SYMBOLIC DYNAMICS

KARL PETERSEN

Mathematics 261, Spring 1998
University of North Carolina at Chapel Hill

NOTETAKERS

Guillaume Bonnet
Suzanne Buchta
David Duncan
Jesse Frey
Kimberly Johnson
Lorelei Koss
Xavier Mela
Corwyn Newman
Kimberly Noonan
Robert Pratt
Kennan Shelton, Coordinating Editor
Sujin Shin
Dan Spitzner
Paul Strack

Contents

1. Introduction

These are notes from a graduate course on symbolic dynamics given at the University of North Carolina, Chapel Hill, in the spring semester of 1998. This course followed one from the previous year which considered sources of symbolic dynamics, especially the construction of Markov partitions for certain smooth systems. The topics included Sturmian and substitution systems, shifts of finite type, codings between systems, sofic systems, some information theory, and connections with topological dynamics and ergodic theory. The author thanks all the students who took notes, wrote them up, and typed them; Kennan Shelton for managing the entire project; and Sarah Bailey Frick for help with corrections.

2. January 13 *(Notes by KJ)*

2.1. **Background and Resources.**

There will be several books on reserve in the Brauer Library which will give necessary background and more details. They are:

- D. Lind and B. Marcus, *An Introduction to Symbolic Dynamics*
- K. Petersen, *Ergodic Theory*
- P. Walters, *An Introduction to Ergodic Theory*
- B. Kitchens, *Symbolic Dynamics*

We also have two handouts: The Spring 1997 Math 261 notes: Attractors and Attracting Measures, and K. Petersen, Lectures on Ergodic Theory.

The background needed for this course is general topology, analysis, and measure theory (especially if we do information theory). Our range of background is large, from first-year students to nth year, from people with no exposure to dynamics to over-exposure to dynamics. It is not necessary to have had a course in ergodic theory, but books and notes will be available for whomever needs them. Occasionally we will have to use a definition or concept from ergodic theory, but there is not time to go into the background and lots of examples, so you may have to do a little reading on the side.

In the Spring 97 version of Math 261 our purpose was to produce symbolic dynamics, to give one place it came from. We set up the geometric machinery of Markov partitions, and this gave us symbolic dynamics. Now we do symbolic dynamics in its own right, knowing what it is based on.

2.2. **General Remarks.**

What is symbolic dynamics, and what are we trying to do in this course? We are not trying to cover all of symbolic dynamics—most of what we cover is in the last chapter of Lind and Marcus.

Symbolic dynamics can be described as the study of *codings* into or between arrays of symbols. We use the term codings to mean three things: the process of quantization of data, the systems that result from quantization, and mappings between systems.

2.2.1. *Quantization.*

Quantization is the process of taking some possibly continuous object and transforming it into something discrete. Here are several examples:

- Measurement takes something continuous (length) and records it to some degree of predetermined accuracy—say 4 significant digits.
- We can take an image and quantize it to an array of pixels with some gray-scale or finite color palette. Digital television is based on this idea.

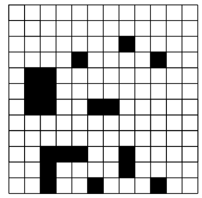

Figure 1. An image quantized into pixels

- An Axiom A dynamical system leads to a subshift of finite type when trajectories are coded according to visits to cells of a Markov partition.

Hadamard and Morse used this method of discretization when they studied geodesics, and it is also useful for studying complex dynamics.

Note that there is an interesting interaction between the continuous and the discrete: Occasionally we take a discrete system, such as population dynamics or a fluid flow, and model it using continuous dynamics, such as PDEs. Then sometimes we solve these by using discrete approximations.

2.2.2. *Systems.* The second sense of the word 'coding' is a system that is the result of quantization. These are shift dynamical systems, often called subshifts. An example of this would be an SFT (subshift of finite type).

We let $X =$ the space of all images under a coding, and let $\sigma =$ the shift, which may be higher-dimensional. In the examples above, a finite decimal measurement is the image of the coding of length. We could also have X be the space of all picture arrays.

The dynamical aspect of the system is that the shift lets us move around. Applying the shift transformation amounts to redirecting our attention to different parts of a recorded string of symbols. Applying horizontal and vertical shifts to an array of pixels lets us move around an image.

The shift can also be time. If we have a string of integers, we can move our field of view around, but we can also consider the string arriving at our computer one unit at a time. Moving one bit to the right corresponds to what arrives at our computer one second later.

The system itself is the result of a coding in the first sense of the word. We study the dynamic properties of the system, in the hopes that they will tell us the properties of what we started with.

2.2.3. *Mappings between systems.* We might call this kind of coding 're-coding.' This is what we do with image processing, for example. If we are given an image in color, we can recode it to be gray-scale.

We can ask how to find these recodings, and what systems allow it. We can also ask what remains after recoding and what information is lost. In information theory, for example, signals are recoded to be preferable in some way. For example, you could have a signal with lots of repeated symbols and irrelevant information. But after compression to get rid of the irrelevancy, you have a very unstable system. So you add in some redundancy in a carefully controlled way to get some error correction, to stabilize your information.

2.3. **Plan of the course.** The first sense of coding is what we studied last spring. That aspect is usually left to the scientists. We will mainly study systems. Here is a proposed plan of the course. It can be altered depending on what people want to see more of (or less of).

I Basic Properties of Subshifts
e.g. ergodicity, mixing, entropy, invariant measures, ...

II Examples
e.g. SFT, sofic systems, adics (Toeplitz), substitution systems (Prouhet-Thue-Morse, Chacon), Sturmian, coded systems, and some countable-state systems such as continued fractions (Gauss map) and beta-shifts.

III Coding (mappings between systems)
Between SFT's (the shift-equivalence problem), and automorphisms of SFTs (cellular automata). This section will draw heavily from Lind and Marcus, and Kitchens.

IV Information Theory
Shannon theorems, connections with Ornstein theory of Bernoulli shifts, complexity.

Please give feedback on what you want to see.

2.4. **Basic Properties of subshifts.** Let A be a finite or countable set. We call A the *alphabet*. We give A the discrete topology and

$$\Sigma(A) = \Pi_{-\infty}^{\infty} A$$

the product topology. Thus $\Sigma(A)$ is defined as

$$\{x = (x_i)_{-\infty}^{\infty} : x_i \in A \text{ for each } i\}. \tag{1}$$

The one-sided shift space is given by

$$\Sigma(A)^+ = \{x = (x_i)_0^{\infty} : x_i \in A \text{ for each } i\}. \tag{2}$$

The shift transformation $\sigma : \Sigma(A) \to \Sigma(A)$ is defined by

$$(\sigma x)_i = x_{i+1} \tag{3}$$

for $-\infty < i < \infty$.

If A has n elements ($n = 2, 3, 4, \ldots$) then we denote $\Sigma(A)$ by Σ_n and call it the *(full) n-shift*. In this case, the topology on Σ_n is compatible with the metric $d(x, y) = 1/(j + 1)$ where $j = \inf\{|k| : x_k \neq y_k\}$. Thus two elements of Σ_n are close if and only if they agree on a long central block

If A is countable, there are many inequivalent metrics compatible with the discrete topology on A. This is quite different than the finite case. For example, consider the metric on 3 points where the distance between the first and second, and the distance between the second and third is one, and the distance between the first and third is two. This gives an equivalent metric to the one where the distance between the second and third points is changed to 1/2. The metrics are equivalent, the topology for each is the discrete topology, and they give the same product topological space (figure 2).

On the other hand, say $A = \mathbb{N} = \{1, 2, 3 \ldots\}$. Then there are at least two natural ways to arrange the natural numbers, one with a limit point and one without, and these give rise to fundamentally different metrics on the product space $\Sigma(A)$ (see figure 3).

Countable alphabets appear more and more, for example in complex dynamics and in non-uniformly hyperbolic dynamical systems, and we need to consider the proper representation. However, for the near future, A will be a *finite* alphabet.

Say $|A| = n$, so $\Sigma(A) = \Sigma_n$. Then Σ_n is compact and $\sigma : \Sigma_n \to \Sigma_n$ is a homeomorphism. These facts are not too hard to see: The first follows by Tychonoff's Theorem and the second follows from observing that σ is a one-to-one onto continuous map. Thus we call the compact metric space with homeomorphism (Σ_n, σ) the n-shift dynamical system.

A *subshift* is a pair (X, σ) where $X \subset \Sigma_n$ (for some n) is a nonempty closed, shift-invariant $(\sigma X = X)$ set. Coding in the second sense given above is basically the study of subshifts.

A *block* or *word* is an element of A^r for some $r = 0, 1, 2, \ldots$, i.e. a finite string on the alphabet A. We denote the *empty block* by ε. For example, if $A = \{0, 1\}$ then $B = 011100$ is a block. We write $l(B)$ for the length of a block B. Note that already we have opportunities to be ridiculously precise: see Gottschalk and Hedlund for more precision.

The *cylinder set* determined by a block B of length r at position $j \in \mathbb{Z}$ is

$$[B]_j = \{x \in X : x_j x_{j+1} x_{j+2} \ldots x_{j+r-1} = B\}. \tag{4}$$

So the cylinder set is the set of all points in the space X which agree with B beginning in the jth place. Some important cylinder sets are those which begin with the 0'th (central) place: when we write $[B]$ we mean $[B]_0$, which begins in the 0'th place. For example, [0] is the set of sequences with 0 in the 0'th place, [01] is the set of sequences with 0 in the 0'th place and 1 in the first place, and $\sigma[01]$ is the set of sequences with 0 in the -1;st place, and 1 in the 0'th place. Note (in the picture) that when we shift the sequence to the left, that has the effect of shifting our attention to the right. This is the same phenomenon that makes taking a step forward in a room equivalent to the room moving backward.

The cylinder sets are open and closed and form a base for the topology of X.

Thus we have that Σ_n is compact, totally disconnected (the only connected sets are single points) and perfect (there are no isolated points). Hence it is homeomorphic to the Cantor middle-thirds set in $[0, 1]$ (as are most of the subshifts we will study). All the base sets we will be working with are the same (up to homeomorphisms). But the mappings on the Cantor sets will be different. If we tried to write out the mappings we are using on the Cantor set on the interval, the definitions would be terrible. But when we change spaces to other subshifts (which are relatively simple) and keep the map the same (the shift) the definitions are simple to work with.

In addition to blocks we have *rays*, which are semi-infinite blocks $(x_i)_{i=m}^{\infty}$ or $(x_i)_{-\infty}^{i=m}$, right-infinite or left-infinite sequences. We say a block B *appears* in a block C if we can find blocks D and E (possibly empty) such that $C = DBE$. We can also make many natural statements about concatenation (as in the last example) which we will assume.

Next time: More basic properties.

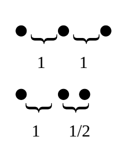

Figure 2. Two equivalent metrics for the finite alphabet shift

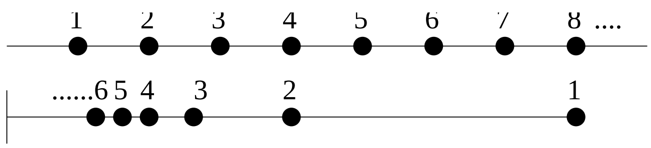

Figure 3. Two non-equivalent definitions of metrics for the countable alphabet shift

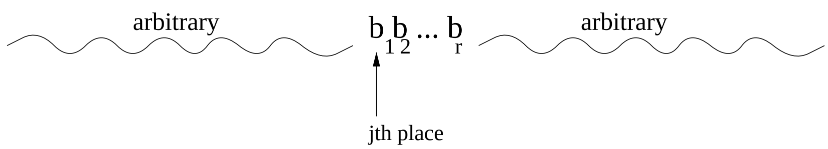

Figure 4. Example of the cylinder set $[B]_j$

3. January 15 *(Notes by KJ)*

3.1. **Basic Properties of Subshifts.** A recap of terminology introduced last time about subshifts.

We start with a finite alphabet $A = \{0, 1, \dots a-1\}$ whose cardinality a is (usually) finite and greater than 2. Then we give $\Sigma_a = A^{\mathbb{Z}}$ the product topology, which makes it a compact metric space. We define $\sigma : \Sigma_a \to \Sigma_a$ by $(\sigma x)_i = x_{i+1}$, that is, the i'th coordinate of the shift of x is the $(i+1)$'st coordinate of the original sequence. We call (Σ_a, σ) the full (two-sided) a-shift.

A *subshift* (X, σ) where $X \subset \Sigma_a$ for some $a \geq 2$ is a nonempty closed σ-invariant set ($\sigma X = X$); X has the subspace topology.

More on the topology (in a full shift): Two sequences x, y are close if they agree on a long central block. If the sequences agree on an even longer block far away from the 0'th place, that doesn't

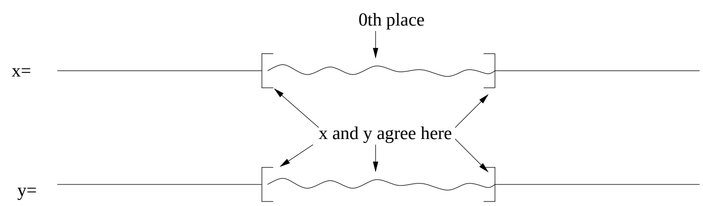

Figure 5. Two points are close if they agree on a long central block

say how close they are, but how close some shift of the two sequences are.

We call $[B]_m = \{x \in \Sigma_a : x_m x_{m+1} \dots x_{m+l(B)-1} = B\}$ (where B is a finite word and $l(B)$ is the length of B) a *cylinder set*. One motivation for the name cylinder set may be in $\mathbb{R}^3$: If you restrict the x and y coordinates but leave the z coordinate free, you get a cylinder-like object.

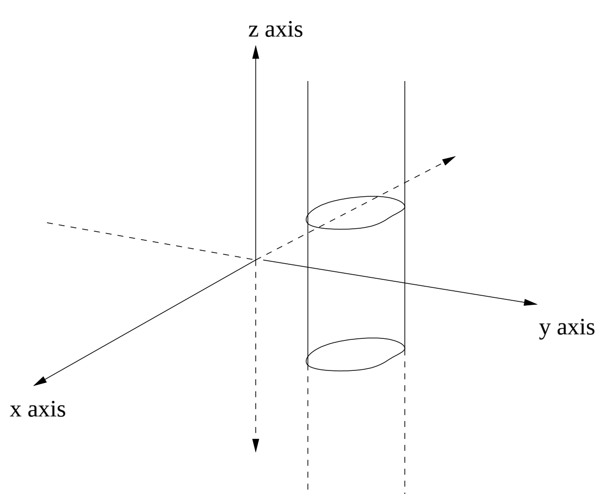

Figure 6. Motivation for the term "cylinder set"

Cylinder sets are open and closed, and they form a base for the topology. To see that they are open, let $x \in [B]_m$, so that B appears in x starting at the m'th place. We want to show that x is in the interior of $[B]_m$. To do this we show that points close enough to x are in the set $[B]_m$ also. Let y be a point close enough to x so that y agrees with x on a long central block including the places beginning with m where B appears. Then y also has B appearing at those places and so must be in $[B]_m$.

To see that every cylinder set is closed, let $x_n \in [B]_m$ be a sequence converging to x. We need to show that $x \in [B]_m$. Notice that after a while, the x_n agree with x on a long central block. If they agree on a long enough block, they agree from m to $m + l(B) - 1$. But since x_n agree with B there, so does x and so $x \in [B]_m$.

Notice that since the base for the topology consists of open and closed sets, we have a difficult time finding connected sets. Thus the only connected sets are singletons, and so subshifts are totally disconnected (zero-dimensional).

3.1.1. *Languages.* Define the following:
$A^0 = \{\varepsilon\}$ = the empty word
$A^1 = A$, $A^2 = AA$ = all words of length 2 (we write AA to mean concatenation), etc.
$A^* = \cup_{n \geq 0} A^n$ = the set of all finite words on the letters in A.

A *language* on the alphabet A is any subset $\mathcal{L} \subset A^*$. This includes all formal languages, programming languages, etc. $\mathcal{L}$ doesn't have to be finite.

Define the *language of a subshift* $(X, \sigma) \subset (\Sigma_a, \sigma)$ to be $\mathcal{L}(X, \sigma)$ = the collection of all finite blocks found in sequences $x \in X$. This is the set $\{B \in A^* :$ there is $x \in X$ such that B appears in $x\}$. Recall B appears in x means that $x = \dots B \dots$. Note that $\mathcal{L}(\Sigma_a, \sigma) = A^*$.

Relating to this definition there is the following theorem:

Theorem 3.1. *If (X, σ) is a subshift, then $\mathcal{L}(X, \sigma)$ is*

(i) factorial*: that is, if $B \in \mathcal{L}(X, \sigma)$ and C is a sub block of B (i.e., $B = pCs$ for some possibly empty blocks $p, s \in A^*$) then $C \in \mathcal{L}(X, \sigma)$;*
(ii) extendable*: that is, if $B \in \mathcal{L}(X, \sigma)$ then there are nonempty blocks $p, s \in A^*$ such that $pBs \in \mathcal{L}(X, \sigma)$.*

Moreover, properties (i) and (ii) characterize the languages of subshifts: Given any $\mathcal{L} \subset A^*$ which is nonempty, factorial, and extendable, there is a unique subshift X such that $\mathcal{L} = \mathcal{L}(X, \sigma)$.

Proof. If B is in the language of a subshift, that means B appears in some sequence x, and so do all its subwords; further, we can extend B to both sides. (Notice that $\mathcal{L}(X, \sigma)$ is never finite, because it contains words of all lengths.)

For the second part of the theorem, given a nonempty, factorial, extendable language $\mathcal{L}$, we define X to be the set of all sequences which do not contain any words that are not in $\mathcal{L}$ and show that it is a subshift. Given an arbitrary language, there might not be anything in X. Since $\mathcal{L}$ is not empty, we can take any word in $\mathcal{L}$ and extend it to the right and left in accordance with item (ii). Thus X is not empty. X is σ-invariant since no one said where the origin was as we extended a word in $\mathcal{L}$.

To show that X is closed, we show that its complement is open. Let x be in the complement. This means that it has some bad word in it (one which is not in $\mathcal{L}$). Words which are close to x (they agree on a long central block) will also have this bad word in it, so the complement of X is open in the metric space topology. $\square$

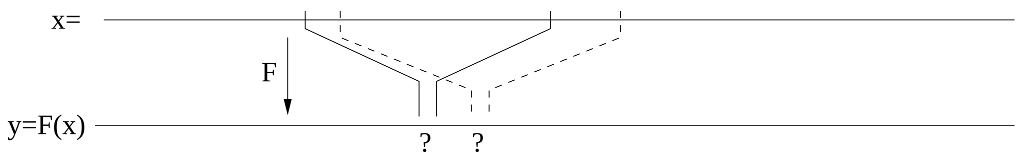

FIGURE 7. A sliding block map

Using this theorem we have the progression

$$X \to \mathcal{L}(X) \to X_{\mathcal{L}(X)} = X,$$

where $X_{\mathcal{L}}$ is the unique subshift given by the extendable factorial language $\mathcal{L}$. Thus there are interesting connections between symbolic dynamics and the study of formal languages, the Chomsky hierarchy, automata theory, and so on.

There was some discussion about whether the languages consisting of the empty word (ε) and the empty language also gave subshifts. The language consisting only of ε is not extendable, so the theorem does not apply. Note that ε is in every factorial language. Lind and Marcus do not restrict the subshift resulting from the language to being nonempty. Unless we use only nonempty languages, it makes sense for us to include the empty subshift also.

3.1.2. *Finite Codes. Finite codes* are also known as *block codes* and *sliding block codes.* They are codes between subshifts. We will describe one of these codes, a map F from Σ_a to Σ_b for some a and b. We will say that $y = F(x)$.

Let $w \geq 0$ be the size of a 'window' or 'range'. We want to figure out what symbol to put in the i'th place of $F(x)$. To do this, we look at a window of symbols in x and use that to decide. We slide the window back and forth to figure out other symbols. This method goes on in the real world in data encoding and receiving.

Formally, let $f : A^w \to B = 0, 1 \dots b-1$ for some $b \geq 1$, where w is the size of the window above. This is a *block map.* Then we define $F : \Sigma_a \to \Sigma_b$ by

$$(Fx)_i = f(x_{i-m}x_{i-m+1}\dots x_{i+n}) \text{ (where } w = m+n+1) \tag{5}$$

for all $i \in \mathbb{Z}$. Then $F : \Sigma_a \to \Sigma_b$, or its restriction to any subshift of Σ_a, is called a *sliding block code* with *window* (or *range*) w, *memory* m and *anticipation* n.

The shift map or a power of the shift is an example of a map which only depends on anticipation, so we may want to have the concept of negative memory. In the shift map, $(\sigma x)_i = f(x_{i+1})$ where f is the identity map on the alphabet. Thus, $i - m = i + n = i + 1$, so $w = 1, m = -1, n = 1$. We could also define the shift map by $(\sigma x)_i = g(x_i x_{i+1})$ where $g(x_0 x_1) = x_1$. In this case, $w = 2, m = 0, n = 1$. For examples, see figure 8.

One important thing to note is that when we slide the sliding window over, we only slide it over one block at a time instead of the width of the window.

We have the following theorem:

Theorem 3.2 (Curtis, Hedlund, Lyndon). *If (X, σ) is a subshift of (Σ_a, σ) and $F : (X, \sigma) \to (\Sigma_b, \sigma)$ is a sliding block code determined by a block map $f : A^w \to B$, then F is a continuous shift commuting map (a* factor map*).*

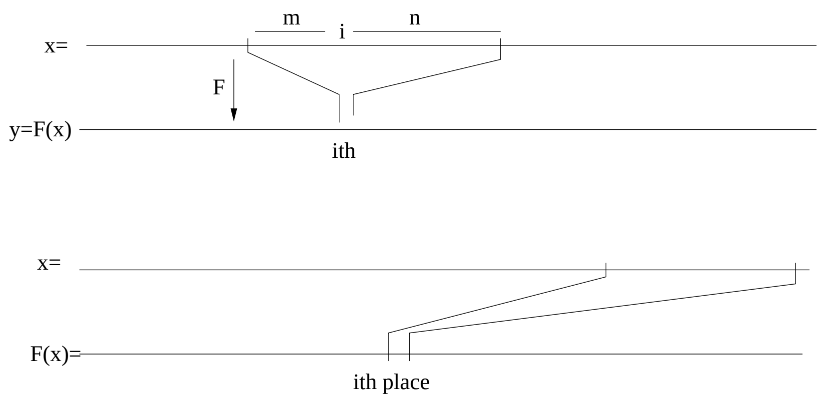

Figure 8. Sliding block map with memory m and anticipation n; sliding block map with negative memory

Conversely, if $(X, \sigma) \subset (\Sigma_a, \sigma)$ *and* $(Y, \sigma) \subset (\Sigma_b, \sigma)$ *are subshifts and* $\phi : (X, \sigma) \to (Y, \sigma)$ *is continuous and commutes with* σ*, then* ϕ *is given by a sliding block code.*

This theorem says that sliding block codes give all continuous shift-commuting maps between subshifts. This is particularly interesting in the viewpoint of category theory, which says we should want to study all maps among things in a category. In shift spaces, we want them to be shift-commuting and continuous (since Σ_a is a topological space), so this theorem says we need not look any farther than sliding block codes.

To help see what this means, we give an example of a simple sliding block code.

Example 3.1. Let $A = \{0, 1\}$. We define a map $f : A^3 \to B = \{\alpha, \beta\}$ pointwise:

$$\begin{aligned} 000 &\to \beta \\ 001 &\to \beta \\ 010 &\to \alpha \\ 011 &\to \beta \\ 100 &\to \alpha \\ 101 &\to \alpha \\ 110 &\to \alpha \\ 111 &\to \beta \end{aligned}$$

If we let F be the induced sliding block code with memory 0, and $x = \ldots 1100001010111\ldots$, then $y = F(x) = \ldots \alpha\alpha\beta\beta\beta\beta\alpha\alpha\alpha\alpha\beta\beta\ldots$.

Proof (of CHL). The first part comes from two observations. The map F is shift-commuting because clearly as you shift the center of the image, you just shift the center of the window along.

The map F is continuous because of the finite range. To see this, we want to show that two points can be as close as we want them to be in the range, provided the points they came from

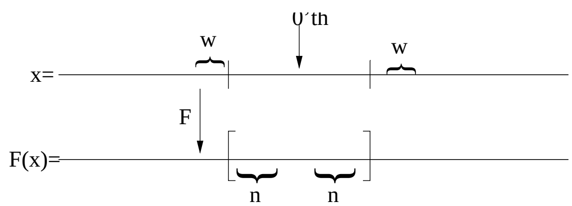

Figure 9. The sliding block map is continuous

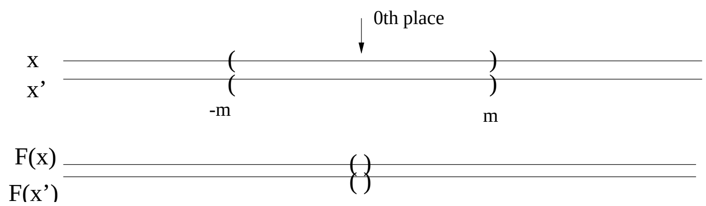

Figure 10. Uniform continuity gives us equivalence classes of $2m+1$ blocks

in the domain are close enough. So, decide how close the points in the range should be—say you want them to agree on a central block of length $2n+1$. Then if you choose $x, x' \in X$ such that they agree on a block of length $2(n+w)+1$, their images must agree on the smaller block: if we have $x'_j = x_j$ for $|j| \leq n+w$ then $F(x)_j = F(x')_j$ for $|j| \leq n$ as in figure 9.

For the converse we use another picture. Let $X \subset \Sigma_a, Y \subset \Sigma_b$ be subshifts, $\phi : X \to Y$ be continuous and such that $\phi\sigma = \sigma\phi$ (here σ is being used both as the shift on X and on Y). The map is continuous and hence is uniformly continuous, since X is compact. (At this point the proof for countable-state alphabets goes to blazes.) Uniform continuity implies that there is m such that if $x_j = x'_j$ for $|j| \leq m$, then $(\phi x)_0 = (\phi x')_0$. To see this, remember that closeness is determined by agreeing on a central place. How close they have to be in the domain is determined by this m from uniform continuity (see figure 10).

The $(2m+1)$-blocks fall into equivalence classes according to which symbol they determine in the image—this gives a block map and hence a sliding block code.

Since ϕ is a shift-commuting map, we have the same story in the j'th place as we do in the 0'th place. The j'th place in x, x' is the 0'th place in $\sigma^j x, \sigma^j x'$. □

This theorem is important and useful because in theory now we know what all the *factor maps* or *homomorphisms* between subshifts are. We give some examples of different types of sliding block codes.

Example 3.2 (1-block code)**.** Let $A = \{0, 1, 2\}$ and $B = \{\alpha, \beta\}$. Then define the block map that sends even numbers to α and odd numbers to β:

$$\begin{aligned} 0 &\rightarrow \alpha \\ 1 &\rightarrow \beta \\ 2 &\rightarrow \alpha \end{aligned}$$

In this example, if $x = \ldots 10112001102 \ldots$ we have $F(x) = \ldots \beta\alpha\beta\beta\alpha\alpha\alpha\beta\beta\alpha\alpha \ldots$. This example collapsed two letters, so information was lost. We could have simply renamed the letters in a one-to-one way and lost no information.

Theorem 3.3. *If $a = |A|$ is prime, then for each $w \geq 1$, any map $f : A^w \rightarrow A$ is given by a polynomial in w variables over $GF(a)$ (the general field with a elements).*

For example, if $a = 2, A = \{0, 1\}$, and $p(x_0, x_1, x_2) = x_1 + x_0x_2$, then we have the block code

$$\begin{aligned} 000 &\rightarrow 0 + 0 \cdot 0 = 0 \\ 101 &\rightarrow 0 + 1 \cdot 1 = 1 \\ 110 &\rightarrow 1 + 1 \cdot 0 = 1 \\ &\vdots \end{aligned}$$

Exercise 1. Find the polynomial for example 3.1 and prove it can always be done. What happens if a is not prime?

Example 3.3 (Higher block code)**.** Let $(X, \sigma) \subset (\Sigma_a, \sigma)$ be a subshift, and fix $r \geq 2$. Take a new alphabet $B = A^r = r -$ blocks on A. Define $F : (X, \sigma) \rightarrow (\Sigma_b, \sigma)$ $(b = a^r)$ as generated by the block map

$$f(x_1 \ldots x_r) = x_1 \ldots x_r \in B \text{ for } x_1, \ldots, x_r \in A. \tag{6}$$

Note that this map is into (Σ_b, σ), but not onto.

For example, we can do the higher block map with $r = 2$ from Σ_2 into Σ_4:

$$\begin{aligned} 00 &\rightarrow \alpha \\ 01 &\rightarrow \beta \\ 10 &\rightarrow \gamma \\ 11 &\rightarrow \delta \end{aligned}$$

Then given a sequence $x = \ldots 110100010 \ldots$, we get $F(x) = \ldots \delta\gamma\beta\gamma\alpha\gamma\beta\gamma \ldots$. When you are taking the image of x, don't shift by two to get the next symbol in $F(x)$. Only shift by one.

Using an r-block code F as above, the image $F(X) \subset \Sigma_b$ is a subshift that is *topologically conjugate* to (X, σ) (it's an isomorphic, i.e., one-to-one onto shift-commuting image); this image is called the *r-block representation* of (X, σ).

Note that the r-block representation of $(X\sigma)$ is different from (X, σ^r). It is true that $(\Sigma_2, \sigma^2) \cong (\Sigma_4, \sigma)$ by mapping each 2-block (they don't overlap) to a separate symbol. However, $(\Sigma_2, \sigma) \ncong (\Sigma_4, \sigma)$ via the higher block map. For example, in the image of the higher block map given above,

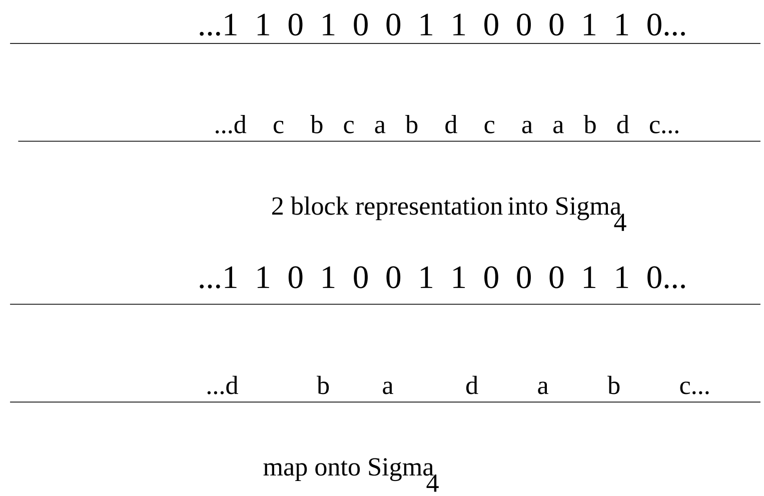

Figure 11. The 2-higher block representation of Σ_2, along with a map from (Σ_2, σ^2) to (Σ_4, σ).

α could never be followed by γ, while in the full 4-shift they could follow each other, as seen in figure 11.

Next time: Dynamic properties such as ergodicity and mixing. Look at the topological dynamics section of Petersen, or page 16 or so in the old 261 notes.

4. January 20 *(Notes by PS)*

4.1. **Topological Dynamical Systems and Maps Between Them.**

Definition 4.1. A *topological dynamical system* is a pair (X, T), where X is a compact Hausdorff space (usually metric) and $T : X \to X$ is a continuous mapping (usually a homeomorphism).

Definition 4.2. A *homomorphism* or *factor mapping* between topological dynamical systems (X, T) and (Y, S) is a continuous onto map $\phi : X \to Y$ such that $\phi T = S\phi$. We say Y is a *factor* of X, and X is an *extension* of Y.

Definition 4.3. A *topological conjugacy* is a one-to-one and onto factor map.

For example, the factor mappings between subshifts are all given by finite window sliding-block codes.

One major problem is to classify subshifts, or even subshifts of finite type, up to topological conjugacy. We need invariants to do so, preferably complete invariants. If this question were solved, the next step would be to construct codings between them:

$$\phi : (\Sigma, \sigma) \to (\Sigma', \sigma)$$

There are various engineering difficulties motivating such recodings. For example, if your original system allowed arbitrarily long sequences of zeros and ones, then slight errors might arise if your hardware had difficulty distinguishing between 6,000,000 zeros and 6,000,001 zeros. It would be better to recode the system so that there was some upper bound on the length of sequences of zeros.

Similar classification problems exist for other classes of topological dynamical systems, and for other kinds of maps with weaker or stronger conditions (for example, measurability).

Definition 4.4. Fix $(X, T) = (\Sigma_a, \sigma)$, a full shift. The factor maps

$$\phi : (\Sigma_a, \sigma) \to (\Sigma_a, \sigma)$$

are *endomorphisms* of (Σ_a, σ) or *cellular automata*. One-to-one endomorphisms are *automorphisms*.

Since $\phi\sigma = \sigma\phi$, we have an action of $\mathbb{Z} \times \mathbb{Z}$ on (Σ_a, σ), namely

$$(m, n)x = \sigma^m \phi^n x$$

for $m, n \in \mathbb{Z}$. If ϕ is not invertible, then we have an action of $\mathbb{Z} \times \mathbb{Z}^+$ on Σ_a.

For example, suppose ϕ is given by the block map

$$(\phi x)_i = x_i + x_{i+1} \bmod 2$$

for $a = 2$, $A = \{0, 1\}$. Consider

$$\begin{aligned} x &= ...101110100010110... \\ \phi x &= ...11001110011101... \\ \phi^2 x &= ...0101001010011 \end{aligned}$$

or

$$\begin{aligned} x &= ...000001000... \\ \phi x &= ...000011000... \\ \phi^2 x &= ...000101000... \\ \phi^3 x &= ...001111000... \\ \phi^4 x &= ...010001000... \\ \phi^5 x &= ...110011000... \end{aligned}$$

In the latter case, we have Pascal's triangle mod 2. In general, 1's that appear within x will try to produce such triangles in the images. Nearby 1's will interfere with each other, creating complex patterns. If there are infinitely many 1's appearing in the sequence x, the pattern becomes nearly impossible to predict. Thus, even seemingly simple maps ϕ can have complex developments.

The system (Σ_a, ϕ) has been studied as a topological dynamical system in its own right, by such individuals as R. Gilman, F. Blanchard and A. Maass.

Similar questions have been asked about topological dynamical systems other than Σ_a, for example subshifts of finite type. Finding all endomorphisms of such systems and studying the properties of those endomorphisms leads to some interesting open problems. For a more in depth study of endomorphisms of subshifts of finite type, see the paper by Boyle, Lind and Rudolph.

4.2. **Dynamical Properties of Topological Dynamical Systems.** Consider a compact metric space X and a homeomorphism $T : X \to X$ (often T continuous is sufficient).

Definition 4.5. For $x \in X$, the *orbit* of x is $\mathcal{O}(x) = \{T^n x : n \in \mathbb{Z}\}$ (alternately, if T is not invertible, $n \in \mathbb{Z}^+$). The *orbit closure* of x is $\overline{\mathcal{O}(x)}$.

Definition 4.6. A set $A \subset X$ is *invariant* if $T(A) \subset A$.

Definition 4.7. We say (X, T) is *topologically ergodic* if it satisfies one of the following equivalent properties:

(1) *Topological transitivity*: There is a point $x \in X$ with a dense orbit, that is $\overline{\mathcal{O}(x)} = X$.
(2) *The set of points with dense orbit is residual.* Recall that a residual set is a set that contains the intersection of countably many dense open sets. Equivalently, a residual set is the complement of a first category set, that is, a set that is contained in the union of countably many nowhere dense sets.
(3) *Regional transitivity*: Given non-empty sets $U, V \subset X$, there is some $n \in \mathbb{Z}$ such that $T^n U \cap V \neq \emptyset$. See Figure 12.
(4) *Every closed invariant set is nowhere dense.* The idea here is a sort of "topological irreducibility", in that the only closed T-invariant subsets of X must be nowhere dense.

For a detailed proof of the equivalence of these properties, see the Spring 1997 Math 261 notes, p. 27. These notes also include an additional equivalent property, Baire ergodic.

Note that the subshift $(\overline{\mathcal{O}(x)}, T)$ is necessarily topologically ergodic.

This leads to the question of how we can determine which *subshifts* are ergodic. Consider closed $X \subset \Sigma_a$ with $\sigma X \subset X$, so that (X, σ) is a subshift. Then (X, σ) is topologically ergodic if and only if it has a dense orbit, which is true if and only if there is an $x \in X$ which contains all the words in $\mathcal{L}(X)$.

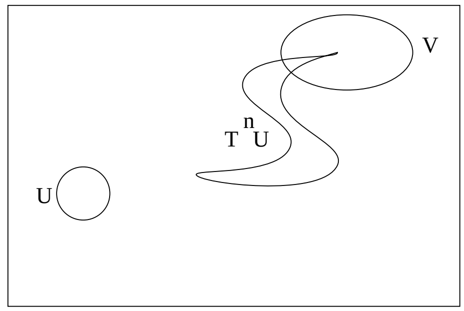

FIGURE 12. Regional Transitivity

To see why this is true, consider $y \in X$, $y = ...[B]...$ where $B \in \mathcal{L}(X)$ is a central block of y. If the word B also appears somewhere in x, then for some n, $\sigma^n x$ will have B as its central block, and therefore will be near y. In this way we can reduce the dynamical question of the ergodicity of (X, σ) to a combinatorial question about x and $\mathcal{L}(X)$.

4.3. **Minimality.**

Definition 4.8. A topological dynamical system is called *minimal* if there is no proper closed invariant set, or, equivalently, if for all $x \in X$ the orbit $\mathcal{O}(x)$ is dense.

Theorem 4.1. *If (X, T) is any (compact) topological dynamical system, then there are proper closed invariant sets $A \subset X$ such that (A, T) is minimal.*

Proof. Order the closed invariant subsets of X by inclusion, and use Zorn's Lemma. □

Let $x \in \Sigma_a$. When is $(\overline{\mathcal{O}(x)}, \sigma)$ minimal? The following property is what we need.

Definition 4.9. A point x is *almost periodic* or *syndetically recurrent* if for every neighborhood U of x, the set of return times

$$R(U) = \{r \in \mathbb{Z} : T^r x \in U\}$$

has *bounded gaps*, in that there is some K such that for all $n \in \mathbb{Z}$,

$$(n - K, n + K) \cap R(U) \neq \emptyset.$$

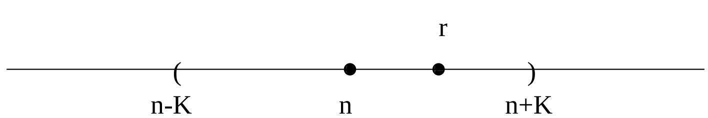

FIGURE 13. Return time $r \in R(U)$

The term *almost periodic* has many different meanings in dynamics, so we prefer the term *syndetically recurrent* instead.

5. January 22 *(Notes by PS)*

Theorem 5.1. *Let (X,T) be a compact topological dynamical system and $x \in X$. Then $(\overline{\mathcal{O}(x)}, \sigma)$ is minimal if and only if x is syndetically recurrent.*

Proof. Suppose x is syndetically recurrent. Let $y \in \overline{\mathcal{O}(x)}$, and let U be a *compact* neighborhood of x. Recall that every compact metric space is locally compact, that is, it has a neighborhood base consisting of compact neighborhoods. Thus, U can be chosen arbitrarily small.

Since $R(U)$ has bounded gaps, there is a K such that

$$\mathcal{O}(x) = \bigcup_{j=-K}^{K} T^j \bigcup_{r \in R(U)} T^r x$$

but

$$\bigcup_{j=-K}^{K} T^j \bigcup_{r \in R(U)} T^r x \subset \bigcup_{j=-K}^{K} T^j(U).$$

Since U is compact and closed, so are $T^j(U)$ and $\cup_{j=-K}^{K} T^j(U)$. Thus

$$\overline{\mathcal{O}(x)} \subset \bigcup_{j=-K}^{K} T^j(U)$$

as well. Since $y \in \overline{\mathcal{O}(x)}$, there is some $j \in [-K, K]$ such that $y \in T^j U$ and $T^{-j} y \in U$.

It follows that $\mathcal{O}(y)$ intersects the neighborhood U of x. Since the neighborhood U can be chosen arbitrarily small, we see that $x \in \overline{\mathcal{O}(y)}$. Thus $\mathcal{O}(x) \subset \overline{\mathcal{O}(y)}$, so that $\overline{\mathcal{O}(x)} = \overline{\mathcal{O}(y)}$ and the orbit of y is dense in $\overline{\mathcal{O}(x)}$, making $\overline{\mathcal{O}(x)}$ be minimal.

Conversely, suppose $\overline{\mathcal{O}(x)}$ is minimal. Let U be a neighborhood of x and $R(U) = \{r \in \mathbb{Z} : T^r x \in U\}$. We want to show that $R(U)$ has bounded gaps. For any $y \in \overline{\mathcal{O}(x)}$, $\overline{\mathcal{O}(x)} = \overline{\mathcal{O}(y)}$ because of the minimality of $\overline{\mathcal{O}(x)}$, so there is a j such that $T^j y \in U$. Therefore,

$$\bigcup_{j=-\infty}^{\infty} T^j(U) \supset \overline{\mathcal{O}(x)}.$$

This is an open covering of $\overline{\mathcal{O}(x)}$, and by the compactness of X, there is an finite subcover. That is, there is a K such that

$$\bigcup_{j=-K}^{K} T^j(U) \supset \overline{\mathcal{O}(x)}.$$

Clearly this implies that $R(U)$ has bounded gaps.

□

This theorem is due to G. D. Birkhoff (Bull. Soc. Math. France, 1912). With appropriate slight modification, it holds for non-invertible T as well.

5.1. **Minimality in** (Σ_a, σ)**.** In a shift space (Σ_a, σ), the orbit closure of a point $x \in \Sigma_a$ will be minimal if and only if every block that appears in x appears with bounded gap. That is, if B is any word in x, then B appears and reappears infinitely many times, with the time between each repeat bounded by some k.

$$x = ...[B]...[B]...[B]...[B]...$$

Note that k will depend on the length of B, with larger k for longer words B.

The full shift itself is *not* minimal, since it has lots of proper closed σ-invariant sets, for example the set consisting of the fixed point

$$\overline{0} = ...00.00....$$

Here, the decimal place marks the 0th place in the sequence $\overline{0}$.

Another example is the cycle of points

$$\begin{aligned} x_1 &= ...0010.01001... \\ x_2 &= ...0100.10010... \\ x_3 &= ...1001.00100... \end{aligned}$$

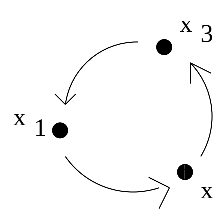

FIGURE 14. The Cycle x_1 to x_3

Both of these examples are also examples of points with minimal orbit closure $\overline{\mathcal{O}(x)}$. All finite cycles are trivially minimal.

5.2. **Ergodicity of** (Σ_a, σ)**.**

Theorem 5.2. *The space* (Σ_a, σ) *is topologically ergodic.*

Proof. We can demonstrate the ergodicity of Σ_a in two ways. First, we show that it has regional transitivity. Let U and V be open sets in (Σ_a, σ). Then both U and V must contain some cylinder sets $[B]_m \subset U$ and $[C]_{m+r} \subset V$. We will construct an $x \in U$ so that $T^n x \in V$, and therefore $T^n(U) \cap V \neq \emptyset$.

Simply choose n large enough so that the word C in the $(m + r + n)$th place and the word B in the mth place do not overlap. Fill in the rest of x with zeros, so that

$$x = ...000[...B...]00...00[...C...]000...$$

Note that with some obvious modification, this proof also demonstrates *strong mixing* for (Σ_a, σ), in that there is some N such that for all $|n| \geq N$, $T^n(U) \cap V \neq \emptyset$.

We can also show that (Σ_a, σ) is topologically transitive, in that it has a dense orbit. To do so, we construct an $x \in \Sigma_a$ which contains *all* words of $\mathcal{L}(\Sigma_a)$. This is like constructing a "Champernowne number". For example, in base 10, a Champernowne number is simply one whose digits consists of all possible integers in order, that is

$$x = .1\ 2\ 3\ 4\ 5\ 6\ 7\ 8\ 9\ 10\ 11\ 12\ 13\ 14\ ...$$

Similarly, for Σ_2, we make a sequence consisting of all possible words base 2, that is

$$x = .0\ 1\ 00\ 01\ 10\ 11\ 000\ 001\ 010\ 011\ 100...$$

If we place a string of zeros before the decimal place, such an x will have a dense orbit in Σ_a.

□

An interesting question related to the topological transitivity of Σ_a is whether or not the decimal expansion of a given irrational number will have a dense orbit in Σ_{10}. Consider π,

$$x = ...0000.314159...$$

It is conceivable that after some point, no 6's appear in this sequence, so that it would *not* have a dense orbit. For examples of applications of questions like this this (involving uniform distribution) to things like random number generators, see the book by Niederreiter.

5.3. **Examples of Non-Cyclic Minimal Systems.** Let X be the unit circle in the complex plane, $X = S^1 = \{z \in \mathbb{C} : |z| = 1\}$, and $\alpha \in (0, 1)$ be an irrational number. Let T be an irrational rotation of X, with $T(z) = e^{2\pi i\alpha}z$. Then every orbit in (X, T) is dense and X is minimal. This is a theorem by Kronecker (and not hard to prove).

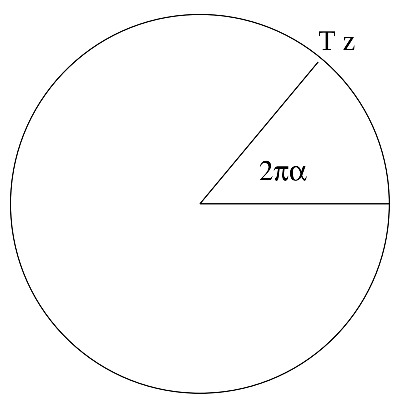

FIGURE 15. Rotation by α

For a minimal subshift of Σ_a, we use an example discovered and rediscovered by Prouhet, Thue, Morse and many others. We construct a sequence ω inductively by using a *substitution* map ζ, with $\zeta(0) = 01$ and $\zeta(1) = 10$. Alternatively, at each stage of the construction, we append the dual of the previous step, where the dual of 0 is 1 and the dual of 1 is 0. Thus

$$\begin{aligned}
\zeta(0) &= 01 \\
\zeta^2(0) &= 01\ 10 \\
\zeta^3(0) &= 0110\ 1001 \\
\zeta^4(0) &= 01101001\ 10010110
\end{aligned}$$

To construct ω, we place the resulting one-sided sequence to the right of the decimal and its reversal to the left, so that

$$\omega = ...01101001\ 1001\ 01\ 10.01\ 10\ 1001\ 10010110...$$

Note first of all that ω contains no 000 or 111 blocks, since ω consists of strings of 01's and 10's. Thus both 0 and 1 appear with a bounded gap of 2. Also note that by recoding ω, changing 01 to a and 10 to b, we simply get ω again on the symbols a and b. In fact, let $a_r = \zeta^r(0)$ and $b_r = \zeta^r(1)$. By recoding with these words as symbols, we again get ω.

Now, let B be a word in ω. Suppose B appears entirely to the right of the decimal, in the initial 2^r-block of ω. Thus B appears in a_r. But by substituting,

$$\omega = .a_r b_r b_r a_r...$$

It follows that a_r, and thus B, appears with bounded gap. The case is similar if B lies entirely to the left of the decimal place. If, on the other hand, the word B overlaps the decimal place, observe that

$$\omega = ...a_r b_r b_r a_r.a_r b_r b_r a_r...$$

or

$$\omega = ...b_r a_r a_r b_r.a_r b_r b_r a_r...$$

depending on r. Either way, B is contained in either $a_r.a_r$ or $b_r.a_r$, both of which appear in a_{r+2}, hence with bounded gap.

To complete the proof, all that remains is to show is that ω is not periodic.

6. January 27 *(Notes by SB)*

Claim: The Prouhet-Thue-Morse (PTM) sequence

$$\omega = \dots\dots.01101001\dots$$

is NOT periodic.

Proof. Suppose the PTM sequence is periodic. Then

$$\omega = \dots BBBB \dots$$

Assume, without loss of generality, that $l(B)$ is ODD. (Grouping ω into 2^r-blocks, a_r and b_r, where $l(B)/2^r$ is integer and odd, again produces a PTM sequence, on the symbols a_r and b_r). Then there exists $r > 0$ such that $2^r \equiv 1 \pmod{l(B)}$. Since $(2, 2^2, 2^3, \dots \pmod{l(B)}$ is an infinite list in a finite set $\{\hat{0}, 1, \dots, l(B) - 1\}$, we have $2^s = 2^{s+r}$ for some $s, r > 0$. But 2^s is relatively prime to $l(B)$, hence it has a multiplicative inverse $\pmod{l(B)}$. This implies $1 \equiv 2^r(\pmod{l(B)})$. Now, looking at the PTM sequence, the sequences starting at the 1st and the 2^r'th place should be the same, since $\omega = \dots BBB \dots$ and $1 \equiv 2^r \pmod{l(B)}$ implies $\sigma\omega = \sigma^{2^r}\omega$

$$\omega = .\overbrace{\underbrace{0}_{0th}\underbrace{1}_{1st}101001\dots}^{a_r}.\overbrace{\underbrace{1}_{2^r th}0010110\dots}^{b_r}\overbrace{10010110\dots}^{b_r}.\overbrace{01101001\dots}^{a_r}\dots\dots$$

however, we can see that this is not the case $(11 \neq 10)$. □

6.1. Applications of the PTM Sequence:

(1) Morse used this sequence to construct recurrent but nonperiodic geodesics on surfaces, using Hadamard's idea of coding geodesics by means of a *partition* on the surface. *Special note:* Hadamard published this result in 1898, making this year the 100'th anniversary of his accomplishment.
(2) Axel Thue used this sequence in his work on questions involving logic and group theory.
(3) The PTM sequence makes it possible to find sequences without too many repetitions, an issue of importance in computer science. The PTM sequence is *cube-free*, i.e., it contains no BBB for any block B, and it can be used to make a *square-free* sequence (i.e., no BB) on three symbols.
(4) *The Burnside Problem:* Decide whether the group G on r generators with relations g^n =identity for all $g \in G$ can be infinite.
(5) K. Mahler's work on other problems in number theory
(6) Prouhet's work, published in 1851.

6.2. Some Other Ways to Generate the PTM Sequence:

(1) For $n \geq 0$, $\omega = \text{sum} \pmod 2$ of binary digits of n: if $n = a_0 + a_1 \cdot 2 + a_2 \cdot 2^2 + \dots + a_r \cdot 2^r$ then $\omega(n) \equiv a_0 + a_1 + \dots a_n \pmod 2$.
(2) *Keane's block multiplication:* Let $0' = 1$, $1' = 0$. If $B = b_1 \dots b_r$ is a block, put $B \times 0 = B$, $b \times 1 = B' = b_1' \dots b_r'$. If $C = c_1 \dots c_n$ is another block, put $B \times C = (B \times c_1)(B \times c_2) \dots (B \times c_n)$

Example 6.1. $(1101) \times (101) = 001011010010$

Example 6.2. $\omega_0\omega_1\omega_2 \dots = 0 \times (01) \times (01) \times (01) \times \dots = 01101001 \dots$

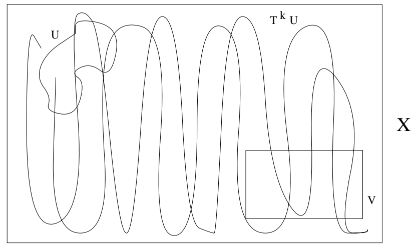

Figure 16. After some time, the set U under the action of T will stay in contact with every sampling set V

Block multiplication leads to "generalized Morse sequences" (Keane '69) which are defined by other infinite block products like:

$$0 \times (001) \times (001) \times (001) \times \ldots = 001001110001001110110110001\ldots$$

Here the code can be interpreted as "starting with 001, take what is written, write it again, then write down its dual (e.g., $001 \to 110$)."

6.3. **Generalizing Properties of the PTM Sequence.** The PTM sequence is the starting point for a number of properties that can be generalized and then applied to a number of different settings.

6.3.1. *Substitutions.* For example, let $\tau 0 = 011$ and let $\tau 1 = 00$. Then

$$\begin{aligned} 0 &= 0 \\ \tau 0 &= 011 \\ \tau^2 0 &= 0110000 \\ \tau^3 0 &= 0110000011011011011 \\ &\vdots \end{aligned}$$

This substitution gives, in the limit, a one-sided sequence. Complete it to the left with all 0's. The set of forward limit points under the shift is a closed invariant set called a *substitution dynamical system.* Among the first to study these were W. Gottschalk, J. Martin and P. Michel. The lecture notes by M. Queffelec summarize much of what is known about these systems.

6.3.2. *Topological Strong Mixing.* Although we have already discussed this property, we have not yet given a formal definition, so we give one now.

Definition 6.1. Let X be a compact metric space and $T : X \to X$ a homeomorphism. We say that (X, T) is *topologically strongly mixing (t.s.m.)* if given non-empty open sets $U, V \subset X$, there exists n such that if $|k| \geq n$, then $T^k U \cap V \neq \emptyset$.

This property is illustrated by Figure 16.

Topological strong mixing for subshifts is easily characterized.

Proposition 6.1. *A subshift (X, σ) of some $(\Sigma_a, \sigma))$ is topologically strongly mixing if and only if given any blocks $B, C \in \mathcal{L}(X)$, there exists n such that if $|k| \geq n$, then there exists $x \in X$ such* *that* $x = \dots\dots \underbrace{B\dots C}_{k} \dots\dots$ *In other words, given enough time, we can get from* B *to* C *within the system* X*.*

Note: The PTM sequence is NOT topologically strongly mixing.

Definition 6.2. We say that (X, T) is *topologically weakly mixing (t.w.m.)* in case the Cartesian square $(X \times X, T \times T)$ is topologically ergodic. (Recall that $((T \times T)(x_1, x_2) = (Tx_1, Tx_2))$.

Notice that topological strong mixing implies topological weak mixing. For given non-empty open A, B, U, V, we want to find n with $(T \times T)^n(A \times B) \cap (U \times V) \neq \emptyset$. This is easily accomplished: find N_1 and N_2 such that $T^n A \cap U \neq \emptyset$ for all $n \geq N_1$ and $T^n B \cap V \neq \emptyset$ for all $n \geq N_2$. This actually shows that (X, T) being t.s.m. implies $(X \times X, T \times T)$ is t.s.m., which in turn implies that $(X \times X, T \times T)$ is ergodic.

There were questions as to why one should consider ergodicity of the Cartesian square. An alternative characterization of t.w.m., for minimal systems, is that there are no nonconstant eigenfunctions: if f is continuous and $f \circ T = \lambda f$ for some constant λ, then f is constant. The definition of measure-theoretic weak mixing in ergodic theory is that

$$\frac{1}{n} \sum_{k=0}^{n-1} \left| \mu(T^k A \cap B) - \mu(A)\mu(B) \right| \longrightarrow 0 \,\text{as}\, n \to \infty$$

for all measurable sets A, B. See p.10 of Petersen's Lectures on Ergodic Theory for this and other equivalent characterizations of measure-theoretic weak mixing.

Measure-theoretic strong mixing is defined by

$$\mu(T^{-n} A \cap B) \to \mu(A)\mu(B)$$

for all measurable sets A, B. Thus, mixing properties concern asymptotic independence. Thinking in terms of *joinings*, we study how or whether the joinings ν_n (measures on $X \times X$) defined by

$$\nu_n(A \times B) = \mu(T^{-n} A \cap B)$$

approach the independent joining

$$\nu(A \times B) = \mu(A)\mu(B).$$

See p. 4 of the Lectures for a discussion of joinings. Anyway, thinking along these lines makes one want to study properties of $T \times T$ on $X \times X$, such as ergodicity.

6.3.3. *Topological Entropy,* h_{top}. For the general definition of topological entropy, see the reference works. For subshifts, the definitions are simpler.

Definition 6.3. Let (X, σ) be a subshift of some (Σ_a, σ). For each $n = 1, 2, \dots$, let $N_n(X)$=card$(\mathcal{L}(X) \cap A^n)$, i.e., the number of n-blocks in $\mathcal{L}(X)$. Then we define the *topological entropy* of X to be

$$h_{\text{top}}(X, \sigma) = \lim_{n \to \infty} \frac{1}{n} \log N_n(X).$$

Note that this limits exists because $N_{n+m} \leq N_n N_m$ so that $\log N_n$ is *subadditive* and therefore $\lim(1/n) \log N_n = \inf\{(1/n) \log N_n\}$ exists.

The question arises, how many allowable words of length $n+m$ can we make by concatenating allowable words of length n with allowable words of length m? Here we can think of $h_{\text{top}}(X,\sigma)$ intuitively as a measure of any of the following equivalent concepts: the exponential growth rate of the number of words in the language, the "concatenability index" of $\mathcal{L}(X)$, the "freedom of speech" permitted in the system, or the possibility of "saying something new." We have $\text{card}(\mathcal{L}(X)\cap A^n) \sim \exp^{n(h_{\text{top}}(X,\sigma))}$. For example, for the full shift, $h_{\text{top}}(\Sigma_2) = \lim \frac{1}{n}\log(2^n) = \log 2$. On the other hand, h_{top}(orbit closure of the PTM sequence)=0. Therefore, there is no "freedom of speech" permitted in this system; in other words, after a while the possibility of seeing something new in the sequence is next to nothing. This makes sense for the PTM sequence, since once we've seen one of the basic blocks of length 2^r, we are locked into one of two possibilities for the next string of 2^r steps.

7. January 29 *(Notes by SB)*

The PTM sequence is syndetically recurrent, but not periodic. In fact, Thue proved that it is *cube free* and Hedlund-Morse proved the even stronger non-periodicity condition that the PTM sequence does not contain BBb for any block $B = b\dots$. The fact that the PTM sequence exists and has this lack of periodicity has a useful application to avoiding infinite looping in iterative procedures. For example, certain chess rules consider a game a draw once a sequence of moves is repeated twice and begun for a third time. However, the *cube-free* property of the PTM sequence demonstrates that it is possible to have an infinite game, i.e., one which never ends, not even by a declared draw.

7.1. **Invariant Measures.**

Definition 7.1. A *measure* is a countably-additive, nonnegative (or sometimes signed) set function.

Definition 7.2. Let X be a compact metric space, $T : X \to X$ a homeomorphism. An *invariant measure* for (X, T) is a probability measure (i.e., $\mu(X) = 1$) defined on the Borel sets $\mathcal{B}(X)$ of X (i.e., the smallest σ-algebra containing the open sets) such that

$$\mu(T^{-1}A) = \mu(A) \text{ for all } A \in \mathcal{B}(X).$$

Proposition 7.1. *There is always at least one invariant measure on* (X, T)*.*

Before we begin the proof of this proposition, we establish a few facts and definitions

Definition 7.3. *Cesaro operators*: $A_n = \frac{1}{n}\sum_{k=0}^{n-1} T^k$.

Given a continuous map $T : X \to X$, we have a map $T : \mathcal{C}(X) \to \mathcal{C}(X)$ such that $(Tf)(x) = f(Tx)$ and its adjoint map $T : \mathcal{C}(X)^* \to \mathcal{C}(X)^*$. Note that each of these maps is called T even though each acts on a different space. Here $\mathcal{C}(X)^*$ refers to the *dual* vector space (i.e., the vector space of all continuous linear maps $\mathcal{C}(X) \to \mathbb{R}$) of the vector space $\mathcal{C}(X)$. By the Riesz Representation Theorem, $\mathcal{C}(X)^* = \mathcal{M}(X)$, the space of signed Borel measures on X. If $\mu \in \mathcal{C}(X)^* = \mathcal{M}$, then $T\mu$ is defined by $T\mu(f) = \mu(Tf) = \int f \circ T \, d\mu$ for all $f \in \mathcal{C}(X)$. Note that $\mathcal{M}(X)$ is a complete metric space with the *weak* topology* defined by $\mu_n \to \mu$ if and only if $\mu_n(f) \to \mu(f)$ for all $f \in \mathcal{C}(X)$.

Proof. We want to make invariant measures using the Cesaro operators. The set of probability measures is a *compact* set in $\mathcal{M}(X)$. Take any probability measure μ_0 on X, say $\mu_0 = \delta_x$ for some $x \in X$: for any continuous function f, $\int f \, d\mu_0 = f(x)$ for all $f \in \mathcal{C}(X)$, or

$$\mu_0(A) = \begin{cases} 0 & \text{if } x \notin A, \\ 1 & \text{if } x \in A. \end{cases}$$

Note that each $A_n\mu_0$ is also a probability measure and that

$$A_n\mu_0(f) = \frac{1}{n}\sum_{k=0}^{n-1}\int f \circ T^k d\mu_0 = \int \frac{1}{n}\sum_{k=0}^{n-1} f \circ T^k d\mu_0.$$

Let μ be a weak* limit point of $\{A_n\mu_0\}$, say $A_n\mu_0 \to \mu$. Then we can show that μ is invariant. Since we can approximate characteristic functions by continuous functions, to show that $\mu(A) =$

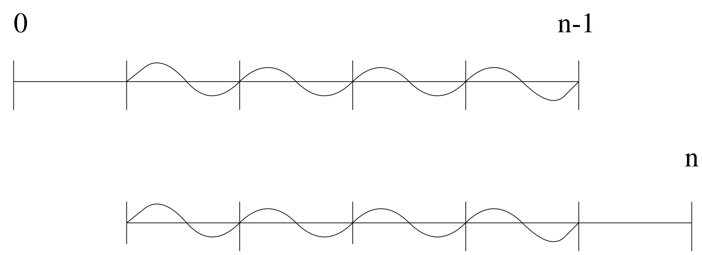

FIGURE 17. Shifting intervals $[0, n-1]$ by small amounts causes heavy overlap.

$\mu(T^{-1}A)$ for all $A \in \mathcal{B}(X)$ it suffices to show that $\int f\, d\mu = \int f \circ T\, d\mu$ for all $f \in \mathcal{C}(X)$. But

$$\int f \circ T d\mu = \mu(Tf) = \lim_j \frac{1}{n_j} \sum_{k=0}^{n_j - 1} \mu_0(fT^{k+1})$$

$$= \lim_j \frac{1}{n_j} \sum_{k=1}^{n_j} \mu_0(fT^k), \text{ while}$$

$$\int f d\mu = \mu(f) = \lim_j \frac{1}{n_j} \sum_{k=0}^{n_j - 1} \mu_0(fT^k),$$

so that $|\mu(Tf) - \mu(f)| = \lim(1/n_j)\, |\mu_0(fT^{n_j}) - \mu_0(f)| = 0$, and hence $\mu(Tf) = \mu(f)$.

□

Remarks 7.1.

(1) It is the nearly abelian property of the acting group that gives invariant measures in this way. More specifically, intervals $[0, n-1]$ in $\mathbb{N}$ form a Følner sequence: a slight translation of one overlaps it heavily. See Figure 17.
(2) If μ is an invariant measure for (X, T), then $(X, \mathcal{B}(X), \mu, T)$ is a *measure-theoretic dynamical system*, one of the fundamental objects of study in *ergodic theory*.

7.1.1. *Some easy examples of invariant measures.*

Example 7.1. *Bernoulli measures on full shifts* $(\Sigma_a, \sigma), a \geq 2$
Put any probability measure on the alphabet $A = \{0, 1, \ldots, a-1\}$, i.e., choose weights $\{p_0, p_1, \ldots, p_{a-1}\}$ for the symbols in A such that $\sum p_j = 1$. Let μ be the corresponding *product measure* on Σ_a. Then the measure of a cylinder set $[B]_m = \{x \in \Sigma_a : x_m x_{m+1} \cdots x_{m+l(B)-1} = B\}$ is $\mu[B]_m = p_{b_1} p_{b_2} \cdots p_{b_r}$, where $B = b_1 b_2 \ldots b_r$. On Σ_2 look at $\mathcal{B}(1/3, 2/3)$, i.e., the measure with weights 1/3 for 0 and 2/3 for 1. Then $\mu(\ldots\ldots.001\ldots) = (1/3)(1/3)(2/3)$. So these measures correspond to finite-state, independent, identically-distributed (stationary) stochastic processes (IID). By stationary, we mean that the measure is shift-invariant, i.e., that the probabilistic laws don't change with time. Independence, for example, in a coin-flipping experiment, means that one outcome does not affect any other one. This is not a property that holds for all types of repeated experiments.

Example 7.2. Rotation on a compact abelian group
Let $G = \{\alpha, \beta, \gamma\} = \mathbb{Z} \pmod 3 = \{0, 1, 2\}$ and let $Tg = g + 1 \pmod 3$. See Figure 19 We can

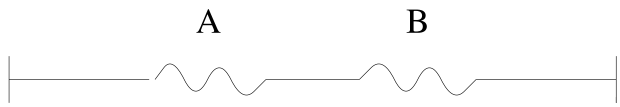

FIGURE 18. $P(A \cap B) = P(A)P(B)$, i.e., the probability that the cylinder sets appears in the places shown is equal to the product of the probability of the cylinder sets appearing on their own.

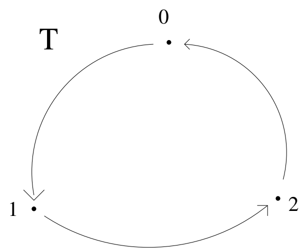

FIGURE 19. Here $0 = \alpha$, $1 = \beta$, $2 = \gamma$.

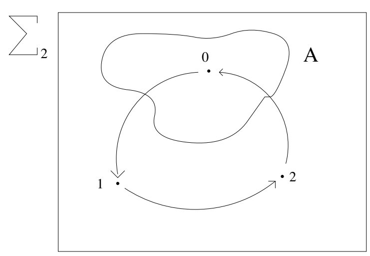

FIGURE 20. Here $0 = \alpha$, $1 = \beta$, $2 = \gamma$.

think of this example in the Σ_2 setting using a periodic sequence having period 3:

$$\alpha = \dots \dot{0}01001001 \dots$$
$$\beta = \dots \dot{0}1001001 \dots$$
$$\gamma = \dots \dot{1}001001 \dots ;$$
$$\text{then}\, \sigma\alpha = \beta, \sigma\beta = \gamma, \sigma\gamma = \alpha.$$

Here there is just one invariant probability measure, the one that puts $\mu(\alpha) = \mu(\beta) = \mu(\gamma) = 1/3$. Define μ on Σ_2 by $\mu = (1/3)\delta_\alpha + (1/3)\delta_\beta + (1/3)\delta_\gamma$, so that for a subset $A \subset X$, $\mu(A) = (1/3) \cdot$ (# of points α, β, γ in A). See Figure 20.

Then μ extends to a σ-invariant measure on Σ_2 by defining $\mu(\Sigma_2 \setminus \{\alpha, \beta, \gamma\}) = 0$.

8. February 3 *(Notes by SS)*

8.1. More Examples of Invariant Measures.

Example 8.1. Let G be a compact group. Then there is a unique Borel probability measure μ on G which is invariant under translation by elements of the group :

$$\mu(gA) = \mu(A) = \mu(Ag) \qquad \text{for any Borel } A \subset G \text{ and for any } g \in G.$$

This measure is called (normalized) *Haar measure*. Note that if G is a topological group, then $(g, h) \mapsto gh^{-1} : G \times G \to G$ is continuous. If g is any fixed element of G and we define a group rotation $T_g : G \to G$ by $T_g h = gh$ for any $h \in G$, then Haar measure μ is an invariant measure for T_g.

For example, if $G = S^1$ = unit circle in $\mathbb{C}$ with multiplication, so that $G \cong [0, 1)$ with addition mod 1, then each map $T_\alpha : [0, 1) \to [0, 1)$ by $T_\alpha x = x + \alpha$ (or $T_\alpha : S^1 \to S^1$ by $T_\alpha z = e^{2\pi i\alpha} z$) preserves Haar measure = Lebesgue measure. If α is irrational, then $([0, 1), T_\alpha)$ is minimal, i.e., every orbit is dense. Once one orbit is dense (i.e., topological ergodicity), then every orbit is dense, because in this case orbits are just translates of one another : $\{x + n\alpha : n \in \mathbb{Z}\} = \{y + n\alpha : n \in \mathbb{Z}\}$ $+(x - y)$. If $\alpha \in \mathbb{Q}$, then $([0, 1), T_\alpha)$ is not necessarily minimal.

More generally, suppose G is a compact group for which there exists $g \in G$ with $\{g^n : n \in \mathbb{Z}\}$ dense (the orbit of the identity is dense under T_g). Such a G is called *monothetic* and is necessarily abelian. For example, $[0, 1)$ with addition mod 1 is monothetic, since $\{n\alpha : n \in \mathbb{Z}\} = [0, 1)$ for any $\alpha \notin \mathbb{Q}$. If G is monothetic and g is a generator, then (G, T_g) is minimal and (normalized) Haar measure is the only T-invariant (Borel probability) measure on G. For, if μ is T_g-invariant, then it is T_{g^n}-invariant for any n, so it is T_h-invariant for any $h \in G$. Hence it is Haar measure. In Example 6.2, μ on G is the unique T-invariant measure. See Figure 19, page 26. Also, $\mathbb{T}^2 \cong [0, 1) \times [0, 1)$ with coordinatewise addition mod 1 is monothetic and so is $\mathbb{T} \times \cdots \times \mathbb{T}$ with coordinatewise addition mod 1.

Example 8.2. Let G be a compact group, μ Haar measure, and $T : G \to G$ a (continuous) endomorphism or automorphism. Then T preserves μ. For, if μ is translation-invariant, then μT^{-1} is also translation invariant and so must be Haar measure. Hence $\mu T^{-1} = \mu$.

For example, let $T : [0, 1) \to [0, 1)$ be $Tx = 2x$ mod 1. Then T is not a translation, but T preserves Lebesgue measure. Also

$$(S^1, z \mapsto z^2, \nu) \cong ([0, 1), T, m) \overset{\phi}{\leftarrow} (\Sigma_2^+, \sigma, \mathcal{B}(\tfrac{1}{2}, \tfrac{1}{2})),$$

where m is Lebesgue measure on $[0, 1)$ and ν is Lebesgue measure on S^1. The map ϕ is one-to-one except on a countable set, where it is two-to-one. For, if $x \in [0, 1)$, then $x = .x_1x_2x_3 \cdots = \sum_{k=1}^{\infty} x_k/2^k$, each $x_k = 0$ or 1, and x corresponds to a point $(x_1x_2x_3 \cdots)$ in Σ_2^+. Then $Tx = 2x$ mod $1 = .x_2x_3x_4 \cdots$ corresponds to a point $\sigma(x_1x_2x_3 \cdots)$ in Σ_2^+, and so T corresponds to the shift σ on Σ_2^+. The expansion $.x_1x_2x_3 \cdots$ of $x \in [0, 1)$ is obtained by following the orbit $x, 2x, 4x, \cdots$ (mod 1) of x and writing down 0 when the point is in $[0, \frac{1}{2})$ and 1 when the point is in $[\frac{1}{2}, 1)$. Similarly, for endomorphisms or automorphisms of a torus, e.g., $T\binom{x}{y} = \left(\begin{smallmatrix} 2 & 1 \\ 1 & 1 \end{smallmatrix}\right)\binom{x}{y}$ on $[0, 1) \times [0, 1)$ (mod 1) preserves Haar measure = Lebesgue measure on $\mathbb{T}^2$.

On (Σ_2, σ), there are many σ-invariant measures; for example, each periodic orbit supports one. In fact, each subshift (X, σ) supports at least one, and there are lots of Bernoulli and Markov measures and so on.

8.2. **Unique Ergodicity.** Let (X, T) be a compact topological dynamical system and $\mathcal{I}(X, T)$ the set of T-invariant (Borel probability) measures on X. Then $\mathcal{I}(X, T)$ is a compact, convex subset of $\mathcal{M}(X) = \mathcal{C}(X)^*$ = vector space of all (signed) Borel measures on X with the weak* topology. For convexity, we can easily see that if $\mu, \nu \in \mathcal{I}(X, T)$, $0 \le t \le 1$, then $t\mu + (1-t)\nu \in \mathcal{I}(X, T)$.

The *extreme points* of $\mathcal{I}(X, T)$ (i.e., the ones which cannot be written as $t\mu + (1-t)\nu$ for some $\mu, \nu \in \mathcal{I}(X, T)$, $\mu \ne \nu$ and $0 < t < 1$) are the *ergodic measures* for (X, T) (i.e., the ones for which $(X, \mathcal{B}(X), T, \mu)$ is ergodic (measure-theoretically), in that every *invariant set* A (i.e., $\mu(A \bigtriangleup T^{-1}A) = 0$) has measure 0 or 1). The theorem of Krein-Milman says that there always exists an extreme point. So, the set of ergodic measures is not empty. If (X, T) has only one invariant (Borel probability) measure, then by the theorem of Krein-Milman, it is an extreme point, hence ergodic.

Proposition 8.1. *If (X, T) has only one invariant (Borel probability) measure, then that measure is ergodic.*

Definition 8.1. If (X, T) has only one invariant, hence ergodic, measure, then we say that (X, T) is *uniquely ergodic.*

For example, $[0, 1)$ with translation mod 1 by an irrational is uniquely ergodic. In Example 6.2, (G, T) is uniquely ergodic. However, if $G = \{a, b, c\} \cong \mathbb{Z}_3$ and T is defined by $T(a) = a, T(b) = c$, and $T(c) = b$, then T is not a group rotation, and it is not uniquely ergodic, since there are two ergodic measures. For, if μ is defined by $\mu(\{a\}) = 1, \mu(\{b\}) = \mu(\{c\}) = 0$, and ν is defined by $\nu(\{a\}) = 0, \nu(\{b\}) = \nu(\{c\}) = \frac{1}{2}$, then they are ergodic, and every other invariant measure is a convex combination $t\mu + (1-t)\nu, 0 < t < 1$, i.e., μ and ν are extreme points. Thus, $\mathcal{I}(X, T) \cong [0, 1]$.

Remark 8.1. (Choquet Simplex Theory) In a certain sense, every point of a compact convex set (in the right sort of space) is a "convex combination" (some kind of integral) of extreme points.

To find out which subshifts are uniquely ergodic, we need the following theorem.

Theorem 8.2. *(Ergodic Theorem) Let $(X, \mathcal{B}, \mu)$ be a probability space and $T : X \to X$ a measure-preserving transformation (i.e., T is 1-1 and onto up to sets of measure 0, $T\mathcal{B} = T^{-1}\mathcal{B} = \mathcal{B}$ and $\mu T^{-1} = \mu T = \mu$) and f an integrable function on X (i.e., $f : X \to \mathbb{R}$ is measurable and $\int_X |f| d\mu < \infty$). Then $A_n f(x) = \frac{1}{n} \sum_{k=0}^{n-1} f(T^k x)$ converges as $n \to \infty$ for a.e. $x \in X$.*

9. February 5 *(Notes by SS)*

9.1. Unique Ergodicity.

Remark 9.1. The description of the invariant and ergodic measures for a topological dynamical system (X, T), uniquely ergodic etc., based on the Ergodic Theorem and Riesz Representation Theorem ($\mathcal{C}(X)^* =$ signed Borel measures on X) is due to Krylov and Bogolioubov. Also see Nemytskii and Stepanov, "Qualitative Theory of Differential Equations" (1960), or J. Oxtoby, "Ergodic sets" (1952), Bull. AMS.

Theorem 9.1. *For a compact topological dynamical system (X, T), the following three conditions are equivalent.*

(1) (X, T) is uniquely ergodic.

(2) For all $f \in \mathcal{C}(X)$, $A_n f(x) = \frac{1}{n}\sum_{k=0}^{n-1} f(T^k x)$ converges as $n \to \infty$ uniformly on X to a constant (the integral of f with respect to the unique invariant measure).

(3) For all $f \in \mathcal{C}(X)$, some subsequence of $\{A_n f(x)\}$ converges pointwise on X to a constant.

Proof. Recall that according to the Ergodic Theorem, if μ is an invariant measure on (X, T) and $f \in \mathcal{C}(X) \subset L^1(X, \mathcal{B}(X), \mu)$, then $A_n f(x)$ converges to some limit function $\bar{f}(x)$ a.e. dμ. If μ is ergodic, then $\bar{f}(x)$ is the constant $\int_X f d\mu$ a.e..

(1) $\Rightarrow$ (2) : Assume (2) does not hold. There is at least one invariant measure for (X, T), call it μ_0. Since (2) doesn't hold, there exists $f \in \mathcal{C}(X)$ such that $A_n f(x)$ does not converge uniformly to $\int_X f d\mu_0$. So, there exist $\delta > 0, x_k \in X, n_k \in \mathbb{N}$ such that

$$|A_{n_k} f(x_k) - \textstyle\int_X f d\mu_0| \geq \delta \qquad \text{for all } k.$$

Then $\lim_{k\to\infty} A_{n_k} h(x_k)$ exists for all $h \in \mathcal{C}(X)$. For, take a countable dense set $\{g_1, g_2, \cdots\}$ in $\mathcal{C}(X)$. By passing to a subsequence of $\{n_k\}$, we may assume that $A_{n_k} g_j(x_k) \to \lambda_j \in \mathbb{R}$ as $k \to \infty$ for all j. Then $A_{n_k} h(x_k) = A_{n_k} g_j(x_k) + A_{n_k}(h - g_j)(x_k)$, and $A_{n_k} g_j(x_k) \to \lambda_j$ as $k \to \infty$; and also if $\|h - g_j\|_\infty < \epsilon$, then $|A_{n_k}(h - g_j)(x_k)| < \epsilon$. Thus, if $g_{j_s} \to h$ in $\mathcal{C}(X)$, then $A_{n_k} h(x_k) \to \lim_{s\to\infty} \lambda_{j_s}$ as $k \to \infty$.

Then, defining $\lambda(h) = \lim_{k\to\infty} A_{n_k} h(x_k)$ for $h \in \mathcal{C}(X)$ defines a positive normalized continuous linear functional on $\mathcal{C}(X)$, i.e., $h \geq 0$ implies $\lambda(h) \geq 0$, $\lambda(1) = 1$, $\lambda(ch) = c\lambda(h)$, and $\lambda(h_1 + h_2) = \lambda(h_1) + \lambda(h_2)$. For continuity, note that if $\|h_1 - h_2\|_\infty$ is small, then $|\lambda(h_1) - \lambda(h_2)|$ is small.

Then by the Riesz Representation Theorem, there exists a unique Borel probability measure μ on X such that $\lambda(h) = \int_X h d\mu$ for all $h \in \mathcal{C}(X)$. The measure μ is T-invariant, since as before, $\lambda(h \circ T) = \lambda(h)$ for all $h \in \mathcal{C}(X)$. But $\mu \neq \mu_0$, since

$$\textstyle\int_X f d\mu = \lambda(f) = \lim_{k\to\infty} A_{n_k} f(x_k)$$

where n_k is now a subsequence of the original one, and so

$$\left|\textstyle\int_X f d\mu - \int_X f d\mu_0\right| \geq \delta.$$

Hence (X, T) is not uniquely ergodic.

(2) $\Rightarrow$ (3) is clear.

(3) ⇒ (1) : Suppose μ, ν are two different ergodic invariant measures for (X, T). Choose $f \in \mathcal{C}(X)$ with $\int_X f d\mu \neq \int_X f d\nu$. By the Ergodic Theorem, $A_n f(x) \to \int_X f d\mu$ for μ-a.e. $x \in X$, say $x \in X_\mu$, and $A_n f(x) \to \int_X f d\nu$ for ν-a.e. $x \in X$, say $x \in X_\nu$. So, $X_\mu \cap X_\nu = \emptyset$, and $\mu(X_\mu) = 1 = \mu(X)$, and $\nu(X_\nu) = 1 = \nu(X)$. Then it is clearly impossible that any subsequence of $\{A_n f(x)\}$ converges pointwise on X to a constant (since $X_\mu \neq \emptyset$ and $X_\nu \neq \emptyset$). □

Remark 9.2. It is enough to check condition (2) for f in a countable dense set in $\mathcal{C}(X)$.

9.1.1. *Interpretation of the unique ergodicity criterion in (2) for subshifts.*

Theorem 9.2. *A topologically transitive subshift* $(\overline{\mathcal{O}(x)}, \sigma)$ *is uniquely ergodic if and only if every block that appears in* x *appears with a uniform limiting frequency: i.e., given any block* B *that appears in* x*, there is* $\lambda(B)$ *such that given* $\epsilon > 0$*, there is* $L > 0$ *such that if* C *is any block that appears in* x *with* $l(C) \geq L$*, then*

$$\left| \tfrac{\nu(B,C)}{l(C)} - \lambda(B) \right| < \epsilon,$$

where $\nu(B, C) =$ *the number of times that* B *appears in* C*, i.e.,* $\text{card}\{j : 1 \leq j \leq l(C) - l(B),\ c_j c_{j+1} \cdots c_{j+l(B)-1} = B\}$.

Proof. Condition (2) of the previous theorem for the characteristic function $\chi_{[B]_m}$ of a basic cylinder set says that $A_n \chi_{[B]_m}$ converges to a constant $\lambda(B)$ uniformly on $\overline{\mathcal{O}(x)}$. Note that

$$A_n \chi_{[B]_0}(x) = \tfrac{1}{n} \textstyle\sum_{k=0}^{n-1} \chi_{[B]_0}(\sigma^k x) = \tfrac{1}{n} \times \text{(the number of } B\text{'s seen in } x_0 \cdots x_{n+l(B)-1}).$$

So, to say $A_n \chi_{[B]_0}$ converges to $\lambda(B)$ uniformly on $\overline{\mathcal{O}(x)}$ says that if C is any long enough block in x, then $\left|\nu(B, C)/l(C) - \lambda(B)\right|$ is small. Now, linear combinations of such characteristic functions with rational coefficients are dense in $\mathcal{C}(X)$, so if (2) holds for the $\chi_{[B]_m}$, it holds for all $f \in \mathcal{C}(X)$. □

9.1.2. *Connection between minimality and unique ergodicity.*

Example 9.1. *A non-minimal uniquely ergodic orbit closure in* (Σ_2, σ)

Let $x = \cdots 000.100 \cdots$ and so $\overline{\mathcal{O}(x)} = \mathcal{O}(x) \cup \{\cdots 000 \cdots\}$. Then the point $z = \cdots 000 \cdots$ fails to have a dense orbit. So $(\overline{\mathcal{O}(x)}, \sigma)$ is not minimal. But δ_z is the unique invariant measure; $\delta_z([B]_m) = 1$ if $B = 0^{l(B)}$, and 0 if $B = 0 \cdots 010 \cdots 0$. Thus $\overline{\mathcal{O}(x)}$ is uniquely ergodic. Also, it is topologically transitive.

For a uniquely ergodic system that is not topologically transitive, let $y = \cdots 000.1100 \cdots$ and $X = \overline{\mathcal{O}(y)} \cup \overline{\mathcal{O}(x)}$ (with $x = \cdots 000.100 \cdots$ as before). Then (X, σ) is uniquely ergodic, but not topologically transitive.

Examples of minimal systems that are not uniquely ergodic were given first by Markov, then by Oxtoby.

10. February 10 *(Notes by KN)*

In the following sections we will mention several more properties of topological dynamical systems and consider for which subshifts these properties hold.

10.1. **Expansive Systems.**

Definition 10.1. A topological dynamical system is called *expansive* if there is $\beta > 0$ such that whenever $x, y \in X$ with $x \neq y$ then there is $n \in \mathbb{Z}$ such that $d(T^n x, T^n y) \geq \beta$. We call β the *expansive constant.*

Notice that every subshift is expansive. This is because if $x \neq y$, then there is some first place where they differ, say $x_n \neq y_n$, and then $d(\sigma^n x, \sigma^n y) = 1$.

Theorem 10.1. *(Reddy) If (X, T) is an expansive topological dynamical system, then (X, T) is a factor of some subshift in a space of sequences on a finite alphabet. In other words, there exists a subshift on a finite alphabet, $(\Sigma, \sigma) \subset (\Sigma_a, \sigma)$ and a factor map $\phi : \Sigma \to X$ so that the following diagram commutes:*

$$\begin{array}{ccc} \Sigma & \xrightarrow{\sigma} & \Sigma \\ {\scriptstyle\phi}\downarrow & & \downarrow{\scriptstyle\phi} \\ X & \xrightarrow{T} & X \end{array}$$

Moreover, if (X, T) is expansive and 0-dimensional (totally disconnected), then (X, T) is topologically conjugate to a subshift on a finite alphabet.

Theorem 10.1 implies that any compact expansive topological dynamical system has finite entropy.

10.2. **Equicontinuous Systems.**

Definition 10.2. A topological dynamical system (X, T) is *equicontinuous* if $\{T^n : n \in \mathbb{Z}\}$ is equicontinuous, i.e., given $\epsilon > 0$ there is $\delta > 0$ such that if $x, y \in X$ and $d(x, y) < \delta$, then $d(T^n x, T^n y) < \epsilon$ for all $n \in \mathbb{Z}$.

10.2.1. *Some examples of equicontinuous systems.*

(1) If $T : X \to X$ is an *isometry* (i.e. preserves distance), then (X, T) is equicontinuous.
(2) If (X, T) is a rotation on a compact (abelian) group, then it is equicontinuous because such a group always has an equivalent invariant metric: $d(h_1, h_2) = d(gh_1, gh_2)$ for all $g, h_1, h_2 \in G$. (See Kelley's General Topology book for more).
(3) $x \to x + \alpha \pmod 1$ on $[0, 1)$.

10.2.2. *Equicontinuous subshifts.* Which subshifts are equicontinuous? As we just saw, subshifts are expansive and it is not easy to be *both* expansive and equicontinuous. We need to find a δ so that whenever $d(x, y) < \delta$, then $d(T^n x, T^n y) < \epsilon$ for all n. Suppose we are given $\epsilon = \beta$, the expansive constant. Then we need to be sure there are no points $x \neq y$ with $d(x, y) < \delta$. This condition requires a finite set of points. An example of an expansive equicontinuous subshift is $(\overline{\mathcal{O}(x)}, \sigma)$, where $x = \ldots 101010 \ldots$. Here, $\beta = 1$. Given any ϵ, we may take $\delta = 1/2$. Then whenever $d(x, y) < \delta = 1/2$, we have $x = y$, so equicontinuity is clear. Equicontinuity also follows because this system is a rotation on a compact group.

10.3. **Distal Systems.**

Definition 10.3. (X,T) is *distal* if there are no *proximal pairs*, i.e. pairs of points $x \neq y$, for which there exists $n_k \to \infty$ or $(-\infty)$ so that $d(T^{n_k}x, T^{n_k}y) \to 0$.

Again, we find that the only distal subshifts are the finite subshifts. Gottschalk and Hedlund actually show that if (X,T) is compact expansive and X is dense in itself (X has no limit points), then there exists an *asymptotic pair* $\{x,y\}$, that is $x \neq y$ with $d(T^n x, T^n y) \to 0$.

10.4. **The structure of the Morse System.** Recall the PTM sequence: Let $\omega = \ldots.01101001\ldots$ and $M = \overline{\mathcal{O}(\omega)}$. Then (M,σ) is the Morse minimal set.

Theorem 10.2. *(M,σ) is (measure-theoretically) a skew product over a group rotation. More precisely, there are a compact abelian group G, an element $\theta \in G$ with $\overline{\{n\theta : n \in \mathbb{Z}\}}$, a uniquely ergodic subshift $(X,\sigma) \subset (\Sigma_2,\sigma)$ which is isomorphic to (G,R_θ) (with Haar measure), and a continuous function $f : X \to \{-1,1\}$ such that (M,σ) is topologically conjugate to the dynamical system $(X \times \{-1,1\}, \hat{T})$, with $\hat{T}$ defined by $\hat{T}(x,\xi) = (\sigma x, f(x)\xi)$.*

Theorem 10.2 implies that the Morse minimal set is *nearly* a group rotation: (M,σ) is a 2-1 extension of a group rotation. (S. Kakutani, 5th Berkeley Symposium, 1969; W.A. Veech, TAMS,1969). From this structure follow easily properties of the Morse minimal set such as unique ergodicity, entropy 0, the spectrum, etc. A similar analysis is possible for more generalized Morse sequences (Goodson, Lemancyk, etc.).

10.4.1. *The group rotation.* The underlying group rotation of (M,σ), denoted by (G, R_{g_o}), is called the *odometer* or *von Neumann-Kakutani adding machine.*

The odometer is defined on $\Sigma_2^+ = \{0,1\}^{\mathbb{N}}$. Let $T(x_1x_2x_3\ldots) = (x_1x_2x_3\ldots) + (1000\ldots)$ with carry to to the right. For example,

$$
\begin{aligned}
T(001100\ldots) &= (101100\ldots)\\
T^2(001100\ldots) &= (011100\ldots)\\
T^3(001100\ldots) &= (111100\ldots)\\
T^4(001100\ldots) &= (000010\ldots)
\end{aligned}
$$

An alternate way to state the rule:

$$\text{if } x = x_1x_2x_3\ldots \text{ and } n = \text{ the first } k \text{ such that } x_k = 0, \text{ then } Tx = \underbrace{0\ldots0}_{n-1}\underbrace{1}_{nth}x_{n+1}x_{n+2}\ldots$$

Notice that the orbit of $g_o = .1000\ldots$ is dense in ${\Sigma_2}^+$:

$$
\begin{aligned}
Tg_o &= 010000\ldots\\
T^2g_o &= 110000\ldots\\
T^3g_o &= 001000\ldots\\
T^4g_o &= 101000\ldots\\
T^5g_o &= 011000\ldots\\
T^6g_o &= 111000\ldots\\
T^7g_o &= 000100\ldots
\end{aligned}
$$

In general,

$$T^n g_o = \alpha_0 \alpha_1 \ldots \alpha_r \text{ where}$$

$n = \alpha_0 + \alpha_1 2 + \ldots + \alpha_r 2^r$. The initial block of $T^n g_o$ is the binary representation of $n+1$. Notice that T is an isometry, i.e. $d(Tx, Ty) = d(x, y)$. Hence, (Σ_2^+, T) is minimal, since topologically ergodic isometries are minimal. To see this, take $g_o = .1000\ldots$ and $g, h \in \Sigma_2^+$. We want some n so that $T^n g$ is arbitrarily close to h, say within some ϵ. We can get this because $T^r g_o \approx h$ for some r and $T^s g_o \approx g$ for some s. Thus, $n = r - s$ gives $T^n g = T^r T^{-s} g \approx h$.

This leads us to a result in the more general case.

Theorem 10.3. *Every topologically ergodic isometry is (topologically conjugate to) a rotation by a generator on a compact monothetic group.*

The odometer may also be represented as a subshift (isomorphic up to countably many points). First we form a *Toeplitz sequence*, τ, as follows: First fill τ in with 0's followed by blanks. Then fill in the remaining blanks with 1's followed by blanks. Repeat this alternating process until all blanks are filled in:

$$\begin{array}{rcl} \tau & = & \ldots\ldots 0\_0\_0\_0\_0\_0\_0\_0\_ \\ & & \ldots\ldots\ 1\ \_\ 1\ \_\ 1\ \_\ 1\ \_ \\ & & \ldots\ldots\quad 0\ \_\quad 0\ \_\quad 0\ \_ \end{array}$$

Let $\mathcal{L}$ be the language consisting of all blocks in the 1-sided sequence, τ, and let X be the 2-sided and X^+ be the 1-sided subshift determined by $\mathcal{L}$. The space (X, σ) is the *Toeplitz subshift.* (K. Jacobs & M. Keane, ZW, 1969). In fact, there exists a whole class of these sequences. For example we could start with any periodic sequence on 0 _ follow it by another one on 1 _, etc. The sequences may vary each time. (J. Neveu, ZW, 1969; E. Eberlein, Diplomarbeit, 1970; S. Williams, Thesis, 1981 and ZW, 1984; Downarowicz; Lemanczyk, Iwanik, ...).

11. February 12 *(Notes by KN)*

We want to discuss the relationship between the odometer and the Toeplitz subshift, and then use these systems to describe the structure of the Morse subshift.

11.1. **The Odometer.** Recall the *odometer* or *von Neumann-Kakutani adding machine* is a rotation on a compact monothetic group, $T : \Sigma_2^+ \to \Sigma_2^+$, defined as follows:

$$T(x_1x_2x_3\dots) = (x_1x_2x_3\dots) + (1000\dots)$$

with carry to to the right. For example,

$$\begin{aligned} T(001100\dots) &= (101100\dots) \\ T^2(001100\dots) &= (011100\dots) \\ T^3(001100\dots) &= (111100\dots) \\ T^4(001100\dots) &= (000010\dots) \end{aligned}$$

There are several equivalent descriptions of the odometer:

(1) The *2-adic integers* are sequences $y = .y_1y_2y_3\dots$ with $y_n \in \{0, 1, \dots, 2^n - 1\}$ and $y_{n+1} \equiv y_n \pmod{2^n}$. Define coordinatewise addition $\pmod{2^n}$ in the nth coordinate) and give it the product topology.

(2) We can think of the odometer in terms of formal series:

$$y = \sum_{j=0}^{\infty} x_j 2^j, \quad \text{where } \ x_j = 0 \text{ or } 1 \text{ and } \ y_n = \sum_{j=0}^{n-1} x_j 2^j.$$

In our analysis of the odometer we will use the definition, that is the description in terms of sequences, $x = x_0x_1x_2\dots$ in Σ_2^+. Notice that coordinatewise addition $\pmod{2^n}$ of the y's corresponds to addition $\pmod 2$ of the x's with carry to the right.

$G = \Sigma_2^+$ with this group operation is a compact, abelian, and monothetic group. $T : G \to G$ is given by $Tx = x + \theta$, and $\overline{\mathcal{O}(\theta)}$ is dense. Thus, (G, T) is a minimal and uniquely ergodic topological dynamical system.

(3) Adic realization of the odometer (Vershik).

We regard a point in Σ_2^+ as an infinite path in the graph in Figure 21(a). For example, the sequences $x = 11101\dots$ and $y = 00011\dots$ are shown in Figure 21(b). We say that two points are *comparable* if they eventually coincide. That is, x and y are *comparable* if there exists a (smallest) k such that $x_n = y_n$ for all $n > k$. Then we define a partial ordering $<$ by $x < y$ if $x_k < y_k$. In Figure 21(b), for example, $x < y$.

$T : \Sigma_2^+ \to \Sigma_2^+$ is defined by $Tx =$ the smallest $y > x$, if there is one, and $T(111\dots) = 000\dots$. We may generalize this adic representation of the odometer in many ways. For example, we can consider odometers with different size wheels. Then the space is $H = \prod_{n=0}^{\infty}\{0, 1, \dots, q_n - 1\}$, where $q_n \geq 1$, instead of just $\Sigma_2^+ = \prod_{n=0}^{\infty}\{0, 1\}$. as in the Figure 22(a). Another variation would be to restrict the paths on the graph, as in Figure 22(b). The advantage of the adic representation is the concrete combinatorial representation of the system, which sometimes provides better and perhaps easier ways of analyzing properties of the system. As we may see later, the system represented in Figure 22(b) is isomorphic to the Fibonacci substitution system, where the substitution ζ is defined by $\zeta(1) = 0$ and $\zeta(0) = 01$.

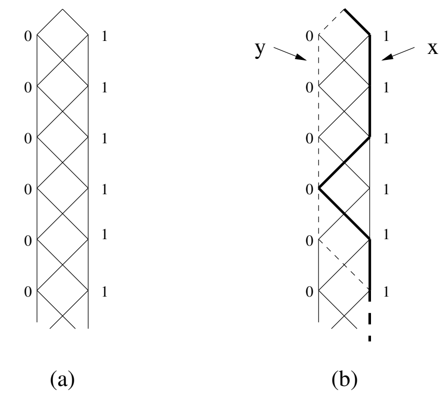

FIGURE 21. (a) The adic graph, (b) The paths $x = 11101\ldots$ and $y = 00011\ldots$.

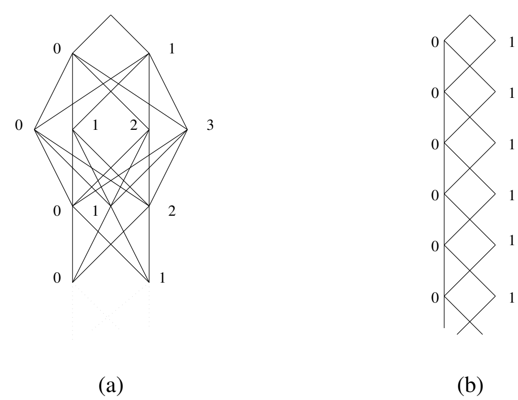

FIGURE 22. (a) Adic graph for H, (b) Adic graph with restricted paths.

11.1.1. *The spectrum of the odometer in* Σ_2^+. The ergodic properties of Σ_2^+ are fairly well understood since it is a minimal uniquely ergodic group rotation. It is known to have a purely discrete spectrum, that is, its eigenfunctions span L^2. In fact, its eigenfunctions are the continuous group characters (a character is a continuous homomorphism, $\phi : G \to \{z \in \mathbb{C} : |z| = 1\}$). The *dual group* or *character group* for Σ_2^+ is the group of dyadic rationals in $[0, 1) \approx S^1$, with the following duality:

$$(x, \hat{x}) = e^{(2\pi i)\hat{x}\sum_{j=0}^{\infty} \epsilon_j 2^j}, \quad \text{where } \hat{x} = \frac{p}{2^N}.$$

Thus if $G = \Sigma_2{}^+$ and m is the Haar measure, then the functions, $f \in L^2(G, m)$ and $(\lambda \in \mathbb{C})$, such that $f(Tx) = \lambda f(x)$ m-a.e. are just the functions $f = \hat{x}$ and $\lambda = \lambda_{\hat{x}} = e^{2\pi i \hat{x}}$ for some $\hat{x}$:

$$\begin{aligned} f(Tx) = \hat{x}(Tx) &= \hat{x}(x+\theta) \\ &= \hat{x}(x)\hat{x}(\theta) \\ &= \underbrace{\hat{x}(\theta)}_{\lambda}\underbrace{\hat{x}(x)}_{f(x)}, \\ \text{and } \hat{x}(\theta) = (\theta, \hat{x}) &= e^{2\pi i \hat{x}}. \end{aligned}$$

Example 11.1. Let $G = S^1$ and $Tz = e^{2\pi i\alpha}z$. Then $\hat{G} = \mathbb{Z}$ with the duality $(n, z) = z^n$. On $[0, 1)$ this is just $Tx = x + \alpha \mod 1$ with $(n, x) = e^{2\pi inx}$. So the eigenfunctions are $\phi_n(x) = e^{2\pi inx}$ with eigenvalues $\lambda_n = e^{2\pi in\alpha}$.

11.2. **Relationship between the Toeplitz subshift and orbit closure of the odometer.** The following theorem states the relationship between the odometer and the Toeplitz sequence. The connection is a near isomorphism between a minimal subshift and a compact group rotation, which allows us to ask and answer many more questions about the structure and properties of each system.

Theorem 11.1. *(Neveu, ZW 1969; E. Eberlein, Thesis) There is a continuous onto map from the (1-sided) Toeplitz subshift, $(X^+, \sigma) \subset (\Sigma_2^+, \sigma)$ to the odometer $(G, T) = (\Sigma_2^+, T)$, which is 1-1 except on a countable set of points on which it is 2-1. Consequently, the Toeplitz subshift is minimal and uniquely ergodic and has purely discrete spectrum.*

Proof. Define "not quite a mapping" $\psi : G \to \Sigma_2^+$ as follows: Let $x = x_1x_2x_3 \ldots \in G = \Sigma_2^+$. We will use this as a program to tell us how to get a Toeplitz-type sequence. Put

$$\psi x = \begin{cases} 0a_00a_10a_2\ldots & \text{if } x_1 = 0, \\ a_00a_10a_20\ldots & \text{if } x_1 = 1, \end{cases}$$

where

$$a_0a_1a_2\ldots = \begin{cases} 1b_01b_11b_2\ldots & \text{if } x_2 = 0\ , \\ b_01b_11b_21\ldots & \text{if } x_2 = 1, \end{cases}$$

and continue in this manner to define ψx.

This does not always give us a point, because it may happen that just one blank does not get filled in. However, if x has infinitely many 0's in it, then ψ is defined, since each 0 tells you to fill in the first blank from the left.

Let $G_o = \{x :$ there is K such that $x_k = 1$ for all $k \geq K\}$. If $x \in G_o$, then just one place in ψx is left undetermined. Fill this one place with both a 0 *and* 1, so that $\psi x = \{u, u'\}$ consists of two points

We will continue the proof next time.

□

12. February 17 *(Notes by RP)*

12.1. **Toeplitz System.** Recall the (one-sided) Toeplitz sequence ω. Let $X^+ \subset \Sigma_2^+$ denote the orbit closure of ω. Let G be Σ_2^+ with group operation addition $\bmod 2$ with carry to the right. Define $T : G \to G$ by $Tx = x + \theta$, where $\theta = 1000\dots$. We defined "not quite a mapping" $\psi : G \to X^+$ last time. If $x \notin G_0 = \{y \in \Sigma_2^+ : y_k = 0 \text{ for only finitely many } k\}$, then $\psi x \in \Sigma_2^+$ is defined. If $x \in G_0$, exactly one entry stays unfilled in ψx; fill it with *both* 0 and 1, so $\psi x = \{u, u'\} \subset \Sigma_2^+$ (two points). Notice that $\psi(000\dots) = \omega$.

ψ is "almost 1:1": Suppose ψx and ψy intersect; we claim that then $x_1 = y_1$. If $x_1 \neq y_1$, then

$$\psi x = 0a_0 0a_1 0\dots$$

$$\psi y = a_0' 0a_1' 0a_2' \dots.$$

Since some a_i or a_i' is 1, $\psi x \neq \psi y$. Now induct, starting with $a = a_0 a_1 \dots$ to get $x_2 = y_2$, etc.

This allows us to *define* $\phi = \psi^{-1} : \psi G \to G$. It also shows that this ϕ is *continuous*: if ψx and ψy agree on a long initial block, then so do x and y. Writing $\psi G = \psi G_0 \cup \psi(G \setminus G_0)$ and $G = G_0 \cup (G \setminus G_0)$, we have that $\phi : \psi G_0 \to G_0$ is 2:1 and $\phi : \psi(G \setminus G_0) \to G \setminus G_0$ is 1:1.

We claim that $T\phi = \phi\sigma$. It suffices to show that $\psi T = \sigma\psi$. For example, take

$$x = 111000\dots.$$

Then

$$Tx = 000100\dots,$$

and

$$\psi Tx = 010001000101\dots.$$

On the other hand,

$$\psi x = 1010001000101\dots,$$

and

$$\sigma\psi x = 010001000101\dots = \psi Tx.$$

Note also that equality holds when $x = 111\dots$, the sequence of all ones. In that case, we have $Tx = 000\dots$ (all zeros), and $\psi Tx = \omega$. On the other hand, $\psi x = z\omega$, where $z \in \{0, 1\}$. So $\sigma\psi x = \omega$ as well:

$$\begin{array}{ccc} (0 \text{ or } 1)|\omega & \xrightarrow{\sigma} & \omega \\ \uparrow{\scriptstyle\psi} & & \uparrow{\scriptstyle\psi} \\ x = 111\dots & \xrightarrow{T} & Tx = 000\dots \end{array}$$

Note that $\phi\psi = \mathrm{id}_G$ and $\psi\phi = \begin{cases} \text{id} & \text{on } \psi(G \setminus G_0), \\ \text{two points} & \text{on } \psi G_0. \end{cases}$

If $u = \psi x \in \psi G$, then $\phi\sigma u = Tx$. This action-commuting property is clarified further in 12.1, below.

Now we show that ψG is closed. Let $u \in \overline{\psi G} \subset \Sigma_2^+$. Say $v_n \in \psi x_n$, $v_n \to u \in \Sigma_2^+$. By compactness, we may assume $x_n \to x$. We claim that $u \in \psi x$ (or $= \psi x$). If x_n and x agree on a long initial block, then so do ψx_n and ψx, except possibly for a single blank, which must occur at the same (say j'th) place. If $x \notin G_0$, there is no such blank. If $x \in G_0$, take $(\psi x)_j = u_j$, since $(v_n)_j = u_j$.

Now $\sigma\psi G = \psi TG = \psi G$, so ψG is a *subshift*. The point ω is in $X^+ = \psi G$, so $X^+ \supset \overline{\mathcal{O}^+(\omega)}$. We claim that $\mathcal{O}^+(\omega)$ is *dense* in ψG, so $X^+ = \psi G = \overline{\mathcal{O}^+(\omega)}$. Let $u \in \psi G$, say $u \in \psi x$. Take n_k with $T^{n_k}\theta = n_k\theta \to x$ in G (odometer (G,T) is minimal). Note that $\psi T^{n_k}\theta = \sigma^{n_k}\psi\theta = \sigma^{n_k}\omega$. Hence by continuity of ψ, $n_k\theta \to x$ implies that $\psi\sigma^{n_k}\omega \to \{u,?\}$:

$$\begin{array}{ccc} \sigma^{n_k}\omega & \longrightarrow & \{u,?\} \\ \uparrow \psi & & \uparrow \psi \\ n_k\theta & \longrightarrow & x \end{array}$$

It is possible to choose n_k so that any residual blank in u is filled in correctly–see 12.1, below. Consequently, (X^+,σ) is minimal and *uniquely ergodic.* m = Haar measure on G is the unique invariant measure for (G,T). Define μ on X by $\mu(A) = m(\phi A)$.

$$\begin{array}{rcc} & \psi G_0 \cup \psi(G \setminus G_0) & \mu, \nu \\ \text{measure } 0 \downarrow & \updownarrow\approx & \updownarrow \\ & G_0 \cup (G \setminus G_0) & m \end{array}$$

If ν on X^+ were also σ-invariant, then $\phi\nu = \lambda$ defined on G by $\lambda(B) = \nu(\phi^{-1}B)$ is T-invariant, hence $= m$ (because of unique ergodicity of (G,T)). In particular, $\nu(\psi G_0) = 0$. So $\mu = \nu$.

Remark 12.1. It's not hard to see directly, using our criteria for subshifts, that $(\overline{\mathcal{O}^+(\omega)},\sigma)$ is minimal and uniquely ergodic.

Every block B that appears in ω appears not only with bounded gap, but along (at least) an arithmetic sequence.

$$\omega = 01\underbrace{0001}_{B}0\_01\underbrace{0001}_{B}\dots$$

Since ω is the limit of periodic sequences on $\{0,1,\_\}$, ω is *regularly almost periodic* in this sense.

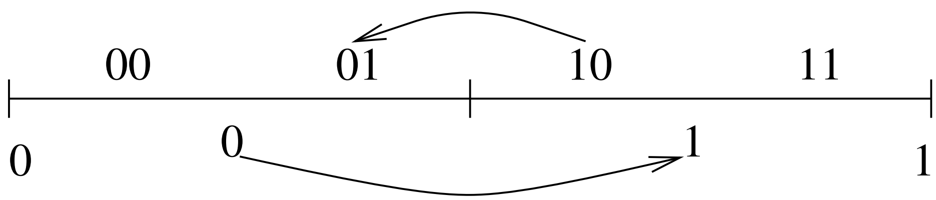

FIGURE 23. Cutting and Stacking the Unit Interval

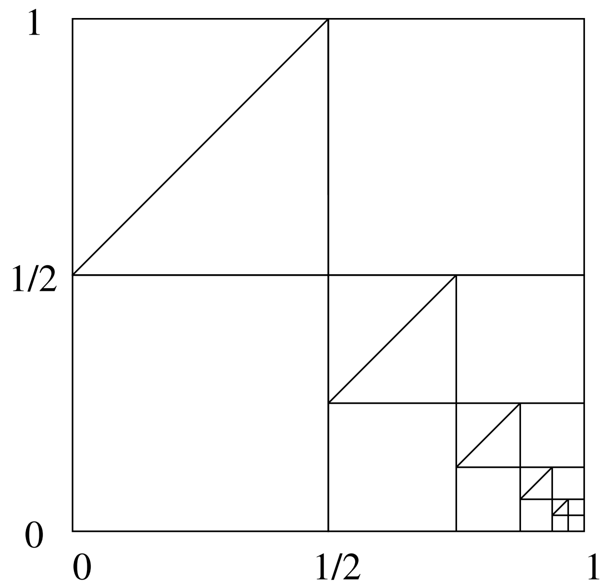

FIGURE 24. Graph of the Cutting and Stacking Function

12.2. **The Odometer as a Map on the Interval–Cutting and Stacking.** Let $x \in [0, 1)$ with dyadic expansion $x_1 x_2 \dots$. Then

$$Tx = 000 \dots 01 x_{k+1} x_{k+2} \dots,$$

where $k = \min\{n : x_n = 1\}$, i.e. k is the first place a 1 appears.

Figure 24 shows the graph of the cutting and stacking function.

Figure 25 illustrates the first three stages of cutting and stacking.

Lebesgue measure is preserved. Lebesgue measure on $[0, 1)$ corresponds to Bernoulli measure $\mathcal{B}(\frac{1}{2}, \frac{1}{2})$ on Σ_2^+, the set of dyadic expansions.

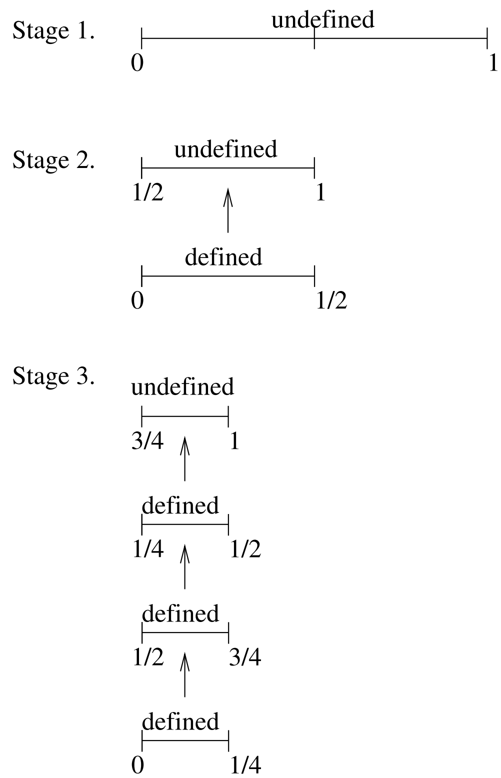

Figure 25. Stages of Cutting and Stacking

13. February 19 *(Notes by RP)*

Recall the odometer (G, T) and the construction and properties (2:1, 1:1, continuity, etc.) of ψ and ϕ. From now on, denote the Toeplitz sequence by s. Then $\phi(000\dots) = s$.

13.1. **Action-Commuting Property.** We show that if $u \in \psi x$, then $\sigma u \in \psi T x$, where $x \in G$. (If $x = 111\dots$, then $\psi x = zs$, where $z \in \{0,1\}$, so $\sigma u = s = \psi(00\dots) = \psi(Tx)$.) Say

$$x = \underbrace{11\dots1}_{k-1}0x_{k+1}x_{k+2}\dots,$$

so that

$$Tx = \underbrace{00\dots0}_{k-1}1x_{k+1}x_{k+2}\dots.$$

Now x prescribes that the first place in ψx is a blank for the first $k-1$ iterations and then gets filled in with 0 or 1, depending on the parity of k. The next blank place in ψx remains blank at this stage. Now consider $\psi T x$. If after the first k steps we insert either a 0 or 1 in front of the sequence (whichever we used in step k), it would be as if for the first $k-1$ steps we skipped it and started with it on step k, i.e., as if we had taken ψx. Hence at this stage $\sigma^{-1}\psi T x = \psi x$, i.e., $\psi T x = \sigma \psi x$. Then afterward Tx and x coincide, hence $\psi T x = \sigma\psi x$, and therefore also $\phi\sigma u = Tx = T\phi u$ (since $\phi\psi =$id and $\phi u = x$). We have shown that $\boxed{\phi\sigma = T\phi}$.

We claim that $\psi G = \overline{\mathcal{O}^+(s)}$. $s \in \psi G$, so from above $\sigma^k s \in \psi G$ for all $k \geq 0$, i.e. $\psi G \supset \mathcal{O}^+(s)$. Since ψG is closed, $\psi G \supset \overline{\mathcal{O}^+(s)}$. We want to show that $\mathcal{O}^+(s)$ is *dense* in ψG. Let $u \in \psi G$, $u = u_1 u_2 \dots$. Find n with $\psi(n\theta)$ agreeing with u on an arbitrarily long initial block. Recall that $\theta = 100\dots \in G$ is a generator, and

$$n\theta = \epsilon_0\epsilon_1\epsilon_2\dots\epsilon_r 00\dots,$$

where

$$n = \epsilon_0 + \epsilon_1 \cdot 2 + \epsilon_2 \cdot 2^2 + \dots + \epsilon_r \cdot 2^r.$$

Thus *every* block is an initial block of $n\theta$ for some n. Say in building up u as ψ(some x), we have $u = u_1 \dots \_ \dots u_m$, a block on $\{0, 1, \_\}$. We can get $n\theta$ to agree with a long initial block of x. Then $\psi(n\theta) = u_1 \dots \_ \dots u_m$ (the same block on $\{0,1,\_\}$). Now $u = u_1 \dots \lambda \dots u_m \dots$, where $\lambda \in \{0,1\}$. Say $n\theta = \epsilon_0\epsilon_1\epsilon_2 \dots \epsilon_r 00 \dots (= x_1 \dots x_{r+1} \dots)$. Append either 0 or 10 to get λ filled in properly. This completes the proof of the following theorem.

Theorem 13.1. *$\phi : (\overline{\mathcal{O}^+(s)}, \sigma) = (X^+, \sigma) \to (G, T)$ is 1:1 except on a countable set, where it is 2:1. It is continuous and satisfies $\phi\sigma = \sigma T$.*

13.2. **Odometer as Cutting and Stacking.** Recall the cutting and stacking approach discussed last time. For $x \in G = \Sigma_2^+$, put $\rho(x) = \min\{n : x_n = 0\} + 1 \bmod 2$, i.e. $\rho(x) = 1-$(parity of the first place in x where you see a zero), and $\rho(111\dots) = 1$. Then $(\rho(T^n x)) = \psi(x)$ for $x \notin G_0$. So as we follow the orbit of $000\dots$ under T and write down the values of ρ, we build the Toeplitz sequence s. We are coding a transformation by a partition.

Definition 13.1 (The *two-sided Toeplitz system*)**.** In the preceding construction, we could just as well have defined $\psi : G \to \Sigma_2$ instead of just $\psi : G \to \Sigma_2^+$:

$$\psi x = \begin{cases} \dots 0a_{-2}0a_{-1}\dot{0}a_0 0a_1\dots & \text{if } x_1 = 0 \\ \dots a_{-2}0a_{-1}0\dot{a_0}0a_1 0\dots & \text{if } x_1 = 1 \end{cases},$$

etc. (Any "unfilled blank" is to the right (left) of the central position if x has finitely many 0's (1's).)

The properties established above still hold, including density of the orbit of each of the two points s_1 and s_2 in $\psi(000\dots)$. In fact, already $\sigma : X^+ \to X^+$ is essentially invertible: one can check that $\sigma^{-1}\{u\}$ is a single point except when $u = s$, in which case $\sigma^{-1}\{s\} = \{0s, 1s\}$ (because $T^{-1}(000\dots) = -\theta = 1111\dots$). Then $X = \psi G = \overline{\mathcal{O}(s_1)} = \overline{\mathcal{O}(s_2)}$ is again a closed subshift which is minimal and uniquely ergodic.

Theorem 13.2 (replaces Theorem 10.2)**.** *Let (M, σ) be the two-sided Morse minimal system $\subset \{-1,1\}^{\mathbb{Z}}$, and (X, σ) the two-sided (minimal, uniquely ergodic) Toeplitz system. Define $f : X \to \{-1,1\}$ by $f(u) = (-1)^{u_0}$ for all $u \in X$. Define $\hat{T} : X \times \{-1,1\} \to X \times \{-1,1\}$ by $\hat{T}(u,\xi) = (\sigma u, f(u)\xi)$, a skew product. Then (M, σ) and $(X \times \{-1,1\}, \hat{T})$ are topologically conjugate.*

Remark 13.1. Note that if $\tau : G \to \mathbb{N}$ is defined by

$$\tau(x_1 x_2 \dots) = \min\{n : x_n = 0\},$$

i.e.

$$\tau(x_1 x_2 \dots) = (\text{first index } k \text{ where } x_k = 0),$$

and

$$\tau(111\dots) = 1,$$

then $(-1)^{\tau(x)+1} = f(\psi x)$ for $x \notin G_0$.

Remark 13.2. Note also that, if we define the *cocycle* $f(u,n)$ by

$$f(u,n) = \begin{cases} 1 & \text{for } n = 0 \\ f(u)f(\sigma u)\cdots f(\sigma^{n-1}u) & \text{for } n > 0 \\ f(\sigma^{-1}u)\cdots f(\sigma^{n}u) & \text{for } n < 0 \end{cases},$$

then $\hat{T}^n(u,\xi) = (\sigma^n u, f(u,n)\xi)$.

Remark 13.3. Consider the Morse sequence

$$\omega = \dots 011010011001011010010110011010 01\dots.$$

Look at the cellular automaton $\eta(x)_n = x_n + x_{n+1} + 1$. Then

$$\eta(\omega) = \dots 010001010100010\_0 \dots = s,$$

the Toeplitz sequence. This is easily seen by grouping ω first into a sequence on the 2-blocks 01 and 10, each of which is sent to 0 by η, then into the 4-blocks 0110 and 1001, etc.

Proof of Theorem 13.2. Let $\hat{X} = X \times \{-1,1\}$. Define $\pi : \hat{X} \to \{-1,1\}^{\mathbb{Z}}$ by $\pi(u,\xi)_n = f(u,n)\xi$ for all $n \in \mathbb{Z}$. (We code the orbit of a point in $\hat{X}$ by the sequence of ± 1's in its second coordinate.) It will turn out that $\pi(s_1, 1)_n = (-1)^{\omega_n}$ for $n \geq 0$ (and of course similarly for s_2–recall that both s_1 and s_2 have s as their right half). ∎

14. February 24 *(Notes by JF)*

We recall that the Morse minimal set M is defined by $M = \overline{\mathcal{O}(\omega)} = \overline{\{\sigma^n\omega : n \in \mathbb{Z}\}}$, and we continue to let X be the orbit closure of the two-sided Toeplitz sequence. Then we have the following theorem, which was stated last time.

Theorem 14.1. *(M, σ) is topologically conjugate to $(\hat{X}, \hat{T})$, where $\hat{X} = X \times \{1, -1\}$, $\hat{T}(u, \xi) = (\sigma u, f(u)\xi)$, and $f(u) = (-1)^{u_0}$.*

Before proving the theorem, we need to make a few comments and introduce additional notation. For $x \in G = {\Sigma_2}^+$ with $x \neq 111\ldots$, we define $\tau(x)$ to be the first positive n such that $x_n = 0$, where $x = x_1x_2x_3\ldots$. When $x = 111\ldots$, we set $\tau(x) = 1$. We claim that τ defined in this way satisfies the equation

$$(-1)^{\tau(x)+1} = f(\psi(x)) \text{ if } x \neq 111\ldots.$$

In fact, as long as there are 1's in x, the right half of $\psi(x)$ is

$$\begin{aligned}
&_0_0_0_0_0_0\ldots \text{(first step)}\\
&_010_010_010\ldots \text{(second step)}\\
&_0100010_010\ldots \text{(third step)}\\
&\vdots
\end{aligned}$$

and the initial blank will be filled in with $1 + n \pmod 2$, where n is the smallest natural number such that $x_n = 0$. Hence the claim follows.

We recall also that if n has a diadic expansion given by

$$n = \varepsilon_0 + \varepsilon_1 2^1 + \varepsilon_2 2^2 + \ldots,$$

then $n\theta = \varepsilon_0\varepsilon_1\varepsilon_2\ldots$. We can now proceed with the proof of the theorem.

Proof. Define a map $\pi : \hat{X} \to \{-1, 1\}^{\mathbb{Z}}$ by $\pi(u, \xi)_n = f(u, n)\xi$, where $f(u, 0) = 1$ and $f(u, n)$ is defined for $n \neq 0$ by

$$f(u,n) = \begin{cases} f(u)f(\sigma u)f(\sigma^2 u)\ldots f(\sigma^{n-1}u) \text{ for } n \geq 1\\ f(\sigma^{-1}u)f(\sigma^{-2}u)\ldots f(\sigma^n u) \text{ for } n \leq -1\end{cases}$$

This $f(u, n)$ is called a *multiplicative cocycle.* Then

$$\pi(u, \xi) = \ldots f(\sigma^{-1}u)\xi, \underbrace{\xi}_{\text{0'th place}}, f(u)f(\sigma u)\xi, \ldots$$

[There was a question about why the term "cocycle" is used here. By way of explanation, consider the following situation: Let $T : X \to X$ and $f : X \to \mathbb{C}$. Then write

$$S_n f(x) = \begin{cases} \sum_{k=0}^{n-1} f(T^k x) \text{ if } n \geq 1\\ 0 \text{ if } n = 0\\ \sum_{k=n}^{-1} f(T^k x) \text{ if } n \leq -1\end{cases}$$

It is then easy to see that $S_{n+m}f(x) = S_nf(x) + S_mf(T^nx)$, and if we let $F(x,n) = S_nf(x)$, we have that $F(x,n+m) = F(x,n)+F(T^nx,m)$, an *additive cocycle* equation. The terminology comes from the cohomology theory of groups, where one finds analogous equations. A skew product such as $\hat{T}$ will in general have a cocycle such as $f(u,n)$ appearing in the second coordinate when it is raised to higher powers.] We can now see that the entries in $\pi(u,\xi)$ are just the *second coordinates* of the expression for $\hat{T}^n(u,\xi)$.

We now show that π is continuous. If (u,ξ) and (v,ζ) are close in $\hat{X}$, then $\xi=\zeta$ and u and v are close in the two-sided shift. Hence u and v agree on a long central block. Specifically, there exists $J\in\mathbb{N}$ such that $(-1)^{u_j} = (-1)^{v_j}$ for $|j|\le J$. But then $\pi(u,\xi)_n = \pi(v,\zeta)_n$ for $|n|\le J$, which means that $\pi(u,\xi)$ and $\pi(v,\zeta)$ are close. This shows that π is continuous.

Notice also that π is one-to-one. If we know $\pi(u,\xi)$, then we know ξ, since $\pi(u,\xi)_0=\xi$. But then once we know ξ, we can determine the value of $f(u)$ and hence of u_0. Continuing in this manner we can determine u_n for all n.

We now make the claim that if s is the one-sided Toeplitz sequence and $n\ge 0$, then $\pi(\ldots .s,1)_n = (-1)^{\omega_n}$, where ω_n is the n'th entry in the Morse sequence. Recall that $\omega_n = \sum_{i=0}^{\infty}\varepsilon_i(n) \pmod 2$, where n has a diadic expansion given by $n=\sum_{i=0}^{r}\varepsilon_i(n)2^i$ and each $\varepsilon_i(n)$ is either 0 or 1. Suppose that n has a diadic expansion such that $\varepsilon_i(n)=1$ for all i such that $|i|<p$ and $\varepsilon_p(n)=0$. Then $n\theta = 111\ldots1 \underbrace{0}_{\text{p'th place}} \ldots$ and $\tau(n\theta)=p$. Hence $f(\sigma^n(\ldots s)) = (-1)^{p+1}$.

Note now that if n has the diadic expansion we described earlier, then

$$n+1 = 0+0\cdot 2+0\cdot 2^2+\cdots+0\cdot 2^{p-1}+1\cdot 2^p\ldots,$$

where the tail of the expansion is unchanged. Therefore

$$\underbrace{1+\sum_{i=0}^{\infty}\varepsilon_i(n)-\sum_{i=0}^{\infty}\varepsilon_i(n+1)}_{\text{both finite sums}} = 1+[p+\varepsilon_{p+1}+\ldots]-[1+\varepsilon_{p+1}+\ldots] = p = \tau(n\theta).$$

Hence

$$\left[\sum_{k=0}^{n-1}\tau(k\theta)\right]+n = \sum_{k=0}^{n-1}\left[1+\sum_{i=0}^{\infty}\varepsilon_i(k)-\sum_{i=0}^{\infty}\varepsilon_i(k+1)\right]+n = 2n+\sum_{i=0}^{\infty}\left[\sum_{k=0}^{n-1}\varepsilon_i(k)-\sum_{k=0}^{n-1}\varepsilon_i(k+1)\right]$$

$$= 2n+\sum_{i=0}^{\infty}(\underbrace{\varepsilon_i(0)}_{0}-\varepsilon_i(n)) = 2n-\sum_{i=0}^{\infty}\varepsilon_i(n) \equiv 0+\sum_{i=0}^{\infty}\varepsilon_i(n) \pmod 2 = \omega_n.$$

We have shown that

$$f(\ldots s,n) = f(s)f(\sigma s)\cdots f(\sigma^{n-1}s) = (-1)^{s_0}(-1)^{s_1}\cdots(-1)^{s_{n-1}} = (-1)^{\tau(0\theta)+1}(-1)^{\tau(1\theta)+1}\cdots(-1)^{\tau((n-1)\theta)+1}.$$

But that demonstrates our claim, for we have shown that $\pi(\ldots s) = \ldots\omega^+$, where ω^+ is the right half of the Morse sequence on $\{-1,1\}$.

Further, it is clear that we have the following relationship:

$$\begin{array}{ccc}(u,\xi) & \xrightarrow{\hat{T}} & (\sigma u, f(u)\xi)\\ \pi\downarrow & & \downarrow\pi\\ (\ldots,\xi,f(u)\xi,\ldots) & \xrightarrow{\sigma} & (\ldots,f(u)\xi,f(\sigma u)f(u)\xi\ldots),\end{array}$$

i.e. $\sigma\pi = \pi\hat{T}$.

The image $\pi(\hat{X}) \subset \{-1,1\}^{\mathbb{Z}}$ is closed and σ-invariant. Since it contains all shifts of a sequence which agrees with λ for $n \geq 0$, where $\lambda_n = (-1)^{\omega_n}$, $\pi(\hat{X})$ contains the orbit closure of λ.

Let N be the orbit closure of λ. Then (N, σ) is topologically conjugate to (M, σ). To determine whether $\pi(\hat{X})$ is contained in N, we need to figure out which words can appear in a sequence in $\pi(\hat{X})$. We claim that the only such words are those which appear in the right half of λ.

If $(u, \xi) \in \hat{X}$, then the image under π of (u, ξ) is just $(\ldots, \xi, f(u)\xi, \ldots, f(u,n)\xi, \ldots)$. Consider the word of length r beginning in the m'th position, namely $(f(u,m)\xi, f(u,m{+}1)\xi, \ldots, f(u,m{+}r{-}1)\xi)$. Since the Morse system is self-dual, we can assume that $\xi = 1$. Then we can find $M \geq 0$ such that $\sigma^M(\ldots s) \approx u$. Take M so large that $f(\sigma^{j+M}(\ldots s)) = f(\sigma^j(u))$ for all $j = 0, \ldots, (m+r-1)$. This completes the proof of the theorem. □

We will use the topological conjugacy we have demonstrated to derive information about the structure of the Morse minimal set and, more importantly, the variants of it.

15. February 26 *(Notes by JF)*

The Morse system (M, σ) is topologically conjugate to $(\hat{X}, \hat{T})$, where $\hat{X} = X \times \{-1, 1\}$, X is the orbit closure of the two-sided Toeplitz sequence, and $\hat{T}(u, \xi) = (\sigma u, f(u)\xi)$, where $f(u) = (-1)^{u_0}$. The isomorphism is given by π, where $\pi(u, \xi)_n = f(u, n)\xi$. Hence $\pi(u, \xi)$ is the sequence of second coordinate values in the orbit of (u, ξ).

15.1. **History:**

(1) Otto Toeplitz in 1928 constructed the Toeplitz sequence as an example of an almost-periodic function on the integers.
(2) Kakutani (1967) discussed the structure and spectrum of the Morse and related sequences. His work was the first place where the skew product came into the discussion.
(3) Veech (1969) dealt with even more generalizations of the structure.
(4) Jacobs and Keane (1969) discussed other Toeplitz systems, and Keane (1969) dealt with generalized Morse systems.
(5) Susan Williams in her 1981 thesis showed how to get non-uniquely ergodic Toeplitz systems, generalizing an example of Oxtoby.

15.2. **Another Theorem.** We deal in this section with a theorem of Furstenberg (1961) which was used by Veech (1969). The theorem, which we will state and prove, relates to the ergodicity of skew product transformations, a topic which was dealt with even earlier by Anzai.

Theorem 15.1. *Let (X, T) be a minimal topological dynamical system. Let $f : X \to \{-1, 1\}$ be a continuous function, and define $\hat{T}$ on $\hat{X} = X \times \{-1, 1\}$ by $\hat{T}(x, \xi) = (Tx, f(x)\xi)$. Then $(\hat{X}, \hat{T})$ is not minimal if and only if there is a nontrivial continuous solution g of the cocycle-coboundary equation $g(Tx) = f(x)g(x)$. Suppose moreover that (X, T) is minimal and uniquely ergodic. Then $(\hat{X}, \hat{T})$ is not uniquely ergodic if and only if there is a nontrivial measurable solution g to the above cocycle-coboundary equation.*

Proof. Suppose that g is continuous on X and satisfies $g(Tx) = f(x)g(x)$ for all $x \in X$. Then $|g|$ is continuous and T-invariant, hence constant. (To see that $|g|$ is constant, look at $\{x : |g(x)| \leq \alpha\}$ for $\alpha \in \mathbb{R}$. This set is closed and T-invariant for every α, hence either empty or X for all α. Alternatively, fix x and note that $|g|$ is constant on $\mathcal{O}(x)$, a dense set.) Then

$$f(x) = \frac{g(Tx)}{g(x)} \frac{g(T^2x)}{g(Tx)} \cdots \frac{g(T^nx)}{g(T^{n-1}x)} = \frac{g(T^nx)}{g(x)}.$$

Hence $\hat{T}^n(x, \xi) = (T^nx, f(x, n)\xi) = (T^nx, \frac{g(T^nx)}{g(x)}\xi)$. If $\{T^{n_k}x\}$ converges to some y, then $\hat{T}^{n_k}(x, \xi) \longrightarrow (y, \frac{g(y)}{g(x)}\xi)$. Thus we cannot have $\hat{T}^{n_k}(x, 1) \longrightarrow (x, -1)$, since $\hat{T}^{n_k}(x, 1) \longrightarrow (x, 1)$. Hence $\mathcal{O}^+(x, 1)$ is not dense in $\hat{X}$ for any x. The result follows similarly for negative powers.

Conversely, suppose that $(\hat{X}, \hat{T})$ is not minimal. Take (x_0, ξ_0) such that $\overline{\mathcal{O}(x_0, \xi_0)}$ is strictly contained in $\hat{X}$.

Lemma. *For every x, there exists a unique $g(x) \in \{-1, 1\}$ such that $(x, g(x)) \in \overline{\mathcal{O}(x_0, \xi_0)} = X_0$.*

Proof. Suppose that for some x both $(x, 1)$ and $(x, -1)$ are contained in X_0. Given any $(y, \xi) \in \hat{X}$, we can choose a sequence $\{n_k\}$ such that $T^{n_k}x \longrightarrow y$ since (X, T) is minimal. By taking a subsequence, we may assume that $\hat{T}^{n_k}(x, 1) \longrightarrow (y, \zeta)$, where $\zeta \in \{-1, 1\}$. But then $\hat{T}^{n_k}(x, -1) \longrightarrow$

$(y,-\zeta)$. Hence $\{(y,1),(y,-1)\} \subset \overline{\mathcal{O}(x,1)} \cup \overline{\mathcal{O}(x,-1)} \subset \overline{\mathcal{O}(x_0,\xi_0)}$, where here we use the fact that if in any system we have $u \in \overline{\mathcal{O}(v)}$, then $\overline{\mathcal{O}(u)} \subset \overline{\mathcal{O}(v)}$. It then follows since (y,ζ) was arbitrary that $\overline{\mathcal{O}(x_0,\xi_0)} = \hat{X}$, a contradiction. This proves the lemma. □

We now demonstrate that the function g given by the lemma is continuous. Suppose that a sequence of points $\{x_n\}$ converges to x. Take a subsequence so that $\underbrace{(x_n, g(x_n))}_{\in \overline{\mathcal{O}(x_0,\xi_0)}} \longrightarrow (x,\xi)$. Then $(x,\xi) \in \overline{\mathcal{O}(x_0,\xi_0)}$, which means that $\xi = g(x)$. Thus $x_n \longrightarrow x$ implies that $g(x_n) \longrightarrow g(x)$. Finally, $\hat{T}(x,g(x)) = (Tx, f(x)g(x))$ implies that $g(Tx) = f(x)g(x)$. This completes the proof of the portion of the theorem which deals with minimality.

Now suppose that μ_X is the unique T-invariant Borel probability measure on X. Define μ on $\hat{X}$ by $\mu = \mu_X \times \frac{\text{counting measure}}{2}$. That is, each of 1 and -1 has measure $\frac{1}{2}$ in the set $\{-1,1\}$. For a set $A \subset X$, $\mu(A \times \{1\}) = \mu(A \times \{-1\}) = \frac{1}{2}\mu_X(A)$.

Lemma. *The system $(\hat{X},\hat{T})$ is uniquely ergodic if and only if $(\hat{X},\hat{T},\mu)$ is ergodic.*

We recall that the system $(\hat{X},\hat{T},\mu)$ is ergodic if every $\hat{T}$-invariant measurable function is constant almost everywhere with respect to μ.

Proof. Suppose μ is not ergodic for $(\hat{X},\hat{T})$. Then $(\hat{X},\hat{T})$ is not uniquely ergodic, since the existence of a nonergodic measure implies the existence of at least two ergodic measures.

Conversely, suppose that μ is ergodic, and let ν be a $\hat{T}$-invariant measure on $\hat{X}$. Then look at the flip transformation $F : \hat{X} \to \hat{X}$ defined by $F(x,\xi) = (x,-\xi)$. It is clear that $F\hat{T} = \hat{T}F$, and the fact that F and $\hat{T}$ commute allows us to show that $F\nu$ is also a $\hat{T}$-invariant measure. Let $\hat{A} \subset \hat{X}$. Then

$$F\nu(\hat{T}^{-1}\hat{A}) = \nu(F^{-1}\hat{T}^{-1}\hat{A}) = \nu(\hat{T}^{-1}F^{-1}\hat{A}) = \hat{T}\nu(F^{-1}\hat{A}) = \nu(F^{-1}\hat{A}) = F\nu(\hat{A}).$$

Having shown that $F\nu$ is $\hat{T}$-invariant, we now claim that $\mu = \frac{1}{2}(\nu + F\nu)$. Both μ and $\frac{1}{2}(\nu + F\nu)$ project to μ_X in the first coordinate. In fact, if π is the projection of $\hat{X}$ onto X, the fact that ν is $\hat{T}$-invariant implies that $\pi\nu$ is T-invariant. Let $A \subset X$. Then

$$\begin{aligned}\frac{1}{2}(\nu + F\nu)(A \times \{\xi\}) &= \frac{1}{2}[\nu(A \times \{\xi\}) + \nu(A \times \{-\xi\})] = \frac{1}{2}[\nu(A \times \{-\xi,\xi\})] \\ &= \frac{1}{2}(\pi\nu)(A) = \frac{1}{2}\mu_X(A) = \mu(A \times \{\xi\}).\end{aligned}$$

Hence $\mu = \frac{1}{2}(\nu + F\nu)$, which contradicts the ergodicity of μ unless $\nu = F\nu$. But then $\mu = \nu$, which proves the lemma. □

Suppose that g is a nontrivial measurable solution to $g(Tx) = f(x)g(x)$. Our method of proof here is to show that if such a g exists, then $\mu = \mu_X \times \frac{\text{counting measure}}{2}$ is not ergodic on $\hat{X}$. Note that $h(x,\xi) = g(x)\xi$ is not constant almost everywhere with respect to μ, and

$$h(\hat{T}(x,\xi)) = h(Tx, f(x)\xi) = g(Tx)\underbrace{f(x)}_{\pm 1\text{-valued}}\xi = g(x)\xi = h(x,\xi).$$

Hence h is nonconstant, measurable, and $\hat{T}$-invariant, which means that μ is not ergodic. Thus the system $(\hat{X},\hat{T})$ is not uniquely ergodic.

We will give the remainder of the proof next time. □

16. March 3 *(Notes by XM)*

To finish the proof of the previous theorem it remains to show that if the skew product is not uniquely ergodic, then the cocycle-coboundary equation has a measurable solution.

Suppose that $(\hat{X}, \hat{T})$ is not uniquely ergodic; then, by the lemma, μ is not ergodic. Hence there exists a $\hat{T}$-invariant function $h : \hat{X} \to \mathbb{C}$ which is not constant a.e. with respect to μ.
Let's define

$$g(x) = \frac{h(x,1) - h(x,-1)}{2}.$$

It is clearly measurable with respect to μ_X, and we will show it satisfies our equation.

Since h is $\hat{T}$-invariant, we have the following :

$$\begin{aligned} h(x,\xi) &= \underbrace{\frac{h(x,1)+h(x,-1)}{2}}_{:=v(x)} + \xi \underbrace{\frac{h(x,1)-h(x,-1)}{2}}_{g(x)} \qquad \text{(simply note that } \xi \in \{-1,1\}) \\ &= h \circ \hat{T}(x,\xi) \\ &= h(Tx, f(x)\xi) \\ &= \underbrace{\frac{h(Tx,1)+h(Tx,-1)}{2}}_{v(Tx)} + f(x)\xi \underbrace{\frac{h(Tx,1)-h(Tx,-1)}{2}}_{g(Tx)} \qquad \mu\text{-a.e.,} \end{aligned}$$

so that

$$v(x) + \xi g(x) = v(Tx) + f(x)\xi g(Tx).$$

For $\xi = 1$ we have

$$v(x) + g(x) = v(Tx) + f(x)g(Tx),$$

and for $\xi = -1$ we have

$$v(x) - g(x) = v(Tx) - f(x)g(Tx).$$

The difference of the two equations gives us $g(x) = f(x)g(Tx)$, and multiplying by$f(x)$ on both sides we see g is a solution of our equation, i.e., $g \circ T = fg$.

It remains to show that g is not identically zero. Suppose it was; then $h(x,1) = h(x,-1) = h_0(x)$ a.e., and since h is $\hat{T}$-invariant, h_0 is T-invariant. By ergodicity of μ_X, h_0 must be constant a.e., which in turns implies h must be constant, so we have a contradiction. □

16.1. **Application to the Morse System.** When the minimal dynamical system (X, σ) is the Toeplitz system (and $f(u) = (-1)^{u_0}$, for $u \in X$), the skew-product defined in the previous theorem is topologically conjugate to the Morse system (*cf.* Theorem 10.2). In order to prove that the Morse system is *minimal* and *uniquely ergodic* we need to show that the cocycle-coboundary equation has no nontrivial measurable solution. The next result gives us a necessary and sufficient condition for this to happen.

First let us recall that we have a continuous map ψ, which is a.e. a homeomorphism, from the odometer (G, R_θ) into the 2-sided Toeplitz system (X, σ), where $G = \Sigma_2^+$, $\theta = (111\ldots)$ and $R_\theta(x) = x + \theta$.

Theorem 16.1. *Let (X,σ) be the 2-sided Toeplitz system and μ_X the unique ergodic measure on X. If there exists a non trivial $g : X \to \mathbb{C}$ satisfying $g(\sigma x) = f(x)g(x)$ a.e. with respect to μ_X, then*

$$\lim_{k\to\infty} \int_X f(u,n_k)d\mu_X(u) = 1 \tag{7}$$

for every sequence $\{n_k\}$ of positive integers such that $n_k\theta \to 0$. (Recall that

$$f(u,n) = f(u)f(\sigma u)f(\sigma^2 u)\dots f(\sigma^{n-1}u)$$

for $n > 0$.) The converse also holds.

Proof. Let g be a nontrivial solution of $g(\sigma x) = f(x)g(x)$. $\Re(g)$ and $\Im(g)$ are also solutions and one of them at least must be nontrivial; therefore we can assume that g is real-valued. Note that $|g\circ\sigma| = |f|\,|g| = |g|$, so $|g|$ is σ-invariant, and hence must be constant almost everywhere with respect to μ_X. Also $|g| \neq 0$ since g is not identically zero. We may then assume g takes values in $\{-1,1\}$ since it is the case for $g/|g|$.

Now for any $u \in X$ and $n > 0$ we have

$$f(u,n) = f(u)f(\sigma u)f(\sigma^2 u)\dots f(\sigma^n u) = \frac{g(\sigma u)}{g(u)}\frac{g(\sigma^2 u)}{g(\sigma u)}\cdots\frac{g(\sigma^n u)}{g(\sigma^{n-1}u)} = \frac{g(\sigma^n u)}{g(u)}.$$

Or, if we look at it on the odometer side,

$$f(\psi t,n) = \frac{g(\sigma^n\psi t)}{g(\psi t)}.$$

Recall that $\sigma\psi = \psi R_\theta$, so that $\sigma^n\psi = \psi R_\theta^n$. Then

$$f(\psi t,n) = \frac{g(\psi R_\theta^n t)}{g(\psi t)} = \frac{g(\psi(t+n\theta))}{g(\psi t)}.$$

Now we calculate that

$$\begin{aligned}\left|\int_X f(u,n_k)d\mu_X(u) - 1\right| &= \left|\int_G f(\psi t,n_k)dm(t) - 1\right| = \left|\int_G \{f(\psi t,n_k)-1\}\,dm(t)\right| \\ &= \left|\int_G \left[\frac{g\psi R_\theta^{n_k}(t)}{g(\psi t)} - 1\right]dm(t)\right| \\ &\le \int_G \left|\frac{g\psi R_\theta^{n_k}(t)}{g(\psi t)} - 1\right|dm(t) = \int_G \left|\frac{g\psi R_\theta^{n_k}(t) - g(\psi t)}{g(\psi t)}\right|dm(t) \\ &\le \|g\circ\psi\circ R_\theta^{n_k} - g\circ\psi\|_{L^1(G,m)} \text{ (since } |g|=1),\end{aligned}$$

and this is small for large k because translation is continuous in $L^1(G,m)$. [Given $\varepsilon > 0$ there exists h continuous over G such that $\|g\circ\psi - h\|_{L^1} < \varepsilon/3$ (because continuous functions are dense in L^1). Then

$$\begin{aligned}\|g\circ\psi\circ R_\theta^{n_k} - g\circ\psi\|_{L^1} &\le \|g\circ\psi\circ R_\theta^{n_k} - h\circ R_\theta^{n_k}\|_{L^1} + \|h\circ R_\theta^{n_k} - h\|_{L^1} + \|h - g\circ\psi\|_{L^1} \\ &\le 2\|g\circ\psi - h\|_{L^1} + \|h\circ R_\theta^{n_k} - h\|_{L^1}.\end{aligned}$$

Since h is uniformly continuous on G we can choose k large enough so that,

$$\sup_{t \in G} |h(t + n_k\theta) - h(t)| < \varepsilon/3,$$

and thus $\|h \circ R_\theta^{n_k} - h\|_{L^1} < \varepsilon/3$. Finally, $\|g \circ \psi \circ R_\theta^{n_k} - g\psi\|_{L^1} < \varepsilon$ for k large enough, which proves 7.]

We don't actually need the converse of this statement so we won't prove it but leave it as an exercise. □

17. March 5 *(Notes by XM)*

Theorem 17.1. *Condition (7) of Theorem 16.1 is not satisfied, hence the Morse system is minimal and uniquely ergodic.*

Remark 17.1. When a system is both minimal and uniquely ergodic we sometimes say it is *strictly ergodic.*

Recall some notation :

$$\tau(x) = \begin{cases} \inf\{n : x_n = 0\} & \text{if } x \neq 11\dots1\dots \\ 1 & \text{if } x = 11\dots1\dots \end{cases}$$

Define $\gamma(x) = f(\psi x) = (-1)^{\tau(x)+1}$ and

$$\gamma(x,n) = \begin{cases} \gamma(x)\gamma(R_\theta x)\dots\gamma(R_\theta^{n-1}x) & \text{if } n > 0 \\ 1 & \text{if } n = 0 \\ \gamma(R_\theta^{-1}x)\gamma(R_\theta^{-2}x)\dots\gamma(R_\theta^{n}x) & \text{if } n < 0. \end{cases}$$

We will show that Condition (7) is not true, i.e. that $\lim_{r\to\infty}\int_G \gamma(x,n_r)dm(x) \neq 1$ for a certain sequence n_r (γ is just the version of f from the odometer side). In order to do that we need a stronger version of the pointwise ergodic theorem, namely :

Theorem 17.2. *Let (X,T) be a uniquely ergodic topological dynamical system and μ a non-atomic T-invariant measure on X. If $h : X \to \mathbb{C}$ is a bounded measurable function with at most finitely many discontinuities, then*

$$A_n h(x) = \frac{1}{n}\sum_{k=0}^{n-1} h\circ T^k x \longrightarrow \int_X h d\mu \tag{8}$$

for all $x \in X$. (What's stronger in this new statement is that the convergence occurs not just almost everywhere).

Proof. If h is continuous we already know (*cf.* Theorem 9.1) that $A_n h$ converges uniformly on X to $\int_X h d\mu$. If we suppose that h is discontinuous only at a finite number of points, say $x_1, x_2, \dots, x_m$, the idea is to show we can find continuous functions u_ε and v_ε such that $u_\varepsilon \le h \le v_\varepsilon$ and which have integrals differing by at most ε. Using standard arguments of measure theory, we may assume h is positive. (To reduce to the real case, write $h = \Re h + i\Im h$; to reduce to the positive case, write $h = h^+ - h^-$, where $h^+ = \max\{h,0\}$ and $h^- = \max\{-h,0\}$).

Let $\varepsilon > 0$ be arbitrary. μ is non-atomic, therefore $\mu(x_i) = 0$ for all i's, and since it is also outer regular (being a Borel measure), we can find disjoint open neighborhoods V_i of the x_i each of measure less than ε/m, so that the total measure of their union is less than ε. Now let U_i be smaller neighborhoods of the x_i with $\overline{U_i} \subset V_i$. Using Urysohn's Lemma (for example), we can build a continuous map χ_ε from X onto $[0,1]$ such that $\chi_\varepsilon = 1$ on U_i and $\chi_\varepsilon = 0$ outside V_i, for all i.

Let's define

$$u_\varepsilon = (1-\chi_\varepsilon)h$$

and

$$v_\varepsilon = (1-\chi_\varepsilon)h + \|h\|_\infty \chi_\varepsilon.$$

The function u_ε is continuous by construction, and v_ε is just the sum of u_ε and of another continuous function, so it is also continuous. Moreover, $u_\varepsilon \leq h$ and $v_\varepsilon = h + (\|h\|_\infty - h).\chi_\varepsilon \geq h$, so that $u_\varepsilon \leq h \leq v_\varepsilon$. Both u_ε and v_ε coincide with h on $\bigcap_{i=1}^m (X \setminus V_i)$, and we can easily verify that $\int_X (v_\varepsilon - u_\varepsilon) \leq \varepsilon \|h\|_\infty$.

The end of the proof is now straightforward: $A_n u_\varepsilon \leq A_n h \leq A_n v_\varepsilon$ for all n, and u_ε and v_ε are continuous, so we have the following:

$$\int_X h - \varepsilon\|h\|_\infty \leq \int_X u_\varepsilon = \underline{\lim}\, A_n u_\epsilon \leq \underline{\lim}\, A_n h \leq \overline{\lim}\, A_n h \leq \overline{\lim}\, A_n v_\varepsilon = \int_X v_\varepsilon \leq \int_X h + \varepsilon\|h\|_\infty.$$

Hence $A_n h$ tends to $\int_X h d\mu$, since ε was arbitrary. □

We now turn to the proof of Theorem 17.1.

Proof (Veech). Take $n_r = 2^r$ with *r even.* The function γ defined previously is continuous except at $-\theta = 11\ldots1\ldots$, therefore $x \mapsto \gamma(x, 2^r)$ has at most finitely many discontinuities. By Theorem 17.2 this brings us to the study of $\lim_{n\to\infty} \frac{1}{n}\sum_{k=0}^{n-1} \gamma(x + k\theta, 2^r)$ for some x that we find convenient. In fact we will estimate the last expression for $x = 0$, and all the difficulty remains in determining what values of τ we see in $\gamma(k\theta, 2^r) = \gamma(k\theta)\gamma((k+1)\theta)\ldots\gamma((k+2^r-1)\theta)$. For this we must examine a fairly arbitrary string of 2^r consecutive multiples of θ in G, i.e. $k\theta, (k+1)\theta, \ldots, (k+2^r-1)\theta$. Consider the dyadic expansions that appear in these multiples, $k, k+1, \ldots, k+2^r-1$. One half of them, that is 2^{r-1} of the total, have $\tau = 1$, so that $\gamma = 1$; one half of the remaining (those that start with a 1), that is 2^{r-2} of the total, have $\tau = 2$, so that $\gamma = -1$; etc. ; 2 of them have $\tau = r-1$, so that $\gamma = 1$. A picture will help.

Here's a string of 2^r multiples of θ in the case where $r = 5$ and $k = 4$.

		τ	γ
4θ	$= 0010000\ldots\ldots\ldots$	1	1
5θ	$= 1010000\ldots\ldots\ldots$	2	-1
6θ	$= 0110000\ldots\ldots\ldots$	1	1
7θ	$= 1110000\ldots\ldots\ldots$	4	-1
8θ	$= 0001000\ldots\ldots\ldots$	1	1
9θ	$= 1001000\ldots\ldots\ldots$	2	-1
10θ	$= 0101000\ldots\ldots\ldots$	1	1
11θ	$= 1101000\ldots\ldots\ldots$	3	1
12θ	$= 0011000\ldots\ldots\ldots$	1	1
13θ	$= 1011000\ldots\ldots\ldots$	2	-1
14θ	$= 0111000\ldots\ldots\ldots$	1	1
15θ	$= 1111000\ldots\ldots\ldots$	5	1
16θ	$= 0000100\ldots\ldots\ldots$	1	1
17θ	$= 1000100\ldots\ldots\ldots$	2	-1
18θ	$= 0100100\ldots\ldots\ldots$	1	1
19θ	$= 1100100\ldots\ldots\ldots$	3	1
20θ	$= 0010100\ldots\ldots\ldots$	1	1
21θ	$= 1010100\ldots\ldots\ldots$	2	-1
22θ	$= 0110100\ldots\ldots\ldots$	1	1
23θ	$= 1110100\ldots\ldots\ldots$	4	-1
24θ	$= 0001100\ldots\ldots\ldots$	1	1
25θ	$= 1001100\ldots\ldots\ldots$	2	-1
26θ	$= 0101100\ldots\ldots\ldots$	1	1
27θ	$= 1101100\ldots\ldots\ldots$	3	1
28θ	$= 0011100\ldots\ldots\ldots$	1	1
29θ	$= 1011100\ldots\ldots\ldots$	2	-1
30θ	$= 0111100\ldots\ldots\ldots$	1	1
$l(k)\theta = 31\theta$	$= 1111100\ldots\ldots\ldots$	6	-1
32θ	$= 0000010\ldots\ldots\ldots$	1	1
33θ	$= 1000010\ldots\ldots\ldots$	2	-1
34θ	$= 0100010\ldots\ldots\ldots$	1	1
35θ	$= 1100010\ldots\ldots\ldots$	3	1

To summarize: $2^{r-1} + 2^{r-3} + \cdots + 1$ have $\gamma = 1$, so their product is 1. $2^{r-2} + 2^{r-4} + \cdots + 2$ have $\gamma = -1$, so their product is 1. All together that makes $2^{r-1} + 2^{r-2} + \ldots + 1 = 2^r - 1$ terms, so that the sign of our expression $\gamma(k\theta, 2^r) = \gamma(k\theta)\gamma((k+1)\theta)\ldots\gamma((k+2^r-1)\theta)$ is determined by the remaining one.

Let $l(k)$ be the integer such that $\gamma(l(k)\theta)$ is the remaining factor whose sign we have to analyze, and let $\beta(k) = \tau(l(k)\theta) + 1$ (so that $\gamma(l(k)\theta) = (-1)^{\beta(k)}$). When $k = 0, 1, \ldots,$ $\tau(l(k)\theta)$ takes the values $r, r+1, r+2, \ldots$ with respective densities $1/2, 1/4, 1/8, \ldots$ (see figure). In other words,

$$\frac{\operatorname{card}\{k : 0 \le k \le n-1, \tau(l(k)\theta) = r+i\}}{n} \to \frac{1}{2^{i+1}} \qquad \text{as } n \to \infty, \text{ for } i = 0, 1, 2, \ldots, \text{ or}$$

$$\frac{\operatorname{card}\{k : 0 \le k \le n-1, \beta(k) = \beta\}}{n} \to \frac{1}{2^{\beta-r}} \qquad \text{as } n \to \infty, \text{ for } \beta \ge r+1.$$

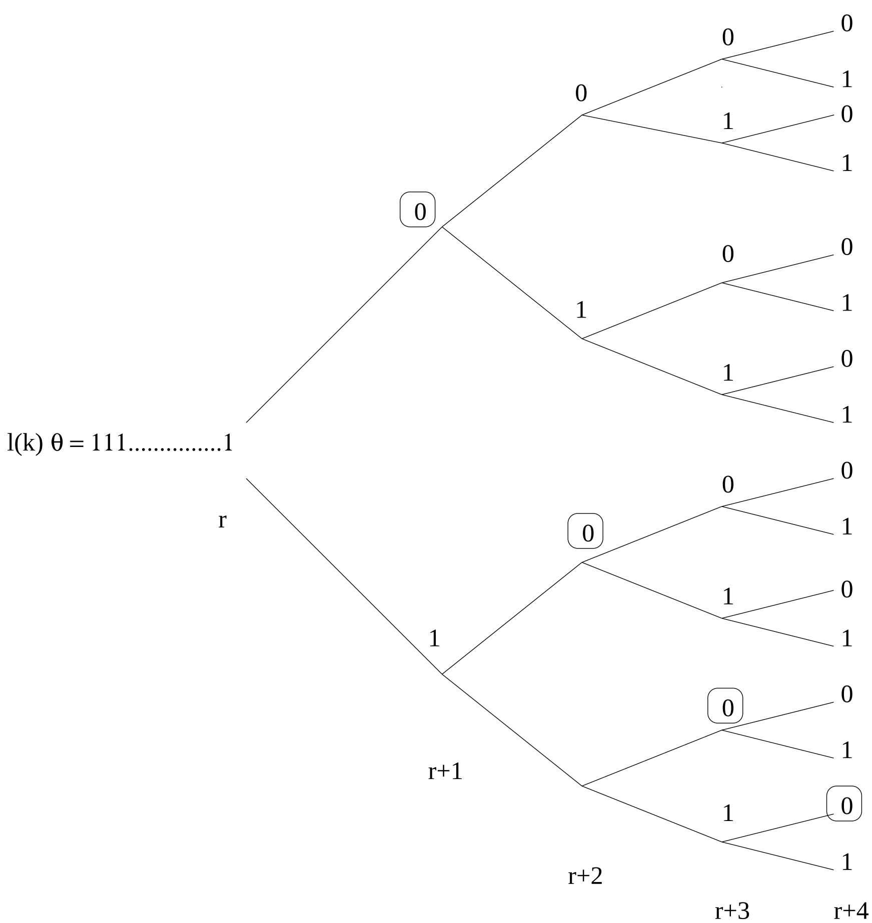

FIGURE 26. Densities of $\tau(l(k)\theta)$.

Hence
$$\frac{1}{n}\sum_{k=0}^{n-1}\gamma(k\theta,2^r)=\frac{1}{n}\sum_{k=0}^{n-1}\gamma(l(k)\theta)=\frac{1}{n}\sum_{\beta=r+1}^{\infty}(-1)^\beta\operatorname{card}\{k:0\le k\le n-1,\beta(k)=\beta\}$$
. Put $f_{n,\beta}=(-1)^\beta\operatorname{card}\{k:1\le k\le n-1,\beta(k)=\beta\}/n$; then $f_{n,\beta}$ tends to $(-1)^\beta/2^{\beta-r}$ and $|f_{n,\beta}|\le 1/2^{\beta-r}$. Hence, as an application of the Lebesgue Dominated Convergence Theorem (for example), we have that $\sum_{\beta\ge r+1} f_{n,\beta}$ tends to $\sum_{\beta\ge r+1}\frac{(-1)^\beta}{2^{\beta-r}}$ as $n\to\infty$.

Conclusion:
$$\begin{aligned}\int_G\gamma(x,2^r)dm(x)&=\lim_{n\to\infty}\frac{1}{n}\sum_{k=0}^{n-1}\gamma(k\theta,2^r)=\sum_{\beta=r+1}^{\infty}\frac{(-1)^\beta}{2^{\beta-r}}\\&=2^r\sum_{\beta=r+1}^{\infty}\left(-\frac{1}{2}\right)^\beta=2^r\left(-\frac{1}{2}\right)^{r+1}\frac{1}{1+\frac{1}{2}}\\&=-\frac{1}{3}\neq 1.\end{aligned}$$

This proves our theorem. □

18. March 17 *(Notes by GB)*

18.1. **The Spectrum of the Morse System.** For reference see S. Kakutani (1966), Fifth Berkeley Symposium in Math Stat., vol II, p. 405.
Let $(M,\sigma) \subset \{-1,1\}^{\mathbb{Z}}$ be the Morse system with unique invariant measure μ_M, and (X,σ) the Toeplitz system. Recall that $\hat{X} = X \times \{-1,1\}$ and $\hat{T} : \hat{X} \to \hat{X}$ is defined by $\hat{T}(u,\xi) = (\sigma u, f(u)\xi)$ and $f(u) = (-1)^{u_0}$. Also, $L^2(M,\mu_M)$ is isomorphic to $L^2(\hat{X},\mu)$ and is a separable Hilbert space.

To any measurable transformation $T : X \to X$ on a measure space $\big(X,\mathcal{B},\mu\big)$ is associated a linear operator U_T on $L^2\big(X,\mathcal{B},\mu\big)$ which is unitary if T is invertible (cf. e.g. Walters, p. 25) in the following way:

$$U_T h(x) = h(T(x)) \quad \text{for all } h \in L^2\big(X,\mathcal{B},\mu\big).$$

Then

$$< U_T^n f, g > = \int_X f \circ T^n \bar{g} d\mu \text{ for all } f,g \in L^2\big(X,\mathcal{B},\mu\big), \text{ and } n \in \mathbb{Z}.$$

For the Morse system, $\sigma : M \to M$ corresponds to $\hat{T} : \hat{X} \to \hat{X}$. The unitary operator $U_{\hat{T}} : L^2(\hat{X},\mu) \to L^2(\hat{X},\mu)$ takes the form

$$U_{\hat{T}} h(x,\xi) = h(\hat{T}(x,\xi)) = h(Tx, f(x)\xi)$$

or more briefly

$$U_{\hat{T}} h = h\hat{T}.$$

The *spectrum* of the system $(\hat{X},\hat{T})$ is defined as the spectrum of the operator $U_{\hat{T}}$ on $L^2(\hat{X},\mu)$.

Recall briefly the definitions of various kinds of spectral measures, spectral types: see Petersen's Lectures p. 10 or book, p. 19 for a quick review. The basic idea is that we can understand the operator by decomposing it as an integral: $U_{\hat{T}} = \int_{-\pi}^{\pi} e^{i\theta} dE(\theta)$, where E is a projection-valued Borel measure on $[-\pi,\pi)$:

$$U_{\hat{T}}^k = \int_{-\pi}^{\pi} e^{ik\theta} dE(\theta), \quad k \in \mathbb{Z}$$

$$< U_{\hat{T}}^k u, u > = \int_{\hat{X}} (\hat{T}^k u)\bar{u} d\mu = \rho(k) = \rho_u(k).$$

One can can check that $\rho(k)$ is positive-definite on $\mathbb{Z}$. Hence by the Bochner-Herglotz Theorem, ρ_u is the Fourier transform of a measure ν_u: $\rho_u(k) = \int_{-\pi}^{\pi} e^{-ik\theta} d\nu_u(\theta)$.

The function of $k < U_{\hat{T}}^k u, v >$ is not positive-definite if $u \neq v$. But then $< U_{\hat{T}}^k u, v >$ are the Fourier coefficients of a complex measure $\lambda_{u,v}$ that is not directly given by Bochner-Herglotz but by polarization: replace first u by $u+v$ and then by $u-v$ in the preceding construction.

The *maximal spectral type of the system* (or of $U_{\hat{T}}$) is the minimal (up to absolute continuity) type that dominates all these measures. The *discrete spectrum* of $U_{\hat{T}}$ is just the set of eigenvalues of $U_{\hat{T}}$ (they lie on the unit circle since it's a unitary operator). These are point masses (or atoms) of E. $E\{\lambda\}$ is the projection onto the eigenspace associated with λ, that is, $E\{\lambda\}$ is the projection on $\{h \in L^2 : \hat{T}h = \lambda h\}$.

Let $\mathcal{K} \subset L^2$ be the closed linear span of the eigenfunctions (note there is a countable number by separability). This is the *Kronecker subspace* of L^2.
If $(X, \mathcal{B}, \mu)$ is a measure space, $T : X \to X$ a measure-preserving transformation, then $\mathcal{K}$ the closed linear span of the eigenfunctions, corresponds to a factor: There are a measure-preserving system $(Y, \mathcal{C}, \nu, S)$ and a factor map $\pi : (X, \mathcal{B}, \mu, T) \to (Y, \mathcal{C}, \nu, S)$ such that:

$$\mathcal{K} = \{h \circ \pi : h \in L^2(Y, \mathcal{C}, \nu)\}.$$

This $(Y, \mathcal{C}, \nu, S)$ is called the *Kronecker factor* of $(X, \mathcal{B}, \mu, T)$. It is the largest factor with purely discrete spectrum. In $L^2(Y, \mathcal{C}, \nu)$ (and in $L^2(X, \mathcal{B}, \mu,))$ the eigenspaces are all one-dimensional and pairwise orthogonal. Actually, then $(Y, \mathcal{C}, \nu, S)$ is isomorphic to an ergodic group rotation with Haar measure, a system like the odometer or an irrational rotation on the circle. If a system does not have purely discrete spectrum, the analysis of its structure is more difficult.

For the Morse system (M, σ) there is a *dualizing* (or *mirroring*) map $\delta : M \to M$ given by

$$\delta x = \bar{x} = \dots, \bar{x}_{-1}, \bar{x}_0, \bar{x}_1, \text{ where } \bar{0} = 1, \bar{1} = 0.$$

δ maps the Morse system $= (\overline{\mathcal{O}(\dots 01101001 \dots)}, \sigma)$ into itself.

Theorem 18.1 (Kakutani)**.** *For the Morse system, represented as $(\hat{X}, \hat{T}, \mu)$ as before,*

$$L^2(\hat{X}, \mu) = V_0 \oplus V_1, \text{ where}$$

$$V_0 = \{h \in L^2(\hat{X}, \mu) : h(x, \xi) = h(x, -\xi)\}$$
$$V_1 = \{h \in L^2(\hat{X}, \mu) : h(x, \xi) = -h(x, -\xi)\}.$$

Moreover, $U_{\hat{T}}$ has discrete spectrum on V_0 with eigenvalues at each diadic rational. On V_1, $U_{\hat{T}}$ has continuous spectrum. In fact there are no eigenfunctions outside V_0, so $V_0 = \mathcal{K} =$ Kronecker subspace of $L^2(\hat{X}, \hat{T}, \mu)$.

Remark 18.1. It can be shown that the maximal spectral type of $U_{\hat{T}}|_{V_1}$ is singular with respect to the Lebesgue measure, and can be given fairly explicitly as a Riesz product and a g-measure. Kakutani further showed that by varying the construction of $\hat{X}$ slightly, one could obtain many examples with pairwise singular maximal spectral types for these $U_{\hat{T}}|_{V_1}$.

There is much subsequent work on spectral types of Morse-like systems. These are steps in the direction of "Banach's problem": find a system with simple, in the sense of multiplicity, Lebesgue spectrum, i.e. $(X, \mathcal{B}, \mu, T)$ such that in $L^2(X, \mathcal{B}, \mu)$ there is a function ϕ such that $1, \phi, \phi T, \phi T^2, \dots$ are pairwise orthogonal and span $L^2(X, \mathcal{B}, \mu)$.

Proof. Recall that there is the factor mapping π:

$$\hat{X} = X \times \{-1, 1\} \xrightarrow{\pi} (X, \sigma) = \text{Toeplitz system} \approx \text{odometer} = (G, T),$$

and that (G, T) is a compact abelian group rotation and so has discrete spectrum.

Let $\hat{G}$ be the character group of G, i.e. the set of all continuous $\gamma : G \to \{z \in \mathbb{C} : |z| = 1\}$ such that $\gamma(g_1 g_2) = \gamma(g_1)\gamma(g_2)$ for all $g_1, g_2 \in G$. Then each $\gamma \in \hat{G}$ is an eigenfunction of T, and $\hat{G}$ spans $L^2(G, m)$.

V_0 does not depend on the second coordinate, so

$$V_0 = \{h \circ \pi_1 : h \in L^2(X, \mu_X)\} \approx L^2(X, \mu_X).$$

We know that the Toeplitz system is isomorphic to (G, T), so we have that $U_{\hat{T}}|_{V_0}$ has purely discrete spectrum with eigenvalues at the diadic rationals. So V_0 is contained in the Kronecker subspace of $L^2(\hat{X}, \hat{T}, \mu)$.

To prove $L^2(\hat{X}, \mu) = V_0 \oplus V_1$, take $u_0 \in V_0, u_1 \in V_1$; then

$$\int_{\hat{X}} u_0 \bar{u}_1 d\mu = \int_{X \times \{1\}} u_0 \bar{u}_1 d\mu + \int_{X \times \{-1\}} u_0 \bar{u}_1 d\mu = \int_{X \times \{1\}} u_0 \bar{u}_1 d\mu + \int_{X \times \{1\}} u_0 (-\bar{u}_1) d\mu = 0,$$

so $V_0 \perp V_1$.

Now we check that V_0 and V_1 span L^2. Take $f \in L^2$ and decompose it as

$$f(x, \xi) = \underbrace{\frac{f(x,\xi) + f(x, -\xi)}{2}}_{= P_{V_0} f \in V_0} + \underbrace{\frac{f(x,\xi) - f(x, -\xi)}{2}}_{= P_{V_1} f \in V_1},$$

which is the sum of an element of V_0 and V_1. Since any $f \in L^2$ has this decomposition, we have $L^2 = V_0 \oplus V_1$.

Now let's see that there are no eigenfunctions except in V_0. Suppose $h \in L^2(\hat{X}, \mu), \lambda \in \mathbb{C}$, and $h\hat{T} = \lambda h$. Then

$$(P_{V_0} h)\hat{T}(x, \xi) = \frac{h\hat{T}(x,\xi) + h\hat{T}(x, -\xi)}{2} = \lambda \frac{h((x,\xi)) + h((x, -\xi))}{2} = \lambda P_{V_0} h.$$

Suppose there exists $v \in V_1$ with $v\hat{T} = \lambda v$. Then $v^2 \in V_0$ and $v^2 \hat{T} = \lambda^2 v^2$. So λ^2 is a diadic rational by the result for V_0, and hence so is λ. Now v is an eigenfunction and the corresponding eigenvalue is a diadic rational. Each eigenspace is one-dimensional by ergodicity, so $v \in V_0$, and therefore $v = 0$ (since also in V_1). This show that $U_{\hat{T}}$ has continuous spectrum on V_1. Now let h be arbitrary in L^2. Then $h = P_{V_0} h + P_{V_1} h$ is an eigenfunction with eigenvalue λ. From above also $P_{V_0} h$ has eigenvalue λ, so does $P_{V_1} h$, but then $P_{V_1} h$ must vanish. □

19. March 19 *(Notes by GB)*

19.1. **Sturmian Systems.** This is a parallel construction to the Toeplitz system. The name comes from the fact that this system is related to the zeros of the classical Sturm-Liouville differential equation: $y'' + f(x)y = 0$ with f periodic in x.
Some references:

- Hedlund and Morse (1938,1940), *Symbolic Dynamics I and II, Am. J. Math 60 and 62*
- Hedlund (1944), *Sturmian minimal sets, Am. J. Math 66*
- Coven and Hedlund (1973), *Sequences with minimal block growth, Math. Syst. Theo. 7*
- Ferenczi (1996), *E.T.D.S. 16*
- Petersen and Shapiro (1973), *Induced flows, Trans. AMS 177*

Let $G = [0,1)$ with addition mod 1, $\alpha \notin \mathbb{Q}$ and $T : G \to G$ the rotation defined by $Tx = x + \alpha$ (mod 1). Consider the following coding.

Let $0 < \beta < 1$ and $\mathcal{P}$ be the partition of $[0,1)$ into two intervals, $\mathcal{P} = \{[0,\beta),[\beta,1)\}$, and $\mathcal{P}'$ the partition without the end-points, $\mathcal{P}' = \{(0,\beta),(\beta,1)\}$. We define the *Sturmian system* $S(\alpha,\beta)$ to be the closure in Σ_2 of the set of all sequences (x_n), where $x_n = \chi_{[0,\beta)}(g + n\alpha)$ for all $n \in \mathbb{Z}$ and some $g \in G$. That is, we code the T-orbits of points in G by their *itineraries*, assigning 0 when they visit the first interval, 1 when they visit the second.

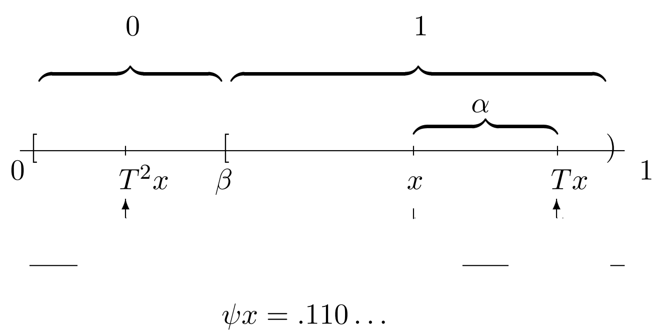

Figure 27. Defining the Sturmian system

Sturmian systems provide examples of $\{0,1\}$ sequences with *minimal block growth*: For $x \in \Sigma_2$ let $N_n(x)$ be the number of different n-blocks that appear in x. If x is a periodic sequence, then $N_n(x)$ is bounded. Even more (exercise): If there is n such that $N_n(x) \leq n+1$, then x is periodic. Some Sturmian systems have $N_n(x) = n+1$ for all $n \in \mathbb{Z}$. So in a sense, Sturmian sequences have minimal complexity among nonperiodic ones (See Coven and Hedlund, Ferenczi).

We give now more precisely the construction of the Sturmian system. As in the Toeplitz situation, define ψ "not quite a map": $\psi : G \to \Sigma_2$, $\psi(x) = 1$ or 2 points for all $x \in G$. This is because we want $\psi G \subset \Sigma_2$ to be closed, shift-invariant, minimal and uniquely ergodic. There is again a bad set where it is harder to define the map, $G_0 = \mathcal{O}(0) \cup \mathcal{O}(\beta)$. On G_0 we will put both 0 and 1 in certain entries, as for the Toeplitz case. Let us first examine the case $x \in G \setminus G_0$. There, $\psi(g) =$ the unique sequence which is the unambiguous coding of G by

$$\psi(x)_n = \chi_{[0,\beta)}(x + n\alpha)$$

Note that in this case, it is the same to consider the partition $\mathcal{P}$ or $\mathcal{P}'$.

Let us define $\psi(x)$ more carefully in case x is in the bad set G_0,

$$\begin{aligned} G_0 &= \mathcal{O}(0) \cup \mathcal{O}(\beta) \\ &= (\mathbb{Z}\alpha \mod 1) \cup (\beta + \mathbb{Z}\alpha \mod 1) \text{ (disjoint union if } \beta \notin \mathbb{Z}\alpha). \end{aligned}$$

If $x \in G_0$ define $\psi(x)$ to be 2 sequences in Σ_2. There are two cases:

(i) The case $\beta \notin \mathbb{Z}\alpha = \mathcal{O}(0)$.

From the decomposition of G_0 as a disjoint union, either $x \in \mathcal{O}(0)$ or $x \in \mathcal{O}(\beta)$, but not both. There exists a unique n such that $T^n(x) = 0$ (or β for the other case) and never hits again one of these two points (since $\beta \notin \mathbb{Z}\alpha$). Define

$$(\psi(x))_m = \begin{cases} \chi_{[0,\beta)}(T^m(x)) & \text{for } m \neq n \\ \text{both 0 and 1} & \text{for } m = n \end{cases}$$

Therefore we have defined ψ in such a way that

$$\psi(x) = \{u, v\}, \quad \text{with } u_m = v_m \text{ for } m \neq n \text{ and } u_n = 0, \quad v_n = 1.$$

(ii) The case $\beta \in \mathbb{Z}\alpha$, say $\beta = j\alpha = T^j 0$.

Suppose $T^n x = 0$, so $T^{n+j}x = \beta$. $T^m x \neq 0$ or β unless $m = n$ or $n + j$. In this case define

$$\psi(x) = \{u, v\}, \quad \text{with } \begin{cases} u_m = v_m = \chi_{[0,\beta)}(T^m(x)) & \text{for } m \neq n \text{ or } n+j, \\ u_n = 0, \quad v_n = 1, \\ u_{n+j} = 1, \quad v_{n+j} = 0 \end{cases}$$

This amounts to splitting each of 0 and β into a "left half" and "right half" and similarly for each point in their orbits.

Now define $\phi : \psi G \to G$ to be $\phi = \psi^{-1}$ (an actual map).

Note that clearly shifting the sequence corresponds to translating the point:

$$\sigma\psi x = \psi T x \quad \text{for all } x \in G;$$

therefore

$$\phi\sigma u = T\phi u \text{ for } u \in \psi G$$

(since $\phi u = x$ implies $u \in \psi x$ so $\sigma u \in \psi T x$, and hence $\phi\sigma u = Tx = T\phi u$). We have just shown that $\boxed{\phi\sigma = T\phi}$.

We claim that ψG is closed. Suppose $u_k \in \psi x_k \in \psi G$ and $u_k \to u \in \Sigma_2$.

Suppose we have a subsequence u_{k_i} such that $x_{k_i} \to y \in G$. We want to show that u is in ψy. This would show that ψG is closed and also that ϕ is *continuous*, since then any convergent subsequence of ϕu_k converges to ϕu.

Suppose $y \notin G_0$. Pick a large k such that u_k and u agree on a large central block $[-J, J]$. If x is in a sufficiently small neighborhood N of y, then $T^j x$ and $T^j y$ are either both in $(0, \beta)$ or both not in $(0, \beta)$, for all $j \in [-J, J]$. We are just making sure that x and y are in the same cell of the partition $\bigvee_{-J}^{J} T^j \mathcal{P}'$. We are taking $N = \bigcap_{-J}^{J} T^{-j} I_j$ with I_j equal to either $(0, \beta)$ or $(\beta, 1)$ depending on which one $T^j y$ is in.

If i is large enough that $k_i \geq k$ and $x_{ki} \in N$, then the itinerary of x_{k_i} from time $-J$ to J is the same as that of y. The central $(2J+1)$-block of u_{k_i} is the same as the central $(2J+1)$-block of u. Since J is arbitrary, $\psi y = \{u\}$.

What if $y \in G_0$? There are at most two times n and $n+j$ when $T^m y = 0$ or β which requires special attention. At these times, for large enough k, all x_k's have to be on the side of 0 or β determined by the n'th coordinate of u. Suppose for example that $y = \beta$ and $u - 0 = 0$. Then for large i, x_{k_i} is near y and to its right. Suppose $\beta \notin \mathbb{Z}\alpha$ so the orbit of y hitS β for time 0 only. Then ψy is two points, one of which is u.

If $\beta \in \mathbb{Z}\alpha$ and again for example $y = \beta$ and $u_0 = 0$, then again for large i x_{k_i} is to the right of β, but also $T^{-1}x_k$ is to the right of 0 so $u_{-1} = 1$. Again $u \in \psi y$.

From above, ψG is shift-invariant, so $(\psi G, \sigma)$ is a *closed subshift* of Σ_2.

Claim: every orbit in ψG is dense. Let $u \in \psi x$ for some $x \in G$. Take $v \in \psi y$ for some $y \in G$. We want to find n such that $\sigma^n u$ agrees with v on a central $(2J+1)$-block. Take a small neighborhood of y such that every point in that neighborhood has the same $\mathcal{P}$ coding for the time interval $[-J, J]$. We hit $\{0, \beta\}$ at most twice with y. Looking at v tells us whether to use $(y-\delta, y)$ or $(y, y+\delta)$. Choose n such that $T^n x$ is in the selected interval. Then σu and v will agree with u on their central $(2J+1)$-blocks.

Corollary 19.1. *The Sturmian systems $(S(\alpha,\beta),\sigma)$ are uniquely ergodic.*

Proof. As for the Toeplitz systems, the factor map ϕ

$$(S(\alpha,\beta),\sigma) \xrightarrow{\phi} (G,T)$$

is 1 to 1 except on a countable set $H_0 = \psi(\mathbb{Z}\alpha \cup (\beta + \mathbb{Z}\alpha))$, that is $\phi H_0 = G_0$. (G,T) is uniquely ergodic with invariant measure $m =$ Lebesgue measure, and G_0 is countable and so has Lebesgue measure 0. Then any σ-invariant measure on $(S(\alpha,\beta),\sigma)$ assigns mass 0 to H_0 and is uniquely determined on $(S(\alpha,\beta),\sigma) \setminus H_0$ by the isomorphism

$$\psi : G \setminus G_0 \to S(\alpha,\beta) \setminus H_0.$$

□

20. March 24 *(Notes by DJS)*

20.1. **Subshifts of Finite Type.** (For reference see Lind and Marcus, and Kitchens books.)

The definition of a subshift of finite type begins with a finite collection of forbidden (finite length) words on a given alphabet, $D = \{0, 1, ..., d\}$ say. Denote this collection by $\mathcal{F}$, and note that $\mathcal{F} \subset \{0, 1, ..., d-1\}^*$. Then define $X_\mathcal{F} \subset D^\mathbb{Z}$ (or $X_\mathcal{F}^+ \subset D^\mathbb{N}$) to be the set of all sequences none of whose sub-blocks are in $\mathcal{F}$. Since $X_\mathcal{F}$ is closed and σ-invariant, $(X_\mathcal{F}, \sigma)$ is a subshift, and this defines a *subshift of finite type (SFT)*. (To see that $X_\mathcal{F}$ is closed, consider a sequence x in its complement. Since x contains some bad word, any sequence sufficiently close to it also contains that word. So there is a neighborhood containing x which does not intersect with $X_\mathcal{F}$, showing that $X_\mathcal{F}$ is closed. The property of σ-invariance is immediate.)

There are many reasons for studying these subshifts, some of which are listed below:

(1) It was shown in Spring '97 Math 261 that SFTs model attractors of Axiom A diffeomorphisms (based on codings of Markov partitions). See Theorem 30.1 on page 63 of those notes.
(2) They are the natural domain for Markov measures.
(3) They model inputs to certain information transmission and storage devices; *e.g.* some devices may not be able to handle quickly alternating blocks of zeros and ones.
(4) They model changing situations in which the choice of future state depend only on the current state.

20.2. **Graph Representations of Subshifts of Finite Type.** We may assume that all of the words in $\mathcal{F}$ are of the same length $m+1$. Then we say that $(X_\mathcal{F}, \sigma)$ is *m-step* or has *memory m*. For example, if $D = \{0, 1\}$ and $\mathcal{F} = \{11, 101\}$, then change $\mathcal{F}$ to $\{110, 011, 111, 101\}$ by extending the shorter word with all possible sub-blocks. Here $(X_\mathcal{F}, \sigma)$ has memory 2.

We represent $(X_\mathcal{F}, \sigma)$ by two kinds of *graph shifts*; namely *edge shifts* and *vertex shifts.* If $(X_\mathcal{F}, \sigma)$ has memory m, construct a directed graph $G_\mathcal{F}$ with vertices $\mathcal{V}(G_\mathcal{F})$ and edges $\mathcal{E}(G_\mathcal{F})$ as follows:

- The vertices are (and are labeled by) the allowed m-blocks in $X_\mathcal{F}$.
- Put an edge from B_2 to B_1 if for $a, z \in D$ and $B \subset D^{m-1}$, $B_2 = aB$ and $B_1 = Bz$. Label this edge as aBz.

This defines higher block codings of $X_\mathcal{F}$ to $X_\mathcal{F}^{[m+1]} \subset \mathcal{E}(G_\mathcal{F})^\mathbb{Z}$ and $X_\mathcal{F}^{[m]} \subset \mathcal{V}(G_\mathcal{F})^\mathbb{Z}$, and the resulting subshifts, $(X_\mathcal{F}^{[m+1]}, \sigma)$ and $(X_\mathcal{F}^{[m]}, \sigma)$, are isomorphic to an *edge shift* and *vertex shift*, respectively, which are defined in the next section. In each of these two ways, by tracking edges or vertices, a bi-infinite walk on the graph is a representation of a sequence in $X_\mathcal{F}$. Theorem 2.3.2 of Lind and Marcus states that any m-step subshift of finite type can be represented by doubly-infinite walks on a graph.

Example 20.1. For $\mathcal{F} = \{110, 011, 111, 101\}$, we have $\mathcal{V}(G_\mathcal{F}) = \{00, 01, 10\}$ and $\mathcal{E}(G_\mathcal{F}) = \{000, 001, 010, 100\}$. $(X_\mathcal{F}, \sigma)$ may be represented by the graph of Figure 28.

A walk along edges $100 \to 000 \to 001$ visits vertices $10 \to 00 \to 00 \to 01$ in that order. The represented block is 10001. We can see the edge-shift representation by examining each 3-block of the sequence:

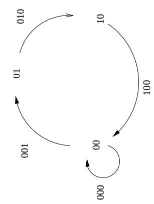

FIGURE 28. Graph representation of $(X_{\mathcal{F}}, \sigma)$ for $\mathcal{F} = \{110, 011, 111, 101\}$

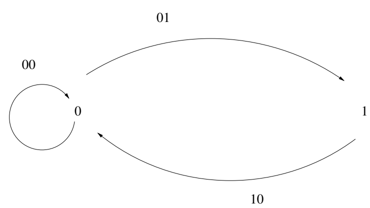

FIGURE 29. Graph representation of $(X_{\mathcal{F}}, \sigma)$ for $\mathcal{F} = \{11\}$

$$\begin{array}{ccccc} 1 & 0 & 0 & . & . \\ . & 0 & 0 & 0 & . \\ . & . & 0 & 0 & 1 \\ \end{array}$$
$$\left[\begin{array}{ccccc} 1 & 0 & 0 & 0 & 1 \end{array} \right]$$

The vertex-shift representation is seen in the 2-blocks of the sequence:

$$\begin{array}{ccccc} 1 & 0 & . & . & . \\ . & 0 & 0 & . & . \\ . & . & 0 & 0 & . \\ . & . & . & 0 & 1 \\ \end{array}$$
$$\left[\begin{array}{ccccc} 1 & 0 & 0 & 0 & 1 \end{array} \right]$$

Example 20.2. *(Golden Mean SFT)*
For $\mathcal{F} = \{11\}$, we have $\mathcal{V}(G_{\mathcal{F}}) = \{0, 1\}$ and $\mathcal{E}(G_{\mathcal{F}}) = \{00, 01, 10\}$. $(X_{\mathcal{F}}, \sigma)$ may be represented by the graph of Figure 29.

20.3. **General Edge and Vertex Shifts.** Let G be a directed graph with edge set $\mathcal{E}(G)$ and vertex set $\mathcal{V}(G)$. The *edge shift*, $(X_{\mathcal{E}}(G), \sigma)$, determined by G is the subshift of $\mathcal{E}(G)^{\mathbb{Z}}$ which consists of all sequences of edges which are the paths of all doubly-infinite walks along edges of G. The *vertex shift*, $(X_{\mathcal{V}}(G), \sigma)$, determined by G is the subshift of $\mathcal{V}(G)^{\mathbb{Z}}$ which consists of all sequences of vertices for such walks on G.

A general vertex shift (on d vertices) is described by a $d \times d$ matrix with entries in $\{0, 1\}$, where a 1 in the $(i, j)^{th}$ entry indicates an edge from vertex i to vertex j. A general edge shift is described by a $d \times d$ (nonnegative) integer matrix where the $(i, j)^{th}$ entry represents the number of edges from vertex i to vertex j. In either case, the matrix is called an *incidence* or *adjacency matrix.*

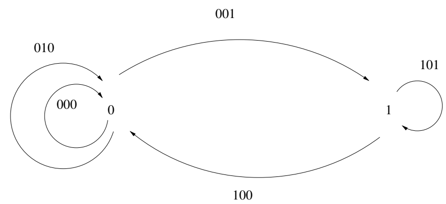

FIGURE 30. Graph of A^2 for Golden Mean SFT

Conversely, any nonnegative integral $d \times d$ matrix A corresponds to some graph G_A (which may have multiple edges between the same pair of vertices) and an edge shift $X_{\mathcal{E}}(G_A) = X_A$. It is interesting and useful to notice that for a transition matrix A, the $(i,j)^{th}$ entry of A^n represents the number of paths *of length* n from vertex i to vertex j.

Example 20.3. The edge shift corresponding to the Golden Mean SFT of Example 20.2 has transition matrix

$$A = \begin{array}{c} \\ 0 \\ 1 \end{array} \begin{array}{c} \begin{array}{cc} 0 & 1 \end{array} \\ \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix} \end{array}.$$

Then

$$A^2 = \begin{array}{c} \\ 0 \\ 1 \end{array} \begin{array}{c} \begin{array}{cc} 0 & 1 \end{array} \\ \begin{bmatrix} 2 & 1 \\ 1 & 0 \end{bmatrix} \end{array},$$

so there are two ways to go from 0 to 0 in two steps. The graph of the matrix A^2, representing paths of length 2 between vertices, is given in Figure 30.

A nonnegative integral $d \times d$ matrix A is called *irreducible* and the corresponding graph *(strongly) connected* if for every $i, j \in \mathcal{V}(G_A)$, there is a $n \in \mathbb{N}$ so that $(A^n)_{ij} > 0$. (*i.e.* there is a path in G_A of length n that starts at i and ends at j.) A is called *primitive* or *aperiodic* if there is a $n \in \mathbb{N}$ so that $(A^n)_{ij} > 0$ for every $i, j \in \mathcal{V}(G_A)$. Notice that aperiodicity implies irreducibility, but the converse does not hold. These terms are also applied to *nonintegral* matrices.

The properties of irreducibility and aperiodicity of A relates to topological properties of the edge shift of G_A. We have that A is irreducible if and only if (X_A, σ) is *topologically ergodic*. Also, A is aperiodic if and only if (X_A, σ) is *topologically (strongly) mixing*. What's more is that the *topological entropy* of the edge shift is determined by A: Set $N_n(A)$ to be the number of n-blocks in the language of the edge shift, (X_A, σ). So

$$N_n(A) = \sum_{i,j} (A^n)_{ij}.$$

Then the topological entropy of (X_A, σ) is

$$h_{top}(X_A, \sigma) = \lim_{n \to \infty} \frac{1}{n} \log N_n(A).$$

It follows from the Perron-Frobenius Theorem that $h_{top}(X_A, \sigma) = \log \lambda_A$, where λ_A is the largest positive eigenvalue of A, called the *Perron-Frobenius eigenvalue*.

To see that irreducibility implies topological ergodicity, consider blocks $B_1, B_2 \in \mathcal{L}(X_{\mathcal{E}}(G))$. Then for some n there is a path of length n starting at B_1 and ending at B_2. Associated with this path is an $x \in X_{\mathcal{E}}(G)$ which contains both B_1 and B_2 separated by n shifts:

$$x = \dots \overbrace{B_1 \dots}^{n} B_2 \dots$$

Then $\sigma^n[B_1] \cap [B_2] \neq \emptyset$, and $(X_{\mathcal{E}}(G), \sigma)$ is topologically ergodic.

21. March 26 *(Notes by DJS)*

From last class we had a theorem relating a transition matrix to the topological entropy of the associated subshift of finite type.

Theorem 21.1. *Let (X_A, σ) be a SFT given by an edge shift with associated (nonnegative) integral irreducible matrix A. Then the topological entropy of the system is*

$$h_{top}(X_A, \sigma) = \lim_{n\to\infty} \frac{1}{n} \log(\#\text{ of } n\text{-blocks in } X_A) = \lim_{n\to\infty} \frac{1}{n} \log \sum_{i,j} (A^n)_{ij} = \log \lambda_A,$$

where λ_A is the positive eigenvalue of A of maximum modulus among all eigenvalues (the Perron eigenvalue).

(Since $n/(n+1) \to 1$ it doesn't matter whether we count blocks on vertices or blocks on edges.)
Theorem 21.1 follows directly from the *Perron-Frobenius Theorem.*

Theorem 21.2. *(Perron-Frobenius Theorem) Let A be a nonnegative irreducible (and not necessarily integral) $d \times d$ matrix. Then A has a strictly positive right (column) eigenvector $\mathbf{r}$ ($r_i > 0$ for every $i = 1, \dots, d$). The corresponding eigenvalue λ_A is positive and it is the unique eigenvalue with these properties. Other properties of λ_A are: if ζ is any other eigenvalue of A, then $\lambda_A \geq |\zeta|$; λ_A is a simple root of the characteristic polynomial of A; its eigenspace is one-dimensional; and there is a $c > 0$ such that*

$$\frac{1}{c}\lambda_A^n \leq \sum_{i,j} (A^n)_{ij} \leq c\lambda_A^n$$

for all $n = 1, 2, \dots$ (so λ_A is unique with this property and further $\lambda_{A^{tr}} = \lambda_A$.) So A also has a positive left (row) eigenvector $\mathbf{l}$, unique up to a constant multiple. (Normalize so that $\mathbf{lr} = 1$.)

If A is aperiodic (primitive), then $\lambda_A > |\zeta|$ for every other eigenvalue ζ of A. Finally, in the aperiodic case

$$(A^n)_{ij} = (l_i r_i + \varepsilon_{ij}(n))\lambda_A^n,$$

where each $\varepsilon_{ij}(n) \to 0$ as $n \to \infty$.

The proof will be given for the aperiodic case, following the one in Seneta's book.

Proof. Suppose A is aperiodic. Let $\mathcal{P} = \{(\text{column vectors})\ \mathbf{x} \in \mathbb{R}^d : x_i \geq 0 \text{ for all } i = 1, \dots, d\}$. Then $A\mathcal{P} = \mathcal{P}$, and we want to find a *fixed direction.*

For an $\mathbf{x} \in \mathcal{P}$ with $\mathbf{x} \neq \mathbf{0}$ (*i.e.* some $x_i > 0$) define

$$S(x) = \min_{i=1,\dots,d} \begin{cases} \dfrac{(A\mathbf{x})_i}{x_i} & \text{if } x_i \neq 0 \\ \infty & \text{if } x_i = 0 \end{cases}$$

Then $0 \leq S(\mathbf{x}) < \infty$ for every $\mathbf{x} \in \mathcal{P}\backslash\{\mathbf{0}\}$, and for every i, $x_i S(\mathbf{x}) \leq (A\mathbf{x})_i$. So

$$\sum_i x_i S(\mathbf{x}) \leq \sum_i (A\mathbf{x})_i;$$

that is

$$S(\mathbf{x})\mathbf{1x} \leq \mathbf{1}A\mathbf{x} = (\mathbf{1}A)\mathbf{x},$$

where $\mathbf{1}$ is the row vector of ones. But $\mathbf{1}A$ is the row vector whose j^{th} entry is $\sum_i A_{ij}$, so setting $M = \max_j \sum_i A_{ij}$ we have

$$S(\mathbf{x})(\mathbf{1x}) \leq M(\mathbf{1x}).$$

Therefore, $S(\mathbf{x})$ is uniformly bounded by M.

Now let

$$\lambda_A = \sup_{\mathbf{x}\in\mathcal{P}\setminus\{\mathbf{0}\}} S(\mathbf{x}).$$

By irreducibility, $\lambda_A \geq S(\mathbf{1}^{tr}) > 0$, so $0 < \lambda_A \leq M$. Also, $S(\mathbf{x}/\|\mathbf{x}\|) = S(\mathbf{x})$, so

$$\lambda_A = \sup_{\mathbf{x}\in\mathcal{P}\,:\,\|\mathbf{x}\|=1} S(\mathbf{x}).$$

Now, $S(\mathbf{x})$ is upper semicontinuous on $\{\mathbf{x}\in\mathcal{P} : \|\mathbf{x}\| = 1\}$ (if $\mathbf{x}^{(k)} \to \mathbf{x}$ then $\limsup S(\mathbf{x}^{(k)}) \leq S(\mathbf{x})$), and since every upper semicontinuous function on a compact set achieves its absolute maximum, there is an $\mathbf{x} \in \{\mathbf{x}\in\mathcal{P} : \|\mathbf{x}\| = 1\}$ with $S(\mathbf{x}) = \lambda_A$. We have that $(A\mathbf{x})_i \geq \lambda_A x_i$ for every i, with equality for some i.

Now let

$$\mathbf{u} = A\mathbf{x} - \lambda_A\mathbf{x} \geq 0,$$

and suppose that $\mathbf{u} \neq 0$. Aperiodicity implies that we can choose n such that $A^n > 0$. Then

$$A^n(A\mathbf{x}) - \lambda_A A^n\mathbf{x} > 0.$$

But then

$$A(A^n\mathbf{x}) - \lambda_A A^n\mathbf{x} > 0,$$

which is impossible since for any $\mathbf{v} \in \mathcal{P}\setminus\{\mathbf{0}\}$ (including $\mathbf{v} = A^n\mathbf{x}$), $(A\mathbf{v})_i/v_i \leq \lambda_A$ for some i. Therefore $\mathbf{u} = 0$, which shows that λ_A is an eigenvalue of A with $A\mathbf{x} = \lambda_A\mathbf{x}$. Furthermore, $A^n\mathbf{x} > 0$ so $A^n\mathbf{x} = \lambda_A^n\mathbf{x}$ implies that $\mathbf{x} > 0$.

Now, for any i and n, $\lambda_A^n x_i = (A^n\mathbf{x})_i = \sum_j (A^n)_{ij}x_j$, so

$$\lambda_A^n = \frac{1}{x_i}\sum_j (A^n)_{ij}x_j.$$

Thus, for $c = (\max x_i)/(\min x_i)$ we have that

$$\frac{1}{c}\lambda_A^n \leq \sum_{i,j}(A^n)_{ij} \leq c\lambda_A^n$$

holds for any $n = 1, 2 \ldots$ □

22. March 31 *(Notes by DD)*

Last week we used the Perron-Frobenius Theorem to see (among other things) that the topological entropy of a subshift of finite type (X_A, σ) (an edge shift given by a square nonnegative integer matrix A) is $\log \lambda_A$, the logarithm of the maximum positive (Perron) eigenvalue of the transition matrix. Thus if $D = \{0, \ldots, d-1\}$ is the alphabet and we let

$$N_n(X_A) = \operatorname{card}(\mathcal{L}(X_A) \cap D^n) \text{ for } n = 1, 2, \ldots$$

denote the number of n-blocks seen in sequences in X_A, then

$$h_{\text{top}}(X_A, \sigma) = \lim_{n \to \infty} \frac{1}{n} \log N_n(X_A) = \log \lambda_A.$$

Topological entropy is a measure of the size or richness of a language, or of how much freedom to concatenate words it allows. Claude Shannon, the founder of information theory, also considered a subshift of finite type to be a type of communication system or *signal generator*, and then its topological entropy is a measure of the number of different messages it can produce in a given time; this could be called the *capacity* of the device. (See the Appendix of the Parry-Tuncel book.) If (X_A, σ) is an edge shift, we consider the vertices of the underlying graph G to be *states* of the device and edges leaving each state (thought of as being labeled by elements of the alphabet D) to represent the *symbols* that the device can print or transmit if it is in that state.

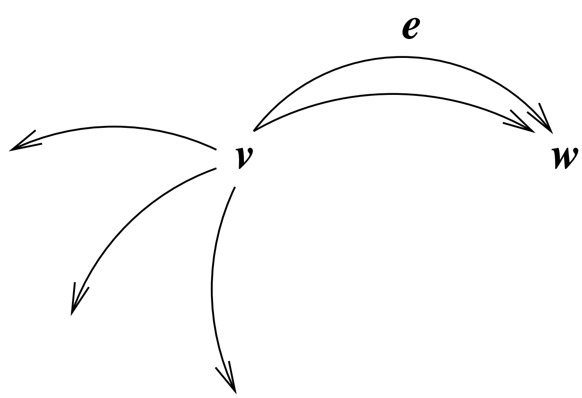

Figure 31. The Shannon Machine G

If there is an edge e from state v to state w, when the device is in state v it can emit symbol e and end up in state w. This is supposed to model some of the practical restrictions that actually occur in communication systems: for example, in Morse code (not the same Morse as before), successive spaces are not allowed (all spaces of any length are considered equivalent), and in ordinary typing we might demand that each period be followed by a space, then a capital letter. We could also have a cost, length, or transmission time associated with each symbol, but here we will keep each symbol having length 1.

It is also of interest to take into consideration the statistics of the messages that are being sent. An actual source (say English text, or its Morse encoding) has its symbols appear with certain frequencies, as do its 2-blocks, 3-blocks, etc. These statistics might be badly matched to the properties of the signal generator. For example, an infrequent letter, like q, might appear as the label of a very popular edge in G, one which forms a sort of hub, in that many edges lead to and from it and so it appears often in many words of $\mathcal{L}(X_A)$. So it seems worthwhile to look for some intrinsic statistics for the signal generator, a sort of guide on *how to operate* the device, that is to say a measure on (X_A, σ) that is somehow optimal; then we could try to recode sources to match these optimal statistics as nearly as possible before trying to make the signal generator produce them. An optimal measure μ on (X_A, σ) is one that has *maximal entropy*:

$$h_\mu(X_A, \sigma) = h_{\text{top}}(X_A, \sigma) = \log \lambda_A.$$

If the signal generator produces messages whose statistics are described by such a measure μ, then the information content per symbol transmitted (or per unit time) is as high as it could possibly be, namely the "capacity" of the device.

To help make sense out of this, we recall the definition of measure-theoretic entropy very briefly—for a more thorough account, see Petersen's book. If α denotes the *time-0 partition* of X_A according to the central symbol, then

$$\alpha_0^{n-1} = \alpha \vee \sigma^{-1}\alpha \vee \cdots \vee \sigma^{-n+1}\alpha$$

is the partition of X_A according to initial n-blocks. Note that this corresponds to the time-0 partition, according to 1-blocks, of the higher block representation $(X_A^{[n]}, \sigma)$. Our (expected) uncertainty of what n-block is to be transmitted is then

$$H_\mu(\alpha_0^{n-1}) = -\sum_B \mu[B]_0 \log \mu[B]_0,$$

the sum being taken over all n-blocks B, i. e. over all cells of the partition α_0^{n-1}; this is the same as the amount of information that is conveyed on average when the initial n-block is received. The uncertainty being removed (the same as the information being conveyed) *per symbol transmitted* (or per unit time) is

$$h_\mu(X_A, \sigma) = \lim_{n\to\infty} \frac{1}{n} H(\alpha_0^{n-1}).$$

Note: The *same* partition α works for all measures μ because the time-0 partition of a shift space is a *universal generator*.

Now our uncertainty about what n-block is to be transmitted cannot be more than if μ were to assign equal probabilities $(1/N_n(X_A))$ to all of the possible n-blocks (this is not only clear intuitively but can be checked by Lagrange multipliers); thus

$$H_\mu(\alpha_0^{n-1}) \leq \sum_{n\text{-blocks } B} \frac{1}{N_n(X_A)} \log\left(\frac{1}{N_n(X_A)}\right) = \log\left(\frac{1}{N_n(X_A)}\right),$$

and hence such a μ would have entropy equal to $h_{\text{top}}(X_A, \sigma)$. Of course it's only in the case of full shifts that we can hope to find measures that actually distribute mass equally among all allowable n-blocks. But it was proved first by Shannon, in a restricted form, and then by Parry in general that on every irreducible subshift of finite type there is such an optimal measure, that it is Markov, and that it is unique. We shift ground slightly to *vertex subshifts* so as to make it easier to describe the Markov measure.

Theorem 22.1. *Let (X_A, σ) be an irreducible vertex subshift of finite type determined by a $d \times d$ $0, 1$ matrix A. Then there is a unique shift-invariant Borel probability measure μ_{SP}, called the* Shannon-Parry measure*, on X_A which has maximum entropy:*

$$h_\mu(X_A, \sigma) = h_{top}(X_A, \sigma) = \log \lambda_A.$$

It is a 1-step Markov measure with fixed probability vector (initial distribution) p given by

$$p_i = l_i r_i, i = 0, \dots, d-1,$$

where l and r are the left and right positive eigenvectors of A corresponding to the eigenvalue λ_A with $\sum l_i r_i = 1$, and stochastic transition matrix P given by

$$P_{ij} = \frac{A_{ij} r_j}{\lambda_A r_i} \quad \text{for } i, j = 0, 1, \dots, d-1.$$

Proof. We have seen above that for any shift-invariant (Borel probability) measure μ on X_A,

$$h_\mu(X_A, \sigma) \le h_{\text{top}}(X_A, \sigma) = \log \lambda_A.$$

So let us proceed to check that the 1-step Markov measure defined in the statement of the Theorem is shift-invariant and achieves this maximum possible entropy. For shift-invariance, we need that $pP = p$. But

$$\begin{aligned}(pP)_j &= \sum_i p_i P_{ij} = \sum_i l_i r_i \frac{A_{ij} r_j}{\lambda_A r_i} \\ &= \frac{r_j}{\lambda_A} \sum_i l_i A_{ij} = \frac{r_j}{\lambda_A} (lA)_j \\ &= \frac{r_j}{\lambda_A} (\lambda_A l)_j = r_j l_j = p_j.\end{aligned}$$

Recall (see Petersen's book, around 5.2.12) that for any shift-invariant measure μ,

$$h_\mu(\alpha, \sigma) = H_\mu(\alpha | \sigma^{-1}\alpha \vee \sigma^{-2}\alpha \vee \dots) = H_\mu(\alpha | \alpha_1^\infty),$$

and for a 1-step Markov μ with fixed vector p and transition matrix P this equals

$$H_\mu(\alpha | \sigma^{-1}\alpha) = -\sum_{i,j} p_i P_{ij} \log P_{ij}.$$

Rather than just computing this out, let's see why it is that the Shannon-Parry measure μ_{SP} gives fairly equal measures (one of a choice of finitely many constant multiples of λ_A^{-n}) to all the allowable n-blocks: for an allowable cylinder $i_0 i_1 \dots i_n$ in X_A,

$$\begin{aligned}\mu_{\text{SP}}[i_0 i_i \dots i_n]_0 &= p_{i_0} P_{i_0 i_1} P_{i_1 i_2} \dots P_{i_{n-1} i_n} \\ &= (l_{i_0} r_{i_0}) \frac{A_{i_0 i_1} r_{i_1}}{\lambda_A r_{i_0}} \frac{A_{i_1 i_2} r_{i_2}}{\lambda_A r_{i_1}} \dots \frac{A_{i_{n-1} i_n} r_{i_n}}{\lambda_A r_{i_{n-1}}} \\ &= l_{i_0} r_{i_n} \lambda_A^{-n}.\end{aligned}$$

¿From this it follows immediately that

$$\begin{aligned} h_\mu(X_A,\sigma) &= \lim_{n\to\infty} \frac{1}{n} H_\mu(\alpha_0^{n-1}) \\ &= -\lim_{n\to\infty} \frac{1}{n+1} \sum_{i_0,\dots,i_n} l_{i_0} r_{i_n} \lambda_A^{-n} \log(l_{i_0} r_{i_n} \lambda_A^{-n}) \\ &= -\lim_{n\to\infty} \left[\frac{1}{n+1} \sum_{i_0,\dots,i_n} l_{i_0} r_{i_n} \lambda_A^{-n} \log(l_{i_0} r_{i_n}) - \frac{n}{n+1} \sum_{i_0,\dots,i_n} l_{i_0} r_{i_n} \lambda_A^{-n} \log(\lambda_A) \right] \\ &= \log \lambda_A, \end{aligned}$$

since

$$\sum_{i_0,\dots,i_n} l_{i_0} r_{i_n} \lambda_A^{-n} \log(l_{i_0} r_{i_n})$$

is bounded and

$$\sum_{i_0,\dots,i_n} l_{i_0} r_{i_n} \lambda_A^{-n} = 1.$$

23. April 2 *(Notes by DD)*

Continuing the proof of Theorem 22.1, it remains to show that the Shannon-Parry measure, defined last time as a 1-step Markov measure on the (topologically transitive) vertex subshift of finite type $X_A, \sigma)$, is the *only* (shift-invariant Borel probability) measure on (X_A, σ) with entropy $\log \lambda_A$.

Suppose that μ is an invariant measure on (X_A, σ) with entropy $\log \lambda_A$. We show first that μ must be a 1-step Markov measure. This is accomplished by considering the 1-step *Markovization* μ_1 of μ, which is defined to be the (unique) 1-step Markov measure on (X_A, σ) which agrees with μ on all cylinder sets defined by blocks of length 2. Thus μ_1 has fixed probability vector

$$q = (\mu([0]_0), \mu([1]_0), \dots, \mu([d-1]_0))$$

and stochastic matrix of transition probabilities Q defined by

$$Q_{ij} = \mu(j|i) = \frac{\mu([ij]_0)}{\mu([i]_0)} = \frac{\mu(A_i \cap \sigma^{-1} A_j)}{\mu(A_i)},$$

where $A_i = \{x \in X_A : x_0 = i\}$ is the i'th cell of the time-0 partition α for $i, j = 0, 1, \dots, d-1$. The important point is that forming the Markovization can only make the entropy go up: by 5.2.5 (2) in Petersen's book,

$$\begin{aligned} \log \lambda_A = h_\mu(X_A, \sigma) = H_\mu(\alpha | \sigma^{-1}\alpha \vee \sigma^{-2}\alpha \vee \dots) = H_\mu(\alpha|\alpha_1^\infty) \\ \le H_\mu(\alpha|\sigma^{-1}\alpha) = H_{\mu_1}(\alpha|\sigma^{-1}\alpha) = h_{\mu_1}(X_A, \sigma) \le \log \lambda_A, \end{aligned}$$

since μ and μ_1 agree on 2-blocks, μ_1 is 1-step Markov, and we know that no measure can have entropy larger than $\log \lambda_A$.

It follows that all expressions in the above chain are equal, in particular

$$H_\mu(\alpha|\alpha_1^\infty) = H_\mu(\alpha|\sigma^{-1}\alpha).$$

This implies that μ is 1-step Markov. For equality can hold here if and only if, given $\sigma^{-1}\alpha$, α is independent of $\sigma^{-2}\alpha \vee \dots$, which says exactly that μ is 1-step Markov: given the present (here time -1), the future (time 0) is independent of the past (times $-2, -3, \dots$).

(This conditional independence is a consequence of the convexity of the function $f(t) = -t \log t$ and Jensen's Inequality—see Petersen's book, 5.2.9, and Smorodinsky's lecture notes, 4.22. Recall that two partitions ξ and η are *independent given a partition* α if on each cell C of α the restricted partitions $\xi|C$ and $\eta|C$ are independent: for each $N \in \eta$ and $Z \in \xi$, $\mu_C(N \cap Z) = \mu_C(N)\mu_C(Z)$, i. e. , $\mu(N \cap Z \cap C)/\mu(C) = [\mu(N \cap C)/\mu(C)][\mu(Z \cap C)/\mu(C)]$. One can now calculate in $\mathbb{R}^d$ to verify that $\mu = \mu_{\text{SP}}$, but Parry's argument using Markovization, entropy, and ergodicity is actually easier as well as more instructive.)

Let μ be any maximal measure for (X_A, σ), i. e. , $h(\mu) = \log \lambda_A$. We know now that μ is 1-step Markov. We also know that the Shannon-Parry measure μ_{SP} is 1-step Markov and has entropy $\log \lambda_A$. Further, μ_{SP} is ergodic, since its fixed vector p is positive and its matrix P of transition probabilities has positive entries wherever A does, and A is irreducible, hence so is P. Form the measure

$$\nu = \frac{1}{2}(\mu + \mu_{\text{SP}}).$$

Because entropy is an affine function of the measure, we have that also

$$h(\nu) = \log \lambda_A.$$

By the preceding paragraph, this implies that also ν is 1-step Markov. Further, ν is ergodic: its matrix of transition probabilities is also irreducible, since ν assigns positive measure to every 2-block to which μ_{SP} does. Since ergodic measures are extreme points of the set of invariant probability measures, we must have $\nu = \mu = \mu_{\text{SP}}$. □

Remark 23.1 (Terminology)**.** A topological dynamical system (X, T) which has a unique measure of maximal entropy is called *intrinsically ergodic*. Thus we have proved that irreducible SFT's are intrinsically ergodic.

Example 23.1. The golden-mean SFT is the vertex shift determined by the matrix

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

Every $p \in (0, 1)$ determines a 1-step Markov measure with matrix of transition probabilities

$$\begin{pmatrix} p & 1-p \\ 1 & 0 \end{pmatrix};$$

but there is only one p that gives the measure of maximum entropy. Putting $\det(I - \lambda A) = 0$, we find the Perron eigenvalue λ_A =the golden mean= $\gamma = (1 + \sqrt{5})/2$. The left and right eigenvectors, fixed vector, and matrix of transition probabilities are

$$l = \left(\frac{1}{\gamma}, \frac{1}{\gamma^2}\right), r = \begin{pmatrix} 1/\gamma \\ 1/\gamma^2 \end{pmatrix}, \tag{9}$$

$$p = (l_i r_i) = \frac{(1/\gamma^2, 1/\gamma^4)}{\sqrt{1/\gamma^4 + 1/\gamma^8}} = \left(\frac{\gamma^2}{1+\gamma^2}, \frac{1}{1+\gamma^2}\right), \tag{10}$$

$$P = \begin{pmatrix} 1/\gamma & 1/\gamma^2 \\ 1 & 0 \end{pmatrix}. \tag{11}$$

24. April 7 *(Notes by LK)*

Last week we showed that an irreducible shift of finite type (X_A, σ) has a unique measure μ_{SP} of maximal entropy; i.e.

$$h(\mu_{SP}) = h_{top}(X_A, \sigma) = \log \lambda_A.$$

We discussed the example of the golden mean subshift of finite type.

24.1. **Generalization to Equilibrium States.** Let (X, T) be a topological dynamical system and $\phi\colon X \to \mathbb{R}$ a continuous function (a "potential function"). An *equilibrium state* for ϕ is any T-invariant Borel probability measure on X which achieves the supremum over all T-invariant Borel probability measures μ of

$$h(\mu) + \int_X \phi d\mu.$$

Thus a measure of maximal entropy is an equilibrium state for the function $\phi = 0$.

We are interested in determining when equilibrium states exist and, if so, when they will be unique. There are various theorems guaranteeing the existence of an equilibrium state when ϕ is "very continuous" in some sense, i.e. is Hölder continuous, or has summable variation. If more than one equilibrium state exists for a function ϕ then we say that there is a *phase transition*. An example of a physical model with a phase transition would be a state where ice and cold water exist simultaneously.

We have already shown that an irreducible subshift of finite type has a unique equilibrium state for the function $\phi = 0$. More generally, if (X_A, σ) is an irreducible subshift of finite type and $\phi\colon X \to \mathbb{R}$ depends on finitely many coordinates (and without loss of generality on two coordinates since we can use a higher block code representation), then there is a unique equilibrium state for ϕ. This measure is one-step Markov and can be found in the same way as the Shannon-Parry measure, except that instead of working with the matrix A we work with $A(\phi)$ where

$$(A(\phi))_{ij} = A_{ij} e^{\phi(ij)}.$$

(Recall that the rows and columns of the $d \times d$ matrix A are indexed by the symbols in the alphabet). From this we obtain the left and right eigenvectors l and r:

$$lA(\phi) = \lambda_A l,$$

$$A(\phi) r = r \lambda_{A(\phi)},$$

$$l \cdot r = 1,$$

and use them to form the probability vector p,

$$p_i = l_i r_i,$$

and the matrix of transition probabilities

$$P(\phi)_{ij} = (A(\phi))_{ij} \frac{r_j}{\lambda_{A(\phi)} r_j}.$$

See page 22 in Parry and Tuncel for the details of this argument or Theorems 39.4 (page 80), 39.3 (page 79), and 39.1 (page 78) in the Spring 1997 261 notes for a more general discussion.

24.2. **Coding between Subshifts.** Let $(X, \sigma) \subset D^{\mathbb{Z}}$ be any subshift where $D = \{0, 1, \dots, d-1\}$, and let $a \in \mathbb{N}$. Recall the higher block representation (a-block) of (X, σ). We take a new alphabet D^a of strings of length a. Define a map

$$c_a \colon X \to (D^a)^{\mathbb{Z}}$$

by

$$(c_a x)_i = x_i x_{i+1} \dots x_{i+a-1}.$$

This map is a continuous, shift-commuting map and gives a topological conjugacy between

$$(X, \sigma) \approx (c_a, \sigma).$$

We note that Lind and Marcus use the notation $(X^{[a]}, \sigma)$ for (c_a, σ). The map c_a is a sliding block code with memory zero and anticipation one. The map ${c_a}^{-1}$ is a one-block code with memory zero and anticipation zero. That is, if $y_0 = b_1 \dots b_{a-1}$ then

$${c_a}^{-1}(y)_0 = b_1.$$

The inverse map ${c_a}^{-1}$ compresses information by taking the first symbol and forgetting the rest.

Recall that if $(X_{\mathcal{F}}, \sigma)$ is an a-step subshift of finite type then $(X^{[a]}, \sigma)$ is a one-step subshift of finite type called the *vertex shift* of $(X_{\mathcal{F}}, \sigma)$, and $(X^{[a+1]}, \sigma)$ is the *edge shift.* If α is the time-0 partition of x, then $\alpha \vee \sigma^{-1}\alpha \vee \dots \sigma^{-n+1}\alpha$ corresponds to the time-0 partition of $X^{[a]}$. Thus higher block codes could let us concentrate on time-0 partitions rather than successive refinements.

24.3. **State Splitting and Amalgamation.** Here we begin to follow closely Lind and Marcus. Let G be an irreducible directed graph, perhaps with multiple edges. Then G has an associated edge shift

$$(X_G, \sigma) \subset \mathcal{E}(G)^{\mathbb{Z}},$$

where the symbols in $\mathcal{E}(G)$ are the edges of G and $\mathcal{E}(G)^{\mathbb{Z}}$ is all strings determined by infinite walks on the graph G. Let $A(G)$ be the nonnegative integer matrix which is indexed by $\mathcal{V}(G)$, the vertices of the graph, and for which $A(G)_{ij}$ is the number of edges from i to j.

Next we describe how to take an elementary *out-splitting* of one vertex $v \in \mathcal{V}(G)$, which has no loops, into two new vertices. First, we need to partition the outgoing edges from v. Let $\mathcal{E}_v = \{\text{edges starting at } v\}$. We will illustrate splitting vertex v on the following graph G.

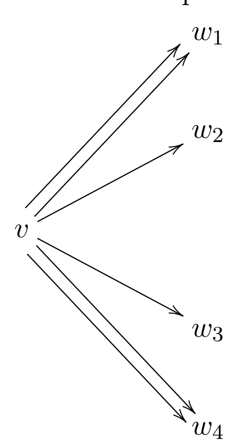

G

Suppose we have a partition $\mathcal{E}_v = \mathcal{E}_v^1 \cup \mathcal{E}_v^2$ where $\mathcal{E}_v^1$ contains the top three edges in the graph of G and $\mathcal{E}_v^2$ contains the bottom three edges. We split the state v into two new vertices v_1 and v_2, where v_1 gets all the outgoing edges from v that are in $\mathcal{E}_v^1$ and v_2 gets all the outgoing edges in $\mathcal{E}_v^2$. Thus the graph of G with v split into v_1 and v_2 can be illustrated as follows:

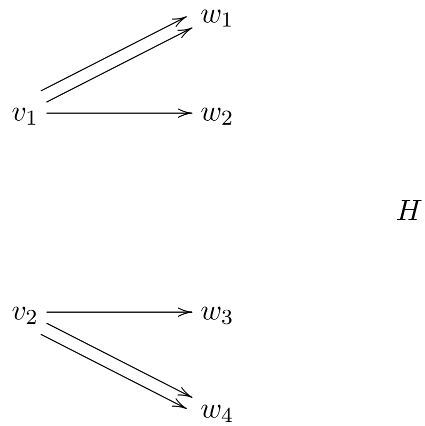

The edges that are coming into v get cloned. That is, if an edge exists from w to v we create a new edge so that we have an edge from w to v_1 and an edge from w to v_2 as follows.

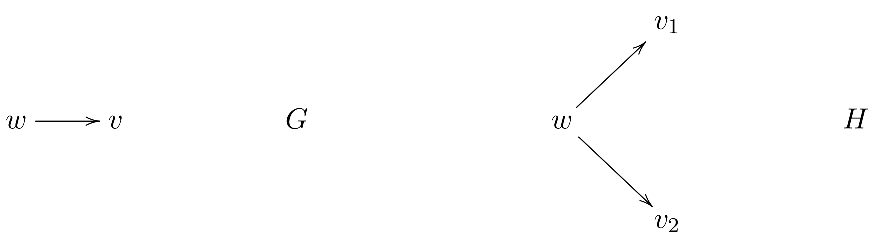

Thus in performing the out-splitting of a vertex we partition the outgoing edges and clone the incoming edges to obtain a new graph H. We also allow H to be replaced by any graph isomorphic to it.

Theorem 24.1. *The edge shifts X_G and X_H are topologically conjugate.*

Proof. We select the particular graph in the isomorphism class of H that is constructed as above and define $\pi\colon X_H \to X_G$ by identifying edges that are clones of each other. That is, if the edges of G are denoted by $\{e, f, \ldots\}$ and the edge e is split into e_1, e_2 in H, and the edge f is split into f_1, f_2 etc. then

$$\pi(e_1) = \pi(e_2) = e, \quad \pi(f_1) = \pi(f_2) = f, \text{ etc.}$$

The map π is one-to-one because we can recover the subscript of an edge (π deletes subscripts) by looking at the next edge and seeing what cell of the partition it lies in.

The map $\pi^{-1}\colon X_G \to X_H$ is a two-block code with memory zero and anticipation one. For example, if $f \in \mathcal{E}_v^k$ and

$$x = \ldots ef \ldots \in X_G,$$

then $\pi^{-1}x = \ldots e_k \ldots$.

□

Next, we discuss state splitting in general, where we *out-split* any number of vertices in a graph which may have loops. As before, our procedure will be to partition outgoing edges and clone incoming edges.

Let G be a graph and (X_G, σ) be its edge shift. For each vertex $v \in \mathcal{V}(G)$ we partition $\mathcal{E}_v = \{\text{edges starting at } v\}$ into

$$\mathcal{E}_v = \mathcal{E}_v^1 \cup \mathcal{E}_v^2 \cup \ldots \cup \mathcal{E}_v^{n(v)}.$$

We denote this partition by $\mathcal{P} = \{\mathcal{E}_v^j : 1 \leq j \leq n(v)\}$. The elements of $\mathcal{P}$ are the vertices of the new graph. That is, $H = H(G, \mathcal{P})$ has vertices which are cells of the partition denoted by $\mathcal{E}_v^j$ or v^j for $1 \leq j \leq n(v)$. We define the edges of H as follows. If an edge connects vertex v to w in G, then find the i such that $e \in \mathcal{E}_v^i$ and put an edge from v_i to each w_j for $1 \leq j \leq n(w)$. Thus each edge is cloned as before according to the number of elements its terminal vertex is split into.

We illustrate this procedure on the following graph G, using the complete partition of $\mathcal{E}(G)$ into singletons.

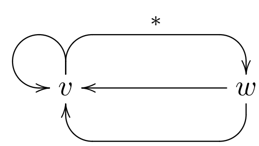

G

We obtain the following graph H. We describe what happens to the edge labeled with a star (*). That edge can be thought of as the part of v that leaves v. We'll assume that it lies in the partition element $\mathcal{E}_v^1$ so in the graph H it will come from the vertex v_1. We must clone it so that it connects the vertex v_1 to the vertices w_1 and w_2 which are created in the graph of H since w has two outgoing edges each belonging to their own atom in the partition of the edges of G.

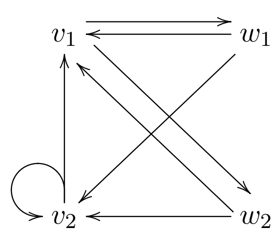

H

We can perform an *in-splitting* of a graph as well. This procedure is similar to out-splitting except that we think of the partition $\mathcal{P}$ of $\mathcal{E}(G)$ as partitioning the edges coming *into* each vertex v. We then clone edges according to which pieces of each (new) vertex they start at.

For example, if we use the golden mean shift with the complete partition,

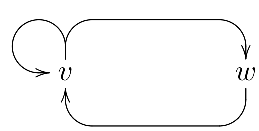

G

we obtain the in-splitting $H = H(G, \mathcal{P})$

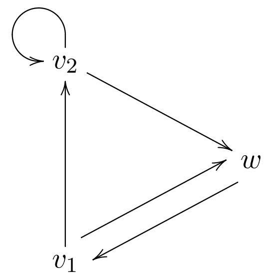

H

If H is an in- or out-splitting of G, then we say that G is an (in- or out-) *amalgamation* of H. Splittings and amalgamations of graphs give rise to topological conjugacies of their edge shifts; these are called *splitting codes* and *amalgamation codes*.

25. April 9 *(Notes by LK)*

Recall that we perform an out-splitting by using a partition of the edge set $\mathcal{E}(G)$. Each vertex $v \in \mathcal{V}(G)$ splits according to the partition of edges leaving v. Each edge into $w \in \mathcal{V}(G)$ is cloned according to fragments of w. The out-splitting is a two-block map with memory zero and anticipation one, and the amalgamation code is a one-block map with memory and anticipation zero.

We want to determine how the matrices for G and H are related. Recall that A_G is the $d \times d$ matrix where

$$(A_G)_{vv'} = \text{ number of edges in } G \text{ from } v \in \mathcal{V}(G) \text{ to } v' \in \mathcal{V}(G).$$

Suppose we have an out-splitting from G to H. We define two useful matrices. The first is the *division matrix* D, which is a $0,1$ matrix determined by a partition $\mathcal{P}$ of $\mathcal{E}(G)$. This is a $|\mathcal{V}(G)| \times |\mathcal{V}(H)|$ matrix, where the entry D_{vu} is

$$\begin{cases} 1 \text{ if } u \text{ is a vertex in } H \text{ that is created by splitting a vertex } v \text{ in } G, \\ 0 \text{ otherwise.} \end{cases}$$

Thus each row of D has at least one 1, and each column of D has exactly one 1.

The second matrix of interest is the nonnegative integer $|\mathcal{V}(H)| \times |\mathcal{V}(G)|$ *edge matrix* E. We let

$$E_{uv} = \text{ the number of edges in } G \text{ that end at } v \text{ and are in the partition element } u \text{ of } \mathcal{P}.$$

The edge and division matrices will give us a good sense of how the matrices A_G and A_H change under splitting or amalgamation.

Theorem 25.1. *If H is an out-splitting of G, and D and E are the corresponding division and edge matrices, then $DE = A_G$ and $ED = A_H$.*

We note that the converse holds as well.

Proof. Let $v, v' \in \mathcal{V}(G)$. Then

$$(DE)_{vv'} = \sum_{u \in \mathcal{V}(H)} D_{vu} E_{uv'}.$$

However, $D_{vu} = 1$ if and only if u is a fragment of v in H. Therefore,

$$(DE)_{vv'} = \sum_{u \textit{ fragment of } v} E_{uv'},$$

where $E_{uv'}$ is the number of edges in G to v' which are in the partition element u (a fragment of v). Thus this sum is equal to the number of edges from v to v' in G, which by definition is $(A_G)_{vv'}$.

If $u, u' \in \mathcal{V}(H)$, then

$$(ED)_{uu'} = \sum_{v \in \mathcal{V}(G)} E_{uv} D_{vu'}.$$

Recall that $D_{vu'}$ is 1 if and only if u' is a fragment of v. Therefore, this sum has only one nonzero term E_{uv}, where u' is a fragment of v. However, E_{uv} is the number of edges in G that are in cell u and end at v. Each edge is cloned according to the fragments of v. That is, if there are n edges in the partition element u, then we put n edges into each u' from u. Therefore, E_{uv} is the number of edges in H from u to u', which by definition is $(A_H)_{uu'}$. $\square$

25.1. **Matrices for In-splittings.** Suppose we want to perform an in-splitting of the graph G. Define the transpose G^{tr} of G by reversing all of the edges of G. Then the adjacency matrix of G^{tr} is the transpose of the adjacency matrix of G:

$$A_{G^{tr}} = (A_G)^{tr}$$

If we perform an out-splitting of G^{tr}, we obtain a new graph denoted H^{tr}, and anything isomorphic to H is an insplitting of G. We can obtain the division and edge matrices D and E for G^{tr}: $DE = (A_G)^{tr}$ and $ED = (A_H)^{tr}$. Then we have $E^{tr}D^{tr} = A_G$ and $D^{tr}E^{tr} = A_H$. Thus we have a theorem similar to Theorem 25.1 for in-splittings.

25.2. **Topological Conjugacy of SFT's and Strong Shift Equivalence.** Suppose we have a finite sequence of splittings and amalgamations taking a graph G and producing end result (isomorphic to) H. Then

(1) The edge subshifts of finite type (X_G, σ) and (X_H, σ) are topologically conjugate.
(2) There exists a sequence of pairs of rectangular non-negative integer matrices $(D_0, E_0), (D_1, E_1), \ldots (D_n, E_n)$ such that

$$A_G = A_0 = D_0E_0, \quad A_1 = E_0D_0,$$
$$A_1 = D_1E_1, \quad A_2 = E_1D_1,$$
$$\ldots$$
$$A_n = D_nE_n, \quad A_{n+1} = E_nD_n = A_H.$$

In each case, D is either an edge or a division matrix depending on whether we are taking an in- or out-splitting or amalgamation. We note that any one of the above rows is called an *elementary equivalence* (e.g. $A_1 \sim A_2$). Elementary equivalence is not transitive – A_G and A_H are related by its transitive closure; we say that the matrices A_G and A_H are *strong shift equivalent.*

Theorem 25.2. *(R. Williams) Every topological conjugacy between two edge shifts X_G and X_H is a composition of finitely many (in- or out-) splittings and (in- or out-) amalgamations. X_G and X_H are topologically conjugate if and only if the matrices A_G and A_H are strong shift equivalent.*

We sketch the proof. The details can be found in Lind and Marcus. We recall that the above equivalence relation allows graph isomorphisms.

Proof. Suppose $\phi\colon X_G \to X_H$ is a topological conjugacy. For $x \in X_G$, we know by the Curtis-Hedlund-Lyndon Theorem (Theorem 3.2) that $(\phi x)_0$ depends on a finite window $[x_{-w}, \ldots, x_w]$. We can thus take a higher block representation $(X_G^{[2w+1]}, \sigma)$ to replace ϕ by a one-block map. Recall that we have a topological conjugacy between $(X_G^{[2w+1]}, \sigma)$ and (X_G, σ). Therefore, we assume that ϕ is a one-block map. In using this representation, though, we don't see any splittings and amalgamations.

However, we can use the following fact: the two-block representation $(X_G^{[2]}, \sigma)$ is topologically conjugate to the out-splitting determined by the complete partition of $\mathcal{E}(G)$ into singletons. We can use this fact to make a finite sequence of complete out-splittings which is topologically conjugate to $(X_G^{[2w+1]}, \sigma)$.

Instead of proving this fact, we illustrate it on the golden mean example.

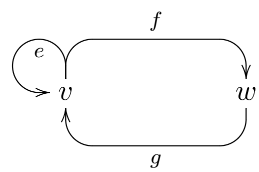

G

The edge shift is all possible sequences of e, f and g on the graph G. The complete out-splitting H is shown next.

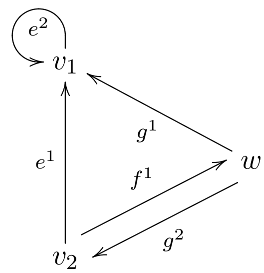

H

We obtain the two-block representation of G as follows:

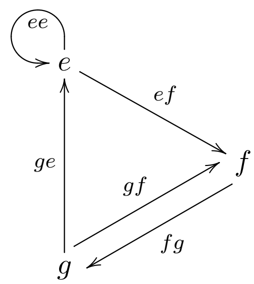

We note that this two-block representation is topologically conjugate to the edge shift of the graph H. $\square$

26. April 14 *(Notes by CN)*

During the previous lecture, we sketched a proof of the theorem of Williams on the topological conjugacy of SFTs and strong shift equivalence. Today we tell how to complete the proof.

Theorem 26.1. *(R. Williams) Let $\phi : X_A \to X_B$ be a topological conjugacy between edge shifts of irreducible nonnegative integer matrices A and B. Then ϕ is a composition of finitely many splitting and amalgamation codes, and the matrices A and B are strong shift equivalent.*

Before proving the theorem, recall that two nonnegative integral matrices are *strong shift equivalent* if there exists a sequence of pairs of nonnegative rectangular integer matrices (D_0, E_0),...,(D_n, E_n) such that

(12)
$$\begin{aligned} A &= D_0E_0, & E_0D_0 &= A_1 \\ A_1 &= D_1E_1, & E_1D_1 &= A_2 \\ &\ \vdots & &\ \vdots \\ A_n &= D_nE_n, & E_nD_n &= B. \end{aligned}$$

(13)

(14)

Ideas of Proof

The complete edge-splitting of a graph corresponds to taking the 2-block code on the edge shift. Repeating this process takes us to a higher block representation of X_A, which corresponds to a sequence of out-splitting codes, we may assume that ϕ is a 1-block code. A 1-block code just relabels the edges.

That ϕ is a 1-block code alone is not sufficient to establish the first result; we must also establish that ϕ^{-1} is a 1-block code. In that case, both ϕ and ϕ^{-1} would be (in- and out-) splitting and (in- and out-) amalgamation codes (which include graph isomorphisms, which simply relabel edges). Let ϕ^{-1} have memory m and anticipation a:

(15)
$$(\phi^{-1}y)_0 = \phi^{-1}(y_{-m}\dots y_{-1}y_0y_1\dots y_a).$$

We will form out-splittings $\tilde{H}$ and $\tilde{G}$ of the graphs G and H (corresponding to edge matrices A and B) Since the out-splittings (let's call them $\psi_{G\tilde{G}}$ and $\psi_{H\tilde{H}}$) are also topological conjugacies, they give rise to a new conjugacy $\tilde{\phi}$:$X_{\tilde{G}} \to X_{\tilde{H}}$. We want to arrange it so that the memory remains m and the anticipation is reduced by to $a - 1$. Then, repeat this process until $a = 0$. Similarly, in-splittings on G and H will give rise to conjugacies with unchanged anticipation and reduced memory. Hence, this process can be used to reduce the memory to $m = 0$. Finally, (after a number of steps depending on a and m) we arrive at a conjugacy ψ which is a 1-block code and whose inverse is also a 1-block code.

Now let $\tilde{H}$ = complete out-splitting of H. (The edge set $\mathcal{E}(H)$ is partitioned into singletons).

Recall: If the graph of H is

$$\underline{H}$$

$$\cdots \qquad v \xrightarrow{e} w \qquad \cdots$$

Then, find the i such that $e \in \mathcal{E}_v^i$, = the i'th partition element of the set of edges starting at v, $i = \{e\}$. We then have

$$\underline{\tilde{H}}$$

$$\cdots \qquad v_i \to w_j \qquad \cdots$$

for each fragment w_j (edge leaving w) of w. There is varied notation for these new edges: call the edge e^j for $v_i \xrightarrow{e^j} w_j$, e^j = (clone of e into fragment) $j = e^{w_j} = e^{\mathcal{E}_w^j} = e^f$ where f is an edge leaving w. (Remember, e is cloned for each edge leaving its destination w.) Now, $\phi : X_G \to X_H$ is a 1-block map, given by a labeling of the edges, so we can assume ϕ comes from a mapping that sends edges of G to edges of H, $\Phi : \mathcal{E}(G) \to \mathcal{E}(H)$ (not necessarily one-to-one).

Now, to form the out-splitting $\tilde{G}$, partition $\mathcal{E}(G)$ according to its elements' images under Φ. (In other words, two elements are in the same partition cell if their images are the same, and we write $\mathcal{E}_G^h = \{g \in \mathcal{E}(G) : \Phi(g) = h\}$.)

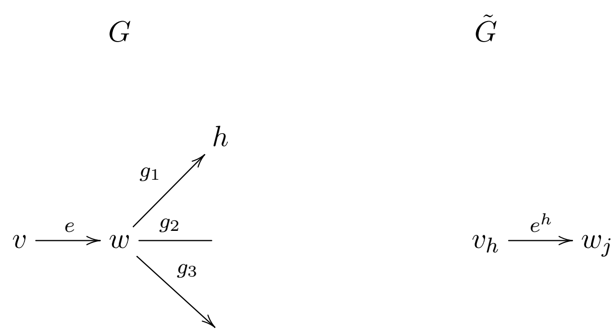

Define $\tilde{\phi}$ on the edges of $\tilde{G}$ by $\tilde{\phi}(e^h) = (\Phi e)^h$. This defines $\tilde{\phi} : \mathcal{E}(\tilde{G}) \to \mathcal{E}(\tilde{H})$ (and hence a 1-block code $X_{\tilde{G}} \to X_{\tilde{H}}$) so that the following diagram commutes:

$$\begin{array}{ccc} X_A = X_G & \underset{\alpha_{\tilde{G}G}}{\overset{\psi_{G\tilde{G}}}{\rightleftarrows}} & X_{\tilde{G}} \\ \phi \downarrow & & \downarrow \tilde{\phi} \\ X_H & \underset{\alpha_{\tilde{H}H}}{\overset{\psi_{H\tilde{H}}}{\rightleftarrows}} & X_{\tilde{H}} \end{array}$$

The actions of these codes work as follows:

$$
\begin{array}{ccc}
X_G & & X_{\tilde{G}} \\
\\
\dots g_{-1}.g_0 g_1 g_2 \dots & \xrightarrow{\psi_{G\tilde{G}}} & \dots g_{-1}^{h_0}.g_0^{h_1} g_1^{h_2} g_2^{h_3} \dots \\
\Big\downarrow \phi & & \Big\downarrow \tilde{\phi} \\
\dots h_{-1}.h_0 h_1 h_2 \dots & \xrightarrow{\psi_{H\tilde{H}}} & \dots h_{-1}^{h_0}.h_0^{h_1} h_1^{h_2} h_2^{h_3} \dots \\
\\
X_H & & X_{\tilde{H}}
\end{array}
$$

We claim that memory hasn't increased, but anticipation has decreased, i.e. $m(\tilde{\phi}^{-1}) = m(\phi^{-1})$, $a(\tilde{\phi}^{-1}) = a(\phi^{-1}) - 1$. For the terms of elements of $X_{\tilde{H}}$ are of the form ${h_0}^{h_1}$ and, thus, the i'th position possesses foreknowledge of the $(i+1)$'st position, one step into the future. For example, if $a = 1$ to determine g_0 via ϕ^{-1}, one must know $\underbrace{\dots h_{-1}.h_0 h_1}_{\text{from } -m \text{ to } 1} h_2 \dots$, but to determine ${g_o}^{h_1}$ via $\tilde{\phi}^{-1}$, one need only know $\underbrace{\dots h_{-1}^{h_0} h_0^{h_1}}_{\text{from } -m \text{ to } 0} h_1^{h_2} h_2^{h_3} \dots$. So if the anticipation $a(\phi^{-1}) = 1$, then $a(\tilde{\phi}^{-1}) = 0$.

26.1. **Shift Equivalence.** We have established strong shift equivalence (SSE) as a necessary and sufficient condition for topological conjugacy between edge shifts. Unfortunately, strong shift equivalence is difficult to check. There is no algorithm for determining strong shift equivalence, and the problem may, in fact, be logically undecidable. It is presently unknown whether it is decidable even for 2×2 matrices. So, mindful of the difficulty of determining strong shift equivalence, R. Williams defined *shift equivalence* (SE) which is decidable and (in theory) can be checked algebraically:

Definition 26.1. Two square nonnegative integral matrices A, B are said to be *shift equivalent* if there exist rectangular nonnegative integral matrices R, S, and an integer l such that

$$AR = RB, SA = BS, \tag{16}$$

$$A^l = RS, B^l = SR. \tag{17}$$

This relation is transitive, and hence is an equivalence relation.

Exercise 2. Prove that the relation is transitive, and, hence, an equivalence relation (Symmetry and reflexivity are clear).

Proposition 26.2. *If A and B are strong shift equivalent (SSE), then they are shift equivalent (SE).*

Proof. Recall the equations from 13. Each row of the equation array, with $A_i = D_iE_i$ and $E_iD_i = A_{i+1}$ represents a shift equivalence, because if $R = D_i$, $S = E_i$, and $l = 1$, then

$$\begin{aligned} A_iR = A_iD_i = D_iE_iD_i = D_iA_{i+1} = RA_{i+1} \\ SA_i = E_iA_i = E_iD_iE_i = A_{i+1}E_i = A_{i+1}S \\ A_i^1 = A_i = D_iE_i = RS \\ A_{i+1}^1 = A_{i+1} = E_iD_i = SR. \end{aligned}$$

□

Since shift equivalence is transitive, this implies that $A \overset{\text{SE}}{\sim} B$.

26.2. **Williams Conjecture.** Shift equivalence was put forth by Williams as a simpler relation which could be checked. It was his hope and belief that this relation was a sufficient condition for conjugacy, and he set out to prove it.
R. Williams Conjecture, a. k. a. the Shift Equivalence Problem:

$$\text{If } A \overset{\text{SE}}{\sim} B, \text{ is } A \overset{\text{SSE}}{\sim} B\ ??$$

26.2.1. *Some Positive Evidence:* Many people tried for a long time without success to find counterexamples. Shift equivalence (SE) also preserves all known (up to 1997) topological conjugacy invariants, including (in particular) zeta functions, and hence topological entropy, the number of periodic points of each period, etc.

If (X,T) is a topological dynamical system, for each $n = 1, 2, \dots$ let $p_n = \text{card}\{x \in X : T^n x = x\}$ and q_n = number of points with least period n, and

$$\zeta(t) = \exp(\sum_{n=1}^{\infty} \frac{p_n}{n} t^n). \tag{18}$$

The function $\zeta(t)$ is an invariant of topological conjugacy as $\{p_n\}$ is.

Exercise 3. For an edge shift (X_A, σ), show that

$$\zeta_A(t) = \frac{1}{\det(1 - tA)} \tag{19}$$

where the denominator above is the characteristic polynomial of A.

In short, it is hard to find a way to tell shift equivalent edge shifts apart! In 1992, however, K. H.

Kim and F. W. Roush found a counterexample for the case of *reducible matrices. And, in 1997, Kim and Roush were able to generate an* irreducible counterexample. The 7×7 counterexample matrices for A, B, S, and R were found using a computerized algebra system and are included in Kim and Roush's paper, *The Williams Conjecture is False for Irreducible Subshifts* which can be found on the internet in *Electronic Research Announcements.* Their achievements were based largely on work by J. Wagoner, M. Boyle, W. Krieger, U. Fiebig, and others.

26.3. **Invariants of Shift Equivalence.** There is a *complete invariant* for shift equivalence, namely the *dimension triple* $(\Delta_A, \Delta_A^+, \delta_A)$ developed by Krieger, Williams, Kim and Roush, Boyle, Handleman, and Marcus, Boyle and Trow, based on ideas from C^*-algebras.

The elements of the dimension triple are as follows:

$A = d \times d$ nonnegative integer matrix,

$\Delta_A =$ *dimension group*,

$\Delta_A^+ =$ *dimension semigroup*,

$\delta_A : \Delta_A \to \Delta_A$ is the *dimension group automorphism.* These things will be defined during the next lecture.

27. April 16 *(Notes by CN)*

Recall that in the previous lecture we defined strong shift equivalence (SSE) and shift equivalence (SE) and proved that strong shift equivalence implied Shift Equivalence 26.2. R. Williams conjectured that the converse was true, but Kim and Roush uncovered a counterexample in their 1997 paper. We neglected to mention, however, that their counterexample matrices, which include negative entries in some places, are acceptable only because of the following:

Proposition 27.1. *If A and B are primitive (aperiodic) matrices, then they are shift equivalent over $\mathbb{Z}^+$ if and only if they are shift equivalent over $\mathbb{Z}$.*

The Kim-Roush example does involve primitive matrices A and B.

27.1. **Invariants of Shift Equivalence.** In the previous lecture, we introduced the dimension triple, a complete invariant for shift equivalence. Now, we can provide definitions.

Definition 27.1. Let A be a $d \times d$ nonnegative integral matrix. The *eventual range of* A, $\Re(A)$, is given by

$$\Re(A) = \bigcap_{k=1}^{\infty} \mathbb{Q}^d A^k, \tag{20}$$

an A-invariant subspace of $\mathbb{Q}^d$, where $\mathbb{Q}^d$ is the set of d-dimensional rationals, i.e. row vectors with rational entries. (Note that here A acts on the right.) The *dimension group*, Δ_A, is given by

$$\Delta_A = \{v \in \Re(A) : \text{there exists } k \geq 0 \text{ such that } vA^k \in \mathbb{Z}^d\}, \tag{21}$$

an additive subgroup of $\Re(A)$. The *dimension group automorphism*, $\delta_A : \Delta_A \to \Delta_A$, is defined by $\delta_A = A \mid \Delta_A$, the restriction of A to Δ_A. The *dimension semigroup*, Δ_A^+, is given by

$\Delta_A^+ = \{ v \in \Re(A)$: there exists $k \geq 0$ such that $vA^k \in (\mathbb{Z}^+)^d\}$.

Remarks 27.1.
The dimension group automorphism, δ_A, preserves the *dimension semigroup*, Δ_A^+.
The complete invariant of shift equivalence over $\mathbb{Z}^+$ is the *dimension triple*, $(\Delta_A, \Delta_A^+, \delta_A)$.
The complete invariant of shift equivalence over $\mathbb{Z}$ is the *dimension pair*, (Δ_A, δ_A).

Theorem 27.2. *For nonnegative integral matrices A and B, the following are equivalent:*

1. *A and B are shift-equivalent (over $\mathbb{Z}^+$).*
2. *There is a group isomorphism $\Delta_A \to \Delta_B$ which maps Δ_A^+ onto Δ_B^+ and commutes with the actions of the two automorphisms δ_A and δ_B.*
3. *The corresponding edge shifts (X_A, σ) and (X_B, σ) are* eventually conjugate *(though not necessarily* conjugate*).*

This theorem, developed by Krieger, R. Williams, Kim and Roush, Boyle and Handleman, and Boyle, Marcus, and Trow, requires the following:

Definition 27.2. Two edge shifts, (X_A, σ) and (X_B, σ), are *eventually conjugate* if, for all large enough n, (X_A^n, σ^n) and (X_B^n, σ^n) are topologically conjugate, where (X_A^n, σ^n) has alphabet consisting of n-blocks in X_A and the shift is by n places each time. (The map $X_A \to X_B$ is a block map that comes from grouping into n-blocks, not a sliding block code).

Example 27.1. Consider the full shift (Σ_2, σ). It has 2-block representation $(\Sigma_2^{[2]}, \sigma)$ over an alphabet of $\underbrace{00}_{a}, \underbrace{01}_{b}, \underbrace{10}_{c}, \underbrace{11}_{d}$. Sequences are recoded as follows:

$$\begin{array}{rcl} \ldots.011010001 & \stackrel{\sigma}{\to} & \ldots.1101000 \\ \ldots.bdcbc & \to & \ldots.dc\ldots \end{array}$$

The map $00 \to a$, etc., is a regular sliding block code. The image is a vertex SFT with graph as follows.

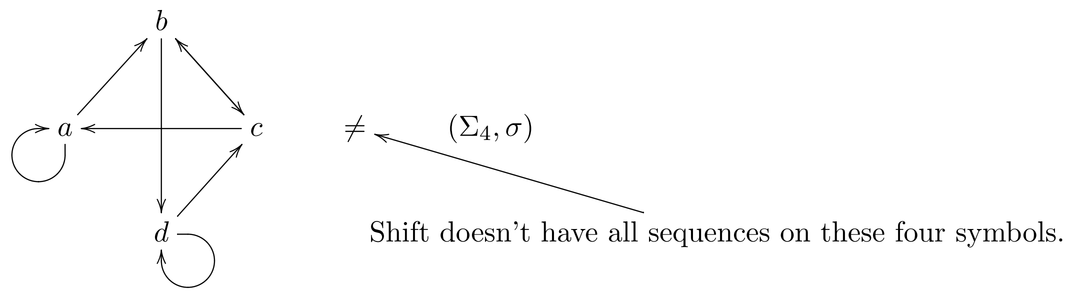

Hence, $(\Sigma_2, \sigma) \stackrel{\text{top.conj.}}{\cong} \Sigma_A \subset \{a, b, c, d\}^{\mathbb{Z}}$, the 4-shift. We have $(\Sigma_2^{[2]}, \sigma) \stackrel{\text{top.conj.}}{\cong} (\Sigma_A, \sigma)$. On the other hand, $(\Sigma_A^2, \sigma^2) \stackrel{\text{top.conj.}}{\cong} (\Sigma_4, \sigma)$. For, consider (Σ_2^2, σ^2), with alphabet $\underbrace{00}_{u}, \underbrace{01}_{v}, \underbrace{10}_{w}, \underbrace{11}_{x}$.

Now, the recoding by grouping works as below:

$$\begin{array}{rcl} \ldots.011010001 & \stackrel{\sigma^2}{\to} & \ldots.1010001\ldots \\ \ldots.vwwu & \to & \ldots.wwu... \end{array}$$

Thus $(\Sigma_2^2, \sigma^2) \cong (\Sigma_4, \sigma)$.

Continuing in the subject of codings between SFTs, we now consider...

27.2. **Embeddings and Factors.** Given two subshifts of finite type, when can you get an embedding? If there is a one-to-one continuous shift-commuting map $\phi : X_A \to X_B$ between irreducible SFTs (X_A, σ) and (X_B, σ) and $\phi(X_A)$ is a proper subset of X_B, then:

(1) $h_{\text{top}}(X_A, \sigma) < h_{\text{top}}(X_B, \sigma)$.
(2) If $q_n(X_A) =$ the number of points in X_A with least period n, we must have $q_n(X_A) \le q_n(X_B)$ for every n. (Or, equivalently, there is a shift-commuting injection $\wp(X_A, \sigma) \hookrightarrow \wp(X_B, \sigma)$, where $\wp(X_A, \sigma)$ is the set of periodic points of (X_A, σ)).

Statement (1) holds because $\phi(X_A)$ is a closed σ-invariant proper subset of X_B and $(\phi(X_A), \sigma) \subset (X_B, \sigma)$. This implies that its complement is nonempty and open. $(\phi(X_A), \sigma)$ is also a SFT, hence it has a unique Shannon-Perry measure, μ, whose entropy $h(\mu) = \log \lambda_A < \log \lambda_B$. (The Shannon-Perry measure on X_B has full support, so it is not equal to μ).

Remark 27.1. Any subshift conjugate to an SFT is an SFT.

Theorem 27.3 (Krieger,1982)**.** *There exists a proper embedding* $\phi : (X_A, \sigma) \to (X_B, \sigma)$ *if and only if*

(1) $h_{\text{top}}(X_A, \sigma) < h_{\text{top}}(X_B, \sigma)$ *and*
(2) *for all* n,$q_n(X_A) \le q_n(X_B)$.

27.2.1. *What about factors?*

Remarks 27.2. If there exists a shift-commuting continuous onto map (i.e., a factor map) $\pi : (X_A, \sigma) \to (X_B, \sigma)$, then

(1) $h_{\text{top}}(X_A, \sigma) \ge h_{\text{top}}(X_B, \sigma)$ (taking a factor map can only cause entropy to decrease), and
(2) If $x \in X_A$ has least period p, then πx has to have period that divides p, since periodic points must be mapped to periodic points, but the actual period may change. If for all $x \in X_A$ with least period p there exists y with least period that divides p, write $\wp(X_A) \searrow \wp(X_B)$. This is equivalent to saying that there exists a shift-commuting map $\wp(X_A, \sigma) \to \wp(X_B, \sigma)$ (not necessarily onto). (Nonperiodic points may be mapped to periodic points).
(3) When factoring is possible in case $h_{\text{top}}(X_A, \sigma) = h_{\text{top}}(X_B, \sigma)$ remains an open question.

Theorem 27.4 (Boyle, 1983)**.** *If* (X_A, σ) *and* (X_B, σ) *are topologically transitive SFTs with* $h_{\text{top}}(X_A, \sigma) > h_{\text{top}}(X_B, \sigma)$, *then there exists a factor map* $(X_A, \sigma) \to (X_B, \sigma)$ *if and only if* $\wp(X_A, \sigma) \searrow \wp(X_B, \sigma)$.

28. Sofic Systems

Definition 28.1. A *sofic system* is a subshift that is a factor of a subshift of finite type.

Equivalently, a sofic system is a subshift consisting of all infinite walks on the edges of a graph whose edges have been labeled (though maybe not in a one-to-one manner). Sofic systems include all SFTs, since every SFT is conjugate to an edge shift ($\subset (\mathcal{E}(G)^{\mathbb{Z}}, \sigma)$).

Consider the following

$$(X_G, \sigma) \subset (\{e_1, e_2, f_1, f_2, f_3, g_1, g_2\}^{\mathbb{Z}}, \sigma)$$
$$\pi(X_G, \sigma) \subset (\{e, f, g\}^{\mathbb{Z}}, \sigma)$$

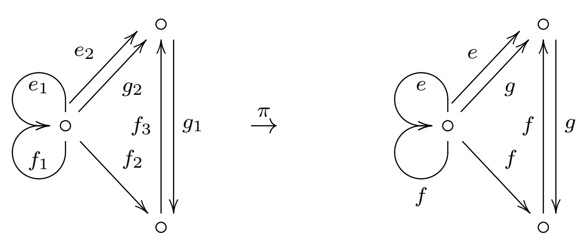

The map π just erases subscripts. The image system may or may not be conjugate to the original system, and it may or may not be an SFT. Image systems all come from relabeling edges in a non-unique way.

Example 28.1. The Golden Mean SFT's edge shift, under non-unique labeling of its edges produces the *even system.*

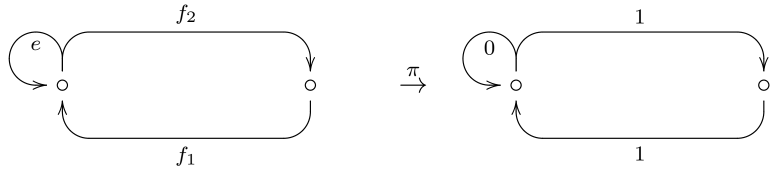

(Γ, σ) (S, σ)

(S, σ) is among the simplest sofic systems. $S =$ all sequences in $\{0,1\}^{\mathbb{Z}}$ such that between any 2 zeros there is an even number of ones. Consequently, (S, σ) is not an SFT (it is not determined by ruling out any finite list of blocks), and hence, by previous remark 27.1 it is not even conjugate to one. For, to see more precisely, suppose the system is m-step. Then if B is a block in $\mathcal{L}(S, \sigma)$ with $l(B) \geq m$, $aB \in \mathcal{L}(S, \sigma)$ and $Bz \in \mathcal{L}(s, \sigma)$, then $aBz \in \mathcal{L}(S, \sigma)$ (a word is *not* in $\mathcal{L}(S, \sigma)$ if and only if it contains a *bad word* of length $\leq m+1$. In (S, σ), however, for arbitrarily large n, $01^{2n+1} \in \mathcal{L}(S, \sigma)$ and $1^{2n+1}0 \in \mathcal{L}(S, \sigma)$, but $01^{2n+1}0 \notin \mathcal{L}(S, \sigma)$.

Remarks 28.1. (1) (S, σ) is *strictly sofic*, i.e. it is not an SFT.
(2) SFTs are also called *topological Markov chains.*
(3) Factors of Markov measure are called *sofic measures.* Furstenberg called them "submarkov processes." B. Weiss coined the term "sofic," a Hebrew term whose meaning conveys finiteness.
(4) Most systems are not sofic. The Morse and Toeplitz sequences, among many others, stand as examples of systems we have studied that are not sofic. The systems fail to meet sofic requirements because for example, they lack periodic points.

28.1. **Shannon's Message Generators.** A Shannon's Message Generator is a directed graph in which the vertices are the 'states' and the edges are labeled (maybe not uniquely) with symbols.

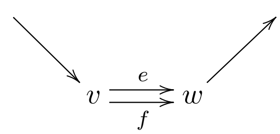

If the machine is in state v, it can emit any symbol written on an edge leaving v (either symbol e or f in the picture) and move to the terminal vertex of that edge (w in the picture). The set of all possible messages is a sofic system.

29. April 21 *(Notes by KN)*

29.1. **Sofic Systems.**

Definition 29.1. A *sofic system* is a subshift which is a factor of a subshift of finite type. A sofic system can also be thought of as a labeled edge-shift of a directed finite graph (with possibly repeated labels).

The word *sofic* comes from a Hebrew word which means finite. As we read a sequence from a sofic system, at each state there are only a finite number of possible futures. This is a generalization of the situation in a 1-step subshift of finite type.

Example 29.1. The Even Sofic Shift. Recall that this shift requires an even number of 1's between any pair of 0's. When you are generating sequences, there are two states determined by the parity of the number of 1's seen since the last 0.

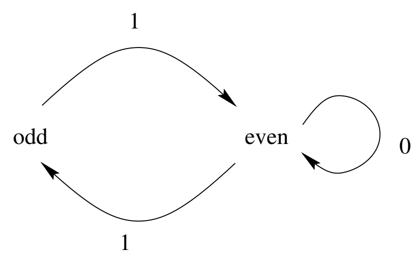

Figure 32. The even sofic shift.

29.1.1. *Characterizations of Sofic Systems.* Consider the following eight different characterizations of a sofic system, (S, σ), which we will prove to be equivalent.

(1) Follower Sets.
Let $(X, \sigma) \subset (D^{\mathbb{Z}}, \sigma)$ be any subshift. Let $w \in \mathcal{L}(X, \sigma)$ be any block that appears in some $x \in X$. Let $D^* = \bigcup_{i=0}^{\infty} D^{\mathbb{Z}}$, the set of all finite words on D. Define the *follower set* of w to be

$$\mathcal{F}(w) = \{ \text{ words } s \in D^* : ws \subset \mathcal{L}(X, \sigma)\}.$$

Then (S, σ) is sofic if and only if $\text{card}\{\mathcal{F}(w) : w \in \mathcal{L}(S, \sigma)\} < \infty$. That is (S, σ) is sofic if and only if the number of follower sets is finite.

Example 29.2. In an SFT or vertex shift, $\mathcal{F}(w_1 \dots w_n)$ is completely determined by w_n.

(2) Predecessor Sets, $\mathcal{P}(w)$.
The predecessor sets are defined in an analogous way to the follower sets.
(3) (S, σ) is a factor of an SFT, the image of a relabeling of the edges of a vertex shift, (1-block map), or 1-block map on an edge shift.
(4) $\mathcal{L}(S, \sigma)$ is a *regular language.*

Definition. A *regular language* consists of all words recognized by a *(deterministic) finite automaton (DFA).*

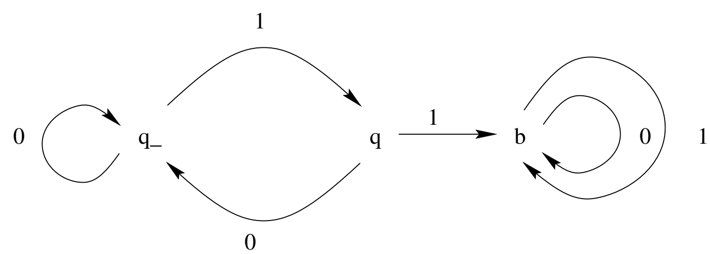

FIGURE 33. A DFA which recognizes the golden mean SFT.

Definition. Let
$D =$ the alphabet (for the words),
$Q =$ the finite set of states or vertices,
$q_{\_} \in Q =$ the initial state,
$Q_+ \subset Q =$ the set of final (good) states,
$\delta : Q \times D \to Q, \delta(q, a) = q \cdot a =$ the state entered by reading symbol a from state q. We require $\delta(q, \epsilon) = q$, where ϵ is the empty word. For words $w \in D^*, w = w_1 \dots w_n$, define $\delta(q, w) = q \cdot w = ((q \cdot w_1) \cdot w_2 \cdot \dots) \cdot w_n$. We say a word w is *accepted* by the DFA if $q \cdot w \in Q_+$.

A *nondeterministic finite automaton (NDFA)* has $\delta : Q \times D \to 2^Q$. That is, an NDFA could have many edges leaving state q with the same label, $a \in D$. It can also have ϵ *-moves*, where $\delta(q_1, \epsilon) = q_2$, but $q_1 \neq q_2$.

A language $\mathcal{L}$ is accepted by a DFA if and only if it is a language accepted by an NDFA with ϵ -moves. Note that there exist regular languages that do not come from subshifts. For example consider the language consisting of just one symbol.

Example 29.3. A DFA that recognizes the golden mean SFT: see Figure 33

(5) $\mathcal{L}(S, \sigma)$ is denoted by a *regular expression*.

Definition. Let D be a finite alphabet. A *regular expression* is a finite string on $\mathbf{D} = D \cup \{+, \cdot, *, \emptyset, \epsilon, (,)\}$. The set of finite regular expressions, $\mathcal{R}(D)$, is the smallest family of finite strings on $\mathbf{D}$ which has the following two properties: (i) $\mathcal{R}(D) \supset D \cup \{\emptyset, \epsilon\}$, and (ii) $\mathcal{R}(D)$ is closed under $+, \cdot$ and $*$, ($r, s \in \mathcal{R}(D)$ implies $r + s, r \cdot s$, and $r^* \in \mathcal{R}(D)$).

For each $r \in \mathcal{R}(D)$ we define a language $\mathcal{L}(r) \subset D^*$ as follows:
$\mathcal{L}(\emptyset) = \emptyset$ (the empty language is accepted by an automaton which has no final states),
$\mathcal{L}(\epsilon) = \{\epsilon\}$,
$\mathcal{L}(a) = \{a\}$, for $a \in D$,
$\mathcal{L}(r + s) = \mathcal{L}(r) \cup \mathcal{L}(s)$,
$\mathcal{L}(r \cdot s) = \mathcal{L}(r) \cdot \mathcal{L}(s) = \{uv : u \in \mathcal{L}(r), v \in \mathcal{L}(s)\}$,
$\mathcal{L}(r^*) = \mathcal{L}(r)^* = \cup_{i=0}^{\infty} \mathcal{L}(r)^i$.

Example 29.4. (a) Let $r = (0 + 1)^*(00)(0 + 1)^*$. Then $\mathcal{L}(r) =$ set of all words on $D = \{0, 1\}$ which contain 00.

(b) Let $r = (1 + 0 + \epsilon)(0 + 11)^*(1 + \epsilon)$. Then $\mathcal{L}(r) =$ all blocks on $\{0, 1\}$ with no $01^{2n+1}0$. So, $\mathcal{L}(r) = \mathcal{L}(S, \sigma)$, where (S, σ) is the even sofic subshift. Note that the even subshift allows words starting or ending with an odd number of 1's.

(6) Semigroup Realization.

There is an injection of D, the alphabet of (S, σ), into a finite multiplicative *semigroup*, S, with an absorbing element $0 \in S$ such that $0s = s0 = 0$ for all $s \in S$. The mapping is defined by $a \to s_a$ for $a \in D$. For any $w = w_1 \dots w_n \in D^*$, we say $w \in \mathcal{L}(S, \sigma)$ if and only if $s_{w_1} s_{w_2} \dots s_{w_n} \neq 0$.

(7) Matrix Semigroup Realization.

The semigroup in (6) can be realized by a semigroup of $d \times d$ matrices on $\{0, 1\}$ with "reduced" matrix multiplication, where any nonzero element is changed to a 1.

(8) $\mathcal{L}(S, \sigma)$ is generated by a *linear phrase structure grammar*.

Let $D =$ *terminals* and $V =$ *variables* be two finite alphabets. Let S be a *start symbol*. A *production* is a pair of words (α, ω) on $V \cup D$, that is $\alpha, \omega \in (V \cup D)^*$. (We also write the production as $\alpha \to \omega$.) We assume we have a *finite set* of such productions. Let $\mathcal{L} =$ all finite words on D which can be made by starting with S and applying a finite sequence of productions. The grammar is *linear* if all productions are of the form $A \to Bw$ or $A \to w$ for some $A, B \in V, w \in D^*$.

30. April 23 *(Notes by PS)*

Theorem 30.1. *The above eight characterizations of sofic subshifts are all equivalent.*

Proof. We give an outline only. We may need to assume some of the subshifts in question are topologically transitive.

30.1. **3 ⇒ 1.** Assume (S,σ) is a relabeling of the edges of a vertex shift (X_A,σ), that defines a factor map

$$\pi : (X_A,\sigma) \to (S,\sigma).$$

For any word $w \in \mathcal{L}(S,\sigma)$, there is some word $w' \in X_A$, not necessarily unique, with $\pi(w') = w$. Consider all possible w', and all possible terminal vertices w'_n of w'. The follower set $\mathcal{F}(w)$ is completely determined by this subset of terminal vertices in X_A. Since there are only finitely many such subsets, there are only finitely many follower sets for (S,σ).

30.2. **1 ⇒ 3.** Suppose (S,σ) is a subshift whose language $\mathcal{L}(S,\sigma)$ is a language on the alphabet D with only finitely many follower sets. We need to construct a vertex SFT whose edge-labeling factors into (S,σ). Make a graph G whose vertices are the follower sets $\mathcal{F}_1, \ldots, \mathcal{F}_n$. Put an edge with label $a \in D$ between $\mathcal{F}_i = \mathcal{F}(w)$ and $\mathcal{F}_j$ if and only if $wa \in \mathcal{L}(S,\sigma)$.

$$\mathcal{F}_i = \mathcal{F}(w) \xrightarrow{a} \mathcal{F}(wa) = \mathcal{F}_j$$

Note that this labeling does not depend on the choice of w. For if w' is another word with $\mathcal{F}(w') = \mathcal{F}(w)$, then $wa \in \mathcal{L}(S,\sigma)$ if and only if $w'a \in \mathcal{L}(S,\sigma)$, and $\mathcal{F}(w'a) = \mathcal{F}(wa)$.

It is clear that all words read off edge paths in G are in $\mathcal{L}(S,\sigma)$: given a path

$$\mathcal{F}(w_1) \xrightarrow{a_1} \mathcal{F}(w_1 a_1) \xrightarrow{a_2} \ldots \xrightarrow{a_{n-1}} \mathcal{F}(w_1 a_1 \ldots a_{n-1})$$

then $a_1 \ldots a_{n-1} \in \mathcal{L}(S,\sigma)$.

Conversely, given $a_1 \ldots a_{n-1} \in \mathcal{L}(S,\sigma)$, by extendability there is an a_0 such that $a_0 a_1 \ldots a_{n-1} \in \mathcal{L}(S,\sigma)$, so that

$$\mathcal{F}(a_0) \xrightarrow{a_1} \mathcal{F}(a_0 a_1) \xrightarrow{a_2} \ldots \xrightarrow{a_{n-1}} \mathcal{F}(a_0 a_1 \ldots a_{n-1})$$

is a path in G.

This gives a 1-block factor map on the edge shift (X_G,σ), which leads to our desired sofic shift.

30.3. **2 ⇔ 3.** The case of finitely many predecessor sets can be handled just as with finitely many follower sets.

30.4. **1 ⇒ 6.** Suppose $\mathcal{L}(S,\sigma)$ has finitely many follower sets. We wish to show that membership in $\mathcal{L}(S,\sigma)$ is determined by a finite semigroup. Define an equivalence relation on $\mathcal{L}(S,\sigma)$ by $w \sim w'$ if and only if $\mathcal{F}(w) = \mathcal{F}(w')$ and $\mathcal{P}(w) = \mathcal{P}(w')$. Define a multiplication of equivalence classes $[u]$ and $[v]$ for $u, v \in \mathcal{L}(S,\sigma)$ by

$$[u][v] = \begin{cases} [uv] & \text{if } uv \in \mathcal{L}(S,\sigma) \\ 0 & \text{otherwise} \end{cases}$$

Furthermore, let $[u]0 = 0[u] = 0$, for all $u \in \mathcal{L}(S,\sigma)$.

Then

$$\Sigma = (\mathcal{L}(S,\sigma)/_{\sim}) \cup \{0\}$$

is the desired semigroup with absorbing element 0. Clearly

$$w_1 w_2 \dots w_n \in \mathcal{L}(S,\sigma) \text{ if and only if } [w_1][w_2]\dots[w_n] \neq 0.$$

Remark: $(\mathcal{L}(S,\sigma)/_{\sim})$ is sometimes called the "syntactic monoid".

30.5. **6 ⇒ 1.** Suppose membership in $\mathcal{L}(S,\sigma)$ is determined by a finite semigroup Σ. Note that $\mathcal{F}(w_1 w_2 \dots w_n)$ depends only on the semigroup element $s = s_{w_1} \dots s_{w_n}$, because

$$u_1 \dots u_n \in \mathcal{F}(w_1 \dots w_n) \text{ if and only if } (s_{w_1} \dots s_{w_n})(s_{u_1} \dots s_{u_n}) \neq 0$$

Since Σ is a finite semigroup, it follows that there are only finitely many follower sets.

30.6. **3 ⇔ 7.** We wish to show (S,σ) is a factor of a vertex subshift of finite type by a 2-block map by relabeling vertices if and only if membership in $\mathcal{L}(S,\sigma)$ is determined by a finite semi- group of $d \times d$ 0,1-matrices.

Suppose we have a 2-block map $\varphi(ij)$ for $i, j \in D$, generating a factor map

$$\varphi : (X_A, \sigma) \to (S, \sigma).$$

Take the adjacency matrix A and decompose it according to the edge labeling φ:

$$A = A_1 + A_2 + \dots + A_k,$$

where 1, 2, …, k are the labels on the edges of the graph of X_A, as follows: put

$$(A_k)_{ij} = \begin{cases} 1 & \text{if } \varphi(ij) = k \\ 0 & \text{otherwise} \end{cases}.$$

Let Σ be the semigroup generated by A_1, …, A_k with reduced matrix multiplication, that is where nonzero entries are replaced by 1 after multiplication. There are only finitely many such matrices, namely

$$\text{card}(\Sigma) \leq 2^{(d^2)}.$$

Note that a word $w = w_1 w_2 \dots w_n$ on the alphabet D of (S,σ) occurs along an edge path if and only if $A_{w_1} A_{w_2} \dots A_{w_n} \neq 0$. If we have

$$x_1 \overset{w_1}{\to} x_2 \overset{w_2}{\to} \dots \overset{w_n}{\to} x_{n+1}$$

then

$$(A_{w_1} A_{w_2} \dots A_{w_n})_{x_1 x_{n+1}} \neq 0$$

and conversely, since a nonzero entry implies the existence of a path.

For example, consider the even shift.

With respect to the the 2-block map $\varphi(12) = \varphi(21) = 1$ and $\varphi(11) = 0$, the matrix A decomposes as

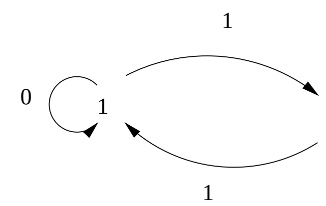

FIGURE 34. The even shift

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

30.7. **3 ⇒ 4.** Suppose we have an edge labeling of a subshift of finite type. We wish to show that the resulting language is recognizable by a DFA, or equivalently an NDFA with ε-moves allowed. Let

Q = the set of states of the automaton = vertexes of the graph (X_A, σ),
Q_+ = the set of good final states = Q,
q_- = starting state with an ε-move to any state in Q,
b = bad final state.
Define moves appropriately. For example, for the even shift

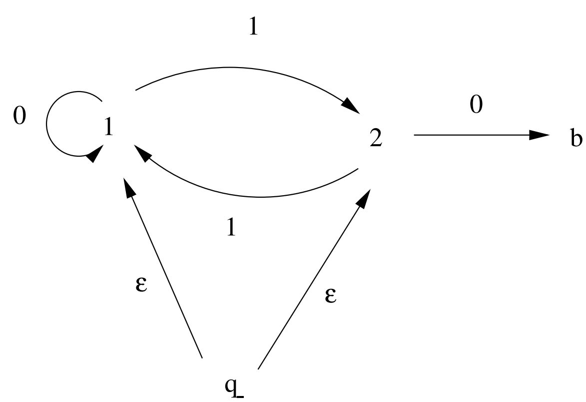

FIGURE 35. The even shift as an NDFA with ε-moves

30.8. **4 ⇒ 3.** Given a DFA, look at all doubly-infinite admissible sequences on the graph of the automaton. This is an edge labeling of a vertex shift and thus a sofic shift.

□

31. April 28 *(Notes by SB)*

Proof (continued) of the equivalence of the eight different characterizations of a sofic system, (S, σ):

4 $\Leftrightarrow$ **8**. Recall:

4. $\mathcal{L}(S, \sigma)$ is a *regular language* (recognized by a DFA).
8. $\mathcal{L}(S, \sigma)$ is generated by a *linear phrase structure grammar*.

We establish an equivalence between the two systems according to the following table:

DFA	phrase structure grammar
edge labels	terminals, D
vertices	variables
initial state q_-	start state, S
each edge (to a good vertex)	a production

Example 31.1. For the golden mean SFT, we have the DFA given by Figure 36. Using the

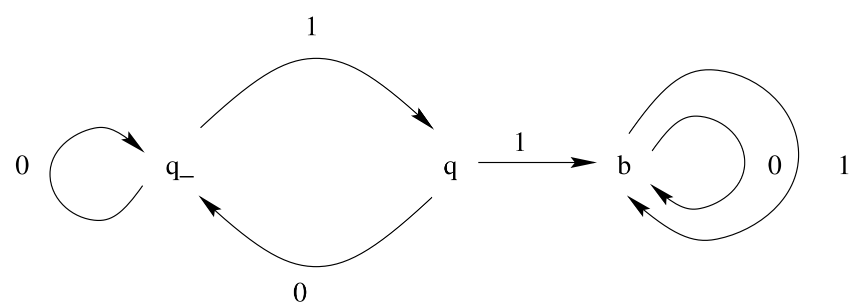

Figure 36. A DFA which recognizes the golden mean SFT

equivalence established in the table above, we have that the terminals for the phrase structure grammar of the golden mean SFT are $D = \{0, 1\}$, and the variables are $V = \{q_-, q, b\}$. To figure out the set of productions we need to decide what q_- and q can go to (since these are the *good* vertices). We have the following:

$q_- \to 0$
$q_- \to 1$
$q_- \to 0q_-$
$q_- \to 1q$
$q \to 0$
$q \to 0q_-$

This gives a right-linear phrase structure grammar, since all of the productions are of the form variable$\to$terminal or variable$\to$terminal$\cdot$variable. To form a word in the language, we start with $S = q_-$ and apply a finite sequence of productions in such a way as to end up with a string of terminals. All possible words in the language are found as the leaves on the *derivation tree*: list all possible productions from the start state S and then list all possible places where each variable could go. For example, suppose $w = 1000101001$. Then part of the derivation tree leading to w would be:

$$S \to 1q \to 10q_- \to 100q_- \to 1000q_- \to 10001q \to 100010.$$

Note: A language $\mathcal{L}$ is generated by a *left-linear* grammar if and only if it is also generated by a right-linear grammar.

To establish the equivalence of (4) and (8) in the other direction, we need to describe how to make a DFA that recognizes the phrase structure grammar. Simply make a graph using the variables as vertices, that is, given $A \to aB$ put an edge $A \xrightarrow{a} B$.

5 ⇔ 4. Recall:

4. $\mathcal{L}(S,\sigma)$ is a *regular language* recognized by a DFA or an NDFA with ϵ-moves.
5. $\mathcal{L}(S,\sigma)$ is denoted by a *regular expression.*

A standard tool for proofs in language theory is to induct on the length of the regular expression r. This is the technique we use to prove $5 \Rightarrow 4$.

Suppose r has length 1. Then

$r \;=\; \epsilon$ or $\emptyset$ or $a \in D$, and correspondingly

$\mathcal{L}(r) \;=\; \epsilon$ or $\emptyset$ or $\{a\}$.

We want NDFA's with ϵ-moves that will recognize these languages. The following one will suffice for the case $r = \epsilon$:

$$b \xleftarrow{a_i} q_- \xrightarrow{\epsilon} q$$

We are allowed to go via ϵ to any *good* state. From q_-, on any other letter that might come in, we go to the bad state b. For the case $r = \emptyset$ we have the following NDFA:

$$q_- \xrightarrow{\epsilon} b \notin Q_+$$
$$q_1 \in Q_+.$$

Here everything is sent straight to a bad state, b. A good state, q_1 can exist in the NDFA, but no edges go into it. The NDFA below recognizes the language denoted by $r = a$:

$$q_- \xrightarrow{a} q \in Q_+.$$

So we have found NDFA's for all cases arising when $r = 1$. The rest follows by induction so long as we can describe how the NDFA will work to be compatible with the $+$, $*$, and $\cdot$ operations for building up regular expressions. We achieve this in the following way:

For $r_1 + r_2$ we have $\mathcal{L}(r_1 + r_2) = \mathcal{L}(r_1) \cup \mathcal{L}(r_2)$. So if $\mathfrak{A}_1$ is the NDFA for r_1 and $\mathfrak{A}_2$ is the NDFA corresponding to r_2, then wire these machines in parallel to obtain the NDFA for $r_1 + r_2$. See Figure 37. If the NDFA for r_1 is $\mathfrak{A}_1$, and the NDFA for r_2 is $\mathfrak{A}_2$, then the NDFA for $r_1 + r_2$ is $\mathfrak{A}_1$ and $\mathfrak{A}_2$ wired in parallel. For $\mathfrak{A}_1$, let q_1 be the initial state and f_1 be a final good state so that every final good move gets sent to f_1 by an ϵ-move. Similarly, for $\mathfrak{A}_2$. Then let q_- be the initial state and f the final good state for $\mathfrak{A}_1$ and $\mathfrak{A}_2$ wired in parallel. Then anything in $\mathcal{L}(r_1 + r_2)$ will be accepted by this machine by going along the correct path. Also, if we can find any path that works (i.e., any path that is accepted by the machine), then the word is in $\mathcal{L}(r_1 + r_2)$.

For concatenation, $r_1 \cdot r_2$, we have $\mathcal{L}(r_1 \cdot r_2) = \mathcal{L}(r_1)\mathcal{L}(r_2)$, so put the machines $\mathfrak{A}_1$ and $\mathfrak{A}_2$ in series, connecting them via ϵ-moves as in Figure 38. There was a question about what would happen if there was a long word for which some first part of it is accepted by machine $\mathfrak{A}_1$ so that the word gets pushed into machine $\mathfrak{A}_2$ too early, causing the word to end up in a bad state. The point is, however, that there is *some* path in machine $\mathfrak{A}_1$ that will accept the entire part of the word that lies in $\mathcal{L}(r_1)$.

For r_*, $\mathcal{L}(r^*) = \bigcup_{i=0}^{\infty}[\mathcal{L}(r)]^*$, so let $\mathfrak{A}$ be the NDFA for r and design the NDFA as in Figure 39. The ϵ-move from the final good state f' of $\mathfrak{A}$ to the initial state q' of $\mathfrak{A}$ allows for iterations. The

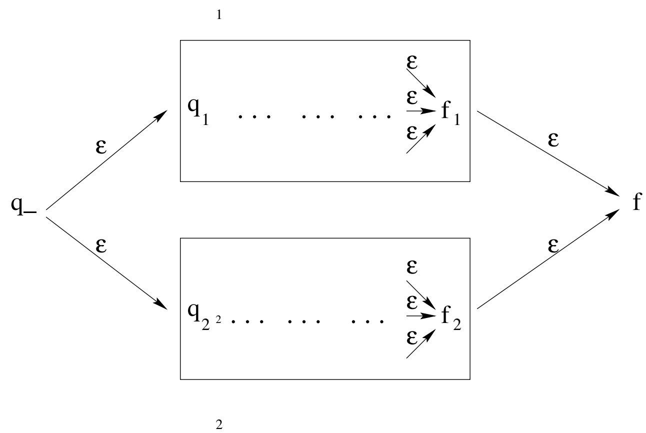

FIGURE 37. To create the NDFA for the sum, wire the two individual machines in parallel

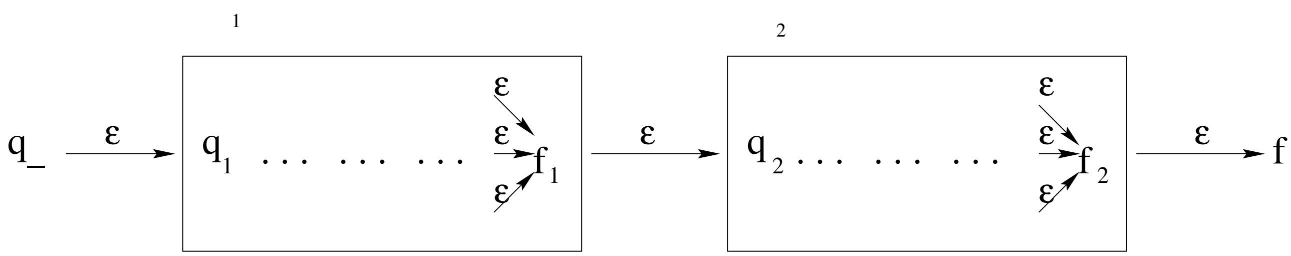

FIGURE 38. The NDFA for $r_1 \cdot r_2$. Again, q_- is the initial state and f the final good state for the whole machine. A word is accepted by the compound machine if and only if the first machine accepts the part of the word in $\mathcal{L}(r_1)$ and the second machine accepts the part of the word in $\mathcal{L}(r_2)$.

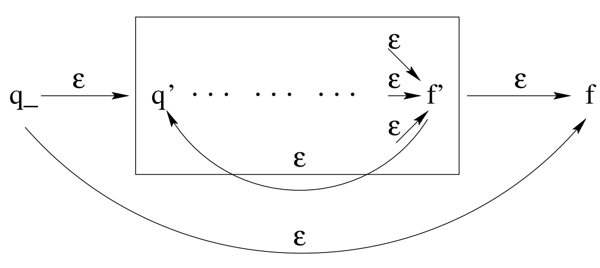

FIGURE 39. The NDFA for r^*.

ϵ-move from the initial state q_- to the final good state f of the whole machine allows for the ϵ word to be in the language. So by induction, as we build up the regular expression, we use our soldering

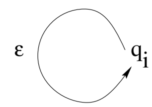

FIGURE 40. The DFA for R^0_{ij}. There are no other ϵ-moves allowed in this DFA.

tool to put smaller machines together, thus building more and more complicated machines until we achieve an NDFA that will recognize the regular expression.

4 ⇒ 5. We may assume that we have a DFA $\mathfrak{A}$ with no ϵ-moves. This narrows our hypothesis, making the argument easier. So suppose the set of vertices of the DFA is $Q = \{q_1, \ldots, q_n\}$. We use a variation of Warshaw's Algorithm (an algorithm used to construct, fairly efficiently, the transitive closure of a relation).

For $k = 0, 1, \ldots, n$ and $i, j = 1, \ldots, n$, put

R^k_{ij} = all strings accepted by paths on the DFA
from the vertex q_i to the vertex q_j
that do not hit any q_m with $m > k$ in between,
except maybe at the ends.

This amounts to restricting to a subgraph or subautomaton so that in between the endpoints, we are only allowed to move within this subgraph. Notice that

$$\mathcal{L}(\mathfrak{A}) = \bigcup_{q_j \in Q^+} R^n_{1j}$$

if $q_- = q_1$. Also $R^k_{ij} \subset R^{k+1}_{ij}$. We induct to show that each R^k_{ij} is denoted by a regular expression.

For the case $k = 0$, R^0_{ij}= all strings accepted by paths from q_i to q_j that hit NO element of Q except maybe at the ends. The only possibilities are $\emptyset$, ϵ, and $a \in D$. For $\emptyset$ we get the empty graph, which has length 0; so $r = \emptyset$. For ϵ we have $\delta(q_i, \epsilon) = q_i$, which corresponds to the graph of length 1 in Figure 40. So we take $r = \epsilon$. For $R^0_{ij} = a \in D$ we get the following graph of length 2:

$$q_i \overset{a}{\longrightarrow} q_j$$

Here q_j could equal q_i.We take $r = a$. So we have proved the base case for the induction.

For $k \geq 1$ we have

$$R^k_{ij} = R^{k-1}_{ik} (R^{k-1}_{kk})^* R^{k-1}_{kj} \cup R^{k-1}_{ij}.$$

To see this, consider a path (not a string), $q_i \ldots q_j$ in which there is nothing between q_i and q_j with index greater than k. There could, however, be q_k's between q_i and q_j. If there are no q_k's between q_i and q_j, then the string is in R^{k-1}_{ij}. If there are q_k's between q_i and q_j, look at the first place where we hit a q_k:

$$q_i \underbrace{\overbrace{\cdots} q_k \overbrace{\cdots} q_k \overbrace{\cdots} q_k \ldots q_k \overbrace{\cdots} q_k \overbrace{\cdots}} q_j$$

Then the first segment $q_i \ldots q_k$ contains no q_k's on the inside. So this part of the path is in R^{k-1}_{ik}. Then we hit a bunch of paths $q_k \ldots q_k$ which have no q_k's on the inside. These all belong to (R^{k-1}_{kk}), so the string of them belongs to $(R^{k-1}_{kk})^*$. Finally, at the last q_k we have a path $q_k \ldots q_j$ which has no q_k's inside, so this path is in R^{k-1}_{kj}. The equivalence follows by induction.

This concludes the theorem that all of the eight characterizations of a sofic system are equivalent.

Theorem 1 (R. Fischer, 1975; also W. Krieger *et.al*)**.** *Let (S, σ) be a topologically transitive ("irreducible") sofic system. Then there exist an irreducible vertex SFT (X_A, σ) and a factor map $\pi : (X_A, \sigma) \to (S, \sigma)$ given by an edge labeling such that:*

(1) *The edge labeling of the graph G of (X_A, σ) is* right resolving*, i.e., the edges leaving any vertex have distinct labels.*
[Definition: If (X, σ) and (Y, σ) are subshifts and $\pi : X \to Y$ is a 1-block map, then π is called right resolving *if whenever $ab, ac \in \mathcal{L}(X, \sigma)$ and $\phi(b) = \phi(c)$ then $b = c$.]*

(2) *Hence $\pi : (X_A, \sigma) \to (S, \sigma)$ is boundedly finite-to-one (i.e., $\operatorname{card} \pi^{-1}(y) \leq m < \infty$ for all $y \in S$) and one-to-one on doubly transitive points (i.e.,points x such that $\overline{\mathcal{O}^+(x)} = \overline{\mathcal{O}^-(x)} = X$. where $\mathcal{O}^+(x) = \{\sigma^n x : n \geq 0\}$). Then π is one-to-one almost everywhere with respect to any ergodic invariant measure with full support and on a residual set.*

(3) *(X_A, σ) is the* smallest *right-resolving extension in that it has the fewest vertices.*
Moreover, any two such extensions (that is, extensions that are right resolving and smallest) have isomorphic labeled graphs. (X_A, σ) is called the right (or future) Fischer cover *of (S, σ).*

Corollary 31.1. *Every topologically transitive sofic system has a unique measure of maximal entropy,(it is intrinsically ergodic), the entropy of which is the logarithm of a Perron number (i.e., a positive algebraic integer that dominates all its conjugates).*

Proof. (Sketch)

(1) Make a lift as in the implication (1) $\Rightarrow$ (3) in the above proof of the equivalent definitions of a sofic system to obtain an SFT that factors onto the sofic system, (S, σ).

(2) This lift might be "too big", that is, some of its vertices might be redundant, so merge the vertices that have the same follower sets (of labeled edges) to get a "tighter" factor map of an SFT onto (S, σ).

(3) In irreducible sofic shifts, there are *synchronizing words* or "Markov magic words," that is, words $\tau \in \mathcal{L}(S, \sigma)$ such that whenever you see τ along the edges of G, the right-resolving graph above, the terminal vertex is always the same. Equivalently, if $w\tau$, $\tau v \in \mathcal{L}(S, \sigma)$, then $w\tau v \in \mathcal{L}(S, \sigma)$. The idea of τ always leading to the same terminal vertex can be described using the analogy of a "road map to Cleveland." The map is a word on the symbols $\{l, r, sa, b, c, d\}$. Roads leaving each junction have been labeled with these four symbols. At each junction, the driver chooses which roads to follow depending on what the next letter in the word tells him to do. Regardless of where the driver starts, the "word road map" takes him or her to Cleveland.

Exercise 4. In an irreducible sofic system, such a word always exists.

(4) Every $y \in S$ that has such a word τ infinitely many times to the left has a singleton preimage.

□

32. April 30 *(Notes by SS)*

32.1. **Shannon Theory.** For reference, see C. Shannon, A. Khinchin, R. Pötschke–F. Sobik, R. Gallager, T.M. Cover, J. Singh, J. Pierce (popular), Martin–England, or I. Csiszar's article.

Definition 32.1. A *source* $[A_0, \mu_0]$ is a finite-state stationary (ergodic) stochastic process, where A_0 is a finite alphabet and μ_0 is a shift-invariant ergodic measure on $A_0^{\mathbb{Z}}$. An example might be a Shannon machine with a probability measure determined by transition probabilities on its edges.

Definition 32.2. A *channel* $[A, \{\nu_x\}, B]$ consists of a finite input alphabet A, a finite output alphabet B, and a family of measures $\{\nu_x : x \in X\}$ defined as follows :

If $X = A^{\mathbb{Z}}$ = all potential input messages and $Y = B^{\mathbb{Z}}$ = all potential output messages, then for each $x \in X$, ν_x is a (Borel probability) measure on Y such that

(1) $\nu_{\sigma x} = \sigma \nu_x = \nu_x \circ \sigma^{-1}$ (so the channel is *stationary*) and

(2) for all measurable $F \subset Y$, the map $x \mapsto \nu_x(F)$ is measurable on X.

Now, given a measure μ on X (the *input measure*), the *input-output measure* λ on $X \times Y$ is defined by

$$\lambda(E \times F) = \int_E \nu_x(F) d\mu(x)$$

for measurable $E \subset X, F \subset Y$.

This represents the probability that the output signal is in F given that the input was in E.

Example 32.1. Let $A = B = \{0, 1\}$, and let $\beta_0 = (1-\epsilon)\delta_0 + \epsilon\delta_1$ and $\beta_1 = \epsilon\delta_0 + (1-\epsilon)\delta_1$ be two measures on A, B. Define, for $x = (x_k) \in X = \{0,1\}^{\mathbb{Z}}$,

$$\nu_x = \prod_{j=-\infty}^{\infty} \beta_{x_j} \qquad \text{on} \quad Y = \{0,1\}^{\mathbb{Z}}.$$

For example, if $E = \{x : x_0 = 0\} \subset X$ and $F = \{y : y_0 = 0\} \subset Y$, then

$$\begin{aligned} \lambda(E \times F) &= \lambda(\{(x, y) : x_0 = 0, y_0 = 0\}) \\ &= \int_{\{x:x_0=0\}} \nu_x(\{y : y_0 = 0\}) d\mu(x) \\ &= (1-\epsilon)\mu(\{x : x_0 = 0\}), \end{aligned}$$

since $\nu_x(\{y : y_0 = 0\}) = \epsilon$ if $x_0 = 1$, and $1 - \epsilon$ if $x_0 = 0$.

Similarly,

$$\lambda(\{(x, y) : x_0 = 1, y_0 = 1\}) = (1-\epsilon)\mu(\{x : x_0 = 1\})$$

and so

$$\lambda(\{(x, y) : x_0 = y_0\}) = 1 - \epsilon.$$

This channel represents a situation in which 0 changes to 1, 1 changes to 0, with probability ϵ, independently in each coordinate. It is called a **DMC**, *discrete memoryless channel.*

32.1.1. *Source coding.* Suppose there are given a source $[A_0, \mu_0]$ and a channel $[A, \{\nu_x\}, B]$. We may wish either to recode the source to compress information before connecting to the channel (Shannon-McMillan-Breiman Theorem) or connect to the channel by means of an *encoder*, a map $e : A_0^{\mathbb{Z}} \to A^{\mathbb{Z}}$ (it is not necessarily one-to-one, continuous, or shift-commuting).

- Three popular types of encoder

(1) Block code : e is determined by a fixed map from A_0^k to A^n, i.e., for $s \in A_0^{\mathbb{Z}}$,

$$s = \cdots\cdots\cdots\cdots . \underbrace{s_0 s_1 \cdots s_{k-1}} \underbrace{s_k s_{k+1} \cdots s_{2k-1}} \cdots$$

$$\downarrow$$

$$e(s) = \cdots\cdots\cdots\cdots . \overbrace{a_0 a_1 \cdots a_{n-1}} \overbrace{a_n a_{n+1} \cdots a_{2n-1}} \cdots .$$

Here $n/k =$ the *rate* of the encoder (e is σ^k, σ^n-commuting).

(2) Sliding block code : For $s \in A_0^{\mathbb{Z}}$, e is defined by

$$(es)_0 = e(s_{-m} \cdots s_{-1} s_0 s_1 \cdots s_a).$$

(3) Arbitrary measurable map $\phi : A_0^{\mathbb{Z}} \to A^{\mathbb{Z}}$.

32.1.2. *Shannon-McMillan-Breiman Theorem.* The possibilities for efficient recoding (or encoding) of a source are given by the Shannon-McMillan-Breiman Theorem :

Theorem 32.1. *(Shannon-McMillan-Breiman) If μ_0 is an ergodic shift-invariant measure on $X = A_0^{\mathbb{Z}}$, then*

$$\underbrace{-\frac{1}{n} \log \mu_0[x_0, \cdots, x_{n-1}]}_{f_n(x)} \longrightarrow h(\mu_0) \quad a.e.$$

and in L^1, i.e., $\int_X \mid f_n(x) - h(\mu_0) \mid d\mu_0(x) \to 0$.

- Two reinterpretations

(1) The *asymptotic entropy equipartition property* : Since convergence in L^1 implies convergence in measure, given $\epsilon > 0$, there exists $N > 0$ such that if $n \geq N$, then

$$\mu_0(\{x : | -\frac{1}{n} \log \mu_0[x_0, \cdots, x_{n-1}] - h(\mu_0)| < \epsilon\}) > 1 - \epsilon.$$

Then we get the *asymptotic entropy equipartition property* : for $n \geq N$, there is a list of "good" n-blocks on A_0 such that
(i) the union of the corresponding cylinder sets has measure $> 1 - \epsilon$,
(ii) each corresponding cylinder set has μ_0-measure between $e^{-n[h(\mu_0)+\epsilon]}$ and $e^{-n[h(\mu_0)-\epsilon]}$,
(iii) there are at least $e^{n[h(\mu_0)-\epsilon]}$ of them.

(2) Source Coding Theorem : Let us switch the base of logarithms to 2 so that we can talk about *bits*. Then the minimum mean number of bits per symbol required to encode an ergodic source is the entropy of the source.

Theorem 32.2. *Let $[A_0, \mu_0]$ be an ergodic stationary source and $A = \{0, 1\}$. Then, given $\epsilon > 0$, if k is large enough, if $n/k > h(\mu_0) + \epsilon$, then there is a block code $e : A_0^k \to A^n$ (so rate $= n/k$) which is one-to-one on a set of k-blocks on A_0 whose associated cylinder sets form a set of input sequences of μ_0-measure $> 1 - \epsilon$.*

Proof. We have $K \leq 2^{k[h(\mu_0)+\epsilon]}$ "good" sequences on A_0^k (that cover $> 1 - \epsilon$ of X). We want to assign to each "good" k-block a different n-block on $\{0, 1\}$. If $n \geq \log K$, then $2^n \geq K$ and we have 2^n n-blocks on $\{0, 1\}$. So the assignment can be made if $n \geq \log K \geq k[h(\mu_0) + \epsilon]$, i.e., $n/k \geq h(\mu_0) + \epsilon$. ∎

Conversely, any code from A_0^* to A^* which is one-to-one on each A_0^k must have expected rate $\geq h(\mu_0)$. We state this more precisely as follows :

If $e : A_0^* \to A^*$ is a map which is one-to one on each A^k, then

$$\int \frac{l(e[x_0 \cdots x_{k-1}])}{k} d\mu_0(x) \geq \frac{H(\alpha_0^{k-1})}{k} \geq h(\mu_0).$$

(Here α is the time-0 partition in $A_0^{\mathbb{Z}}$ and $l(B)$ denotes the length of a block B.)

H. White showed that in fact

$$\liminf_{k\to\infty} \frac{l(e[x_0 \cdots x_{k-1}])}{k} \geq h(\mu_0) \quad \text{a.}e.\, d\mu_0.$$

33. May 5 *(Notes by KJ and RP)*

The Shannon-McMillan-Breiman Theorem of the last section gives a coding-theoretic understanding of the entropy of the source. There might be some part of the source which is irrelevant of infrequent, so you don't need to reserve all the extra blocks to code efficiently.

Remark 33.1. This sort of source coding becomes more effective in rate distortion theory. This theory begins with a measure of cost or distortion associated with each pair

$$u \in A_0^*, e(u) \in A^*. \tag{22}$$

Whereas in the applications of the Shannon-McMillan-Breiman Theorem we either code a block or we don't, in the *Source Coding Rate Distortion Theorem* we minimize total distortion.

33.1. **Connecting the Source to the Channel.** Recall that we have a source $[A_0, \mu_0]$ and a channel $[A, \{\nu_x\}, B]$ where A and B are the input and output alphabets, respectively, and each ν_x is a measure on $Y = B^{\mathbb{Z}}$ for each $x \in A^{\mathbb{Z}}$. We also have an input-output measure λ defined by

$$\lambda(E \times F) = \int_E \nu_x(F) d\mu(x) \tag{23}$$

where $E \subset X = A^{\mathbb{Z}}$, $F \subset Y = B^{\mathbb{Z}}$ and μ is the input measure on X. As discussed before, the channel can be seen as a wire which transmits the message or maybe something else, depending on chance.

We also have an encoder $e : A_0^{\mathbb{Z}} \to A^{\mathbb{Z}}$, a measurable function which prepares the message to be sent across the channel. For example, e could be a block or a sliding block code, the cases we will focus on.

The encoder determines an input measure $\mu = e\mu_0$ on $A^{\mathbb{Z}}$, i.e. $\mu(E) = \mu_0(e^{-1}E)$. We also get an output measure ν on $B^{\mathbb{Z}} = Y$ defined by $\nu F = \lambda(\pi_Y^{-1} F)$, where $\pi_Y : X \times Y \to Y$ is the projection. The measure ν gives the statistics of messages coming out of the channel if they come in with statistics given by the input measure.

In addition to the encoder, we may also have a decoder d, a measurable mapping from Y to $A_0^{\mathbb{Z}}$. After the process of encoding, transmission across the channel, and decoding, we hope that the messages which come out give an idea of the messages that went in.

In some of what follows we simplify the situation, assuming (for example) that the encoding map is the identity map, so that we can concentrate on the channel.

33.2. **Mutual Information and Capacity.** Capacity is a measure of how much information one can hope to push across the channel, analogous to how much water per-unit time can flow through a pipe.

Let $(X, \mathcal{B}, \mu)$ be a probability space and $T : X \to X$ a measure-preserving transformation. In our case, X is a subshift and T is the shift. If α and β are measurable partitions of X, then the

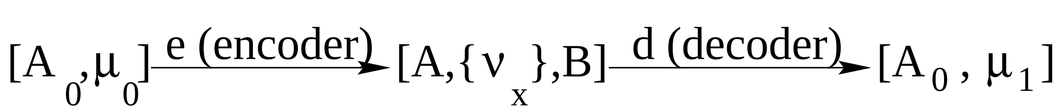

Figure 41. Encoding, transmission and decoding of a message

mutual information of α and β is

$$H(\alpha;\beta) = H(\alpha) - H(\alpha|\beta) \tag{24}$$

$$= H(\beta) - H(\beta|\alpha) \tag{25}$$

$$= H(\alpha) + H(\beta) - H(\alpha \vee \beta) \tag{26}$$

(The last equation comes from page 21 of the ergodic theory notes, Proposition 6.3 which gives $H(\alpha \vee \beta) = H(\alpha) + H(\beta|\alpha)$).

We can interpret $H(\alpha|\beta)$ as the amount of extra information you get from knowing what cell of α you are in, given that you know what cell of β you are in.

From the definitions of entropy,

$$H(\alpha;\beta) = -\sum_{A\in\alpha} \lambda(A) \log \lambda(A) - \sum_{A\in\alpha, B\in\beta} \lambda(A)\lambda(B) \log \lambda(A|B), \tag{27}$$

(where the log in question may be base 2).

When we are talking about a channel, we use this definition on

$$\alpha = \pi_X^{-1}(\text{time-0 partition of } X) \tag{28}$$

$$\beta = \pi_Y^{-1}(\text{time-0 partition of } Y). \tag{29}$$

This gives the mutual information of one symbol going in and another coming out. If we wish, we can use a higher block representation and get mutual information based on initial n-blocks: instead of α and β above, use

$$\alpha_0^{n-1} = \pi_X^{-1}(\text{time 0 partition of } X \text{ into } n \text{ blocks) and} \tag{30}$$

$$\beta_0^{n-1} = \pi_Y^{-1}(\text{time 0 partition of } Y \text{ into } n \text{ blocks).} \tag{31}$$

If μ is an input measure on $X = A^{\mathbb{Z}}$ for a channel $[A, \{\nu_x\}, B]$, define its *transmission rate* to be

$$R(\mu) = \lim_{n\to\infty} \frac{1}{n} H_\lambda(\alpha_0^{n-1};\beta_0^{n-1}) \tag{32}$$

$$= \lim_{n\to\infty} \frac{1}{n}[H_\lambda(\alpha_0^{n-1}) - H_\lambda(\alpha_0^{n-1}|\beta_0^{n-1})] \tag{33}$$

$$= \lim_{n\to\infty} \frac{1}{n}[H_\lambda(\beta_0^{n-1}) - H_\lambda(\beta_0^{n-1}|\alpha_0^{n-1})] \tag{34}$$

$$= \lim_{n\to\infty} \frac{1}{n}[H_\lambda(\alpha_0^{n-1}) + H_\lambda(\beta_0^{n-1}) - H_\lambda(\alpha \vee \beta)_0^{n-1}] \tag{35}$$

$$= h_\mu(\sigma_X) + h_\nu(\sigma_Y) - h_\lambda(\sigma_{X\times Y}). \tag{36}$$

(Recall that λ projects to μ on the first coordinate, and α depends only on the first coordinate, and similarly for ν and β).

To interpret these equations, remember that H represents the average information per symbol, and that gaining information is like losing uncertainty. So in equation 33 we see that the amount of information coming across per unit time is the information put in minus the uncertainty remaining about what was put in (α_0^{n-1}) given what we received (β_0^{n-1}). Also, we can interpret it (see 34) as the amount of information coming out minus the extra uncertainty due to noise, i.e., the uncertainty about what would come out (β_0^{n-1}) even if we knew what had been put in (α_0^{n-1}). Finally (in 35), the entropy of the input process plus the entropy of the output process minus entropy of the joint process reflects the information resulting from the connection between the two processes.

Given an input measure μ, $R(\mu)$ tells how much useful information is coming out of the other end of the channel per unit time.

The *capacity* of the channel is

$$C = \sup_{\mu} R(\mu), \tag{37}$$

the supremum being taken over all stationary ergodic measures μ.

In talking about capacity, engineers like to consider "operational definitions", discussing what you can actually do, as opposed to this more theoretical treatment. For other variations, we can restrict or open up the types of input statistics allowed (for example, consider non-ergodic or even non-stationary input measures). In principle, these definitions of capacity might be essentially different.

33.3. **Shannon's Channel Coding Theorem.** We will consider a "good" channel (which will be defined later) with capacity C.

Lemma 33.1 (Feinstein's Lemma)**.** *Given $\epsilon > 0$, for large enough n there are at least $N \geq 2^{n(c-\epsilon)}$ ϵ-distinguishable code words u_i in A^n: that is, there exist disjoint sets $V_1, V_2, \dots, V_N$ of n-blocks in B^n such that for every i,*

$$\nu_{u_i}(V_i) = \lambda(u_i \textit{ as input} \mid \textit{output was in } V_i) \tag{38}$$

$$= \lambda\{(x,y) : x_0 \dots x_{n-1} = u_i | y_0 \dots y_{n-1} \in V_i\} \tag{39}$$

$$\geq 1 - \epsilon. \tag{40}$$

Remark 33.2. So with probability $1-\epsilon$ we can decode the message: if a block in V_i is received, then with probability $1-\epsilon$ we can determine the block that it came from.

The proof involves a jazzed-up Shannon-McMillan-Breiman Theorem, again a random-coding, i.e. block-counting, argument. How you code depends on the channel statistics: we add check bits to protect against noise in the channel. There is a give and take between recoding the source for efficiency (noiseless coding, compression) and adding bits back to make sure the message gets across. Feinstein's Lemma says that it is theoretically possible to do this.

Finally, we have

Theorem 33.2 (Channel Coding Theorem)**.** *Consider a "good" channel $[A, \{\nu_x\}, B]$ with capacity C. Let $[A_0, \mu_0]$ be a stationary ergodic source with $h(\mu_0) < C$. Then given $\epsilon > 0$, for large enough n there exists a block code $e : A_0^n \to A^n$ and a decoder $d : B^n \to A_0^n$ such that when messages from the source are sent across the channel and decoded, the probability of error is $< \epsilon$, i.e.*

$$\lambda\{(x,y) : e^{-1}(x_0 \dots x_{n-1}) \neq d(y_0 \dots y_{n-1})\} < \epsilon, \tag{41}$$

and also $R(e\mu_0) > h(\mu) - \epsilon$.

Conversely, if $h(\mu_0) > C$, this is not possible for every $\epsilon > 0$.

The idea is that, for example, when $h(\mu_0) < 1$ and $A = \{0, 1\}$, then there are about $2^{nh(\mu_0)}$ good n-blocks in the source. For large n, we find $N \geq 2^{n(C-\epsilon)}$ ϵ-distinguishable n-blocks in A^n. Assign to each good n-block an ϵ-distinguishable block.

If $C - \epsilon > h(\mu_0)$, then this can be done using Feinstein's Lemma and the Shannon-McMillan-Breiman Theorem. We conclude that for all but ϵ of the source involved, blocks are distinguished with probability $> 1 - \epsilon$.

When this can be done for every $\epsilon > 0$, the source is called *block transmissible.*

33.4. **Good Channels.** What is a "good" channel (i.e., good enough that the conclusion of the preceding Channel Coding Theorem should hold)? It should produce an *ergodic* output process (Y, ν) for each ergodic input (X, μ). What if it's *nonanticipating* and with *finite memory*? That is, $\nu_x\{y_0 = b\}$ depends only on $(x_{-m} \dots x_{-1} x_0)$. This is *not enough.*

One way we can get a good channel is to impose an additional condition called *Nakumura ergodicity*: For all cylinder sets $U, V \subset X$ and $W, Z \subset Y$,

$$\frac{1}{n} \sum_{k=0}^{n-1} \int_{\sigma^k U \cap V} \left| \nu_x(\sigma^k W \cap Z) - \nu_x(\sigma^k W) \nu_x(Z) \right| \, d\mu(x) \to 0.$$

Alternatively, we can replace all three conditions with R. L. Adler's *output weakly mixing* (do *not* assume finite memory, nonanticipating): For all cylinder sets $F, F' \subset Y$ and for all $x \in X$,

$$\frac{1}{n} \sum_{k=0}^{n-1} \left| \nu_x(\sigma^k F \cap F') - \nu_x(\sigma^k F) \nu_x(F') \right| \to 0.$$

33.5. **Sliding Block Code Versions (Ornstein, Gray, Dobrushin, Kieffer).** Basic idea: approximate a block code $\phi : A^n \to B^n$ by a factor map $\psi : X \to Y$ (where $X = (A^{\mathbb{Z}}, \mu)$ and $Y = (B^{\mathbb{Z}}, \nu)$) by using Rokhlin towers: Given $\epsilon > 0$ and $n \in \mathbb{N}$, we can find $F \subset X$ with $F, \sigma F, \sigma^2 F, \dots, \sigma^{n-1} F$ disjoint and union having measure $> 1 - \epsilon$.

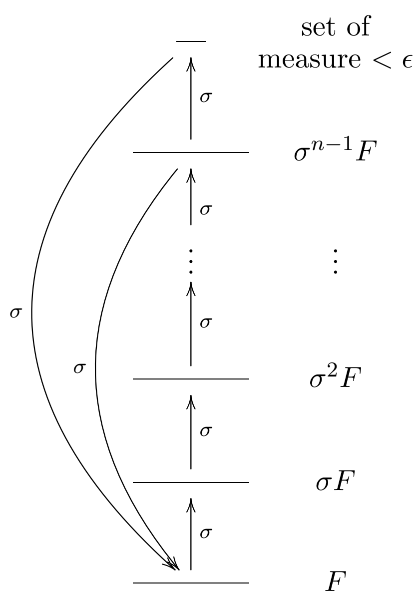

Let $x \in X$,

$$x = \dots x_r x_{r+1} \dots x_{r+n-1} \dots x_s x_{s+1} \dots x_{s+n-1} \dots,$$

where $\sigma^r x, \sigma^s x \in F$. Then define

$$\psi x = \dots 00 \dots 0\{\phi(x_r \dots x_{r+n-1})\}00 \dots 0\{\phi(x_s \dots x_{s+n-1})\} \dots.$$

ψ is shift-commuting, measurable, and approximates ϕ. The $00 \dots 0$ spacers have frequency $< \epsilon$. Then approximate ψ by a sliding block code by approximating F by cylinder sets. This construction leads to the following ergodic-theoretic channel coding theorems, which hold for different kinds of "good" channels.

Definition 33.1. A channel $[A, \{\nu_x\}, B]$ is

(1) *weakly continuous* if $\mu_n \to \mu$ weakly on $A^{\mathbb{Z}}$ (i.e., weak *, $\mu_n(E) \to \mu(E)$ for all cylinder sets E) implies that the corresponding input-output measures $\lambda_n \to \lambda$ (defined by μ_n and μ) weakly on $X \times Y = A^{\mathbb{Z}} \times B^{\mathbb{Z}}$;
(2) $\overline{d}$-*continuous* if

$$\sup_{\substack{E \text{ cylinder sets} \\ \text{of length } n \text{ in } X}} \sup_{x, x' \in E} \overline{d_n}(\nu_x^{(n)}, \nu_{x'}^{(n)}) \to 0 \text{ as } n \to \infty,$$

where

$$\overline{d_n}(\nu_1^{(n)}, \nu_2^{(n)}) = \inf_{\substack{\text{joinings } (u,v) \\ \text{(with measure } m) \\ \text{of } \nu_1 \text{ and } \nu_2}} \frac{1}{n} \sum_{k=0}^{n-1} m\{u_k \neq v_k\}.$$

Recall that a *joining* (or *coupling*) of two spaces, systems, or processes, is a system that factors onto both of them:

$$\begin{array}{ccc} & (Y_1 \times Y_2, m) & \\ {}^{\pi_{Y_1}}\swarrow & & \searrow^{\pi_{Y_2}} \\ (Y_1, \nu_1) & & (Y_2, \nu_2) \end{array} \qquad (u, v) \in Y_1 \times Y_2$$

(See Petersen's notes on ergodic theory, p.4).

Proposition 33.3. *Finite memory, nonanticipating* $\Rightarrow$ $\overline{d}$-*continuous* $\Rightarrow$ *weakly continuous.*

Theorem 33.4. *Suppose* $[A, \{\nu_x\}, B]$ *is a weakly continuous (stationary) channel.*

(1) *Then an ergodic source* $[A_0, \mu_0]$ *is block transmissible (in the above sense) over the channel if* $h(\mu_0) < C$ *and not block transmissible if* $h(\mu_0) > C$ *(the case* $h(\mu_0) = C$ *is not settled).*
(2) *It's sliding-block transmissible if and only if* $h(\mu_0) \leq C$.
(3) *If* $h(\mu_0) < C$, *then the source is* 0-error transmissible *across the channel: There exist measurable* $e : A_0^{\mathbb{Z}} \to A^{\mathbb{Z}}$ *and* $d : B^{\mathbb{Z}} \to A_0^{\mathbb{Z}}$ *such that* $\lambda\{(x, y) : e^{-1}x \neq dy\} = 0$.

33.6. **Further Topics.**

(1) "ergodic decompositions" of channels
(2) different definitions of capacity
(3) different kinds of channels (parallel, feedback, etc.)
(4) rate distortion theory (see McEliece)
(5) construction of codes

List of Figures