Chapter 1
Introduction to hierarchical tiling dynamical systems

Natalie Priebe Frank*

Abstract This paper is about the tiling dynamical systems approach to the study of aperiodic order. We compare and contrast four related types of systems: ordinary (one-dimensional) symbolic systems, one-dimensional tiling systems, multidimensional $\mathbb{Z}^d$-systems, and multidimensional tiling systems. Aperiodically ordered structures are often hierarchical in nature, and there are a number of different yet related ways to define them. We will focus on what we are calling "supertile construction methods": symbolic substitution in one and many dimensions, S-adic sequences, self-similar and pseudo-self-similar tilings, and fusion rules. The techniques of dynamical analysis of these systems are discussed and a number of results are surveyed. We conclude with a discussion of the spectral theory of supertile systems from both the dynamical and diffraction perspectives.

1.1 Introduction

The central objects in these lecture notes are tilings constructed via a variety of methods that together we call *supertile methods*. These tilings display hierarchical structure that is highly ordered yet not periodic. Their study is truly multidisciplinary, having originated in fields as disparate as logic, chemistry and geometry. To motivate the topic we offer three examples from the history of the field that are relevant to these lectures.

Natalie Priebe Frank, Professor of Mathematics
Department of Mathematics and Statistics, Vassar College, e-mail: nafrank@vassar.edu

* The author would like to thank Michael Baake, Franz Gähler, E. A. Robinson, Jr., Dan Rust, Lorenzo Sadun, and Boris Solomyak for their comments on drafts of this work.

First, imagine square tiles whose edges come in given combinations of colors, and you are only allowed to put two tiles next to each other if the edge colors match. Can you make an infinite tiling of the plane with these tiles? This is the question logician Hao Wang was considering in 1961 [105]. In particular he was thinking about the decidability of what is now known as the domino problem: "Given a finite set of tiles in the plane, can it be made to form an infinite tiling?" The answer depended on whether an *aperiodic prototile set* exists, i.e. a set of tiles that are able to form an infinite tiling of the plane, but every tiling they make must be nonperiodic. The question was proved to be undecidable by Robert Berger [22] with the discovery of an aperiodic set of prototiles. That prototile set had over 20000 tiles in it, but in 1971 Raphael Robinson published an aperiodic set with only 6 tiles [94]. In Robinson's version the hierarchical structure is clearly evident and in fact drives the proof of aperiodicity. In this volume in [63] we find four proofs of undecidability, including how to construct the aperiodic tile set(s).

A second development, which entered the public consciousness through a *Scientific American* article by Martin Gardner [54], was Penrose's 1974 discovery of an aperiodic set of two tiles. In the middle of the 20th century, Roger Penrose began to develop an interest in tiling questions in part because of Hilbert's Problem 18. The interest intensified as Penrose and his father developed a collaboration with M. C. Esher (see the foreword to [11]). Penrose was trying to create a hierarchical tiling and found his original tiling (which in that foreword he tells us is [82, Fig. 4]) by experimentation.

There are a number of versions of Penrose tilings, all of which can be generated by a supertile method. In figures 1.1 and 1.2 we show a *pseudo-self-similar* version (see section 1.4.4.2), for which the tiles also form an aperiodic tile set. Figure 1.1 shows the rule for inflating and replacing the tiles, and figure 1.2 shows the result of inflating and replacing a central patch twice.

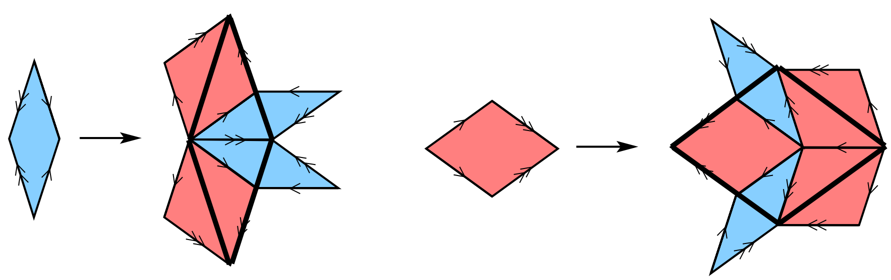

Fig. 1.1 The Penrose rhombuses and their inflate-and-subdivide rules.

The third and possibly most invigorating development we mention here is the discovery of quasicrystals in 1982 [97]. This earned Dan Shechtman the Wolf Prize in Physics in 1999 and the Nobel Prize in Chemistry in 2011 [73]. In his laboratory in what is now the U. S. National Institute of Standards and

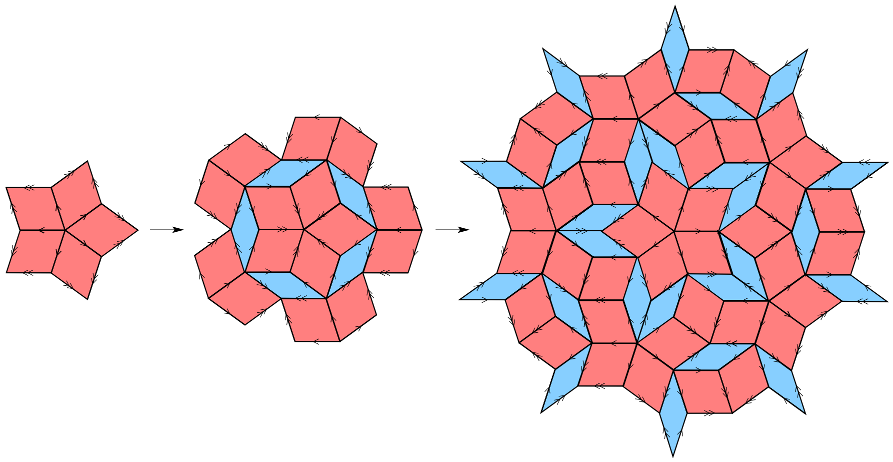

Fig. 1.2 1-, 2-, and 3-supertiles for the Penrose rhombus tiling.

Technology, Shechtman analyzed an aluminum-magnesium alloy and found that its diffraction image revealed contradictory properties: it had bright spots indicative of a periodic atomic structure, but had symmetries impossible for such a structure. The discovery went against all conventional wisdom at the time, but eventually the scientific community accepted that there was no mistake, this alloy did indeed display 'quasi'-crystalline structure. In some of the images in figure 1.3 one can see the 'forbidden' 10-fold rotational symmetries.

Coincidentally, in 1982 Alan Mackay [77] published the diffraction pattern of a Penrose tiling, shown in its original form in figure 1.4. Once Shechtman's diffraction pattern was published, it did not take long for similarities between it and Mackay's to be noticed. This established Penrose tilings and highly structured tilings like them (including some generated by supertile methods), as mathematical models of quasicrystals. It is apparent that spectral methods, then, are an interesting way to study aperiodic tilings. Spectral analysis, including mathematical diffraction, has proved to be an effective tool for the study of tilings generated by supertile methods. The last three sections of these notes discuss spectral theory from dynamical and diffraction perspectives.

1.1.1 Outline of the paper

In section 1.2 we begin by giving general definitions of the four types of structures of interest and the basic relationships between them. Specifics of why and how the dynamical systems viewpoint is used appears in section 1.3. In

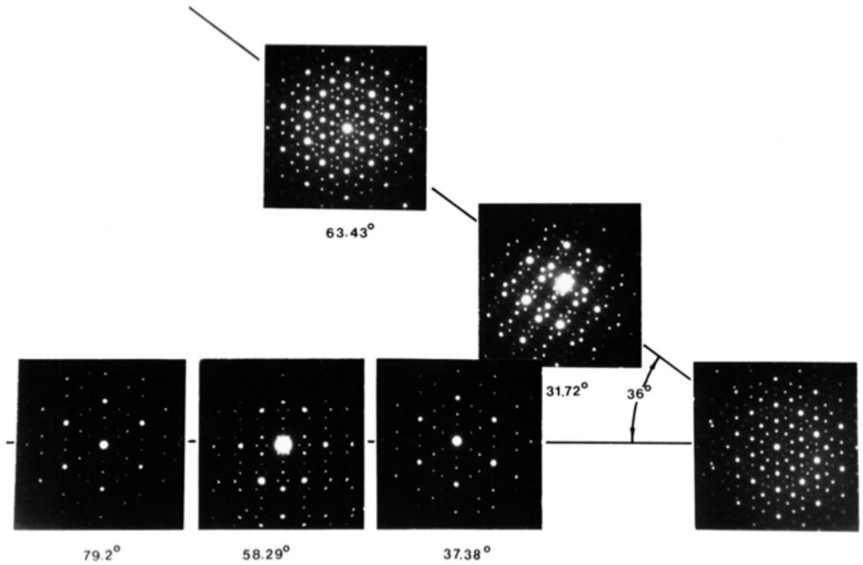

FIG. 2. Selected-area electron diffraction patterns taken from a single grain of the icosahedral phase. Rotations match those in Fig. 1.

Fig. 1.3 The quasicrystal diffraction images as they appear in the original paper [97].

this section we compare and contrast how the metrics are related, show how standard dynamical properties like minimality can be interepreted, and talk about invariant measures and their connection to the idea of frequency. In section 1.4 we learn about the various supertile construction methods and give examples of many of them. In section 1.5 we introduce the idea of transition matrices and how their properties allow us to extend dynamical results to supertile systems. Section 1.6 is devoted to the dynamical spectrum of supertile systems, while section 1.7 is devoted to their diffraction spectra. Section 1.8 presents results on the connection between the two types of spectrum. We conclude in section 1.9 with a selection of references.

1.1.2 Not covered

The field of aperiodic order and tiling dynamical systems spans a broad range of topics and we have not attempted to give a complete survey. Topics we do not discuss include tiling cohomology, matching rules, the projection method, tilings with infinite local complexity, K-theory of C^*-algebras for tilings, spectral triples, decidability and tiling problem questions. Moreover

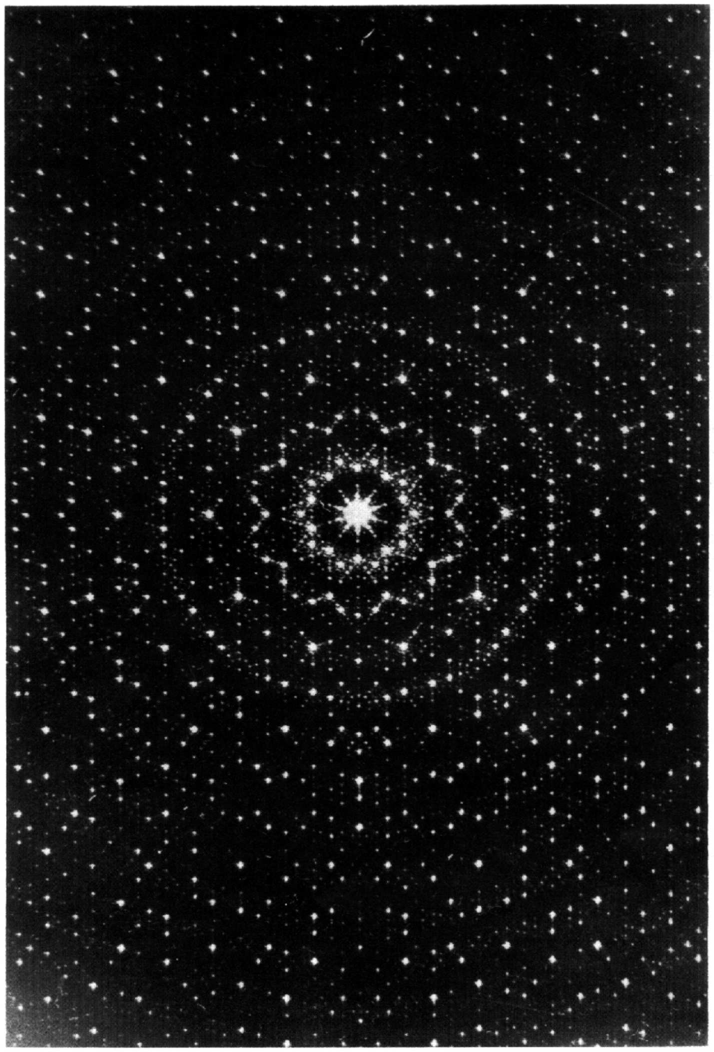

Fig. 5. The optical transform of the pattern of fig. 4. The annular objects show in the circular strong and weak modulations of the transform. The pattern itself exhibits local ten-fold symmetry and repeats the shape of the quasi-lattice cells which gave rise to it.

Fig. 1.4 The Penrose diffraction image as it appears in the original paper [77].

we do not consider tilings of hyperbolic space or other spaces, or anything about the spectrum of Schrodinger operators modeled on tilings.

1.2 The fundamental objects

1.2.1 Motivation: shift spaces

The way that we study tiling spaces is a generalization of symbolic dynamics, a large branch of dynamical systems theory. Thus we begin by describing the basic setup in this situation.

Let $\mathcal{A}$ be a finite set we will call our *alphabet*. A *sequence* is a function $\mathbf{x} : \mathbb{Z} \to \mathcal{A}$ and we denote the set of all sequences to be $\mathcal{A}^{\mathbb{Z}}$. We equip the space with a metric that defines the product topology, as follows. Let $N(\mathbf{x}, \mathbf{y}) = \min\{n \geq 0 \text{ such that } \mathbf{x}(j) \neq \mathbf{y}(j) \text{ for some } |j| = n\}$, and define $d(\mathbf{x}, \mathbf{y})$ to be $\exp(-N(\mathbf{x}, \mathbf{y}))$. That is, $\mathbf{x}$ and $\mathbf{y}$ are very close if they agree on a large ball centered at the origin.

For each $j \in \mathbb{Z}$, we can shift a sequence $\mathbf{x}$ by j, yielding the sequence $\mathbf{x} - j$ defined by $(\mathbf{x} - j)(k) = \mathbf{x}(k + j)$. This is known as the shift map, and is a $\mathbb{Z}$-action on sequences. (Notice that $(\mathbf{x} - j)(0) = \mathbf{x}(j)$, meaning that $\mathbf{x}$ has been shifted so that what was at j is now at the origin.) The space $\mathcal{A}^{\mathbb{Z}}$ along with the shift map is known as the *full shift* on $|\mathcal{A}|$ symbols.[2] The shift map allows us to investigate the long-range structure of sequences by moving distant parts 'into view' of the origin. This perspective is consistent with our choice of metric topology.

There is already some dynamics to study for the full shift, but things get much more interesting when we restrict our attention to closed, nonempty, shift-invariant subsets $\Omega \subset \mathcal{A}^{\mathbb{Z}}$. We call such an Ω a *shift space* or *subshift of* $\mathcal{A}^{\mathbb{Z}}$ and use the terminology *shift dynamical system* for $(\Omega, \mathbb{Z})$. Subshifts are handy tools for encoding the dynamics of many types of systems, and they also arise in natural processes. Readers interested in diving into the vast literature on this subject might find [6, 67, 74, 76, 42] in their libraries.

Example 1 *Let Ω consist of the periodic sequences* ...0101.0101... *and its shift* ...1010.1010..., *where the decimal point is there to denote where the origin is. Shifting by $j \in \mathbb{Z}$ just moves the decimal point j units to the right (or left, if j is negative). One sees quickly that Ω is shift-invariant and that the dynamical system is periodic with period 2.*

Because periodic systems like these are completely understood we will tend to assume that the sequences in our sequence spaces are not periodic. Instead, supertile construction methods generate sequences with just the right amount of long-range order to be tractable for analysis.

---

[2] Ordinarily in the literature the shift map is given with notation like $\sigma(\mathbf{x})$, so that $\mathbf{x} - j = \sigma^j(\mathbf{x})$. We use the notation "$\mathbf{x} - j$" instead to be consistent with the more general case.

1.2.2 Straightforward generalization: sequences in $\mathbb{Z}^{\boldsymbol{d}}$

Let $\mathcal{A}$ be a finite alphabet and consider $\mathcal{A}^{\mathbb{Z}^d}$ to be the set of all sequences in $\mathbb{Z}^d$, that is, functions from $\mathbb{Z}^d$ to $\mathcal{A}$. Given $\mathbf{x}, \mathbf{y} \in \mathcal{A}^{\mathbb{Z}^d}$, let $N(\mathbf{x}, \mathbf{y}) = \min\{n \geq 0$ such that $\mathbf{x}(\boldsymbol{\jmath}) \neq \mathbf{y}(\boldsymbol{\jmath})$ for some $|\boldsymbol{\jmath}| = n\}$, where $|\boldsymbol{\jmath}|$ is the largest absolute value of the components of $\boldsymbol{\jmath}$. Then $d(\mathbf{x}, \mathbf{y}) = \exp(-N(\mathbf{x}, \mathbf{y}))$ provides an origin-centric metric as before.

Translation by elements of $\mathbb{Z}^d$ is defined as before and provides a way to analyze the structure of multidimensional sequences. There are complications and considerations due to the extra dimensions that we will discuss as we encounter them.

1.2.3 Straightforward generalization: tilings of $\mathbb{R}$

We choose a closed interval for each symbol in $\mathcal{A}$. For any element $\mathbf{x} \in \mathcal{A}^{\mathbb{Z}}$ make a tiling by placing the interval corresponding to $\mathbf{x}(0)$ with its left endpoint at 0, and placing copies of all the other symbols of $\mathbf{x}$ in the corresponding order, with overlap at the interval endpoints. In this perspective a tile is a closed interval labelled by an element of $\mathcal{A}$. Tiles and tilings can be translated by elements of $\mathbb{R}$ and there is an origin-centric tiling metric that we will describe in the general situation in the next section.

Figure 1.5 depicts a tiling of $\mathbb{R}$ with two tile types, a longer interval pictured in dark blue and a shorter interval pictured in light blue. (The colors represent the labels). The patch shown corresponds to the symbolic sequence *...abbbaaaabbbabbbabbba...*

Fig. 1.5 A patch of a one-dimensional tiling with two tile lengths.

1.2.4 Geometric generalization: tilings of $\mathbb{R}^{\boldsymbol{d}}$

The finite alphabet $\mathcal{A}$ is replaced by a finite *prototile set* $\mathcal{P}$. A prototile $p \in \mathcal{P}$ is a closed topological disk in $\mathbb{R}^d$ carrying a label (for instance, a color). The closed set is known as the *support* of p (denoted $\text{supp}(p)$) and the label is there to distinguish any tiles that may have congruent shapes. We can apply any self-map of $\mathbb{R}^d$ to a prototile by applying it to the support and carrying the label along. Although it is common to use some or all elements of the Euclidean group to move tiles around, we restrict our attention to translations only. A $\mathcal{P}$-*tile* or just *tile* is any translate of a prototile from $\mathcal{P}$. Two tiles are *equivalent* if their supports are translates of each other and they carry the same label.

Consider some fixed set of prototiles $\mathcal{P}$. A $\mathcal{P}$-*patch* (or *patch* when the prototile set is understood) is a set of tiles that intersect at most on their boundaries that is supported on a connected set in $\mathbb{R}^d$. For technical reasons it is often assumed that the supports form a topological disk. A $\mathcal{P}$-*tiling* or just *tiling* of $\mathbb{R}^d$ is a collection of tiles that 'cover' $\mathbb{R}^d$ in the sense that the union of the tile supports is $\mathbb{R}^d$, but also 'pack' $\mathbb{R}^d$ in the sense that any two supports intersect only on their boundaries. Let $\Omega_{\mathcal{P}}$ denote the space of all $\mathcal{P}$-tilings. As with the full shift $\mathcal{A}^{\mathbb{Z}}$, in which an element is an infinite sequence, elements of $\Omega_{\mathcal{P}}$ are infinite tilings of $\mathbb{R}^d$.

Like tiles, patches and tilings can be translated by elements of $\mathbb{R}^d$. We write $\mathcal{T} - \mathbf{v}$ to denote the tiling obtained by translating the support of every tile of $\mathcal{T}$ by $\mathbf{v}$. Note that the origin in $\mathcal{T} - \mathbf{v}$ corresponds to $\mathbf{v}$ in $\mathcal{T}$, so this translation brings the neighborhood of $\mathbf{v}$ into view of the origin.

Analogous to the simpler cases, we say a tiling $\mathcal{T}$ is *nonperiodic* if there is no $\mathbf{v}$ for which $\mathcal{T} - \mathbf{v} = \mathcal{T}$. In higher dimensions it is possible to be periodic in some directions but not *fully periodic*: the directions of periodicity must form a basis for $\mathbb{R}^d$ for full periodicity.

Now geometry plays a fundamental role, and there is the possibility that tiles in a tiling can be adjacent in many different ways. Consider the tiling in figure 1.6, which is constructed from a set of four rectangular tiles, with side lengths given by 1 and $\frac{1+\sqrt{17}}{2}$. There are many offsets where vertices meet edges, and the number of those offsets will go to infinity as we consider larger and larger patches.

Definition 1. We say a tiling $\mathcal{T} \in \Omega_{\mathcal{P}}$ has *finite local complexity (FLC)* if it contains only finitely many two-tile patches up to translation. A subset of tilings $\Omega \subset \Omega_{\mathcal{P}}$ is said to have finite local complexity if there are only finitely many two-tile patches up to translation in Ω.

For the purposes of this work, we assume finite local complexity in all tilings and tiling spaces unless otherwise stated.[3]

It is convenient to have notation for the patch of a tiling that intersects a subset of $\mathbb{R}^d$. Let $\mathcal{T}$ be a tiling of $\mathbb{R}^d$ and let $B \subset \mathbb{R}^d$. The patch of tiles in $\mathcal{T}$ whose supports intersect B is denoted $\mathcal{T} \cap B$. One could say that $\mathcal{T}$ has finite local complexity if the set of patches

$$\left\{\mathcal{T} \cap \{x\} \text{ such that } x \in \mathbb{R}^d\right\}$$

is finite up to translation.

[3] FLC is a common restriction, but if you want to learn about the infinite local complexity case see [47] and references within.

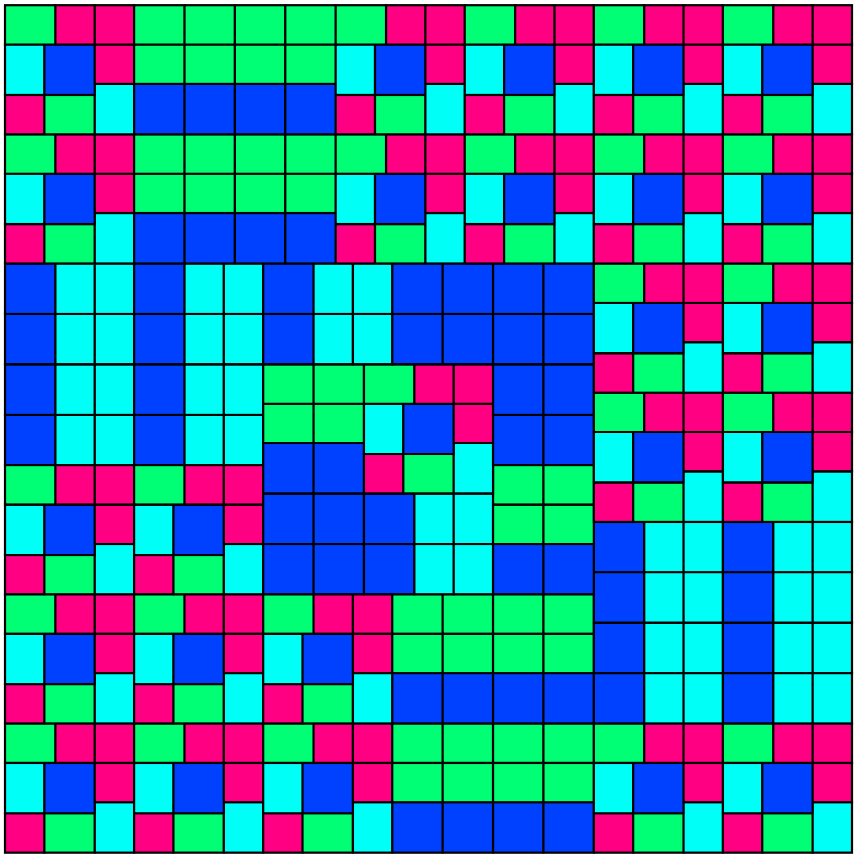

Fig. 1.6 A patch of a tiling with four prototiles and infinite local complexity.

1.3 The dynamical systems viewpoint

1.3.1 Tiling spaces

A tiling space is a translation-invariant subset Ω of the full tiling space $\Omega_{\mathcal{P}}$ that is closed in the metric we describe next. We form a dynamical system by letting an additive subgroup G of $\mathbb{R}^d$ act on it by translation. Ordinarily G is just $\mathbb{R}^d$ itself, but occasionally it might be $\mathbb{Z}^d$ or some other full rank subgroup. We use the notation (Ω, G) to denote the tiling dynamical system.

1.3.1.1 The "big ball" metric

The metric used in dynamical systems theory for tilings is modeled on the metric for shift spaces, and therefore is also origin-centric. The definition of the metric becomes technical because translation is a continuous action and because the prototiles can have interesting geometry. Still the basic idea is that two tilings are close if they very nearly agree on a ball around the origin.

We give the definition of metric for tilings of finite local complexity. Let $\mathcal{T}$ and $\mathcal{T}'$ be tilings of $\mathbb{R}^d$ from a prototile set $\mathcal{P}$. Informally, we say $\mathcal{T}$ and $\mathcal{T}'$ are within ϵ of one another if they agree on a ball of radius $1/\epsilon$, except for a small translation. Here is a formal definition.

Definition 2. Let $R(\mathcal{T},\mathcal{T}')$ be the supremum of all $r \geq 0$ such that there exists $\mathbf{x},\mathbf{y} \in \mathbb{R}^d$ with

1. $|\mathbf{x}| < 1/2r$ and $|\mathbf{y}| < 1/2r$, and
2. On the ball of radius r around the origin, $(\mathcal{T} - \mathbf{x}) \cap B_r(0) = (\mathcal{T}' - \mathbf{y}) \cap B_r(0)$.

We define
$$d(\mathcal{T},\mathcal{T}') := \min\left\{\frac{1}{R(\mathcal{T},\mathcal{T}')}, 1\right\}.$$

There are various versions in the literature; this version parallels [100].

1.3.1.2 Two common ways to construct tiling spaces

There are two main ways that tiling spaces are constructed. One is to make a closed, translation-invariant space around a given tiling; in this situation the space is called the 'hull' of the tiling. The other is to specify an 'atlas' of allowed patches and include tilings that contain these patches only.

For the first method, suppose there is some tiling $\mathcal{T} \in \Omega_{\mathcal{P}}$ that is of particular interest. We can construct the *hull of the tiling* $\mathcal{T}$ as the orbit closure of $\mathcal{T}$:
$$\Omega_{\mathcal{T}} = \overline{\{\mathcal{T} - \mathbf{v} \text{ for all } \mathbf{v}\}}.$$
By definition it is closed and it is not difficult to show that it is translation-invariant.

The second method is akin to making a shift space from a language of allowed words. Let $\mathcal{R}$ be a set of $\mathcal{P}$-patches. We say that $\mathcal{T} \in \Omega_{\mathcal{P}}$ is *allowed* by $\mathcal{R}$ if every patch in $\mathcal{T}$ is translation-equivalent to a subpatch of an element of $\mathcal{R}$. The tiling space $\Omega_{\mathcal{R}}$ is the set of all allowed tilings.

We will ignore all questions of which types of rules $\mathcal{R}$ admit non-trivial tiling spaces, referring them to theoretical computer scientists and/or logicians. But it should be clear that if nontrivial, $\Omega_{\mathcal{R}}$ should be translation-invariant and closed in the big ball metric.

1.3.1.3 Cylinder sets and the metric topology

Because both are important to our further analysis we discuss these topics for both the motivating symbolic case and for the tiling situation.

In symbolic dynamics the fundamental sets are *cylinder sets*. Consider a shift space Ω and suppose w is a finite word in $\mathcal{A}^*$, where $\mathcal{A}^*$ is the set of non-empty words on $\mathcal{A}$. The *cylinder set Ω_w generated by w* is given by $C_w = \{\mathbf{x} \in \Omega \text{ such that } \mathbf{x}(U) = w\}$; it is the set of all sequences that contain the word w in the location given by U, with no other restrictions. One can check that cylinder sets are both closed and open in the metric topology. One can also check that for any $\epsilon > 0$ and any $\mathbf{x} \in \Omega$, the ball of radius ϵ around x is a cylinder set for a word around the origin in $\mathbf{x}$. Thus cylinder sets form a basis for the topology in shift spaces. When we are considering sequences in $\mathbb{Z}^d$ the cylinder sets are completely analogous.

The situation becomes somewhat more complicated for tilings of $\mathbb{R}^d$ when the translation group G is uncountable. Let P be a $\mathcal{P}$-patch, let $U \subset \mathbb{R}^d$, and let $\Omega \subset \Omega_{\mathcal{P}}$ be a tiling space. The cylinder set $\Omega_{P,U}$ is the set of all tilings in Ω that contain a copy of P translated by an element of U. That is,

$$\Omega_{P,U} = \{\mathcal{T} \in \Omega \text{ such that } P - u \text{ is a } \mathcal{T}\text{-patch for some } u \in U\}.$$

The reader can check that if ϵ is sufficiently small and $U = B_\epsilon(0)$ (the open ball of radius ϵ around the origin), the cylinder set is open. In [93, p. 13] we see how to get a countable basis for the topology by discretizing $\epsilon_n \to 0$, since there are only a countable number of patches of any size up to translation.

The basic fact of compactness is proved in several works, see for example [92]. We include a short argument here for the tiling situation.

Lemma 1. *If $\Omega \subset \Omega_{\mathcal{P}}$ is closed and of finite local complexity, then Ω is complete and compact.*

Proof. Let $\{\mathcal{T}_n\}$ be a Cauchy sequence in Ω and fix some $K \in \mathbb{Z}$. Consider $\epsilon > 0$ for which $1/\epsilon > K$. There is some M such that for $n, m \geq M$, $d(\mathcal{T}_n, \mathcal{T}_m) < \epsilon$. This means that the patches $\mathcal{T}_n \cap B_{1/\epsilon}(0)$ and $\mathcal{T}_m \cap B_{1/\epsilon}(0)$ agree up to translation by at most ϵ. Thus the patches $\mathcal{T}_n \cap B_K(0)$ and $\mathcal{T}_m \cap B_K(0)$ agree up to translation $< \epsilon$. As $\epsilon \to 0$ there is a patch P_K covering $B_K(0)$ that is the limit of the patches $\mathcal{T}_n \cap B_K(0)$. We obtain a nested sequence of patches P_K and therefore there is a tiling $\mathcal{T}$ such that $P_K \subset \mathcal{T}$ for all $K \in \mathbb{Z}$, and this tiling must be the limit of the Cauchy sequence. All of the tiles in $\mathcal{T}$ belong in a P_K for some K and so are $\mathcal{P}$-tiles, thus $\mathcal{T} \in \Omega_{\mathcal{P}}$. Since Ω is closed we know $\mathcal{T} \in \Omega$, proving sequential compactness.

Under many conditions, for instance topological transitivity, Ω is connected. Each tiling in Ω defines a path component that is a continuous embedding of $\mathbb{R}^d$. In general there are uncountably many distinct $\mathcal{P}$-tilings up to translation, and therefore there are uncountably many path components.

1.3.2 *Notions of equivalence for symbolic and tiling dynamical systems*

Suppose we have a sequence in $\mathcal{A} = \{0, 1\}$. In what way does it change if we make every 0 into an a and every 1 into a b? What about if we had a checkerboard tiling with black and white squares, and split each black square horizontally into two rectangles? In the symbolic case there are local maps called "sliding block codes" which determine factor maps and topological conjugacies between shift spaces. The tiling equivalent is local derivablility through local mappings.

1.3.2.1 Sliding block codes

We follow [74]. Let $\mathcal{A}$ and $\mathcal{A}'$ be finite alphabets and suppose Ω is a shift space in $\mathcal{A}^{\mathbb{Z}}$. Choose nonnegative integers m and n and let $B_{m,n}$ denote the set of all words of length $m + n + 1$ that appear in Ω. Let $\Phi : B_{m,n} \to \mathcal{A}'$ be any map. Then the *sliding block code* $\phi : \Omega \to (\mathcal{A}')^{\mathbb{Z}}$ is defined by this map via

$$y_i = \Phi(x_{i-m}x_{i-m+1} \cdots x_{i+n-1}x_{i+n}) = (\phi(\mathbf{x}))_i.$$

Thus we see that a sliding block code will convert every sequence $\mathbf{x}$ to a sequence $\mathbf{y}$ entry by entry, examining the block in $\mathbf{x}$ around x_i and using it to determine the value of y_i. It is not difficult to check that sliding block codes are continuous. This powerful theorem tells us that sliding block codes are the only maps on shift spaces that are both continuous and shift-commuting:

Theorem 1.3.1 (Curtis-Lyndon-Hedlund, see [74]) *Suppose Ω and Ω' are shift spaces, not necessarily on the same alphabet, and let $\theta : \Omega \to \Omega'$. Then θ is a sliding block code if and only if it is shift-commuting and continuous.*

In particular this means that topological conjugacies between shift dynamical systems are invertible sliding block codes and vice versa.

1.3.2.2 Local derivability

Local mappings are the analogue to sliding block codes for tilings of $\mathbb{R}^d$. We give a brief definition here; a full exposition appears in section 5.2 of [11].

Definition 3. A continuous surjective mapping between tiling spaces $Q : \Omega \to \Omega'$ is a *local mapping* if there is an $r > 0$ such that for any $x \in \mathbb{R}^d$ and $\mathcal{T}_1, \mathcal{T}_2 \in \Omega$, if $\mathcal{T}_1 \cap B_r(x) = \mathcal{T}_2 \cap B_r(x)$, then $Q(\mathcal{T}_1) \cap \{x\} = Q(\mathcal{T}_2) \cap \{x\}$.

That is to say, the patch in $\mathcal{T}$ containing the ball $B_r(x)$ completely determines the tile at the center of the ball in $Q(\mathcal{T})$. If such a local mapping exists we say $Q(\mathcal{T})$ is *locally derivable* from $\mathcal{T}$. If Q is invertible we say $\mathcal{T}$ and $Q(\mathcal{T})$ are *mutually locally derivable*, and we also use this terminology for their tiling spaces. It is not difficult to show that any local mapping is continuous in the big ball topology.

Lemma 2. *If Ω and Ω' are mutually locally derivable tiling spaces, then their dynamical systems are topologically conjugate.*

If there were to be a tiling analogue of the Curtis-Lyndon-Hedlund theorem, it would mean that the only continuous translation-commuting maps between tiling spaces are local mappings. That is, the above lemma would be an "if and only if". The fact that it is not was first shown in [85] and [91].

Nonlocal homeomorphisms for tilings tend to require information from far distances in $\mathcal{T}$ to settle the precise location of the origin in $Q(\mathcal{T})$. In example 13 of section 1.4.7 we describe how to make a nonlocal homeomorphism between two tiling spaces generated by a related pair of supertile methods.

1.3.3 Repetitivity and minimality

Recall that a dynamical system is called *transitive* if there is a dense orbit and *minimal* if every orbit is dense.

Definition 4. A tiling $\mathcal{T}$ is said to be *repetitive*[4] iff for every finite patch P in $\mathcal{T}$ there is an $R = R(P) > 0$ such that $\mathcal{T} \cap B_R(x)$ contains a translate of P for every $x \in \mathbb{R}^d$. It is *linearly repetitive* iff there is a $C > 0$ such that for any $\mathcal{T}$-patch P there is a translate of P in any ball of radius $C\,diam(P)$ in $\mathcal{T}$.

In other words, a tiling is repetitive if for every patch P there is some radius R such that every ball of that radius contains a copy of P. Moreover, it is linearly repetitive if R can be taken to be $C\,diam(P)$, that is, the radius depends only linearly on the size of P. In [35, 36] it is shown that a symbolic system is linearly repetitive if and only if it is "primitive and proper" S-adic (a supertile method discussed in section 1.4.5.2).

Standard arguments show the following, stated here using tiling terminology but applicable to symbolic spaces as well.

Lemma 3 (See e.g. [92, 100, 93]). *Let $\mathcal{T} \in \Omega_{\mathcal{P}}$ and let $\Omega_{\mathcal{T}}$ denote its hull. The tiling dynamical system $(\Omega_{\mathcal{T}}, \mathbb{R}^d)$ is minimal if and only if $\mathcal{T}$ is repetitive.*

Large classes of supertile methods produce sequences or tilings that are repetitive, and therefore their dynamical systems are minimal.

[4] Also known in the literature as $\mathcal{T}$ being *uniformly recurrent*, *almost periodic*, and having the *local isomorphism property.*

1.3.4 Invariant and ergodic measures

Let Ω be a shift or tiling space with topology given by the appropriate metric, and let G be the group of translations defining its dynamical system. A Borel probability measure μ on Ω is said to be *invariant* with respect to translation if $\mu(A - g) = \mu(A)$ for all Borel measurable sets A and all $g \in G$. We say μ is *ergodic* with respect to translation if whenever A is a translation-invariant set, then $\mu(A)$ equals 0 or 1. The set of invariant Borel probability measures is convex and its extremal elements are ergodic (see for example [84] or [104] for the general theory of ergodic measures).

A dynamical system is said to be *uniquely ergodic* if it possesses only one ergodic measure. Because the set of all invariant measures is convex and the ergodic measures are the extremal measures from that set, this implies that the ergodic measure is also the only invariant measure.

Let P be a $\mathcal{P}$-patch and let U be a fixed and very small ball so that if $\mathcal{T} \in \Omega_{P,U}$, then $P - g \in \mathcal{T}$ for at most one $g \in U$. Denote by $\mathbb{I}_{P,U}$ the indicator function for $\Omega_{P,U}$ and suppose μ is some ergodic measure for translation. Then from elementary measure theory along with the ergodic theorem[5] we know that for μ-almost every $\mathcal{T}_0 \in \Omega$,

$$\mu(\Omega_{P,U}) = \int_{\Omega} \mathbb{I}_{P,U}(\mathcal{T}) d\mu(\mathcal{T}) = \lim_{r\to\infty} \frac{1}{Vol(B_r(0))} \int_{B_r(0)} \mathbb{I}_{P,U}(\mathcal{T}_0 - x)dx.$$

Consider the integral on the right. For every copy of P in $\mathcal{T}_0 \cap B_r(0)$ that isn't too close to the boundary of $B_r(0)$ the indicator function will be 1 over a set of size $Vol(U)$. Any copy of P that is too close to the boundary will only yield a portion of that, but as $r \to \infty$ this effect is negligible. Letting the notation $\#(P \in \mathcal{T}_0 \cap B_r(0))$ mean the number of copies of P in $\mathcal{T}_0 \cap B_r(0)$, it is straightforward to show that the term on the right therefore becomes $\lim_{r\to\infty} \frac{\#(P\in\mathcal{T}_0\cap B_r(0))}{Vol(B_r(0))} Vol(U)$. For μ-almost every $\mathcal{T}_0$ we get the same answer and so we can say that μ defines a frequency measure on the set of $\mathcal{P}$-patches as:

$$\mu(\Omega_{P,U}) = Vol(U) freq_{\mu}(P).$$

1.4 Supertile construction techniques

When we consider a sequence or tiling space Ω under the action of translation it is not particularly interesting if the elements of Ω are periodic. Considering the dynamics on the full shift $\mathcal{A}^{\mathbb{Z}^d}$ or full tiling space $\Omega_{\mathcal{P}}$ is more interesting, since the spaces have many properties including carrying many different

---

[5] For a more indepth discussion of the meaning of the word 'frequency' and the appropriate ergodic theorem for this setting, see [50, Section 3.3].

measures, having many possible letter/tile frequencies, and having nontrivial positive topological and measure-theoretic entropies, for instance. But in this study we wish to apply the theory to spaces whose elements all have common properties that arise from given construction techniques. In particular we look at sequences and tilings constructed via *substitution* or *fusion*, which we are generically terming "supertile constructions".

1.4.1 Motivation: symbolic substitutions

Introduced as examples of symbolic dynamical systems by Gottschalk in [57], these are the fundamental (and simplest) objects on which our other supertile methods are based. Much is known about substitution sequences and the books [89, 42] are devoted to results on the subject. We will expose many of these results and see when they have generalizations or fail to generalize to higher-dimensional structures.

Given a finite alphabet $\mathcal{A}$, a *substitution* is a map $\sigma : \mathcal{A} \to \mathcal{A}^*$, where $\mathcal{A}^*$ is the set of non-empty words on $\mathcal{A}$. The substitution can be applied to words by concatenating the substitution of the letters in the word. We use the terminology *n-superword* to mean a word of the form $\sigma^n(a)$ for some $a \in \mathcal{A}$.

A sequence $\mathbf{x} \in \mathcal{A}^{\mathbb{Z}}$ is said to be *admitted* by σ if every subword of $\mathbf{x}$ is a subword of a superword of some size. We define $\Omega_\sigma \subset \mathcal{A}^{\mathbb{Z}}$ to be the set of all sequences admitted by σ. It is clear that Ω_σ is a shift-invariant subset of $\mathcal{A}^{\mathbb{Z}}$ and the reader can check that it is also closed in the metric topology defined in section 1.2.1.

Example 2 *(A constant-length substitution) Let $\mathcal{A} = \{a, b\}$ and let $\sigma(a) = abb$ and $\sigma(b) = aaa$. The first few level-n blocks of type a are:*

$$a \to abb \to abb\,aaa\,aaa \to abb\,aaa\,aaa\;abb\,abb\,abb\;abb\,abb\,abb \to \cdots,$$

where the spaces are there to help the reader see the level-n subblocks within. In this example each letter is substituted by a block of the same length (in this case 3), which is why the substitution is known as having constant length.

Clearly one could construct a substitution of constant length on any size of alphabet and any given integer length. The family in the next example contains the most famous and well-studied example of a substitution of *non-constant length.*

Example 3 *('Noble Means' substitutions and the Fibonacci substitution)*
Let $\mathcal{A} = \{a, b\}$ and choose a positive integer k. Define $\sigma(a) = a^k b$ and $\sigma(b) = a$, where by 'a^k' we mean the word composed of k consecutive 'a's. For example, let $k = 1$. In this case the first several level-n blocks of type a are:

$$a \to ab \to ab\,a \to ab\,a\;ab \to ab\,a\;ab\;ab\,a \to ab\,a\,ab\;ab\,a\;\;ab\,a\;ab \to \cdots,$$

where again we've included spaces to help the reader distinguish the level-n subblocks. Note that the superword lengths are Fibonacci numbers, and in fact all of the superwords for $k = 1$ share this property. That is why this case is called the Fibonacci substitution*.*

Of course Fibonacci numbers are closely related to the golden mean, and in fact it is the larger eigenvalue of a matrix associated with the substitution (see section 1.5.1). When k is changed we obtain other 'noble means' (silver if $k = 2$) from this matrix. All noble means substitutions have dynamical, spectral, and geometric properties in common with one another and therefore with the Fibonacci tiling, which is well-studied (see [89, 42, 11] and references therein).

The next family also contains the Fibonacci substitution, but in this class Fibonacci is the outlier, having few properties in common with the other elements.

Example 4 *Let $\mathcal{A} = \{a, b\}$, choose a positive integer k, and let $\sigma(a) = ab^k$ and $\sigma(b) = a$. The first few supertiles in the case where $k = 3$ are*

$$a \to abbb \to abbb\,a\,a\,a \to abbb\,a\,a\,a\,abbb\,abbb\,abbb \to \cdots.$$

A corresponding tiling of the line that yields well to spectral and dynamical analysis is discussed next.

1.4.2 Generalization: one-dimensional self-similar tilings

We know from section 1.2.3 that we can take any sequence in $\mathcal{A}^{\mathbb{Z}}$ and convert it into a tiling of $\mathbb{R}$ by choosing interval lengths for each element of $\mathcal{A}$. We certainly can do this for substitution sequences, and if we do it artfully we get tilings with some geometry to exploit. The process of doing it artfully leads naturally to the idea of inflation rules and self-similar tilings. It is worth doing in the context of an example first.

Example 5 *Consider tiles t_a and t_b that are intervals of length $|t_a|$ and $|t_b|$ labelled by a and b. For a positive integer k, we can use the symbolic substitution $\sigma(a) = ab^k$ and $\sigma(b) = a$ to define a tile substitution $\mathcal{S}$ so that the patch $\mathcal{S}(t_a)$ is the tile t_a followed by k t_b's and the patch $\mathcal{S}(t_b)$ is just t_a. In that case the lengths of the supertiles are $|\mathcal{S}(t_a)| = |t_a| + k|t_b|$ and $|\mathcal{S}(t_b)| = |t_a|$.*

The ideal situation, geometrically, would be if there was an "inflation factor" $\lambda > 1$ such that $|\mathcal{S}(t_a)| = \lambda|t_a|$ and $|\mathcal{S}(t_b)| = \lambda|t_b|$. A quick calculation yields that this λ would have to satisfy the equation $k = \lambda^2 - \lambda$. As expected, *in the Fibonacci case when $k = 1$, we obtain that λ is the golden mean. For larger values of k we find that λ is either 'strongly non-Pisot' or, occasionally, an integer. Later we will discuss how the algebraic properties of λ affect the dynamics of the system. The case where $k = 3$ yields $\lambda = \frac{1+\sqrt{13}}{2}$ and the rule $\mathcal{S}$ is depicted in figure 1.7.*

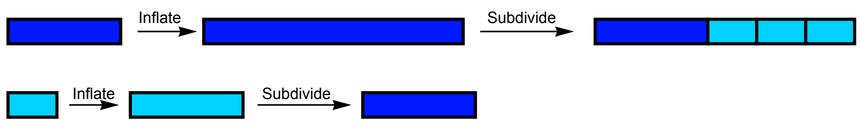

Fig. 1.7 Inflation and subdivision for the case $k = 3$.

The $k = 3$ case is fully analyzed from a diffraction standpoint in [9] and as the basis for a two-dimensional tiling with infinite local complexity in [48]. The diffraction spectrum for all values of k is given a preliminary analysis in [13] and thorough treatment in [14].

Suppose now that σ is a symbolic substitution on a general finite alphabet $\mathcal{A}$. As before it is possible to find the expansion factor and natural tile lengths (more on that later). Suppose t_e is the tile corresponding to the symbol $e \in \mathcal{A}$. Then we define $\mathcal{S}(t_e)$ to be the patch of tiles corresponding to $\sigma(e)$, supported on the interval $\lambda \operatorname{supp}(t_e)$. Often, $\mathcal{S}$ is referred to as an 'inflation rule' or an 'inflate-and-subdivide rule'.

We can extend $\mathcal{S}$ to be a map on the space of all tilings $\Omega_{\mathcal{P}}$ as follows. Let $\mathcal{T} \in \Omega_{\mathcal{P}}$ be a tiling and let $t \in \mathcal{T}$ be any tile. We define $\mathcal{S}(t)$ to be the patch given by the substitution of the prototile of t, translated so that it occupies the set $\lambda \operatorname{supp}(t)$. Then $\mathcal{S}(\mathcal{T})$ is the tiling obtained by applying $\mathcal{S}$ to all tiles in $\mathcal{T}$ simultaneously. For most $\mathcal{T} \in \Omega_{\mathcal{P}}$, $\mathcal{S}(\mathcal{T})$ is not equal to $\mathcal{T}$. However there will be fixed or periodic points for $\mathcal{S}$. A fixed point for $\mathcal{S}$ is known as a *self-similar tiling.*

1.4.3 More tricky generalization: Multidimensional constant-length substitutions in $\mathbb{Z}^d$

Symbolic substitutions of constant length generalize directly to substitutions of constant size in $\mathbb{Z}^d$. We choose a 'rectangular' shape in d dimensions, and every letter of the alphabet is substituted with a block of letters in that shape. There is no problem with iteration of the substitution because all of the blocks fit together along every dimension so concatenation of the blocks happens naturally. To generalize symbolic substitutions of non-constant length to $\mathbb{Z}^d$ we will use the fusion paradigm.

Fix *lengths* $l_1, l_2, ...l_d$, positive integers with each $l_i > 1$, and define the *location set* $\mathcal{I}^d$ to be the 'rectangle' given by

$$\mathcal{I}^d = \{\boldsymbol{\jmath} = (j_1, ...j_d) \text{ such that } j_i \in 0, 1, ..., l_i - 1 \text{ for all } i = 1, ..d\}. \quad (1.1)$$

A *block substitution* $\mathcal{S}$ is defined to be a map from $\mathcal{A} \times \mathcal{I}^d$ into $\mathcal{A}$. Then for any $e \in \mathcal{A}$ we denote by $\mathcal{S}(e)$ a block of letters; we call it a *1-superblock* or *1-supertile*.

In any particular example it is not hard to build an n-superblock through concatenation, but notation describing it precisely obscures this fact. Since the notation is not needed elsewhere in these notes we omit it. Instead we define n-superblocks inductively, using the relative positions of the letters in $\mathcal{S}(e)$ to determine the relative positions of their $\mathcal{S}^{n-1}$ blocks in $\mathcal{S}^n(e)$. Because all of the substituted blocks have the same dimensions, if two letters were adjacent it is clear that their substitutions will fit next to one another properly.

Any position $\mathbf{k} \in \mathcal{I}^d$ represents a location in a 1-superblock and we can think of $\mathcal{S}$ restricted to $\mathbf{k}$ as a map from $\mathcal{A}$ to itself. Indeed it can be useful to think of $\mathcal{S}$ as a block of maps $(p_{\mathbf{k}})_{\mathbf{k} \in \mathcal{I}^d}$. The nature of these maps is key to the dynamics of the system and they are used to compute the cocycle for the skew product representation of the system [89, 45]. An important subclass is defined as follows.

Definition 5. Let the substitution $\mathcal{S}$ as defined in this section be written as $\mathcal{S} = (p_{\mathbf{k}})_{\mathbf{k} \in \mathcal{I}^d}$. We say $\mathcal{S}$ is *bijective* if and only if each $p_{\mathbf{k}}$ is a bijection from $\mathcal{A}$ to itself.

Example 6 *A two-dimensional version of the Thue-Morse substitution uses* $\mathcal{A} = \{0, 1\}$ *with* $d = 2$ *and* $l_1 = l_2 = 2$.

$$\mathcal{S}(0) = \begin{matrix} 1 & 0 \\ 0 & 1 \end{matrix}, \qquad \mathcal{S}(1) = \begin{matrix} 0 & 1 \\ 1 & 0 \end{matrix}, \quad (1.2)$$

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

$$\mathcal{S}(*, \mathcal{I}^2) = (p_{\mathbf{k}})_{\mathbf{k} \in \mathcal{I}^2} = \begin{matrix} g_1 & g_0 \\ g_0 & g_1 \end{matrix}. \quad (1.3)$$

For example we see that $p_{(0,0)} = g_0$ *and* $p_{(0,1)} = g_1$. *Also we note that this example is bijective. The first few superblocks of type* 0 *are shown in figure* 1.8.

$$
0 \to \begin{array}{cc} 1 & 0 \\ 0 & 1 \end{array} \to \begin{array}{cc|cc} 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ \hline 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 \end{array} \to \begin{array}{cccc|cccc} 1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 & 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \\ \hline 0 & 1 & 1 & 0 & 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 & 0 & 0 & 1 \end{array}
$$

Fig. 1.8 The first three superblocks of type 0. The lines emphasize $(n-1)$-superblocks inside the n-superblocks.

1.4.4 Geometric generalization: Self-affine, self-similar, and pseudo-self-similar tilings

The geometric structure evident in the tilings in this section is governed by expanding linear maps. This makes it particularly amenable to study from a number of viewpoints, and therefore these are the most widely studied of the tilings considered in these notes. There are many examples in Chapter 6 of [11] and we try not to repeat too many of those here. Sometimes tilings created using other supertile methods can be transformed into self-similar tilings and the results that exist for them can be used. Sometimes they can't.

The earliest definition of self-similar tilings that seems to appear in print is in [102], which is a set of AMS colloquium lecture notes by William Thurston. However, the author tells us the ideas in the lectures are not all his own and refers in an informal way to a number of places where the subject was beginning to be studied.

1.4.4.1 Self-affine and self-similar tilings

We first follow the definitions laid out in [100], and use terminology from there, [46], and [11]. We also give a simpler but more restrictive version of the definition that the reader will find in [11] and lots of other places.

Definition 6. Let $\phi : \mathbb{R}^d \to \mathbb{R}^d$ be a linear transformation all of whose eigenvalues are greater than one in modulus. A tiling $\mathcal{T}$ is called *self-affine with expansion map* ϕ if

1. for each tile $t \in \mathcal{T}$, $\phi(\mathrm{supp}(t))$ is the support of a union of $\mathcal{T}$-tiles, and
2. t and t' are equivalent up to translation if and only if $\phi(\mathrm{supp}(t))$ and $\phi(\mathrm{supp}(t'))$ support equivalent patches of tiles in $\mathcal{T}$.

If ϕ is a similarity, the tiling is called *self-similar*. For self-similar tilings of $\mathbb{R}$ or $\mathbb{R}^2 \cong \mathbb{C}$ there is an *inflation constant* λ for which $\phi(z) = \lambda z$.[6]

[6] In the literature (notably [102, 100]) it is taken as given that ϕ is orientation preserving, which can be assumed by squaring any substitution that is not.

There are a few differences between our definition and the one in [100] upon which it is based. One is that ϕ is not required to be diagonalizable, and the other is that a self-affine tiling is not required to be repetitive. Proofs on the algebraic nature of the expansion constant originally required diagonalizability [66] but the condition was recently removed in [68]. A nonrepetitive tiling satisfying our definition of self-affine would be called "ϕ-subdividing" in [100]. Although we have taken finite local complexity to be a blanket assumption throughout this paper, note that our definition can be used in the infinite local complexity case as well.

What is inconvenient about this definition is the fact that one must begin by already having the self-affine tiling at hand. To be more consistent with the way we think about symbolic substitutions we can define an inflation rule on prototiles first.

Definition 7. Let $\mathcal{P}$ be a finite prototile set in $\mathbb{R}^d$ and let $\phi : \mathbb{R}^d \to \mathbb{R}^d$ be a diagonalizable linear transformation all of whose eigenvalues are greater than one in modulus. A function $\mathcal{S} : \mathcal{P} \to \mathcal{P}^*$ is called a *tiling inflation rule*[7] *with inflation map* ϕ if for every $p \in \mathcal{P}$,

$$\phi(\text{supp}(p)) = \text{supp}(\mathcal{S}(p)).$$

The linear map ϕ makes it easy to extend $\mathcal{S}$ to tiles, patches, and tilings. The substitution of a tile $t = p + x$, for $p \in \mathcal{P}$ and $x \in \mathbb{R}^d$, is the patch $\mathcal{S}(t) := \mathcal{S}(p) + \phi(x)$. The substitution of a patch or tiling is the substitution applied to each of its tiles. This means that we can consider $\mathcal{S}$ as a self-map on the full tiling space $\Omega_{\mathcal{P}}$. If a tiling $\mathcal{T}$ is invariant under $\mathcal{S}$ we call it a *self-affine tiling.* We use the term *n-supertile* to mean a patch of the form $\mathcal{S}^n(t)$.

We can use either of the methods in section 1.3.1.2 to construct a tiling space for $\mathcal{S}$. If there is a self-similar tiling $\mathcal{T}$, we can make its hull by taking the orbit closure under translation. Or, we could consider the set $\mathcal{R}$ of all n-supertiles, for all n and all prototile types, and use that as our set of admissible patches. Often the resulting spaces are identical, though not in the following example.

Example 7 *(Danzer's T2000 tiling.) Figure 1.9 gives an example from the tilings encyclopedia [1] attributed to Ludwig Danzer. It uses a total of 24 tiles, two sizes of triangles in twelve orientations. We show the inflation rule on the two sizes only; the inflation rule of the rotations are the corresponding rotations of these. The expansion map is* $\phi(x, y) = (\sqrt{3}x, \sqrt{3}y)$.

Figure 1.10 shows the second and third iteration of the larger triangle, and figure 1.11 shows a large patch of an infinite tiling. The sharp-eyed viewer will notice that the tiling appears to use only 6 rotations each of the small and

[7] These are also known as inflate-and-subdivide rules and tiling substitution rules.

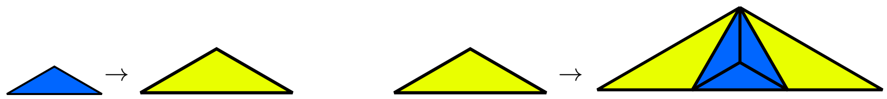

Fig. 1.9 The T2000 inflate-and-subdivide rule.

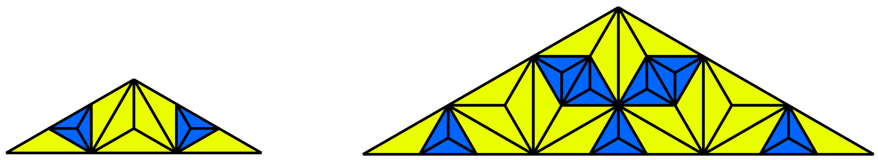

Fig. 1.10 Two- and three-supertiles.

large triangles, not 12. This is because the substitution is not 'primitive' in the sense of section 1.5.2: there is no number N such that each N-supertile contains all 24 tile types. They will always have exactly 12 when N is sufficiently large. A side effect is that the two methods for producing tiling spaces are different in that the hull of any self-similar tiling is a connected space, while the space of admissible tilings has two connected components, one a rotation of the other by 60 *degrees. In this example we could "fix" this problem by restricting our attention to an appropriate* 12*-prototile set and using two iterations of the substitution.*[8]

For the same reason there will be no self-similar tilings for the substitution as it is shown in figure 1.9, but there are period-two tilings. By replacing the substitution as shown with its square, we obtain self-similar tilings with expansion factor 3 instead of $\sqrt{3}$.

1.4.4.2 Pseudo-self-similar tilings

In this situation we still have an expanding linear map ϕ acting on our tilings, but we no longer have that $\phi(\text{supp}(p))$ is exactly the support of a patch of tiles. Instead it may only approximate the shape of $\text{supp}(p)$. Well-known examples of such substitutions are the Penrose tilings using rhombuses and/or 'kites and darts' [54, 83] (or also [58, 90, 46]) and the "binary" tilings [55] (or see [46, p. 307] or [11, p. 217]). In these examples there is a substitution rule that still 'fits' to ultimately form a tiling, but not exactly on top of the expanded tiles. Examples of this nature appear in abundance in the Tilings Encyclopedia [1] as they occur in projection tilings constructed in a similar way to Penrose tilings.

For the definition we must make precise what we mean by expanding a tiling $\mathcal{T}$ to obtain the tiling $\phi(\mathcal{T})$. For every tile t in $\mathcal{T}$, $\phi(t)$ is defined

---

[8] In general, non-primitive substitutions can have more complicated structure.

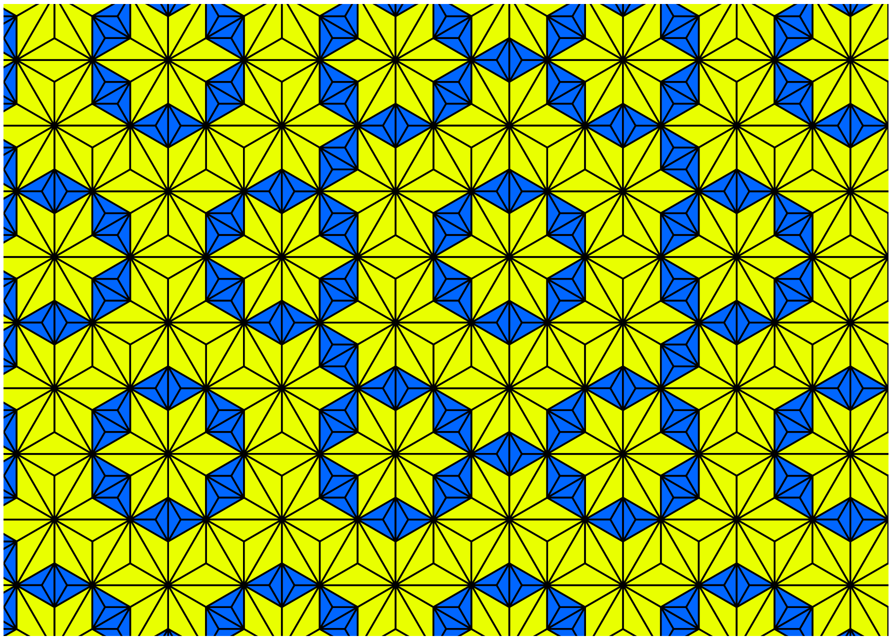

Fig. 1.11 A patch from a T2000 tiling.

to be a tile supported on $\phi(\mathrm{supp}(t))$ that carries the label of t. We define $\phi(\mathcal{T}) := \bigcup_{t\in\mathcal{T}} \phi(t)$.

Definition 8 (See [88, 99]). Let $\mathcal{P}$ be a finite prototile set in $\mathbb{R}^d$ and let $\phi : \mathbb{R}^d \to \mathbb{R}^d$ be a diagonalizable linear transformation all of whose eigenvalues are greater than one in modulus. We say a tiling $\mathcal{T} \in \Omega_{\mathcal{P}}$ is *pseudo-self-similar with expansion* ϕ if $\mathcal{T}$ is locally derivable from $\phi(\mathcal{T})$.

Example 8 *(Variation on Thurston's hexagonal example) In [102] a fractal "rep-tile" is constructed that makes a periodic tiling on the hexagonal lattice. The example is based on the observation that a regular hexagon is approximated by a patch of seven hexagons. To make a nonperiodic version we use two colors of hexagons in our inflation rule.*

The inflation map ϕ *is given by the matrix* $\begin{pmatrix} 5/2 & \sqrt{3}/2 \\ -\sqrt{3}/2 & 5/2 \end{pmatrix}$. *Pictured in figure 1.12 is what happens to the hexagonal prototiles when inflated by this map and 'subdivided' into a patch of tiles at the original scale. The location of the origin is marked with a point in each tile and patch. To see how each supertile is spatially related to the inflated prototile we have shown their overlap in figure 1.13. Figure 1.14 shows the 1-, 2-, and 3-supertiles for the blue hexagon.*

First in [88] for $\mathbb{R}^2$ and later [99] for $\mathbb{R}^d$ it is proved that every pseudo-self-similar tiling is mutually locally derivable from a self-similar tiling. In $\mathbb{R}^2$

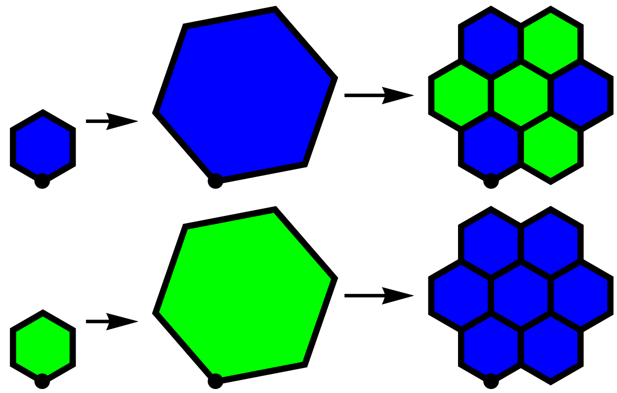

Fig. 1.12 The inflate-and-subdivide rule for a hexagonal pseudo-self-similar tiling.

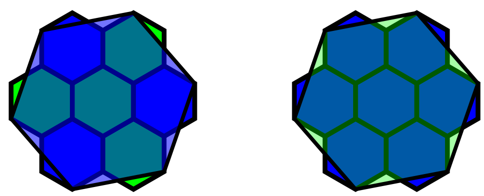

Fig. 1.13 The inflated blue tile and its patch, left; the inflated green tile and its patch, right.

the argument ends up using ideas from iterated function systems, but in $\mathbb{R}^d$ other methods are required.

1.4.5 Fusion: A general viewpoint

Symbolic substitutions and tiling inflation rules can be seen as a sort of 'cellular' model: the supertiles grow, level by level, because each symbol or tile within them has expanded to become a word or patch. Fusion takes an 'atomic' model: symbols or tiles are like atoms that come together to form 'molecules' (our 1-supertiles) that then assemble themselves into larger structures (2-supertiles) that continue to merge into larger supertiles level by level.

The sets of n-supertiles obtained by symbolic or tiling substitutions can be seen as fusion rules since it is possible (and natural) to see an n-supertile as being a union of $(n-1)$-supertiles just as easily as seeing it as the union of lots and lots of 1-supertiles. In this viewpoint the $(n-1)$-supertiles in $\mathcal{S}^n(a)$ are concatenated as prescribed by the original substitution rule on a. One could also imagine creating supertiles by applying different substitutions

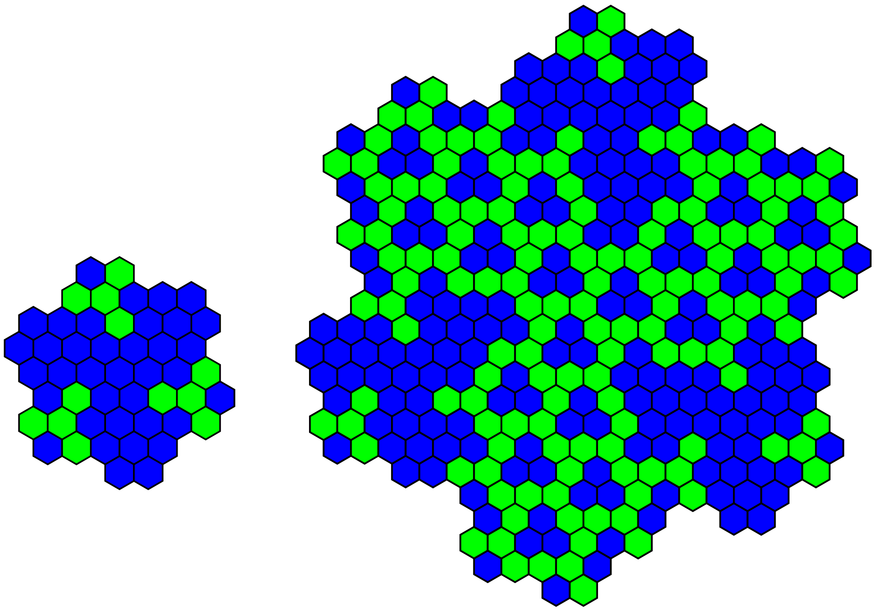

Fig. 1.14 2- and 3-supertiles for the blue prototile.

or tile inflations at each stage (if geometry permits); this is a more general situation captured in the symbolic case by S-adic systems and in the tiling case by fusion rules.

Fusion does not require an underlying linear map ϕ but instead takes a combinatorial approach. There have been other combinatorial approaches to generalizing symbolic substitutions. An early attempt to understand substitutions on graphs and in particular the dual graph for the Penrose inflation appears in [86]. A definition of substitution for the dual graphs of planar tilings is discussed in [46]. Combinatorial substitutions are defined a little bit differently in [40], and a related notion called "local rule" substitutions are defined in [39]. A definition of "topological substitutions" is given in [20]. Separately, an extremely successful program on "generalized" or "dual" substitutions began with [8]; see [7] for results tying substitutions to Markov partitions of hyperbolic toral automorphisms, complex numeration systems and β-expansions. We follow [50] for the fusion definition we give here and note that although we use the context of tilings in $\mathbb{R}^d$ the definitions are appropriate for (multidimensional) sequences as well.

Suppose that we have a finite prototile set $\mathcal{P}$ in $\mathbb{R}^d$. Given two $\mathcal{P}$-patches P_1 and P_2 and two translations x_1 and x_2, if the patches $(P_1 - x_1)$ and $(P_2 - x_2)$ intersect only on their boundaries to form a patch with connected support we call $(P_1 - x_1) \cup (P_2 - x_2)$ a *fusion* of P_1 to P_2. Of course there could be many different ways two given patches can be fused, and we could make the fusion of any finite number of patches inductively. Patch fusion is our tiling analogue to concatenation for symbols.

We form our fusion rule by defining the sets of supertiles as follows. The set $\mathcal{P}_0$ is just the prototile set $\mathcal{P}$. The set $\mathcal{P}_1$ is our set of 1-supertiles and is defined to be a finite set of finite $\mathcal{P}$-patches. We use the notation $\mathcal{P}_1 = \{P_1(1), P_1(2), ..., P_1(j_1)\}$. The set $\mathcal{P}_2$ is defined to be some finite set of finite patches that are fusions of patches from $\mathcal{P}_1$. The elements of $\mathcal{P}_2$ are our 2-supertiles and we use the notation $\mathcal{P}_2 = \{P_2(1), P_2(2), ..., P_2(j_2)\}$. One could think of the patches in $\mathcal{P}_2$ as either $\mathcal{P}$-patches, or as $\mathcal{P}_1$-patches (i.e., as patches made from 1-supertiles).

We continue inductively, forming $\mathcal{P}_3$ as a finite set of patches that are fusions of 2-supertiles, and in general letting $\mathcal{P}_n$ be a finite set of patches that are fusions of $(n-1)$-supertiles. We use the notation $\mathcal{P}_n = \{P_n(1), ..., P_n(j_n)\}$ and think of our n-supertiles both as patches of ordinary tiles and also as patches of k-supertiles for any $k < n$. We collect all of our supertiles together into an atlas of patches called our *fusion rule* $\mathcal{R}$, that is

$$\mathcal{R} = \bigcup_{n \in \mathbb{N}} \mathcal{P}_n = \{P_n(j) \mid n \in \mathbb{N} \text{ and } 1 \leq j \leq j_n\}.$$

The fusion rule can be used as a pre-language for our tiling space as defined in section 1.3.1.2. That means that tilings are fusion tilings by this rule if and only if all of their patches are seen somewhere in a patch in $\mathcal{R}$.

Remarks

1. In general we will assume that some sequence of n-supertiles grows to cover $\mathbb{R}^d$ so that there are tilings of $\mathbb{R}^d$ that are allowed by the fusion rule. That is, we take as a standing assumption that our fusion tiling spaces are nonempty.
2. The number j_n of supertiles can vary from level to level.
3. When $d = 1$, if we consider translations by elements of $\mathbb{Z}$ with all tiles having unit length, fusion tilings correspond to Bratteli-Vershik systems (except for edge sequences that have no predecessors or no successors). See [21] for more about the relationship between tilings and Bratteli-Vershik systems.
4. As stated currently, the definition of fusion rule allows for every $\mathcal{P}$-tiling to be seen as a fusion tiling. Construct the fusion rule $\mathcal{R}$ by letting the set of n-supertiles contain every possible patch of $\mathcal{P}$-tiles containing n or fewer tiles. All of $\Omega_{\mathcal{P}}$ is contained in this fusion tiling space.

Example 9 The Chacon transformation [29]. *This example was the first to show that there exist transformations with weakly but not strongly mixing dynamical systems. It was originally constructed using the "cut-and-stack" method*[9] *and can be seen as a substitution as well as a fusion tiling.*

[9] Actually, it is possible to see the process of fusion as a cutting and stacking process.

To see the Chacon transformation as a fusion rule for tilings of the line, let l_a *and* l_b *to be two positive numbers, let* a *denote a prototile with support* $[0, l_a]$*, and let* b *denote a prototile with support* $[0, l_b]$*. We let the symbols* a *and* b *also serve as the labels of the tiles if those are needed.*

For each n *there are two* n*-supertiles, which we consider being of types* a *and* b*. We define* $P_1(a) = a \cup (a + l_a) \cup (b + 2l_a) \cup (a + 2l_a + l_b)$ *and* $P_1(b) = b$*. We think of* $P_1(a)$ *as "aaba", and it has length* $3l_a + l_b$*. The length of* $P_1(b)$ *is just* l_b*, and we will let* $P_n(b) = b$ *for all* n*.*

To make $P_2(a)$ *we fuse three copies of* $P_1(a)$ *and one copy of* $P_1(b)$ *together in the same order as we did for* $P_1(a)$*, and we let* $P_2(b) = b$*. The length of the new* a *supertile is three times that of the previous* a *supertile plus the length of* b*.*

In general, we have:

$$\begin{aligned} \mathcal{P}_{n+1} &= \{P_{n+1}(a), P_{n+1}(b)\} \\ &= \{P_n(a)P_n(a)P_n(b)P_n(a)\, , P_n(b)\} \\ &= \{P_n(a)P_n(a)\, b\, P_n(a)\, , b\}. \end{aligned}$$

This is an example where not all supertiles expand. In the original formulation b *is seen as a 'spacer', and the offsets between* a*'s it provides are the cause of the weak but not strong mixing.*

Example 10 *(A direct product substitution.) Let* $\mathcal{A} = \{a, b\}$ *and define* $\sigma(a) = abb, \sigma(b) = aa$*. We take the direct product of this substitution with itself, with alphabet* $(a, a), (a, b), (b, a), (b, b)$*. We use the convention that the substitution on the first letter runs horizontally and the substitution on the second letter goes upwards. With that we obtain*

$$\mathcal{S}((a,a)) = \begin{matrix} (a,b) & (b,b) & (b,b) \\ (a,b) & (b,b) & (b,b) \\ (a,a) & (b,a) & (b,a) \end{matrix}, \qquad \mathcal{S}((a,b)) = \begin{matrix} (a,a) & (b,a) & (b,a) \\ (a,a) & (b,a) & (b,a) \end{matrix},$$

$$\mathcal{S}((b,a)) = \begin{matrix} (a,b) & (a,b) \\ (a,b) & (a,b) \\ (a,a) & (a,a) \end{matrix}, \qquad \mathcal{S}((b,b)) = \begin{matrix} (a,a) & (a,a) \\ (a,a) & (a,a) \end{matrix}.$$

It is better to visualize the substitution as a tiling, so we show the prototiles and 1*-supertiles in figure 1.15. The first row is* $\mathcal{P}_0 = \{(a, a), (a, b), (b, a), (b, b)\}$ *and the second is*

$$\mathcal{P}_1 = \{P_1\left((a,a)\right), P_1\left((a,b)\right), P_1\left((b,a)\right), P_1\left((b,b)\right)\}.$$

It is possible to iterate this as a substitution, concatenating in two dimensions in much the same way as we would do in one dimension. We choose instead to think of it as a fusion, where the $n+1$*-supertiles are constructed using the same concatenation of* n*-supertiles at every level.This concatena-*

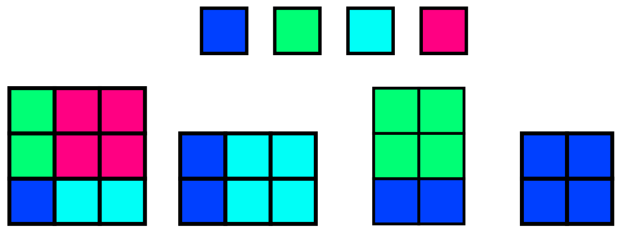

Fig. 1.15 The prototiles and 1-supertiles for the direct product substitution.

tion is diagrammed in figure 1.16 and the 2-supertiles are shown in figure 1.17.

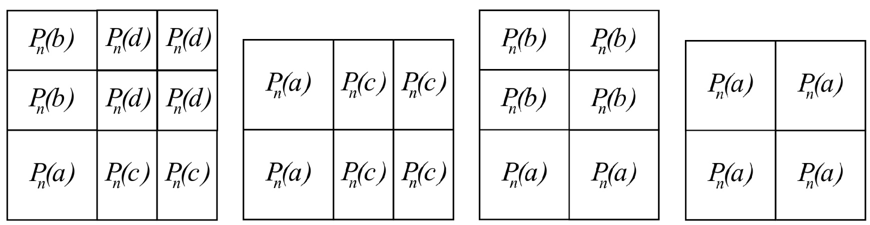

Fig. 1.16 Direct product fusion rule making $(n+1)$-supertiles from n-supertiles.

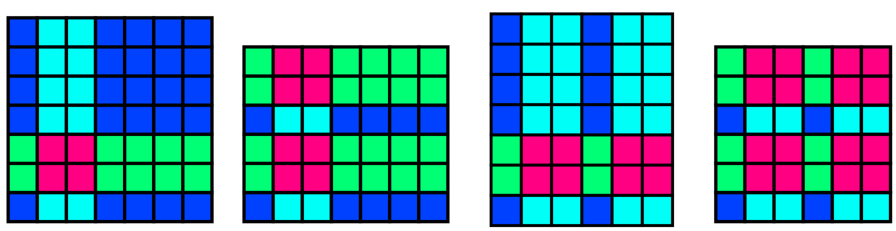

Fig. 1.17 The 2-supertiles for the direct product fusion.

The $\mathbb{Z}^2$ dynamical system associated with the direct product acts the same as the direct product of the one-dimensional systems. To get something new we rearrange the substitution on some of the letters to break the direct product structure, obtaining "Direct Product Variation" (DPV) tilings. Varying the structure is easily done but care must be taken so that the DPV substitution can be iterated to form legitimate patches and tilings.

Example 11 *(A variation on the direct product.) In this example we have chosen to rearrange only the tiles in the first supertile (compare figures 1.15* *and 1.18). The requisite care was taken to ensure that the 1-supertiles fit together to form 2-supertiles supported on topological disks, and that this nice situation will continue in perpetuity.*

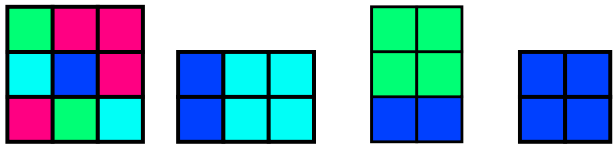

Fig. 1.18 The 1-supertiles for the DPV.

Unlike the DP case, if we try to iterate it as a substitution it becomes problematic: it is not clear how to concatenate the substituted tiles. Each supertile may be in a different relationship to its neighbors than the original tile was. In some examples it is possible to determine how to fit the supertiles together by looking at bounded patches around the original tiles. Those are the kinds of examples that have prompted definitions 'combinatorial' or 'local rules' substitutions [86, 43, 39, 40]. For this example, however, concatenation of individual 1*-supertiles inside large patches cannot be determined by local information and so the fusion paradigm is necessary. (See [46] for a discussion of the origin of these nonlocal problems and their consequences.)*

Figure 1.19 gives the general template for concatenating the n*-supertiles to make the* $(n+1)$*-supertiles, and figure 1.20 shows us the set of* 2*-supertiles. Already we can see that the direct product structure has been disrupted.*

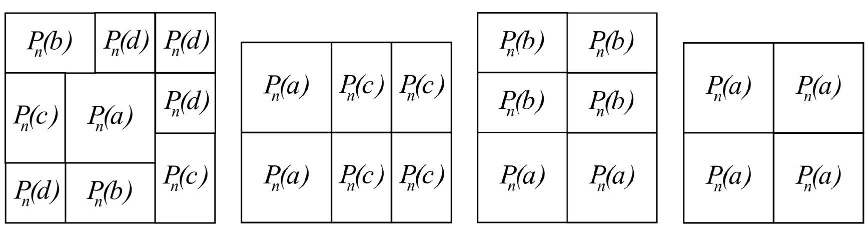

Fig. 1.19 DPV fusion rule making $(n+1)$-supertiles from n-supertiles.

For your entertainment we include a comparison of the 3*-supertiles of type* (a, a) *for the DP and DPV substitution in figure 1.21. Direct product tilings have a distinct appearance with horizontal and vertical bands clearly visible. The DPV can be compared to the introductory figure 1.6, which is a version of the DPV with 'natural' tile sizes.*

The topology, and in particular the cohomology, of a DPV based on a strongly non-Pisot substitution in product with the substitution $1 \to 11$

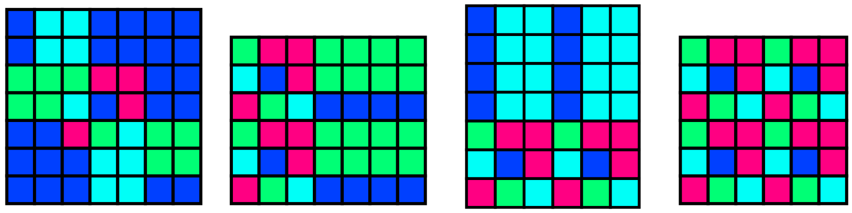

Fig. 1.20 The 2-supertiles for the DPV fusion.

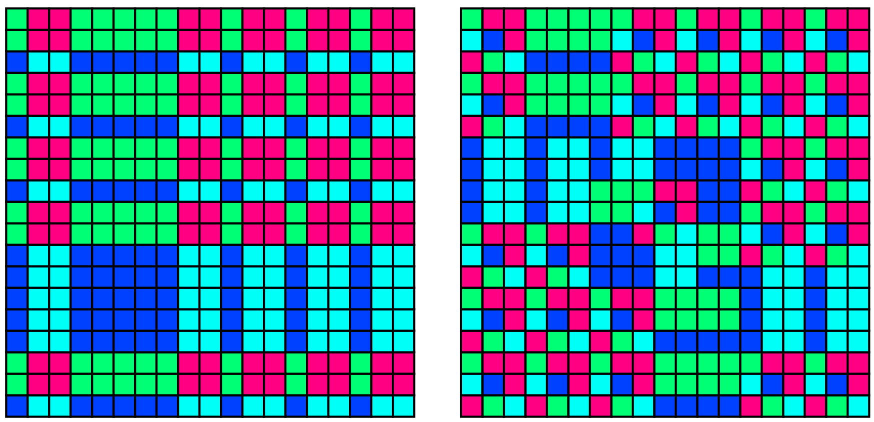

Fig. 1.21 The 3-supertiles of type (a, a) for the DP (left) and DPV (right).

(which gives a “solenoid” system) is analyzed in [49]. In that example the DPV uses natural tile sizes and has infinite local complexity.

1.4.5.1 Special classes of fusion tilings

A fusion rule $\mathcal{R}$ is called *prototile-regular* if each $\mathcal{P}_n$ has the same number of elements. A prototile-regular fusion rule is called *transition-regular* if the number of each tile type in each supertile type doesn’t change from level to level. The DP and DPV examples shown were both prototile- and transition-regular, and they are also ‘algorithmic’ in the sense of the next example.

Example 12 *(A “uniform shape substitution”.) This is an example of the type of substitution found in [39]. We call it* algorithmic *because a simple computer algorithm can be written to describe the formation of the n-supertiles. The algorithm is iterative, accepting n-supertiles and fundamental level-n translations and returning $(n+1)$ versions of these. Interesting for these* *examples is that there may be more than one possible input (prototile set and fundamental translations) that leads to a tiling of* $\mathbb{R}^2$*.*

Because it isn't obvious how to make a simplified figure like we did for DPVs that describe the combinatorics of how to put the n*-supertiles together to make the* $(n+1)$*-supertiles, we give the algorithm instead. Let* A_n *and* B_n *denote* n*-supertiles and let* $\mathbf{k}_n$ *and* $\mathbf{l}_n$ *represent fundamental translation vectors at the* n*th level, and let* $L = \begin{pmatrix} 2 & 1 \\ -1 & 1 \end{pmatrix}$*. Then*

$$A_{n+1} = A_n \cup (B_n + \mathbf{k}_n) \cup (B_n + \mathbf{l}_n) \qquad \text{and} \qquad B_{n+1} = B_n \cup (A_n + \mathbf{k}_n) \cup (A_n + \mathbf{l}_n).$$

The level-$(n+1)$ *translations are* $\mathbf{k}_{n+1} = L\mathbf{k}_n$ *and* $\mathbf{l}_{n+1} = L\mathbf{l}_n$*.*

To run the algorithm, put in a prototile set and some initial vectors and see what happens when the algorithm is iterated. There are at least three distinctly-shaped prototile sets and corresponding initial vectors that lead to tilings of $\mathbb{R}^2$*; we include only one here.*

One possible prototile set is colored unit squares with their lower left endpoints at the origin. With this input set it is necessary to set $\mathbf{k}_0 = (1,0)$ *and* $\mathbf{l}_0 = (0,1)$*. Figure 1.22 shows the first six supertiles of the blue (A) prototile. It is essential to note that each successive supertile is shown at a scale smaller than it actually is: all the tiles should be the same size as the first one. The rescaling is just there to display the supertiles together. Also note the dot in each image: it represents the location of the origin and allows us to see the rotational aspect of this fusion rule.*

It is the 'shape' of the substitution that matters: it is 'uniform' in the sense that the shapes of the n*-supertiles of either type are the same and it is only the coloring that differs. In our example we have chosen a bijective coloring, using the word 'bijective' in the same way as we used it for substitutions in* $\mathbb{Z}^d$*.*

Another prototile set that works is a pair of colored hexagons with vertex set

$$\{(1,0),(0,1),(-1,1),(-1,0),(0,-1),(1,-1)\},$$

in which case it is necessary to set $\mathbf{k}_0 = (2,-1)$ *and* $\mathbf{l}_0 = (1,1)$*.*

These tilings turn out to be pseudo-self-similar with expansion map L*. Moreover, figure 1.22 provides convincing evidence of the existence of a fractal-shaped tile that could be used as the 'uniform shape' and which makes a self-similar tiling.*

1.4.5.2 *S*-adic systems

S-adic systems are both generalizations of symbolic substitutions and specializations of fusion rules. There are substitutions, but they can change from level to level and thus have to be applied in reverse order, effectively turning them into a fusion rule. The term "S-adic" and basic definitions are proposed

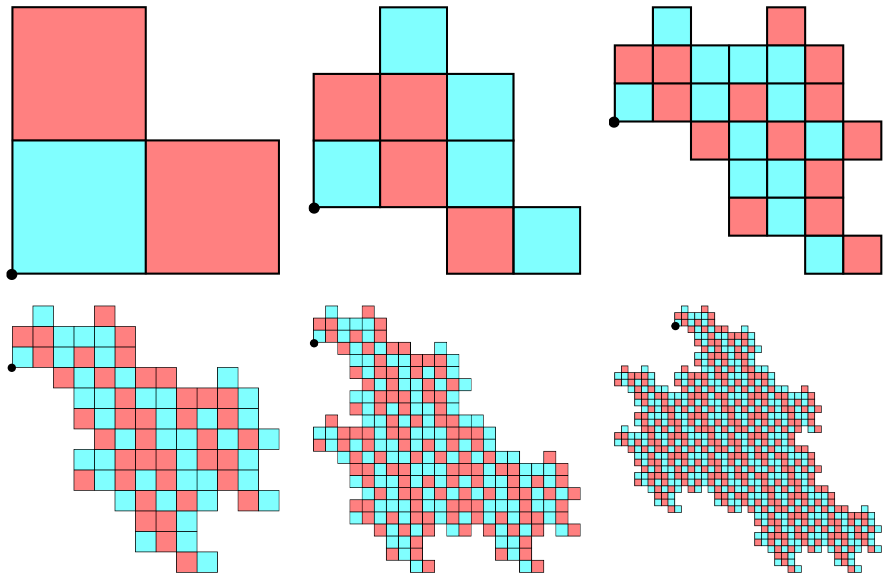

Fig. 1.22 Level-n tiles of type A for $n = 1, ...6$, beginning with square tiles. Each square tile is the same size despite the image rescaling, which is there solely for display purposes.

in [38], as part of a larger study of symbolic systems of low complexity. There are many reasons why this generalization is useful, as it intersects with continued fractions and interval exchange transformations, and has been very interesting in the study of combinatorics on words. The topic and its connections to numerous areas is surveyed in [103]. Recently, a few generalizations of S-adic systems to higher dimensions have been made that do not use the fusion paradigm [51, 53].

We follow the notation of [24] first and then explain how this well-studied family of systems fits into the fusion paradigm and can be seen as a supertile construction method. Let $\mathcal{A}_0, \mathcal{A}_1, \mathcal{A}_2, ...$ be a family of finite alphabets, and, for each n, let $\sigma_n : \mathcal{A}_{n+1} \to \mathcal{A}_n^*$ be a map taking an element from $\mathcal{A}_{n+1}$ to a nonempty word in the alphabet $\mathcal{A}_n$.[10] Let $\{a_n\}_{n=0}^{\infty}$ represent a sequence for which $a_n \in \mathcal{A}_n$ for all $n \in \mathbb{N}$. An infinite word $\mathbf{x} \in \mathcal{A}_0^{\mathbb{N}}$ admits the *S-adic expansion* $\{(\sigma_n, \mathcal{A}_n)\}_{n=0}^{\infty}$ if

$$\mathbf{x} = \lim_{n\to\infty} \sigma_0\sigma_1 \cdots \sigma_{n-1}(a_n).$$

[10] There is a separate lexicon in which what are known to some as "substitutions" are known to others as "non-erasing morphisms" and the set $\mathcal{A}^*$ is called the "free monoid" instead of the set of all finite words from $\mathcal{A}$ [6, 76].

We have the usual two options to make a sequence space, either as the hull of $\mathbf{x}$ or as the set of all sequences admitted by the set of n-superwords.

(The theoretical computer science community has established terminology for the set of all allowed words: a 'language'. In this terminology the directive sequence $\{\sigma_n\}$ has a language associated with it given by

$$L = \bigcap_{n \in \mathbb{N}} \overline{\sigma_0 \sigma_1 \cdots \sigma_{n-1}(\mathcal{A}_n^*)},$$

where the notation $\overline{M}$ denotes the smallest language containing the set M. The S-adic system given by $\{\sigma_n\}$ can then be studied through the shift space admitted by this language.)

Let us see how this fits into the fusion paradigm. The prototile set is $\mathcal{A}_0$, which could be seen as labelled unit intervals if we prefer a tiling to a sequence. The 1-supertiles are constructed using the map $\sigma_0 : \mathcal{A}_1 \to \mathcal{A}_0$. For each $a \in \mathcal{A}_1$, $\sigma_0(a)$ is a word in $\mathcal{A}_0$ which by abuse of notation we might think of as a patch instead. In either case we call it a 1 supertile. The 1-supertiles are given by

$$\mathcal{P}_1 = \{\sigma_0(a) \text{ such that } a \in \mathcal{A}_1\}$$

The 2-supertiles are given by $\sigma_0(\sigma_1(a)))$, where a is now an element of $\mathcal{A}_2$, and we need to see why those are fusions of 1-supertiles. Notice that $\sigma_1(a)$ is a word in $\mathcal{A}_1^*$, and so we can apply σ_0 to each of its letters. Thus one can see $\sigma_0(\sigma_1(a))$ as the fusion of blocks of the form $\sigma_0(a')$ in the order prescribed by $\sigma_1(a)$. Thus the set of 2-supertiles is

$$\mathcal{P}_2 = \{\sigma_0 \sigma_1(a) \text{ such that } a \in \mathcal{A}_2\}$$

Now the 3-supertiles are given by $\sigma_0 \sigma_1 \sigma_2(a)$, where now $a \in \mathcal{A}_3$. To see these are fusions of 2-supertiles, suppose $\sigma_2(a) = b_1 b_2 \cdots b_k$, which is in $\mathcal{A}_2^*$. Then $\sigma_0 \sigma_1(\sigma_2(a)) = \sigma_0 \sigma_1(b_1) \sigma_0 \sigma_1(b_2) \cdots \sigma_0 \sigma_1(b_k)$, which is a fusion of 2-supertiles. Clearly, then, the n-supertiles take the form

$$\mathcal{P}_n = \{\sigma_0 \sigma_1 \cdots \sigma_{n-1}(a) \text{ such that } a \in \mathcal{A}_n\},$$

and can be seen as fusions of $(n-1)$-supertiles as desired. So in this work we consider S-adic constructions to be supertile constructions as well.

1.4.6 Tiling spaces from supertile methods

As we discussed in section 1.3.1.2, tiling spaces are often given as either the hull of a specific tiling or as the tilings admitted by a specific language. It is reasonable to try both approaches for supertile methods.

In the case of general symbolic substitutions in $\mathbb{Z}$, constant-length substitutions in $\mathbb{Z}^d$, and self-affine or -similar tilings of $\mathbb{R}^d$, we have noted that we can consider the substitution as an action on the sequence or tiling space. Thus if $\mathcal{T}$ is a fixed point of the substitution we can study the supertile rule by studying the hull of $\mathcal{T}$. This doesn't work for arbitrary fusion rules, since they don't define actions on the full tiling space.

For fusion rules (and therefore all supertile methods) one can consider the set of all n-supertiles to be the set $\mathcal{R}$. Thinking of $\mathcal{R}$ as a pre-language (which may be an abuse of terminology), we see that if there are infinite tilings in $\Omega_{\mathcal{P}}$ that are admitted by $\mathcal{R}$, then the space $\Omega_{\mathcal{R}}$ is nontrivial. We do not attempt to determine precise conditions that enforce nontriviality, but that is easy to check in examples. In [50] a blanket assumption is that the boundaries of the n-supertiles become arbitrarily small compared to their interiors as $n \to \infty$. (Such a sequence of sets is called a "van Hove sequence", and we call fusions of this sort van Hove.) We clearly need some sort of growth condition for supertiles in order for the tiling space to be nontrivial.

Often the Ωs you get by either method are identical, but there are exceptions. One notable exception is if the invariant $\mathcal{T}$ contains a 'defective' patch that is not allowed by the substitution but is stable under the substitution rule. The Danzer T2000 example was another exception arising from not having primitivity.

1.4.7 Recognizability or the unique composition property

Suppose you are given an element of a tiling (or sequence) space given by some supertiling method. All you see are the tiles in the tiling. Can you determine uniquely how the tiles group into supertiles? If so, the tiling (or sequence) is recognizable.

To make this definition it is convenient to introduce the notion of "supertiling spaces". Let Ω be a (nonempty) tiling space defined for some supertile method. Fix an n and choose some $\mathcal{T} \in \Omega$. For any $x \in \mathbb{R}^d$ the patch $\mathcal{T} \cap \{x\}$ must be contained in some n-supertile, either from the generating tiling or from $\mathcal{R}$. The n-supertile might not be unique, but there are only finitely many possibilities. A diagonalization argument can be made to extrapolate that all tiles in $\mathcal{T}$ itself can be composed into n-supertiles that overlap only on their boundary. A tiling $\mathcal{T}_n$ obtained by this composition, i.e. where the prototile set is considered at $\mathcal{P}_n$ rather than $\mathcal{P}$, is called an *n-supertiling* of $\mathcal{T}$. The space of all n-supertilings of Ω is denoted Ω_n and is a translation-invariant subspace of the tiling space $\Omega_{\mathcal{P}_n}$.

Since each n-supertile is constructed from $(n-1)$-supertiles by definition, there is a unique *decomposition map* f_n taking Ω_n to Ω_{n-1}. It is possible that the tiling $\mathcal{T}$ can be composed in more than one way into a tiling in Ω_n. In this case the supertile rule is “not recognizable” or does not have the “unique composition property”.

Definition 9. (See [50]) A supertile rule is said to be *recognizable* if the decomposition map from Ω_n to Ω_{n-1} is invertible for all n.

This definition looks at recognizability as a sort of global property determined by the connection between supertiling spaces. It can, however, be convenient to think of it locally as converting patches of n-supertiles into $(n+1)$-supertiles. If a supertile rule is recognizable then every tiling in Ω can be unambiguously expressed as a tiling with n-supertiles for every n. It is not difficult to show that the decomposition maps are uniformly continuous, and if they are invertible the inverse is also uniformly continuous. Thus there exists a family of *recognizability radii* r_n $(n = 1, 2, \ldots)$, such that, whenever two tilings in Ω have the same patch of radius r_n around a point $\mathbf{v} \in \mathbb{R}^d$, then the n-supertiles intersecting $\mathbf{v}$ in those two tilings are identical.

The terminology original to the symbolic substitutions case is ‘recognizability’, and it is shown in [81] that recognizability and nonperiodicity are equivalent in that case. Solomyak [100] gave the concept the name ‘unique composition property’ for self-affine tilings. He proved in [98] that unique composition and nonperiodicity are equivalent. In [50] the notion is defined for fusions, where it is shown by example that there are no general results connecting nonperiodicity to recognizability.

Recognizability turns out to be essential for many arguments and is almost always assumed. It is central in the following construction.

Example 13 (How to make a nonlocal homeomorphism) *We use two tiling spaces associated with the Fibonacci substitution of Example 3. Let Ω be the space with labelled unit-length tiles and Ω' to be the space with natural tile lengths, normalized so that the n-supertiles in both spaces asymptotically converge in length. These spaces are easily shown to be recognizable.*

We can make an invertible local map Q by requiring that $\mathcal{T}$ and $Q(\mathcal{T})$ have the same underlying sequence of a's and b's and then determining the precise location of 0 in $Q(\mathcal{T})$. This location is determined from the location of 0 in $\mathcal{T}$ using the supertile structure: one considers the sequence of n-supertiles in $\mathcal{T}$ containing 0, and translates $Q(\mathcal{T})$ a little bit for each n so that, say, the left endpoints of the n-supertiles line up. Since the lengths converge asymptotically the adjustments at each stage will go to 0 and the precise location of $Q(\mathcal{T})$ can be determined.

1.5 Ergodic-theoretic and dynamical analysis of supertile methods

1.5.1 Transition (a.k.a. incidence, substitution, abelianization, or subdivision) matrices

We can obtain basic geometric and statistical information by associating a matrix or matrices to a supertile rule. Transition matrices keep track of how many n-supertiles of each type there are in each $(n+1)$-supertile. They go by many names in the literature but we will use 'transition' as our terminology.

These matrices are fundamental to their supertile rules, among other things helping to compute frequencies and ergodic measures. Since we can count how many times p_i appears in $\mathcal{S}^n(p_j)$, and we can count the total number of tiles in $\mathcal{S}^n(p_j)$, we can estimate the relative frequency of p_i. The Perron-Frobenius theory of matrices allows us to draw conclusions about the frequency statistics in our hulls accordingly. In particular we will see how to use transition matrices to construct ergodic measures.

The matrices for self-similar tilings of $\mathbb{R}^d$, substitutions in $\mathbb{Z}$, constant-length substitutions of $\mathbb{Z}^d$, and stationary (or even transition-regular) fusion rules in $\mathbb{R}^d$ or $\mathbb{Z}^d$ are all obtained the same way, so we use tiling terminology to refer to all cases. We assume the prototile set has been given some arbitrary order $\mathcal{P} = \{p_1, ...p_{|\mathcal{P}|}\}$, which we shall keep fixed. Then the transition matrix for $\mathcal{S}$ is the $|\mathcal{P}| \times |\mathcal{P}|$ matrix M whose (i, j) entry M_{ij} is the number of tiles of type p_i in $\mathcal{S}(p_j)$. It is not hard to check that the tile population information for $\mathcal{S}^n$ is given by M^n.

Fusion rules have a somewhat more complicated transition matrix situation. This is due to two facts: first, that the fusion rules used to construct the n-supertiles may be completely unrelated to the rules used to construct the m-supertiles when $m \neq n$, and second, that the number of supertiles can vary from level to level. Thus we need an infinite family of (possibly non-square) transition matrices in order to give us the population information we seek.

Recall that $\mathcal{P}_n$ is the set of n-supertiles. We define the *transition matrix* $M_{n-1,n}$ to be the $|\mathcal{P}_{n-1}| \times |\mathcal{P}_n|$ matrix whose (k, l) entry is the number of supertiles equivalent to $P_{n-1}(k)$ (i.e. the number of $(n-1)$ supertiles of type k) in the supertile $P_n(l)$. Notice that the matrix product $M_{n,N} = M_{n,n+1}M_{n+1,n+2} \cdots M_{N-1,N}$ is well-defined when $N > n$. The entries of $M_{n,N}$ reveal the number of n-supertiles of every type of N-supertile. If there is more than one fusion of $\mathcal{P}_{n-1}$-supertiles that can make $P_n(l)$, we fix a preferred one to be used in all computations.

1.5.2 Primitivity

Generally speaking, a supertile method is primitive if one finds every type of n-supertile in every type of N-supertile, provided N is sufficiently large. This assumption is useful in obtaining minimality, repetitivity, and unique ergodicity results and ensures that our hulls have a certain level of homogeneity.

For this definition, suppose M or $M_{n,n+N}$ are the transition matrices for the supertile rules.

Definition 10. A symbolic or tiling substitution rule is defined to be *primitive* if and only if there is an N such that all of the entries of M^N are strictly positive. A fusion rule is defined to be *primitive* if and only if for every $n \in \mathbb{N}$ there exists an N such that the entries of $M_{n,n+N}$ are strictly positive.

In this latter situation it is possible that N varies depending on n.

1.5.2.1 General result: Primitivity implies minimality

Recall that a topological dynamical system (Ω, G) is said to be *minimal* if and only if Ω is the orbit closure of any of its elements.

Proposition 1.5.1 [11] *Let Ω be the space of tilings allowed by a supertile construction and let G be its group of translations. If the supertile construction is primitive, then (Ω, G) is minimal.*

Note that this proposition is not an if and only if: the Chacon substitution is an example of a supertile construction that is minimal even though it is not primitive. Although the most study has been done on primitive systems, progress has been made in the non-primitive case as well [25, 26, 78]. We should note, however, that primitivity and minimality are 'morally' the same in the sense that minimal nonprimitive systems in one dimension can be transformed into primitive ones using a return word procedure [78]. A way to do it for a Chacon DPV in two dimensions is shown in [46].

1.5.2.2 Result for substitution systems: Primitivity implies unique ergodicity

Because it is easy to write down the transition matrix for a substitution rule, it can often easily be determined that a substitution dynamical system is uniquely ergodic. The general theorem that makes this possible in all cases is the Perron-Frobenius theorem. The situation for fusions is more subtle and will be discussed later.

[11] See [64, 89] for symbolic substitutions, [87] for self-affine tilings, and [50] for fusions.

Part of the Perron-Frobenius theorem requires matrices that are *irreducible* in the sense that for each index (i, j) there is an $n \in \mathbb{N}$ such that $M^n_{ij} > 0$. This condition is weaker than primitivity because the entries are not required to be simultaneously positive for any n. Clearly, a primitive matrix is irreducible. We cite the portions of the Perron-Frobenius theorem that are relevant to our study as a combination of statements from [74, p. 109], [67, p. 16], and [100, p. 704].

Theorem 1.5.2 (Perron-Frobenius Theorem) *Let M be an irreducible matrix. Then M has positive left and right eigenvectors $\mathbf{l}$ and $\mathbf{r}$ with corresponding eigenvalue $\theta > 0$ that is both geometrically and algebraically simple. If θ' is another eigenvalue for M then $|\theta'| \leq \theta$. Any positive left or right eigenvector for M is a multiple of $\mathbf{l}$ or $\mathbf{r}$.*

Moreover, if M is primitive and $\mathbf{l}$ and $\mathbf{r}$ are normalized so that $\mathbf{l} \cdot \mathbf{r} = 1$, it is true that

$$\lim_{n\to\infty} \frac{M^n}{\theta^n} = \mathbf{rl}.$$

The eigenvector $\mathbf{l}$ and $\mathbf{r}$ are ordinarily called the Perron eigenvectors and θ is always called the *Perron eigenvalue* for M.

In order to find the result that primitivity implies unique ergodicity for symbolic substitutions, a good reference is [89, Ch. V.4]. There we find this stronger result:

Theorem 1.5.3 ([89], Theorem V.13) *If a symbolic substitution system is minimal, then it is uniquely ergodic.*

It is stronger because primitivity is not the only way for a symbolic substitution system to be minimal. For example, the Chacon system is minimal and therefore uniquely ergodic despite it not being primitive.

The situation for self-affine tilings is as follows.

Corollary 1.5.4 (Folklore; Corollary 2.4 of [100]) *Suppose $\mathcal{T}$ is a self-affine tiling with expansion map ϕ for which the transition matrix M is primitive. Then the Perron eigenvalue of M is $|det\phi|$. Writing the prototile set as $\mathcal{P} = \{p_1, p_2, ..., p_m\}$, the left eigenvector can be obtained by $\mathbf{l} = (Vol(p_j))_{j=1}^m$. Moreover,*

$$\lim_{n\to\infty} |det\phi|^{-n} M^n_{ij} = r_i Vol(p_j).$$

The last equation proves particularly useful in computing frequencies and ergodic measures as we will show in the case study, next. Moreover, one can show that the n-supertile frequencies are given by $\frac{1}{|det(\phi)|^n}\mathbf{r}$ and we can get the frequencies of everything else from that information.

The adaptation of the previous result to multidimensional constant-length symbolic substitutions is carried out in [45].

1.5.2.3 Case study: ergodic measures for constant-length $\mathbb{Z}^d$ substitutions

By results in [45], we can use Corollary 1.5.4 to make an instructive example. We consider a tiling model of sequences in Ω where each tile is a unit 'cube' in $\mathbb{R}^d$ labeled by the appropriate element of $\mathcal{A}$.

If $\mathcal{S}$ is a primitive, nonperiodic substitution with size $l_1 \cdot l_2 \cdots l_d = K$ and ϕ is its natural expanding map, corollary 1.5.4 implies that the largest eigenvalue of M must be equal to $|\det \phi| = K$. The left Perron eigenvector $\mathbf{l}$ must the the tile volumes, which are all 1. This implies that $\sum_{i=1}^{|\mathcal{A}|} r_i = 1$ and that

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

We know that the dynamical system $(\Omega, \mathbb{Z}^d)$ is uniquely ergodic, and by our discussion of the connection between ergodic measures and frequencies in section 1.3.4 we know that frequencies exist and must be equal for almost every element of Ω. One can show that it suffices to compute frequencies in larger and larger supertiles, rather than in arbitrary balls of expanding radius. In fact, by primitivity it doesn't matter which type of supertiles we look at, so we will just look at $\mathcal{S}^n(a_1)$ as $n \to \infty$. We will use the notation $N_{a_i}(B)$ to denote the number of occurrences of the letter a_i in a block B. Then

$$freq(a_i) = \lim_{n\to\infty} \frac{N_{a_i}(\mathcal{S}^n(a_1))}{K^n}, \tag{1.1}$$

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^n_{i1}. Thus we have that $freq(a_i) = \lim_{n\to\infty} K^{-n} M^n_{i1} = r_i$, and so computation of $freq(a_i)$ reduces to computation of the right eigenvector for M.

Sometimes the computation comes out particularly nice. For instance, we have the following proposition:

Proposition 1.5.5 *Let $\mathcal{S}$ be a primitive and nonperiodic substitution of constant length $l_1 \cdot l_2 \cdots l_d = K$ in $\mathbb{Z}^d$. Then M has the property that $\sum_{j=1}^{|\mathcal{A}|} M_{ij} = K$ for all $i \in 1, 2, ...|\mathcal{A}|$ 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 a right eigenvector for M is given by $\mathbf{r} = (1/|\mathcal{A}|, ..., 1/|\mathcal{A}|)$. Since $\mathbf{l} \cdot \mathbf{r} = 1$ it must be the (unique) right Perron eigenvector for M. Since $freq(a_i) = r_i$ the result follows.

Conversely, the vector $\mathbf{r}$ defined as above again is a right eigenvector and we have that

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

and the result follows.

Bijective substitutions (defined in section 1.4.3) automatically satisfy the former condition, and so do the the Rudin-Shapiro-like substitutions seen in [44].

Corollary 1.5.6 *If $\mathcal{S}$ is a primitive, nonperiodic, bijective substitution of constant length in $\mathbb{Z}^d$, then the frequency of any letter $a_i \in \mathcal{A}$ is $1/|\mathcal{A}|$.*

Proof (Sketch). The row sum for row i is the number of times we see a_i in the substitution of any tile. Because the substitution is bijective, for any given location in the substitution we know that a_i appears exactly once. That means that the number of times a_i appears is the number of spots in the substitution, which is K.

1.5.3 Ergodic measures for fusions

Let us suppose that $\mathcal{R}$ is a primitive, recognizable, van Hove fusion rule in $\mathbb{R}^d$ that admits a nontrivial tiling space Ω. What are the possibilities for translation-invariant measures? We suppose for convenience[12] that the group action for dynamics is $G = \mathbb{R}^d$ and follow [50], section 3.4.

We cannot use the Perron-Frobenius theorem in this situation because our transition matrices change from level to level, so we adapt it to work here. Because of the relationship between S-adic systems, Bratteli diagrams, and fusion tilings, our analysis is closely related to that in [41, 27] and others. (Our work takes the analysis into the continuous dynamics situation.)

By recognizability we know that every tiling $\mathcal{T} \in \Omega$ has a unique n-supertiling $\mathcal{T}_n \in \Omega_n$. Consider a particular n-supertile $P_n(j) \in \mathcal{P}_n$. We denote its frequency in $\mathcal{T}$ as an n-supertile as $\tilde{f}_{P_n(j)}$, if it exists. We know that $P_n(j)$ as a patch of ordinary tiles may have a larger frequency $\bar{f}_{P_n(j)}$ in $\mathcal{T}$. In fact, recognizability gives us a finite list of patches $S_1, S_2, ...S_q$ that appear if and only if $P_n(j)$ appears as a supertile. That means that we can compute $\bar{f}_{P_n(j)} = \sum_{i=1}^q \bar{f}_{S_i}$ if the latter frequencies exist in $\mathcal{T}$.

In the symbolic or tiling substitution case we found a right eigenvector $\mathbf{r}$ that represented the prototile frequencies and satisfied $\mathbf{r} \cdot \mathbf{l} = 1$, where $\mathbf{l}$ is the vector of tile volumes. We might say that $\mathbf{r}$ is ‘volume-normalized’, and the useful thing about that is that it makes the ergodic measure a probability measure. We need to extend this concept to fusion tilings.

In the substitution case, the n-supertile frequencies are all given by the vector $\frac{1}{|det\phi|^n}\mathbf{r}$. For fusion rules the supertile frequencies are not as simple.

[12] How to adapt the analysis appears in [50], section 3.7.

We let a nonnegative vector $\rho_n = (\rho_n(1), ..., \rho_n(j_n)) \in \mathbb{R}^{j_n}$ represent the n-supertile frequencies.

Definition 11. Let ρ be a sequence of vectors $\{\rho_n\}$ described above. We say that ρ is *volume-normalized* if for all n we have $\sum_{i=1}^{j_n} \rho_n(i) Vol(P_n(i)) = 1$. We say that it has *transition consistency* if $\rho_n = M_{n,N}\rho_N$ whenever $n < N$. A transition-consistent sequence ρ that is normalized by volume is called a sequence of *well-defined supertile frequencies*.

As before, volume-normalization is there to ensure that the measure is a probability measure. The transition-consistency requirement ensures that the measure is additive: it is necessary that the frequency of N-supertiles be related to the frequencies of the n-supertiles they are composed of. This property was automatically satisfied before because $\mathbf{r}$ was an eigenvector. For fusion rules, it turns out that the invariant measures are completely determined by sequences of well-defined supertile frequencies:

Theorem 1.5.7 *[50] Let $\mathcal{R}$ be a recognizable, primitive, van Hove fusion rule with tiling dynamical system $(\Omega, \mathbb{R}^d)$. There is a one-to-one correspondence between the set of all invariant Borel probability measures on $(\Omega, \mathbb{R}^d)$ and the set of all sequences of well-defined supertile frequencies with the correspondence that, for all patches P,*

$$freq_\mu(P) = \lim_{n\to\infty} \sum_{i=1}^{j_n} \#\left(P \;\, in \;\, P_n(i)\right) \rho_n(i). \tag{1.2}$$

Thus one could, given a sequence of well-defined supertile frequencies, construct a translation-invariant ergodic measure μ as follows. Given any patch P, get the frequency $freq_\mu(P)$ as in equation (1.2). Then, the measure of a cylinder set $\Omega_{P,U}$ will be $freq_\mu(P)Vol(U)$, provided U is not too large. Since the cylinder sets form a basis for the topology, we can now measure any Borel measurable set.

1.5.4 General result: Substitution systems are not strongly mixing

Recall that a measure-preserving system is *strongly mixing* if for all measurable sets A, B and for any sequence of vectors $\mathbf{v}_n$ whose lengths increase without bound it is true that $\lim_{n\to\infty} \mu\left((A \cap (B - \mathbf{v}_n)\right) = \mu(A)\mu(B)$. There is a standard argument proving that a substitution system isn't strongly mixing, appearing for substitution sequences in [34] and for self-similar tilings in [100]. Here is a general result.

Theorem 1.5.8 ([50]) *The dynamical system of a strongly primitive van Hove fusion rule with a constant number of supertiles at each level and bounded transition matrices, and with group $G = \mathbb{Z}^d$ or $\mathbb{R}^d$, cannot be strongly mixing.*

For tiling or symbolic substitutions neither the transition matrix nor the number of supertiles changes from level to level, so these are never strongly mixing. Indeed, we see that constructing a strongly mixing supertile method requires a certain degree of unboundedness, either in the number of supertiles at each level or in the entries of the transition matrices between consecutive levels.

1.5.5 Fusion rules with various properties

The fusion paradigm can be used to construct interesting examples where the standard results from substitution systems need not apply. Here are some that appear in [50].

1. Example 3.7 provides us with an example of a one-dimensional, prototile- but not transition-regular fusion rule that has a minimal but not uniquely ergodic dynamical system.
2. Again prototile-regular, example 3.8 shows how a measure arising from a sequence of supertiles can fail to be ergodic.
3. Example 4.4, the "scrambled Fibonacci", is based on the Fibonacci substitution/fusion, but is systematically altered at occasional levels. The alterations are enough to eliminate topological point spectrum, but the system still has measurable eigenvalues. Thus we have a system that is topologically weakly mixing but measure-theoretically pure point. It appears in [65], along with a related example where the measurable and topological eigenvalues differ.
4. Example 4.8 gives us an example of a uniquely ergodic symbolic fusion system that has 'coincidence with finite waiting' but is not pure point spectrum. This contrasts with Dekking's classical result on coincidence for constant-length symbolic substitutions [33], where coincidence implies pure pointedness.
5. Example 4.9 provides a one-dimensional symbolic fusion that is not prototile-regular, where not only is the system not uniqely ergodic, but also the ergodic measures can have different spectral types.
6. Example 4.11 is an example of a symbolic fusion rule that is strictly ergodic and yet has positive entropy.

1.6 Spectral analysis of supertile methods: Dynamical spectrum

For the remainder of these lectures we discuss the two main spectral methods for analyzing tiling spaces: dynamical and diffraction. Spectral theory of dynamical systems is widely used and provides a measure-theoretic tool that standardizes spaces acted on by a group $G \subseteq \mathbb{R}^d$ by comparing them to Lebesgue measure on the dual group of G. This is achieved using what is now called the Koopman representation, which represents G as unitary operators on $L^2(\Omega, \mu)$, and applying the spectral theorem for unitary operators. A nice development of the subject for $G = \mathbb{Z}$ or $\mathbb{Z}^+$ can be found in [84, 104]; also for $G = \mathbb{Z}$ [56] provides a well-contextualized historic overview and survey of results up to 1999. The book [89] is entirely devoted to analyzing symbolic substitutions in one dimension using spectral theory, as is Chapter 7 of [42]. The paper [19] extends many of these results to the multidimensional constant-length $\mathbb{Z}^d$ substitution case. The dynamical spectrum framework for self-similar tilings of $\mathbb{R}^d$ is laid out in [100]. One can see [11, Appendix B] for a brief discussion of how the theory for $\mathbb{Z}$ actions can be extended to the higher-dimensional and/or continuous case, but the author is not aware of any survey of the spectral theory of tiling dynamical systems.

Strong motivation for studying the diffraction spectrum of tilings comes from the quasicrystal model. Indeed, Shechtman's Nobel prize-winning discovery of quasicrystals arose from a diffraction experiment [97]. A lovely and somewhat underappreciated early book on crystals, quasicrystals, tilings, diffraction, and the history of the subject is [96], which bridges the gap between physics and mathematics for those looking for perspective. Fourier analysis is used to define virtual diffraction experiments on tilings, and [11] is an excellent source to learn about how it works. The method was originally proposed in [60], and there is a fundamental argument in [37] that allows us to see that the diffraction spectrum is related to the dynamical spectrum. There has been special emphasis on the discrete (a.k.a. point or atomic) part of the diffraction spectrum of tilings because it represents the bright spots appearing on a diffraction image known as "Bragg peaks". In particular, extensive work has been done to determine conditions under which there exists such point spectrum, and when the spectrum is composed solely of it. We will discuss a selection of results in this direction as well as what is known about the connection between the two types of spectral analysis.

1.6.1 The Koopman representation

Let (Ω, G) be a dynamical system with G representing a group of translations. Suppose μ is a translation-invariant ergodic Borel probability measure for the system. The function space $L^2(\Omega, \mu)$ is a Hilbert space and one often looks at it when trying to analyze a system. From a physics perspective one could consider a function as taking measurements or running experiments on the tilings in the tiling space.

For each $\jmath \in G$ there is a unitary operator $U^\jmath : L^2(\Omega,\mu) \to L^2(\Omega,\mu)$ defined by

$$U^\jmath(f)(\mathcal{T}) = f(\mathcal{T} - \jmath).$$

This family of operators is sometimes called the Koopman operator and it is a representation of G. Since $L^2(\Omega,\mu)$ is a separable Hilbert space the tools of operator theory are available. The spectrum of the Koopman operator is called the *dynamical spectrum* of Ω (or of the supertile method that generated it).

Every $f \in L^2(\Omega,\mu)$ has a spectral measure associated with it. No matter which construction method was used, all of the cases look like this: for $\jmath \in G$ we define $\hat{f}(\jmath) = \int_\Omega f(\mathcal{T} - \jmath)\overline{f(\mathcal{T})}d\mu(\mathcal{T})$. One can think of comparing the values of f at two spots in $\mathcal{T}$, separated by $\jmath$, and averaging the result over all of Ω. In each of our situations this satisfies the appropriate notion of positive definiteness so that the appropriate version of Bochner's theorem [13] guarantees us the existence of a positive real-valued measure σ_f on $\mathbb{T}^d$ with these same Fourier coefficients. That is,

$$\hat{f}(\jmath) = \int_\Omega f(\mathcal{T} - \jmath)\overline{f(\mathcal{T})}d\mu(\mathcal{T}) = \int_{\hat{G}} z^\jmath d\sigma_f(z), \tag{1.1}$$

where $z^\jmath := z_1^{j_1} \cdots z_d^{j_d}$ and the dual group to G is denoted $\hat{G}$.

The *spectral type* of a function f, then, is determined by how σ_f decomposes with respect to Lebesgue measure. Is it discrete, singular continuous, or absolutely continuous? Or perhaps a mix? The supertile rule that determines Ω ultimately determines what is possible for these spectral types and is our primary interest in this topic.

Each $f \in L^2(\Omega,\mu)$ generates a cyclic subspace of L^2 given by its closed linear span:

$$Z(f) = \overline{\text{span}\,\{U^\jmath f \text{ such that } \jmath \in G\}}.$$

Part of the spectral analysis of operators involves finding generating functions $f_i, i = 1, 2, ...$ for which

$$L^2(\Omega,\mu) = \bigoplus Z(f_i).$$

This decomposition is not exactly unique but the number and spectral types of the functions are. It is possible to find functions $f_1, f_2, f_3, ...$ such that $L^2(\Omega,\mu) = \bigoplus Z(f_i)$ and for which $\sigma_{f_1} \gg \sigma_{f_2} \gg \sigma_{f_3} \gg ....$ Again the func-

[13] The general results on spectral theory of dynamical systems in this section can be found in many places, for instance [56, 71]; specialization to the tiling case first appears in [100].

tions are not unique, but their spectral types are and the spectral type of f_1 is known as the *maximal spectral type* of the system. This decomposition determines the Koopman operator up to unitary equivalence.

1.6.2 Eigenfunctions

To begin thinking more about the spectrum of $U^{\boldsymbol{\jmath}}$ we can investigate its eigenvalues and eigenvectors.The easiest case is when $G = \mathbb{Z}$. In that case we are looking at any functions $f \in L^2(\Omega, \mu)$ for which there is some λ such that $Uf = \lambda f$. In general f is known as a *measurable eigenfunction*, and if it happens to be continuous then it is called a *topological eigenfunction*. Notice that since U is unitary it must be that $|\lambda| = 1$ and we write $\lambda = e^{2\pi i \alpha}$ for some $\alpha \in \mathbb{R}$. Obviously once we have an eigenvalue/eigenfunction pair then for any $n \in \mathbb{Z}$ we have that $U^n f = \lambda^n f = e^{2\pi i \alpha n} f$.

Now let's generalize to a continuous one-dimensional action, i.e. when $G = \mathbb{R}$. An *eigenfunction* is a function for which there exists an $\alpha \in \mathbb{R}$ such that for all $x \in \mathbb{R}$ it is true that $U^x f = e^{2\pi i \alpha x} f$. That is, for all $\mathcal{T} \in \Omega$ it is true that $f(\mathcal{T} - x) = e^{2\pi i \alpha x} f(\mathcal{T})$.

When G is $\mathbb{Z}^d$ or $\mathbb{R}^d$ with $d > 1$, the eigenvalues themselves live in higher dimensions and the inner product becomes necessary. In this situation we say f is an eigenfunction if there exists an $\boldsymbol{\alpha} \in \mathbb{R}^d$ for which $U^{\boldsymbol{\jmath}} f = \exp(2\pi i \boldsymbol{\alpha} \cdot \boldsymbol{\jmath}) f$ for all $\boldsymbol{\jmath} \in G$. That is, for all $\mathcal{T} \in \Omega$ and all $\boldsymbol{\jmath} \in G$ we find $f(\mathcal{T} - \boldsymbol{\jmath}) = \exp(2\pi i \boldsymbol{\alpha} \cdot \boldsymbol{\jmath}) f(\mathcal{T})$.

An important technical point, which seems to be a point of contention, is whether it is λ or α that is considered the eigenvalue for f. Both perspectives have merit. The argument for using λ is that it is standard usage in functional analysis for the spectrum of the unitary operator. The argument for α is that it resides in the dual group of G and therefore is more directly relevant to diffraction analysis and abstract harmonic analysis. We will allow either to count as the eigenvalue, being more precise when necessary.

There is a difference between the situation where $G = \mathbb{Z}^d$ and $G = \mathbb{R}^d$ that also manifests itself when we allow $G = \mathbb{R}^d$ but the tiling is a suspension of a $\mathbb{Z}^d$ action. When $G = \mathbb{Z}^d$ or the tiling space is a suspension of a $\mathbb{Z}^d$ action, then every $\boldsymbol{\alpha} \in \mathbb{Z}^d$ is an eigenvalue of the system; that is, in these cases there is always some discrete spectrum. In fact when $G = \mathbb{Z}^d$, the dual group is $\mathbb{T}^d$ and we have that if $\boldsymbol{\alpha}$ is any eigenvalue, then $\boldsymbol{\alpha} + \boldsymbol{\jmath}$ is also an eigenvalue for the same eigenfunction, so the spectrum is only considered on $\mathbb{T}^d$.

When $G = \mathbb{R}^d$ but the system is a suspension of a $\mathbb{Z}^d$-action, all elements of $\mathbb{Z}^d$ continue to be eigenvalues but we get that $\boldsymbol{\alpha}$ and $\boldsymbol{\alpha} + \boldsymbol{\jmath}$ have different eigenfunctions for $\boldsymbol{\jmath} \in \mathbb{Z}^d$. However, the eigenfunctions are closely related: if we let f denote the eigenfunction for $\boldsymbol{\alpha}$ and $f_{\boldsymbol{\jmath}}$ that of $\boldsymbol{\alpha} + \boldsymbol{\jmath}$, then $f_{\boldsymbol{\jmath}}(\mathcal{T}) = \exp(2\pi i \boldsymbol{\jmath} \cdot \mathbf{x}(\mathcal{T})) f(\mathcal{T})$, where $\mathbf{x}(\mathcal{T})$ is the location in $\mathbb{R}^d$ of any vertex[14] of $\mathcal{T}$. Put another way, the eigenfunction for $\boldsymbol{\alpha} + \boldsymbol{\jmath}$ is the product of the eigenfunction for $\boldsymbol{\alpha}$ with a function that keeps track of where the tiling is relative to the suspension.

When $G = \mathbb{R}^d$ and the space cannot be seen as the suspension over a lattice we must consider all possible values of the dual group of $\mathbb{R}^d$, which is still $\mathbb{R}^d$. Thus spectral images in this case are not restricted to a torus.

Example 14 *Let's compute the spectral measure of an eigenfunction f of (Ω, μ) with eigenvalue α. For $\boldsymbol{\jmath} \in G$ we have*

$$\hat{f}(\boldsymbol{\jmath}) = \int_\Omega \exp(2\pi i \alpha \cdot \boldsymbol{\jmath}) \overline{f(\mathcal{T})} f(\mathcal{T}) d\mu(\mathcal{T}) = \exp(2\pi i \alpha \cdot \boldsymbol{\jmath}),$$

since eigenfunctions are of almost everywhere constant modulus that can be taken to be 1. Thus the spectral measure of f is a measure σ_f on $\mathbb{T}^d$ with these Fourier coefficients. One can check that the measure on $\mathbb{T}^d$ with these coefficients is the atomic measure supported on α, and so $\sigma_f = \delta_\alpha$.

1.6.2.1 Conditions for presence of discrete spectrum

It is not surprising that spectral properties were investigated soon after the dynamical systems approach to substitution sequences was introduced. Coven and Keane [32] investigated a class of examples and their approach was generalized by Martin [80]. Dekking [33] generalized these results using different methods, completely determining the point part of the dynamical spectrum of constant-length substitutions in one dimension.

Theorem 1.6.1 ([33], quoted from Section 7.3 of [42]) *Let σ be a non-periodic (symbolic) substitution of constant length n. Let u be a periodic point for σ. We call the* height *of the substitution the greatest integer m which is coprime with n and divides all the strictly positive ranks of occurrence of the letter u_0 in u. The height is less than the cardinality of the alphabet. The maximal equicontinuous factor*[15] *of the substitutive dynamical system associated with σ is the addition of $(1,1)$ on the abelian group $\mathbb{Z}_n \times \mathbb{Z}/m\mathbb{Z}$.*

One would expect the discrete spectrum of a symbolic substitution of non-constant length to have a connection to the expansion factor of the system, and Host's result [62] gives a criterion that we quote here, following [42, Chapter 7]. We leave undefined the term "coboundary", which we will not be using again.

Theorem 1.6.2 ([62], quoted from Section 7.3 of [42]) *Let σ be a not shift-periodic and primitive substitution over the alphabet $\mathcal{A}$. A complex num-*

[14] One can show that this term is independent of the choice of vertex.

[15] Recall that this is the largest topological factor of the dynamical system that is a rotation of a compact group.

ber λ of modulus one is an eigenvalue of $(\Omega, \mathbb{Z})$ if and only if there exists $p > 0$ such that for every $a \in \mathcal{A}$, the limit $h(a) = \lim_{n\to\infty} \lambda^{|\sigma^{pn}(a)|}$ is well defined, and h is a coboundary of σ.

The constant function 1 is always a coboundary, making it simpler to check:

Theorem 1.6.3 (Corollary 7.3.17 of [42]) *Let σ be a not shift-periodic and primitive substitution over the alphabet $\mathcal{A}$. If there exists p such that $\lambda \in \mathbb{C}$ satisfies $\lim_{n\to\infty} \lambda^{|\sigma^{pn}(a)|} = 1$ for every $a \in \mathcal{A}$, then λ is an eigenvalue of the substitutive dynamical system associated with σ.*

One of the main results in [100] is the characterization of eigenvalues and eigenfunctions for self-similar tiling systems. The method of proof is constructive in that given the eigenvalue condition, an eigenfunction is constructed whose value for a tiling $\mathcal{T}$ depends on special points derived from the supertile structure of $\mathcal{T}$.

In [100], general results on the presence or absence of eigenfunctions and therefore on weak mixing are determined. There are results for self-affine tilings of $\mathbb{R}^d$ that are made stronger in the $d = 1$ and 2 cases. The statement presented here is essentially quoted from that paper, except that the result of [98] is taken into account. Note that the set of translations between tiles in $\mathcal{T}$ is given by $\Xi(\mathcal{T})$, where $x \in \Xi(\mathcal{T})$ if and only if x is a translation taking a tile $t \in \mathcal{T}$ to an equivalent tile in $\mathcal{T}$, i.e. $t - x \in \mathcal{T}$.

Theorem 1.6.4 (Theorem 5.1 of [100]) *(i) Let $\mathcal{T}$ be a nonperiodic self-affine tiling of $\mathbb{R}^d$ with expansion map ϕ. Then $\alpha \in \mathbb{R}^d$ is an eigenvalue of the measure-preserving system (Ω, μ) if and only if*

$$\lim_{n\to\infty} e^{2\pi i(\phi^n(x)\cdot\alpha)} = 1 \quad \text{for all} \quad x \in \Xi(\mathcal{T}). \tag{1.2}$$

Moreover, if equation 1.2 holds, the eigenfunction can be chosen continuous[16].
(ii) Let $\mathcal{T}$ be a self-similar tiling of $\mathbb{R}$ with expansion constant λ. The tiling dynamical system is not weakly mixing if and only if $|\lambda|$ is a real Pisot number. If λ is real Pisot and $\mathcal{T}$ is nonperiodic, there exists nonzero $a \in \mathbb{R}$ such that the set of eigenvalues contains $a\mathbb{Z}[\lambda^{-1}]$.
(iii) Let $\mathcal{T}$ be a nonperiodic self-similar tiling of $\mathbb{R}^2 \equiv \mathbb{C}$ with expansion constant $\lambda \in \mathbb{C}$. Then the tiling dynamical system is not weakly mixing if and only if λ is a complex Pisot number. Moreover, if λ is a non-real Pisot number there exists nonzero $a \in \mathbb{C}$ such that the set of eigenvalues contains $\{(\alpha_1, \alpha_2) : \alpha_1 + i\alpha_2 \in a\mathbb{Z}[\lambda^{-1}]\}$.

[16] This is also true for substitution sequences [62], but not for general fusions.

1.6.3 Pure discrete dynamical spectrum

Dekking's work in [33] was the first to define a notion now known as the "coincidence condition" for a symbolic substitution σ. It is that there are numbers k and l such that the image of any letter of the alphabet under σ^k has the same lth letter. This combinatorial condition is easy to check in any given example and eliminates any spectrum that is not discrete. The idea has been generalized to non-constant length symbolic substitutions and to self-similar tiling systems with the goal of characterizing purely discrete spectrum in those cases. Algorithms have been developed such as the "balanced pair algorithm" for substitution sequences (see [79] and references therein or the original source [75], stated for adic transformations on Markov compacta). For the multidimensional case there is a series of papers [4, 5, 3] on the "overlap algorithm" making checkable conditions; the original version of this for tilings appears in [100]. In terms of spectral analysis of supertile systems, spaces with a purely discrete dynamical spectrum are the most well understood.

1.6.3.1 Symbolic substitutions and the Pisot substitution conjecture

Let us begin with Dekking's original result.

Theorem 1.6.5 ([33], quoted from Section 7.3 of [42]) *Let σ be a substitution of constant length and of height 1. The substitutive dynamical system associated with σ has a purely discrete spectrum if and only if the substitution σ satisfies the condition of coincidence.*

The situation is not settled in the non-constant length substitution case. A substitution σ satisfies the *strong coincidence condition* if there are integers k and l such that for every $a, b \in \mathcal{A}$, the substitutions $\sigma^k(a)$ and $\sigma^k(b)$ not only have the same lth letter, but the prefixes of length $l-1$ in each have the same number of letters of each type. The latter part of this condition ensures that $\sigma^{(k+j)}(e)$ will have coinciding j-supertiles at the corresponding location for all $j \geq 1$.

It is thought, but not known, that algebraic properties of the transition matrix can determine coincidences. In particular, a substitution is said to be of *Pisot type* if all of the eigenvalues of its transition matrix except for the Perron-Frobenius eigenvalue have modulus strictly between 0 and 1. It is said to be *irreducible* if the characteristic polynomial is irreducible.

There are a family of conjectures collectively known as "Pisot substitution conjectures", that all more or less say, "the substitution dynamical system has pure discrete spectrum if it is of irreducible Pisot type". These conjectures are in place not only for one-dimensional symbolic substitutions, but also for one-dimensional tiling substitutions and a few other sorts of substitutions;

the situation is nicely summarized in [2]. Immediately relevant to our work following two conjectures cited there that are equivalent by [31].

Conjecture 1 (Pisot substitution conjecture: symbolic substitutive case). If σ is an irreducible Pisot substitution then the substitutive system $(\Omega, \mathbb{Z})$ has pure discrete spectrum.

Conjecture 2 (Pisot substitution conjecture: one-dimensional tiling case). If $\mathcal{S}$ is an irreducible Pisot substitution for one-dimensional tilings, then its tiling dynamical system $(\Omega, \mathbb{R})$ has pure discrete spectrum.

Progress has recently been made to settle the second conjecture: in [17] the conjecture is confirmed for substitutions that are injective on initial letters and constant on final letters. Closely related is the question of whether the substitution satisfies the strong coincidence condition. Again from [2]:

Conjecture 3 (Strong coincidence conjecture). Every irreducible Pisot substitution satisfies the strong coincidence condition.

The following two theorems together settle the Pisot substitution conjecture for two-letter substitutions:

Theorem 1.6.6 ([61]) *Let σ be a substitution of Pisot type over a two-letter alphabet which satisfies the coincidence condition. Then the substitution dynamical system associated with σ has a purely discrete spectrum.*

Theorem 1.6.7 ([18]) *Any substitution of Pisot type over a two-letter alphabet satisfies the coincidence condition.*

The question of whether a substitution sequence has purely discrete dynamical system if and only if its expansion factor is Pisot remains open for alphabets of size 3 or higher and has proved remarkably difficult to resolve. The methods used in the two-letter case don't apply in these cases. A survey of the state of results as of 2015 (not including [17]) appears in [2].

1.6.3.2 Purely discrete spectrum for supertile methods in $\mathbb{R}^d$

There is a sufficient condition given in [100] that works in all dimensions, and then a specialized version for $\mathbb{R}^2$ we will talk about.

Theorem 1.6.8 *[100, Theorem 6.1] Let $\mathcal{T}$ be a self-affine tiling of $\mathbb{R}^d$ with expansion map ϕ. If there exists a basis B for $\mathbb{R}^d$ such that for all $x \in B$,*

$$\sum (1 - dens(D_{\phi^n(x)}) < \infty, \tag{1.3}$$

then the tiling dynamical system $(\Omega_{\mathcal{T}}, \mathbb{R}^d, \mu)$ has pure discrete spectrum.

The strongest result is obtained for $\mathbb{R}$ and $\mathbb{R}^2$.

Theorem 1.6.9 *[100, Theorem 6.2] Let $\mathcal{T}$ be a self-similar tiling of $\mathbb{R}^d$, $d \leq 2$, with expansion constant λ. The tiling dynamical system $(\Omega_{\mathcal{T}}, \mathbb{R}^d, \mu)$ has pure discrete spectrum if and only if λ is Pisot (real or non-real) and*

$$\lim_{n\to\infty} dens(D_{\lambda^n x}) = 1, \qquad x \in \Xi[\mathcal{T}]. \tag{1.4}$$

An *overlap algorithm* is defined in [100, p. 721-724] that determines whether a self-similar tiling has a pure discrete spectrum. We give only the general idea of the algorithm here. One makes a graph $G_O(\mathcal{T}, x)$ out of all overlaps one can see by comparing the tile(s) in $\mathcal{T}$ at some $z \in \mathbb{R}^2$, with the tile(s) in $\mathcal{T} - x$ at z. Most of those overlaps will be two tiles of different types, but sometimes they are of the same type, and on occasion x may be a return vector for the tiles at this particular z. When that happens we are said to have a *coincidence*, properly defined as $\mathcal{T} \cap \{z\} = (\mathcal{T} - x) \cap \{z\}$. An overlap is a vertex in the overlap graph, and edges connect overlaps that are related via substitution.

Theorem 1.6.10 *[100, Proposition 6.7] Suppose that $\mathcal{T}$ is a self-similar tiling of the plane with expansion constant λ a non-real Pisot number, and $x \in \Xi[\mathcal{T}]$. The following are equivalent:*

(i) the tiling dynamical system $(\Omega_{\mathcal{T}}, \mathbb{R}^2, \mu)$ has pure discrete spectrum,
(ii) from any vertex of $G_O(\mathcal{T}, x)$ there is a path leading to a coincidence, and
(iii) $dens(D_{\lambda^n x}) \to 1$ as $n \to \infty$.

1.6.4 The continuous part of the spectrum

The results given above mean that we have a pretty good understanding of the discrete part of the spectrum for supertile systems. In particular we understand when there is point spectrum, and where the point masses are. We also understand when there must be a continuous component to the spectrum. What remains is to understand the nature of the continuous part. For instance, is it singular or absolutely continuous with respect to Lebesgue measure?

There are two classic examples of one-dimensional constant-length substitution sequences[17] with mixed spectrum: the Thue-Morse sequence and the Rudin-Shapiro sequence. Thue-Morse is given by the bijective substitution $0 \to 01, 1 \to 10$ which, lacking coincidences, was long known to have some continuous spectrum. A variety of results over the years proved that the continuous part is singular with respect to Lebesgue measure. It has been thought that the continuous part of the spectrum of a bijective substitution is always singular, but that question remains open. However, if the alphabet

[17] These sequences were actually defined using number-theoretic constructions, but have simple substitution rules also.

only has two symbols then singularity has been proved in [12], in one and several dimensions.

The original Rudin-Shapiro sequence is not bijective but it lacks coincidences and therefore has a continuous spectral component that has been known to be absolutely continuous for some time (see [89, 42]). Generalizations to higher dimensions were developed in [44], where the continuous part of the dynamical spectrum was shown to be absolutely so. Recent work [30] gives a different generalization and shows the continuous part of the diffraction spectrum to be absolutely continuous. The "twisted silver mean" substitution, introduced in [10], uses a 'bar swap' method seen in some of the Rudin-Shapiro constructions, but on a non-constant length substitution. The result is a mixed spectrum, which is analysed fully in [10] and found to have a mixed but singular spectrum.

For substitutions of length q, a necessary condition for the presence of absolutely continuous spectrum had been conjectured: The transition matrix should have an eigenvalue of modulus $\sqrt{q}$. This conjecture was verified in [23] with the result that if the transition matrix has no eigenvalue of modulus $\sqrt{q}$ then the dynamical spectrum is singular.

The results and techniques of [89] are given an important generalization to multidimensional constant-length substitutions in [19]. Without requirements on primitivity or height, Bartlett is able to show that "abelian" bijective substitutions have only singular continuous spectrum, settling the longstanding conjecture in these cases. A general algorithm for computing the spectrum for constant-length multidimensional substitutions in $\mathbb{Z}^d$ is also given.

Example 15 *(Explicit computations for the Thue-Morse substitution.) Let* Ω_{TM} *be the subshift for the Thue-Morse symbolic substitution* $0 \to 01, 1 \to 10$. *We will show how to make all of the eigenfunctions for the dynamical system, and also how to make a function whose measure is (singular) continuous with respect to Lebesgue measure. Together these functions generate all of* L^2.

To make an example of a nontrivial eigenfunction we consider the 2-supertiles in $\mathcal{T}$. *Since 2-supertiles have length 4 in this example, there are four locations the origin can occupy in either type of 2-supertile. Let us call them 0, 1, 2, 3 as we go from left to right. Define* $\mathcal{O}(\mathcal{T}) = i$ *if the origin occupies the* i*th location of its 2-supertile. Now define* $f(\mathcal{T}) = \exp(2\pi i \mathcal{O}(\mathcal{T})/4)$.

Notice that if $\mathcal{O}(\mathcal{T}) = i$ *and* $i < 3$ *then* $\mathcal{O}(\mathcal{T} - 1) = i + 1$. *If* $\mathcal{O}(\mathcal{T}) = 3$ *then* $\mathcal{O}(\mathcal{T} - 1) = 0$.[18] *Thus if* $\mathcal{O}(\mathcal{T}) < 3$ *we get that* $U(f)(\mathcal{T}) = f(\mathcal{T} - 1) = \exp(2\pi i(\mathcal{O}(\mathcal{T}) + 1)/4)$ *and if* $\mathcal{O}(\mathcal{T}) = 3$ *then* $U(f)(\mathcal{T}) = 1$. *This means that* $U(f) = \exp(2\pi i/4)f$, *and so* f *is an eigenfunction with eigenvalue* $\alpha = 1/4$.

---

[18] This is indicative of the 'odometer'-like structure of constant-length substitutions: the shift map augments until a fixed number and then resets to 0, augmenting elsewhere. In general the supertile structure of any constant-length substitution looks like an odometer, see for example [45].

We compute the Fourier coefficient $\hat{f}(n)$*, as defined in equation (1.1). We know*

$$U^n(f)(\mathcal{T}) = \exp(2\pi i(\mathcal{O}(\mathcal{T} - n))/4)$$
$$= \exp(2\pi i(\mathcal{O}(\mathcal{T}) + n)/4) = \exp(2\pi in/4)f(\mathcal{T}),$$

and so

$$f(n) = (U^n(f), f) = \int_{\Omega_{TM}} \exp(2\pi in/4)f(\mathcal{T})\overline{f(\mathcal{T})}d\mu(\mathcal{T})$$
$$= \int_{\Omega_{TM}} \exp(2\pi in/4)d\mu(\mathcal{T}) = \exp(2\pi in/4)$$

for all n*. These are the Fourier coefficients of a Dirac* δ*-function with its peak at* $\exp(2\pi in/4)$*. Since* σ_f *is unique this makes it equal to this Dirac delta function.*

Variations of the function f *are easy to construct by looking at different sizes of supertiles. If* f *is based on the location of the origin in an* n*-supertile, it will have an eigenvalue with a* 2^n *in the denominator. Since this substitution has height 1, this implies that the eigenvalues of the Koopman operator for the Thue-Morse substitution are* $\mathbb{Z}[1/2]$*.*

None of the functions we just constructed depend on the actual letter at the origin. We can supplement these with a function that only knows what letter is at the origin in $\mathcal{T}$*: Define* $g(\mathcal{T}) = \begin{cases} 1 & \text{if } \mathcal{T}(0) = 1 \\ -1 & \text{if } \mathcal{T}(0) = 0 \end{cases}$*. It is shown in [45] that together this and the eigenfunctions span all of* $L^2(\Omega_{TM}, \mu)$*.*

Notice that g *is orthogonal to the eigenfunction* f *since*

$$< f, g >= \int_{\Omega_{TM}} f(\mathcal{T})\overline{g(\mathcal{T})}d\mu(\mathcal{T}) = \int_{\Omega^+_{TM}} f(\mathcal{T})d\mu(\mathcal{T}) - \int_{\Omega^-_{TM}} f(\mathcal{T})d\mu(\mathcal{T}),$$

where Ω^+_{TM} *(resp.* Ω^-_{TM}*) are the set of all tilings with a 1 (resp. 0) at the origin. Each of the two integrals on the right are equal because they depend only on the supertile structure of* $\mathcal{T}$ *and not the letter at the origin. Thus the inner product of* g *and* f *is 0. We have obtained:*

$$L^2(\Omega_{TM}, \mu) = H_0 \bigoplus Z(g), \tag{1.5}$$

where H_0 *denotes the span of the eigenfunctions.*

The function space L^2 of general bijective substitutions in $\mathbb{Z}^d$ breaks into a direct sum of pieces that are discrete or continuous analogously to this example. If the bijections comprising the substitution commute with translation, it is possible to explicitly define the generators of the continuous spectral pieces [45, Theorem 4.2]. The nature of the continuous part of the spectrum continues to be investigated.

1.7 Spectral analysis of supertile methods: Diffraction spectrum

The diffraction spectrum of tilings is motivated by physics. In this viewpoint we consider the tiling as representing the atomic structure of a solid and we wish to mathematically simulate what happens in a diffraction experiment on the solid. That is, one passes x-rays or electrons through the solid, where they will bounce off atoms and interfere constructively and destructively, ultimately creating an image that represents something about the structure they passed through. Fourier analysis turns out to be the right mathematical analogue for this. We describe the situation for symbolic dynamics first, then generalize to $\mathbb{R}^d$.

1.7.1 Autocorrelation for symbolic sequences

Consider a sequence $\mathcal{T}_0 \in \Omega \subset \mathcal{A}^{\mathbb{Z}}$, where for convenience we assume that $\mathcal{A}$ is a finite subset of the complex numbers. We know that constructive and destructive interference depends on the repetition at various distances in $\mathcal{T}_0$. For instance, if $\mathcal{T}_0$ was periodic then there would be total agreement at distances that are multiples of the period, leading to strong constructive interference at those distances.

A reasonable way to measure the extent to which $\mathcal{T}_0$ agrees with itself at a distance $k \in \mathbb{Z}$ is to consider the global average of $\mathcal{T}_0(n-k)\overline{\mathcal{T}_0(n)}$ over all n. Thus we define a *correlation function* $C : \mathbb{Z} \to \mathbb{C}$ to be a cluster point of the sequences

$$\left\{ \frac{1}{N} \sum_{n=0}^{N-1} \mathcal{T}_0(n-k)\overline{\mathcal{T}_0(n)} \right\}.$$

A diagonalization argument shows that for a given C there will be some sequence $\{N_j\}$ such that

$$C(k) = \lim_{j \to \infty} \frac{1}{N_j} \sum_{n=0}^{N_j - 1} \mathcal{T}_0(n-k)\overline{\mathcal{T}_0(n)}.$$

For most examples of interest from primitive supertile methods, notably the uniquely ergodic ones, it turns out that $C(k)$ is unique, and it is useful to assume that. (The general situation is described in detail starting on page 74 of [89]).

Suppose μ is an ergodic measure on Ω and consider the continuous function $\mathcal{O} : \Omega \to \mathbb{C}$ given by $\mathcal{O}(\mathcal{T}) = \mathcal{T}(0)$. Then for μ-almost every $\mathcal{T} \in \Omega$ we have that

$$C(k) = \lim_{j\to\infty} \frac{1}{N_j} \sum_{n=0}^{N_j-1} \mathcal{T}(n-k)\overline{\mathcal{T}(n)} = \int_{\Omega} \mathcal{O}(\mathcal{T}-k)\overline{\mathcal{O}(\mathcal{T})} d\mu(\mathcal{T}).$$

That means $C(k) = \hat{\mathcal{O}}(k)$ from a dynamical spectrum perspective.

On the other hand, $C(k)$ can be shown to be positive definite and so there is a positive measure on the torus that has $C(k)$ as its Fourier coefficients. This measure, which we denote $\hat{\gamma}$, is known in [89] as the correlation measure of the sequence $\mathcal{T}$; it is the analogue of the diffraction measure (In general the diffraction measure is the Fourier transform of the autocorrelation, which is defined similar to $C(k)$). One can see from this analysis that the diffraction spectrum should be subordinate to the dynamical spectrum.

1.7.2 Diffraction in $\mathbb{R}^d$

1.7.2.1 Overview

A tiling $\mathcal{T}$ is a model for the atomic structure of matter, where the atoms or molecules occupy locations given by the tiles. In our simulation of a diffraction experiment, we imagine that waves of some appropriate wavelength are sent through the tiling, where they interfere constructively and destructively as determined by relative distances between the tiles. The diffracted waves form an image where we see bright spots of intense constructive interference (our "Bragg peaks") and a greyscale spectrum were the interference ranges from constructive to destructive.

The mathematics of diffraction has a long development that is based on Fourier analysis. Because our tilings are infinite there are technicalities that have to be handled using tempered distributions and translation-bounded measures. It was Hof in [60] who first advocated using this overall method to approach the diffraction of aperiodic structures and Dworkin [37] who noticed the connection between diffraction and dynamical spectrum; [15] provides a recent and quite accessible survey. An early computation for self-similar tilings is [52]. A serious treatment of the details as well as the history behind mathematical diffraction appears in Chapters 8 and 9 of [11], along with numerous references. A more condensed and self-contained description of the diffraction spectrum appears in [69], and we loosely follow that development here.

Before we begin, consider this intuitive description of the mathematics of diffraction that appears in [72, Section 5]. It clearly shows why the Fourier transform is central to the theory.

> "When modeling diffraction, the two basic principles are the following: Firstly, each point x in the solid gives rise to a wave $\xi \mapsto \exp(-ix\xi)$. The overall wave w is the sum of the single waves. Secondly, the quantity measured in an

experiment is the intensity given as the square of the modulus of the wave function.

"We start by implementing this for a finite set $F \subset \mathbb{R}^d$. Each $x \in F$ gives rise to a wave $\xi \mapsto \exp(-ix\xi)$ and the overall wavefunction w_F induced by F is accordingly

$$w_F(\xi) = \sum_{x \in F} \exp(-ix\xi).$$

Thus, the intensity I_F is

$$I_F(\xi) = \sum_{x,y \in F} \exp(-i(x-y)\xi) = \widehat{\sum_{x,y \in F} \delta_{x-y}}."$$

1.7.2.2 Diffraction via Delone sets

It is natural to consider diffraction theory on discrete sets in $\mathbb{R}^d$ called Delone sets, so we need to convert our tiling $\mathcal{T}$ into a point set $\mathbf{\Lambda}$ that represents the locations and types of atoms in the solid $\mathcal{T}$ represents.[19] We recall that a subset $\Lambda \subset \mathbb{R}^d$ is called a Delone set if there exist $0 < r \leq R$ such that every ball of radius r contains at most one point of Λ and every ball of radius R contains at least one point of Λ.

A *Delone multiset* is a set $\mathbf{\Lambda} = \Lambda_1 \times \Lambda_2 \times \cdots \times \Lambda_m$, where each Λ_i is a Delone set in $\mathbb{R}^d$ and the set $\bigcup_{i \leq m} \Lambda_i$, which by abuse of notation we also denote by $\mathbf{\Lambda}$, is Delone. An obvious way to turn $\mathcal{T}$ into a Delone multiset is to mark a special point in the prototile of type i for each $i = 1, .., m$, and let Λ_i be the Delone set of all copies of that point in $\mathcal{T}$.

So $\mathbf{\Lambda}$ represents our set of scatterers from $\mathcal{T}$ and we have kept track of the type of each scatterer. To account for different scattering strengths choose $a_i \in \mathbb{C}$ for $i \leq m$. Using the notation δ_x to represent the Dirac delta function at x thought of as a probability measure with support concentrated at x, we have the *weighted Dirac comb*

$$\omega = \sum_{i \leq m} a_i \delta_{\Lambda_i} = \sum_{i \leq m} a_i \sum_{x \in \Lambda_i} \delta_x.$$

This is a point measure on $\mathbb{R}^d$ that is not bounded, but is *translation bounded* in the sense that $\sup_{x \in \mathbb{R}^d} |\omega|(x+K) < \infty$ for all compact K.

The autocorrelation is defined to be the convolution of ω with the weighted Dirac comb $\tilde{\omega} = \sum_{i \leq m} \overline{a_i} \delta_{-\Lambda_i}$. Because convolutions are necessarily defined on measures with bounded support we end up with a limit that yields the *autocorrelation measure*[20]

---

[19] This volume contains an overview of the history and development of tilings and Delone sets in [101].

[20] This is also known as the "natural" autocorrelation measure because the averaging sets used are balls centered at the origin as opposed to an arbitrary van Hove sequence.

$$\gamma_\omega = \lim_{R\to\infty} \frac{1}{Vol(B_R(0))} \left(\omega|_{B_R(0)} * \tilde{\omega}|_{B_R(0)}\right) = \sum_{i,j\le m} a_i\overline{a_j} \sum_{z\in\Lambda_i-\Lambda_j} freq(z)\delta_z,$$

where the freqency is computed as the limit, if it exists, as the average number of times z is a return vector per unit area:

$$freq(z) = \lim_{R\to\infty} \frac{1}{Vol(B_R(0))} \#\{x \in \Lambda_i \cap B_R(0) \text{ and } x - z \in \Lambda_j\}.$$

Since $\mathcal{T}$ has finite local complexity, $\Lambda_i - \Lambda_j$ is a discrete set and that makes γ_ω a point measure also.

Definition 12. If the autocorrelation measure γ_ω exists, the *diffraction measure* of $\mathcal{T}$ is the Fourier transform $\widehat{\gamma_\omega}$.

From a physical perspective where we are running a diffraction experiment on a solid modeled by $\mathcal{T}$, the measure $\widehat{\gamma_\omega}$ tells us how much intensity is scattered into a given volume. We decompose $\widehat{\gamma_\omega}$ into its pure point, singular continuous, and absolutely continuous parts with respect to Lebesgue measure on $\mathbb{R}^d$:

$$\widehat{\gamma_\omega} = (\widehat{\gamma_\omega})_{pp} + (\widehat{\gamma_\omega})_{sc} + (\widehat{\gamma_\omega})_{ac}.$$

The pure point part tells us the location of the Bragg peaks that are so characteristic of the diffraction images of crystals and quasicrystals. The degree of disorder in the solid is quantified by the continuous parts. The singular continuous part is rarely observed in physical experiments [11, Remark 9.3].

Example 16 *The left column of figure 1.23 is a series of increasingly complex examples of self-similar tilings with two colors of square tiles. These tilings are examples of the sort analyzed in [44, 45]. Simulations of the corresponding diffraction images are also shown. Each tiling is a substitution of constant length* 2×2 *or* 4×4*. The tiling in the last row is a two-letter factor of a substitution on 8 letters; the other three are simple two-letter substitutions. In all cases there are point measures concentrated on* $\mathbb{Z}[1/2]\times\mathbb{Z}[1/2]$.

The top tiling has a purely discrete spectrum because its substitution satisfies the strong coincidence condition. The tiling in the middle of the figure is made from a bijective substitution and thus has a continuous component to its spectral measure. Because it is a constant-length symbolic substitution on two letters the continuous portion of the measure is singularly continuous with respect to Lebesgue measure [19].

The tiling on the right is a generalized Rudin-Shapiro tiling [44]. The original substitution is on eight letters and although it has no coincidence, it is not bijective. The tiling shown in the figure is a two-tile factor that is locally derived from the 8-letter substitution (and in fact the local derivability is mutual, so the factor makes no difference dynamically). The continuous portion of the spectral measure for this tiling is absolutely continuous.

It is interesting to simulate the diffraction images of these tilings in light of these theoretical results. Anyone who produces sample images of any sort

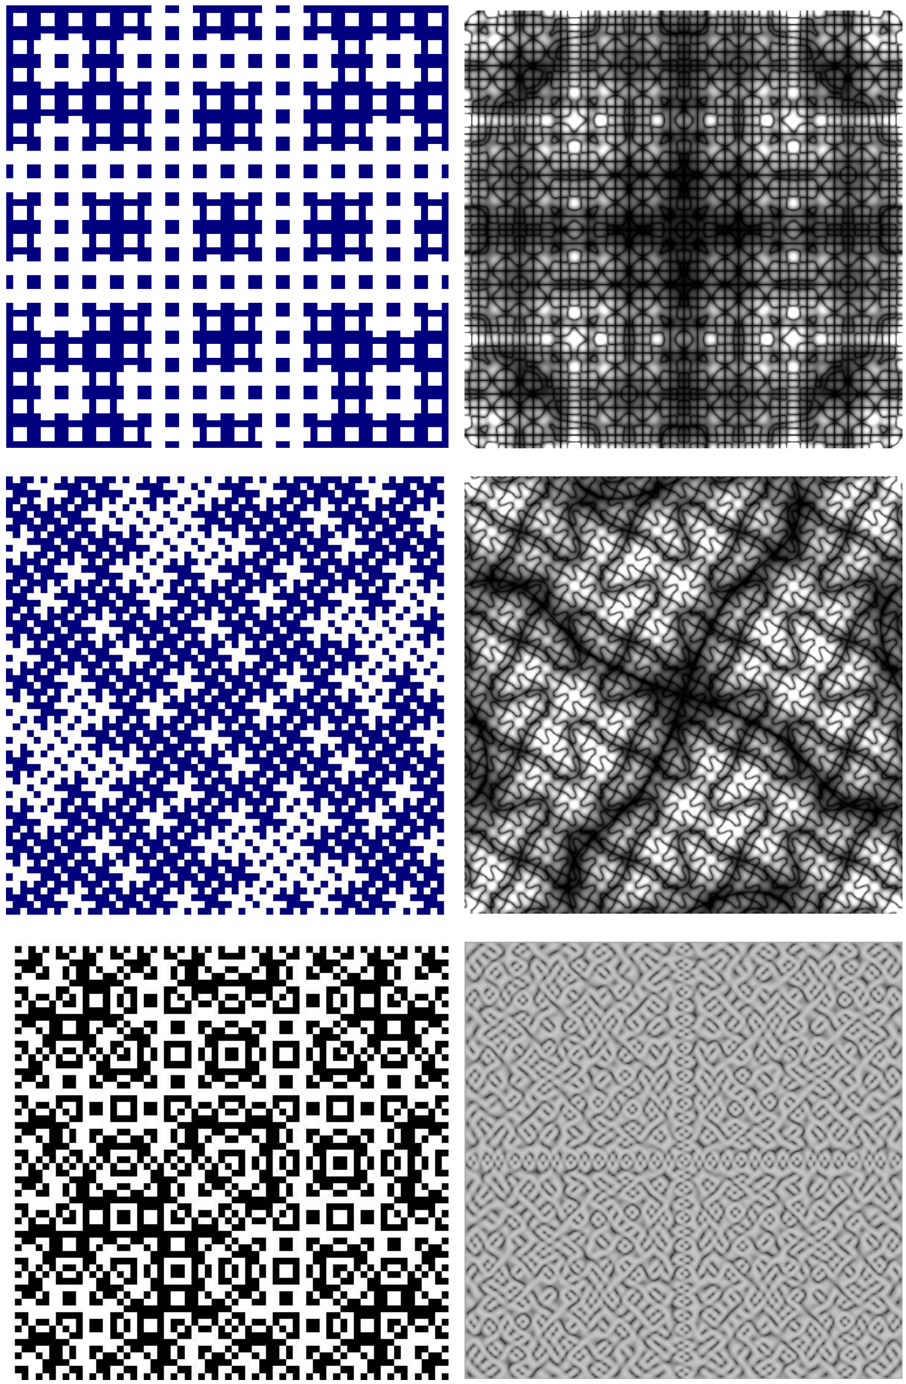

Fig. 1.23 The top tiling is substitutive with coincidence and it has a purely discrete spectrum. The substitution for the middle tiling is bijective and the spectrum is mixed, with a singular continuous part. The bottom tiling is not bijective but has a mixed spectrum with absolutely continuous part.

probably knows that there are usually parameters that can be altered to enhance the images. In our case such parameters include the weights on the dirac comb, the maximum intensity, and scaling functions. Tinkering with the parameters on a local scale does not change the overall qualitative appearance too much, and the apparent difference between the absolutely continuous diffraction (on the right) and the other two is persistent. The diffraction images for tilings with pure discrete spectrum and those with a singular component consistently appear similar throughout a wide range of parameters, with areas of extreme brightness and darkness. The absolutely continuous spectral images are notable for their lack of these extremes.

1.7.3 Intensities

How bright are the Bragg peaks? A preliminary formula was asserted in [28] and became known after a while as the "Bombieri/Taylor conjecture" (see also [60]). The formula, given below, is in terms of a limit. The convergence of this limit has been studied in many different situations, surveyed in [72]. In that paper Lenz shows that the formula is correct for a wide swath of aperiodic structures, including tilings generated through substitution and through projection, as well as those that are linearly repetitive. The setting in [72] is as follows.

Given a Delone set Λ, an element $\xi \in \mathbb{R}^n$, and a subset $B \subset \mathbb{R}^n$

$$c_B^{\xi}(\Lambda) = \frac{1}{Vol(B)} \sum_{x \in \Lambda \cap B} \exp(-2\pi i \xi \cdot x).$$

For the cases under consideration the intensity at $\xi \in \mathbb{R}^d$ is shown to exist and is given by

$$\widehat{\gamma(\xi)} = \lim_{n \to \infty} |c_{C_n}^{\xi}(\Lambda)|^2,$$

where C_n is the cube of side length $2n$ centered at the origin and $\xi \in \mathbb{R}^d$. In many of the situations discussed in [72] it is also proved that the eigenfunctions for the Koopman operator are continuous.

1.8 Connection between diffraction and dynamical spectrum

A recent survey of this topic is [15], which unifies the various notions of diffraction and dynamical spectrum, explains what was known up until 2016, and provides numerous references. Done in the context of Delone sets with finite local complexity, it applies to tilings of $\mathbb{R}^d$ and their dynamical systems.

In particular it explains the notions of diffraction for individual sets $\mathbf{\Lambda}$ as well as their hulls, and explicitly shows how to map from the Schwarz space of test functions under the diffraction measure to the Koopman representation of the dynamical system. Through this mapping they note that "the diffraction measure completely controls a subrepresentation of T", thus making explicit the connection between dynamical and diffraction spectrum.

The original paper connecting diffraction to dynamical spectrum is [37]. In it, Dworkin makes an argument showing how to deduce pure point diffraction spectrum if pure point dynamical spectrum has been established.

For a general system it may be that the diffraction spectrum does not contain as much information as the dynamical spectrum, but in the case of pure point spectrum it is known that the two classes are identical as long as there is unique ergodicity. The result, proved in [69], is in the context of Delone multisets.

Theorem 1.8.1 ([69]) *Suppose that a Delone multiset $\mathbf{\Lambda}$ has finite local complexity and uniform cluster frequencies. Then the following are equivalent:*

(i) *$\mathbf{\Lambda}$ has pure point dynamical spectrum;*
(ii) *The measure $\nu = \sum_{i\leq m} a_i \delta_{\Lambda_i}$ has pure point diffraction spectrum, for any choice of complex numbers $(a_i)_{i\leq m}$;*
(iii) *The measures δ_{Λ_i} have pure point spectrum, for $i \leq m$.*

The condition of $\mathbf{\Lambda}$ having "uniform cluster frequencies" is equivalent to the fact that its hull is uniquely ergodic, which we know is the case for many tilings constructed using supertile methods. It would be remiss not to mention [70], the companion work to [69]. It includes the result that for lattice substitution multiset systems[21], being a regular model set is equivalent to having pure point spectrum.

1.8.1 When the diffraction is not pure point

Recent work in [16] attempts to understand the dynamical spectrum when it is larger than the diffraction spectrum. An idea has been around for a while that factors of a system can give nuance to the diffraction spectrum. That is, "the missing parts of the dynamical spectrum could be reconstructed from the diffraction measures of suitable factors of the original system". In the uniquely ergodic case, the authors of [16] are able to show (see Corollary 9 for technical details) that (i) the diffraction measure of a factor is a spectral measure for the Koopman operator, and (ii) the set of diffraction measures of factors of a system are dense in the set of all spectral measures for the system.

[21] Not a particularly restrictive subclass according to Section 5.1 of [70].

In [59] it is shown that there exist substitutions which require infinitely many factors to reconstruct the pure point dynamical spectrum from the respective diffraction. There it is noted that it is not true that the maximal spectral measure of a subshift can be realized as the fundamental diffraction of a subshift factor.

As is true for the dynamical spectrum, one of the major areas of study is to determine the nature of the continuous part of the diffraction spectrum. In [9] it is shown that the continuous part of the spectrum of $a \to abbb, b \to a$ is singularly continuous with respect to Lebesgue measure. The general case $a \to ab^k, b \to a$ is considered in [14]. The analysis is based on a 'renormalization' process wherein the substitution structure of the self-similar tiling is used to find recursion relations for the autocorrelation measure. This method was also applied to the twisted silver mean in [10].

1.9 For further reading

A good primary source for fundamental results on tiling dynamical systems is B. Solomyak's "The dynamics of self-similar tilings" [100]. This paper lays out the basic definitions and takes an ergodic theoretic approach to the systems. A fundamental resource in elementary tiling theory is B. Grunbaum and G. C. Shephard's *Tilings and Patterns* [58], which catalogs nearly everything that is known about periodic tilings and more. It contains an enormous number of examples, and does include a few nonperiodic tilings such as the Penrose, Robinson, and Ammann tilings. Good general ergodic theory references for $\mathbb{Z}$-actions are K. Petersen's *Ergodic Theory* and P. Walters' *An Introduction to Ergodic Theory* [84, 104]. Fundamental symbolic dynamics references are D. Lind and B. Marcus' *An Introduction to Symbolic Dynamics and Coding* and Bruce Kitchens' *Symbolic Dynamcs* [74, 67]. Symbolic substitutions are surveyed up to 2002 in the collectively written *Substitutions in Dynamics, Arithmetics, and Combinatorics* [42]. A recent survey of S-adic expansions appears in V. Berthé and V. Delecroix's "Beyond substitutive dynamical systems: S-adic expansions" [24]. The definitive volume for the study of aperiodic order is M. Baake and U. Grimm's *Aperiodic Order* [11]. It takes a physical perspective and is full of examples of every sort, many analyzed fully.

There are a few other expositions of tilings and tiling spaces that are worth mentioning here. For a rigorous dynamical introduction to the theory, with multidimensional actions surveyed up to 2004 see E. A. Robinson, Jr.'s "Symbolic dynamics and tilings of $\mathbb{R}^d$" [93]. Radin's AMS Student Mathematical Library notes *Miles of Tiles* [90] introduces readers to the dynamics and ergodic theory with a strong physical motivation. At a university student level, it carries the additional interest of treating tilings with infinitely many tile rotations such as the pinwheel tiling. Substitutions on the graphs of tilings are considered in the author's "A primer on substitution tilings of Euclidean space" [46], which includes several examples of such combinatorial substitutions and their associated self-similar tilings. The topology of tiling spaces is the subject of L. Sadun's *Topology of Tiling Spaces* [95], which takes the reader through self-similar tiling constructions with and without rotations, shows tiling spaces are inverse limits, and does cohomology in the tiling context. There are many more topics we have not even mentioned, so the reader is encouraged to find a compelling topic to pursue.

References

1. *Tilings encyclopedia*, http://tilings.math.uni-bielefeld.de/, Accessed December 2018.
2. S. Akiyama, M. Barge, V. Berthé, J.-Y. Lee, and A. Siegel, *On the Pisot substitution conjecture*, Mathematics of aperiodic order, Progr. Math., vol. 309, Birkhäuser/Springer, Basel, 2015, pp. 33–72. MR 3381478
3. Shigeki Akiyama, *Strong coincidence and overlap coincidence*, Discrete Contin. Dyn. Syst. **36** (2016), no. 10, 5223–5230. MR 3543543
4. Shigeki Akiyama and Jeong-Yup Lee, *Algorithm for determining pure pointedness of self-affine tilings*, Adv. Math. **226** (2011), no. 4, 2855–2883. MR 2764877
5. ———, *Overlap coincidence to strong coincidence in substitution tiling dynamics*, European J. Combin. **39** (2014), 233–243. MR 3168528
6. Jean-Paul Allouche and Jeffrey Shallit, *Automatic sequences*, Cambridge University Press, Cambridge, 2003, Theory, applications, generalizations. MR 1997038
7. Pierre Arnoux, Maki Furukado, Edmund Harriss, and Shunji Ito, *Algebraic numbers, free group automorphisms and substitutions on the plane*, Trans. Amer. Math. Soc. **363** (2011), no. 9, 4651–4699. MR 2806687
8. Pierre Arnoux and Shunji Ito, *Pisot substitutions and Rauzy fractals*, Bull. Belg. Math. Soc. Simon Stevin **8** (2001), no. 2, 181–207, Journées Montoises d'Informatique Théorique (Marne-la-Vallée, 2000). MR 1838930
9. Michael Baake, Natalie Priebe Frank, Uwe Grimm, and E. Arthur Robinson Jr., *Spectral analysis of a family of binary inflation rules*, Studia Math. **Online First** (2018).
10. Michael Baake and Franz Gähler, *Pair correlations of aperiodic inflation rules via renormalisation: some interesting examples*, Topology Appl. **205** (2016), 4–27. MR 3493304
11. Michael Baake and Uwe Grimm, *Aperiodic order. Vol. 1*, Encyclopedia of Mathematics and its Applications, vol. 149, Cambridge University Press, Cambridge, 2013, A mathematical invitation, With a foreword by Roger Penrose. MR 3136260
12. ———, *Squirals and beyond: substitution tilings with singular continuous spectrum*, Ergodic Theory Dynam. Systems **34** (2014), no. 4, 1077–1102. MR 3227148
13. Michael Baake and Uwe Grimm, *Diffraction of a binary non-pisot inflation tiling*, Journal of Physics: Conference Series **809** (2017), no. 1, 012026.
14. Michael Baake, Uwe Grimm, and Neil Mañibo, *Spectral analysis of a family of binary inflation rules*, Lett. Math. Phys. **108** (2018), no. 8, 1783–1805. MR 3814725

15. Michael Baake and Daniel Lenz, *Spectral notions of aperiodic order*, Discrete Contin. Dyn. Syst. Ser. S **10** (2017), no. 2, 161–190. MR 3600642
16. Michael Baake, Daniel Lenz, and Aernout van Enter, *Dynamical versus diffraction spectrum for structures with finite local complexity*, Ergodic Theory Dynam. Systems **35** (2015), no. 7, 2017–2043. MR 3394105
17. Marcy Barge, *Pure discrete spectrum for a class of one-dimensional substitution tiling systems*, Discrete Contin. Dyn. Syst. **36** (2016), no. 3, 1159–1173. MR 3431249
18. Marcy Barge and Beverly Diamond, *Coincidence for substitutions of Pisot type*, Bull. Soc. Math. France **130** (2002), no. 4, 619–626. MR 1947456
19. Alan Bartlett, *Spectral theory of* $\mathbb{Z}^d$ *substitutions*, Ergodic Theory Dynam. Systems **38** (2018), no. 4, 1289–1341. MR 3789166
20. Nicolas Bédaride and Arnaud Hilion, *Geometric realizations of two-dimensional substitutive tilings*, Q. J. Math. **64** (2013), no. 4, 955–979. MR 3151599
21. Jean Bellissard, Antoine Julien, and Jean Savinien, *Tiling groupoids and Bratteli diagrams*, Ann. Henri Poincaré **11** (2010), no. 1-2, 69–99. MR 2658985
22. Robert Berger, *The undecidability of the domino problem*, Mem. Amer. Math. Soc. No. **66** (1966), 72. MR 0216954
23. Artemi Berlinkov and Boris Solomyak, *singular substitutions of constant length*, ArXiv:1705.00899v2, 2017.
24. Valérie Berthé and Vincent Delecroix, *Beyond substitutive dynamical systems: S-adic expansions*, Numeration and substitution 2012, RIMS Kôkyûroku Bessatsu, B46, Res. Inst. Math. Sci. (RIMS), Kyoto, 2014, pp. 81–123. MR 3330561
25. S. Bezuglyi, J. Kwiatkowski, and K. Medynets, *Aperiodic substitution systems and their Bratteli diagrams*, Ergodic Theory Dynam. Systems **29** (2009), no. 1, 37–72. MR 2470626
26. S. Bezuglyi, J. Kwiatkowski, K. Medynets, and B. Solomyak, *Invariant measures on stationary Bratteli diagrams*, Ergodic Theory Dynam. Systems **30** (2010), no. 4, 973–1007. MR 2669408
27. ———, *Finite rank Bratteli diagrams: structure of invariant measures*, Trans. Amer. Math. Soc. **365** (2013), no. 5, 2637–2679. MR 3020111
28. E. Bombieri and J. E. Taylor, *Which distributions of matter diffract? An initial investigation*, J. Physique **47** (1986), no. 7, Suppl. Colloq. C3, C3–19–C3–28, International workshop on aperiodic crystals (Les Houches, 1986). MR 866320
29. R. V. Chacon, *A geometric construction of measure preserving transformations*, Proc. Fifth Berkeley Sympos. Math. Statist. and Probability (Berkeley, Calif., 1965/66), Vol. II: Contributions to Probability Theory, Part 2, Univ. California Press, Berkeley, Calif., 1967, pp. 335–360. MR 0212158
30. Lax Chan and Uwe Grimm, *Spectrum of a Rudin-Shapiro-like sequence*, Adv. in Appl. Math. **87** (2017), 16–23. MR 3629260
31. Alex Clark and Lorenzo Sadun, *When size matters: subshifts and their related tiling spaces*, Ergodic Theory Dynam. Systems **23** (2003), no. 4, 1043–1057. MR 1997967
32. Ethan M. Coven and Michael S. Keane, *The structure of substitution minimal sets*, Trans. Amer. Math. Soc. **162** (1971), 89–102. MR 0284995
33. F. M. Dekking, *The spectrum of dynamical systems arising from substitutions of constant length*, Z. Wahrscheinlichkeitstheorie und Verw. Gebiete **41** (1977/78), no. 3, 221–239. MR 0461470
34. F. M. Dekking and M. Keane, *Mixing properties of substitutions*, Z. Wahrscheinlichkeitstheorie und Verw. Gebiete **42** (1978), no. 1, 23–33. MR 0466485
35. Fabien Durand, *Linearly recurrent subshifts have a finite number of non-periodic subshift factors*, Ergodic Theory Dynam. Systems **20** (2000), no. 4, 1061–1078. MR 1779393

36. ______, *Corrigendum and addendum to: "Linearly recurrent subshifts have a finite number of non-periodic subshift factors" [Ergodic Theory Dynam. Systems* **20** *(2000), no. 4, 1061–1078; MR1779393 (2001m:37022)]*, Ergodic Theory Dynam. Systems **23** (2003), no. 2, 663–669. MR 1972245
37. Steven Dworkin, *Spectral theory and x-ray diffraction*, J. Math. Phys. **34** (1993), no. 7, 2965–2967. MR 1224190
38. Sébastien Ferenczi, *Rank and symbolic complexity*, Ergodic Theory Dynam. Systems **16** (1996), no. 4, 663–682. MR 1406427
39. Thomas Fernique, *Local rule substitutions and stepped surfaces*, Theoret. Comput. Sci. **380** (2007), no. 3, 317–329. MR 2331001
40. Thomas Fernique and Nicolas Ollinger, *Combinatorial substitutions and sofic tilings*, Journées Automates Cellulaires 2010 (Turku).
41. Albert M. Fisher, *Nonstationary mixing and the unique ergodicity of adic transformations*, Stoch. Dyn. **9** (2009), no. 3, 335–391. MR 2566907
42. N. Pytheas Fogg, *Substitutions in dynamics, arithmetics and combinatorics*, Lecture Notes in Mathematics, vol. 1794, Springer-Verlag, Berlin, 2002, Edited by V. Berthé, S. Ferenczi, C. Mauduit and A. Siegel. MR 1970385
43. Natalie Priebe Frank, *Detecting combinatorial hierarchy in tilings using derived Voronoï tesselations*, Discrete Comput. Geom. **29** (2003), no. 3, 459–476. MR 1961011
44. ______, *Substitution sequences in* $\mathbb{Z}^d$ *with a non-simple Lebesgue component in the spectrum*, Ergodic Theory Dynam. Systems **23** (2003), no. 2, 519–532. MR 1972236
45. ______, *Multidimensional constant-length substitution sequences*, Topology Appl. **152** (2005), no. 1-2, 44–69. MR 2160805
46. ______, *A primer of substitution tilings of the Euclidean plane*, Expo. Math. **26** (2008), no. 4, 295–326. MR 2462439
47. ______, *Tilings with infinite local complexity*, Mathematics of aperiodic order, Progr. Math., vol. 309, Birkhäuser/Springer, Basel, 2015, pp. 223–257. MR 3381483
48. Natalie Priebe Frank and E. Arthur Robinson, Jr., *Generalized* β*-expansions, substitution tilings, and local finiteness*, Trans. Amer. Math. Soc. **360** (2008), no. 3, 1163–1177. MR 2357692
49. Natalie Priebe Frank and Lorenzo Sadun, *Topology of some tiling spaces without finite local complexity*, Discrete Contin. Dyn. Syst. **23** (2009), no. 3, 847–865. MR 2461829
50. ______, *Fusion: a general framework for hierarchical tilings of* $\mathbb{R}^d$, Geom. Dedicata **171** (2014), 149–186. MR 3226791
51. Dirk Frettlöh, *More inflation tilings*, Aperiodic order. Vol. 2, Encyclopedia Math. Appl., vol. 166, Cambridge Univ. Press, Cambridge, 2017, pp. 1–37. MR 3791847
52. F. Gähler and R. Klitzing, *The diffraction pattern of self-similar tilings*, The mathematics of long-range aperiodic order (Waterloo, ON, 1995), NATO Adv. Sci. Inst. Ser. C Math. Phys. Sci., vol. 489, Kluwer Acad. Publ., Dordrecht, 1997, pp. 141–174. MR 1460023
53. Franz Gähler, Eugene E. Kwan, and Gregory R. Maloney, *A computer search for planar substitution tilings with n-fold rotational symmetry*, Discrete Comput. Geom. **53** (2015), no. 2, 445–465. MR 3316232
54. Martin Gardner, *Mathematical games*, Scientific American **236** (1977), no. 1, 110–121.
55. C. Godrèche and F. Lançon, *A simple example of a non-Pisot tiling with five-fold symmetry*, J. Physique I **2** (1992), no. 2, 207–220. MR 1185612
56. G. R. Goodson, *A survey of recent results in the spectral theory of ergodic dynamical systems*, J. Dynam. Control Systems **5** (1999), no. 2, 173–226. MR 1693318

57. W. H. Gottschalk, *Substitution minimal sets*, Trans. Amer. Math. Soc. **109** (1963), 467–491. MR 0190915
58. Branko Grünbaum and G. C. Shephard, *Tilings and patterns*, W. H. Freeman and Company, New York, 1987. MR 857454
59. Joseph L. Herning, *Spectrum and Factors of Substitution Dynamical Systems*, ProQuest LLC, Ann Arbor, MI, 2013, Thesis (Ph.D.)–The George Washington University. MR 3167382
60. A. Hof, *On diffraction by aperiodic structures*, Comm. Math. Phys. **169** (1995), no. 1, 25–43. MR 1328260
61. Michael Hollander and Boris Solomyak, *Two-symbol Pisot substitutions have pure discrete spectrum*, Ergodic Theory Dynam. Systems **23** (2003), no. 2, 533–540. MR 1972237
62. B. Host, *Valeurs propres des systèmes dynamiques définis par des substitutions de longueur variable*, Ergodic Theory Dynam. Systems **6** (1986), no. 4, 529–540. MR 873430
63. Emmanuel Jeandel and Pascal Vanier, *The undecidability of the domino problem*, This volume, Lecture Notes in Mathematics, Springer, 2019.
64. Teturo Kamae, *Spectrum of a substitution minimal set*, J. Math. Soc. Japan **22** (1970), 567–578. MR 0286092
65. Johannes Kellendonk and Lorenzo Sadun, *Meyer sets, topological eigenvalues, and Cantor fiber bundles*, J. Lond. Math. Soc. (2) **89** (2014), no. 1, 114–130. MR 3174736
66. Richard Kenyon and Boris Solomyak, *On the characterization of expansion maps for self-affine tilings*, Discrete Comput. Geom. **43** (2010), no. 3, 577–593. MR 2587839
67. Bruce P. Kitchens, *Symbolic dynamics*, Universitext, Springer-Verlag, Berlin, 1998, One-sided, two-sided and countable state Markov shifts. MR 1484730
68. Jarosł aw Kwapisz, *Inflations of self-affine tilings are integral algebraic Perron*, Invent. Math. **205** (2016), no. 1, 173–220. MR 3514961
69. J.-Y. Lee, R. V. Moody, and B. Solomyak, *Pure point dynamical and diffraction spectra*, Ann. Henri Poincaré **3** (2002), no. 5, 1003–1018. MR 1937612
70. Jeong-Yup Lee, Robert V. Moody, and Boris Solomyak, *Consequences of pure point diffraction spectra for multiset substitution systems*, Discrete Comput. Geom. **29** (2003), no. 4, 525–560. MR 1976605
71. Mariusz Lemańczyk, *Spectral theory of dynamical systems*, Mathematics of complexity and dynamical systems. Vols. 1–3, Springer, New York, 2012, pp. 1618–1638. MR 3220776
72. Daniel Lenz, *Continuity of eigenfunctions of uniquely ergodic dynamical systems and intensity of Bragg peaks*, Comm. Math. Phys. **287** (2009), no. 1, 225–258. MR 2480747
73. Sven Lidin, *The discovery of quasicrystals: Scientific background on the nobel prize in chemistry 2011*, The Royal Swedish Academy of Sciences, 2011.
74. Douglas Lind and Brian Marcus, *An introduction to symbolic dynamics and coding*, Cambridge University Press, Cambridge, 1995. MR 1369092
75. A. N. Livshits, *On the spectra of adic transformations of Markov compact sets*, Uspekhi Mat. Nauk **42** (1987), no. 3(255), 189–190. MR 896889
76. M. Lothaire, *Algebraic combinatorics on words*, Encyclopedia of Mathematics and its Applications, vol. 90, Cambridge University Press, Cambridge, 2002, A collective work by Jean Berstel, Dominique Perrin, Patrice Seebold, Julien Cassaigne, Aldo De Luca, Steffano Varricchio, Alain Lascoux, Bernard Leclerc, Jean-Yves Thibon, Veronique Bruyere, Christiane Frougny, Filippo Mignosi, Antonio Restivo, Christophe Reutenauer, Dominique Foata, Guo-Niu Han, Jacques Desarmenien, Volker Diekert, Tero Harju, Juhani Karhumaki and Wojciech Plandowski, With a preface by Berstel and Perrin. MR 1905123

77. Alan L. Mackay, *Crystallography and the Penrose pattern*, Phys. A **114** (1982), no. 1-3, 609–613. MR 678468
78. Gregory R. Maloney and Dan Rust, *Beyond primitivity for one-dimensional substitution subshifts and tiling spaces*, Ergodic Theory Dynam. Systems **38** (2018), no. 3, 1086–1117. MR 3784255
79. Brian F. Martensen, *Generalized balanced pair algorithm*, Topology Proc. **28** (2004), no. 1, 163–178, Spring Topology and Dynamical Systems Conference. MR 2105455
80. John C. Martin, *Substitution minimal flows*, Amer. J. Math. **93** (1971), 503–526. MR 0300261
81. Brigitte Mossé, *Puissances de mots et reconnaissabilité des points fixes d'une substitution*, Theoret. Comput. Sci. **99** (1992), no. 2, 327–334. MR 1168468
82. R. Penrose, *The rôle of aesthetics in pure and applied mathematical research*, Bull. of the Institute of Math. and its Applications **10** (1974), 266–271.
83. ———, *Pentaplexity: a class of nonperiodic tilings of the plane*, Math. Intelligencer **2** (1979/80), no. 1, 32–37. MR 558670
84. Karl Petersen, *Ergodic theory*, Cambridge Studies in Advanced Mathematics, vol. 2, Cambridge University Press, Cambridge, 1989, Corrected reprint of the 1983 original. MR 1073173
85. ———, *Factor maps between tiling dynamical systems*, Forum Math. **11** (1999), no. 4, 503–512. MR 1699171
86. J. Peyrière, *Frequency of patterns in certain graphs and in Penrose tilings*, J. Physique **47** (1986), no. 7, Suppl. Colloq. C3, C3–41–C3–62, International workshop on aperiodic crystals (Les Houches, 1986). MR 866322
87. Brenda Praggastis, *Numeration systems and Markov partitions from self-similar tilings*, Trans. Amer. Math. Soc. **351** (1999), no. 8, 3315–3349. MR 1615950
88. N. Priebe and B. Solomyak, *Characterization of planar pseudo-self-similar tilings*, Discrete Comput. Geom. **26** (2001), no. 3, 289–306. MR 1854103
89. Martine Queffélec, *Substitution dynamical systems—spectral analysis*, Lecture Notes in Mathematics, vol. 1294, Springer-Verlag, Berlin, 1987. MR 924156
90. Charles Radin, *Miles of tiles*, Student Mathematical Library, vol. 1, American Mathematical Society, Providence, RI, 1999. MR 1707270
91. Charles Radin and Lorenzo Sadun, *Isomorphism of hierarchical structures*, Ergodic Theory Dynam. Systems **21** (2001), no. 4, 1239–1248. MR 1849608
92. Charles Radin and Mayhew Wolff, *Space tilings and local isomorphism*, Geom. Dedicata **42** (1992), no. 3, 355–360. MR 1164542
93. E. Arthur Robinson, Jr., *Symbolic dynamics and tilings of* $\mathbb{R}^d$, Symbolic dynamics and its applications, Proc. Sympos. Appl. Math., vol. 60, Amer. Math. Soc., Providence, RI, 2004, pp. 81–119. MR 2078847
94. Raphael M. Robinson, *Undecidability and nonperiodicity for tilings of the plane*, Invent. Math. **12** (1971), 177–209. MR 0297572
95. Lorenzo Sadun, *Topology of tiling spaces*, University Lecture Series, vol. 46, American Mathematical Society, Providence, RI, 2008. MR 2446623
96. Marjorie Senechal, *Quasicrystals and geometry*, Cambridge University Press, Cambridge, 1995. MR 1340198
97. D. Shechtman, I. Blech, D. Gratias, and J. W. Cahn, *Metallic phase with long-range orientational order and no translational symmetry*, Phys. Rev. Lett. **53** (1984), 1951–1953.
98. B. Solomyak, *Nonperiodicity implies unique composition for self-similar translationally finite tilings*, Discrete Comput. Geom. **20** (1998), no. 2, 265–279. MR 1637896
99. ———, *Pseudo-self-affine tilings in* $\mathbb{R}^d$, Zap. Nauchn. Sem. S.-Peterburg. Otdel. Mat. Inst. Steklov. (POMI) **326** (2005), no. Teor. Predst. Din. Sist. Komb. i Algoritm. Metody. 13, 198–213, 282–283. MR 2183221

100. Boris Solomyak, *Dynamics of self-similar tilings*, Ergodic Theory Dynam. Systems **17** (1997), no. 3, 695–738. MR 1452190
101. ———, *Lecture on delone sets and tilings*, This volume, Lecture Notes in Mathematics, Springer, 2019.
102. William Thurston, *Groups, tilings, and finite state automata*, AMS Colloquium Lecture Notes, American Mathematical Society, 1989.
103. Jörg Thuswaldner, *S-adic sequences. a bridge between dynamics, arithmetic, and geometry*, This volume, Lecture Notes in Mathematics, Springer, 2019.
104. Peter Walters, *An introduction to ergodic theory*, Graduate Texts in Mathematics, vol. 79, Springer-Verlag, New York-Berlin, 1982. MR 648108
105. Hao Wang, *Proving theorems by pattern recognition, ii*, Bell Systems Tech. J. (1961), 1–41.