_1^n and M_2^m. Hint: compare this result with Lemma 7.7.

Exercise 7.6. Let S be the set of accumulation points of the set

$$\big\{(p_1^n q_1^m \alpha_1, p_2^n q_2^m \alpha_2) \in \mathbb{T}^2 \colon n, m \in \mathbb{N}\big\}.$$

Show that if $(0,0) \in X$, then one of the following holds:

(1) $(0,0)$ is isolated in S.

(2) S contains the whole x-axis or the whole y-axis.

(3) For some $c > 0$, S contains the curve $y = cx^\rho$, for $x > 0$, where $\rho = \log p_2 / \log p_1 = \log q_2 / \log q_1$.

Conclude that either $(0,0)$ is isolated in S or that $\{x + y \colon (x,y) \in S\} = \mathbb{T}$.

Exercise 7.7.

Exercise 7.8. Combine the preceeding exercises to complete the proof of Theorem 7.12.

Exercise 7.9. Construct multiplicatively independent integers $p, q > 1$ and irrationals α and β that are independent over $\mathbb{Q}$ such that the orbit closure

$$\overline{\{(p^n q^m \alpha, p^n q^m \beta) \colon m, n \in \mathbb{N}\}}$$

consists of exactly two lines. Hint: consider $p = 10$ and $q = 20$. Construct α and β as sums $\sum_i 1/10^{n_i}$ and $\sum_i 1/20^{m_i}$, where the n_i and m_i are sequences that grow quickly with $n_1 < m_1 < n_2 < m_2 < \ldots$.

Chapter 8

Van der Waerden's Theorem

8.1. Translating van der Waerden's Theorem to dynamics

Our goal in this chapter is to prove Van der Waerden's Theorem (Theorem 1.4):

Theorem 8.1. *If $\mathbb{N} = C_1 \cup C_2 \cup \ldots \cup C_r$ is a finite partition, then some C_j, $j \in \{1, 2, \ldots, r\}$, contains arbitrarily long arithmetic progressions.*

We start with some definitions designed to formalize the link between recurrence and coloring theorems.

Definition 8.2. Define $P \subseteq \mathbb{N}^k$ to be a *van der Waerden collection* if given any finite partition $\mathbb{N} = C_1 \cup C_2 \cup \ldots \cup C_r$, there exist $(s_1, s_2, \ldots, s_k) \in P$ and $j \in \{1, 2, \ldots, r\}$ such that $\{s_1, s_2, \ldots, s_k\} \subset C_j$.

More generally, this definition can be made in any semigroup, not just in $\mathbb{N}^k$. Since our interest lies in finding arithmetic progressions in the integers, we focus on the group $\mathbb{N}$.

In our terminology, van der Waerden's Theorem becomes:

Theorem 8.3. *In $\mathbb{N}$, the arithmetic progressions of length k form a van der Waerden collection.*

The version of van der Waerden's Theorem stated in Theorem 8.1 follows immediately: by the pigeonhole principle, some piece of any finite partition contains arithmetic progressions of length k for infinitely many k, and therefore for arbitrary k.

Definition 8.4. Define $P \subseteq \mathbb{N}^k$ to be a *Birkhoff collection* if for any dynamical system (X, T), $x \in X$ and $\varepsilon > 0$, there exists $(s_1, s_2, \ldots, s_k) \in P$ such that $T^{s_1}x, T^{s_2}x, \ldots, T^{s_k}x$ all lie within ε of each other.

In this definition, we have implicitly assumed that (X, T) is metrizable. This notion can also be formulated in purely topological terms: if U is a nonempty open set, there exists $(s_1, s_2, \ldots, s_k) \in P$ such that

$$T^{s_1}U \cap T^{s_2}U \cap \ldots \cap T^{s_k}U \neq \varnothing.$$

Although this is slightly more general, in order to clarify the ideas in the proof we assume metrizability and leave general reformulations to the exercises.

In spite of the seemingly different formulations, the definitions of a Birkhoff collection and of a van der Waerden collection are equivalent. This makes it possible to translate a combinatorial problem, such as van der Waerden's Theorem, into a dynamical question:

Theorem 8.5. *The subset $P \subseteq \mathbb{N}^k$ is a Birkhoff collection if and only if P is a van der Waerden collection.*

Proof. Assume that P is a Birkhoff collection and that $\mathbb{N} = C_1 \cup C_2 \cup \ldots \cup C_r$ is a finite partition. Let $X = \{1, 2, \ldots, r\}^{\mathbb{N}}$ and let $T\colon X \to X$ be the shift map. Choose a metric d on X such that $d(x, y) < 1$ if and only if $x_0 = y_0$. Define the point $x \in X$ by $x_n = i$ if $n \in C_i$. By assumption, there exists $(s_1, s_2, \ldots, s_k) \in P$ such that $T^{s_1}x, T^{s_2}x, \ldots, T^{s_k}x$ all lie within ε of each other. By choice of the metric, we have that

$$(T^{s_1}x)_0 = (T^{s_2}x)_0 = \ldots = (T^{s_k}x)_0$$

and so $x_{s_1} = x_{s_2} = \ldots = x_{s_k}$. But this means that $(s_1, s_2, \ldots, s_k) \subset C_j$ for some $j \in \{1, 2, \ldots, r\}$.

Conversely, assume that P is a van der Waerden collection. Let (X, T) be a dynamical system, $x \in X$ and $\varepsilon > 0$. Since X is compact, given $\varepsilon > 0$, we can partition X into a finite number of pieces such that the diameter of each piece is bounded by ε. Thus we can write $X = Y_1 \cup Y_2 \cup \ldots \cup Y_r$ so that the diameter of each Y_j, $j = 1, 2, \ldots, r$ is smaller than ε. Define a partition of $\mathbb{N}$ by setting $n \in C_j$ if $T^n x \in Y_j$. Since P is a van der Waerden collection, there exists $(s_1, s_2, \ldots, s_k) \in P$ with $\{s_1, s_2, \ldots, s_k\} \subset C_j$ for some $j \in \{1, 2, \ldots, r\}$. Thus $T^{s_1}x, T^{s_2}x, \ldots, T^{s_k}x$ all lie in the same Y_j and so all lie within ε of each other. $\square$

8.2. Dynamical van der Waerden Theorem

Using the correspondence of Theorem 8.5, van der Waerden's Theorem follows from the following "dynamical" van der Waerden Theorem:

Theorem 8.6. *Assume that* (X, T) *is a dynamical system,* $k \in \mathbb{N}$*,* $x \in X$ *and* $\varepsilon > 0$*. There exist* $n, m \in \mathbb{N}$ *with* $n \geq 1$ *such that*

$$T^m x, T^{n+m} x, T^{2n+m} x, \ldots, T^{kn+m} x$$

all lie within ε *of each other.*

The proof follows from a generalization of Birkhoff's Recurrence Theorem. Recall that Birkhoff's Theorem states that any dynamical system (X, T) contains a recurrent point, meaning that there exists $x \in X$ such that for all $\varepsilon > 0$, there exists $n \in \mathbb{N}$ with $d(T^n x, x) < \varepsilon$. Our interest lies in the existence of a multiply recurrent point:

Definition 8.7. If (X, T) is a dynamical system, the point $x \in X$ is *k-multiply recurrent* if for all $\varepsilon > 0$, there exists $n \in \mathbb{N}$ satisfying $d(T^{jn} x, x) < \varepsilon$ for $j = 1, 2, \ldots, k$. If $x \in X$ is k-multiply recurrent for all $k \in \mathbb{N}$, then we says that x is *multiply recurrent.*

The proof of Theorem 8.6 is carried out in several steps. First we show (Lemma 8.9) that in a minimal system, so long as some point is multiply recurrent, then we have a dense set of points with strong recurrence properties. We use this to show (Theorem 8.10) that any system contains some point with these recurrence properties. Finally we use such points to complete the proof of Theorem 8.6.

Notation 8.8. Let $B(x; \varepsilon)$ denote the ball around x of radius ε, meaning that

$$B(x; \varepsilon) = \{y \in X : d(y, x) < \varepsilon\}.$$

Lemma 8.9. *Let* (X, T) *be a minimal dynamical system and let* $k \in \mathbb{N}$*. Assume that for each* $\varepsilon > 0$*, there exist* $x \in X$ *and* $n \in \mathbb{N}$ *such that* $d(T^{jn} x, x) < \varepsilon$ *for* $j = 1, 2, \ldots, k$*. Then there is a dense set* $Y \subseteq X$ *such that for each* $y \in Y$*, there exists* $n \in \mathbb{N}$ *with* $d(T^{jn} y, y) < \varepsilon$ *for* $j = 1, 2, \ldots, k$*.*

Proof. Let $U \subseteq X$ be a nonempty open set and let $B_0 \subset U$ be a ball of radius ε. By Proposition 4.9, there exists $J \in \mathbb{N}$ such that

$$X = \bigcup_{j=1}^{J} T^{-j} B_0.$$

By Lemma F.14, there exists $\delta > 0$ (the Lebesgue number of the covering) such that any set of diameter bounded by δ is contained in some $T^{-j} B_0$ where $1 \leq j \leq J$.

By hypothesis, there exist $x \in X$ and $n \in \mathbb{N}$ such that $d(T^{jn}x, x) < \delta$ for $j = 1, 2, \ldots, k$. Thus

$$T^n x, T^{2n}x, \ldots, T^{kn}x \in B(x; \delta).$$

The choice of δ implies that $B(x; \delta) \subset T^{-j}(B_0)$ for some $j \in \{1, \ldots, J\}$. Then $T^j(B(x; \delta)) \subset B_0$ and so setting $y = T^j x$, we have

$$y, T^n y, T^{2n}y, \ldots, T^{kn}y \in B(y; \varepsilon).$$

Thus we have produced $y \in U$ and $n \in \mathbb{N}$ such that $d(T^{jn}y, y) < \varepsilon$ for $j = 1, 2, \ldots, k$. Since U is an arbitrary nonempty open set, the set of points satisfying this property is dense. $\square$

We use this to show that any dynamical system contains a multiply recurrent point:

Theorem 8.10. *If (X, T) is a dynamical system, $k \in \mathbb{N}$, and $\varepsilon > 0$, then there exist $x \in X$ and $n \in \mathbb{N}$ such that $d(T^{jn}x, x) < \varepsilon$ for $j = 1, 2, \ldots, k$.*

Proof. Since any dynamical system contains a minimal set, by replacing X by this minimal system, we can assume that X is minimal.

We proceed by induction on k. For $k = 1$, the existence of a recurrent point follows from Birkhoff's Recurrence Theorem. (In fact this holds for all points. Minimality implies that $\overline{\mathcal{O}_T^+(x)} = X$ for any $x \in X$ and so in particular, for any $\varepsilon > 0$ there exists $n \in \mathbb{N}$ such that $d(T^n x, x) < \varepsilon$. Thus every point is 1-recurrent.)

Assume that the statement holds for some $k \geq 1$, meaning that for any $\varepsilon > 0$, there exist $x \in X$ and $n \in \mathbb{N}$ such that $d(T^{jn}x, x) < \varepsilon$ for $j = 1, 2, \ldots, k$. By Lemma 8.9, for each $\varepsilon > 0$, there is a dense set $Y \subseteq X$ such that for all $y \in Y$, there exists $n \in \mathbb{N}$ with $d(T^{jn}y, y) < \varepsilon$ for $j = 1, 2, \ldots, k$. We show that the same conclusion holds for $k + 1$.

Fixing $\varepsilon > 0$, we can choose $x_0 \in X$ and an integer $n_0 \in \mathbb{N}$ such that $d(T^{jn_0}x_0, x_0) < \varepsilon/2$ for $j = 0, 1, 2, \ldots, k$. A minimal transformation is surjective and so we can choose $x_1 \in X$ such that $T^{n_0}x_1 = x_0$. Then for $j = 1, 2, \ldots, k$,

$$d(T^{(j+1)n_0}x_1, x_0) = d(T^{jn_0}T^{n_0}x_1, x_0) = d(T^{jn_0}x_0, x_0) < \varepsilon/2.$$

This means that for $j = 1, 2, \ldots, k + 1$,

$$d(T^{jn_0}x_1, x_0) < \varepsilon/2.$$

Since T is continuous, the same conclusion holds in some neighborhood of x_1, Thus we can choose ε_1 with $0 < \varepsilon_1 < \varepsilon$ such that $d(T^{jn_0}y, x_0) < \varepsilon/2$ for $j = 1, 2, \ldots, k + 1$ and for all $y \in B(x_1; \varepsilon_1)$. By the inductive assumption, there exist a point $y_1 \in B(x_1; \varepsilon_1/2)$ and $n_1 \in \mathbb{N}$ such that

$d(T^{jn_1}y_1, y_1) < \varepsilon_1/2$ for $j = 1, 2, \ldots, k$. This means that y_1 and $T^{jn_1}y_1$, for $j = 1, 2, \ldots, k$, lie in $B(x_1, \varepsilon_1)$. Thus for $j = 1, 2, \ldots, k+1$,

$$d\big(T^{jn_0}(T^{(j-1)n_1}y_1), x_0\big) < \varepsilon/2.$$

Taking any point $x_2 \in X$ such that $T^{n_1}x_2 = y_1$, we have

$$d(T^{jn_1}x_2, x_1) < \varepsilon_1/2 < \varepsilon/2$$

for $j = 1, 2, \ldots, k+1$, as well as:

$$d(T^{j(n_1+n_0)}x_2, x_0) < \varepsilon/2.$$

Inductively, we find $x_0, x_1, x_2, \ldots \in X$ and $n_1, n_2, n_3, \ldots \in \mathbb{N}$ such that for any $i \in \mathbb{N}$ and for $j = 1, 2, \ldots, k+1$,

$$d(T^{jn_{i-1}}x_i, x_{i-1}) < \varepsilon/2,$$

$$d(T^{j(n_{i-1}+n_{i-2})}x_i, x_{i-2}) < \varepsilon/2,$$

and

$$d(T^{j(n_{i-1}+\ldots+n_0)}x_i, x_0) < \varepsilon/2.$$

By compactness of X, there exist integers $0 < m < l$ such that $d(x_l, x_m) < \varepsilon/2$. Thus

$$d(T^{j(n_{l-1}+\ldots+n_m)}x_l, x_l) \le d(x_l, x_m) + d(x_m, T^{j(n_{l-1}+\ldots+n_m)}x_l) < \varepsilon$$

for $j = 1, 2, \ldots, k+1$. Taking $x = x_l$ and $n = n_{l-1} + \ldots + n_m$, we have produced a point $x \in X$ such that $d(T^{jn}x, x) < \varepsilon$ for $j = 1, 2, \ldots, k+1$. □

Several corollaries are immediate:

Corollary 8.11. *If (X, T) is a minimal dynamical system, $k \in \mathbb{N}$, and $\varepsilon > 0$, then there exists a dense set $X_0 \subseteq X$ such that for each $x \in X_0$, there exists $n \in \mathbb{N}$ with $d(T^{jn}x, x) < \varepsilon$ for $j = 1, 2, \ldots, k$.*

Proof. By Theorem 8.10, some point satisfies the conclusion and so by Lemma 8.9, we have a dense set points satisfying the conclusion. □

One can strengthen this to show that in any minimal system, one can find a dense set of points with this recurrence property, meaning a dense set of k-multiply recurrent points (Exercise 8.9).

Corollary 8.12. *Assume that (X, T) is a dynamical system, $k \in \mathbb{N}$, and $\varepsilon > 0$. If $X_0 \subseteq X$ is dense, then for some $x \in X_0$ and $n \in \mathbb{N}$, $d(T^{jn}x, x) < \varepsilon$ for $j = 1, 2, \ldots, k$.*

Proof. By Theorem 8.10, there exist $x \in X$ and $n \in \mathbb{N}$ with $d(T^{jn}x, x) < \varepsilon$ for $j = 1, 2, \ldots, k$. By continuity of T, any point in a sufficiently small neighborhood of X satisfies the same conclusion. □

Lastly, we use Theorem 8.10 to complete the proof of Theorem 8.6:

Proof. (of Theorem 8.6) Fix $k \in \mathbb{N}$, $x \in X$, and $\varepsilon > 0$ and let $Y = \overline{\mathcal{O}_T^+(x)}$. By Theorem 8.10, there exist $y \in Y$ and $n \in \mathbb{N}$ such that

$$y, T^n y, T^{2n} y, \ldots, T^{kn} y$$

all lie within ε of each other. Since y lies in the orbit closure of x, choosing m such that $T^m x$ stays close to y for kn iterates, we have the same conclusion for $T^m x$. Thus

$$T^m x, T^{n+m} x, T^{2n+m} x, \ldots, T^{kn+m} x$$

all lie within ε of each other. □

There are numerous generalizations of this result. In Chapter 9, we give an alternate proof of Theorem 8.6 as a corollary of a stronger multiple recurrence theorem, proving the recurrence theorem for a commuting set of transformations, and we also show that one can also place restrictions on the return times.

Notes

Solving a conjecture of Schur, van der Waerden proved Theorem 8.1 in [**61**] via combinatorial methods. The translation to the dynamical formulation of Theorem 8.6 and its proof was given by Furstenberg and Weiss in [**30**].

Exercises

Exercise 8.1. State the finite version of van der Waerden's Theorem (see Chapter 1.2 for the meaning). Show that it is equivalent to van der Waerden's Theorem.

Exercise 8.2. Show that the conclusion of van der Waerden's Theorem does not hold for infinite length progressions, meaning show that there exists a finite partition of $\mathbb{N}$ such that no piece contains an infinite length arithmetic progression.

Exercise 8.3. Show that van der Waerden's Theorem is equivalent to the following statement: if $\mathbb{N} = C_1 \cup C_2 \cup \ldots \cup C_r$ and $F \subset \mathbb{N}$ is finite, then there exists some $a \in \mathbb{N}$, $n \in \mathbb{N}$ and $j \in \{1, 2, \ldots, r\}$ so that C_j contains $\{a + nx : x \in F\}$. (Such a set is called an *affine* image of the set F.)

Exercise 8.4. Show that a syndetic set in $\mathbb{N}$ contains arbitrarily long arithmetic progressions.

Exercise 8.5. Show that a piecewise syndetic set in $\mathbb{N}$ contains arbitrarily long arithmetic progressions.

Exercise 8.6. Show that if a set $A \subseteq \mathbb{N}$ contains arbitrarily long arithmetic progressions and $A = A_1 \cup \ldots \cup A_r$ is a finite partition, then some A_j, $1 \le j \le r$ also contains arbitrarily long arithmetic progressions.

Exercise 8.7. Use van der Waerden's Theorem to show that for all $\alpha \in \mathbb{R}$ and all $\varepsilon > 0$, there exist $m, n \in \mathbb{N}$ such that $|n^2\alpha - m| < \varepsilon$. Generalize this for any polynomial $p(n)$ with integer coefficients such that $p(0) = 0$.

Exercise 8.8. Show that Theorem 8.6 is equivalent to the following: if (X, T) is a dynamical system, $k \in \mathbb{N}$ and $U \subseteq X$ is a nonempty open set, then there exists $n \in \mathbb{N}$ such that $U \cap T^n U \cap T^{2n} U \cap \ldots \cap T^{kn} U \neq \varnothing$.

Exercise 8.9. Show that for any $k \in \mathbb{N}$, there is a dense set of k-multiply recurrent points in any minimal dynamical system.

Exercise 8.10. Construct a dynamical system (X, T) and $x \in X$ such that x is recurrent under T, x is recurrent under T^2, but (x, x) is not recurrent under $T \times T^2$. (Hint: use a shift space.)

Exercise 8.11. Show that not all points in an arbitrary system satisfy the conclusion of Theorem 8.10. Show that not all points in a minimal system satisfy the conclusion of Theorem 8.10. (The first part is easy, but the second is not.)

Exercise 8.12. Show that if (X, T) is a dynamical system, $k \in \mathbb{N}$, $\varepsilon > 0$, and $a_1, a_2, \ldots, a_k \in \mathbb{N}$, then there exists $x \in X$ and $n \in \mathbb{N}$ such that $d(T^{a_j n} x, x) < \varepsilon$ for $j = 1, 2, \ldots, k$. Show that in a minimal system, this holds (with the same n) for a dense set of x.