Multidimensional constant-length substitution sequences

Natalie Priebe Frank

Vassar College, Department of Mathematics, Box 248, Poughkeepsie, NY 12604, USA

Received 30 September 2003

Abstract

We consider multidimensional substitutions of constant length in a primarily expository setting, explaining how results from both symbolic dynamics and tiling dynamical systems can be applied. We focus in particular on ergodic and spectral theoretic concepts in an analysis that includes results and proofs extending what is known to our case. Tools such as the frequency measure, spectral measures, and the multidimensional odometer are used. We investigate several examples, among them the class of bijective substitutions. Bijective substitutions are of particular interest due to their mixed dynamical spectrum and because they are skew products over multidimensional odometers. For these, a condition is given allowing a full decomposition of the spectral measures.
© 2004 Elsevier B.V. All rights reserved.

MSC: primary 37B50, secondary 37A30

Keywords: Self-similar tilings; Substitution sequences; Dynamical spectrum

E-mail address: nafrank@vassar.edu (N.P. Frank).

1. Substitution dynamical systems

1.1. Sequences, the tiling model, and shift dynamical systems

Let $\mathcal{A}$ be a finite set (*alphabet*) and let a *sequence in* $\mathbb{Z}^d$ be a function $\mathcal{T}: \mathbb{Z}^d \to \mathcal{A}$. The collection of all such sequences is $\mathcal{A}^{\mathbb{Z}^d}$; the action of *translation* by an element $\vec{j} \in \mathbb{Z}^d$ on $\mathcal{T} \in \mathcal{A}^{\mathbb{Z}^d}$ yields the sequence $\mathcal{T} - \vec{j}$ whose $\vec{k}$th element is

$$(\mathcal{T} - \vec{j})(\vec{k}) = \mathcal{T}(\vec{k} + \vec{j}). \tag{1}$$

This is also called a *shift* because it shifts the sequence so that the element that used to be at $\vec{j}$ is now at $\vec{0}$.

We will write elements of $\mathbb{Z}^d$ as $\vec{j} = (j_1, j_2, \ldots, j_d)$; the sup norm of such a vector will be written $|\vec{j}| = \max_{i \in 1,2,\ldots,d} |j_i|$. A block B is a map from some finite subset of $\mathbb{Z}^d$ into $\mathcal{A}$; a subblock of $\mathcal{T}$ is the mapping $\mathcal{T}$ restricted to a finite subset. In this paper, we will in general assume that our sequences are *almost periodic*: there is a radius $R(B) > 0$ such that for any $\vec{j} \in \mathbb{Z}^d$, a copy of the block B appears in $\mathcal{T}$ somewhere within the set $\{\vec{j} + \vec{k}: |\vec{k}| \leqslant R(B)\}$. (We refer to a copy of the block B since any block has a fixed domain in $\mathbb{Z}^d$ that is usually not in the set $\{\vec{j} + \vec{k}: |\vec{k}| \leqslant R(B)\}$.) We will also assume in general that our sequences are *nonperiodic*: $\mathcal{T} - \vec{j} = \mathcal{T}$ if and only if $\vec{j} = \vec{0}$.

Sequences in $\mathbb{Z}^d$ may be thought of as *tilings of* $\mathbb{R}^d$ by considering $\mathcal{T}$ to be composed of unit cubic tiles that are colored or labeled by elements of $\mathcal{A}$. Everything appearing in this paper can be framed in this more general environment, and we will use fundamental results from papers such as [17,21,22] and references therein for this perspective. In addition we refer to the sources [3,16] and references therein for the discrete dynamical system perspective.

For any sequences $\mathcal{T}$ and $\mathcal{T}' \in \mathcal{A}^{\mathbb{Z}^d}$ with $\mathcal{T} \neq \mathcal{T}'$ we write $N(\mathcal{T}, \mathcal{T}') = \inf\{n \geqslant 0$ such that $\mathcal{T}(\vec{j}) \neq \mathcal{T}'(\vec{j})$ for some $|\vec{j}| = n\}$, and we define the metric $d(\mathcal{T}, \mathcal{T}') = \exp(-N(\mathcal{T}, \mathcal{T}'))$. For $\mathcal{T} = \mathcal{T}'$ the distance is defined to be zero. This metric is similar to those seen in [16,17], and since this metric yields the product topology we find that $\mathcal{A}^{\mathbb{Z}^d}$ is compact. Considering all blocks B, the metric topology has a basis given by the *cylinder sets* $[B] = \{\mathcal{T} \in \mathcal{A}^{\mathbb{Z}^d}$ such that B is a subblock of $\mathcal{T}\}$. We denote the Borel σ-algebra corresponding to this topology by $\mathcal{B}$.

Fixing an almost periodic, nonperiodic $\mathcal{T}_0 \in \mathcal{A}^{\mathbb{Z}^d}$, we define the sequence space

$$X = X_{\mathcal{T}_0} = \overline{\{\mathcal{T}_0 - \vec{j} \text{ such that } \vec{j} \in \mathbb{Z}^d\}}. \tag{2}$$

This space is invariant under the action of translation and is compact because it is closed. The set X along with the action of translation by elements of $\mathbb{Z}^d$ compose the dynamical system $(X, \mathbb{Z}^d)$. Putting a translation-invariant Borel probability measure μ on $\mathcal{B}$, we have the measure-theoretic dynamical system $(X, \mathbb{Z}^d, \mu)$, which is sometimes called a *shift dynamical system* or a *subshift of* $\mathcal{A}^{\mathbb{Z}^d}$.

1.2. Substitution sequences in $\mathbb{Z}^d$

These generalize one-dimensional substitutions sequences where the substitution is constant length. They can also be seen as self-similar tilings of $\mathbb{R}^d$, with particularly simple tile

geometry. The general notion is that every letter in $\mathcal{A}$ is assigned a replacement rule that is a d-dimensional 'rectangular' array of letters, and iterated application of the replacement rule results in an infinite sequence.

Fix a *dimension d* and *lengths* $l_1, l_2, \ldots, l_d$, positive integers with each $l_i > 1$. The *location set* for the d-dimensional substitution arrays is denoted $\mathcal{I}^d$ where

$$\mathcal{I}^d = \big\{ \vec{j} = (j_1 \ldots, j_d), \text{ such that } j_i \in 0, 1, \ldots, l_i - 1 \text{ for all } i = 1, \ldots, d \big\}. \tag{3}$$

A substitution $\mathcal{S}$ is a map from $\mathcal{A} \times \mathcal{I}^d$ into $\mathcal{A}$. For each element $e \in \mathcal{A}$, it assigns a map which we denote $\mathcal{S}_e : \mathcal{I}^d \to \mathcal{A}$. For $\vec{k} \in \mathcal{I}^d$, the mapping $\mathcal{S}$ restricted to the element $\vec{k}$ is a mapping from $\mathcal{A}$ to $\mathcal{A}$ which we denote $p_{\vec{k}}$. We will frequently think of $\mathcal{S}$ as a block of maps $(p_{\vec{k}})_{\vec{k} \in \mathcal{I}^d}$.

Example 1. An example of a substitution rule on the alphabet $\mathcal{A} = \{0, 1\}$ with $d = 2$ and $l_1 = l_2 = 5$ is given here:

$$\mathcal{S}_0(*) = \begin{pmatrix} 0 & 1 & 1 & 1 & 0 \\ 1 & 1 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 & 1 \\ 1 & 1 & 0 & 1 & 1 \\ 0 & 1 & 1 & 1 & 0 \end{pmatrix}, \qquad \mathcal{S}_1(*) = \begin{pmatrix} 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 \end{pmatrix}, \tag{4}$$

where both blocks are located in $\mathbb{Z}^2$ as prescribed by $\mathcal{I}^2$, with their lower left corners at the origin. If instead we wish to see $\mathcal{S}$ as a matrix $(p_{\vec{k}})_{\vec{k} \in \mathcal{I}^2}$ of maps on $\mathcal{A}$, denote by g_0 the identity map and g_1 the map switching 0 and 1, we obtain:

$$\mathcal{S}(*, \mathcal{I}^2) = (p_{\vec{k}})_{\vec{k} \in \mathcal{I}^2} = \begin{pmatrix} g_0 & g_1 & g_1 & g_1 & g_0 \\ g_1 & g_1 & g_0 & g_1 & g_1 \\ g_1 & g_0 & g_0 & g_0 & g_1 \\ g_1 & g_1 & g_0 & g_1 & g_1 \\ g_0 & g_1 & g_1 & g_1 & g_0 \end{pmatrix}. \tag{5}$$

For example, we see that $p_{(0,0)} = g_0$ and $p_{(3,1)} = g_1$.

Since the location set $\mathcal{I}^d$ has a total of $K = l_1 \cdot l_2 \cdots l_d$ elements, we can consider the *size* of the substitution to be K. Moreover, there is a natural *expansion map* given by the linear map $\phi : \mathbb{R}^d \to \mathbb{R}^d$ that acts on the basis vector $\vec{e}_i$ by multiplication by l_i, so that $\phi(\vec{e}_i) = l_i \vec{e}_i$. Since $K = \det(\phi)$, we can call it the *expansion constant* of the substitution. Every element of $\mathbb{Z}^d$ can be expressed in the form

$$\vec{w} = \phi(\vec{j}) + \vec{k}$$

for some $\vec{j} \in \mathbb{Z}^d$ and $\vec{k} \in \mathcal{I}^d$, so we can consider the substitution $\mathcal{S}$ as an action from $\mathcal{A}^{\mathbb{Z}^d}$ into itself by assigning

$$\mathcal{S}\mathcal{T}\big(\phi(\vec{j}) + \vec{k}\big) = \mathcal{S}_{\mathcal{T}(\vec{j})}(\vec{k}). \tag{6}$$

The substitution can be considered to act on any subblock of $\mathcal{T}$ in the obvious manner.

Definition 1.1. A *substitution sequence for* $\mathcal{S}$ is a nonperiodic, almost periodic sequence invariant under the action of $\mathcal{S}^k$ for some positive integer k.

Fig. 1. Part of a substitution sequence.

Example 1 (*continued*). In Fig. 1, we have used the tiling model to depict a portion of a substitution sequence given by (4). The origin is at the lower left of the image, 0 and 1 are white and black unit squares, respectively, and we see three iterations of the substitution applied to 0. An interesting optical effect of the construction is that this image is a stereogram that can resolve itself in many ways, one of which shows three smaller bars from the tiling in the foreground and the rest dissolves into the background. Some people find the image uncomfortable to look at, and that may be because of the interference between many copies of positive and negative images.

By (6) we can consider the action of $\mathcal{S}^k$ on sequences. It is natural to consider $\mathcal{S}^k$ as a substitution mapping as well. When we do this we consider the block $\mathcal{S}^k(a)$ as being defined on the set $(\mathcal{I}^d)^k = \{\vec{j}: 0 \leqslant j_i \leqslant l_i^k - 1\}$. A substitution $\mathcal{S}$ is called *primitive* if there is a positive integer k such that for each $a \in \mathcal{A}$, $\mathcal{S}^k(a)$ contains all of the elements of $\mathcal{A}$. Conditions for primitivity for sequences are discussed in [16] and for tilings in [15]. If $\mathcal{S}$ does not admit any periodic substitution sequences, then we refer to it as a *nonperiodic* substitution. Note that the substitution in Example 1 is both primitive (with $k = 1$) and nonperiodic.

Considering a substitution sequence $\mathcal{T}_0$ for a nonperiodic, primitive substitution $\mathcal{S}$, we construct the tiling space X for $\mathcal{T}_0$ as described in (2). The reader should note that there is an alternative method for constructing X given as follows. Consider the *language* of the

substitution to be the set $\mathcal{L}(\mathcal{S})$ comprised of all subblocks of blocks of the form $\mathcal{S}^k(a)$, $k \in \mathbb{Z}^+$ and $a \in \mathcal{A}$. We may define X to be the set of all sequences $\mathcal{T} \in \mathcal{A}^{\mathbb{Z}^d}$ such that every block in $\mathcal{T}$ is a translate of some block in $\mathcal{L}(\mathcal{S})$. By primitivity, $\mathcal{T}_0$ contains every possible $\mathcal{S}^k(a)$, so the space created via (2) must contain that created by using $\mathcal{L}(\mathcal{S})$. And since $\mathcal{T}_0$ is contained in the space created by $\mathcal{L}(\mathcal{S})$ the two spaces must coincide. In this case we refer to $(X, \mathbb{Z}^d)$ as the *dynamical system associated to $\mathcal{S}$*, and it will be minimal.

Definition 1.2. We say the substitution $\mathcal{S}$ is a *bijective substitution* if each map $p_{\vec{k}}$ is a bijection of the alphabet $\mathcal{A}$.

The substitution rule given in (4) is an example of a bijective substitution. We will see that bijective substitutions generate dynamical systems that are isomorphic to skew products over odometer transformations. We will also see that if the substitution is bijective, then the frequency distribution of letters from $\mathcal{A}$ is uniform.

1.3. Examples

Example 2 (*Chair tiling*). The tiles of this tiling are *L-triominoes*; three squares attached in an L shape. There is a well-known substitution given by:

It was shown in [18] that this tiling can be recoded as a $\mathbb{Z}^2$ substitution on a four-letter alphabet as follows:

$$0 \to \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix}, \quad 1 \to \begin{pmatrix} 1 & 3 \\ 0 & 1 \end{pmatrix}, \quad 2 \to \begin{pmatrix} 2 & 3 \\ 0 & 2 \end{pmatrix}, \quad 3 \to \begin{pmatrix} 2 & 3 \\ 3 & 1 \end{pmatrix}.$$

This is not a bijective substitution and in fact has "coincidences" (see Section 3.2), implying that it has purely discrete spectrum.

Example 3 (*Table tiling*). The tiles of this tiling are simply *dominoes*: two squares attached along an edge. There is a well-known substitution given by:

It was shown in [18] that this tiling can be recoded as a $\mathbb{Z}^2$ substitution on a four-letter alphabet as follows:

$$0 \to \begin{pmatrix} 2 & 3 \\ 0 & 0 \end{pmatrix}, \quad 1 \to \begin{pmatrix} 1 & 1 \\ 2 & 3 \end{pmatrix}, \quad 2 \to \begin{pmatrix} 0 & 2 \\ 1 & 2 \end{pmatrix}, \quad 3 \to \begin{pmatrix} 3 & 0 \\ 3 & 1 \end{pmatrix}.$$

This is a bijective substitution and so, as discussed in Section 3.3, we have that there is a continuous spectrum component to its dynamical system. A complete spectral analysis of the system appears in [18].

Fig. 2. Factoring the sequence from Example 4 onto two letters.

Example 4. We present one from a family of examples that was introduced in [4]. Letting $l_1 = l_2 = d = 2$ and $\mathcal{A} = \{1, 2, 3, 4, \bar{1}, \bar{2}, \bar{3}, \bar{4}\}$, we assign the substitution rule as follows.

$$S(1) \to \begin{array}{|c|c|} \hline 3 & \bar{4} \\ \hline 1 & 2 \\ \hline \end{array} \qquad S(2) \to \begin{array}{|c|c|} \hline \bar{3} & 4 \\ \hline 1 & 2 \\ \hline \end{array} \qquad S(3) \to \begin{array}{|c|c|} \hline 3 & \bar{4} \\ \hline 1 & \bar{2} \\ \hline \end{array} \qquad S(4) \to \begin{array}{|c|c|} \hline 3 & \bar{4} \\ \hline \bar{1} & 2 \\ \hline \end{array}$$

$$S(\bar{1}) \to \begin{array}{|c|c|} \hline \bar{3} & 4 \\ \hline \bar{1} & 2 \\ \hline \end{array} \qquad S(\bar{2}) \to \begin{array}{|c|c|} \hline 3 & \bar{4} \\ \hline \bar{1} & \bar{2} \\ \hline \end{array} \qquad S(\bar{3}) \to \begin{array}{|c|c|} \hline \bar{3} & \bar{4} \\ \hline \bar{1} & \bar{2} \\ \hline \end{array} \qquad S(\bar{4}) \to \begin{array}{|c|c|} \hline \bar{3} & \bar{4} \\ \hline 1 & \bar{2} \\ \hline \end{array}$$

Note that the substitution for a letter a and its barred counterpart are opposite. It is apparent that the substitution is not bijective, but there are no coincidences (see Section 3.2) and it is proved in [4] that there is an absolutely continuous component to the spectrum. Since what really matters is the placement of the barred elements in the tiling, in Fig. 2 we show six iterations of the letter 1 with the barred elements drawn black and the unbarred drawn white.

1.4. Unique ergodicity and the frequency measure

Self-similar tiling dynamical systems are uniquely ergodic (see [21] and references therein) with the frequency measure being the unique translation-invariant measure. This result applies in our situation because the tiling model uses tiles that have area 1, and so the frequencies for the $\mathbb{Z}^d$ action are the same as those for the tiling action. For a cylinder set $[B]$ corresponding to a fixed block B, the measure $\mu([B])$ represents the frequency of occurrence of the block B in any sequence in X [21]. We will discuss this measure as it is used in our context. Then, using the tiling model and relying heavily on Sections 2 and 3 of [21], we will give a condition forcing uniform frequency of letters.

Let $\mathcal{S}$ be a primitive, nonperiodic substitution with size $l_1 \cdot l_2 \cdots l_d = K$ and let ϕ be the natural expanding map. Then $\mathcal{S}$ has an $|\mathcal{A}| \times |\mathcal{A}|$ *subdivision matrix M* defined by

letting M_{ij} be the number of letters of type a_i in $\mathcal{S}(a_j)$. Since $\mathcal{S}$ is primitive, Corollary 2.4 in [21] implies that the largest eigenvalue of M must be equal to $|\det \phi| = K$. The Perron–Frobenius theorem (see [20]) states that there are unique (up to scalar multiplication) strictly positive right and left *Perron eigenvectors* $\vec{r}$ and $\vec{l}$ for which

$$\vec{l} \cdot \vec{r} = 1, \quad \text{and} \quad \lim_{n \to \infty} K^{-n} M^n = \vec{r} \cdot \vec{l}. \tag{7}$$

Additionally, we know from [21] that $\vec{l}$ should be taken to be the vector given by the volumes of the $|\mathcal{A}|$ tiles under consideration, which in our case are all 1. This implies that $\sum_{i=1}^{|\mathcal{A}|} r_i = 1$ and that

$$\lim_{n \to \infty} K^{-n} M^n = \begin{pmatrix} r_1 & r_1 & \cdots & r_1 \\ r_2 & r_2 & \cdots & r_2 \\ \vdots & \vdots & & \vdots \\ r_{|\mathcal{A}|} & r_{|\mathcal{A}|} & \cdots & r_{|\mathcal{A}|} \end{pmatrix}.$$

Following [21], in order to find the frequency of a letter a_i in an arbitrary sequence in X, it suffices to compute the frequency in larger and larger substituted blocks. By primitivity it does not matter which type of block we substitute so we will look at iterations of a_1. Denote by $N_{a_i}(B)$ the number of occurrences of the letter a_i in a block B. Then

$$freq(a_i) = \lim_{n \to \infty} \frac{N_{a_i}(\mathcal{S}^n(a_1))}{K^n}, \tag{8}$$

since K^n is the volume of the substituted block $\mathcal{S}^n(a_1)$. The numerator is easily computed since it is simply M_{i1}^n. Thus we have that $freq(a_i) = \lim_{n \to \infty} K^{-n} M_{i1}^n = r_i$, and so computation of $freq(a_i)$ reduces to computation of the right eigenvector for M.

Proposition 1.1. *Let $\mathcal{S}$ be a primitive, nonperiodic substitution. Then M has the property that $\sum_{j=1}^{|\mathcal{A}|} M_{ij} = K$ for all $i \in 1, 2, \ldots, |\mathcal{A}|$ if and only if the frequency of any letter $a_i \in \mathcal{A}$ is $1/|\mathcal{A}|$.*

Proof. If $\sum_{j=1}^{|\mathcal{A}|} M_{ij} = K$, then the vector $\vec{r}$ with each coordinate equal to $1/|\mathcal{A}|$ is a right eigenvector for M. Since $\vec{l} \cdot \vec{r} = 1$ it must be the right Perron eigenvector for M satisfying (7), and since $freq(a_i) = r_i$ the result follows. Conversely, if $r_i = 1/|\mathcal{A}|$ for all $i = 1, 2, \ldots, |\mathcal{A}|$, then since $\vec{r}$ is a right eigenvector we have that

$$(M\vec{r})_i = \sum_{j=1}^{|\mathcal{A}|} M_{ij}/|\mathcal{A}| = (K\vec{r})_i = K/|\mathcal{A}|$$

and the result follows. □

Corollary 1.2. *If $\mathcal{S}$ is a primitive, nonperiodic, bijective substitution, then the frequency of any letter $a_i \in \mathcal{A}$ is $1/|\mathcal{A}|$.*

Proof. The row sums of M represent how many times a letter a_i appears in the substitutions of all letters in $\mathcal{A}$ taken together. Since $\mathcal{S}$ is bijective, a_i can appear in any given spot

in the substitution blocks no more than once, and so we see that $\sum_{j=1}^{|\mathcal{A}|} M_{ij} \leqslant K$. But if there is a spot in which a_i never appears, then this would imply that some other letter from $\mathcal{A}$ had to appear in that spot two times, and this would contradict the bijectivity of $\mathcal{S}$. So it must be that $\sum_{j=1}^{|\mathcal{A}|} M_{ij} \geqslant K$, and we see that for a bijective substitution $\mathcal{S}$, the subdivision matrix satisfies $\sum_{j=1}^{|\mathcal{A}|} M_{ij} = K$. $\quad\square$

The proposition also applies to the substitutions seen in [4] such as Example 4.

2. The odometer representation

Assume fixed a primitive, nonperiodic substitution $\mathcal{S}$ on $\mathbb{Z}^d$ with lengths $l_1, l_2, \ldots, l_d$, and let $(X, \mathbb{Z}^d, \mu)$ denote the dynamical system of $\mathcal{S}$. We will show that this dynamical system factors onto a direct product of d Kakutani odometers (i.e., adic transformations) of lengths $l_1, l_2, \ldots, l_d$, respectively. To show this, each $\mathcal{T} \in X$ will be represented by a one-sided sequence of values from $\mathcal{I}^d$ that codes the location of the origin with respect to levels of the hierarchical structure of $\mathcal{T}$. When $\mathcal{S}$ is a bijective substitution, we will show that its dynamical system is measure-theoretically isomorphic to a skew product over a d-dimensional odometer.

The representation shown here follows standard arguments done by Kakutani [8], Keane [9,10], Goodson [6], Ferenczi [2], Kwiatkowski [11], and others. However, since the generalization to multidimensional actions is somewhat nontrivial, it seems worthwhile to present the result. We begin with a brief description of a one-dimensional odometer.

2.1. von Neumann–Kakutani odometers

The description of the classical one-dimensional odometer is based on [14]. Consider the set of digits $\mathcal{D} = \{0, 1, 2, \ldots, l-1\}$ for some $l \in \mathbb{Z}$. The odometer is a transformation on the space $\mathcal{D}^{\mathbb{N}}$ that acts on a sequence by increasing the first element that is less that $l-1$, resetting the previous ones to zero, and leaving the rest alone. One should think of their automobile's odometer with $l = 10$. For each $x \in \mathcal{D}^{\mathbb{N}}$ for which it makes sense, define

$$\eta(x) = \min\{m \text{ such that } x_m < l-1\}. \tag{9}$$

For $x \neq (l-1, l-1, l-1, \ldots)$, the odometer action is:

$$\bigl(V(x)\bigr)_n = \begin{cases} 0 & n < \eta(x), \\ x_n + 1 & n = \eta(x), \\ x_n & n > \eta(x) \end{cases} \tag{10}$$

and we define $V(l-1, l-1, l-1, \ldots) = (0, 0, 0, \ldots)$. Putting the product topology on $\mathcal{D}^{\mathbb{N}}$ and considering the Bernoulli $(1/l, 1/l, \ldots, 1/l)$ measure gives us an invertible measure-preserving map.

When we are concerned with the skew product representation of a substitution we will need to restrict $\mathcal{D}^{\mathbb{N}}$ to the set Σ of all sequences that are not eventually identically 0 or identically $l-1$. On Σ the action V is well-defined and invertible, and Σ is a set of full measure. It is the measure-preserving system (Σ, V, μ_Σ) that forms the base of the skew

product representation for bijective substitutions (described below) and for the dynamical
systems in [4].

2.2. Factoring substitution systems onto odometers

To establish the factor map from the substitution dynamical system to the odometer, we
need the notion of *unique composition* [22] of a sequence $\mathcal{T} \in X$ (called *recognizability*
[13] in the one-dimensional case). The basic idea is that any sequence in X must look like
the image under $\mathcal{S}$ of a unique sequence that is also in X. That is to say, the letters in
$\mathcal{T}$ can be put together in a unique way to form composite blocks that are the images of
letters under $\mathcal{S}$. The work of Solomyak [22] can be applied, considering our sequences as
tilings; or the work of Mossé [13] can be generalized to arrive at the definition of *unique
composition* for $\mathbb{Z}^d$ substitutions: for every $\mathcal{T} \in X$, there exists a unique $\mathcal{T}^1 \in X$ and a
unique $\vec{j} \in \mathcal{I}^d$ such that $\mathcal{T} = \mathcal{S}(\mathcal{T}^1) - \vec{j}$. Whenever $\mathcal{S}$ is nonperiodic and primitive the
result in [22] implies that $\mathcal{S}$ has the unique composition property.

Unique composition implies that for any $M \in \mathbb{Z}^+$ there exists a unique $\mathcal{T}^M \in X$ and
$\vec{j} \in \mathbb{Z}^d$ with $j_i \in 0, 1, 2, \ldots, l_i^M - 1$ for all $i \in 1, 2, \ldots, d$ such that $\mathcal{T} = \mathcal{S}^M(\mathcal{T}^M) - \vec{j}$.
For any letter $a \in \mathcal{T}^M$, we call $\mathcal{S}^M(a) - \vec{j}$ a *level-M block of $\mathcal{T}$*, of *type a*. Each level-M
block B of $\mathcal{T}$ is composed of K level-$(M-1)$ blocks of $\mathcal{T}$ whose positions inside of B are
indexed by $\mathcal{I}^d$ in the obvious manner. We define functions $\mathcal{O}_M : X \to \mathcal{I}^d$ as follows: for
$\mathcal{T} \in X$, $\mathcal{O}_M(\mathcal{T})$ is the position of the level-$(M-1)$ block of $\mathcal{T}$ containing $\vec{0}$ in its level-M
block. We write $\mathcal{T}^0 = \mathcal{T}$, so that $\mathcal{O}_1(\mathcal{T})$ is the position of $\mathcal{T}(\vec{0})$ in its level-1 block.

With this map, we are ready to define our coding of X into an odometer space. Let
$\Sigma_0 = (\mathcal{I}^d)^{\mathbb{Z}^+}$ denote the set of all one-sided sequences of vectors from $\mathcal{I}^d$. Let $\Sigma \subset \Sigma_0$ be
the set of all sequences such that for all $i = 1, 2, \ldots, d$, the ith coordinate is not eventually
identically 0 or $l_i - 1$. Put the product topology on Σ_0 and denote by μ_Σ the Borel prob-
ability measure assigning the measure $1/K$ to cylinder sets with one coordinate fixed, so
that Σ is a set of full measure. Then we have a measurable map $\Theta : (X, \mu) \to (\Sigma_0, \mu_\Sigma)$
given by

$$\Theta(\mathcal{T}) = \left\{ \mathcal{O}_M(\mathcal{T}) \right\}_{M=1}^{\infty}.$$

It is clear that Θ maps onto Σ_0 since there will always be a tiling situated according to any
given sequence in Σ_0.

Example 5. Let $d = 2$ and $l_1 = l_2 = 2$, so that $\mathcal{I}^2 = \{(0,0), (1,0), (0,1), (1,1)\}$ and
suppose that $\mathcal{S}$ is a nonperiodic primitive substitution with these parameters. In Fig. 3,
we diagram the coding for an arbitrary sequence $\mathcal{T} \in X$. We have outlined the begin-
ning of the level-M skeleton around the origin $(0,0)$ for $\mathcal{T}$, and we see that $\Theta(\mathcal{T}) =
\{(1,0), (1,1), (1,0), (0,1), \ldots\}$.

Now let us examine what happens to the coding of $\mathcal{T}$ when it has been translated by
the standard basis element $\vec{e}_i$. A look at Fig. 4 will convince the reader that $\Theta(\mathcal{T} - \vec{e}_i)$ and
$\Theta(\mathcal{T})$ are identical in all but the ith coordinate sequence. Unless $\mathcal{T}$ is situated in a very
special way, there will be some smallest M for which both $\mathcal{T}$ and $\mathcal{T} - \vec{e}_i$ are in the same
level-M tile. The Mth coordinate of the ith sequence of $\Theta(\mathcal{T})$ will be increased by one to

Fig. 3. A sketch of how a typical sequence $\mathcal{T} \in X$ codes into Σ_0.

Fig. 4. How the shift affects the coding.

obtain that of $\Theta(\mathcal{T} - \vec{e}_i)$, and all previous coordinates will be reset to 0. But this is exactly the odometer action, which we now make precise.

Use the notation $\{\vec{x}_m\} \in \Sigma$, leaving implicit that m runs from 1 to ∞, and write $\vec{x}_m \in \mathcal{I}^d$ as $\vec{x}_m = (x_m^1, \dots, x_m^d)$. Thinking of $\{\vec{x}_m\}$ as d infinite sequences, we will have d odometers, each defined as in Section 2.1. We define $\eta_i : \Sigma \to \mathbb{Z}^+$ which keeps track of the first time an $\vec{x}_n$ can be augmented in the ith coordinate and remain in $\mathcal{I}^d$:

$$\eta_i(\{\vec{x}_m\}) = \min\{m \text{ such that } x_m^i < l_i - 1\}. \tag{11}$$

We define the odometer on the ith coordinate to be:

$$\left(V_i(\{\vec{x}_m\})\right)_n = \begin{cases} (x_n^1, \dots, x_n^{i-1}, 0, x_n^{i+1} \dots x_n^d) & n < \eta_i(\{\vec{x}_m\}), \\ (x_n^1, \dots, x_n^{i-1}, x_n^i + 1, x_n^{i+1} \dots x_n^d) & n = \eta_i(\{\vec{x}_m\}), \\ (x_n^1, \dots, x_n^{i-1}, x_n^i, x_n^{i+1} \dots x_n^d) & n > \eta_i(\{\vec{x}_m\}). \end{cases} \tag{12}$$

Returning to the sequence $\mathcal{T}$ of Example 5, we have that $\eta_1(\Theta(\mathcal{T})) = 4$ and that $\Theta(\mathcal{T} - \vec{e}_1)$ will have the coding $\{(0, 0), (0, 1), (0, 0), (1, 1), \dots\}$ as represented in Fig. 4.

So we see that Θ intertwines the action of translation on X with the action of V on Σ so that we have a measure-theoretic factor map from $(X, \mathbb{Z}^d, \mu)$ onto $(\Sigma_0, V, \mu_\Sigma)$.

2.3. Bijective substitutions as skew products over the odometer

Give the alphabet $\mathcal{A}$ the discrete topology with counting measure $\mu_{\mathcal{A}}(a) = 1/|\mathcal{A}|$ for all $a \in \mathcal{A}$. Every $\mathcal{T} \in X$ is coded using the map $\Psi : X \to \Sigma_0 \times \mathcal{A}$ by

$$\Psi(\mathcal{T}) = \left(\left\{\mathcal{O}_M(\mathcal{T})\right\}_{M=1}^{\infty}, \mathcal{T}(\vec{0})\right). \tag{13}$$

The map Ψ is one-to-one on a set of full measure in X composed of tilings that encode in the first coordinate to Σ. This is due to the bijectivity of the substitution: once we know what the symbol at the origin is and its position in its level-1 tile, we know the rest of the letters in the level-1 tile. But knowing the level-1 tile along with its position in its level-2 tile allows us to fill in the rest of the level-2 tile. We can continue filling in the letters of $\mathcal{T}$ in this fashion, and as long as the origin is not eventually always on the edge of its level-m blocks, we will uniquely specify the letters in $\mathcal{T}$ for all of $\mathbb{Z}^d$. The set of all sequences that are eventually always on the edge of their level-m blocks has measure zero and maps onto $\Sigma_0 - \Sigma$.

The map Ψ is a measure-theoretic isomorphism between (X, μ) and $(\Sigma \times \mathcal{A}, \mu_{\Sigma} \times \mu_{\mathcal{A}})$. To check this, we can show that a generating set for the Borel σ-algebras have their measures preserved by Ψ. Let $[B]$ be the cylinder set in Σ given by fixing the first n coordinates $[B] = \{\vec{b}_1, \vec{b}_2, \ldots, \vec{b}_n, *, *, *, \ldots\}$. The collection of all sets $([B], a) \subset \Sigma \times \mathcal{A}$ for all B and all $a \in \mathcal{A}$ will generate the Borel σ-algebra. Naturally $\mu_{\Sigma} \times \mu_{\mathcal{A}}([B], a) = 1/K^n \cdot 1/|\mathcal{A}|$. The cylinder $[B]$ not only fixes the location of the level-n block at the origin precisely inside its level-$(n + 1)$ block, it fixes all of the blocks down to the level-0 block. That is, we know the coding of any sequence in $\Psi^{-1}([B], a)$ precisely out to the nth place, and therefore the frequency of this block is $1/K^n$. Fixing a tells us which of the $|\mathcal{A}|$ possibilities our sequence is in, and so we find that $\mu(\Psi^{-1}([B], a))$ must be $1/K^n \cdot 1/|\mathcal{A}|$.

We are ready to define the cocycle maps ϕ_i that keep track of the change on the second coordinate of $\Sigma \times \mathcal{A}$ resulting from translation. Since the formula seems complicated, it will be useful to describe the cocycle for the example from Figs. 3 and 4 first, then write down the general formula. We know that $\Psi(\mathcal{T}) = (\{\vec{x}_m\}, \mathcal{T}(0, 0)) = ((1, 0), (1, 1), (1, 0), (0, 1), \ldots\}, \mathcal{T}(0, 0))$ and that $\Psi(\mathcal{T} - (1, 0)) = (V_1(\{\vec{x}_m\}), (\mathcal{T} - (1, 0))(0, 0))$ $= (\{(0, 0), (0, 1), (0, 0), (1, 1), \ldots\}, \mathcal{T}(1, 0))$. But in order to define an action on $\Sigma \times \mathcal{A}$ that commutes with translation, we must be able to determine the second coordinate without reference to $\mathcal{T}(1, 0)$.

For each $\{\vec{x}_m\} \in \Sigma$, we figure out the value of $\mathcal{T}(1, 0)$ inductively. Knowing $\mathcal{T}(0, 0)$ and knowing $\vec{x}_1 = (1, 0)$ means that we can use the inverse of the substitution map $p_{(1,0)}$ to figure out the type of the level-1 block at the origin: it must have type $p_{(1,0)}^{-1}(\mathcal{T}(0, 0))$. (Note that the invertibility of the substitution map follows from its bijectivity.) Knowing the type of the level-1 block and knowing that $\vec{x}_2 = (1, 1)$ allows us to figure out the type of the level-2 block at the origin: it must have type $p_{(1,1)}^{-1} p_{(1,0)}^{-1}(\mathcal{T}(0, 0))$. We can continue in this fashion until we have figured out the type of the level-4 block at the origin, and for this sequence that is going to be enough because our shift to the left occurs inside this block (which is to say, $\eta_1(\{\vec{x}_m\}) = 4$). So we know the type of the level-4 block at the origin is determined by the bijection:

$$\left(\mathcal{P}_1\left(\{(1, 0), (1, 1), (1, 0), (0, 1), \ldots\}\right)\right)^{-1} = p_{(0,1)}^{-1} p_{(1,0)}^{-1} p_{(1,1)}^{-1} p_{(1,0)}^{-1}.$$

Now that we know the type of the level-4 block at the origin, we may use the maps given by $V_1(\{\vec{x}_m\}) = \{(0,0), (0,1), (0,0), (1,1), \ldots\}$ to work our way back down to the type of the letter $\mathcal{T}(1,0)$. The type of the level-3 block at the origin in $\mathcal{T} - (1,0)$ must be $p_{(1,1)}$ applied to the type of the level-4 block, then the type of the level-2 block must be $p_{(0,0)}$ applied to that, then the type of the level-1 block is $p_{(0,1)}$ applied to that, and finally the type of the level-0 block is $p_{(0,0)}$ applied to that. That is, we can define a map that works its way back down to the level-0 block as:

$$\mathcal{P}_1\big(V_1\big(\{(1,0), (1,1), (1,0), (0,1), \ldots\}\big)\big) = p_{(0,0)} \, p_{(0,1)} \, p_{(0,0)} \, p_{(1,1)}.$$

Thus we can specify the type of $\mathcal{T}(1,0)$ entirely by knowing the coding $\{\vec{x}_m\}$ and by knowing $\mathcal{T}(0,0)$ by writing

$$\mathcal{T}(1,0) = \mathcal{P}_1\big(V_1\big(\{\vec{x}_m\}\big)\big) \circ \big(\mathcal{P}_1\big(\{\vec{x}_m\}\big)\big)^{-1} \mathcal{T}(0,0) = \phi_1\big(\{\vec{x}_m\}\big) \mathcal{T}(0,0).$$

Now we can give the general definition of the cocycle maps ϕ_i that act on the second coordinate of the coding to produce the correct letter change as we shift by $\vec{e}_i$. We define $\phi_i : \Sigma \to \mathcal{S}_\mathcal{A}$ in such a way that the map $V_i \times \phi_i : \Sigma \times \mathcal{A} \to \Sigma \times \mathcal{A}$ commutes with the action of translation by $\vec{e}_i$ on the sequence space X. In order to figure out the type of the next letter in the $\vec{e}_i$ direction in the tiling coded by $\Psi(\mathcal{T}) = (\{\vec{x}_m\}, \mathcal{T}(\vec{0}))$, we use the first $\eta_i(\{\vec{x}_m\})$ terms of $\{\vec{x}_m\}$ to figure out the type of the level-$\eta_i(\{\vec{x}_m\})$ block using the inverses of the substitution maps, then use the first $\eta_i(\{\vec{x}_m\})$ terms of $V_i(\{\vec{x}_m\})$ to work our way back down to the level-0 block we want. The iterative step in both cases can be defined in terms of a map

$$\mathcal{P}_i\big(\{\vec{x}_m\}\big) = p_{\vec{x}_1} \, p_{\vec{x}_2} \cdots p_{\vec{x}_{\eta_i\{\vec{x}_m\}}}. \tag{14}$$

So that we define the cocycle to be:

$$\phi_i\big(\{\vec{x}_m\}\big) = \mathcal{P}_i\big(V_i\big(\{\vec{x}_m\}\big)\big) \circ \big(\mathcal{P}_i\big(\{\vec{x}_m\}\big)\big)^{-1}. \tag{15}$$

We can define the $\mathbb{Z}^d$ action on $\Sigma \times \mathcal{A}$ to be generated by the action

$$V_i \times \phi_i\big(\{\vec{x}_m\}, a\big) = \big(V_i\big(\{\vec{x}_m\}\big), \phi_i\big(\{\vec{x}_m\}\big)a\big), \tag{16}$$

for each $i \in \{1, 2, \ldots, d\}$. In this case it is clear by construction that the actions of the generators commute with the action of translation by $\vec{e}_i$, and we have that

$$\begin{aligned} V_i \times \phi_i\big(\Psi(\mathcal{T})\big) &= V_i \times \phi_i\big(\{\mathcal{O}_m(\mathcal{T})\}, \mathcal{T}(\vec{0})\big) \\ &= \big(V_i\big(\{\mathcal{O}_M(\mathcal{T})\}\big), \phi_i\big(\{\mathcal{O}_M(\mathcal{T})\}\big)\big(\mathcal{T}(\vec{0})\big)\big) \\ &= \big(\{\mathcal{O}_M(\mathcal{T} - \vec{e}_i)\}, \mathcal{T}(\vec{e}_i)\big) = \Psi(\mathcal{T} - \vec{e}_i). \end{aligned}$$

So Ψ is a measure-theoretic isomorphism between the dynamical systems $(X, \mathbb{Z}^d, \mu)$ and $(\Sigma \times \mathcal{A}, \mathbb{Z}^d, \mu_\Sigma \times \mu_\mathcal{A})$.

We have shown that our shift dynamical system is an almost $|\mathcal{A}|$-point extension of the product of d odometers. It should be noted that in this form, it is not a group extension. Nonetheless, we will see in Section 4 a spectral decomposition which looks quite like that for a standard group extension. In fact, one can refer to [16] to see a description of how to derive a group extension from this which is in some cases isomorphic and in all cases closely related to our system.

3. Spectral theory and bijective substitution sequences in $\mathbb{Z}^d$

Consider the unitary $\mathbb{Z}^d$-action on a Hilbert space given by $U^{\vec{j}}: L^2(X, \mu) \to L^2(X, \mu)$ with $U^{\vec{j}}(f(\mathcal{T})) = f(\mathcal{T} - \vec{j})$ for all $\vec{j} \in \mathbb{Z}^d$. We can analyze the action of $\mathbb{Z}^d$ on X by consideration of the action of $U^{\vec{j}}$ on $L^2(X, \mu)$. The *spectral coefficients* of an $L^2(X, \mu)$ function are given, for each $\vec{j} \in \mathbb{Z}^d$, by

$$\hat{f}(\vec{j}) = \langle U^{\vec{j}} f, f \rangle = \int_X U^{\vec{j}} f(\mathcal{T}) \overline{f(\mathcal{T})} \, d\mu(\mathcal{T}). \tag{17}$$

It is known that these coefficients form a positive definite sequence and that therefore there is a unique measure σ_f on the d-torus [19] with:

$$\hat{f}(\vec{j}) = \int_{\mathbb{T}^d} z^{\vec{j}} \, d\sigma_f(z), \tag{18}$$

where $z^{\vec{j}} = z_1^{j_1} \cdots z_d^{j_d}$. For a fixed $f \in L^2(X, \mu)$, we consider the *cyclic subspace* generated by the closed linear span of f as $Z(f) = span\overline{\{U^{\vec{j}}(f): \vec{j} \in \mathbb{Z}^d\}}$. The action of U restricted to $Z(f)$ is unitarily equivalent to the action $V^{\vec{j}}: L^2(\mathbb{T}^d, \sigma_f) \to L^2(\mathbb{T}^d, \sigma_f)$ given by $V^{\vec{j}}(g(\vec{z})) = z^{\vec{j}} g(\vec{z})$. A survey of a wide variety of spectral results in the context of dynamical systems appears in [7].

Any Borel measure σ on the torus can be decomposed into at most three mutually singular parts: a discrete part corresponding to purely atomic measure, a singular continuous part that is nonatomic but singular with respect to Lebesgue measure, and a part that is absolutely continuous with respect to Lebesgue measure. It follows from [21] that every substitution sequence in $\mathbb{Z}^d$ has functions whose spectral measures are purely discrete. This is because the expansion constant K of a $\mathbb{Z}^d$ substitution must be an integer, which is a Pisot number, and that is the condition precluding weak mixing. We have included examples of substitution sequences having mixed spectrum: bijective substitutions, substitutions like Example 4 (see [4]), and the "table" substitution of Example 3 (see [18]).

3.1. Eigenvalues and eigenfunctions

We refer to a constant $\vec{\alpha} \in \mathbb{R}^d$ as an *eigenvalue* of the action $U^{\vec{j}}$ if there is a function $f \in L^2(X, \mu)$ for which

$$U^{\vec{j}} f = \exp\big(2\pi i (\vec{\alpha} \cdot \vec{j})\big) f \tag{19}$$

for all $\vec{j} \in \mathbb{Z}^d$. (Here $\vec{\alpha} \cdot \vec{j}$ denotes the usual dot product in $\mathbb{R}^d$.) One can check that the spectral measure of such an eigenfunction is the atomic measure supported on $\exp(2\pi i \vec{\alpha}) = (\exp(2\pi i \alpha_1), \ldots, \exp(2\pi i \alpha_d)) \in \mathbb{T}^d$. Every function with an atomic spectral measure is in the linear span of the eigenfunctions, which we denote H_D and call the *discrete component* of the spectrum. Since the spectral measures of distinct eigenfunctions are mutually singular, a nontrivial result from spectral theory is that H_D can be written as a single cyclic subspace.

The discrete component of $L^2(X, \mu)$ for any substitution sequence of the type constructed in this paper must contain eigenfunctions given by the odometer coding. The eigenvalues for the odometer system take the form

$$\vec{\alpha} = \left(\frac{m_1}{l_1^{n_1}}, \dots, \frac{m_d}{l_d^{n_d}} \right),$$

where m_i and n_i are in $0, 1, 2, \dots$ (i.e., d-tuples of the l_i-adic numbers). One can specify the value of the associated eigenfunction g evaluated at $\mathcal{T} \in X$ by knowing the location of $\vec{0}$ in the level-N tile of $\mathcal{T}$ for an appropriate choice of N. Since each eigenvalue can have only one eigenfunction, we see that $g(\mathcal{T})$ depends only on the coding of $\mathcal{T}$ into Σ_0 and not on $\mathcal{T}(\vec{0})$.

In many cases the odometer transformation forms the maximal equicontinuous factor of the system. In some cases there may be additional eigenfunctions that arise from an underlying periodicity of a sort we describe next. In the one-dimensional, constant-length case this happens when the "height" of the substitution is nontrivial; a complete characterization is proved in [1] and summarized in [3]. In the one-dimensional case, the height h of a substitution of constant length q is defined to be:

$$h = \max\{n \geqslant 1, \ (n, q) = 1, \ n \mid \gcd\{a: u(a) = u(0)\}\}, \tag{20}$$

where u is a fixed point of the substitution. That is, h is the largest number that divides all return times to the letter $u(0)$ and is relatively prime to q.

To generalize this notion to higher dimensions, it is necessary to translate these concepts relative to sublattices of $\mathbb{Z}^d$. We begin with a generalization of greatest common divisor. If $\mathcal{L}_1$ and $\mathcal{L}_2$ are sublattices of $\mathbb{Z}^d$, we define $(\mathcal{L}_1, \mathcal{L}_2)$ to be the smallest lattice $\mathcal{L}$ containing both $\mathcal{L}_1$ and $\mathcal{L}_2$, where the word "smallest" means that any other lattice $\mathcal{L}'$ containing $\mathcal{L}_1$ and $\mathcal{L}_2$ must contain $\mathcal{L}$; it is the lattice generated by $\mathcal{L}_1$ and $\mathcal{L}_2$. Note that if one considers the sublattices $n\mathbb{Z}$ and $m\mathbb{Z}$ of $\mathbb{Z}$, our definition gives the sublattice $k\mathbb{Z}$, where $k = (n, m)$. Thus if n and m are relatively prime, we have that $(n\mathbb{Z}, m\mathbb{Z}) = \mathbb{Z}$.

Following [21], choosing any $\mathcal{T} \in X$ we define the set of return times

$$\varXi = \{\vec{j} \in \mathbb{Z}^d: \text{ there exists } \vec{k} \in \mathbb{Z}^d \text{ with } \mathcal{T}(\vec{k} + \vec{j}) = \mathcal{T}(\vec{k})\}. \tag{21}$$

Note that this is well-defined independent of $\mathcal{T}$ by minimality. Define the lattice $\mathcal{L}(\varXi)$ to be the smallest lattice containing $\varXi$. Then the *height lattice* $\varLambda$ is the smallest lattice containing $\mathcal{L}(\varXi)$ for which $(\varLambda, \phi(\mathbb{Z}^d)) = \mathbb{Z}^d$. If $\varLambda = \mathbb{Z}^d$, the height is said to be *trivial*. Examples of substitutions with nontrivial height appear in Section 5.

Letting $\varLambda^*$ represent the dual lattice of $\varLambda$ (the lattice of all elements of $\mathbb{R}^d$ that have integer inner product with all elements of $\varLambda$), the eigenvalue group is given by $\bigcup_{n \geqslant 1} \phi^{-n} \varLambda^*$. Thus the height is trivial if and only if the odometer is the maximal equicontinuous factor.

3.2. Coincidences and purely discrete spectrum

Recall that the location set for a k-times substituted letter is given by $(\mathcal{I}^d)^k = \{\vec{j} \in \mathbb{Z}^d: 0 \leqslant j_i < l_i^k\}$. We say a substitution $\mathcal{S}$ admits a *coincidence* if there is a $k \in 1, 2, \dots$ and a $\vec{j} \in (\mathcal{I}^d)^k$ such that $\mathcal{S}^k(a, \vec{j}) = \mathcal{S}^k(b, \vec{j})$ for all $a, b \in \mathcal{A}$. That is, if one iterates the substitution enough times, there will be a location in which all of the letters agree.

In [1] Dekking showed that for substitutions of constant length (the $d = 1$ case in this paper), height 1 implies that coincidence is equivalent to purely discrete spectrum. If the height is nontrivial, Dekking characterizes purely discrete spectrum in terms of coincidences in the "pure base" of the substitution. In Section 6 of [21], Solomyak adapts this result to tilings of $\mathbb{R}^1$ and $\mathbb{R}^2$. Moreover, Theorem 6.1 in the same paper gives a sufficient condition for pure discrete spectrum for tilings of $\mathbb{R}^d$ that applies to the sequences given in this paper. One can check that the hypotheses of Theorem 6.1 hold when a $\mathbb{Z}^d$ substitution admits a coincidence, and thus we have purely discrete spectrum in this case.

3.3. Continuous spectrum

Bijective substitution sequences in $\mathbb{Z}^d$ do not have a purely discrete spectrum. One way to see this is to exhibit a function in $L^2(X, \mu)$ that is orthogonal to the eigenfunctions. In the trivial height case, this is a simple matter: take the alphabet $\mathcal{A}$ and any one-to-one map $F : \mathcal{A} \to \{1, 2, \ldots |\mathcal{A}|\}$. Define the function

$$f(\mathcal{T}) = \exp\left(2\pi \mathrm{i} \frac{F(\mathcal{T}(\vec{0}))}{|\mathcal{A}|}\right). \tag{22}$$

For any eigenfunction g with eigenvalue α, we check:

$$\langle U^{\vec{j}} g, f \rangle = \int_X U^{\vec{j}} g(\mathcal{T}) f(\mathcal{T}) \, \mathrm{d}\mu(\mathcal{T}) = \int_X \mathrm{e}^{2\pi \mathrm{i}\alpha \cdot \vec{j}} g(\mathcal{T}) f(\mathcal{T}) \, \mathrm{d}\mu(\mathcal{T})$$

$$= \sum_{a \in \mathcal{A}} \int_{X_a} \mathrm{e}^{2\pi \mathrm{i}\alpha \cdot \vec{j}} g(\mathcal{T}) \mathrm{e}^{2\pi \mathrm{i}\frac{F(a)}{|\mathcal{A}|}} \, \mathrm{d}\mu(\mathcal{T})$$

$$= \sum_{a \in \mathcal{A}} \mathrm{e}^{2\pi \mathrm{i}\frac{F(a)}{|\mathcal{A}|}} \left(\int_{X_a} \mathrm{e}^{2\pi \mathrm{i}\alpha \cdot \vec{j}} g(\mathcal{T}) \, \mathrm{d}\mu(\mathcal{T})\right),$$

where X_a is the cylinder set of all sequences with $\mathcal{T}(\vec{0}) = a$. Since g is the eigenfunction of a substitution of trivial height, it only depends on the odometer codings of tilings in X_a. Because of bijectivity each X_a has exactly one tiling for each coding in Σ; moreover the measure of any subset of X_a on which g is constant is independent of a. This shows that each of the integrals are equal to a constant $C(g)$. So we have that

$$\langle U^{\vec{j}} g, f \rangle = C(g) \sum_{a \in \mathcal{A}} \mathrm{e}^{2\pi \mathrm{i}\frac{F(a)}{|\mathcal{A}|}} = 0.$$

Since f is orthogonal to all of the eigenfunctions, its spectral measure is purely continuous. Because of its specialized nature, it can be used to compute correlation measures of sequences $\mathcal{T} \in X$.

If the height is nontrivial, one must be somewhat more careful with the definition of the function f to ensure that it is also orthogonal to the extra functions in the discrete spectrum that arise apart from the odometer.

3.4. Correlation measures

When one has a one-sided sequence $\{u_n\}$ of real or complex numbers, one can consider the *correlation measures* [16] given by measures on the circle with spectral coefficients

$$\gamma(k) = \lim_{N \to \infty} \frac{1}{N} \sum_{n < N} u_{n+k} \overline{u_n} \tag{23}$$

provided the limit exists. If the limit does not exist, one proceeds either by taking the lim sup or by going along subsequences. In the case of a primitive substitutive sequence, the shift dynamical system generated by $\{u_n\}$ is uniquely ergodic [12], and this implies that the correlation measures will exist and be dominated by the maximal spectral type of the shift dynamical system (see [16] for details). We can analyze a specific substitution sequence even if it is not on an alphabet of complex numbers by considering a bijection from the alphabet $\mathcal{A}$ onto the $|\mathcal{A}|$th roots of unity as in (22). By considering every such bijection, we will have a family of correlation measures which should in some sense encompass every autocorrelation possible for the sequence. Indeed, work in the one-dimensional case in [16], which we do not generalize in this paper, indicates that the maximal spectral type of the system is probably dominated by the sum of these correlation measures.

In order to create a limit that would apply to substitution sequences in $\mathbb{Z}^d$, we can consider a *Van Hove sequence* of subsets $\{A_n\}$ of $\mathbb{R}^d$, following the usage in [21]. Define $\partial(A)^{+r} = \{\vec{x}: \text{dist}(\vec{x}, \partial(A)) \leqslant r\}$, the set of all points in $\mathbb{R}^d$ within r of the boundary of the set A. A sequence of subsets will be a Van Hove sequence if $\lim_{n \to \infty} \text{Vol}(\partial(A_n))^{+r} / \text{Vol}(A_n) = 0$.

Consider a primitive substitution sequence $\mathcal{T}$ and any bijection f from $\mathcal{A}$ into the $|\mathcal{A}|$th roots of unity as given in Eq. (22). Then $\mathcal{T}_f$ can be defined to be the sequence given by $\mathcal{T}_f(\vec{j}) = f(\mathcal{T}(\vec{j}))$ for all $\vec{j} \in \mathbb{Z}^d$, so that $\mathcal{T}_f$ is a sequence on a complex alphabet with correlation coefficients

$$\gamma_f(\vec{k}) = \lim_{n \to \infty} \frac{1}{\text{Vol}(A_n)} \sum_{\vec{j} \in A_n} \mathcal{T}_f(\vec{j} + \vec{k}) \overline{\mathcal{T}_f(\vec{j})}. \tag{24}$$

If $\mathcal{A} \subset \mathbb{C}$, then we can find the correlation measure γ without reference to f.

We prove that the correlation coefficients exist whenever $\mathcal{T}$ is a primitive substitution sequence by relying again on [21]. Note that the product $\mathcal{T}_f(\vec{j} + \vec{k}) \overline{\mathcal{T}_f(\vec{j})}$ can assume only a finite number of values that depend entirely on $\mathcal{T}(\vec{j} + \vec{k})$ and $\mathcal{T}(\vec{j})$. Given $a, b \in \mathcal{A}$, denote by P_{ab} the set of all sequences $\mathcal{T}' \in X$ with $\mathcal{T}'(\vec{k}) = a$ and $\mathcal{T}'(\vec{0}) = b$, and let χ_{ab} denote the indicator function of P_{ab}. We can rewrite the correlation coefficient:

$$\begin{aligned}
\gamma_f(\vec{k}) &= \lim_{n \to \infty} \sum_{(a,b) \in \mathcal{A} \times \mathcal{A}} \frac{1}{\text{Vol}(A_n)} \sum_{\vec{j} \in A_n} \chi_{ab}(\mathcal{T} - \vec{j}) f(a) \overline{f(b)} \\
&= \sum_{(a,b) \in \mathcal{A} \times \mathcal{A}} f(a) \overline{f(b)} \lim_{n \to \infty} \frac{1}{\text{Vol}(A_n)} \sum_{\vec{j} \in A_n} \chi_{ab}(\mathcal{T} - \vec{j}) \\
&= \sum_{(a,b) \in \mathcal{A} \times \mathcal{A}} f(a) \overline{f(b)} \mu(P_{ab}),
\end{aligned}$$

where the last step follows from the unique ergodicity proved in [21]: the interior limit is the integral of χ_{ab} over X, which is simply the frequency measure of P_{ab}.

In unpublished computations the author found that for two specific examples on a two-letter alphabet, the only continuous spectrum is singular. It seems that this will be the case for many (if not all) bijective substitutions. For instance, we have the following proposition whenever the substitution is generated by a cyclic group of order N.

Proposition 3.1. *Let $N > 1$ be a positive integer, $w = \exp(2\pi i/N)$, and let $\mathcal{A} = \{1, w, w^2, \ldots, w^{N-1}\}$. Suppose $\mathcal{S}$ has the property that $\mathcal{S}(w^a, \vec{l}) = w^a \mathcal{S}(1, \vec{l})$ for all $a \in \mathbb{Z}$ and all $\vec{l} \in \mathcal{I}^d$. Then if $\mathcal{S}$ is a primitive substitution we have that*

$$\gamma\big(\phi(\vec{k})\big) = \gamma(\vec{k})$$

for all $\vec{k} \in \mathbb{Z}^d$.

Proof. Let A_n be any Van Hove sequence, to be used in the computation of $\gamma(\vec{k})$, and assume that $\mathcal{T}$ is a substitution sequence for $\mathcal{S}$ invariant under $\mathcal{S}^1$ by renaming $\mathcal{S}$ to be a power of itself if necessary. The sequence $\phi(A_n)$ is also a Van Hove sequence and we will use this to compute $\gamma(\phi(\vec{k}))$.

Fixing A_n, let $\vec{j} \in A_n$ and $\vec{l} \in \mathcal{I}^d$. We know that $\mathcal{T}(\vec{j}) = w^a$ and $\mathcal{T}(\vec{j} + \vec{k}) = w^b$ for some $a, b \in \mathbb{Z}$. Since $\mathcal{S}(w^b) = w^{b-a}\mathcal{S}(w^a)$ by hypothesis, we now know that $\mathcal{T}(\phi(\vec{j}) + \phi(\vec{k}) + \vec{l}) = w^{b-a}\mathcal{T}(\phi(\vec{j}) + \vec{l})$. So for all $\vec{l} \in \mathcal{I}^d$ we have that

$$\mathcal{T}\big(\phi(\vec{j}) + \phi(\vec{k}) + \vec{l}\big)\overline{\mathcal{T}\big(\phi(\vec{j}) + \vec{l}\big)}$$
$$= w^{b-a}\mathcal{T}\big(\phi(\vec{j}) + \vec{l}\big)\overline{\mathcal{T}\big(\phi(\vec{j}) + \vec{l}\big)} = w^{b-a}$$
$$= \mathcal{T}(\vec{j} + \vec{k})\overline{\mathcal{T}(\vec{j})}.$$

Hence we see that

$$\frac{1}{\text{Vol}(A_n)} \sum_{\vec{j} \in A_n} \mathcal{T}(\vec{j} + \vec{k})\overline{\mathcal{T}(\vec{j})} = \frac{1}{\text{Vol}(\phi(A_n))} \sum_{\vec{j} \in \phi(A_n)} \mathcal{T}(\vec{j} + \vec{k})\overline{\mathcal{T}(\vec{j})}$$

and the result follows.   □

In this case if there is any nonzero vector $\vec{k}$ such that $\gamma(\vec{k}) \neq 0$, then immediately we have that the correlation measure is singular continuous. The known examples with Lebesgue components to the spectrum are not bijective [4].

4. Spectral decomposition of bijective substitutions

For some substitutions, the spectrum $L^2(X, \mu)$ of the dynamical system associated to $\mathcal{S}$ can be decomposed into orthogonal subspaces that are invariant under the action of translation. The condition making the decomposition possible is the existence of translation-commuting *letter inversions*.

4.1. Letter inversions

Since almost every sequence $\mathcal{T} \in X$ is determined by its coding over the odometer along with the letter at the origin, we see that changing $\mathcal{T}(\vec{0})$ will result in a change of all the letters of $\mathcal{T}$ in a specified manner, and we will arrive at a new sequence $\mathcal{T}' \in X$ which has the same odometer coding as $\mathcal{T}$ but different letters. Depending on the nature of the bijections $p_{\vec{j}}$ constituting $\mathcal{S}$, there may exist translation-commuting *letter inversions* $\sigma \in \mathcal{S}_{\mathcal{A}}$ defined by the equation $\sigma(\mathcal{T}(\vec{k})) = \mathcal{T}'(\vec{k})$ for all $\vec{k} \in \mathbb{Z}^d$. So for each $\vec{e}_i$, $i \in 1, 2, \ldots, d$, we would have that $\sigma(\mathcal{T} - \vec{e}_i(\vec{0})) = \mathcal{T}' - \vec{e}_i(\vec{0})$. When passed to the skew product this would require that:

$$V_i \times \phi_i\big(\{\vec{x}_m\}, \sigma a\big) = \big(V_i\big(\{\vec{x}_m\}\big), \sigma\phi_i\big(\{\vec{x}_m\}\big)a\big).$$

Since $V_i \times \phi_i(\{\vec{x}_m\}, \sigma a) = (V_i(\{\vec{x}_m\}), \phi_i(\{\vec{x}_m\})\sigma a)$, we see that $\phi_i(\{\vec{x}_m\})\sigma a = \sigma\phi_i(\{\vec{x}_m\})a$ for all possible $\{\vec{x}_m\}$. So we obtain

$$\sigma^{-1}\phi_i\big(\{\vec{x}_m\}\big)\sigma = \phi_i\big(\{\vec{x}_m\}\big). \tag{25}$$

So there is a translation commuting letter inversion if and only if there is a σ satisfying the above equation for all possible i and $\{\vec{x}_m\}$.

It is not difficult to check that letter inversions form a group G of bijections of $\mathcal{A}$ and that if $\sigma \in G$ fixes any letter in $\mathcal{A}$, then it is the identity. Moreover, if any $p_{\vec{j}}$ is the identity, then every substitution bijection $p_{\vec{j}}$ must commute with each $\sigma \in G$. If the substitution bijections generate an abelian group, then they are clearly letter inversions.

Fix a letter $a \in \mathcal{A}$, and consider any sequence $\mathcal{T} \in X$ with $\mathcal{T}(\vec{0}) = a$. For $\sigma \in G$, suppose that $\sigma a = a'$. If $a' = a$, then σ must be the identity map. If not, then for any letter $b \in \mathcal{A}$, we can determine σb, since b must appear in $\mathcal{T}$ and the value of σb is uniquely determined because σ commutes with translation. Thus every letter inversion is completely determined by where it sends the letter a, and this shows that $|G| \leqslant |\mathcal{A}|$.

Proposition 4.1. *Let $\mathcal{S}$ be given by $(p_{\vec{k}})_{\vec{k} \in \mathcal{I}^d}$. Then $\sigma \in G$ iff σ commutes with all products $p_{\vec{k}_1} p_{\vec{k}_2} \cdots p_{\vec{k}_N}$, where $\vec{k}_i \in \mathcal{I}^d$ and N is the order of any $p_{\vec{k}}$.*

Proof. Fix an N so that $p_{\vec{k}}^N = id$ for some $\vec{k} \in \mathcal{I}^d$. One can consider the substitution map given by $\mathcal{S}^N$, which, when $\mathcal{S}$ is primitive and nonperiodic, will generate the same dynamical system $(X, \mathbb{Z}^d, \mu)$ that $\mathcal{S}$ does. That is because if $\mathcal{T}_0$ is a substitution sequence for $\mathcal{S}$ which is invariant under $\mathcal{S}^k$, then $\mathcal{T}_0$ is invariant under $\mathcal{S}^{Nk}$, making it a substitution sequence for $\mathcal{S}^N$ as well. (This is also easily seen by comparing the languages $\mathcal{L}(\mathcal{S}^N)$ and $\mathcal{L}(\mathcal{S})$.) Note that the substitution bijections defining $\mathcal{S}^N$ encompass all possible products of N of the $p_{\vec{j}}$ permutations. Since by assumption one of these products is the identity, it is clear that the letter inversions for $\mathcal{S}^N$ must commute with all such products.

We claim that letter inversions for $\mathcal{S}^N$ are letter inversions for $\mathcal{S}$. The dynamical system $(X, \mathbb{Z}^d, \mu)$ codes into skew products for both substitutions, and the codings must be isomorphic. The cocycle actions will look different, but they must be identical because for any given $\mathcal{T} \in X$, a translation-commuting letter inversion applied to $\mathcal{T}(\vec{0})$ will be independent of the coding. $\quad\square$

4.2. Spectral decomposition in the presence of letter inversion

We now assume that our substitution $\mathcal{S}$ has a nontrivial group G of letter inversions. In the event that the $p_{\vec{j}}$'s generate an abelian group, that group is automatically the group of letter inversions for the system. We will use the representation of the dynamical system as a skew product over an odometer to decompose $L^2(X, \mu)$ in terms of functions that are well-behaved with respect to letter inversion. In Section 5, we list some examples satisfying the hypotheses of the following theorem and examine the consequences of it.

It should be noted that the result is very similar to an analogous result that holds for compact group extensions, even though our skew product formulation is not quite in this form. In the case that $\{p_{\vec{k}}, \vec{k} \in \mathcal{I}^d\}$ generates an abelian group, one can rewrite the system as a group extension that is equivalent to our product. If it does not, a group extension exists that dominates our system [16]. We include the result here so that the reader can see how standard arguments about functional decomposition can be applied to this situation.

Theorem 4.2. *Let $G \subset S_{\mathcal{A}}$ denote the group of translation-commuting letter inversions on X, and let $\widehat{G}$ denote the continuous group homomorphisms from G to the complex unit circle (i.e., the group characters of G). If $\sum_{\chi \in \widehat{G}} \chi(\sigma) = 0$ for all $\sigma \in G$ with $\sigma \neq$ identity, then*

$$L^2(\Sigma \times \mathcal{A}, \mu_\Sigma \times \mu_\mathcal{A}) = H_0 \oplus H_{\chi_1} \oplus \cdots \oplus H_{\chi_{|\widehat{G}|-1}} = \bigoplus_{\chi \in \widehat{G}} H_\chi, \qquad (26)$$

where $f \in H_\chi$ if $f(\{\vec{x}_m\}, \sigma a) = \chi(\sigma^{-1}) f(\{\vec{x}_m\}, a)$ for all $\sigma \in G$. The decomposition is invariant under the action of $\mathbb{Z}^d$ on $\Sigma \times \mathcal{A}$.

Note that the condition $\sum_{\chi \in \widehat{G}} \chi(\sigma) = 0$ for any $\sigma \in G$ is automatically satisfied if G is an abelian group.

Proof. Let $f \in L^2(\Sigma \times \mathcal{A}, \mu_\Sigma \times \mu_\mathcal{A})$. Our goal is to decompose f into a sum of functions that behave nicely with respect to letter inversions. We begin by noting that

$$\frac{1}{|\widehat{G}|} \sum_{\chi \in \widehat{G}} \chi(id) f(\{\vec{x}_m\}, a) = f(\{\vec{x}_m\}, a),$$

since $\chi(id)$ is always equal to 1. Whenever $\sigma \neq id$ we can write

$$\frac{1}{|\widehat{G}|} \sum_{\chi \in \widehat{G}} \chi(\sigma) f(\{\vec{x}_m\}, \sigma a) = \frac{1}{|\widehat{G}|} f(\{\vec{x}_m\}, \sigma a) \sum_{\chi \in \widehat{G}} \chi(\sigma) = 0.$$

Letting σ range through G we see that

$$f(\{\vec{x}_m\}, a) = \frac{1}{|\widehat{G}|} \sum_{\chi \in \widehat{G}} \sum_{\sigma \in G} \chi(\sigma) f(\{\vec{x}_m\}, \sigma a). \qquad (27)$$

Write

$$f_\chi(\{\vec{x}_m\}, a) = \sum_{\sigma \in G} \chi(\sigma) f(\{\vec{x}_m\}, \sigma a)$$

so that the above decomposition is given by $f = \frac{1}{|G|} \sum_{\chi \in \widehat{G}} f_\chi$. We show that $f_\chi \in H_\chi$. Fixing $\rho \in G$, we have that $\sigma\rho \in G$ for all $\sigma \in G$, and furthermore $\sigma\rho$ ranges through G as σ does. Thus

$$f_\chi\big(\{\vec{x}_m\}, \rho a\big) = \sum_{\sigma \in G} \chi(\sigma) f\big(\{\vec{x}_m\}, \sigma\rho a\big) = \sum_{\sigma \in G} \chi(\sigma\rho)\chi\big(\rho^{-1}\big) f\big(\{\vec{x}_m\}, \sigma\rho a\big)$$

$$= \chi\big(\rho^{-1}\big) \sum_{\sigma\rho \in G} \chi(\sigma\rho) f\big(\{\vec{x}_m\}, \sigma\rho a\big) = \chi\big(\rho^{-1}\big) f_\chi\big(\{\vec{x}_m\}, a\big).$$

Next we show that if $\chi \neq \zeta$ for $\chi, \zeta \in \widehat{G}$, then H_χ is orthogonal to H_ζ. Let $f \in H_\chi$ and $g \in H_\zeta$, and let $\sigma \in G$ such that $\chi(\sigma) \neq \zeta(\sigma)$. Using change of variables we have

$$\int\limits_{\Sigma \times \mathcal{A}} f\big(\{\vec{x}_m\}, a\big) \overline{g\big(\{\vec{x}_m\}, a\big)} \, \mathrm{d}(\mu_\Sigma \times \mu_\mathcal{A})$$

$$= \int\limits_{\Sigma \times \mathcal{A}} f\big(\{\vec{x}_m\}, \sigma a\big) \overline{g\big(\{\vec{x}_m\}, \sigma a\big)} \, \mathrm{d}(\mu_\Sigma \times \mu_\mathcal{A})$$

$$= \int\limits_{\Sigma \times \mathcal{A}} \chi(\sigma^{-1}) f\big(\{\vec{x}_m\}, a\big) \overline{\zeta\big(\sigma^{-1}\big) g\big(\{\vec{x}_m\}, a\big)} \, \mathrm{d}(\mu_\Sigma \times \mu_\mathcal{A}).$$

Since $\overline{\zeta(\sigma^{-1})} = \zeta(\sigma)$, we have shown that $\langle f, g \rangle = \chi(\sigma^{-1})\zeta(\sigma)\langle f, g \rangle$, and by choice of σ this can only happen if $\langle f, g \rangle = 0$. Thus f and g are orthogonal and we have shown that the decomposition is a direct sum.

Finally we show that each subspace H_χ is invariant under the action of $\mathbb{Z}^d$. For $f \in H_\chi$, and $i \in \{1, 2, \ldots, d\}$, we show that the function $\widehat{U}_i f$ is in H_χ, where $\widehat{U}_i f(\{\vec{x}_m\}, a) = f(V_i(\{\vec{x}_m\}), \phi_i(\{\vec{x}_m\})a)$. Since σ commutes with translation, we have that

$$\widehat{U}_i f\big(\{\vec{x}_m\}, \sigma a\big) = f\big(V_i\big(\{\vec{x}_m\}\big), \phi_i\big(\{\vec{x}_m\}\big)\sigma a\big) = f\big(V_i\big(\{\vec{x}_m\}\big), \sigma\phi_i\big(\{\vec{x}_m\}\big)a\big)$$

$$= \chi\big(\sigma^{-1}\big) f\big(V_i\big(\{\vec{x}_m\}\big), \phi_i\big(\{\vec{x}_m\}\big)a\big) = \chi\big(\sigma^{-1}\big)\widehat{U}_i f\big(\{\vec{x}_m\}, a\big).$$

Thus $\widehat{U}_i f \in H_\chi$, which finishes the proof. $\quad\square$

5. Examples: Applications of Theorem 4.2

The condition for decomposition of the spectrum of the dynamical system of a bijective substitution $\mathcal{S}$ in Theorem 4.2 can manifest itself in a variety of ways. Interesting finite groups from the symmetric group on $\mathcal{A}$ can be used to form the substitution permutation matrix for $\mathcal{S}$. We will show only a few examples in the two- and four-letter cases.

5.1. Two letter alphabets

It is convenient to let $\mathcal{A} = \{1, -1\}$, so that the two bijections in $S_\mathcal{A}$ can be seen as multiplication by 1 (denoted g_0) or -1 (denoted g_1). The permutations $p_{\vec{j}}$ of the substitution

matrix must include both of these elements to ensure that substitution is nonperiodic. In this case the translation commuting letter inversions are $G = \{g_0, g_1\}$ and the character group $\widehat{G}$ is:

	g_0	g_1
χ_0	1	1
χ_1	1	-1

Assuming it is nonperiodic, the substitution $\mathcal{S}$ will satisfy the conditions of Theorem 4.2. Therefore the isomorphic skew product space will have spectral decomposition $L^2(\Sigma \times \mathcal{A}, \mu_\Sigma \times \mu_\mathcal{A}) = H_0 \oplus H_1$. Any $f \in L^2(\Sigma \times \mathcal{A}, \mu_\Sigma \times \mu_\mathcal{A})$ will be written as the sum

$$f(\{\vec{x}_m\}, a) = \frac{f(\{\vec{x}_m\}, a) + f(\{\vec{x}_m\}, -a)}{2} + \frac{f(\{\vec{x}_m\}, a) - f(\{\vec{x}_m\}, -a)}{2}.$$

Of course functions in $L^2(X, \mu)$ can be decomposed by considering their counterparts on the skew product space. Thus one can consider very basic functions such as indicator functions $\mathbb{I}_1$ or $\mathbb{I}_{-1}$ and find that their decompositions into discrete and continuous parts are:

$$\mathbb{I}_1 = 1/2 + (\mathbb{I}_1 - 1/2), \qquad \mathbb{I}_{-1} = 1/2 + (\mathbb{I}_{-1} - 1/2).$$

In unpublished work on two specific $\mathbb{Z}^2$ substitutions, we have found that the spectral measures associated to the continuous spectrum functions $\mathbb{I}_1 - 1/2$ and $\mathbb{I}_{-1} - 1/2$ are singular with respect to Lebesgue measure.

5.2. Four-letter alphabets

The next four examples involve bijective substitutions with $\mathcal{A} = \{0, 1, 2, 3\}$. The difference in their construction lies in the letter inversion groups allowed by the set of bijections $\{p_{\vec{j}}, \vec{j} \in \mathcal{I}^d\}$. In each example we will have $d = 2$ and $l_1 = l_2 = 3$, and we keep the substitution permutation matrix $(p_{\vec{j}})_{\vec{j} \in \mathcal{I}^2}$ fixed, changing the actual bijections in each case. The substitution permutation matrix will be:

$$(p_{\vec{j}})_{\vec{j} \in \mathcal{I}^2} = \begin{pmatrix} g_0 & g_1 & g_0 \\ g_1 & g_2 & g_1 \\ g_0 & g_1 & g_0 \end{pmatrix}. \tag{28}$$

Examples 6–9 will use the following four choices for the permutations, respectively (g_0, g_1, g_2 expressed in cycle notation):

$$P_1 = \{id, (01)(23), (02)(13)\},$$
$$P_2 = \{id, (0123), (02)(13)\},$$
$$P_3 = \{id, (01)(23), (0123)\},$$
$$P_4 = \{id, (01)(23), (023)\}.$$

The interested reader can check that, when P_1 is used, the substitution matrices on the individual letters become:

Fig. 5. With P_1 there are four translation-commuting letter inversions.

Fig. 6. With P_2 the letter inversions form an order-four cyclic group.

$$0 \to \begin{pmatrix} 0 & 1 & 0 \\ 1 & 2 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \quad 1 \to \begin{pmatrix} 1 & 0 & 1 \\ 0 & 3 & 0 \\ 1 & 0 & 1 \end{pmatrix}, \quad 2 \to \begin{pmatrix} 2 & 3 & 2 \\ 3 & 0 & 3 \\ 2 & 3 & 2 \end{pmatrix}, \quad 3 \to \begin{pmatrix} 3 & 2 & 3 \\ 2 & 1 & 2 \\ 3 & 2 & 3 \end{pmatrix}.$$

Using grey-scale tiles to denote the letters, from 0 being white to 3 being black, we can look at the tilings that result from these choices. In each case Figs. 5–8 show four iterations of the 0 tile.

Fig. 7. Using P_3 there are only two letter inversions.

Fig. 8. When P_4 is used there are no nontrivial letter inversions.

Example 6. When P_1 is used it is not difficult to check that $g_1 \circ g_2 = g_3 = (03)(12)$, and that these four bijections form an abelian subgroup of the symmetric group on $\mathcal{A}$ of order four. Since that is the order of $\mathcal{A}$, we know that the group of translation-commuting letter inversions must be $G = \{g_0, g_1, g_2, g_3\}$.

In order to use Theorem 4.2 we need the group characters for G:

	g_0	g_1	g_2	g_3
χ_0	1	1	1	1
χ_1	1	1	-1	-1
χ_2	1	-1	1	-1
χ_3	1	-1	-1	1

This substitution has nontrivial height lattice Λ generated by the vectors $(2, 0)$ and $(1, 1)$. One point of particular interest involves the factors onto two-letter sequences. We can choose two letters to become 0 and let the other two become 1, and for two such choices the resulting factor sequence will be a substitution sequence with a bijective substitution. The remaining choice is periodic and that is due to the nontrivial height of this substitution.

Example 7. We begin by noting that the permutations in P_2 form a cyclic group of order four by adding $g_1 \circ g_2 = g_3 = (0321)$, so this forms the group G of letter inversions. In order to use Theorem 4.2, we again compute the group characters for G:

	g_0	g_1	g_2	g_3
χ_0	1	1	1	1
χ_1	1	i	-1	$-i$
χ_2	1	-1	1	-1
χ_3	1	$-i$	-1	i

This substitution also has nontrivial height lattice Λ generated by the vectors $(2, 0)$ and $(1, 1)$. With the choices of substitution permutation matrix and specific permutations made in this example, our system is the factor of a direct product of one-dimensional substitutions. To see this, set $\mathcal{S}'$ to be the one-dimensional substitution taking $0 \to 010$, $1 \to 121$, $2 \to 232$, and $3 \to 303$. Taking the alphabet $\mathcal{A}'$ to be the set of ordered pairs (j, k), where $j, k \in \{0, 1, 2, 3\}$, we can construct a two-dimensional substitution on 16 letters as the direct product of $\mathcal{S}'$ with itself. This substitution factors onto $\mathcal{S}$ by taking any letter (j, k) onto $j + k \bmod 4$. One should note that a different choice of substitution permutation matrix may avoid this problem.

Example 8. We begin by noting that the bijections in P_4 form a nonabelian group of order 8. The only nontrivial group element from the symmetric group on $\mathcal{A}$ that commutes with g_1 and g_2 is $\sigma_1 = (02)(13)$, and so the group G is composed of $\sigma_0 = identity$ and σ_1. Using the substitution permutation matrix (28), we again get a primitive, nonperiodic substitution, and so Theorem 4.2 does apply. The character table for G is given by:

	σ_0	σ_1
χ_0	1	1
χ_1	1	-1

This example is intriguing because the continuous component should be more complex than the individual components in the previous two examples. Further analysis should be done to determine the spectral multiplicity of the system.

Example 9. With P_4 as the choice of permutations there are no elements of the symmetric group on $\mathcal{A}$ that commute with both g_1 and g_2, and so there can be no translation-commuting letter inversions for $\mathcal{S}$. There is no character table to compute for this substitution and Theorem 4.2 does not apply.

It is unclear to what extent the qualitative differences between these four examples are due to the presence, absence, and nature of translation-commuting letter inversions.

6. Related questions

What accounts for the striking differences in appearance among substitution sequences that have the same spectral decomposition as determined by Theorem 4.2? One can experiment with the exact same set of bijections $\{p_{\bar{j}}\}$, but place them differently in $\mathcal{I}^d$, and find that the resulting substitution sequences can range from having large connected components and/or obvious, coherent patterns to appearing quite randomly disordered. All the while, the spectral information and frequency measures are identical! The same basic question also extends to a discussion of the diffraction (spectral) images of substitution sequences, which are related to the correlation measures. The interested reader can use our MATLAB program [5] to experiment with $\mathbb{Z}^2$ substitution sequences (up to $|\mathcal{A}| = 10$); the program allows manipulation of many aspects of the systems, some of which are beyond the scope of this paper.

We have obtained, in Theorem 4.2, a decomposition of the spectrum in the event that there exist a nontrivial group of letter inversions. Even so, it is unclear exactly what the nature of the invariant pieces are. The spectral multiplicity is almost certainly bounded by $|\mathcal{A}|$. In the event that the group formed by the substitution bijections $p_{\bar{j}}$ is abelian, it may be that each invariant piece is simple, and may have singular spectrum (this is what happens in the one-dimensional case). If the situation is more like that in Example 8, where $|G| < |\mathcal{A}|$, it would appear that the invariant continuous component(s) are probably more complicated. And in examples like Example 9, our theorem yields no information whatsoever—there is one continuous invariant component that seems to contain all of the continuous spectrum functions. It should be noted that the functions in the continuous spectrum part with the form of (22) are in a sense transverse to the decomposition in Theorem 4.2, and further investigation into their nature would certainly help answer the question about the nature of the spectrum. The work in [16] should be used as a guide. Finally, it would be interesting to give examples of constant-length substitution sequences with continuous spectrum containing both the Lebesgue and singular types.

References

[1] F.M. Dekking, The spectrum of dynamical systems arising from substitutions of constant length, Z. Wahr. 41 (1978) 221–239.

[2] S. Ferenczi, Tiling and local rank properties of the Morse sequence, Theoret. Comput. Sci. 129 (2) (1994) 369–383.

[3] N. Pytheas Fogg, Substitutions in Dynamics, Arithmetics, and Combinatorics, Lecture Notes in Mathematics, vol. 1794, Springer, Berlin, 2002.

[4] N.P. Frank, Substitution sequences in $\mathbb{Z}^d$ with a nonsimple Lebesgue component in the spectrum, Ergodic Theory Dynamical Systems 23 (2) (2003) 519–532.

[5] N.P. Frank, B.J. Rinaldi, Tiling computer program, MATLAB program allowing experimentation with substitution sequences in $\mathbb{Z}^2$. Available by request, 2003.

[6] G.R. Goodson, On the spectral multiplicity of a class of finite rank transformations, Proc. Amer. Math. Soc. 93 (1985) 303–306.

[7] G.R. Goodson, A survey of recent results in the spectral theory of ergodic dynamical systems, J. Dyn. Control Systems 5 (1999) 173–226.

[8] S. Kakutani, Ergodic theory of shift transformations, in: Proc. 5th Berkeley Symp., vol. 2, 1967, pp. 405–414.

[9] M. Keane, Generalized Morse sequences, Z. Wahrsch. Verw. Gebiete 10 (1968) 335–353.

[10] M. Keane, Strongly mixing *g*-measures, Invent. Math. 16 (1972) 309–324.

[11] J. Kwiatkowski, Spectral isomorphism of Morse dynamical systems, Bull. Pol. Acad. Sci. 29 (1981) 105–114.

[12] P. Michel, Stricte ergodicite d'ensembles minimaux de substitutions, C. R. Acad. Sci. Paris 278 (1974) 811–813.

[13] B. Mossé, Reconnaissabilité des substitutions et complexité des suites automatiques, Bull. Soc. Math. France 124 (1996) 329–346.

[14] K. Petersen, Ergodic Theory, Cambridge University Press, Cambridge, 1990.

[15] B. Praggastis, Markov partitions for hyperbolic toral automorphisms, Ph.D. Thesis, University of Washington, 1992.

[16] M. Queffelec, Substitution Dynamical Systems—Spectral Analysis, Lecture Notes in Mathematics, vol. 1294, Springer, Berlin, 1987.

[17] C. Radin, M. Wolff, Space tilings and local isomorphism, Geom. Dedicata 42 (1992) 355–360.

[18] E.A. Robinson Jr., On the table and the chair, Indag. Math. 10 (4) (1999) 581–599.

[19] W. Rudin, Fourier Analysis on Groups, Interscience, New York, 1962.

[20] E. Seneta, Non-Negative Matrices and Markov Chains, Springer, New York, 1981.

[21] B. Solomyak, Dynamics of self-similar tilings, Ergodic Theory Dynamical Systems 17 (1997) 695–738.

[22] B. Solomyak, Non-periodicity implies unique composition for self-similar transationally-finite tilings, J. Discrete Comput. Geometry 20 (2) (1998) 265–279.