If $n=1$, shift invariance implies that only one letter from the alphabet is being used, hence X is a singleton.
So suppose that $n>1$. By minimality, we have that $p_X(n-1)=p_X(n)=n$. Let $\mathcal{L}_n=\{w_1,\dots,w_n\}\subseteq A^n$ be the words of length n appearing in X. Let

$$L: \mathcal{L}_n \longrightarrow \mathcal{L}_{n-1}$$
$$w_i=(a_0,\dots,a_{n-1}) \longmapsto (a_1,\dots,a_{n-1}),$$
$$R: \mathcal{L}_n \longrightarrow \mathcal{L}_{n-1}$$
$$w_i=(a_0,\dots,a_{n-1}) \longmapsto (a_0,\dots,a_{n-2}).$$

Claim:
| L, R are bijections

We only prove it for L.

Surjectivity: Let $v\in\mathcal{L}_{n-1}$ and let $x\in X$ s.t. v appears in x. By shift-invariance we can assume w.l.o.g. that $x_1=v_1,\dots,x_{n-1}=v_{n-1}$.
Then $(x_0,v_1,v_2,\dots,v_{n-1})\in\mathcal{L}_n$ and this is mapped to v under L.

Injectivity: Since $p_X(n)=|\mathcal{L}_n|=|\mathcal{L}_{n-1}|$ as argued above, surjectivity implies that L is a bijection.

Using the claim, we know that the composition

$$\mathcal{L}_{n-1} \xrightarrow{R^{-1}} \mathcal{L}_n \xrightarrow{L} \mathcal{L}_{n-1}$$
$$v \longmapsto \underbrace{(v,*)}_{\text{uniquely determined by } v} \longmapsto (v_2,\dots,v_{n-1},*)$$

is a bijection. By shift-invariance of X, it follows that every element in X is completely determined by the first $n-1$ symbols, therefore $|X|<\infty$. ∎

End of lecture 9

Definition.
| A Sturmian shift space is a shift X with $p_X(n)=n+1$ for all n (i.e., just not finite).

Remark: One can similarly define the complexity of a sequence by letting X be the shift space defined by a single sequence. On calls a sequence Sturmian if this associated shift space is Sturmian.
Similarly, one defines Sturmian one-sided shifts.
The Fibonacci numbers form a Sturmian sequence (Maybe exercise).

Constructions of Sturmian shifts

Example 0: $C_0 \longrightarrow \infty$

Example 1: Let $\alpha\in\mathbb{R}-\mathbb{Q}$. Define $I_1=[0,1-\alpha)\equiv[0,-\alpha) \bmod 1$ and
$I_2=[1-\alpha,1)\equiv[-\alpha,0) \bmod 1$.

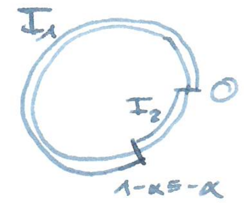

We define the shift space by defining the language. To this end, we call a word w in $\{1,2\}^n$ allowed if

$$I_w := I_{w_1}\cap R_\alpha^{-1} I_{w_2}\cap\dots\cap R_\alpha^{-(n-1)} I_{w_n}\neq\emptyset,$$

i.e., there exists $x\in\mathbb{T}$ s.t.

$$\forall\, 0\le j<n \quad R_\alpha^j x\in I_{w_{j+1}}.$$

Claim:
| There are exactly $n+1$ allowed words of length n.

To this end it is useful to introduce the notation

$$P_k := \{R_\alpha^{-k} I_1, R_\alpha^{-k} I_2\},$$

i.e., for all $k\in\mathbb{Z}$ the set $P_k\subseteq 2^{\mathbb{T}}$ is a partition of $\mathbb{T}$. We define for any two partitions $\mathcal{A},\mathcal{B}$ of $\mathbb{T}$ the common refinement

$$\mathcal{A}\vee\mathcal{B} := \{A\cap B : A\in\mathcal{A}, B\in\mathcal{B}\}.$$

Then $\mathcal{A}\vee\mathcal{B}$ is again a partition of $\mathbb{T}$.
Define inductively $Q_0 := P_0$ and for $k\in\mathbb{N}$

$$Q_k := Q_{k-1}\vee P_k.$$

Claim:
| $\{I_w : w\in\{1,2\}^n \text{ is allowed}\} = Q_{n-1}$

This is pretty much the definition.

Claim:
| $\{w\in\{1,2\}^n \text{ is allowed}\} \longrightarrow Q_{n-1}$
| $w \longmapsto I_w$
| is injective

Later

Claim:

right-closed, left-open counterclockwise

$\mathcal{Q}_n$ consists of the half-open intervals defined by

$$\{0, -\alpha \bmod 1, \dots, -(n+1)\alpha \bmod 1\}$$

Proof: $n=0$: ✓

$n \mapsto n+1$: $\{R_\alpha^{-(n+1)} I_1, R_\alpha^{-(n+1)} I_2\} = \mathcal{P}_{n+1} = \{[-(n+1)\alpha, -(n+2)\alpha), [-(n+2)\alpha, -(n+1)\alpha)\}$ and $-(n+2)\alpha \bmod 1$ is the only end-point not yet contained among the endpoints of $\mathcal{Q}_n$. Thus the refinement

$$\mathcal{Q}_{n+1} = \mathcal{Q}_n \vee \mathcal{P}_{n+1}$$

splits exactly one of the intervals at $-(n+2)\alpha \bmod 1$.

Corollary

$|\mathcal{Q}_n| = n+2$

This is immediate (again) from the irrationality of α.

It remains to prove that $\omega \mapsto I_\omega$ is injective.

Let $n \in \mathbb{N}$ and $\omega, \omega' \in \{1,2\}^n$ allowed. If $n=1$, then it is clear that

$$I_\omega = I_{\omega'} \iff \omega = \omega'.$$

So suppose $n \in \mathbb{N}$ and we know injectivity for n. Let $\omega, \omega' \in \{1,2\}^{n+1}$ allowed and suppose $I_\omega = I_{\omega'}$ but $\omega \neq \omega'$. Note that $I_\omega \cap I_{\omega'} \neq \emptyset$

Denote by $p: \{1,2\}^{n+1} \longrightarrow \{1,2\}^n$, $\omega \longmapsto {}_p\omega$ the map which forgets the last component. Then

$$I_\omega = I_{{}_p\omega} \cap R_\alpha^{-n} I_{\omega_{n+1}},$$

$$I_{\omega'} = I_{{}_p\omega'} \cap R_\alpha^{-n} I_{\omega'_{n+1}}.$$

We know that $I_{{}_p\omega} = [a_\omega, b_\omega)$, $I_{{}_p\omega'} = [a_{\omega'}, b_{\omega'})$, where $a_\omega, a_{\omega'}, b_\omega, b_{\omega'}$ are contained in $\{0, \dots, -n\alpha \bmod 1\}$, and that

$$R_\alpha^{-n} I_{\omega_{n+1}}, R_\alpha^{-n} I_{\omega'_{n+1}} \in \{[-n\alpha, -(n+1)\alpha), [-(n+1)\alpha, -n\alpha)\}.$$

Hence we can do a case distinction. If $I_{{}_p\omega} = I_{{}_p\omega'}$, then ${}_p\omega = {}_p\omega'$, thus $\omega_{n+1} \neq \omega'_{n+1}$ and, since $R_\alpha^{-n} I_{\omega_{n+1}}$ and $R_\alpha^{-n} I_{\omega'_{n+1}}$ are disjoint, it follows that

$$I_\omega \cap I_{\omega'} = \emptyset,$$

which is absurd.

Hence we know that $I_{{}_p\omega} \neq I_{{}_p\omega'}$, but this implies that $I_{{}_p\omega} \cap I_{{}_p\omega'} = \emptyset$,

hence

$$I_\omega \cap I_{\omega'} \subseteq I_{{}_p\omega} \cap I_{{}_p\omega'} = \emptyset \quad ⚡$$

This proves that the set $\mathcal{L}_n$ of allowed words has cardinality $n+1$. We let X be the shift space given by the set

$$X = \{x \in \{1,2\}^{\mathbb{Z}} : \text{every subword of } x \text{ is allowed}\}.$$

III – Measure preserving dynamics & Ergodic Theory

In what follows, we will always assume that $(X,\mathcal{B},\mu)$ is a probability space, i.e., X is a set, $\mathcal{B}$ is a σ-algebra on X, and $\mu:\mathcal{B}\longrightarrow[0,1]$ is a probability measure.

In many examples, X is also a compact metric space and $\mathcal{B}$ is the Borel σ-algebra generated by the open sets.

Definition:

Let $(X,\mathcal{B},\mu)$ a probability space and let $(Y,\mathcal{C})$ a measurable space. If $T:X\longrightarrow Y$ measurable, we denote by $T_*\mu$ the push-forward, i.e.,

$$\forall C\in\mathcal{C}\quad T_*\mu(C)=\mu(T^{-1}C).$$

A measurable map $T:X\longrightarrow X$ is <u>measure preserving</u> if $T_*\mu=\mu$. In this case, we call μ a <u>T-invariant</u> measure and the system $(X,\mathcal{B},\mu,T)$ is a (probability) measure preserving system.

Lemma:

Let $(X,\mathcal{B},\mu)$ and $(Y,\mathcal{C},\sigma)$ probability spaces. The following are equivalent.

(i) $\sigma=T_*\mu$

(ii) $\forall f\geq 0$ measurable (from Y to $[0,\infty]$), we have

$$\int_Y f\,d\sigma=\int_X f\circ T\,d\mu.$$

(iii) $\forall f\in L^1_\sigma(Y)$

$$\int_Y f\,d\sigma=\int_X f\circ T\,d\mu.$$

Proof: Exercise (measure-theoretic induction).

Examples: (i) Let $(X,\mathcal{B},\mu)=(\mathbb{T},\mathcal{B}(\mathbb{T}),\mathrm{Leb})$ and consider $T:X\longrightarrow X$ given by the rotation R_α (for any $\alpha\in\mathbb{R}$). Then $(X,\mathcal{B},\mu,T)$ is a (probability) measure preserving system.

(ii) Let $p\in\mathbb{Z}$, then the multiplication by p defines a pmps $(\mathbb{T},\mathcal{B},\mathrm{Leb},T_p)$ (p non-zero).

(iii) Let G a compact metric group. Then G admits a (Borel) Haar measure m_G characterized uniquely by the following properties:

(H1) $\forall g\in G\ \forall B\in\mathcal{B}(G)$

$$m_G(gB)=m_G(B),$$

(H2) $\forall \mathcal{O}\subset G$ open non-empty

$$m_G(\mathcal{O})>0,$$

(H3) $m_G(G)=1$. (this makes it unique)

Let $T:G\longrightarrow G$ be a surjective endomorphism. Then $(G,\mathcal{B}(G),m_G,T)$ is a probability measure preserving system. (Exercise)

(iv) Let $A=\{0,\dots,p-1\}$ and $\underline{p}=(p_0,\dots,p_{p-1})$ a (fully supported) probability vector, i.e., $\forall 0\leq k<p$ $p_k\in(0,1]$ and $\sum_{k=0}^{p-1}p_k=1$.

Then $\underline{p}$ defines a probability measure $\mu_{\underline{p}}$ on $\mathcal{B}(A)=2^A$ by

$$\forall B\in\mathcal{B}(A)\quad \mu_{\underline{p}}(B)=\sum_{k\in B}p_k.$$

By the Kolmogorov extension theorem, there exists a unique probability measure $\mu_{\underline{p}}$ on $(A^{\mathbb{Z}},\mathcal{B}(A^{\mathbb{Z}}))$ s.t.

$\forall \ell\in\mathbb{N}\ \forall j_1,\dots,j_\ell\in\mathbb{Z}$ pairwise distinct

$$\mu_{\underline{p}}\Big(\underbrace{\{(x_n)_{n\in\mathbb{Z}}\in A^{\mathbb{Z}}:\forall 1\leq i\leq \ell\ x_{j_i}=a_i\}}_{{}_{j_1}[a_1]\cap\dots\cap{}_{j_\ell}[a_\ell]}\Big)=\prod_{i=1}^{\ell}\mu_{\underline{p}}(\{a_i\})=\prod_{i=1}^{\ell}p_{a_i}.$$

The latter determines $\mu_{\underline{p}}$ uniquely. Note that

$$\mu_{\underline{p}}(\sigma^{-1}({}_{j_1}[a_1]\cap\dots\cap{}_{j_\ell}[a_\ell]))=\mu_{\underline{p}}(\sigma^{-1}({}_{j_1}[a_1])\cap\dots\cap\sigma^{-1}({}_{j_\ell}[a_\ell]))$$

$$=\mu_{\underline{p}}({}_{j_1+1}[a_1]\cap\dots\cap{}_{j_\ell+1}[a_\ell])$$

$$=\prod_{i=1}^{\ell}p_{a_i}=\mu_{\underline{p}}({}_{j_1}[a_1]\cap\dots\cap{}_{j_\ell}[a_\ell])$$

Hence $(A^{\mathbb{Z}},\mathcal{B}(A^{\mathbb{Z}}),\mu_{\underline{p}},\sigma)$ is a probability measure preserving system.

One similarly obtains a pmps $(A^{\mathbb{N}},\mathcal{B}(A^{\mathbb{N}}),\mu^+_{\underline{p}},\sigma)$. These are called the two- and one-sided Bernoulli shifts on p symbols (for probability vector $\underline{p}$).

End of lecture 10

(v) More generally, suppose that for all $\ell\in\mathbb{N}\cup\{0\}$ we are given a map

$$p^{(\ell)}:A^{\ell+1}\longrightarrow[0,1]$$

such that (i) $\sum_{a\in A}p^{(0)}(a)=1$

(ii) $\forall a_0,\dots,a_\ell\in A\quad p^{(\ell)}(a_0,\dots,a_\ell)=\sum_{a\in A}p^{(\ell+1)}(a_0,\dots,a_\ell,a)$.