TOPOLOGICAL DYNAMICS

STEFAN GESCHKE

1. Introduction

Definition 1.1. A *topological semi-group* is a topological space S with an associative continuous binary operation $\cdot$. A topological semi-group $(S,\cdot)$ is a *topological group* if $(S,\cdot)$ is a group and taking inverses is continuous.

Now let S be a topological semi-group and X a topological space. Given a map $\varphi : S\times X \to X$, we write sx for $\varphi(s,x)$. φ is a *continuous semi-group action* if it is continuous and for all $s,t\in S$ we have $s(tx) = (st)x$.

A *dynamical system* is a topological semi-group S together with a compact space X and a continuous semi-group action of S on X. X is the *phase space* of the dynamical system. Sometimes we say that X together with the S-action is an *S-flow*.

Exercise 1.2. Let S be a topological semi-group acting continuously on a compact space X. Show that for all $s\in S$ the map $f_s : X \to X; x\mapsto sx$ is continuous.

Exercise 1.3. Let G be a topological group acting continuously on a compact space X. Show that for all $g\in G$ the map $f_g : X\to X; x\mapsto gx$ is a homeomorphism.

Example 1.4. Let X be a compact space and let $f : X \to X$ be a continuous map. We define a semi-group action of $(\mathbb{N},+)$ on X by letting $nx = f^n(x)$ for each $n\in\mathbb{N}$. This action is continuous with respect to the discrete topology on $\mathbb{N}$.

If $f : X\to X$ is a homeomorphism, we can define $nx = f^n(x)$ for all $n\mathbb{Z}$ and obtain a group action of $(\mathbb{Z},+)$ on X. This action is continuous with respect to the discrete topology on $\mathbb{Z}$.

2. The Cantor space

Definition 2.1. The *middle thirds Cantor set* is the subspace of the unit interval $[0,1]$ that is defined as follows: Let $I_0 = [0,1]$. Suppose we have define $I_n \subseteq [0,1]$ and I_n is the union of 2^n disjoint closed intervals I_n^i, $1 \leq i \leq 2^n$. Let I_{n+1} be obtained by removing from each of the intervals I_n^i the open middle third. Now I_{n+1} is the union of 2^{n+1} disjoint closed intervals. The Cantor set is the intersection $\bigcap_{n\in\mathbb{N}} I_n$.

Recall that $2 = \{0,1\}$. 2 carries the discrete topology. The *Cantor cube* is the space $2^{\mathbb{N}}$ with the product topology. The Cantor cube is a compact metric space with respect to the metric defined as follows: for $x, y \in 2^{\mathbb{N}}$ let

$$d(x,y) = \begin{cases} 0, & \text{if } x = y \\ 2^{-n}, & \text{if } n \in \mathbb{N} \text{ is minimal with } x(n) \neq y(n) \end{cases}.$$

Exercise 2.2. Show that the Cantor cube is indeed a complete metric space with respect to the metric given in the previous definition. Show that in fact, with this metric the Cantor cube is *ultra-metric*, that is, it satisfies the following stronger form of the triangle inequality:

$$\forall x, y, z (d(x,z) \leq \max(d(x,y), d(y,z))$$

Definition 2.3. A topological space X is *zero-dimensional* iff every open set is the union of sets that are clopen, i.e., both open and closed. A point $x \in X$ is *isolated* if the set $\{x\}$ is open. $D \subseteq X$ is *dense* if it intersects every non-empty open subset of X. X is *separable* if it has a countable dense set.

Lemma 2.4. *Every compact metric space is separable.*

Proof. Let X be a compact metric space. For each $n \in N$ the open balls of radius 2^{-n} cover X. Since X is compact, finitely many open 2^{-n}-balls actually cover X. Let D_n be a finite subset of X such that the open 2^{-n}-balls centered at the elements of F_n cover X. Now for every $x \in X$ there is $y \in D_n$ such that the distance of x in y is smaller than 2^{-n}.

Let $D = \bigcup_{n\in\mathbb{N}} D_n$. Now for all $\varepsilon > 0$ and all $x \in X$ there is $y \in D$ such that the distance of x and y is less than ε. This shows that D is dense in X. Clearly, D is countable. □

Definition 2.5. Let Σ be a finite set, the *alphabet*. Σ^* is the set of all finite sequence of element of Σ, the set of *words* over Σ. Note that there is a word of length 0, the empty word λ. Given two words $v, w \in \Sigma^*$, by $v^\frown w$ we denote the concatenation of the two words, i.e., the word v followed by the word w. $(\Sigma^*, {}^\frown, \lambda)$ is a monoid, the *free monoid* over the set Σ. Formally, a word $v \in \Sigma$ is a function $f : \{0, \dots, n-1\} \to \Sigma$, where n is the *length* of v. We follow the set theoretic definition of the natural numbers, where the natural number n is just the set $\{0, \dots, n-1\}$. With this convention a word v of length n over Σ is just a function from n to Σ. We denote the set of functions from n to Σ, i.e., the set of words of length n over the alphabet Σ by Σ^n. With this notation we have $\Sigma^* = \bigcup_{n\in\mathbb{N}} \Sigma^n$.

Lemma 2.6. *The Cantor cube is compact, zero-dimensional, and has no isolated points.*

Proof. We first study the topology on $2^{\mathbb{N}}$. Given a finite binary sequence $s \in 2^n$, by $[s]$ we denote the set

$$\{x \in 2^{\mathbb{N}} : x \restriction n = s\},$$

i.e., the set of all infinite binary sequences that extend s.

A quick computation shows that $[s]$ is precisely the closed 2^{-n} ball centered at any of its elements. It is also the open $\frac{1001}{1000} \cdot 2^{-n}$-ball centered at any of its elements. Hence $[s]$ is both open and closed. It also follows that a set $O \subseteq 2^{\mathbb{N}}$ is open if and only if it is a union of sets of the form $[s]$. Therefore, $2^{\mathbb{N}}$ is zero-dimensional. Also, the sets $[s]$ are infinite. Hence, $2^{\mathbb{N}}$ has no isolated points.

Tychonov's theorem states that products of compact spaces are compact. The topology on $2^{\mathbb{N}}$ is in fact just the product topology coming from the discrete topology on 2. Since 2 is compact, so is $2^{\mathbb{N}}$.

Let us give a direct proof of the compactness of the Cantor cube. Suppose $\mathcal{O}$ is a familiy of open subsets of $2^{\mathbb{N}}$ such that $2^{\mathbb{N}} = \bigcup \mathcal{O}$. Let

$$\mathcal{S} = \{[s] : s \text{ is a finite binary sequence and there is } O \in \mathcal{O} \text{ such that } [s] \subseteq O\}.$$

Since every open set is the union of sets of the form $[s]$,

$$\bigcup \mathcal{S} = \bigcup \mathcal{O} = 2^{\mathbb{N}}.$$

Clearly, if $2^{\mathbb{N}}$ is the union of a finite subfamily of $\mathcal{S}$, then it is the union of a finite subfamily of $\mathcal{O}$. Hence we only have to show that $2^{\mathbb{N}}$ is covered by finitely many elements of $\mathcal{S}$.

We assume that this is not the case, i.e., no finitely many elements of $\mathcal{S}$ cover all of $2^{\mathbb{N}}$. First we drop the square brackets. Let

$$S = \{s \in 2^* : [s] \in \mathcal{S}\}.$$

Consider the collection T of all finite binary sequences that have no initial segment in S. Clearly, T is closed under taking initial segments. Also, T is infinite. If not, the elements of T have a maximal length n. There are 2^{n+1} binary sequences of length $n+1$. All of them have an initial segment in S. Hence, there is a finite subset S_0 of S such that every binary sequence of length at least 2^{n+1} has an initial segment in S_0. Now the finitely many sets $[s]$, $s \in S_0$, cover $2^{\mathbb{N}}$, contradicting our assumption. It follows that T is indeed infinite.

Recursively, we choose a sequence $(t_i)_{i\in\mathbb{N}}$ of elements of T such that for all $i \in \mathbb{N}$ the following hold:

(1) t_i is of length i
(2) The set of all extensions of t_i in T is infinite.

For t_0 we have to choose the empty sequence. Suppose we have chosen t_i. Since t_i has infinitely many extensions in T, $t_i{}^\frown 0$ or $t_i{}^\frown 1$ have infinitely many extensions in T. Let t_{i+1} be either $t_i{}^\frown 0$ or $t_i{}^\frown 1$, but in such a way that t_{i+1} has infinitely many extensions in T.

Finally, let $x = \bigcup_{i\in\mathbb{N}} t_i$. This infinite binary sequence has no initial segment in S and hence is not an element of $\bigcup \mathcal{S}$. This contradicts our assumptions on $\mathcal{S}$. □

Theorem 2.7. *Every zero-dimensional compact metric space without isolated points is homeomorphic to the Cantor cube* $2^{\mathbb{N}}$.

Proof. Let X be a zero-dimensional compact space without isolated points. We define a strictly increasing sequence $(n_i)_{i\in\mathbb{N}}$ of natural numbers > 0 and for each $i \in \mathbb{N}$ a family

$$\langle U_s : s \text{ is a word of length } n_i \text{ over the alphabet } 2\rangle$$

of nonempty clopen subsets of X such that $U_\lambda = X$ and for all $i \in \mathbb{N}$ the following hold:

(1) If $s, t \in 2^{n_i}$ are different, then $U_s \cap U_t = \emptyset$.

(2) For all $s \in 2^{n_i}$, $U_s = \bigcup_{t \in 2^{n_{i+1}-n_i}} U_{s^\frown t}$.
(3) For all $s \in 2^{n_i}$, U_s is of diameter at most 2^{-i}.

Suppose we have already chosen n_i and the sets U_s for all $s \in 2^{n_i}$. Fix $s \in 2^{n_i}$ for a moment. U_s is a compact subset of X. Therefore it is the union of finitely many nonempty clopen sets of diameter not more than $2^{-(i+1)}$. Since X has no isolated points, every nonempty clopen set can be split into two nonempty clopen subsets. Hence, for sufficiently large $n > 0$, U_s is the union of 2^n pairwise disjoint nonempty clopen subset.

We can now choose some $n \in \mathbb{N}$ that works for every $s \in 2^{n_i}$ at the same time: Each U_s, $s \in 2^{n_i}$, is the union of 2^n pairwise disjoint clopen subsets of diameter not more than $2^{-(i+1)}$. We denote these sets by $U_{s^\frown t}$, $t \in 2^n$. Finally, let $n_{i+1} = n_i + n$. This finishes the definition of the sequence $(n_i)_{i \in \mathbb{N}}$ and of the sets U_s.

We now define a homeomorphism h from the Cantor cube $2^{\mathbb{N}}$ onto X. For each $s \in 2^{\mathbb{N}}$ let $h(s)$ be the unique element of $\bigcap_{i \in \mathbb{N}} U_{s \restriction n_i}$. It is easily checked that h is continuous, 1-1 and onto. Hence h is a homeomorphism. □

Corollary 2.8. *The middle thirds Cantor set is homeomorphic to the Cantor cube $2^{\mathbb{N}}$. Given a finite alphabet Σ, the space $\Sigma^{\mathbb{N}}$ is homeomorphic to $2^{\mathbb{N}}$.*

Exercise 2.9. In order to prove this corollary, we have to check that the middle thirds Cantor set is a zero-dimensional compact metric space. Verify this.

Definition 2.10. Let Σ be a finite alphabet. We define the *shift action* of $\mathbb{Z}$ on $\Sigma^{\mathbb{Z}}$ as follows. For $x \in \Sigma^{\mathbb{Z}}$ and $n \in \mathbb{Z}$ let $nx \in \Sigma^{\mathbb{Z}}$ be defined such that for all $m \in \mathbb{Z}$ we have $(nx)(m) = x(m+n)$. The same definition gives a semi-group action of $\mathbb{N}$ on $\Sigma^{\mathbb{N}}$ that we also call the *shift action.*

3. Orbits and invariant sets

Definition 3.1. Let X be an S-flow for some semi-group S. For every $x \in X$, the *orbit* of x is the set $Sx = \{sx : s \in S\}$. We also consider the *orbit closure* of x, the closure $\mathrm{cl}(Sx)$ of the orbit Sx. A set $A \subseteq X$ is *closed under S* if $SA = \{sa : s \in S \wedge a \in A\} \subseteq A$. We call a nonempty set $A \subseteq X$ that is topologically closed and closed under S a *subflow* of X. A is S*-invariant* if for all $s \in S$ we have $sA = \{sa : a \in A\} = A$.

Note that every subflow is itself a flow with respect to the restricted semi-group action.

Exercise 3.2. Let S be a semi-group and let X be an S-flow.

a) Show that if S is a group, then the S-invariant subsets of X are precisely the sets that are closed under S.

b) Let X be an S-flow for some semi-group S. Show that if $A \subseteq X$ is closed under S, then the closure of A is also closed under S.

In particular, for every point x in an S-flow X, the orbit closure $\mathrm{cl}(Sx)$ is a subflow of X.

Example 3.3. Let $x \in 2^{\mathbb{N}}$ be such that every finite binary sequence occurs in x. This is the case if for every finite binary sequence s there is some $n \in \mathbb{N}$ such that s is an initial segment of nx, i.e., $nx \in [s]$. It follows that the orbit of x intersects every nonempty open subset of $2^{\mathbb{N}}$. It is dense. Hence the orbit closure of x is all of $2^{\mathbb{N}}$.

The two sequences that are constantly 0 or 1 are the only fixed points of the action of $\mathbb{N}$ on $2^{\mathbb{N}}$. Their orbit closures only consist of a single point. The subflows of $2^{\mathbb{Z}}$ and of $2^{\mathbb{N}}$ are frequently called *subshifts*.

Exercise 3.4. Let $x, y \in 2^{\mathbb{N}}$ Show that x is in the orbit closure of y if and only if every initial segment of x occurs in y.

Definition 3.5. An S-flow X is minimal if X has no nonempty closed subset that is closed under S.

Note that a flow is minimal if and only if every point has a dense orbit.

Theorem 3.6. *Every S-flow X has a minimal subflow.*

Proof. A family $\mathcal{F}$ of sets has the *finite intersection property* (*fip*) if the intersection of any finitely many sets from $\mathcal{F}$ is non-empty. Recall that in a compact space X every family $\mathcal{C}$ of closed sets that has the finite intersection property has a nonempty intersection.

The proof of the theorem is an application of Zorn's Lemma. Let $\mathcal{P}$ denote the collection of all topologically closed non-empty subsets of

X that are closed under S. $\mathcal{P}$ is partially ordered by reverse inclusion. Let $\mathcal{C} \subseteq \mathcal{P}$ be a chain with respect to this partial order. $\mathcal{C}$ has the finite intersection property: if $\mathcal{C}_0 \subseteq \mathcal{C}$ is finite, then it has a largest element C with respect to the partial order. With respect to inclusion, C is the smallest element of $\mathcal{C}_0$ and hence $\bigcap \mathcal{C}_0 = C \neq \emptyset$.

It follows that $\bigcap \mathcal{C}$ is a nonempty closed subset of X. Clearly, $\bigcap \mathcal{C}$ is closed under S. Hence $\bigcap \mathcal{C} \in \mathcal{P}$. It follows that the chain $\mathcal{C}$ has an upper bound in $\mathcal{P}$, namely $\bigcap \mathcal{C}$.

By Zorn's Lemma, $\mathcal{P}$ has a maximal element C. With respect to inclusion, C is a minimal nonempty subset of X that is topologically closed and closed under S. Ie., C is a minimal subflow of X. □

In some sense, minimal flows are the flow-equivalents of prime numbers.

Example 3.7. We give an explicit example of a minimal $\mathbb{N}$-flow (even $\mathbb{Z}$-flow). The example is the rotation on the unit circle by an irrational angle. For simplicity, instead of the circle we consider the unit interval with the endpoints identified, the one-dimensional torus T^1. We write the elements of T^1 as reals in the interval $[0, 1)$. The distance of two points is the minimal length of the arc from one point to the other.

Let α be an irrational number. On T^1 we consider the map

$$f : T^1 \to T^1; x \mapsto (x + \alpha) \mod 1.$$

Since f is a homeomorphism of the 1-torus, it give rise to an $\mathbb{Z}$-action on T^1. We will for the moment only consider the $\mathbb{N}$-action induced by f.

Let $x \in T^1$. We consider the orbit $\mathbb{N}x$. If the orbit is finite, then there are $n, m \in \mathbb{N}$ such that $nx = mx$ and $n < m$. It follows that $(x + n\alpha) \mod 1 = (x + m\alpha) \mod 1$. Hence $(m - n)\alpha$ is an integer, contradicting the assumption that α is irrational. This argument not only shows that $\mathbb{N}x$ is infinite, but in fact that for $n < m$ we have $nx \neq mx$.

Now consider the sequence $(nx)_{n\in\mathbb{N}}$. Since T^1 is a compact metric space, the sequence has a convergent subsequence. It follows that for every $\varepsilon > 0$ there are $n, m \in \mathbb{N}$ such that $n < m$ and the distance between nx and mx is less than ε. It follows that the distance between $(m - n)\alpha$ and 0 on the 1-torus is less than ε. But this implies that every point on the 1-torus has distance $< \varepsilon$ to some point of the form $(x+k(m-n)\alpha) \mod 1$, where $k \in \mathbb{N}$. Since ε was arbitrary, this shows that $\mathbb{N}x$ is dense in T^1.

Hence every orbit is dense in T^1. Therefore, T^1 is a minimal $\mathbb{N}$-flow. This also implies that the corresponding $\mathbb{Z}$-flow is minimal.

We investigate a subshifts of $2^{\mathbb{Z}}$ that are closely related to the rotations in Example 3.7, the *Sturmian subshifts*.

Definition 3.8. Let $G = \mathbb{N}$ or $G = \mathbb{Z}$. A *Sturmian word* is a word $x \in \{0,1\}^G$ such that there are two real numbers, the *slope* α and the *intercept* ρ, with $\alpha \in [0,1)$ irrational such that for all $i \in G$ we have

$$x(i) = 1 \quad \Leftrightarrow \quad (\rho + i \cdot \alpha) \mod 1 \in [0, \alpha).$$

In the context of $\mathbb{N}$-flows, we consider Sturmian words in $\{0,1\}^{\mathbb{N}}$ and when we talk about $\mathbb{Z}$-flows, we consider Sturmian words in $\{0,1\}^{\mathbb{Z}}$.

It turns out that the orbit closure $C_x = \mathrm{cl}(Gx)$ of a Sturmian word x with the restriction of the shift is a minimal G-flow. In fact, this orbit closure consist of all Sturmian words that have the same slope as x.

If $x \in \{0,1\}^{\mathbb{Z}}$ is a Sturmian word of slope α, then for all y in the orbit closure of x the limit

$$\lim_{n\to\infty} \frac{|x^{-1}(1) \cap \{-n, \dots, n\}|}{2n+1}$$

exists and equals α. Similarly, if $x \in \{0,1\}^{\mathbb{N}}$ is a Sturmian word of slope α, then for all y in the orbit closure of x the limit

$$\lim_{n\to\infty} \frac{|y^{-1}(1) \cap \{0, \dots, n-1\}|}{n}$$

exists and equals α.

It follows that for different irrational numbers $\alpha, \beta \in [0,1)$, Sturmian words of slope α and β have different (even disjoint) orbit closures. We call the orbit closure of a Sturmian word together with the restriction of the shift a *Sturmian subshift*. A Sturmian subshift is a G-flow.

Given a Sturmian subshift $X \subseteq 2^G$, we denote the common slope of all Sturmian words that generate X by $\alpha(X)$. We will revisit Sturmian dynamics once we have more tools at hand.

4. Flow maps, factors, and embeddings

Definition 4.1. . We fix a semi-group S and S-flows X and Y. A map $f : X \to Y$ is a *flow map* if it is continuous and for all $s \in S$ and all $x \in X$ we have $f(sx) = sf(x)$. Maps that commute with actions in this way are also called *equivariant* or *S-equivariant*.

Y is a *factor* of X with *factor map* $f : X \to Y$ if f is a flow map that is onto.

A flow map $f : X \to Y$ is an *embedding* if it is 1-1.

Flow maps are the morphisms in the category of S-flows for a fixed semi-group S.

Example 4.2. Let T^1 be the 1-torus with a $\mathbb{Z}$-action that comes from a rotation by some irrational number α as in Example 3.7. We consider the cylinder $Z = T^1 \times [0, 1]$ with the $\mathbb{Z}$-action defined by $n(x_1, x_2) = (nx_1, x_2)$ for all $x_{\in}T^1$ and $x_2 \in [0, 1]$. Then $f : T_1 \to Z; x \mapsto (x, 0)$ is an embedding and $g : Z \to T^1; (x_1, x_2) \mapsto x_1$ is a factor map.

Exercise 4.3. a) Let X be a minimal S-flow for some semi-group S. Let Y be a factor of X. Show that Y is a minimal S-flow.

b) Let X be an S-flow and let Y be a minimal S-flow. Show that every flow map $f : X \to Y$ is a factor-map.

c) Let Y be a minimal S-flow. If there there is a flow map $f : X \to Y$, then Y is a factor of a minimal subflow of X.

Theorem 4.4 (Anderson)**.** *Let G be a countable group with the discrete topology. Then every metrizable G-flow is a factor of a 0-dimensional metrizable G-flow.*

Proof. We may assume that X is infinite. Otherwise it is zero-dimensional and we have nothing to prove. Let $(U_n, V_n)_{n\in\mathbb{N}}$ be an enumeration of all pairs of nonempty, disjoint open subsets of X from a fixed basis of the topology on X. Consider the space $2^{\mathbb{N}\times G}$ with the following continuous G-action:

For all $s \in 2^{\mathbb{N}\times G}$, $g, h \in G$, and $n \in \mathbb{N}$ let

$$(hs)(n, g) = s(n, h^{-1}g).$$

Note that G acts continuously on $X \times 2^{\mathbb{N}\times G}$ by the coordinate wise action.

Now let $Y \subseteq X \times 2^{\mathbb{N}\times G}$ consist of all pairs (x, s) such that for all $n \in \mathbb{N}$ and all $g \in G$ we have

(1) if $x \in g[U_n]$, then $s(n, g) = 0$ and
(2) if $x \in g[V_n]$, then $s(n, g) = 1$.

We first observe that Y is a closed subset of $X \times 2^{\mathbb{N}\times G}$. Suppose $(x, s) \in X \times 2^{\mathbb{N}\times G}$ is not in Y. Assume there are n and g violating (1). So $x \in g[U_n]$ and $s(n, g) = 1$. Now the pairs (x', s') with $x' \in g[U_n]$ and $s'(n, g) = 1$ form an open neighborhood of (x, s) that is disjoint from Y. If there are n and g violating (2), we obtain an open neighborhood of (x, s) that is disjoint from Y in the same way.

Y is zero-dimensional since any two of its elements can be separated by a clopen set. Namely, let $(x, s), (x', s') \in Y$ be distinct. If $s \neq s'$ there is a clopen subset A of $2^{\mathbb{N}\times G}$ such that $s \in A$ and $s' \notin A$. Now $Y \cap (X \times A)$ is a clopen subset of Y that contains (x, s) but not (x', s'). If $x \neq x'$, then there is n such that $x \in U_n$ and $x' \in V_n$. By the definition of Y, this implies that $s \neq s'$ and hence (x, s) and (x', s') are separated by a clopen subset of Y.

Also, Y is closed under the action of G. To see this, let $(x, s) \in Y$ and $h \in G$. We have to show $(hx, hs) \in Y$. Let $n \in \mathbb{N}$ and $g \in G$. Suppose $hx \in g[U_n]$. Then $x \in h^{-1}g[U_n]$. Since $(x, s) \in Y$, $0 = s(n, h^{-1}g) = (hs)(n, g)$. Hence (hx, hs) satisfies condition (1) for n and g. The argument for condition (2) is the same.

It follows that Y is a metrizable zero-dimensional G-flow. Let $\pi_1 : X \times 2^{\mathbb{N}\times G} \to X$ be the projection onto the first coordinate and let φ be its restriction to Y. It is clear that φ is continuous, G-equivariant, and onto. This finishes the proof of the theorem. □

We observe that if X is a minimal metric G-flow for some countable group and X is a factor of a zero-dimensional metric G-flow Y, then we can choose a minimal subflow Z of Y, which still maps onto X since X is minimal. Hence, Anderson's theorem also holds for minimal flows:

Every minimal metric G-flow is a factor of a G-flow that is minimal, metric, and zero-dimensional.

We now return to the particular case of irrational rotations on the 1-torus and Sturmian subshifts. Let $G = \mathbb{Z}$. The arguments for the case $G = \mathbb{N}$ are similar. Fix an irrational number $\alpha \in [0, 1)$ and consider the rotation

$$f : T^1 \to T^1; x \mapsto (x + \alpha) \mod 1.$$

The irrational number α splits the 1-torus into two disjoint sets, $V_1 = [0, \alpha)$ and $V_0 = [\alpha, 1)$. (Recall that we identify the elements of T^1 with the elements of $[0, 1)$.) To each $\rho \in T^1$ we assign the Sturmian sequence $x_\rho \in 2^{\mathbb{Z}}$ as in Definition 3.8.

Let $\rho, \sigma \in T^1$ be distinct. We have already observed that every $\mathbb{G}$-orbit is dense in T^1. In fact, our argument shows that the *backward orbit* $(-\mathbb{N})\alpha$ of α is dense in T^1. It follows that there is $n \in \mathbb{N}$ such that $\rho \in f^{-n}[V_1]$ and $\sigma \in f^{-n}[V_0]$. For this n we have $x_\rho(n) = 1$ and $x_\sigma(n) = 0$. This shows that the map $T^1 \to 2^{\mathbb{Z}}; \rho \mapsto x_\rho$ is 1-1.

Exercise 4.5. Show that $\rho \mapsto x_\rho$ is not continuous.

Clearly, for all $n, m \in \mathbb{Z}$ and all $\rho \in T^1$ we have

$$x_{n\rho}(m) = 1 \Leftrightarrow (\rho + n\alpha + m\alpha) \mod 1 \in [0, \alpha) \Leftrightarrow nx_\rho(m) = 1.$$

In other words, $\rho \mapsto x_\rho$ is $\mathbb{Z}$-equivariant.

Fix $\rho \in T^1$. We define a map $\pi : \mathrm{cl}(\mathbb{Z}x_\rho)$ to T^1 as follows. Let $x \in \mathrm{cl}(\mathbb{Z}x_\rho)$. Then every word of the form $(x(-n), \dots, x(n))$ occurs in x_ρ.

Consider the set

$$\bigcap_{n \in \mathbb{Z}} f^{-n}[\mathrm{cl}(V_{x(n)})] \subseteq T^1.$$

Since T^1 compact, the intersection is nonempty if no finite intersection

$$\bigcap_{i=-n}^{n} f^{-i}[\mathrm{cl}(V_{x(i)})]$$

is empty.

But for every $n \in \mathbb{N}$ there is $m \in \mathbb{Z}$ such that $(x(-n), \dots, x(n)) = (x_\rho(m-n), \cdots, x_\rho(m+n))$. By the definition of x_ρ,

$$\rho \in \bigcap_{i=m-n}^{m+n} f^{-i}[\mathrm{cl}(V_{x_\rho(i)})].$$

Therefore

$$f^m(\rho) \in \bigcap_{i=-n}^{n} f^{-i}[\mathrm{cl}(V_{x(i)})].$$

It follows that $\bigcap_{n \in \mathbb{Z}} f^{-n}[\mathrm{cl}(V_{x(n)})]$ is indeed nonempty. Also, since the orbits of 0 and α in T^1 are dense, $\bigcap_{n \in \mathbb{Z}} f^{-n}[\mathrm{cl}(V_{x(n)})]$ has not more than one element. This unique element we call $\pi(x)$. The whole construction is $\mathbb{Z}$-equivariant, so $\pi : \mathrm{cl}(\mathbb{Z}x) \to T^1$ is a $\mathbb{Z}$-equivariant map.

We show that π is continuous. Let $x \in \mathrm{cl}(\mathbb{Z}x_\rho)$ and let $\sigma = \pi(x)$. Let $U \subseteq T^1$ be an open set containing σ. The closed set $T^1 \setminus U$ is covered by the open sets

$$T^1 \setminus \bigcap_{i=-n}^{n} f^{-i}[\mathrm{cl}(V_{x(i)})]$$

and hence by finitely many of them. Since these sets are increasing with n, one such set is enough to cover $T^1 \setminus U$. It follows that there is $n \in \mathbb{N}$ such that $\bigcap_{i=-n}^{n} f^{-i}[\mathrm{cl}(V_{x(i)})] \subseteq U$.

Now all $y \in \mathrm{cl}(\mathbb{Z}x_\rho)$ that agree with x on the set $\{-n, \dots, n\}$ are mapped by π into the set U. This shows the continuity of π.

Finally, we show that $\mathrm{cl}(\mathbb{Z}x_\rho)$ is a minimal subshift. It is enough to show that x_ρ is in the orbit closure of every $x \in \mathrm{cl}(\mathbb{Z}x_\rho)$. Let $x \in \mathrm{cl}(\mathbb{Z}x_\rho)$. We show that every word of the form $(x_\rho(-n), \dots, x_\rho(n))$ occurs in x.

Let $n \in \mathbb{N}$. Consider the set $\bigcap_{i=-n}^{n} f^{-i}[V_{x_\rho(i)}]$. By the definition of x_ρ, ρ is an element of this set, so the set is non-empty. Since the set is an intersection of half-open intervals, it contains a nonempty open interval (a, b). Since every orbit in T^1 is dense, there is $m \in \mathbb{Z}$ such that $\pi(mx) = m\pi(x) \in (a, b)$. By the definition of $\pi(mx)$, $\pi(mx) \in \bigcap_{i=-n}^{n} f^{-i}[\mathrm{cl}(V_{mx(i)})]$. It follows that for every $i \in \{-n, \dots, n\}$ we have $\pi(mx) \in f^{-i}[\mathrm{cl}(V_{mx(i)})]$. Since $\pi(mx) \in (a, b) \subseteq f^{-i}[\mathrm{cl}(V_{x_\rho(i)})]$, $\pi(mx)$ is not one of the endpoints of $f^{-1}[V_1]$ and $f^{-1}[V_0]$.

It follows that for every $i \in \{-n, \dots, n\}$, $\pi(mx) \in f^{-i}[V_{x_\rho(i)}]$ iff $\pi(mx) \in f^{-i}[V_{mx(i)}]$. Since the sets $f^{-i}[V_0]$ and $f^{-1}[V_1]$ are disjoint, this implies that mx and x_ρ agree on the set $\{-n, \dots, n\}$. In other words,

$$(x(m-n), \dots, x(m+n)) = (x_\rho(-n), \dots, x_\rho(n)),$$

which is what we wanted to prove.

We have thus shown the following:

Theorem 4.6. *Let $\alpha \in [0, 1)$ be irrational and let $f : T^1 \to T^1; x \mapsto (x + \alpha) \mod 1$. Then the orbit closure of every Sturmian word $x \in 2^{\mathbb{Z}}$ of slope α is a minimal subshift of $2^{\mathbb{Z}}$ and T^1 with the $\mathbb{Z}$-action induced by f is a factor of the Sturmian subshift $\mathrm{cl}(\mathbb{Z}x)$.*

The analog holds true for the $\mathbb{N}$-action induced by f.

5. Recurrence

Theorem 5.1 (Simple recurrence theorem for open covers)**.** *Let f be a homeomorphism on a compact space X. Let $\mathcal{U}$ be an open cover of X. Then there is an open set $U \in \mathcal{U}$ such that for infinitely many many $n \in \mathbb{Z}$, $U \cap f^n[U] \neq \emptyset$.*

Proof. Since X is compact, we may assume that $\mathcal{U}$ is finite. Consider the action $(n,x) \to nx = f^n(x)$ of $\mathbb{Z}$ on X. By the infinite pigeonhole principle, for every $x \in X$ there is $U \in \mathcal{U}$ such that for infinitely many $n \in \mathbb{Z}$, $nx \in U$. In other words, the recurrence set $S = \{n \in \mathbb{Z} : nx \in U\}$ is infinite. Let $n_0 \in S$. Now $n_0 x \in U$ and for every $n \in S$, $(n - n_0)(n_0 x) = nx \in U$ and therefore $f^{n-n_0}[U] \cap U \neq \emptyset$. □

Definition 5.2. Let (X,d) be a compact metric space and let f be a homeomorphism of X. As above, we denote $f^n(x)$ by nx. A set $A \subseteq \mathbb{Z}$ is *syndetic* if there is $m \in \mathbb{N}$ such that for all $a \in \mathbb{Z}$, $A \cap \{a, a+1, \dots, a+m\} \neq \emptyset$, i.e., if A has *bounded gaps.* Let $x \in X$

1. x is *invariant* or a *fixed point* if $f(x) = x$.
2. x is *periodic* if $nx = x$ for some $n > 0$.
3. x is *almost periodic* if for every $\varepsilon > 0$ the set $\{n \in \mathbb{Z} : d(nx,x) < \varepsilon\}$ is syndetic.
4. x is recurrent if for every $\varepsilon > 0$ the set $\{n \in \mathbb{Z} : d(nx,x) < \varepsilon\}$ is infinite.

Equivalently, x is recurrent if there is a sequence $(n_i)_{i\in\mathbb{N}}$ of integers such that $\lim_{i\to\infty} n_i x = x$.

Theorem 5.3. *Let X be a compact space and let f be a homeomorphism on X such that the corresponding $\mathbb{Z}$-flow is minimal. Then every element of X is almost periodic and in particular recurrent.*

Proof. Suppose the point $x \in X$ is not almost periodic. Then there is $\varepsilon > 0$ such that the set $\{n \in \mathbb{Z} : d(nx,x) < \varepsilon\}$ is not syndetic. Thus, for every $m \in \mathbb{N}$ there is $n_m \in \mathbb{Z}$ such that for all $n \in \{n_m - m, n_m + m\}$, $d(nx,x) \geq \varepsilon$. Since X is compact, the sequence $(n_m x)_{m\in\mathbb{N}}$ has a convergent subsequence $(n_{m_i} x)_{i\in\mathbb{N}}$. By the continuity of the metric and of f, for every $k \in \mathbb{Z}$ we have

$$d(ky, x) = \lim_{i\to\infty} d((n_{m_i} + k)x, x) \geq \varepsilon.$$

But this shows that the orbit closure of y does not contain x, contradicting the minimality of X. □

Since every $\mathbb{Z}$-flow contains a minimal $\mathbb{Z}$-flow, this theorem implies:

Corollary 5.4 (Birkhoff recurrence theorem)**.** *Let X be a compact space with a homeomorphism f. Then X contains an almost periodic point.*

We will now use topological dynamics to prove a Ramsey theoretic statement that does not seem mention group actions at all.

Theorem 5.5 (van der Waerden)**.** *Let $\mathbb{Z} = P_1 \cup \dots \cup P_n$ be a partition of $\mathbb{Z}$ into finitely many classes. Then at least one of the classes contains arbitrarily long arithmetic progressions, i.e., sets of the form $\{a, a + r, \dots, a + (k-1)r\}$.*

We will derive this theorem from the following:

Theorem 5.6 (Multiple recurrence theorem in open covers)**.** *Let f be a homeomorphism of a compact space X. If $\mathcal{U}$ is an open cover of X, then there is $U \in \mathcal{U}$ such that for all $k \geq 1$ there are infinitely many $r \in \mathbb{Z}$ such that*

$$U \cap f^{-r}[U] \cap \dots \cap f^{-(k-1)r}[U] \neq \emptyset.$$

In order to see the connection between Theorem 5.5 and Theorem 5.6, we introduce the Čech-Stone compactification of the integers.

Definition 5.7. Let M be a nonempty set. A *filter* on M is a nonempty family $p \subseteq \mathcal{P}(M)$ such that for all $A, B \subseteq M$ the following hold:

(1) $\emptyset \notin p$
(2) If $A \in p$ and $A \subseteq B$, then $B \in p$.
(3) If $A, B \in p$, then $A \cap B \in p$.

A filter on M is a *ultrafilter*, if for all $A \subseteq M$, either $A \in p$ or $M \setminus A \in p$.

Lemma 5.8. *A family p of subsets of a nonempty set M is an ultrafilter if and only if it is a maximal family of subsets of M with the finite intersection property.*

Proof. Let p be an ultrafilter. Then p is closed under finite intersections and does not contain the empty set. It follows that it has the finite intersection property. Since for each set $A \subseteq M$, p contains either A or $M \setminus A$, it is maximal with respect to the finite intersection property.

Now let p be a maximal family of subsets of M with the finite intersection property. Because of the finite intersection property, $\emptyset \notin p$. If $A \in p$ and $B \subseteq M$ is such that $A \subseteq B$, then $p \cup \{B\}$ also has the finite intersection property and hence, by the maximality of p, $B \in p$. Similarly, if $A, B \in p$, then the family $p \cup \{A \cap B\}$ has the finite intersection property and hence, again by the maximality of p, $A \cap B \in p$. □

Definition 5.9. For every nonempty set M let βM denote the collection of all ultrafilters over the set M with the topology gererated by the sets of the form $\hat{A} = \{p \in \beta M : A \in p\}$, where $A \subseteq M$.

Observe that for every $m \in M$, $\{A \subseteq M : m \in A\}$ is an ultrafilter, the *principal* ultrafilter generated by m. Identifying the principal ultrafilters generated by the elements of M with the respective elements of M itself, we can consider M as a subset of βM.

Lemma 5.10. *βM is a compact Hausdorff space and M is a dense subset of βM with the discrete topology. Every map f from M into a compact space X has a unique continuous extension $\beta f : \beta M \to X$.*

Proof. We first show the Hausdorffness of βM. Let p and q be two distinct ultrafilters on M. Then there is $A \subseteq M$ which is contained in only one of the two ultrafilters, say in p. The set $B = M \setminus A$ is then contained in q. Now $\hat{A}$ and $\hat{B}$ are disjoint neighborhoods of p and q, respectively.

The density of M in βM is easily checked. First note that $\hat{\emptyset}$ is empty. If $A \subseteq M$ is nonempty, we can choose $m \in A$ and now $\hat{A}$ contains the principal ultrafilter generated by m. This shows the density of M.

For the compactness of βM let $\mathcal{U}$ be an open cover of βM. We may assume that all the cover consists of sets of the form $\hat{A}$ for some $A \subseteq M$. Suppose $\mathcal{U}$ does not have a finite subcover. Consider the family

$$\mathcal{F} = \{A \subseteq M : (M \setminus A)\hat{} \in \mathcal{U}\}.$$

This family has the finite intersection property and hence extends to an ultrafilter p. Now $p \notin \bigcup \mathcal{U}$, a contradiction.

Finally, let X be compact and let $f : M \to X$ be any function. Given $p \in \beta M$ we define $\beta f(p)$ as the unique element of $\bigcap_{A\in P} \mathrm{cl}(f[A])$. It is easily checked that $\beta f \restriction M = f$ and that βf is continuous. Since M is dense in βM, βf is unique. $\square$

Exercise 5.11. Let M be an infinite set. Show that no sequence $(m_n)_{n\in\mathbb{N}}$ of pairwise distinct elements of M converges to a point in βM. In particular, βM is not metrizable.

Actually, one can prove that βM has no nontrivial convergent sequences at all.

We now show that Theorems 5.6 and 5.5 are equivalent.

First assume Theorem 5.6 and let $\mathbb{Z} = P_1 \cup \dots \cup P_n$ be a partition of $\mathbb{Z}$. Also, let $k \geq 1$. Consider $f : \mathbb{Z} \to \mathbb{Z}; m \mapsto m+1$ and let $\beta f : \beta\mathbb{Z} \to \beta\mathbb{Z}$ be the unique extension of f to $\beta\mathbb{Z}$. Note that $\hat{P}_1, \dots, \hat{P}_n$ is an open cover of $\beta\mathbb{Z}$.

By Theorem 5.6, there are $i \in \{1, \dots, n\}$ and $r \in \mathbb{Z}$ such that

$$\hat{P}_i \cap \beta f^{-r}[\hat{P}_i] \cap \dots \cap \beta f^{-(k-1)r}[\hat{P}_i] \neq \emptyset.$$

It is clear that $\beta f^{-m}[\hat{P}_i] = (f^{-m}[P_i])\hat{}$. Hence

$$P_i \cap f^{-r}[P_i] \cap \dots \cap f^{-(k-1)r}[P_i] \neq \emptyset.$$

Choose $a \in P_i \cap f^{-r}[P_i] \cap \dots \cap f^{-(k-1)r}[P_i]$. Now we have $a, a+r, \dots, a+r(k-1) \in P_i$. This shows Theorem 5.5.

Now assume Theorem 5.5 holds. Let X be a compact space with a homeomorphism f and let $\mathcal{U}$ be an open cover of X. Also, let $k \geq 1$. Let $U_1, \dots, U_n$ be a finite subcover of $\mathcal{U}$. Choose any $x \in X$. For $i \in \{1, \dots, n\}$ let $P_i = \{m \in \mathbb{Z} : f^m(x) \in U_i\}$. The P_i might not be disjoint, but together they cover $\mathbb{Z}$ and hence, by Theorem 5.5, one of them contains an arithmetic progression $a, a+r, \dots, a+(k-1)r$. Now let $y = f^a(x)$. We have $y, f^r(y), \dots, f^{(k-1)r}(y) \in U_i$. This shows

$$y \in U_i \cap f^{-r}[U_i] \cap \dots \cap f^{-(k-1)r}[U_i]$$

and Theorem 5.6 follows.

Lemma 5.12. *Let X be a compact space and let $f : X \to X$ be a homeomorphism such that the induced $\mathbb{Z}$-flow is minimal. Then for every nonempty open set $U \subseteq X$, X is the union of finitely many sets of the form $f^n[U]$, $n \in \mathbb{N}$.*

Proof. Consider the set $Y = X \setminus \bigcup_{n\in\mathbb{Z}} f^n[U]$. Y is closed and closed under the $\mathbb{Z}$-action. Since U is nonempty, $Y \neq X$. Since X is minimal, Y must be empty. It follows that $X = \bigcup_{n\in\mathbb{Z}} f^n[U]$. By the compactness of X, finitely many sets of the form $f^n[U]$ cover it. □

Theorem 5.6 and hence Theorem 5.5 follow from the following result.

Theorem 5.13. *Let X be a compact space and let $f : X \to X$ be a homeomorphism such that the induced $\mathbb{Z}$-flow is minimal. Let $U \subseteq X$ be a non-empty open set and let $k \geq 1$. Then U contains an* arithmetic progression *$x, f^r(x), \dots, f^{(k-1)r}(x)$ for some $x \in X$ and $r \geq 1$.*

Proof. Since X is a minimal $\mathbb{Z}$-flow, finitely many translates $f^m[U]$, $m \in \mathbb{Z}$, cover X. We fix an open cover $\mathcal{U}$ of X by finitely many translates of U.

We prove the theorem by induction on k. It is trivially true for $k = 1$. Now let $k > 1$.

We prove the following claim first:

Claim 5.14. Then for every $J \geq 0$ there exists a sequence $x_0, \dots, x_J$ of points in X, a sequence $U_0, \dots, U_J$ of sets in the open cover (not necessarily distinct), and a sequence $r_1, \dots, r_J$ of positive integers such that $f^{i(r_{a+1}+\dots+r_b)}(x_b) \in U_a$ for all a, b with $0 \leq a \leq b \leq J$ and all i with $1 \leq i \leq k-1$.

The proof of the claim proceeds by induction on J. The case $J = 0$ is trivial. Now let $J \geq 1$ and suppose we have already constructed $x_0, \dots, x_{J-1}$, $U_0, \dots, U_{J-1}$, and $r_1, \dots, r_{J-1}$ with the required properties. Now let V be a suitably small neighbourhood of x_{J-1} (depending on all the above data) to be chosen later. By Theorem 5.13 for k-1, V contains an arithmetic progression $y, f^{r_J}(y), \dots, f^{(k-2)r_J}(y)$ of length $k-1$. Now set $x_J := f^{-r_J}(y)$ and let U_J be an arbitrary set in the open cover containing x_J. We observe that

$$f^{i(r_{a+1}+\dots+r_J)}(x_J) = f^{i(r_{a+1}+\dots+r_{J-1})}(f^{(i-1)r_J}(y)) \in f^{i(r_{a+1}+\dots+r_{J-1})}[V]$$

for all $a \in \{0, \dots, J-1\}$ and $i \in \{1, \dots, k-1\}$. If V is a sufficiently small neighbourhood of x_{J-1}, we thus see (from the continuity of the $f^{i(r_{a+1}+\dots+r_{J-1})}$) that we verified all the required properties needed to close the induction. This proves the claim.

We apply the above claim with J equal to the number of sets in the open cover. By the pigeonhole principle, we can thus find a, b with $0 \leq a < b \leq J$ such that $U_a = U_b$. If we then set $x = x_b$ and $r = r_{a+1} + \ldots + r_b$ we obtain Theorem 5.13 as required. □

6. More on the Čech-Stone compactification of a discrete semigroup

We extend the multiplication on a discrete semigroup to its Čech-Stone compactification.

Definition 6.1. Let $(S, \cdot)$ be a semigroup. For each $s \in S$ the left multiplication $\lambda_s : S \to S; x \mapsto sx$ can be considered as a map from S to βS and therefore has a continuous extension $\beta\lambda_s : \beta S \to \beta S$. Now for each $x \in \beta S$ we have a map $\rho_x : S \to \beta X; s \mapsto \beta\lambda_s(x)$. This has a unique continuous extension $\beta\rho_x : \beta S \to \beta S$. For $x, y \in \beta S$ we define $x \cdot y = \beta\rho_y(x)$.

A semigroup $(S, \cdot)$ carrying a topology is *right topological* if for every $t \in S$ the right multiplication $S \to S; s \mapsto s \cdot t$ is continuous.

Lemma 6.2. *For every discrete semigroup* $(S, \cdot)$, $(\beta S, \cdot)$ *is a compact right-topological super-semigroup of* $(S, \cdot)$.

Proof. From the definition of $\cdot$ on βS it follows that $\cdot$ extends the multiplication on S. Also, for each $y \in \beta S$ the right multiplication $x \mapsto xy$ is just $\beta\rho_y$, which is continuous by definition.

For all $x, y, z \in \beta S$ we have

$$(xy)z = \beta\rho_z(xy) = \beta\rho_z(\beta\rho_y(x)) = (\beta\rho_z \circ \beta\rho_y)(x)$$

and $x(yz) = \beta\rho_{yz}(x)$. Hence, in order to show that the multiplication on βS is associative, we have to prove that $\beta\rho_z \circ \beta\rho_y = \beta\rho_{yz}$.

Clearly, $\beta\rho_z \circ \beta\rho_y$ is a continuous function from βS to βS. By the uniqueness of $\beta\rho_{yz}$, it is enough to show that $\beta\rho_z \circ \beta\rho_y$ agrees with ρ_{yz} on S. Let $s \in S$. Then $\rho_{yz}(s) = \beta\lambda_s(yz) = \beta\lambda_s(\beta\rho_z(y))$. On the other hand, $(\beta\rho_z \circ \beta\rho_y)(s) = \beta\rho_z(\beta\lambda_s(y))$. The two functions $\beta\lambda_s \circ \beta\rho_z$ and $\beta\rho_z \circ \beta\lambda_s$ are both continuous functions from βS to βS. In order to show that they are equal, it is enough to show that they agree on the dense subset S of βS.

Let $t \in S$. Then $(\beta\lambda_s \circ \beta\rho_z)(t) = \beta\lambda_s(\beta\lambda_t(z))$ and

$$(\beta\rho_z \circ \beta\lambda_s)(t) = \beta\rho_z(\lambda_s(t)) = \beta\rho_z(st) = \beta\lambda_{st}(z).$$

It remains to show that $\beta\lambda_s \circ \beta\lambda_t$ is the same as $\beta\lambda_{st}$. Again, it is enough to verify this on S.

Let $r \in S$. Then $(\beta\lambda_s \circ \beta\lambda_t)(r) = \beta\lambda_s(tr) = s(tr)$. On the other hand, $\beta\lambda_{st}(r) = (st)r$. Finally, $(st)r = s(tr)$ since $\cdot$ is associative on S. It follows that $\cdot$ is associative on βS. □

Definition 6.3. Let S be a semigroup. A nonempty subset I of S is a *right* (*left*) *ideal* if for all $t \in I$ and all $s \in S$ we have $t \cdot s \in I$ ($s \cdot t \in I$). I is an *ideal* of S if it is both a left and a right ideal. $s \in S$ is *idempotent* if $s \cdot s = s$.

Lemma 6.4. *Let S be a compact right topological semigroup. Then S has a minimal left ideal I.*

Proof. First observe that for all $t \in S$, $St = \{s \cdot t : s \in S\}$ is a closed left ideal of S. The closedness comes from the fact that St is the continuous image of the compact space S under the continuous map right multiplication with t. Now, if T is a left ideal of S, then for every $t \in T$, $St \subseteq T$. If follows that every left ideal of S contains a closed left ideal that is generated by a single element.

An easy application of Zorn's Lemma shows that S has a minimal closed left ideal I. By the previous remark, I is also a minimal among all left ideals. □

Lemma 6.5. *Let S be a compact right topological semigroup. Then S has an idempotent element.*

Proof. Very much like in the previous lemma, we can use Zorn's Lemma to find a minimal closed subsemigroup T of S. Let $s \in T$. Then $Ts \subseteq T$ and Ts is closed. Also, given $t, r \in T$, we have $tsrs = (tsr)s \in Ts$. It follows that Ts is a closed subsemigroup of T and hence $Ts = T$.

Now pick $t \in T$and let $U = \{u \in T : ut = t\}$. It is easily checked that U is a subsemigroup of T. Being the preimage of a singleton under the continuous map right multiplication with t, U is closed. Hence $U = T$. It follows that $ut = t$ for all $u \in T$ and hence $tt = t$. This shows that t is idempotent and $T = \{t\}$. □

This lemma ist quite remarkable: From the compactness of a right topological semigroup we get the existence of elements with particular algebraic properties.

6.1. **The Hales-Jewett theorem and van der Waerden revisited.** We use an abstract theorem of Koppelberg to deduce two classical theorems in Ramsey theory.

Definition 6.6. Let S be a semigroup and T a subsemigroup of S. We call T a *nice* subsemigroup if $R = S \setminus T$ is an ideal of S, i.e., $R \cdot S$ and $S \cdot R$ are subsets of R. Note that T is nice iff for all $x, y \in S$ we have $xy \in T$ iff $x \in T$ and $y \in T$.

A semigroup homomorphism $\sigma : S \to T$ is a *retraction* (from S to T) if $\sigma(t) = t$ for all $t \in T$.

Theorem 6.7 (Koppelberg)**.** *Let S be a semigroup and T a proper nice subsemigroup of S. Let Σ be a finite set of retractions from S to T and let $(B_1, \dots, B_n)$ be a partition of T. Then there are $j \in \{1, \dots, n\}$ and $r \in R = S \setminus T$ such that for all $\sigma \in \Sigma$, $\sigma(r) \in B_j$.*

Proof. It is easily checked that $\hat{T}$ is isomorphic to a βT and therefore a subsemigroup of βS. Also, $\hat{R}$ is equal to $\beta S \setminus \hat{T}$ and is an ideal of βS. Finally, for each $\sigma \in \Sigma$, $\beta\sigma$ is a retraction from βS to $\hat{T}$.

Let L be a minimal left ideal of $\hat{T}$ and let $q \in L$ be an idempotent. Let I be a minimal left ideal in the left ideal $\beta S \cdot q$ of βS and choose an idempotent $i \in I$. Let $p = qi$. Now $p \in I$.

Note that $I \subseteq \hat{R}$ since $\hat{R}$ is an ideal of βS. It follows that $p \in \hat{R}$ and thus $R \in p$. Since $i \in I \subseteq \beta S \cdot q$ and q is an idempotent, $iq = i$. Now $qp = qqi = qi = p$, $pq = qiq = qi = p$ and $pp = qiqi = qii = qi = p$. Hence

$$(*) \qquad p = p^2 = pq = qp.$$

Let $\sigma \in \Sigma$ and $u = \beta\sigma(p)$. Clearly, $u \in \hat{T}$. Also $q \in \hat{T}$. We apply $\beta\sigma$ to equation $(*)$ and obtain

$$u = u^2 = uq = qu.$$

In particular, $u = uq \in L$. Since L is a minimal left ideal of $\hat{T}$, $L = \hat{T} \cdot u$. Hence $q \in \hat{T} \cdot u$. Note that u is an idempotent. It follows that $qu = q$. This shows that $\beta\sigma(p) = q$ for every $\sigma \in \Sigma$.

Recall that q is actually an ultrafilter on S such that $T \in q$. It follows that there is $j \in \{1, \ldots, n\}$ such that $B_j \in q$. For every $\sigma \in \Sigma$ we have $B_j \in q = \beta\sigma(p)$. It follows that $\hat{B}_j$ intersects $\sigma[A]$ for every $A \in p$. In other words, for every $A \in p$, $\sigma[A]$ contains an ultrafilter that contains the set B_j. But $\sigma[A]$ consists of ultrafilters that correspond to elements of S. Identifying these ultrafilters with the corresponding elements of S, we see that for all $A \in p$, $\sigma[A]$ intersects B_j. Hence $\sigma^{-1}[B_j]$ intersects every set $A \in p$. Since p is an ultrafilter, this implies $\sigma^{-1}[B_j] \in p$.

Since $R \in p$, also the set $D = R \cap \bigcap_{\sigma \in \Sigma} \sigma^{-1}[B_j]$ is in p, and hence nonempty. Every $r \in D$ works for the theorem. □

We can now use Koppelberg's theorem to give an alternative proof of van der Waerden's theorem (Theorem 5.5), which we restate for convenience.

Theorem 6.8 (van der Waerden)**.** *Assume $(A_1, \ldots, A_n)$ is a partition of ω into finitely many pieces and $m \in \omega$. Then there are $j \in \{1, \ldots, n\}$ and natural numbers a and $d > 0$ such that*

$$\{a, a+d, a+2d, \ldots, a+md\} \subseteq A_j,$$

i.e., A_j contains an arithmetic progression of length $m+1$.

Proof. Consider the semigroup $S = \omega \times \omega$ and let $T = \omega \times \{0\}$. Then T is a nice subsemigroup of S. For each $j \in \{1, \ldots, n\}$ let $B_j = A_j \times \{0\}$. Now $(B_1, \ldots, B_n)$ is a partition of T. For each $k \leq m$ and all $a, d \in \omega$ let $\sigma_k(a, d) = (a + kd, 0)$. Each σ_k is a retraction from S to T.

Hence, by Koppelberg's theorem there are $j \in \{1, \ldots, r\}$ and $(a, d) \in S \setminus T$ such that for all $k \leq m$, $\sigma_k(a, d) \in B_j$. Now by the definition of σ_k and of B_j, for all $k \leq m$ we have $a + kd \in B_j$. This finishes the proof of the theorem. □

Definition 6.9. Let M be a finite set, the alphabet.

Let x be a variable not in M. A *variable word* over M is a word over $M \cup \{x\}$ with at least one occurrence of x. Given a word w over $M \cup \{x\}$ and $u \in M$, let $w(u)$ denote the word over M obtained by replacing every occurence of x by u.

A *combinatorial line* over M is a set of the form $\{w(u) : u \in M\}$, where w is a variable word over M.

Theorem 6.10 (Hales-Jewett). *Let M be a finite alphabet and let $(A_1, \dots, A_n)$ be a partition of M^*. Then there is $j \in \{1, \dots, n\}$ such that A_j includes a combinatorial line.*

Proof. Let S be the semigroup $(M \cup \{x\})^*$ with the concatenation of words as multiplication. Let $T = M^*$. Then T is a nice subsemigroup of S. For each $u \in M$ and $w \in S$ let $\sigma_u(w) = w(u)$. Each σ_u is a retraction from S to T.

Hence there are some $w \in S \setminus T$, i.e., a variable word over M, and some $j \in \{1, \dots, n\}$, such that for all $u \in M$, $w(u) \in A_j$. In other words, A_j includes the combinatorial line generated by w. □

Exercise 6.11. Derive van der Waerden's theorem directly from the Hales-Jewett theorem.

We derive a finite version of the Hales-Jewett theorem from the infinite version above. For $m \in \omega$, $M^{\leq m}$ denotes the set of all words over M of length at most m. M^m is the set of words of length m.

Theorem 6.12 (Hales-Jewett, finite version). *Let M be a finite alphabet. For every $n \in \omega$ there is $m > 0$ such that whenever $C_1, \dots, C_n$ is a partition of M^m, then there is $j \in \{1, \dots, n\}$ such that C_j includes a combinatorial line.*

We derive the theorem from the following lemma.

Lemma 6.13. *Let M and n be as in Theorem 6.12. There is $m > 0$ such that whenever $C_1, \dots, C_n$ is a partition of $M^{\leq m}$, then there is $j \in \{1, \dots, n\}$ such that C_j includes a combinatorial line.*

Proof. Suppose there is no such m. Consider the collection T of all n-tuples $(C_1, \dots, C_n)$ such that for some $m \in \omega$, $C_1, \dots, C_n$ is a partition of $M^{\leq m}$ such that no C_j includes a combinatorial line. Note that for technical reasons we also consider partitions of M^0. For $A, B \in T$, $A = (A_1, \dots, A_n)$, $B = (B_1, \dots, B_n)$, let $A \sqsubset B$ if for all $j \in \{1, \dots, n\}$, $A_j \subsetneq B_j$.

Clearly, T is a tree. By our assumption, T is infinite. Moreover, whenever $(A_1, \dots, A_n) \in T$ is a partition of $M^{\leq m}$ and $k \leq m$, then $(A_1 \cap M^{\leq k}, \dots, A_n \cap M^{\leq k}) \in T$. It follows that each level of T consists of partitions of $M^{\leq m}$ for a fixed $m > 0$. Since for each m the set $M^{\leq m}$

is finite, each level of T is finite. Hence, by König's Lemma, T has an infinite branch $\mathcal{B}$. For each $j \in \{1, \dots, n\}$ let

$$C_j = \bigcup\{A : \exists (B_1, \dots, B_n) \in \mathcal{B}(A = B_j)\}.$$

Now $(C_1, \dots, C_n)$ is a partition of M^* such that no C_j includes a combinatorial line. This contradicts Theorem 6.10. □

Proof of Theorem 6.12. Let $m > 0$ be a minimal witness of Lemma 6.13. Let $C_1, \dots, C_n$ be a partition of M^m. By the minimality of m, there is a partition $A_1, \dots, A_n$ of $M^{\leq m-1}$ such that no A_j includes a combinatorial line. Now $A_1 \cup C_1, \dots, A_n \cup C_n$ is a partition of $M^{\leq m}$. By the choice of m, there is $j \in \{1, \dots, n\}$ such that $C_j \cup A_j$ includes a combinatorial line. Since A_j does not include a combinatorial line, the combinatorial line included in $A_j \cup C_j$ consists of words of length m. Hence A_j includes a combinatorial line. □

7. Universal minimal flows

Definition 7.1. Let G be a topological group and let X be a minimal G-flow. X is *universal* if every minimal G-flow Y is a factor of X.

Theorem 7.2. *Let G be a discrete group. Then there is a universal minimal G-flow X.*

Proof. Extending the left multiplication with $g \in G$ to βG, we obtain an action of G on βG. Let X be a minimal subflow of βG. Let Y be any minimal G-flow. Fix $y \in Y$ and consider the continuous map $f : G \to Y; g \mapsto gy$. This map is clearly equivariant. It extends to an equivariant continuous map $\beta f : \beta G \to Y$. Since Y is minimal, the restriction $\beta f \restriction X$ is onto Y, showing that Y is a factor of X. □

We extend this construction of a universal minimal flow to groups that are metrizable, but not necessarily discrete.

Definition 7.3. Let G be a topological group. A G-ambit is a G-flow X with a distinguished point $x \in X$ whose orbit is dense in X.

Theorem 7.4. *For every (Hausdorff) topological group G there is a universal minimal G-flow.*

Proof. Let G_d denote the group G with the discrete topology. It is clear that every G-flow is also a G_d-flow. Given a G-ambit (X, x) we consider the equivariant map $f : G_d \to X; g \mapsto gx$ and its unique continuous extension $\beta f : \beta G_d \to X$. This shows that every G-ambit, considered as a G_d-flow, is a factor of βG_d.

Up to isomorphism we can realize every factor of βG_d as a quotient of βG_d by a suitable equivalence relation. This shows that there is a set $\mathcal{A}$ of G-ambits such that every G-ambit is isomorphic to one in $\mathcal{A}$. Given a G-ambit A, let x_A denote the distinguished point of A.

Let $Y = \prod_{A \in \mathcal{A}} A$ equipped with the coordinate wise action of G. Let $y = (x_A)_{A \in \mathcal{A}}$. Finally, let $S(G)$ be the orbit closure of y in Y, with the distinguished point y. $S(G)$ is the greatest G-ambit in the sense that every G-ambit is a factor of $S(G)$, namely just the projection of $S(G)$ to a suitable coordinate.

Let X be a minimal subflow of $S(G)$. Then X is a universal minimal G-flow. Namely, let Z be any minimal G-flow. Choose a point $z \in Z$. The G-ambit (Z, z) is a factor of $S(G)$ by some factor map $f : S(G) \to Z$. The restriction $f \upharpoonright X$ is still onto Z since Z is minimal. This show the universality of X. $\square$

Exercise 7.5. Let A and B be G-ambits that have all G-ambits as factors, by factor maps that map the distinguished point to the distinguished point. Show that A and B are isomorphic. In particular, $S(G)$ is uniquely determined by its universal property.

Definition 7.6. Let G be a topological group. A *fixed point* of a G-flow X is a point x that is not moved by any $g \in G$. G is *extremely amenable* if its universal minimal flow is a singleton or equivalently, every continuous action of G on a compact space has a fixed point.

We finish this section by proving criteria for the existence of fixed points in flows.

Lemma 7.7. *Let G be a topological group acting continuously on a compact space X. Then the following are equivalent:*

(1) *The G-flow X has a fixed point.*
(2) *For every $n \geq 1$, continuous function $f : X \to \mathbb{R}^n$, $\varepsilon > 0$ and finite set $F \subseteq G$, there is $x \in X$ such that for all $g \in F$, $|f(x) - f(gx))| < \varepsilon$.*

Proof. (1)⇒(2): Any fixed point of X works for x.

(2)⇒(1): We use a compactness argument to obtain a fixed point in X. Given F, ε, and f as in (2), let

$$A_{f,F,\varepsilon} = \{x \in X : \forall g \in F(|f(x) - f(gx)| \leq \varepsilon\}.$$

By the continuity of f, $A_{f,F,\varepsilon}$ is closed and by (2) it is nonempty. It is easily checked that the intersection of finitely many sets of the form $A_{f,F,\varepsilon}$ is nonempty. By the compactness of X, the intersection of all sets of the form is nonempty. Clearly, this intersection consists of fixed points of X. □

Theorem 7.8 (Kechris, Pestov, Todorcevic)**.** *Let $\mathcal{S}_\infty$ be the group of all permutations of $\mathbb{N}$ equipped with the topology inherited from the product topology on $\mathbb{N}^\mathbb{N}$, the* topology of pointwise convergence. *If $G \subseteq \mathcal{S}_\infty$ is a closed subgroup, then the following are equivalent:*

1. *G is extremely amenable.*
2. *For every open subgroup V of G, every coloring $c : G/V \to \{1, \ldots, k\}$ of the set of left cosets hV of V, and every finite set $A \subseteq G/V$ of left cosets of V, there is $g \in G$ and $i \in \{1, \ldots, k\}$ such that c is constant on $g \cdot A = \{ghV : hV \in A\}$.*

Proof. (1)⇒(2): Let V, k, c, and A be as in (2). We consider the set G/V of left cosets of V in G with the natural G-action $(g, hV) \mapsto (gh)V$. Let $Y = \{1, \ldots, k\}^{G/V}$ with the shift action defined by $(gy)(hV) = y(g^{-1}hV)$ for all $g, h \in G$ and $y \in Y$. Since $\{1, \ldots, k\}$ is finite, Y is compact. The action of G on Y is continuous. This is not completely trivial and depends on the fact that V is open in G (Exercise).

Clearly, $c : G/v \to \{1, \ldots, k\}$ is an element of Y. Let X be the orbit closure of c in Y. By (1), there is a fixed point $d \in X$. Since the action of G on G/V is *transitive*, i.e., G moves every point in G/V to every other point, the fixed point d has to be constant with some value $i \in \{1, \ldots, k\}$. Since d is in the orbit closure of c, there is $g \in G$ such that $g^{-1}c \restriction A = d \restriction A$. This shows that on A, $g^{-1}c$ has the constant value i. I.e., for each $hV \in A$ we have $c(ghV) = i$. This shows that c is constant on gA.

(2)⇒(1): The statement (2) is equivalent to the corresponding statement about the set $V \setminus G$ of right cosets Vh of V in G on which G acts by $g(Vh) = Vhg^{-1}$. By Lemma 7.7 it is enough to show that for every G-flow X, continuous $f : X \to \mathbb{R}^n$, $\varepsilon > 0$, and finite $F \subseteq G$ there is

$x \in X$ such that for all $h \in F$, $|f(x) - f(hx)| \leq \varepsilon$. We fix such X, f, ε, and F. Assume that $1_G \in F$.

For each $x \in X$, we choose an open neighborhood $U_x \subseteq X$ of x and $V_x \subseteq G$ of the identity of G such that for all $y \in U_x$ and $g \in V_x$ we have $|f(x) - f(gx)| \leq \varepsilon/3$. Finitely many of the U_x cover X. Let V be the intersection of the corresponding V_x. Then V is an open subset of G containing the identity such that for all $g \in V$ and all $x \in X$ we have $|f(x) - f(gx)| \leq \varepsilon/3$. Since G is a closed subgroup of $\mathcal{S}_\infty$, V has a basic open subset that is a neighborhood of the identity. I.e., there is a finite set $S \subseteq \mathbb{N}$ such that for all $g \in G$ with $g \restriction S = \mathrm{id}_S$, $g \in V$. Clearly, the set of all $g \in G$ that are the identity on S is an open subgroup of G. Hence we may assume that V is already equal to this open subgroup of G.

We partition the compact set $f[X] \subseteq \mathbb{R}^n$ into finitely many sets A_1,...,A_k of diameter $\leq \varepsilon/3$. Fix $x_0 \in X$ and let $U_i = \{g \in G : f(gx_0) \in A_i\}$ for each $i \in \{1, \dots, k\}$. Now let $V_i = VU_i$, a union of right cosets of V. We consider V_i as a subset of $V \setminus G$ in the natural way. Observe that $V \setminus G = \bigcup_{i=1}^k V_i$. Hence there is $c : V \setminus G \to \{1, \dots, k\}$ such that for each i, $V_i \subseteq c^{-1}(i)$.

By (2), there are $i \in \{1, \dots, k\}$ and $g \in G$ with $Fg \subseteq V_i = VU_i$. We show that $x = gx_0$ works. Let $h \in F$. Let $v \in V$ be such that $vhg \in U_i$. Now $f(vhgx_0) = f(vhx) \in A_i$. Since $|f(vhx) - f(hx)| \leq \varepsilon/3$, $f(hx)$ is no further from the set A_i than $\varepsilon/3$. Since $1_G \in F$, $|f(x) - f(hx)| \leq \varepsilon$. □

8. Fraïsse theory, Ramsey theory, and the theorem of Kechris, Pestov, and Todorcevic

Theorem 8.1. *Let G be a closed subgroup of $\mathbb{S}_\infty$. Then the following are equivalent:*

1. *G is extremely amenable.*
2. *G is the automorphism group of a structure A which is the Fraïsse limit of a Fraïsse order class with the Ramsey propoerty.*

It follows from this theorem that the automorphism groups of $(\mathbb{Q}, <)$ and the ordered Random graph are extremely amenable.

Department of Mathematics, University of Hamburg
E-mail address: stefan.geschke@uni-hamburg.de