Lecture 1: An introduction to multidimensional shifts of finite type

Nishant Chandgotia

Hebrew University of Jerusalem

September

The full shift

Let $\mathbb{A}$ be a finite set.

We will give it the discrete topology and consider the product space $\mathbb{A}^{\mathbb{Z}^d}$, the set of all functions from $\mathbb{Z}^d$ to $\mathbb{A}$.

This space is compact. In fact, it is homeomorphic to a Cantor set. This is called the full shift.

An element of $\mathbb{A}^{\mathbb{Z}^d}$ is called a configuration and will usually be denoted by letters like $x, y, \ldots$. The values taken by configurations will be denoted by

$$x_{\vec{i}} := x(\vec{i}).$$

Intuition for the topology

In this topology, two points x and y are close if they agree on a large ball around the origin. A natural metric on the space is

$$d(x,y) := \max\{2^{-n} \;:\; x|_{\{-n,\dots,n-1,n\}^d} \neq y|_{\{-n,\dots,n-1,n\}^d}\}.$$

Given a finite set $B \subset \mathbb{Z}^d$ and a function $c : B \to \mathbb{A}$ let

$$[c]_B := \{x \in \mathbb{A}^{\mathbb{Z}^d} \;:\; x|_B = c\}.$$

Sets of the type $[c]_B$ are called cylinder sets and form the base of the topology.

Intuition for the topology

Thus a sequence x^j converges to $x \in \mathbb{A}^{\mathbb{Z}^d}$ if for all n there exists m such that for all $j > m$,

$$x^j|_{\{-n,\dots,n-1,n\}^d} = x|_{\{-n,\dots,n-1,n\}^d}.$$

In other words, the sequence x^j starts agreeing with x on larger and larger blocks.

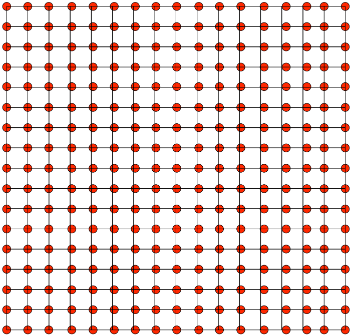

x^1

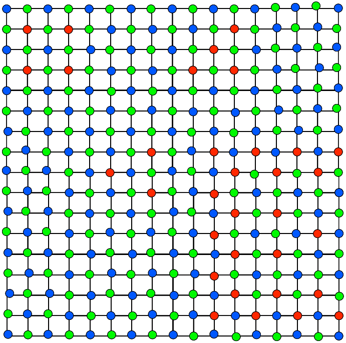

x

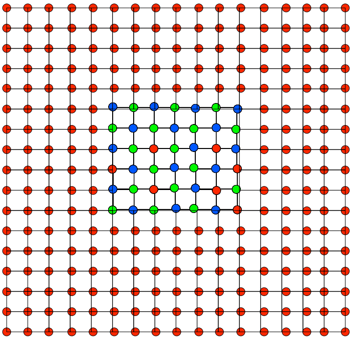

x^2

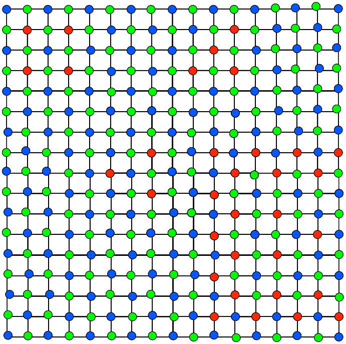

x

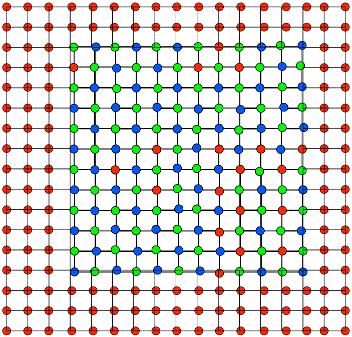

x^3

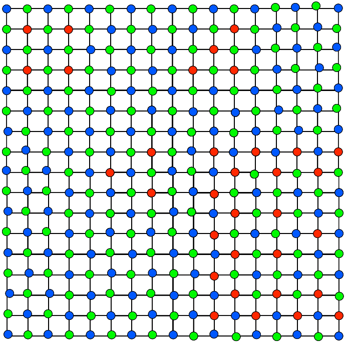

x

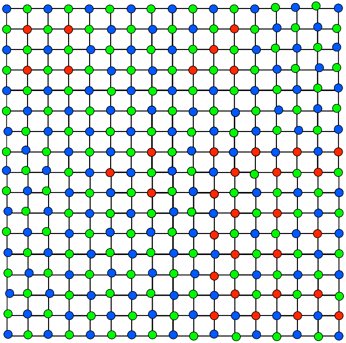

x^4

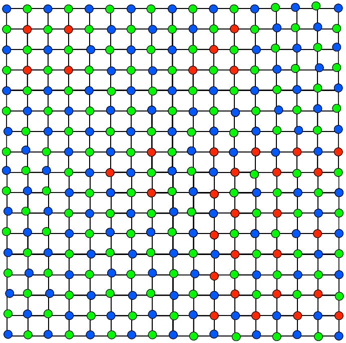

x

The shift map

$\mathbb{Z}^d$ acts naturally on the space $\mathbb{A}^{\mathbb{Z}^d}$ by homeomorphisms called shifts: $\sigma : \mathbb{Z}^d \times \mathbb{A}^{\mathbb{Z}^d} \to \mathbb{A}^{\mathbb{Z}^d}$ given by

$$\sigma(\vec{i}, x)_{\vec{j}} = x_{\vec{i}+\vec{j}}.$$

We will use the notation $\sigma^{\vec{i}}(x)$ to mean $\sigma(\vec{i}, x)$.

The pair $(\mathbb{A}^{\mathbb{Z}^d}, \sigma)$ is called a full shift. Since the action remains the same we will often drop it from the notation.

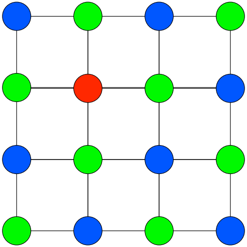

Figure : Shift action.

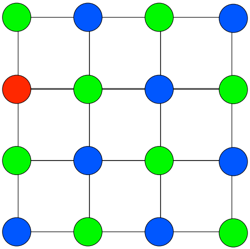

Figure : Moving left, $\sigma^{(1,0)}$

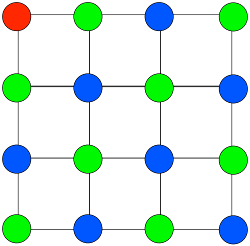

Figure : Moving up, $\sigma^{(0,-1)}$

Subshifts

A subshift is a closed subset $X \subset \mathbb{A}^{\mathbb{Z}^d}$ for some finite set $\mathbb{A}$ which is invariant under the shift action σ.

Before getting into examples let me introduce an equivalent definition for subshifts.

Given a finite set $B \subset \mathbb{Z}^d$ a pattern on B is a function $c : B \to \mathbb{A}$ while a configuration is an element of $\mathbb{A}^{\mathbb{Z}^d}$. Given a set of patterns $\mathcal{F}$ we define $X_{\mathcal{F}}$ to be the set of configurations which avoid translates of elements of $\mathcal{F}$.

$$X_{\mathcal{F}} := \{x \in \mathbb{A}^{\mathbb{Z}^d} \ : \ \sigma^{\vec{i}}(x)|_B \notin \mathcal{F} \text{ for all } \vec{i} \in \mathbb{Z}^d \text{ and finite sets } B \subset \mathbb{Z}^d\}.$$

Subshifts

$$X_{\mathcal{F}} := \{x \in \mathbb{A}^{\mathbb{Z}^d} \ : \ \sigma^{\vec{i}}(x)|_B \notin \mathcal{F} \text{ for all } \vec{i} \in \mathbb{Z}^d \text{ and finite sets } B \subset \mathbb{Z}^d\}.$$

Clearly if $x \in X_{\mathcal{F}}$ then $\sigma^{\vec{i}}(x) \in X_{\mathcal{F}}$ for all $\vec{i} \in \mathbb{Z}^d$.

Also if $x^j \in X_{\mathcal{F}}$ converges to $x \in \mathbb{A}^{\mathbb{Z}^d}$ then since larger and larger blocks of x eventually agree with x^j we must have that x also avoids patterns from $\mathcal{F}$ and is an element of $X_{\mathcal{F}}$.

Thus $X_{\mathcal{F}}$ is both shift-invariant and closed; it is a subshift.

An alternative definition of subshifts

Theorem

$X \subset \mathbb{A}^{\mathbb{Z}^d}$ is a subshift if and only if there exists a set of patterns $\mathcal{F}$ such that $X = X_{\mathcal{F}}$.

We have already seen that $X_{\mathcal{F}}$ are always subshifts. Now consider a subshift X and let $\mathcal{F}$ be the set of all patterns which do not appear in elements of X. Then $X = X_{\mathcal{F}}$:

Clearly $X \subset X_{\mathcal{F}}$ since $\mathcal{F}$ is precisely the set of patterns which do not appear in elements of X.

Conversely if $x \in X_{\mathcal{F}}$ then large blocks of it do not belong to $\mathcal{F}$ and hence must agree with some element of X. Since X is closed it follows that $x \in X$.

Thus subshifts arise as a set of configurations which are obtained by forbidding patterns.

Examples: Hard-core shift

Here $\mathbb{A} = \{0, 1\}$ and adjacent symbols can't both be one.

1	0	0	0	0
0	0	0	0	0
1	0	1	0	0
0	0	0	1	0
0	1	0	0	0

Examples: k-colorings

Here $\mathbb{A} = \{1, 2, 3, \ldots, k\}$ and adjacent symbols can't both be the same.

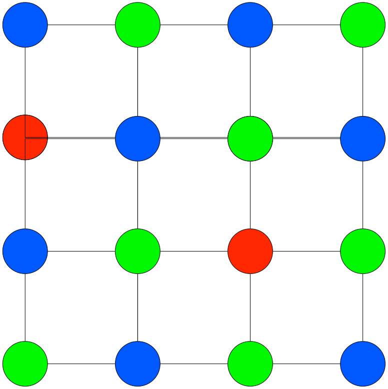

Figure : A 3-colouring

Examples: Dimer tilings

These are tilings of $\mathbb{Z}^d$ by dimers, that is, rectangular parallelepipeds of which one side is length two and the rest are one.

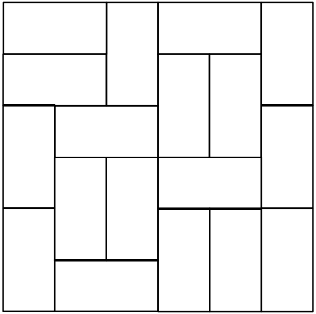

Why does this fit the frame work of subshifts?

Examples: Dimer tilings

It can be seen as a shift space by placing symbols on integer vectors which identify the kind of dimer covering them.

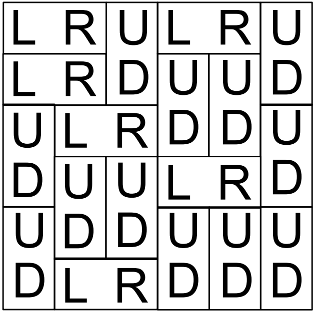

The symbols L and R identify the horizontal dimers while the symbols U and D identify the vertical ones. The forbidden list here ensures that L must appear to the left of R, U must appear above D etc.

Examples: The Even-shift

Fix $d = 1$. Here $\mathbb{A} := \{0, 1\}$. Here the connected components of 0's must be even.

$$\ldots 1001000010000001100 1 \ldots$$

In the previous examples, one could choose a finite forbidden list $\mathcal{F}$ such that $X = X_{\mathcal{F}}$.

This is not possible for the even shift. For the sake of contradiction assume that such a finite forbidden list exists. Then the domain of the patterns in $\mathcal{F}$ must be contained in $\{-n, \ldots, n\}$ for some n. But then such a list can't forbid

$$\ldots 0^{4n+1} 1 0^{4n+1} 1 0^{4n+1} \ldots$$

Subshift of finite type

A subshift X is called a subshift of finite type (SFT) if it can be obtained by forbidding finitely many patterns.

The hard core shift was obtained by forbidding adjacent appearances of 1s.

k-colorings were obtained by forbidding adjacent appearances of the same symbol.

Dimer tilings can be obtained by ensuring that each symbol has exactly one more adjacent symbol which forms a dimer with it.

These are all SFTs.

The even shift is not an SFT.

Motivation

We will be primarily interested in the study of SFTs when $d > 1$. We will now discuss the main motivations for studying these objects:

1. Statistical physics models
2. Questions arising in theoretical computer science
3. As a higher-dimensional analogue of one dimensional SFTs.

Let us discuss these one by one.

Statistical physics models

Ising model

In the 1920's, Ising investigated a simple model which was supposed to exhibit properties of a magnet. The model is the following:

$2n + 1$ atoms would be placed in a line and each one could take a random spin value $x_i = 1$ or -1. Fix $\beta > 0$ (called the inverse temperature)

$$\mathbb{P}(x_{-n}, x_{-n+1}, \dots, x_n) \propto \prod_{i=-n}^{n-1} e^{\beta x_i x_{i+1}}.$$

Under this distribution there is a penalty for x_i being different from x_{i+1}.

Ising model

Ising wanted to know whether fixing x_{-n} and x_n significantly influences x_0 as n goes to infinity.

Much to Ising's dismay, the spin values on the boundary did not influence the central spin and he concluded that no such influence existed in higher dimensions either (as published in his thesis 1924).

Ising model

However very soon it was discovered by Peierl in 1936 that the two dimensional analogue of this model did exhibit this influence. In the model, atoms would be arranged in the square lattice in a box B and their random spins $\sigma_{\vec{i}}$ would be distributed as

$$\mathbb{P}(x_{\vec{i}}; \vec{i} \in B) \propto \prod_{(\vec{i},\vec{j})\text{adjacent in the box}} e^{\beta x_{\vec{i}} x_{\vec{j}}}.$$

For certain values of the parameter β it was discovered that fixing the spins on the boundary of the box had significant influence on the spin at the centre even as the box increased to $\mathbb{Z}^d$.

Ising model

The precise value at which the model changed its behaviour was calculated by Onsager in 1944.

Such behaviour can also be captured by SFTs where you have two finite collection of symbols A and B of equal cardinality and we make a forbidden list such that transitioning from A to B or vice versa can be made comparatively harder as compared to staying among the same set of symbols.

For instance $A := \{-m, -m+1, \ldots, -1\}$ and $B := \{1, 2, \ldots, m\}$ and the only constraint is that adjacent symbols must have the same sign unless they are 1 and -1. This is a famous example studied by Burton and Steif. We will return to this in the third lecture.

-2	-3	-1	-1	-2	-1	1	1	-1	-2	-2
-1	-1	-2	-2	-1	1	4	1	-1	-2	-2
-5	-2	-6	-1	1	5	3	1	-1	-2	-2
-1	-2	-1	-1	1	3	6	1	-1	-2	-2
1	-1	1	1	2	5	1	-1	-2	-2	-2
2	1	2	2	3	1	-1	-2	-2	-2	-2
5	6	7	3	4	1	-1	-2	-2	-2	-2
1	1	1	4	5	1	-1	-2	-2	-2	-2
-1	-1	-1	1	1	-1	-2	-2	-2	-2	-2
-2	-2	-2	-1	-1	-2	-2	-2	-2	-2	-2
-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2

Many such models like the hard-core model and the Potts model have their analogues in SFTs.

The study of their statistical behaviour is equivalent to studying the measures of maximal entropy (or more generally equilibrium states) of the corresponding SFTs.

This has been a very rich interface of work and we will introduce some of these terms in the second lecture/exercise session and discuss in the third and the fourth lecture.

Questions arising in theoretical computer science

An SFT can be described by its finite forbidden list $\mathcal{F}$ but it is not obvious that using the forbidden list we can always find out the properties of the corresponding SFT.

Theorem

It is undecidable whether an SFT is non-empty. In other words, there is no algorithm which will take the finite forbidden list $\mathcal{F}$ as input and tell us whether or not $X_{\mathcal{F}}$ is non-empty.

This was proved by Robert Berger in 1966 answering a question of his advisor Hao Wang arising from logic.

There is however an algorithm which will stop if indeed $X_{\mathcal{F}}$ is empty.

The algorithm will keep trying to build patterns on larger and larger shapes. If indeed $X_{\mathcal{F}}$ is empty there will be a shape on which no valid pattern avoiding the forbidden list.

However if $X_{\mathcal{F}}$ is non-empty then we need to check this for all (and infinitely many) shapes that there is a valid patten on the shape avoiding it.

In specific cases it might be easy to check: For instance if we manage to build a periodic pattern.

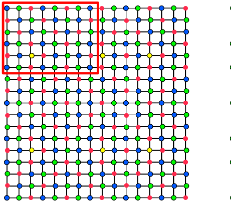

Figure : Once we have the pattern in the red square we can tile the entire plane with it

One can use this to build an algorithm to check whether a given SFT is non-empty in one dimension since they always have periodic points (exercise).

But there are SFTs in higher dimensions which have no periodic pattern.

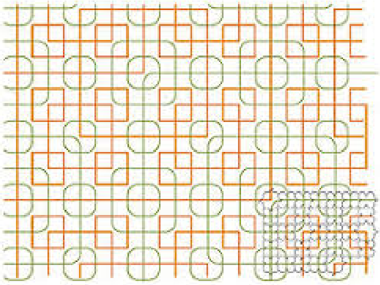

Figure : This picture is from a beautiful construction by Robinson (1972)

Let us look at another very important aspect.

Topological entropy

If we know that the subshift is non-empty then we can try and measure the growth rate of the number of allowed patterns. This gives rise to the most important invariant attached to subshifts called the topological entropy. Let $B_n := \{1, 2, \dots, n\}^d$.

Given a subshift X and $B \subset \mathbb{Z}^d$ let

$$\mathcal{L}(X, B) := \{x|_B \ : \ x \in X\}.$$

The topological entropy of X is given by

$$h_{\text{top}}(X) := \lim_{n \to \infty} \frac{1}{|B_n|} \log(|\mathcal{L}(X, B_n)|).$$

Topological entropy can be defined in a more general fashion for actions by homeomorphisms of the group $\mathbb{Z}^d$ (in fact amenable groups) on compact metric spaces using open covers but we do not need this generality in this series.

Topological entropy

$\mathcal{L}(X, B_n) := \{x|_{B_n} \ : \ x \in X\}$.
$h_{\text{top}}(X) := \lim_{n\to\infty} \frac{1}{|B_n|} \log(|\mathcal{L}(X, B_n)|)$.

Clearly if $a \in \mathcal{L}(X, B_{kn}))$ then $a|_{\vec{i}+B_n} \in \mathcal{L}(X, \vec{i} + B_n)$ for all $\vec{i}$.
The set B_{kn} can be tiled by k^d copies of B_n. Hence

$$|\mathcal{L}(X, B_{kn})| \leq |\mathcal{L}(X, B_n)|^{k^d}.$$

Using this and subadditivity arguments one can prove that the limit in h_{top} exists and is equal to

$$\inf_n \frac{1}{|B_n|} \log(|\mathcal{L}(X, B_n)|).$$

Given a finite set $\mathcal{F}$ can we find $h_{\mathrm{top}}(X_{\mathcal{F}})$?

Given a finite set $\mathcal{F}$ can we find $h_{\text{top}}(X_{\mathcal{F}})$?

Depends on what we mean by find.

Topological entropy

Theorem (Lind (1984))

The set of entropies of SFTs when $d = 1$ *are precisely the logarithm of the Perron numbers, that is,* $\log \alpha$ *where* α *is an algebraic integer strictly greater than the modulus of its Galois conjugates.*

Topological entropy

Theorem (Lind (1984))

The set of entropies of SFTs when $d = 1$ *are precisely the logarithm of the Perron numbers, that is,* $\log \alpha$ *where* α *is an algebraic integer strictly greater than the modulus of its Galois conjugates.*

The situation is drastically different in higher dimensions.

Theorem (Hochman and Meyerovitch (2010))

The set of entropies of SFTs when $d \geq 2$ *are precisely the non-negative right recursively enumerable numbers, that is, numbers for which there exists algorithms approximating it from above.*

Thus when $d \geq 2$, in general, given any algorithm which tries to approximate the entropy will fail to say how close it is to the actual value. This and many further results linked theoretical computer science with the study of SFTs.

As an analogue of one dimensional SFTs

The suggestion here is that one dimensional SFTs are much better behaved and we try to prove in higher dimensions results analogous to those in one dimension.

One dimensional SFTs

The study of one dimensional SFTs started with applications in data recording and as means of encoding complicated dynamics in a more combinatorial way. One of the first examples of the latter was due to Adler and Weiss (1970).

Consider the automorphism of the torus $\begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix} : \mathbb{R}^2/\mathbb{Z}^2 \to \mathbb{R}^2/\mathbb{Z}^2$. The torus can be partitioned into three parts, Red, Blue and Green such that the following holds:

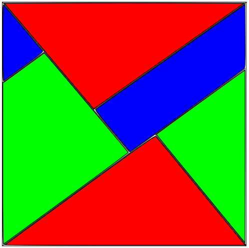

Encoding of a toral automorphism

The torus can be partitioned into three parts, Red, Blue and Green such that the following holds:

Given each point $x \in \mathbb{R}^2/\mathbb{Z}^2$ we get a sequence of partitions visited by $\left(\begin{smallmatrix} 1 & 1 \\ 1 & 0 \end{smallmatrix}\right)^i (x)$ as i varies over $\mathbb{Z}$.

This map from points of the torus to sequences in Red, Blue and Green is essentially injective.

One dimensional SFTs

It gives rise to sequences that follow the constraints given by the edges absent in the following graph

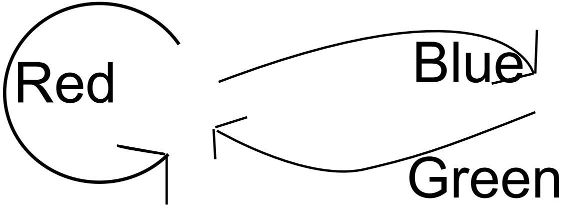

Such encodings are extremely useful and an important area of study where nowadays the alphabet considered is often countably infinite.

Sofic shifts

We leave the discussion with some important open questions of the field which intrigue me but we will not get to discuss.

A subshift Y is a factor of a subshift X if there exists a continuous surjective map $\phi : X \to Y$ which commutes with the shift map, that is,

$$\sigma^{\vec{i}} \circ \phi = \phi \circ \sigma^{\vec{i}}.$$

The factors of SFTs are called sofic shifts.

Even shifts are sofic

Sequences arising by

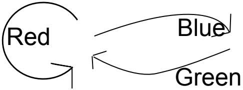

gives us an SFT while

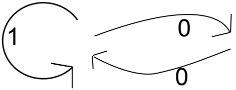

gives us the even shift proving that it is sofic. One can check that they have the same entropy.

Sofic shifts conjecture

It is not difficult to see that if Y is a factor of X then $h_{\text{top}}(Y) \leq h_{\text{top}}(X)$. (exercise)

Question (Benjamin Weiss)

Is every sofic shift Y a factor of an SFT X with the same entropy?

Here are some partial results.

1. Yes when $d = 1$. (Coven and Paul 1975)
2. For $d > 1$, for all $\epsilon > 0$ there exists a sofic shift X such that $h_{\text{top}}(X) < h_{\text{top}}(Y) - \epsilon$. (Desai 2006)
3. The set of entropies for sofic shifts is the same as the set of entropies for SFTs. (Hochman and Meyerovitch 2010)

Nivat's conjecture

Fix $d = 2$. Let X be a subshift such that for some $n, m \in \mathbb{N}$

$$|\mathcal{L}(X, \{1, 2, \ldots, n\} \times \{1, 2, \ldots, m\})| \leq nm.$$

Prove all configurations in X are periodic.

For $d = 1$ it is a famous result by Morse and Hedlund.

Recently, Kari and Moutout (2019) proved that X must contain a periodic configuration.

Periodic Tiling conjecture

Let $X(F)$ be the set of tilings of $\mathbb{Z}^d$ by a finite set $F \subset \mathbb{Z}^d$. Prove that if $X(F)$ is non-empty then $X(F)$ must contain periodic points.

For $d = 1$ this follows from ideas by Morse and Hedlund.

For $d = 2$ this was resolved recently by Bhattacharya (2016). This is wide open in higher dimensions.

References

1. For the Ising model in two dimensions you can look at Chapter 6 of “Gibbs Measures and Phase Transitions” by Georgii.
2. The paper by Burton and Steif paper referred to is “Non-uniqueness of measures of maximal entropy for subshifts of finite type”.
3. For a good introduction to the undecidability of tiling problem look at “Undecidability and Nonperiodicity for Tilings of the Plane” by Robinson.
4. The paper by Lind being referred to is “The entropies of topological Markov shifts and a related class of algebraic integers.”
5. The paper by Hochman and Meyerovitch referred to is “A Characterization of the Entropies of Multidimensional Shifts of Finite Type”.

References

1. The construction by Adler and Weiss can be found in "Similarity of automorphisms of the torus."(1970)
2. Coven and Paul's paper is titled "Sofic systems"
3. The paper by Desai is titled "Subsystem entropy for $\mathbb{Z}^d$ sofic shifts."
4. Kari and Moutout's paper being referred to is titled "Decidability and Periodicity of Low Complexity Tilings". Also look at the paper by Kari titled "Low-Complexity Tilings of the Plane"
5. Bhattacharya's paper is titled "Periodicity and decidability of tilings of $\mathbb{Z}^2$"