SYMBOLIC DYNAMICS AND SUBSHIFTS OF FINITE TYPE

Contents

1. Outline, Notation

To begin, here is a very brief and rough outline of what the course intends to cover. It will focus on shifts of finite type, with an eye toward the problem of classifying shifts of finite type up to topological conjugacy (usually called the Classification Problem).

Date: December 2018.

At forest's edge Basic things about subshifts, morphisms of symbolic systems, basic things about shifts of finite type including matrix presentations, Perron-Frobenius theorem, entropy, zeta functions.

Into the forest Strong shift equivalence, Williams' Theorem, shift equivalence, Williams' Problem. Dimension groups, and algebraic invariants.

The deep forest Wagoner's spaces, the triangle identities, positive cocycles, polynomial matrix land, the construction of the obstruction map Φ_n taking values in $K_2(\mathbb{Z}[t]/(t^n)$ and basic properties about K_2.

1.1. **Notation.** A pair (X, f) denotes a space together with a self-map $f\colon X \to X$, which is always assumed to be continuous. Almost always, we will also assume f is a homeomorphism. If $A \subset X$ is f-invariant, we let (A, f) denote $(A, f|_A)$. The f-orbit of a point $x \in X$ is the set $\{f^n(x) \mid n \in \mathbb{Z}\}$ and is denoted $\mathcal{O}_f(x)$.

By a map between systems $(X, f) \xrightarrow{h} (Y, g)$ we mean a continuous map $h\colon X \to Y$ such that $hf = gh$. A factor map $\pi\colon (X, f) \to (Y, g)$ is a map such that $\pi\colon X \to Y$ is surjective, and a topological conjugacy (usually just conjugacy) is a map $X, f) \xrightarrow{h} (Y, g)$ such that h is a homeomorphism.

2. Subshifts, SFT's, examples

Let $\mathcal{A}$ denote a finite collection of symbols (i.e. an alphabet). Let $\mathcal{A}^n$ denote the set of all strings of length n over $\mathcal{A}$, and let $\mathcal{A}^*$ denote the set of all finite strings formed from elements of $\mathcal{A}$. We denote the length of a word $w \in \mathcal{A}^*$ by $|w|$.

Formally, the space $\mathcal{A}^{\mathbb{Z}}$ consists of all functions $\mathbb{Z} \to \mathcal{A}$. However, we can identify $\mathcal{A}^{\mathbb{Z}}$ with $\prod\limits_{n\in\mathbb{Z}} \mathcal{A}$, which I find easier to think of, and consider elements $x \in \prod\limits_{n\in\mathbb{Z}} \mathcal{A}$ as bi-infinite strings of symbols from $\mathcal{A}$. I will let x_i denote the ith component of a point x (so $x_i \in \mathcal{A}$). Given $k \in \mathbb{Z}$ and $i \geq 0$, we let $x_{[k,k+i]}$ denote the finite string $x_k x_{k+1} \cdots x_{k+i}$ in $\mathcal{A}^*$.

With the product topology (using the discrete topology on $\mathcal{A}$), the space $\prod\limits_{n\in\mathbb{Z}} \mathcal{A}$ is compact. We can also make $\prod\limits_{n\in\mathbb{Z}} \mathcal{A}$ into a metric space by defining

$$d(x, y) = \inf\{\frac{1}{2^i} \mid i \text{ an integer such that } x_j = y_j \text{ for all } |j| \leq i\}$$

. The topology induced by this metric coincides with the product topology.

Exercise: Verify this is actually a metric on $\prod_{n\in\mathbb{Z}} \mathcal{A}$, and that the topology induced by it coincides with the product topology.

Given a word $w = a_0 a_1 \dots a_i \in \mathcal{A}^*$, there are associated *cylinder sets* $C_w^k \subset \prod_{n\in\mathbb{Z}} \mathcal{A}$ defined by

$$C_w^k = \{x \in \prod_{n\in\mathbb{Z}} \mathcal{A} \mid x_{[k,k+i]} = w\}.$$

I will often refer to *the* cylinder set associated to w, by which I will mean C_w^0, and denote this more simply by C_w.

The space $\prod_{n\in\mathbb{Z}} \mathcal{A}$ is homeomorphic to the Cantor set: it is compact, totally disconnected (the connected components are points), and perfect (has no isolated points).

Exercise: Show that the collection of cylinder sets form a basis for the topology on $\prod_{n\in\mathbb{Z}} \mathcal{A}$, and that $\prod_{n\in\mathbb{Z}} \mathcal{A}$ is a Cantor set.

The *shift map* $\sigma\colon \prod_{n\in\mathbb{Z}} \mathcal{A} \to \prod_{n\in\mathbb{Z}} \mathcal{A}$ is defined by $\sigma(x)_i = x_{i+1}$; one can think of σ as taking a bi-infinite string and shifting it left one unit. The map σ is a homeomorphism.

The fundamental object of study in symbolic dynamics is that of a *subshift*.

Definition 1. *A subshift consists of the pair (X, σ), where $X \subset \prod_{n\in\mathbb{Z}} \mathcal{A}$ is a closed σ-invariant subset.*

Sidequest: A system (X, f) is called *expansive* if there exists $\epsilon > 0$ such that, for $x \neq y$ in X, $\sup d(f^n x, f^n y) > \epsilon$. Show that a system (X, f) is topologically conjugate to a subshift if and only if X is totally disconnected and (X, f) is expansive.

Given a subshift $X \subset \prod_{n\in\mathbb{Z}} \mathcal{A}$, we say a word $w \in \mathcal{A}^*$ is X-*admissible* if $C_w \subset X$. We define the *language* of X, denoted $\mathcal{L}(X)$, to be the subset of $\mathcal{A}^*$ consisting of all X-admissible words, and let $\mathcal{L}_n(X) = \mathcal{L}(X) \cap \mathcal{A}^n$ denote the set of X-admissible words of length n.

A word w is X-*forbidden* if it is not contained in the language of X. One way to define a subshift is to stipulate the set of forbidden words. Given a set of words

$\mathcal{F} \subset \mathcal{A}^*$, define the subshift $X(\mathcal{F})$ to be the subset of $\prod_{n\in\mathbb{Z}} \mathcal{A}$ whose set of forbidden words is precisely $\mathcal{F}$.

Exercise: Show that a subset $X \subset \prod_{n\in\mathbb{Z}} \mathcal{A}$ is a subshift if and only if there exists a countable set $\mathcal{F}$ of forbidden words such that $X = X(\mathcal{F})$.

The case $X = \mathcal{A} = \{0, 1, \ldots, n-1\}$ coincides with $X(\emptyset)$, and is called the *full shift* (on n symbols).

Here is a rather cheap way to construct many examples of subshifts. Let $x \in \prod_{n\in\mathbb{Z}} \mathcal{A}$, and let $X(x) = \overline{\mathcal{O}_f(x)}$. One can check that this is σ-invariant, and obviously closed, so is a subshift. When the point x is *recurrent* (define), the orbit closure $X(x)$ is minimal (exercise).

Example 2. *Let $\mathcal{A} = \{0, 1\}$, and let $x \in \prod_{n\in\mathbb{Z}} \mathcal{A}$ denote the point constructed iteratively as follows: (define Thue-Morse point).*

A fundamental class of subshifts are the *subshifts of finite type*, obtained by restricting only a finite set of words.

Definition 3. *A subshift $X \subset \prod_{n\in\mathbb{Z}} \mathcal{A}$ is called a subshift of finite type (or usually just a shift of finite type) if there exists a finite set of words $\mathcal{F}$ such that $X = X(\mathcal{F})$.*

Some examples (and non-examples):

(1) The full shift is a shift of finite type, corresponding to $\mathcal{F} = \emptyset$.

(2) When $\mathcal{F} = \{11\}$, $X(\mathcal{F})$ is a shift of finite type called the golden mean shift.

(3) Let $\mathcal{F} = \{10^{2n+1}1 \mid n \geq 0\}$. The subshift $X(\mathcal{F})$ consists of all sequences such that between any two 1's there are an even number of 0's. $X(\mathcal{F})$ is *not* a shift of finite type.

Exercise: Show the even shift is not a shift of finite type.

3. Morphisms, higher block presentations

We naturally want to consider maps between subshifts. As one might guess, the maps should really be maps of pairs $(X_1, \sigma_1) \to (X_2, \sigma_2)$, i.e. maps of subshifts which intertwine the shift maps. Thus we define a morphism between subshifts $f\colon (X_1, \sigma_1) \to (X_2, \sigma_2)$ to be a continuous map $f\colon X_1 \to X_2$ such that $f\sigma_1 = \sigma_2 f$. The following result, due to Curtis, Hedlund, and Lyndon gives a symbolic interpretation of a morphism of subshifts.

Given alphabets $\mathcal{A}, \mathcal{B}$, an (m,n)-*block code* is a map $\alpha\colon \mathcal{A}^{m+n+1} \to \mathcal{B}$ for some $m, n \geq 0$. An r-block code is a map $\alpha\colon \mathcal{A}^{2r+1} \to \mathcal{B}$. A block code α induces a continuous map $f_\alpha\colon \mathcal{A}^{\mathbb{Z}} \to \mathcal{B}^{\mathbb{Z}}$ by defining

$$f_\alpha(x)_i = \alpha(x_{i-m} \dots x_i \dots x_{i+n}). \tag{1}$$

Note that f_α satisfies $f_\alpha \sigma = \sigma f_\alpha$. It follows that for any subshift $X \subset \prod_{n \in \mathbb{Z}} \mathcal{A}$ and block code α, the map $f_\alpha|_X$ is a morphism between X and its image. A continuous map f_α induced by a block code α is often called a *sliding block code.*

Theorem 4. *Any morphism of subshifts $f\colon (X, \sigma_X) \to (Y, \sigma_Y)$ is induced by a block code. In other words, given $f\colon (X, \sigma_X) \to (Y, \sigma_Y)$ there exists a block code $\alpha(f)$ such that $f = f_{\alpha(f)}$.*

Examples:

(1) Define a 2-block code on the 2-shift by $\alpha(x_0 x_1) = x_0 + x_1 \text{ mod} 2$. Then $f_\alpha\colon \{0,1\}^{\mathbb{Z}} \to \{0,1\}^{\mathbb{Z}}$ is onto, and hence a factor map, but not a conjugacy.

(2) Consider the 0-block code $\beta\colon \{0,1,2\} \to \{a,b\}$ defined by

$$\beta\colon \begin{cases} 0 \mapsto a \\ 1 \mapsto a \\ 2 \mapsto b. \end{cases}$$

The induced map $f_\beta\colon \{0,1,2\}^{\mathbb{Z}} \to \{a,b\}^{\mathbb{Z}}$ is a factor map taking the full 3-shift onto the full 2-shift.

(3) The 2-block code γ defined by $\gamma(00) = 1, \gamma(01) = 0, \gamma(11) = 0$ induces a continuous map $f_\gamma\colon X_{\text{golden mean}} \to \{0,1\}^{\mathbb{Z}}$.

Exercise: Show that the image of f_γ is the even shift defined in Example ??.

(4) Consider the shift of finite type Y defined on $\mathcal{A} = \{a,b,c,d,e,f\}$ whose set of *allowable* two-letter words is the following: $aa, ab, bc, cd, ce, cf, dc, ee, ef, fa, fb$.

There is a 1-block code

$$\delta\colon \begin{cases} a \mapsto 0 \\ b \mapsto 0 \\ c \mapsto 1 \\ d \mapsto 0 \\ e \mapsto 1 \\ f \mapsto 0 \end{cases}$$

inducing a continuous map $f_\delta\colon Y \to \{0,1\}^{\mathbb{Z}}$.

Exercise: Show that f_δ induces a topological conjugacy between Y and the full 2-shift, and find a block code inducing the inverse map.

The last example shows that, in general, it is far from clear when a given block code induces a topological conjugacy.

3.1. **Higher Block Presentations.** It is sometimes convenient to recode a shift using its admissible words of length n. To write this down, consider a subshift $X \subset \prod_{n\in\mathbb{Z}} \mathcal{A}$. For any $n \in \mathbb{N}$ there is a block code $\Phi_n\colon X \to \mathcal{L}_n(X)^{\mathbb{Z}}$ given by $\Phi(x)_i = x_i x_{i+1} \dots x_{i+n-1}$. The image of X under Φ_n is called the *n-block presentation* of X. By an *higher block presentation* of X we mean an n-block presentation of X for some $n \geq 1$.

Exercise: Show that for any subshift $X \subset \prod_{n\in\mathbb{Z}} \mathcal{A}$, a higher block presentation of X is conjugate to X.

Exercise: Find a presentation (with forbidden words) of the 2-block presentation of the golden mean shift.

4. Graph and matrix presentations of SFT's

We now focus on shifts of finite type.

Definition 5. *An SFT (X, σ_X) is called M-step if it can be presented using a set of forbidden words each of whose length is $M+1$.*

Of course, a 0-step SFT is just a full shift. The following gives our first characterization of when a subshift is a shift of finite type.

Proposition 6. *A subshift X is an M-step SFT if and only if it satisfies the following property: if $uv, vw \in \mathcal{L}(X)$ and $|v| \geq M$, then $uvw \in \mathcal{L}(X)$.*

Exercise: Prove the proposition.

Proposition 7. *If X and Y are conjugate subshifts and Y is an SFT then X is an SFT.*

Proof. Let $\alpha\colon X \to Y$ be a conjugacy induced by a block code Φ_α and let $\Phi_{\alpha^{-1}}$ be a block code inducing α^{-1}. We can assume both Φ_α and $\Phi_{\alpha^{-1}}$ have range r. The previous proposition gives $N \in \mathbb{N}$ such that if two Y-admissible words overlap by at least N then they can be glued along their overlap to form a new Y-admissible word. Then setting $M = N + 4r$, one can check that X satisfies the overlap gluing condition with length M. □

Exercise: Fill in the rest of the details for the proof.

Note that Example ?? shows that SFT's are *not* closed under factors. We'll revisit this problem later.

4.1. **Graph presentations.** Given an M-step SFT (X, σ_X), we can always recode using a higher block presentation to get a 2-step SFT over a possibly (and usually) larger alphabet. While this may seem like a rather useless procedure, there is an alternative way to carry this out, using labeled directed graphs, which is extremely profitable. In a way, this is the key step which underlies much of our investigation into SFT's and gives the connection between SFT's and algebra+positivity. In short, SFT's are associated to labeled directed graphs through the edge shift construction, and labeled directed graphs are linked with non-negative matrices through adjacency matrices.

4.2. **$\mathbb{Z}_+$-matrices, directed graphs, and edge shifts:** A finite directed graph Γ consists of a finite collection of vertices (often referred to as *states*) $V(\Gamma)$ together with a finite collection of edges $E(\Gamma)$. Each edge $e \in E(\Gamma)$ has an initial vertex and a terminal vertex. Multiple edges between two vertices are allowed, as are self loops. Often, we will assume our graph Γ also comes with a chosen labeling of the vertices.

Let Γ be a directed graph with k labeled vertices. Suppose there are $e(i,j)$ edges between vertex i and vertex j. Define the $k \times k$ *adjacency matrix* A_Γ by $A_{i,j} = e(i,j)$.

Exercise: Show that if Γ is a graph and A_1 and A_2 are adjacency matrices associated to Γ using two labelings of vertices, then A_1 and A_2 are conjugate via a permutation matrix.

On the other hand, given a $k \times k$ matrix A over $\mathbb{Z}_+$, the directed graph Γ_A associated to A has k vertices indexed by the rows (equivalently columns) of A and A_{ij} edges from vertex i to vertex j.

These constructions are inverses of each other, in the following sense:

$$\Gamma_{A_\Gamma} = \Gamma, \quad A_{\Gamma_A} = A.$$

picture

Given a graph Γ, we may associate an *edge shift* SFT $(X_\Gamma, \sigma_\Gamma)$ as follows. Label each edge in Γ, giving an alphabet $\mathcal{A}_\Gamma$, and define the 1-step SFT over $\mathcal{A}_\Gamma$

$$X_\Gamma = \{x \mid \text{ terminal vertex of } x_i = \text{ initial vertex of } x_{i+1}\}.$$

Like before, we have defined X_Γ by stipulating the set of admissible length two words.

Definition 8. *Given a matrix A over $\mathbb{Z}_+$ we denote by (X_A, σ_A) the edge shift of finite type associated to the graph $\Gamma(A)$.*

Examples:

(1) The matrix $A = (n)$ presents the full shift on n symbols.
(2) The graph associated to the matrix $B = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$ looks like:

Exercise: Show that, for B as in the last example, (X_B, σ_B) is simply the 2-block presentation of the golden mean shift, and hence is conjugate to it.

4.3. **Vertex Shift.** Suppose A over $\mathbb{Z}_+$ is actually over $ZO = \{0, 1\}$. Instead of the associated edge shift defined above, we could instead consider the *vertex shift* over the alphabet $Vert(\Gamma_A)$ of vertices of Γ_A

$$X_A^{vert} = \{x \in \prod_{n \in \mathbb{Z}} \mathcal{A} \mid \text{for all } i \text{ there is an edge from } x_i \text{ to } x_{i+1}\}.$$

For example, the vertex shift associated to the $k \times k$ matrix of all 1's presents the full k-shift. This example shows how much more economical $\mathbb{Z}_+$ presentations are than ZO presentations.

Exercise: Let A be a ZO matrix. Show that (X_A^{vert}, σ) and (X_A, σ_A) are conjugate.

Theorem 9. *For any SFT (X, σ_X) there is a graph $\Gamma(X)$ such that (X, σ_X) is conjugate to $(X_\Gamma, \sigma_\Gamma)$.*

Proof. If (X, σ_X) is k-step, define a graph Γ as follows: the vertex set is given by $B_k(X)$, the set of allowed k-blocks, and if $I = a_1 \dots a_k$, $J = b_1 \dots b_k$ are two vertices in Γ, then we draw an edge from I to J (labeled $a_1 \dots a_k b_k$) if $a_2 \dots a_k = b_1 \dots b_{k-1}$. Then a quick check shows that $(X_\Gamma, \sigma_\Gamma)$ is conjugate to the $k+1$-step presentation of (X, σ_X). □

We will work almost exclusively with edge shifts. The reason is the incredible utility provided by the association of the edge shift with its matrix. For example, we have the following.

:pathcount **Proposition 10.** *Let A be a $k \times k$ $\mathbb{Z}_+$-matrix, and let $m \geq 0$. The number of paths in $\Gamma(A)$ of length m from vertex I to vertex J is given by $A^m_{I,J}$.*

Proof. **Exercise.** □

Some definitions:

(1) **Graph side:**
   (a) A vertex i is stranded if either no edges leave i or no edges end at i. We call a graph *essential* if it has no stranded vertices.
   (b) A graph Γ is *irreducible* (or sometimes *strongly connected* if for every pair of vertices i, j there exists a path in Γ starting at i and ending at j.
(2) **Matrix side:**
(3) A $\mathbb{Z}_+$-matrix A is *essential* if it has no row or column consisting of all zeros.
(4) A $k \times k$ $\mathbb{Z}_+$-matrix A is *irreducible* if for any $i, j \leq k$ there exists $n(i, j)$ such that $A^n_{i,j} > 0$.

It is straightforward to check now that the a graph Γ is essential (resp. irreducible) if and only if the adjacency matrix $A(\Gamma)$ is essential (resp. irreducible).

Disclaimer: From here on out, unless otherwise stated, graphs are assumed to be essential, as are $\mathbb{Z}_+$-matrices.

5. Transitivity, Mixing, Irreducibility and Primitivity

Recall a system (X, f) is topologically transitive if for any pair of (non-empty) open sets $U, V \subset X$ there exists $n > 0$ such that $f^n(U) \cap V \neq \emptyset$.

Exercise: Show that if (X, f) is compact metric and topologically transitive, then it has a dense (forward) orbit. Show that, if one assumes X has no isolated points and

(X, f) has a dense forward orbit, then (X, f) is topologically transitive.

By the previous exercise, an SFT (X, σ_X) is topologically transitive if and only if it has a dense forward orbit. In the symbolic setting, often the term *irreducible* is used instead of topologically transitive.

Definition 11. *A subshift (X, σ_X) is* irreducible *if for any pair of words $u, v \in \mathcal{L}(X)$ there exists $w \in \mathcal{L}(X)$ such that $uwv \in \mathcal{L}(X)$.*

The following shows how a dynamical property like irreducibility corresponds to the corresponding property of the associated presenting matrix.

Proposition 12. *For an edge SFT (X_A, σ_A), the following are equivalent:*

(1) *(X_A, σ_A) is topologically transitive.*
(2) *(X_A, σ_A) is irreducible.*
(3) *A is irreducible.*

Proof. □

Exercise: Show for an irreducible SFT that the periodic points are dense.

Recall a system (X, f) is *topologically mixing* if for any pair of (non-empty) open sets $U, V \subset X$ there exists $N \in \mathbb{N}$ such that $f^n(U) \cap V \neq \emptyset$ for all $n \geq N$. In subshift land this condition takes the following form, and is often referred to as just *mixing*.

Definition 13. *A subshift (X, σ_X) is mixing if for every pair $u, v \in \mathcal{L}(X)$ there exists $N \in \mathbb{N}$ such that for each $n \geq N$ there exists a word $w \in \mathcal{L}_n(X)$ such that $uwv \in \mathcal{L}(X)$.*

The matrix condition corresponding to a mixing SFT is known as *primitivity*.

Definition 14. *A $\mathbb{Z}_+$-matrix A is primitive if there exists $n \geq 1$ such that $A^n > 0$.*

Proposition 15. *An edge SFT (X_A, σ_A) is mixing if and only if A is primitive.*

Proof. □

Often (though sometimes with some work), results about irreducible SFT's reduce to the mixing case. Irreducible systems have cyclic permuted mixing pieces and a notion of period governing the order of this cyclic structure (see LM). I won't discuss this much more, and will for the most part stay in the mixing case.

6. Interlude

So now we know what a shift of type is, and how each such shift can be presented (up to conjugacy) by a $\mathbb{Z}_+$-matrix. One of our main goals is to get a sense of what is known about the following question:

Classification Problem: When are two shifts of finite type topologically conjugate?

Toward this end, we'll start by surveying some useful conjugacy invariants. Apart from being useful for distinguishing SFT's, often the invariants are important in their own right: they contain important information about an SFT, and are very often computable. This last point is especially interesting; later on we will give a complete matrix invariant of the conjugacy class of an SFT, but its not clear whether it is more computable or not. A more precise case of the Classification Problem then could be stated as follows:

Decidability of conjugacy of SFT's: Is there an algorithm to determine when two SFT's are topologically conjugate?

We'll come back to this later. For now, we'll start with by showing how to compute an invariant, topological entropy, which is also of tremendous dynamical importance in its own right.

7. Topologicaly entropy of an SFT

For the general definition of topological entropy of a system (X, f), see ??. Exercise 6.3.8 in LM shows that, for subshifts, the following definition is equivalent.

Definition 16. *For a subshift (X, σ_X), the (topological) entropy is defined to be*

$$h(X) = \lim_{n\to\infty} \frac{1}{n} \log |\mathcal{L}_n(X)|.$$

Note that sometimes $\log_2$ is used instead of log. I guess for historical reasons. The definition is designed around the idea that the numbers $|\mathcal{L}_n(X)|$ grow like $e^{h \cdot n}$ for some h (h of course turns out to be the entropy).

Example: For the full shift on k symbols (X_k, σ_k) we have $|\mathcal{L}_n(X)| = k^n$, so $h(\sigma_k) = \log k$.

To compute the topological entropy of a shift of finite type, we'll need the Perron-Frobenius Theorem.

7.1. **Perron-Frobenius Theorem.**

The Perron-Frobenius Theorem is a key tool in the theory of non-negative matrices. The theorem has two parts to it: the primitive case (originally due to Perron), and the irreducible case (proved later by Frobenius). I'll give here a proof of the primitive case, which is used to prove the irreducible case. The proof here is based on one found by Brin in the 90's (though the idea of using Brouwer's goes back before that), and it could be very similar to older ones (I do not know).

Recall the spectral radius $\rho(A)$ of a matrix A is the maximum of the moduli of the eigenvalues of A.

Theorem 17 (Perron)**.** *Let A be a primitive matrix over $\mathbb{Z}_+$ with spectral radius λ. Then λ is a simple eigenvalue of A, and $\lambda > |\mu|$ for any other eigenvalue μ of A. Moreover, any eigenvector for λ has strictly positive values.*

Before beginning the proof, we need a quick lemma.

Lemma 18. *Let T be a linear map of a finite dimensional vector space V, and P be a polyhedron such that the origin lies in the interior of P. Suppose a power of T maps P into its interior. Then $\rho(T) < 1$.*

Proof. It's not hard to see that $\rho(T) \leq 1$. Suppose then that T has an eigenvector v with eigenvalue λ satisfying $|\lambda| = 1$. If $\lambda = \pm 1$ then a power of T fixes a point on the boundary of P, which is impossible, so we may suppose $\lambda \neq \pm 1$. Then there is a two-dimensional subspace W on which T acts by a rotation, and we can choose a point $x \in W$ which lies on the boundary of P. Since T acts by a rotation on W, x is a limit point of the sequence $T^n x$. But this is also impossible, since there is a power of T which maps P into its interior. □

Proof of the Theorem. Suppose A is $n \times n$, and consider the unit simplex

$$\Delta^n = \{x \in \mathbb{R}^{n+1} \mid \sum_i x_i = 1 \text{ and } x_i \geq 0 \text{ for all } i\}$$

in $\mathbb{R}^{n+1}$. Define a map $T_A \colon \Delta^n \to \Delta^n$ by $x \mapsto \frac{Ax}{||Ax||}$. Since Δ^n is a topological disk, fixed point theory (e.g. Brouwer's) implies the map T_A has a fixed point $v \in \Delta^n$. The vector v is then an eigenvector for A, and since A^k is positive for some k (by primitivity), v must be strictly positive. Letting λ denote the eigenvalue corresponding to v, it follows that λ is also positive.

Let $S = \operatorname{diag}(v)$ be the diagonal matrix whose i,ith entry is given by v_i (the ith entry of v), and consider

$$B = \frac{1}{\lambda} S^{-1} A S.$$

Note that B is also primitive, and that B has the vector of all 1's as an eigenvector with eigenvalue 1. Since λB and A are conjugate they have the same spectrum, so it suffices to show that 1 is a simple eigenvalue for B such that $|\mu| < 1$ for any other eigenvalue of B.

Consider $T_{B^{tr}} \colon \mathbb{R}^{n+1} \to \mathbb{R}^{n+1}$ given by $T_{B^{tr}} x = B^{tr} x$. The map $T_{B^{tr}}$ leaves Δ^n invariant, and so again using a fixed point theorem, must have a fixed point. Since B^{tr} is also primitive, a power of $T_{B^{tr}}$ actually maps Δ^n into its interior, so there must in fact be a fixed point, say w, in the interior of Δ^n. Consider the polyhedron $P = \Delta^n - w$ which contains the origin in its interior. A power of $T_{B^{tr}}$ maps P into its interior, so by the lemma the restriction of $T_{B^{tr}}$ to span$P$ has spectral radius less than one. Since span$P$ is a codimension one $T_{B^{tr}}$ subspace, this completes the proof. □

Theorem 19 (Frobenius)**.** *Suppose A be an irreducible $\mathbb{Z}_+$-matrix. Then A has a strictly positive eigenvector v with corresponding eigenvalue λ such that $\lambda = \rho(A)$. Any other nonnegative eigenvector of A is a multiple of v, and there exists $p \geq 1$ such that the set of eigenvalues of A whose modulus is λ is precisely $\{\lambda, \lambda\zeta^i, \dots, \lambda\zeta^{p-1}\}$ where ζ is a primitive pth root of unity.*

Proof. Use the cyclic structure and the primitive case. See ?? for details. □

For an irreducible or primitive matrix A, we let λ_A denote the Perron eigenvalue.

Now we can compute the entropy of an irreducible edge shift of finite type. Since any irreducible SFT is conjugate to an edge SFT, this lets us compute the entropy of any irreducible SFT. For the extension to the reducible case, see LM.

Theorem 20. *Let A be a $k \times k$ irreducible $\mathbb{Z}_+$-matrix. Then $h(\sigma_A) = \log \lambda_A$.*

Proof. By Proposition prop:pathcount 10, $\mathcal{L}_n(X_A) = \sum_{i,j=1}^k A_{i,,j}^n$, so we want to understand the growth of $\sum_{i,j=1}^k A_{i,,j}^n$. By Perron-Frobenius, A has a positive eigenvector v, and we let $a = \min\{v_i\}, b = \max\{v_i\}$. A simple calculation shows that

$$\sum_{j=1}^k A_{i,j}^n v_j = \lambda_A^n v_i.$$

Then for any $1 \leq i \leq k$ we have

$$a\lambda_A^n \leq v_i \lambda_A^n = \sum_{j=1}^k A_{i,j}^n v_j \leq b \sum_{j=1}^k A_{i,j}^n \leq b \sum_{i,j=1}^k A_{i,j}^n$$

so

$$\frac{a}{b} \lambda_A^n \leq \sum_{i,j=1}^k A_{i,j}^n.$$

Similarly, we have

$$a\sum_{j=1}^{k} A^n_{i,j} \leq \sum_{j=1}^{k} A^n_{i,j} v_j = \lambda_A^n v_i \leq b\lambda_A^n$$

and summing this over i gives

$$a\sum_{i,j=1}^{k} A^n_{i,j} \leq \sum_{i=1}^{k} b\lambda_A^n = kb\lambda_A^n$$

so that

$$\sum_{i,j=1}^{k} A^n_{i,j} \leq \frac{kb}{a}\lambda_A^n.$$

Altogether we have

$$\frac{a}{b}\lambda_A^n \leq \sum_{i,j=1}^{k} A^n_{i,j} \leq \frac{kb}{a}\lambda_A^n.$$

Applying log, dividing by n and taking $n \to \infty$ gives $h(\sigma_A) = \log \lambda_A$. □

Example: The golden mean shift is presented by the matrix $A = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$, so the entropy of the golden mean shift is $\log \phi$ where $\phi = \frac{1+\sqrt{5}}{2}$ is the golden mean.

For irreducible SFT's, the entropy also counts the asymptotic growth rate of the number of periodic points. This phenomenon occurs for more general 'irreducible' systems (i.e. Smale systems?). Recall for a system (X, f), $p_n(X)$ denotes the set of periodic points of least period n.

Theorem 21. *If (X, σ_X) is a subshift then*

$$\limsup_{n\to\infty} \frac{1}{n} \log |p_n(X)| \leq h(X).$$

If (X, σ_X) is a shift of finite type then

$$\limsup_{n\to\infty} \frac{1}{n} \log |p_n(X)| = h(X).$$

Proof. Exercise - see LM. □

8. Periodic Points and the zeta function

Exercise: Show that if (X, f) is expansive, then for $|P_n(f)|$ is finite for any $n \in \mathbb{N}$.

Definition 22. *Suppose (X, f) is a system such that $|P_n(f)|$ is finite for every $n \in \mathbb{N}$. The zeta function (sometimes called the Artin-Mazur zeta function) associated to (X, f) is defined by*

$$\zeta_f(t) = \exp\left(\sum_{n=1}^{\infty} \frac{|p_n(f)|}{n} t^n\right).$$

The zeta function, by definition, encodes the data of the periodic points of the system (X, f). Indeed, the sequence $|P_n(f)|$ and $\zeta_f(t)$ determine each other, as shown by the following exercise.

Exercise: Suppose $\zeta_f(t)$ has radius of convergence $R > 0$. Show that

$$\frac{d^n}{dt^n} \log \zeta_f(t)|_{t=0} = (n-1)!|P_n(f)|.$$

Why not encode the periodic data more simply, say with the generating function $\sum_{n=1}^{\infty} |p_n(f)| t^n$? There are advantages to using $\zeta_f(t)$, a few of which we'll see. For example, we have the Euler product formula.

Proposition 23 (Euler Product Formula)**.** *Let (X, f) be a system such that $|P_f|$ is finite for all $n \in \mathbb{N}$, and let $Q(f)$ denote the set of periodic orbits of (X, f). For a periodic orbit $\mathcal{O} \in Q(f)$ let $|\mathcal{O}|$ denote the length of the orbit. Then*

$$\zeta_f(t) = \prod_{\mathcal{O} \in Q} \left(1 - t^{|\mathcal{O}|}\right)^{-1}.$$

Proof. **Exercise.** □

For shifts of finite type, there is a remarkable formula for the zeta function. This formula, originally due to Bowen and Lanford, led to further results showing closed forms for the zeta functions of certain systems (see Manning). As one might guess, in the shift of finite type case, the heart of it is in using prop:pathcount 10.

For a matrix A over $\mathbb{Z}_+$, define polynomials

$$ch_A(t) = \det(tI - A)$$

$$\chi_A(t) = \det(I - tA).$$

Lemma 24. *Suppose A is a $\mathbb{Z}_+$-matrix and $ch_A(t)$ is degree d. Then*

$$t^d \cdot ch_A(t^{-1}) = \chi_A(t).$$

Furthermore, if $\{\lambda_i\}$ is the collection of non-zero eigenvalues of A, then

$$\chi_A(t) = \prod_i (1 - \lambda_i t).$$

Proof. **Exercise.** □

Theorem 25. *Let A be a $k \times k$ $\mathbb{Z}_+$-matrix and (X_A, σ_A) the associated shift of finite type. Then*

$$\zeta_{\sigma_A}(t) = \frac{1}{\chi_A(t)} = \frac{1}{\det(I - tA)}.$$

Proof. This again relies on pretty much just 10 and basic linear algebra. Let $\{\lambda_i\}_{i=1}^k$ denote the set of eigenvalues of A. We know by 10 that for any $n \geq 1$, $|P_n(\sigma_A)| = trA^n$. Since

$$trA^n = \sum_{i=1}^k \lambda_i^n$$

we have

$$\begin{aligned}
\zeta_{\sigma_A}(t) = \exp\left(\sum_{n=1}^\infty \frac{trA^n}{n} t^n\right) &= \exp\left(\sum_{n=1}^\infty \frac{\sum_{i=1}^k \lambda_i^n}{n} t^n\right) \\
= \exp\left(\sum_{i=1}^k \sum_{n=1}^\infty \frac{\lambda_i^n}{n} t^n\right) &= \prod_{i=1}^k \exp\left(\sum_{n=1}^\infty \frac{(\lambda_i t)^n}{n}\right) \\
= \prod_{i=1}^k \frac{1}{1 - \lambda_i t} &= \frac{1}{\chi_A(t)}
\end{aligned}$$

where the last equality comes from the lemma above. □

Given a matrix A, we denote the non-zero spectrum of A by $\mathrm{sp}^\times(A)$. So $\mathrm{sp}^\times(A)$ is the collection of non-zero eigenvalues of A.

Proposition 26. *Let A be a $\mathbb{Z}_+$-matrix. Any of the following items determines the other two.*

(1) $\{|P_n(\sigma_A)|\}_{n \geq 1}$.
(2) $\zeta_{\sigma_A}(t)$.
(3) $sp^\times(A)$.

Proof. Combine all the previous results. □

Note in particular something kind of interesting: in terms of periodic data, the SFT (X_A, σ_A) does not care about the zero spectrum of A.

Note also that for an SFT (X_A, σ_A), the zeta function $\zeta_{\sigma_A}(t)$ is an analytic function on the open disc of radius $\frac{1}{\rho(A)}$ centered at the origin.

Sidequest: What can the zeta function of a mixing SFT be? This is, by the above, very related to the question of what the spectrum of a primitive integral matrix can be. A set of necessary conditions were conjectured by Boyle and Handelman to be actually sufficient, becoming known as the Spectral Conjecture. The spectral conjecture was proved for many subrings of $\mathbb{R}$ by Boyle-Handelman in their Annals paper, and then later settled in the $\mathbb{Z}$ case by Kim-Ormes-Roush. Thus, the possible zeta functions are determined by this set of conditions, and can be found in the Kim-Ormes-Roush paper.

9. Strong shift equivalence and Williams' Theorem

Theorem 27. *Let A, B be $\mathbb{Z}_+$-matrices. The shifts of finite type (X_A, σ_A) and (X_B, σ_B) are topologically conjugate if and only if A and B are SSE-$\mathbb{Z}_+$.*

To prove this, we'll first do the "conjugacy $\implies A \underset{}{\overset{\text{SSE-}\mathbb{Z}_+}{\sim}} B$" direction, as it is harder then the "$A \overset{\text{SSE-}\mathbb{Z}_+}{\sim} B \implies$ conjugacy" direction.

Proof that conjugacy $\implies A \overset{\textbf{SSE-}\mathbb{Z}_+}{\sim} B$:

Proof.

Definition 28. *Given a graph Γ, an out-splitting of Γ is a graph Γ^{out} obtained roughly as follows:*

We say Γ_1 is an out-amalgamation of Γ_2 if Γ_2 is an out-splitting of Γ_1.
Analogously, there are in-splittings, and in-amalgamations.
A conjugacy $\phi\colon (X, \sigma_X) \to (Y, \sigma_Y)$ is a splitting (resp. amalgamation) if it is induced by either an out-splitting or an in-splitting (resp. an in-amalgamation or an in-amalgamation).

Lemma 29 (Decomposition Lemma)**.** *Any $\phi\colon (X_A, \sigma_A) \to (X_B, \sigma_B)$ is a composition of splittings and amalgamations.*

Proof. Sketch: first recode (X_A, σ_A) to a higher block presentation so that the induced ϕ' on this higher block presentation becomes a 0-block code. The conjugacies induced by passing to a higher block presentation are a composition of splittings, and the resulting 0-block code ϕ' is also a splitting. One now employs the following lemma.

Lemma 30. *Suppose* $\alpha\colon (X,\sigma_X) \to (Y,\sigma_Y)$ *is given by a 0-block code, and the inverse* α^{-1} *is given by a block code of range* (m,n) *with* $n \geq 1$. *Then there are outsplittings* $(X',\sigma_{X'}), (Y',\sigma_{Y'})$ *and a diagram*

$$\begin{tikzcd} (X,\sigma_X) \arrow[r, "\cong"] \arrow[d, "\alpha"] & (X',\sigma'_X) \arrow[d, "\beta"] \\ (Y,\sigma_Y) & (Y',\sigma'_Y) \arrow[l, "\cong"] \end{tikzcd} \tag{2}$$

such that β *is a 0-block code conjugacy whose inverse* β^{-1} *is a block code of range* $(m,n-1)$.

Proof. See Lemma 7.1.3 in LM. □

Continuing the proof, we repeatedly apply the lemma above to replace ϕ' with ϕ'_1 which is a 0-block code conjugacy whose inverse is a block code of range $(m,0)$. Then we apply the lemma, but now working on the dual graph to replace ϕ'_1 with ϕ'_2 which is a 0-block code conjugacy whose inverse is also a 0-block code, which finishes the proof. □

Lemma 31. *Suppose* $\Gamma_1 \to \Gamma_2$ *is an outsplitting. Then* $A(\Gamma_1) = RS, A(\Gamma_2) = SR$ *for some matrices* $\mathbb{Z}_+$*-matrices* R, S.

Proof. Consider the example given above. □

This finishes the proof of this direction. □

Proof that $A \overset{\textbf{SSE-}\mathbb{Z}_+}{\sim} B \implies$ **conjugacy:**
As promised, this direction is easier. It is enough to suppose $A = RS, B = SR$ for some R, S over $\mathbb{Z}_+$.

10. Shift equivalence and the dimension group

Consider the matrices $A = \begin{pmatrix} 1 & 2 \\ 2 & 3 \end{pmatrix}$ and $B = \begin{pmatrix} 1 & 1 \\ 4 & 3 \end{pmatrix}$. Are the SFT's (X_A, σ_A) and (X_B, σ_B) conjugate? A quick calculation shows that $\chi_A(t) = \chi_B(t)$, so $\zeta_{\sigma_A}(t) = \zeta_{\sigma_B}(t)$, and it follows then from the previous sections that they have the same entropy and periodic point counts. Nevertheless, they are not conjugate, and to see this, we'll use a new invariant which we introduce now.

This invariant has a couple of different forms, and we'll start by introducing two of these forms, and then showing that they are the same.

Definition 32. *Square matrices A, B over $\mathbb{Z}_+$ are* shift equivalent over $\mathbb{Z}_+$ *(SE-$\mathbb{Z}_+$), denoted $A \overset{\text{SE-}\mathbb{Z}_+}{\sim} B$, if there exist $\mathbb{Z}_+$ matrices R, S (not necessarily square) and $l \in \mathbb{N}$ such that all of the following hold:*

$$A^l = RS, \;\; B^l = SR$$
$$AR = RB, \;\; BS = SA.$$

Exercise: Show that SE-$\mathbb{Z}_+$ is an equivalence relation on square matrices over $\mathbb{Z}_+$.

Obviously matrices which are SSE-$\mathbb{Z}_+$ are SE-$\mathbb{Z}_+$, so we get the following.

Proposition 33. *If A, B are matrices over $\mathbb{Z}_+$ such that (X_A, σ_A) and (X_B, σ_B) are conjugate, then $A \overset{\text{SE-}\mathbb{Z}_+}{\sim} B$.*

There are a couple of immediate questions one might ask.

(1) If $A \overset{\text{SE-}\mathbb{Z}_+}{\sim} B$, what is the relationship between (X_A, σ_A) and (X_B, σ_B)?
(2) If $A \overset{\text{SE-}\mathbb{Z}_+}{\sim} B$, is $A \overset{\text{SSE-}\mathbb{Z}_+}{\sim} B$?
(3) Is SE-$\mathbb{Z}_+$ at least more understandable, and maybe computable, then SSE-$\mathbb{Z}_+$?
(4) Along the lines (3), is there a more 'familiar' algebraic interpretation of the SE-$\mathbb{Z}_+$-equivalence class of a matrix A?

We'll answer the first of these of these relatively soon, and the third we'll address in time. The fourth is our immediate goal. The second question was a long open problem in the field, which we restate for emphasis.

Williams' Problem: If A and B are square $\mathbb{Z}_+$ matrices and $A \overset{\text{SE-}\mathbb{Z}_+}{\sim} B$, are $A \overset{\text{SSE-}\mathbb{Z}_+}{\sim} B$?

The answer to this, it turns out, is **no**. We'll spend a fair amount of time exploring how this came to be answered.

To answer the first question above, we consider the following relation on systems.

Definition 34. *Systems (X, f) and (Y, g) are* eventually conjugate *if there exists some $N \in \mathbb{N}$ such that (X, f^n) and (Y, g^n) are conjugate for all $n \geq N$.*

Theorem 35. *For matrices A, B over $\mathbb{Z}_+$, the SFT's (X_A, σ_A) and (X_B, σ_B) are eventually conjugate if and only if $A \overset{\text{SE-}\mathbb{Z}_+}{\sim} B$.*

10.1. **Dimension Groups.** Knowing the SE-$\mathbb{Z}_+$-equivalence class of a matrix A turns out to be equivalent to knowing the isomorphism class of a certain ordered $\mathbb{Z}[t,t^{-1}]$-module built from A. This object, known as the *ordered dimension module* associated to A, plays an important role in the study of SFT's overall. We will see later that the underlying abelian group (called the *dimension group* built from A) plays the role of a homology group one can associate to the SFT (X_A, σ_A).

The dimension module has many incarnations. We'll begin with the matrix presentation, since it is usually most amenable to calculation. We note again two things:

(1) Our matrices act on row vectors
(2) We abuse notation and let A also denote the linear map associated to the matrix A (oh goodness).

Let A be $k \times k$ over $\mathbb{Z}$. The *eventual range* of A is

$$ER(A) = \cap_{i=1}^{\infty} \mathbb{Q}^k A^i.$$

This is an A-invariant rational subspace of $\mathbb{Q}^k$ on which A acts invertibly. We define the *dimension group* associated to A to be the additive abelian group

$$\mathcal{G}_A = \{v \in ER(A) \mid vA^k \in \mathbb{Z}^k \text{ for some } k \geq 1\}.$$

Since A acts invertibly on $ER(A)$, there is an automorphism (of abelian groups)

$$\begin{aligned} \delta_A &\colon \mathcal{G}_A \to \mathcal{G}_A \\ \delta_A &\colon v \mapsto vA. \end{aligned} \tag{3}$$

It is important not just to record the abelian group $\mathcal{G}_A$, but also the induced action of δ_A. We make $\mathcal{G}_A$ into a (right) $\mathbb{Z}[t,t^{-1}]$-module then by defining $v \cdot t = \delta_A^{-1}(v)$ (the reason for t acting by δ_A^{-1} and not δ_A will be more clear later on).

Definition 36. *The* dimension module *associated to A is defined to be the $\mathbb{Z}[t,t^{-1}]$-module $\mathcal{G}_A$.*

One can avoid the module language by just working with pairs $(\mathcal{G}_A, \delta_A)$, but it turns out to be useful (and more 'functorial') to instead work in the language of modules. In fact, this phenomenon of passing from endomorphisms (in this case the action of A) to an associated $\mathbb{Z}[t]$-module is part of a larger framework, which we'll explore later.

The matrices we are concerned with are not just integral, but over $\mathbb{Z}_+$. In this case, there is an additional bit of structure we can put on $\mathcal{G}_A$, which is actually quite important. Given A over $\mathbb{Z}_+$, the *positive cone* is the semigroup contained in $\mathcal{G}_A$ defined by

$$\mathcal{G}_A^+ = \{v \in \mathcal{G}_A \mid vA^k \in \mathbb{Z}_+^k \text{ for some } k \geq 1\}.$$

We call the pair $(\mathcal{G}_A, \mathcal{G}_A^+)$ the *ordered dimension module* associated to A. A morphism between ordered dimension modules $(\mathcal{G}_A, \mathcal{G}_A^+)$ and $(\mathcal{G}_B, \mathcal{G}_B^+)$ is a $\mathbb{Z}[t, t^{-1}]$-module map $\Psi \colon \mathcal{G}_A \to \mathcal{G}_B$ such that $\Psi(\mathcal{G}_A^+) \subset \mathcal{G}_B^+$. Note that δ_A is actually an automorphism of ordered dimension modules, since $\mathcal{G}_A^+$ is invariant under δ_A.

The connection between the dimension modules and shift equivalence is given by the following.

Theorem 37. *For any $\mathbb{Z}_+$-matrices A and B, each of the following hold:*

1. $A \underset{}{\overset{SE\text{-}\mathbb{Z}_+}{\sim}} B$ *if and only if the ordered dimension modules $(\mathcal{G}_A, \mathcal{G}_A^+)$ and $(\mathcal{G}_B, \mathcal{G}_B^+)$ are isomorphic.*
2. $A \overset{SE\text{-}\mathbb{Z}}{\sim} B$ *if and only if the dimension modules $\mathcal{G}_A$ and $\mathcal{G}_B$ are isomorphic.*

Proof. Suppose the SE-$\mathbb{Z}_+$ is implemented by R, S, and consider the map $ER(A) \to ER(B)$ given by $v \mapsto vR$. $\square$

10.2. **Shift equivalence and the Jordan form away from zero.** It is not hard to check now that if $A \overset{\text{SE-}\mathbb{Z}}{\sim} B$, then A and B have the same Jordan form (over $\mathbb{Q}$) away from zero, i.e. ignoring the set of zero eigenvalues (equivalently, consider their actions on their respective eventual ranges). An interesting question is, for a given Jordan form away from zero, how many SE-$\mathbb{Z}$ classes refine it?

Theorem 38 (Boyle)**.** *There are only finitely many SE-$\mathbb{Z}$ classes with a given Jordan form away from zero.*

It will follow from the results in the next section that the same holds for primitive matrices if one replaces $\mathbb{Z}$ in the above with $\mathbb{Z}_+$.

11. From SE-$\mathbb{Z}$ to SE-$\mathbb{Z}_+$

It is clear that $A \overset{\text{SE-}\mathbb{Z}_+}{\sim} B$ implies $A \overset{\text{SE-}\mathbb{Z}}{\sim} B$. For primitive matrices, it turns out the converse also holds. Put another way, in the primitive case, the positive cone data of $\mathcal{G}_A$ is redundant. This is an important result, and one of the first cases where primitivity is very useful.

Theorem 39. *Let A and B be primitive $\mathbb{Z}_+$-matrices. Then $A \overset{SE\text{-}\mathbb{Z}_+}{\sim} B$ if and only if $A \overset{SE\text{-}\mathbb{Z}}{\sim} B$.*

Before proving this, we'll record some lemmas. Both are actually interesting in their own right. The first describes the limiting behavior of A^k when A is primitive. Then, using this, we'll show that in the primitive case, the order structure of $\mathcal{G}_A$ is determined by the Perron subspace.

Lemma 40. *Suppose A is primitive, let u be a positive left PF-eigenvector and v be a positive right PF-eigenvector, chosen such that $u \cdot v = 1$. Then $\left((\frac{1}{\lambda_A})A\right)^n$ converges to the matrix vu.*

Proof. The matrix $(\frac{1}{\lambda_A})A$ fixes both u and v. If w is any (generalized) eigenvector for $\alpha \neq \lambda_A$, then $(\frac{1}{\lambda_A}A)^n w \to 0$, likewise for row vectors. It follows that $(\frac{1}{\lambda_A}A)^n$ converges to a matrix P with the property that $uP = u, Pv = v$, and any non-Perron generalized eigenvector of A is contained in the kernel of P. It is easy to check now that P must be a projection onto the subspace span$v$, and since $uP = u$ and $Pv = v$, the only such matrix is $P = vu$. □

Lemma 41. *Suppose A is primitive and v is a right Perron eigenvector for A. If w is any (row) vector, then wA^k is positive for some k if and only if $w \cdot v > 0$.*

Proof. One direction is clear, since v has strictly positive entries. So suppose $w \cdot v > 0$. By the previous lemma, we can write

$$A^k = \lambda_A^k vu + \lambda_A^k \cdot \text{error}$$

where error$\to 0$ as $k \to \infty$. Thus

$$wA^k = \lambda_A^k wvu + \lambda_A^k w \cdot \text{error} = \lambda_A^k (wvu + w \cdot \text{error}) > 0$$

for large enough k. □

Proof of Theorem ?? One direction is trivial. Suppose then that $A \overset{\text{SE-}\mathbb{Z}}{\sim} B$, using matrices R, S over $\mathbb{Z}$. Note that for any $k \geq 0$, RB^k and SA^k also enact a shift equivalence from A to B. We want to show that, for some $k \geq 0$, RB^k and SA^k are either both positive or both negative. Then use either RB^k, SA^k or $-RB^k, -SA^k$.

Suppose now that v is a right Perron eigenvector for A, so $Av = \lambda_A v$. Then

$$BSv = SAv = S\lambda_A v = \lambda_A Sv.$$

Since $A \overset{\text{SE-}\mathbb{Z}_+}{\sim} B$, we know from ?? that A and B have the same non-zero spectrum, so $\lambda_A = \lambda_B$. It follows then that Sv (or its negative) is a positive right Perron eigenvector for B. Now for each $1 \leq i \leq r$, we have

$$e_i Rw = e_i RSv = e_i A^l v = \lambda_A^l v_i > 0.$$

From Lemma ?? this implies, for each $1 \leq i \leq r$, there exists $k(i)$ such that $e_i RB^k$ is positive, so choosing k large enough, we have that RB^k is either eventually positive or negative. Likewise for SA^k. Finally, RB^k and SA^k must be the same sign, since

$$0 \leq A^{2k+l} = RSA^{2k} = RB^k SA^k.$$

□

We say two systems (X, f) and (Y, g) are *eventually conjugate* if there exists $N \geq 1$ such that (X, f^n) and (Y, g^n) are conjugate for all $n \geq N$.

Theorem 42. *Suppose A and B are $\mathbb{Z}_+$-matrices. Then $A \xrightarrow{\text{SE-}\mathbb{Z}_+} B$ if and only if the SFT's (X_A, σ_A) and $X_B, \sigma_B)$ are eventually conjugate.*

Proof. One direction is obvious: if $A \xrightarrow{\text{SE-}\mathbb{Z}_+} B$, then (X_A, σ_A) and (X_B, σ_B) are eventually conjugate. Suppose then that (X_A, σ_A^n) and (X_B, σ_B) are conjugate for all $n \geq N$. Then the dimension modules are isomorphic, and it is easy to check that $\mathcal{G}_{A^n}, \delta_A^n = \mathcal{G}_A, \delta_A^n$, so we really have an isomorphism

$$\mathcal{G}_A, \delta_A^n \xrightarrow{\Psi_n} \mathcal{G}_B, \delta_B^n.$$

Consider $\alpha_n = \Psi_n^{-1} \delta_B \Psi_n$. We know $\alpha_n^n = \delta_A^n$, and we would like $\alpha_n = \delta_A$ for *some* $n \geq N$. We will use the lemma:

Lemma: Suppose matrices E, F are non-singular and satisfy $E^n = F^n$. Suppose also that when λ is an eigenvalue of E and μ is an eigenvalue of F, and $\lambda^n = \mu^n$, then $\lambda = \mu$. Then $E = F$.

The proof of this lemma is a (fun) exercise.

Continuing the proof, choose m large enough such that for each pair of non-zero eigenvalues λ, μ of A and B, either $\lambda = \mu$ or $\lambda^m = \mu^m$ (this is also an **exercise**). Note $\text{spec} B = \text{spec} \alpha_n$. Then the lemma applied to $\delta_A^{\mathbb{Q}}$ and $\alpha_m^{\mathbb{Q}}$ imply $\delta_A^{\mathbb{Q}} = \alpha_m^{\mathbb{Q}}$. But this implies $\delta_A = \alpha_m$. □

11.1. **Bowen-Franks and the polynomial presentation of $\mathcal{G}_A$.** Still, are the matrices $A = \begin{pmatrix} 1 & 2 \\ 2 & 3 \end{pmatrix}$ and $B = \begin{pmatrix} 1 & 1 \\ 4 & 3 \end{pmatrix}$ shift equivalent over $\mathbb{Z}_+$? Since they are both primitive, it's enough to determine whether they are SE-$\mathbb{Z}$. It is possible to argue directly that they are not. However, let's use a different invariant to distinguish them.

Proposition 43. *Let A be a $k \times k$ $\mathbb{Z}_+$ matrix, and consider the map*

$$\begin{aligned} \mathbb{Z}[t]^k &\xrightarrow{I-tA} \mathbb{Z}[t]^k \\ v &\mapsto v \cdot (I - tA). \end{aligned} \tag{4}$$

The cokernel module $\text{coker}(I - tA)$ *is isomorphic (as a $\mathbb{Z}[t]$-module) to $\mathcal{G}_A$.*

Proof. Let $\delta_{A,\mathbb{Q}}$ denote the extension of δ_A to $ER(A)$. Consider the map of $\mathbb{Z}[t]$-modules

$$\pi\colon \mathbb{Z}[t]^k \to \mathcal{G}_A$$

defined by

$$w_0 + w_1 t + \cdots + w_n t^n = w \mapsto w_0 A^k \delta_A^{-k} + w_1 A^k \delta_A^{-k} \delta_A^{-1} + \cdots + w_n A^k \delta_A^{-k} \delta_A^{-n}. \tag{5}$$

It is easy to check that π lands in $\mathcal{G}_A$. Furthermore, note that if $w = v(I - tA)$, then $\pi(w) = 0$, so π descends to a map $\pi\colon \operatorname{coker}(I - tA) \to \mathcal{G}_A$. To see that π is onto, consider $v \in ER(A)$ such that $vA^j \in \mathbb{Z}^k$. Then $v = wA^k$ for some $w \in \mathbb{Q}^k$, and $wA^{k+j} \in \mathbb{Z}^k$. Then a small calculation shows that $wA^{k+j}t^j$ maps to $wA^{2k+j}\delta_A^{-k-j} = wA^k = v$.

To show that π is injective, it is enough to show that if $\pi(w) = 0$ then $w = u(I - tA)$ for some $u \in \mathbb{Z}[t]$. This is left as an exercise. □

The above proposition is part of a much larger idea, which is broadly contained in the following picture:

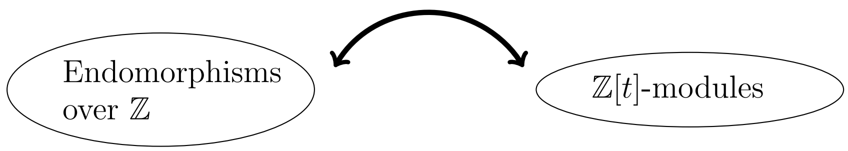

In fact, the place on the right should not be all of $\mathbb{Z}[t]$-modules, but only special ones. Which ones depends on how you want to study endomorphisms. For the study of shift equivalence, the right hand side should really be the set of S-torsion $\mathbb{Z}[t]$-modules of projection dimension ≤ 1, where S is the set of polynomials with constant term 1. We'll say more about this later.

What about the positive cone $\mathcal{G}_A^+$ in the polynomial setting? We'll explore later, but for now, we want to emphasize the following picture:

mgrouppres

$$0 \to \mathbb{Z}[t]^k \xrightarrow{I - tA} \mathbb{Z}[t]^k \to \operatorname{coker}(I - tA) \to 0. \tag{6}$$

11.2. **The Bowen-Franks group.** Finally, let's show that $A = \begin{pmatrix} 1 & 2 \\ 2 & 3 \end{pmatrix}$ and $B = \begin{pmatrix} 1 & 1 \\ 4 & 3 \end{pmatrix}$ are not SE-$\mathbb{Z}$.

Definition 44. *Given a $k \times k$ A over $\mathbb{Z}_+$, the* Bowen-Franks group *associated to A is the abelian group* $\operatorname{coker}(I - A)$*; that is, it is the cokernel of the map*

$$\mathbb{Z}^k \xrightarrow{I - A} \mathbb{Z}^k.$$

We won't go into much detail at the moment about why this group has a name. For now, we'll just mention the following beautiful theorem.

Theorem 45 (Parry-Sullivan,Bowen-Franks,Franks)**.** *Let A and B be irreducible $\mathbb{Z}_+$-matrices with spectral radii > 1. Then (X_A, σ_A) and (X_B, σ_B) are flow equivalent if and only if $BF(A) = BF(B)$ and det$(I - A)$ and det$(I - B)$ have the same sign.*

By definition, the Bowen-Franks group sits in the sequence

eqn:bfpres

$$\mathbb{Z}^n \xrightarrow{I-A} \mathbb{Z}^n \to BF(A). \tag{7}$$

eqn:dimgrouppres

The reason we write it this way is to contrast with the dimension module sequence (6). Using this point of view, the following becomes immediate.

Theorem 46. *Suppose A and B are $\mathbb{Z}_+$-matrices. If $A \overset{\text{SE-}\mathbb{Z}}{\rightsquigarrow} B$ then $BF(A) \cong BF(B)$.*

Proof. Set $t \mapsto 1$. ☐

Exercise: Consider $A = \begin{pmatrix} 5 & 3 \\ 3 & 5 \end{pmatrix}$ and $B = \begin{pmatrix} 6 & 8 \\ 1 & 4 \end{pmatrix}$. Show that A and B have the same Jordan form away from zero, and the same Bowen-Franks group, but are not shift equivalent. (Hint: Consider the $t \mapsto -1$ invariant.)

12. From SE-$\mathbb{Z}$ to SSE-$\mathbb{Z}$.

It is clear that SSE-$\mathbb{Z}$ implies SE-$\mathbb{Z}$. That the converse holds is originally due to Williams, with an additional key piece later provided by Effros. Our argument here uses the Effros idea, though presented in a different way. The presentation is meant to emphasize a key aspect of the relationship between SE and SSE in a general framework: the structure of nilpotent matrices over the ring in question. The argument here works more generally, but for the time being, we state it over $\mathbb{Z}$.

Definition 47. *Let A be over $\mathbb{Z}$.*

(1) *A* right zero extension *of A is a matrix of the form $\begin{pmatrix} A & U \\ 0 & 0 \end{pmatrix}$ where U is some rectangular matrix.*
(2) *A* left zero extension *of A is a matrix of the form $\begin{pmatrix} 0 & U \\ 0 & A \end{pmatrix}$ where U is some rectangular matrix.*

By a zero extension of A *we mean either a right or left zero extension.*

Definition 48. *Let A be over $\mathbb{Z}$.*

(1) *A* right nilpotent extension *of A is a matrix of the form $\begin{pmatrix} A & U \\ 0 & N \end{pmatrix}$ where U is some rectangular matrix and N is a square nilpotent matrix.*

(2) *A* left nilpotent extension *of A is a matrix of the form* $\begin{pmatrix} N & U \\ 0 & A \end{pmatrix}$ *where U is some rectangular matrix and N is a square nilpotent matrix.*

By a nilpotent extension of *A we mean either a right or left nilpotent extension.*

Recall two $n \times n$ matrices A, B over $\mathbb{Z}$ are said to be isomorphic if there exists $U \in GL_n(\mathbb{Z})$ such that $U^{-1}AU = B$. We write $A \overset{\text{iso-}\mathbb{Z}}{\sim} B$ when A and B are isomorphic, and note that this is an equivalence relation.

Proposition 49. (1) *Strong shift equivalence over $\mathbb{Z}$ is the equivalence relation generated by the following relations:*
 (a) $A \overset{\text{iso-}\mathbb{Z}}{\sim} B$.
 (b) $A \sim B$ *where B is any zero extension of A.*

(2) *Shift equivalence over $\mathbb{Z}$ is the equivalence relation generated by the following relations:*
 (a) $A \overset{\text{iso-}\mathbb{Z}}{\sim} B$.
 (b) $A \sim B$ *where B is any nilpotent extension of A.*

Proof. To be added. □

Lemma 50. *Let N be a nilpotent matrix over $\mathbb{Z}$. Then N is isomorphic over Z to a strictly upper triangular matrix.*

Theorem 51. *If* $A \overset{\text{SE-}\mathbb{Z}}{\sim} B$ *then* $A \overset{\text{SSE-}\mathbb{Z}}{\sim} B$.

Proof. By the proposition above, it is enough to show that A is strong shift equivalent over $\mathbb{Z}$ to any nilpotent extension of it. We'll show this for a right nilpotent extension, as the other case is analogous. Thus consider a right nilpotent extension $\begin{pmatrix} A & U \\ 0 & N \end{pmatrix}$ where N is nilpotent. Using the lemma above, choose V such that $VNV^{-1} = N_1$ is strictly upper triangular. Then, using $\begin{pmatrix} 1 & 0 \\ 0 & V \end{pmatrix}$, we have $\begin{pmatrix} A & U \\ 0 & N \end{pmatrix}$ is isomorphic over $\mathbb{Z}$ to $\begin{pmatrix} A & U' \\ 0 & N_1 \end{pmatrix}$. Since N_1 has zero on and below the diagonal, it is easy to check that $\begin{pmatrix} A & U' \\ 0 & N_1 \end{pmatrix}$ can be obtained from A by a sequence of right zero extensions. □

Interesting Problem: The spectral radius λ_A of A determines the entropy of the SFT (X_A, σ_A). What is a dynamical interpretation of the other eigenvalues of A? There is an idea of what this might be, related to signed measures, and signed traces, which will be mentioned later. To summarize, the other eigenvalues should give information regarding the error rates in the convergence to the entropy in its limit definition.