7 Chaos

Throughout these notes the notation I will denotes a closed interval in $\mathbb{R}$, and the map $F : I \to \mathbb{R}$ will be assumed to be continuous even if this is not stated.

Proofs included here but not given in lectures are included in brackets: [** $\cdots$ **] since they are certainly not examinable.

7.1 Introduction

In this chapter we will investigate the generation of complicated dynamics in the simplest possible setting: discrete time maps of the interval.

We take the point of view that 'complicated dynamics' means 'orbit complexity' i.e. the guaranteed generation of large numbers of periodic orbits (recall we use the terminology N*-cycles* for orbits of least period N). We will begin by discussing two particular motivating examples. Along with the investigation of the examples we will develop the definitions we need, and prove various straightforward results.

Then we will define a more abstract class of dynamical systems: the shift map acting on spaces of sequences of symbols. Such 'symbolic dynamics' turns out to provide a very good model for 'complicated dynamics': in particular we can count the numbers of N-cycles that are guaranteed to arise. We can then quite easily relate the symbolic dynamics results to investigate any given continuous map of an interval. We can prove the existence of N-cycles in the map using the various properties of the symbolic dynamical systems. This is a powerful and much more general idea than we will have time to explore in the course. Our reasoning will be largely 'topological' in nature, and in the nicest cases we will look at we can actually demonstrate a topological conjugacy between the symbolic dynamics and iteration of the map. Even in cases where the map dynamics are not topologically conjugate, we can relate the dynamics to symbol sequences with important and useful consequences.

We begin with a motivating example which gives a first glimpse of the idea of symbolic dynamics, and enables us to define two much-loved properties of maps displaying complicated dynamics.

Example: The Sawtooth Map

This is the map $F : [0, 1] \to [0, 1]$ defined by

$$x_{n+1} \quad = \quad 2x_n \mod 1$$

which appears to have interesting dynamics, and many periodic points: for example $x = 0$ is a fixed point, $\{\frac{1}{3}, \frac{2}{3}\}$ is a 2-cycle, $\{\frac{1}{5}, \frac{2}{5}, \frac{4}{5}, \frac{3}{5}\}$ is a 4-cycle. All these seem to be unstable since the gradient $F'(x)$ is always 2.

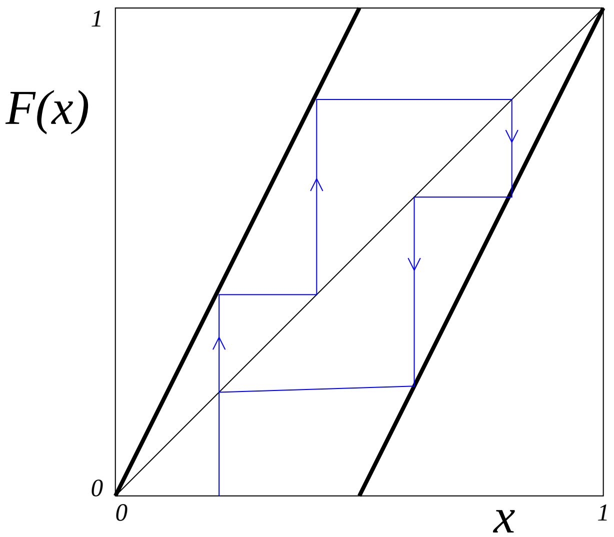

The Sawtooth Map $x_{n+1} = 2x_n \bmod 1$. The 4-cycle is indicated by the blue arrows.

A very nice way to represent the dynamics is to write points $x \in [0,1]$ in terms of their base-2 (binary) expansions, e.g.:

$$\frac{1}{3} \quad = \quad 0 \cdot \overset{\frac{1}{2}}{0}\,\overset{\frac{1}{4}}{1}\,\overset{\frac{1}{8}}{0}\,\overset{\frac{1}{16}}{1}\,\overset{\frac{1}{32}}{0}\,\overset{\frac{1}{64}}{1}\ \cdots = \frac{1}{4} + \frac{1}{16} + \frac{1}{64} = \frac{\frac{1}{4}}{1-\frac{1}{4}} \ = \frac{1}{3}.$$

Then the $x \to 2x \bmod 1$ map corresponds to shifting the sequence to the left and discarding the leading '1' (if any). Checking this explicitly:

$$F\left(\frac{1}{3}\right) = \frac{2}{3} \quad = \quad 0 \cdot \overset{\frac{1}{2}}{1}\,\overset{\frac{1}{4}}{0}\,\overset{\frac{1}{8}}{1}\,\overset{\frac{1}{16}}{0}\,\overset{\frac{1}{32}}{1}\,\overset{\frac{1}{64}}{0}\ \cdots = \frac{1}{2} + \frac{1}{8} + \frac{1}{32} = \frac{\frac{1}{2}}{1-\frac{1}{4}} \ = \frac{2}{3}.$$

After a little reflection, it looks as if any periodic binary expansion $0 \cdot a_0 a_1 a_2 \cdots$ corresponds to a periodic point $x \in [0,1]$ for F.

Note: binary expansions are not unique for points of the form $p/2^n$ since $0 \cdot 0111111 \cdots$ codes for the same point ($x = \frac{1}{2}$) as $0 \cdot 10000 \cdots$, etc.

The Sawtooth map has many nice properties of which two are often taken as defining 'chaos' (but not here!). We will now define these two properties in the general setting common to the whole of this chapter: let $F : I \to \mathbb{R}$ be a continuous map of a a closed bounded interval $I \subset \mathbb{R}$ into $\mathbb{R}$. Let Λ be an invariant subset of I.

Definition (SDIC): A map $F : I \to \mathbb{R}$ has <u>sensitive dependence on initial conditions</u> (SDIC) on an invariant subset $\Lambda \subseteq I$ if $\exists\ \delta > 0$ such that for any $x \in \Lambda$ and $\varepsilon > 0$ there exists $y \in \Lambda$ and $n > 0$ such that $|x - y| < \varepsilon$ and $|F^n(x) - F^n(y)| > \delta$.

i.e. near any x there is always some point that separates to at least a distance δ away. This is a formalised version of the notion that even the smallest errors in initial conditions (or in a numerical integration scheme) inevitably grow to become as large as the true value of the state, meaing that it becomes impossible to predict the future behaiour of the system, even though it remains entirely deterministic. This is often referred to as the 'butterfly effect' and makes prediction of, for example, the weather, extremely uncertain once one looks 5-10 days ahead.

Definition (TT): A map $F : I \to \mathbb{R}$ is <u>topologically transitive</u> (abbreviated to TT) on Λ if for all nbhds U, V which intersect Λ, there exists $n > 0$ such that $F^n(U) \cap V \neq \emptyset$.

i.e. even the smallest open sets eventually intersect and 'mix together' under iteration of the map. Hence Λ cannot be decomposed into smaller disjoint open invariant sets.

These two properties are independent: for example the irrational rotation $\theta_{n+1} = \theta_n + 2\pi\omega$ for ω irrational is TT on the circle $S^1 \simeq [0, 2\pi]/\sim$, but points do not move away from (or towards) each other, so this rotation map does not have SDIC. On the other hand, the map $x_{n+1} = 2x_n$ on $\mathbb{R}$ displays SDIC (the distance between points doubles with each iteration) but is not TT.

Remark: An equivalent property to TT is the existence of a dense orbit for F, i.e. a trajectory that comes arbitrarily close to every point in Λ. Since it is usually much easier to establish the existence of a dense orbit we will do this to prove TT.

Returning to the Sawtooth Map we now prove it has SDIC and is TT on $\Lambda = [0, 1]$.

Proof (SDIC): Set $\delta = \frac{1}{4}$. Given $\varepsilon > 0$ and $x \in [0, 1]$, pick n such that $2^{-n-1} < \varepsilon$. Then construct the binary expansion for x, say $0 \cdot a_0 a_1 a_2 \cdots a_{n-1} a_n a_{n+1} \cdots$. Take y to be the point with symbol sequence $0 \cdot a_0 a_1 a_2 \cdots a_{n-1} \bar{a}_n a_{n+1} \cdots$ where $\bar{a}_n = 1 - a_n$ means change the symbol a_n from a '0' to a '1' or vice-versa as appropriate. Then we see that $|x - y| = 2^{-n-1} < \varepsilon$ but $|F^n(x) - F^n(y)| = \frac{1}{2} > \delta$. □

Proof (TT): We construct a point x which has a dense orbit, i.e. comes arbitrarily close to any given point $y \in [0, 1]$. Let x be the point given by the binary expansion

$$0 \cdot \underbrace{0\ 1}\ \underbrace{00\ 01\ 10\ 11}\ \underbrace{000\ 001\ 010\ 011\ 100\ 101\ 110\ 111}\ \cdots$$

taking all blocks of lengths $1, 2, 3, \ldots$ in order. Then for any point y we can compute the corresponding symbol sequence $0 \cdot a_0 a_1 a_2 \cdots a_{n-1} a_n a_{n+1} \cdots$. Then for any $n > 0$ there exists a $k > 0$ such that the binary expansion of $F^k(x)$ agrees with the expansion of y on at least the first n places, implying $|F^k(x) - y| < 2^{-n}$, i.e. the forward orbit of x comes arbitrarily close to any point y and this is a dense orbit. □

Remark: As we have seen, the nicest features of a dynamical system are preserved under topological conjugacy, enabling us to understand more complicated problems in terms of simpler ones (e.g. nonlinear flows near a fixed point in terms of the linearised flow). It turns out that SDIC is not preserved under topological conjugacy (see the reelvant starred question on example sheet 4) which in part motivates our search for a better defining characteristic of 'chaos'.

7.2 Symbolic Dynamics

In this section we will define a new class of dynamical systems and prove that they have nice properties. We will then set up a topological conjugacy between this nice dynamical system and the logistic map when μ is large enough. This enables us to understand completely the dynamics of the logistic map when μ is large enough.

Notice that, in our investigation of the sawtooth map we used binary sequences and the action of F was equivalent to shifting the binary sequence along one place. This motivates the definition:

Definition (sequence space on N symbols): Let

$$\Sigma_N = \{ \mathbf{a} = (a_0 a_1 a_2 \cdots) : a_i \in \{0, 1, \ldots, N-1\} \ \forall i \geq 0 \ \}$$

be the sequence space on N symbols; the collection of ∞ sequences of symbols, each drawn from the set $\{0, 1, \ldots, N-1\}$.

Points in Σ_N are symbol sequences. There is a natural distance measure (metric) on Σ_N: two sequences are close together if they agree on a long initial segment. We define the distance measure

$$d(\mathbf{a},\mathbf{b}) = \sum_{n=0}^{\infty} \frac{\gamma(a_n,b_n)}{3^n},$$

where $\gamma(p,q)=0$ if $p=q$ and $\gamma(p,q)=1$ if $p\neq q$.

Suppose, for example, that $\mathbf{a},\mathbf{b}\in\Sigma_N$ and $a_i=b_i$ for $0\le i<m$ and then $a_m\neq b_m$. Then we can compute directly that

$$3^{-m}\le d(\mathbf{a},\mathbf{b}) \le \frac{3}{2}3^{-m}.$$

The natural evolution operator on the state space is the shift map σ:

Definition (shift map): The shift map $\sigma:\Sigma_N\to\Sigma_N$ acts by $\sigma(a_0a_1a_2\cdots)=(a_1a_2a_3\cdots)$.

Properties of $\sigma:\Sigma_N\to\Sigma_N$

1. σ is continuous.

   Proof:

$$\begin{aligned} d(\mathbf{a},\mathbf{b}) &= |a_0-b_0|+\frac{1}{3}d(\sigma(\mathbf{a}),\sigma(\mathbf{b})) \\ \Rightarrow d(\sigma(\mathbf{a}),\sigma(\mathbf{b})) &= 3d(\mathbf{a},\mathbf{b}) \end{aligned}$$

   so as $\mathbf{a}\to\mathbf{b}$ we can guarantee $\sigma(\mathbf{a})\to\sigma(\mathbf{b})$. □

2. σ^k has N^k fixed points.

   Proof: $\sigma^k(\mathbf{a})=\mathbf{a} \iff a_{k+j}=a_j\ \forall j\ge 0$ so we need only choose the initial block $(a_0\cdots a_{k-1})$ to determine $\mathbf{a}$. There are clearly N^k distinct blocks of length k. □

3. The set of periodic points of σ, $\mathrm{Per}(\sigma)$, is dense in Σ_N (i.e. periodic points exist arbitrarily close to any given symbol sequence).

   Proof: Given $\mathbf{a}\in\Sigma_N$ and $\varepsilon>0$, take n such that $\frac{3}{2}3^{-n}<\varepsilon$. Then let $\mathbf{b}=(a_0a_1\cdots a_{n-1}a_0a_1\cdots a_{n-1}a_0\cdots)$. Then we see that $d(\mathbf{a},\mathbf{b})<\varepsilon$ and $\mathbf{b}$ is clearly a periodic symbol sequence. □

4. $\sigma:\Sigma_N\to\Sigma_N$ is TT because there exists a point $\mathbf{a}$ with a dense orbit.

   Proof: Let $\mathbf{a}$ be the symbol sequence given by listing all blocks of length 1, then all blocks of length 2, and so on, e.g. for $N=2$:

$$\mathbf{a} = (\underbrace{0\ 1}\ \underbrace{00\ 01\ 10\ 11}\ \underbrace{000\ 001\ 010\ 011\ 100\ 101\ 110\ 111}\ \cdots).$$

   then, given a point $\mathbf{b}\in\Sigma_N$ and $\varepsilon>0$ there exists $n>0$ such that $\sigma^n(\mathbf{a})$ agrees with $\mathbf{b}$ in the first k places for any k. Taking k large enough that $\frac{3}{2}3^{-k}<\varepsilon$ we then have that $d(\sigma^n(\mathbf{a}),\mathbf{b})\le\frac{3}{2}3^{-k}<\varepsilon$. □

5. $\sigma:\Sigma_N\to\Sigma_N$ has SDIC.

   Proof: Take $\delta=1$. Given $\mathbf{a}\in\Sigma_N$ and $\varepsilon>0$ there exists $\mathbf{b}\in\Sigma_N$ and $n>0$ such that $3^{-n}<d(\mathbf{a},\mathbf{b})<\varepsilon$, i.e. $\mathbf{a}$ and $\mathbf{b}$ differ first in the n^{th} place in the symbol sequence. Then $d(\sigma^n(\mathbf{a}),\sigma^n(\mathbf{b}))>1$ since the symbol sequence $\sigma^n(\mathbf{a})$ differs from $\sigma^n(\mathbf{b})$ in the first place in the sequence. □

Conjugacy between the logistic map and Σ_2

Having proved various nice properties of $\sigma : \Sigma_N \to \Sigma_N$ (a rather abstract dynamical system), we now make a direct link between the abstract and the 'real' dynamics of a continuous map F. The idea is to be able to translate our results about σ acting on Σ_N back into results about F acting on I.

Recall that a map $h : \Lambda \to Y$ is surjective (or 'onto') if for all $y \in Y$ there exists a point $x \in \Lambda$ such that $h(x) = y$. Also recall that $h : \Lambda \to Y$ is injective (or '1–1') if $h(x_1) = h(x_2) \Rightarrow x_1 = x_2$.

We will also (for the proof of surjectivity) need the following standard result:

Lemma 0.1 (Cantor Intersection Theorem) *The intersection $S_\infty = \cap_{i=0}^{\infty} S_i$ of an infinite sequence $S_0 \supseteq S_1 \supseteq S_2 \supseteq \cdots$ of nested non-empty closed bounded subsets $S_i \subset \mathbb{R}^n$ is non-empty.*

Definition (semiconjugacy): Let $F : I \to \mathbb{R}$ be a cts map of the interval, and let $\Lambda \subseteq I$ be an invariant set. Let $G : Y \to Y$ be a cts map on a (metric) space Y. If there exists a cts surjection $h : \Lambda \to Y$ such that $h \circ F = G \circ h$ then we say that F is semiconjugate to G via h.

Definition (conjugacy): If in addition h is injective (i.e. $1 - 1$) and h^{-1} is cts, so that h is a homeomorphism, then F is conjugate to G via h.

Theorem 1 *The logistic map $F(x) = \mu x(1 - x)$, when $\mu > 2 + \sqrt{5}$, has an invariant set $\Lambda \subset [0, 1]$ on which $F|_\Lambda$ is conjugate to $\sigma|_{\Sigma_2}$.*

The invariant set Λ is the collection of points that, when iterated under F, remain in $[0, 1]$ for all time, i.e. $\Lambda = \{x \in I : F^n(x) \in I \ \forall \ n \geq 0\}$.

Before we present a proof of the theorem, we give a key ingredient in the form of an easily-proved lemma that shows where the value $2+\sqrt{5}$ comes from: for μ greater than this value the magnitude of the slope of the logistic map is greater than unity within the parts of $[0, 1]$ that are mapped inside $[0, 1]$. The proof of the lemma is left as a straightforward exercise.

Lemma 1.1 *If $\mu > 2 + \sqrt{5}$ then there exists $\lambda > 1$ such that $|F'(x)| > \lambda$ for all $x \in I \cap F^{-1}(I)$.*

Proof of Theorem

To prove the theorem we need to

- construct a map $h : \Lambda \to \Sigma_2$
- show h is injective $(1 - 1)$
- show h is surjective (for every element $\mathbf{a}$ of Σ_2 there exists a point $x \in \Lambda$ such that $h(x) = \mathbf{a}$)
- show h is continuous with a continuous inverse

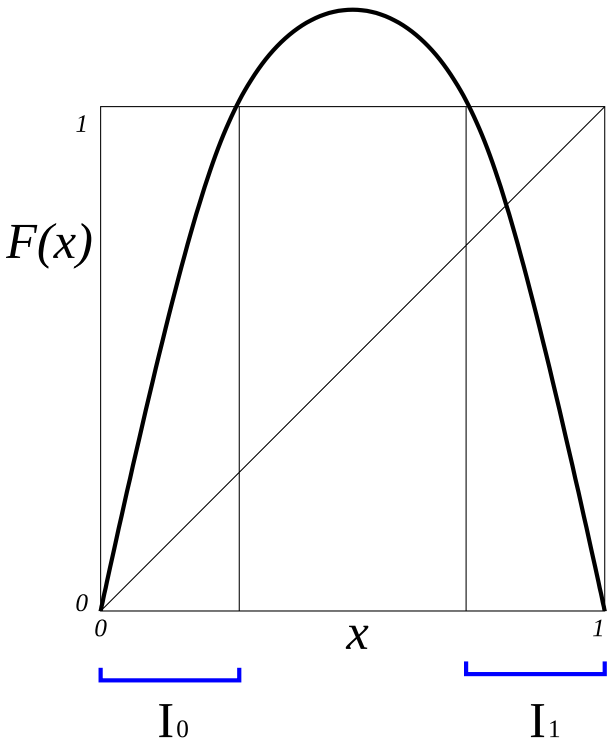

The Logistic Map $x_{n+1} = \mu x_n(1 - x_n)$ for $\mu > 2 + \sqrt{5}$.

Proof (construction): We can check that the set $I \cap F^{-1}(I)$ is a disjoint union of two closed intervals $I_0 = [0, x_-]$ and $I_1 = [x_+, 1]$ where $x_\pm = (1 \pm \sqrt{1 - 4/\mu})/2$ are the points at which $F(x) = 1$. In the open interval (x_-, x_+), F maps points above $x = 1$ so they cannot be part of the invariant set (in fact, iterates move rapidly off to $-\infty$). We define the symbol sequence $\mathbf{a}$ which will correspond to a point $x \in \Lambda$ by setting

$$\begin{aligned} a_j &= 0 \text{ if } F^j(x) \in I_0 \\ a_j &= 1 \text{ if } F^j(x) \in I_1 \end{aligned}$$

Then the map $h : \Lambda \to \Sigma_2$ is defined by setting $h(x) = \mathbf{a}$. □

We now prove each of the desired properties of h in turn.

[** **Proof (h is injective):** Suppose there exists $x, y \in \Lambda$ with $x \neq y$ and $h(x) = h(y) = \mathbf{a}$. Then (from the definition of $\mathbf{a}$) we see that $F^j(x)$ and $F^j(y)$ are always on the same side of $x = \frac{1}{2}$ as each other (since they are always in the same interval I_0 or I_1). This implies $|F^j(x) - F^j(y)| < \frac{1}{2}\ \forall\ j$.

But, since $|F'(x)| > \lambda > 1$ we have $|F^j(x) - F^j(y)| > \lambda^j |x - y|$ for all j, and the right-hand side eventually becomes greater than $\frac{1}{2}$ so there is a contradiction here unless $|x - y| = 0$ which means than in fact we must have $x = y$. □ **]

[** **Proof (h is surjective):** Given a symbol sequence $\mathbf{a} = (a_0 a_1 a_2 \cdots)$ we need to show there exists $x \in \Lambda$ such that $F^j(x) \in I_{a_j}\ \forall\ j$.

Let $J \subset I$ be a closed interval, then we will use the notation $F^{-1}(J) = \{x : F(x) \in J\}$ to denote the preimage of J. From the graph of F, $F^{-1}(J)$ is the disjoint union of a pair of closed subintervals, one in each of I_0 and I_1.

Define the set

$$\begin{aligned} I_{a_0 a_1 \cdots a_n} &= \{x : x \in I_{a_0},\ F(x) \in I_{a_1}, \ldots, F^n(x) \in I_{a_n}\} \\ &= I_{a_0} \cap F^{-1}(I_{a_1}) \cap F^{-2}(I_{a_2}) \cap \cdots \cap F^{-n}(I_{a_n}). \end{aligned}$$

Now, just by combining the definitions we see that, also

$$I_{a_0 a_1 \cdots a_n} \quad = \quad I_{a_0} \cap F^{-1}(I_{a_1 a_2 \cdots a_n})$$

because these are both exactly the sets of points for which $F^j(x) \in I_{a_j}$ for $0 \leq j \leq n$.

You should check that the statement $I_{a_0 a_1} = I_{a_0} \cap F^{-1}(I_{a_1})$ agrees with the labelling indicated on the sketch of F given in lectures, and possibly appearing earlier on in these notes.

Now we use induction to assert that, if $I_{a_1 \cdots a_n}$ is a nonempty closed interval then $F^{-1}(I_{a_1 \cdots a_n})$ consists of a pair of closed intervals which implies $I_{a_0} \cap F^{-1}(I_{a_1 \cdots a_n})$ is exactly one closed interval. Hence, inductively, $I_{a_0 a_1 \cdots a_n} = I_{a_0} \cap F^{-1}(I_{a_1 \cdots a_n})$ is a closed interval for all n.

Moreover,

$$I_{a_0 \cdots a_n} \quad = \quad I_{a_0 \cdots a_{n-1}} \cap F^{-n}(I_{a_n}) \subset I_{a_0 \cdots a_{n-1}}$$

so the intervals $I_{a_0 \cdots a_n}$ are nested closed intervals and the Cantor Intersection Theorem implies that

$$I_{\mathbf{a}} \quad \stackrel{\text{def}}{=} \quad \cap_{n=0}^{\infty} I_{a_0 \cdots a_n}$$

is not empty. So there exists $x \in I_{\mathbf{a}}$ with the property that $F^j(x) \in I_{a_j}$ for all j. Since h is $1-1$ we must have that $I_{\mathbf{a}}$ contains exactly one point. We can also see this since the length of the intervals $I_{a_0 \cdots a_n}$ tends to zero at least as fast as λ^{-n} as $n \to \infty$ because of the condition that $|F'(x)| > \lambda > 1$ everywhere. $\square$ **]

[** **Proof (h and h^{-1} are continuous):** For any pair of points $x, y \in \Lambda$ we have established that there exist unique symbol sequences $\mathbf{a} = h(x)$ and $\mathbf{b} = h(y)$. Now, $x \to y$, $\iff$ (since $|F'(x)| > \lambda > 1$ everyhwhere) the iterates $F^j(x)$ and $F^j(y)$ remain on the same side of $\frac{1}{2}$ for longer $\iff$ $\mathbf{a}$ and $\mathbf{b}$ agree on a longer initial segment of symbols $\iff$ $d(\mathbf{a}, \mathbf{b}) \to 0$. $\square$ **]

Remarks:

1. In showing that $h \circ F|_{\Lambda} = \sigma \circ h$ we now know that the dynamics of $F|_{\Lambda}$ has all the properties of $\sigma|_{\Sigma_2}$, including dense periodic points, TT and SDIC.

2. The invariant set Λ is an example of a Cantor set: it is closed, contains no intervals (a set with this property is called 'totally disconnected') and every point in Λ is a limit point of a sequence of points in Λ (a set with this property is called 'perfect'). The classic Cantor set is the 'middle-third' construction which arises directly in a closed related problem, and is explored on example sheet 4.

3. Conjugacy between $F|_{\Lambda}$ and $\sigma|_{\Sigma_2}$ holds in fact for any $\mu > 4$, but the proof required more careful estimates of the rates of separation of nearby points so that we can guarantee that h is $1-1$.

7.3 Subshifts of Finite Type (SSFT)

A very useful refinement of the sequence space Σ_N is the case where there are rules about which symbols allowed to follow each other. Allowed symbol sequences are encoded by an $N \times N$ transition matrix A.

Definition (transition matrix): an $N \times N$ matrix A is a transition matrix if $A_{ij} = 1$ whenever symbol j is allowed to follow symbol i and $A_{ij} = 0$ whenever symbol j cannot follow symbol i.

Note that the matrix indices now range over $0 \le i,j \le N-1$ rather than $1 \le i,j \le N$ - this will cause no notational or other difficulties in what follows but it is worth pointing out.

Remark: This idea of transitions between states is very similar to that developed in the theory of discrete-time Markov chains.

Example: Let $N = 2$. If only the sequences 01, 10 and 11 are allowed and 00 is not, then

$$A \quad = \quad \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}.$$

The set of allowed sequences forms a closed subset of Σ_N which we denote $\Sigma_{N,A}$:

$$\Sigma_{N,A} \quad = \quad \{\mathbf{a} \in \Sigma_N \ : \ A_{a_n a_{n+1}} = 1 \ \forall n \ge 0 \ \}.$$

Clearly $\Sigma_{N,A}$ is invariant under the shift σ.

Definition (SSFT): the action of $\sigma : \Sigma_{N,A} \to \Sigma_{N,A}$ defines a dynamical system called a subshift of finite type (SSFT). We write σ_A as shorthand for $\sigma|_{\Sigma_{N,A}}$.

Properties of σ_A

1. The number $N_{ij}^{(n)}$ of allowed sequences $ia_1a_2\cdots a_{n-1}j$ of length $n+1$ from symbol i to symbol j is given by $(A^n)_{ij}$.

   Proof: the product $A_{ia_1}A_{a_1a_2}\cdots A_{a_{n-1}j} = 1$ if and only if $ia_1a_2\cdots a_{n-1}j$ is an allowed sequence. So

   $$N_{ij}^{(n)} \quad = \quad \sum_{a_1,\ldots,a_{n-1}} A_{ia_1}A_{a_1a_2}\cdots A_{a_{n-1}j} = (A^n)_{ij}$$

   □

2. The number of (not necessarily least) period-n orbits P_n is given by $\mathrm{tr}(A^n)$.

   Proof: with the above notation, period-n orbits are exactly those sequences where $i = j$, so

   $$P_n = \text{no. of period}-n \text{ orbits} \quad = \quad \sum_i N_{ii}^{(n)} = \sum_i (A^n)_{ii} = \mathrm{tr}(A^n).$$

   □

3. Let N_q be the number of q-cycles (i.i. periodic orbits of least period q). Then $P_n = \sum_{q|n} qN_q$. Proof: each q-cycle contributes q points to every set P_n for which n contains q as a factor.

   Remark: We can compute $\mathrm{tr}(A^n)$ from the recurrence relation given by the Cayley–Hamilton Theorem.

   Example: $A = \left(\begin{smallmatrix} 0 & 1 \\ 1 & 1 \end{smallmatrix}\right)$ has characteristic polynomial $P(\lambda) = \lambda^2 - \lambda - 1$ which implies $A^2 - A - I = 0$.

   $$\Rightarrow \mathrm{tr}(A^{m+2}) \quad = \quad \mathrm{tr}(A^{m+1}) + \mathrm{tr}(A^m),$$

and hence we can compute

$$\begin{array}{lllll} P_1 = 1 & P_2 = 3 & P_3 = 4 & P_4 = 7 & P_5 = 11 \\ N_1 = 1 & N_2 = 1 & N_3 = 1 & N_4 = 1 & N_5 = 2 \end{array}$$

So there are exactly two distinct 5-cycles for σ_A, i.e. exactly 2 distinct least period-5 allowed symbol sequences. They can be easily seen to be $(01011\cdots)$ and $(01111\cdots)$.

Definition (irreducible): the transition matrix A is <u>irreducible</u> if, for all i, j there exists $n \geq 0$ such that $(A^n)_{ij} = 1$ i.e. for all i, j there exists an allowed symbol sequence of some length that starts with i and ends with j.

4. If A is irreducible then σ_A is TT.

   [** **Proof:** We construct a dense orbit. As before, we list all allowed symbol sequences of length 1, then of length 2, and so on, putting 'transition' sequences between each if necessary. The existence of such 'transition' sequences is guaranteed always, by the irreducibility of A. $\square$ **]

   Definition (non-trivial): the transition matrix A is <u>non-trivial</u> if, for some i there exists j_1 and j_2 $(j_1 \neq j_2)$ such that the sequences ij_1 and ij_2 are allowed.

   Note: this excludes permutation matrices such as

$$\begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix}$$

   which are 'not non-trivial' (maybe even 'trivial'?).

5. If A is irreducible and non-trivial then $\sigma_A : \Sigma_{N,A} \to \Sigma_{N,A}$ has SDIC.

   [** **Proof:** Given a sequence $\mathbf{a} = a_0 a_1 \cdots \in \Sigma_{N,A}$ and $\varepsilon > 0$ we choose $M \geq 0$ such that $\frac{1}{2}3^{-M} < \varepsilon$. Construct $\mathbf{b} = a_0 a_1 \cdots a_M b_{M+1} b_{M+2} \cdots b_K$ where $a_M b_{M+1} \cdots b_K$ is an allowed sequence from a_M to $b_K = i$ (this exists by irreducibility) and ij_1 ij_2 are allowed sequences (by non-triviality). Then $\mathbf{b}$ differs from $\mathbf{a}$ either at some place $n < K+1$ or at place $n = K+1$ since we can choose $b_{K+1} = j_1$ or j_2 whichever is not equal to a_{K+1}.

   So, by construction, $d(\mathbf{a}, \mathbf{b}) \leq \frac{1}{2}3^{-M}$ and $d(\sigma_A^n(\mathbf{a}), \sigma_A^n(\mathbf{b})) \geq 1$ and so σ_A has SDIC. $\square$ **]

7.4 Continuous Maps $F : I \to \mathbb{R}$

In this subsection we use our results on the abstract SSFT dynamical systems to prove results about the dynamics of maps on the real line.

Definition (invariant set): Let $\Lambda = \{x \in I : \ F^n(x) \in I \ \forall \ n \geq 0\}$ be the invariant set for F.

Let $\{I_i : \ i = 0, \ldots, N-1\}$ be a collection of **disjoint** closed bounded intervals in $\mathbb{R}$. We now consider the dynamics of F restricted to $\cup_{i=0}^{N-1} I_i$.

Definition (F-covering): we say I_i F-covers I_j (written $I_i \to I_j$) if $F(I_i) \supseteq I_j$.

There is a directed graph Γ (on N vertices) indicating the F-covering relations.

Example:

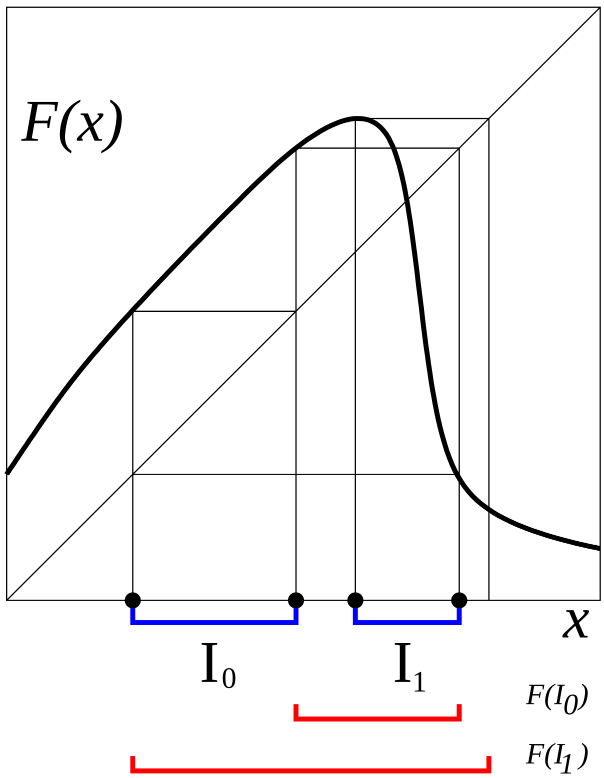

A map $F : I \to \mathbb{R}$ containing 2 disjoint intervals I_0 and I_1 which have the F-covering relations given by the graph:

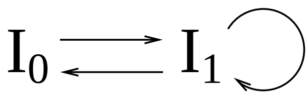

So I_0 F-covers I_1 and I_1 F-covers I_0 and I_1.

There is a natural correspondance between the F-covering relations and the transition matrix A defined by $A_{ij} = 1$ if $I_i \to I_j$ and $A_{ij} = 0$ if not.

Remark: if $I_0 \xrightarrow{F} I_1 \xrightarrow{F} I_2$ then $I_0 \xrightarrow{F^2} I_2$ by continuity of F. This construction is called the induced graph for F^2.

Lemma 1.2 *Let I, J be closed intervals. If I F-covers J then there exists a closed subinterval $K \subseteq I$ such that $F(K) = J$. i.e.*

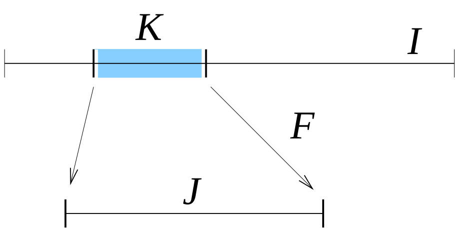

A map $F : I \to \mathbb{R}$ F-covers an interval J, and there is a closed interval K such that $F(K) = J$.

Proof: Let $J = [p, q]$ be a closed interval. Then $F^{-1}(p)$ and $F^{-1}(q)$ are closed and nonempty. Discarding parts of $F^{-1}(p)$ and$F^{-1}(q)$ if they are disconnected, we may choose points $u \in F^{-1}(p)$ and $v \in F^{-1}(q)$ such that $(u, v) \cap \left(F^{-1}(p) \cup F^{-1}(q)\right) = \emptyset$:

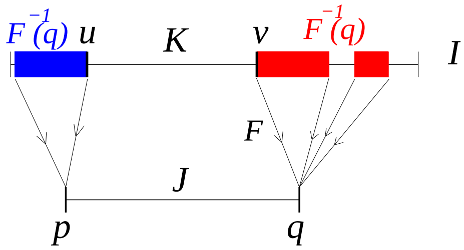

If the picture looks like this, we set $u = \sup\{F^{-1}(p)\}$, $v = \inf\{F^{-1}(q)\}$. There is another case where all of $F^{-1}(q)$ lies to the left of all of $F^{-1}(p)$; in this case we take $u = \sup\{F^{-1}(q)\}$ and $v = \inf\{F^{-1}(p)\}$ instead, but it is essentially equivalent.

Then set $K = [u, v]$, a closed interval, and check, using the Intermediate Value Theorem, that the map is surjective:

For any c: $a < c < b$ let $g(x) = F(x) - c$. Then

$$\begin{aligned} g(u) &= F(u) - c = a - c < 0 \\ g(v) &= F(v) - c = b - c > 0 \end{aligned}$$

and g is continuous. Hence there exists $\hat{x} \in K$ such that $g(\hat{x}) = 0$ i.e. $F(\hat{x}) = c$, so c has a preimage in K. □

Theorem 2 $F|_\Lambda$ *is semiconjugate to* σ_A*.*

Remark: from this theorem it follows that for any symbol sequence $\mathbf{a}$ in $\Sigma_{N,A}$ there is a corresponding point $x \in \Lambda$ whose iterates follow the symbol sequence $\mathbf{a}$. For example, $F|_\Lambda$ has as many periodic orbits as σ_A does. Of course, it may be that a given F has more periodic orbits than this - we are concerned here with those orbits that are guaranteed by the semiconjugacy.

Proof: define the semiconjugacy $h : \Lambda \to \Sigma_{N,A}$ by $x \mapsto \mathbf{a}$ where $\mathbf{a}$ is the symbol sequence of $\{F^n(x)\}_{n\geq 0}$ such that $F^n(x) \in I_{a_n}$ for all $n \geq 0$. We need to show (i) h is continuous and (ii) h is surjective.

[** (i) h is continuous because, given $\varepsilon > 0$ we can choose an M such that $\frac{1}{2}3^{-M} < \varepsilon$, i.e. $d(h(x), h(y)) < \varepsilon \iff$ the symbol sequences $h(x)$ and $h(y)$ agree in the first M places. This can be guaranteed by taking $|x - y|$ small enough since the intervals I_i are **disjoint** and closed, so there exists $\delta > 0$ such that $|x - y| < \delta$ implies $F^n(x), F^n(y) \in I_{a_n}$ for $0 \leq n \leq M$, i.e. x and y remain in the same intervals as each other for (at least) the first M iterates. **]

(ii) h is surjective because, given an allowed sequence $\mathbf{a} = a_0 a_1 a_2 \cdots \in \Sigma_{N,A}$ we claim that there exists a nested sequence of closed intervals

$$I_{a_0} \supseteq I_{a_0 a_1} \supseteq I_{a_0 a_1 a_2} \supseteq \cdots \supseteq I_{a_0 a_1 \cdots a_M}$$

such that $x \in I_{a_0 a_1 \cdots a_M}$ implies $F^n(x) \in I_{a_n}$ for all $0 \leq n \leq M$. This is true by repeated use of the lemma above.

Looking at the first few iterates of the map, by the lemma we can assert the existence of a closed interval $I_{a_0 a_1} \subseteq I_{a_0}$ such that $F(I_{a_0 a_1}) = I_{a_1}$ and also that there exists a closed interval $I_{a_1 a_2} \subseteq I_{a_1}$ such that $F(I_{a_1 a_2}) = I_{a_2}$. We may now apply the lemma again to assert the existence of a smaller closed subset $I_{a_0 a_1 a_2} \subseteq I_{a_0 a_1} \subseteq I_{a_0}$ and so on, inductively, to obtain $I_{a_0 a_1 \cdots a_M}$. This is illustrated in the figure below:

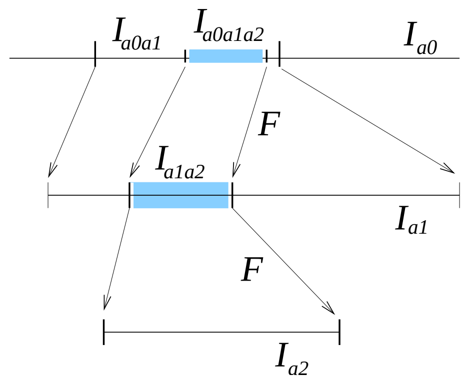

Nested intervals $I_{a_0a_1a_2} \subseteq I_{a_0a_1} \subseteq I_{a_0}$ and their images under F.

Now we apply the Cantor Intersection Theorem to assert that

$$S_{\mathbf{a}} \stackrel{\text{def}}{=} \cap_{M=0}^{\infty} I_{a_0a_1\cdots a_M}$$

is non-empty, so there exists a point $x \in \Lambda$ such that $h(x) = \mathbf{a}$, i.e. h is surjective. Notice that the iterates of x must lie in exactly the sequence of intervals $\{I_{a_n}\}$. $\square$

Remarks:

1. a crucial difference from the example in theorem 1 is that we might not have $\mathrm{diam}(S_{\mathbf{a}}) \to 0$ for all $\mathbf{a}$, so that it might happen that $S_{\mathbf{a}}$ contains at least one interval of points x in which case h could not a a conjugacy.

2. $F|_{\Lambda}$ semiconjugate to σ_A implies that $F|_{\Lambda}$ has a periodic orbit $\{x_0, x_1, \ldots, x_{N-1}\}$ corresponding to every periodic symbol sequence $(a_0a_1, \cdots a_{N-1}a_0a_1\cdots)$ in $\Sigma_{N,A}$. Moreover $x_j \in I_{a_j}$ for all $0 \le j < N$ and hence this periodic orbit corresponds exactly to a closed path $I_{a_0} \to I_{a_1} \to \cdots \to I_{a_{N-1}} \to I_{a_0}$ in the graph Γ.

3. $F|_{\Lambda}$ semiconjugate to σ_A with A irreducible and nontrivial **does not imply** that $F|_{\Lambda}$ has SDIC or TT. This is because F could have intervals on which it is non-expanding, see the example below. If, however, we have additional information that shows F is suitably 'expanding' then we can show, in a similar fashion to the proof of theorem 1, that h is a conjugacy which would imply that $F|_{\Lambda}$ has SDIC and is TT.

Example: A map $F(x)$ for which $\Lambda = [0, 1]$ is an invariant set, and for which the map $h : \Lambda \to \Sigma_2$ is surjective but which does not have SDIC or TT.

Define $\tilde{F}(x)$ to be the piecewise-linear map

$$\tilde{F}(x) = \begin{cases} x + \frac{1}{2} & \text{if } 0 \le x \le \frac{1}{2}, \\ \frac{3}{2} - x & \text{if } \frac{1}{2} \le x \le 1 - \frac{1}{K}, \\ \left(\frac{K}{2} + 1\right)(1 - x) & \text{if } 1 - \frac{1}{K} \le x. \le 1 \end{cases}$$

Fix $K > 4$, e.g. for concreteness think about $K = 5$. The graph of $\tilde{F}$ is as in the figure below:

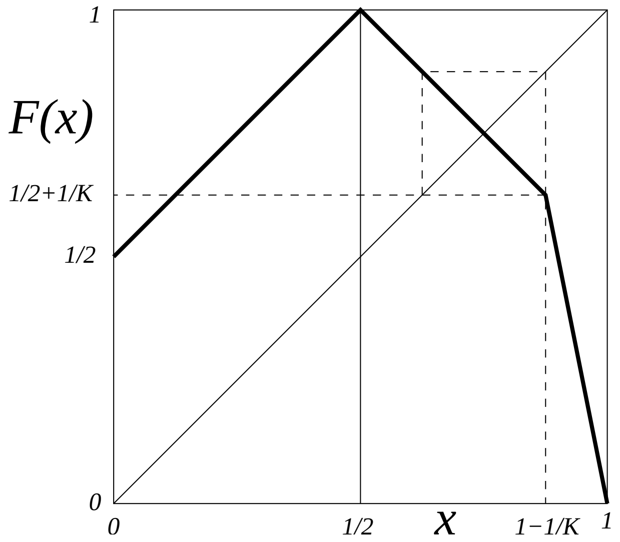

Graph of the piecewise-linear function $\tilde{F}(x)$.

Then, all points in the interval $\frac{1}{2}+\frac{1}{K} \le x \le 1-\frac{1}{K}$ lie on (nonhyperbolic) 2-cycles (except $x=\frac{3}{4}$ which is clearly a fixed point!), so these points do not move apart from, or towards, each other (so $\tilde{F}|_\Lambda$ does not have SDIC) and the 2-cycles do not 'mix' under iteration, so there is no TT in this part of $[0,1]$ and hence $\tilde{F}|_\Lambda$ is not TT.

It looks like the preimages of the interval $\frac{1}{2}+\frac{1}{K} \le x \le 1-\frac{1}{K}$ are dense in $[0,1]$, so almost all initial conditions end up on a 2-cycle, as the following figure shows:

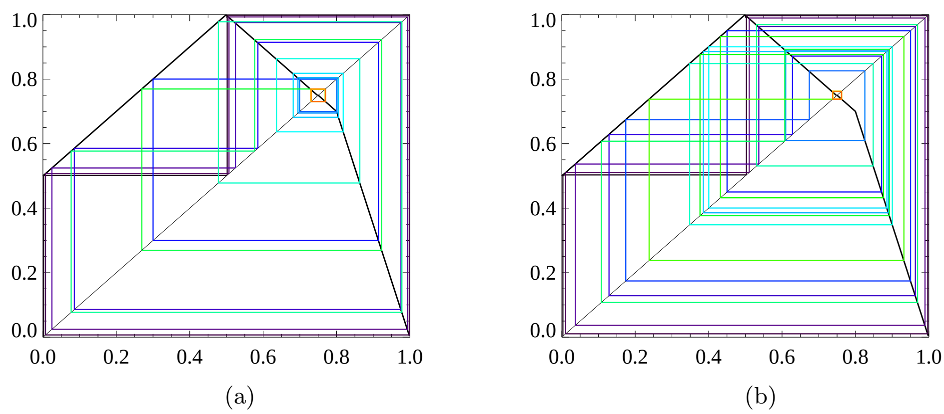

The first 48 (or so) iterates of the map $\tilde{F}$, taking $K=5$, starting from (a) $x_0=0.002$ and (b) $x_0=0.003$. The colour of the iterates x_n darkens as n increases. Eventually, in both cases, the iterates settle to a (nonhyperbolic) 2-cycle near $x=\frac{3}{4}$.

Note, however, that there exists a smaller invariant subset $\tilde{\Lambda} \subset \Lambda$ on which F does have SDIC and TT (remove all the preimages of the set of nonhyperbolic 2-cycles from $[0,1]$ and we are left with a Cantor set of N-cycles and their preimages).

Note also that h is not continuous here so the map is not a semiconjugacy. This arises from the fact that the intervals $[0,\frac{1}{2}]$ and $[\frac{1}{2},1]$ are not disjoint as we assumed the collection $\{I_i\}$ was at the start of this section.

7.5 Horseshoes and N-cycles

As we have already hinted, chaotic dynamics is better described in terms of complicated orbit structure rather than the properties of SDIC and TT. Indeed, since SDIC is not preserved under conjugacy, it is not as robust as we would like. A good definition of chaos can be constructed from a fundamental topological property.

Definition (horseshoe): $F : I \to \mathbb{R}$ has a <u>horseshoe</u> if there exists a closed interval $J \subseteq I$ and closed subintervals $I_0, I_1 \subset J$ with disjoint interiors such that $F(I_0) = F(I_1) = J$.

Remark: This implies that the graph Γ contains the subgraph:

$$\begin{tikzcd} I_0 \arrow[loop left] \arrow[r, shift left] & I_1 \arrow[l, shift left] \arrow[loop right] \end{tikzcd}$$

Definition (chaotic): F is <u>chaotic</u> if F^n has a horseshoe for some $n \geq 1$.

Period 3 Implies Chaos

This, at first sight very surprising, result is the title of a paper by T.-Y. Li and J.A. Yorke (American Mathematical Monthly **82**, pp985–992, 1975).

Theorem 3 *Let $F : I \to \mathbb{R}$ be a continuous map on a closed bounded interval $I \subset \mathbb{R}$. If F has a 3-cycle then*

- *F^2 has a horseshoe and hence F is chaotic,*
- *F has an N-cycle for all $N \geq 1$.*

Proof: let $x_0 < x_1 < x_2$be the 3-cycle, with $x_{i+1} = F(x_i)$ (i is taken mod 3). [Note that the only alternative ordering to this is $x_2 < x_1 < x_0$ in which case we consider the map $G(x) = -F(-x)$ which is clearly conjugate to F].

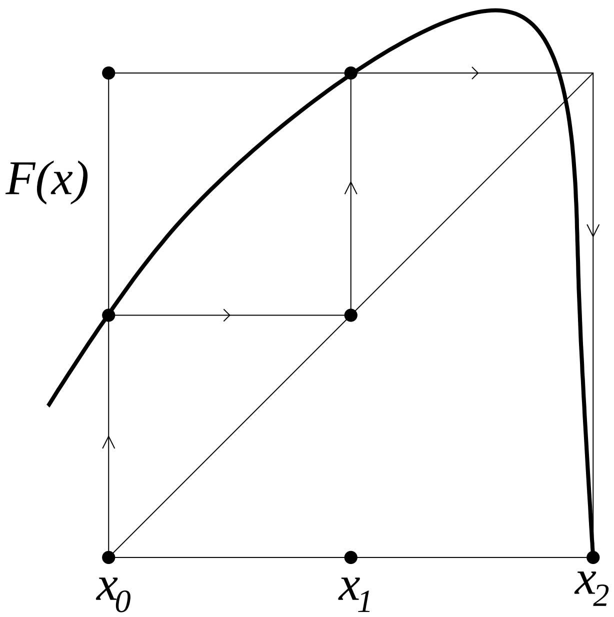

An example of a continuous map F that has a 3-cycle.

Let $I_0 = [x_0, x_1]$ and $I_1 = [x_1, x_2]$, then $F(I_0) \supseteq I_1$ and $F(I_1) \supseteq I_0 \cup I_1$ so we have (at least) the F-covering relations indicated by the graph:

$$\begin{tikzcd} I_0 \arrow[r, shift left] & I_1 \arrow[l, shift left] \arrow[loop right] \end{tikzcd}$$

Hence we have (at least) the graph

$$\circlearrowleft\; I_0 \rightleftarrows I_1 \;\circlearrowright$$

of F^2-covering relations. So F^2 has a horseshoe and F is chaotic.

Moreover, the symbol sequence

$$\mathbf{a}^{(N)} \quad = \quad \underbrace{01111\cdots1}\underbrace{01111\cdots1}\underbrace{01111\cdots1}0\cdots$$

is allowed, for any $N \geq 2$ and so there exists a point $x^{(N)}$ which corresponds to the symbol sequence $\mathbf{a}^{(N)}$ and which lies on an N-cycle, for any $N \geq 2$. Notice that, because I_0 is visited only once per period of the period-N symbol sequence, the orbit must have least period N.

Finally, for $N = 1$ we observe that the allowed sequence $\mathbf{a}^{(1)} = 111\cdots$ implies the existence of a fixed point in I_1. □

In fact, this result of Li & Yorke is only a special case of a much more general result which was, in fact, proved earlier in time (around 1964).

Theorem 4 (Sharkovsky's Theorem) *Consider the ordering of the integers define by*

$$\begin{aligned}
1 \lhd 2 \lhd 4 \lhd 2^3 \lhd \cdots \lhd 2^n \lhd 2^{n+1} \lhd \cdots \\
\cdots \lhd 2^{n+1}.9 \lhd 2^{n+1}.7 \lhd 2^{n+1}.5 \lhd 2^{n+1}.3 \lhd \cdots \\
\cdots \lhd 2^n.9 \lhd 2^n.7 \lhd 2^n.5 \lhd 2^n.3 \lhd \cdots \\
\cdots \lhd 9 \lhd 7 \lhd 5 \lhd 3.
\end{aligned}$$

Let $F : I \to \mathbb{R}$ be a continuous map of the interval. If F has an N-cycle than F has a k-cycle for all $k \lhd N$ in the above ordering.

Proof: not in course, see Glendinning, pages 329–334. □

Remarks: Sharkovsky's theorem implies the following:

1. Existence of a 3-cycle implies the existence of an N-cycle for all $N \geq 1$ (i.e. the Li & Yorke result).

2. Existence of a 4-cycle implies the existence of a 1-cycle (fixed point) and a 2-cycle but maybe no more (see example sheet 4).

3. Let m be an odd integer. The existence of a $2^n.m$-cycle implies the existence of a $2^{n+1}.3$-cycle which implies that $F^{2^{n+1}}$ has a 3-cycle, which implies F is chaotic.