

Fractal dual substitution tilings

Natalie Priebe Frank, Samuel B. G. Webster, and Michael F. Whittaker¹

Abstract. Starting with a substitution tiling, we demonstrate a method for constructing infinitely many new substitution tilings. Each of these new tilings is derived from a graph iterated function system and the tiles have fractal boundary. We show that each of the new tilings is mutually locally derivable to the original tiling. Thus, at the tiling space level, the new substitution rules are expressing geometric and combinatorial, rather than topological, features of the original. Our method is easy to apply to particular substitution tilings, permits experimentation, and can be used to construct border-forcing substitution rules. For a large class of examples we show that the combinatorial dual tiling has a realization as a substitution tiling. Since the boundaries of our new tilings are fractal we are led to compute their fractal dimension. As an application of our techniques we show how to compute the Čech cohomology of a (not necessarily border-forcing) tiling using a graph iterated function system of a fractal tiling.

Mathematics Subject Classification (2010). Primary: 37D40; Secondary: 05B45, 52C20.

Keywords. Fractals, graph iterated function systems, nonperiodic tilings, substitution tilings.

¹ This research was partially supported by the following: the Australian Research Council (DP130100490), the Institute of Mathematics and its Applications at University of Wollongong, the German Research Council (DFG) within the CRC 701, and the Babette Rogol '61 Memorial Fund of Vassar College.

Contents

1	Introduction	267
2	Tiling definitions	272
3	The combinatorics of tilings	275
4	Alternate views of tiling substitutions	280
5	Constructing fractal tilings from self-similar tilings	283
6	The existence of fractal dual and quasi-dual tilings	287
7	Edge contraction map and Hausdorff dimension of $\partial T^{(\infty)}$	297
8	Čech cohomology of a tiling space	301
A	Examples	309
	References	315

List of figures

1	The two-dimensional Thue–Morse substitution rule	270
2	A recurrent pair on the 2-dimensional Thue–Morse substitution leading to a fractal graph	270
3	A patch of the 2-dimensional Thue–Morse tiling with a fractal dual tiling overlaid	271
4	A portion of a self-similar 2DTM tiling, with the origin at the center	275
5	A patch from the third substitution of the chair tiling (see Example 4.3). The natural vertices marked with a \bullet , and the vertices in the T -consistent set generated by the natural vertices are marked with an \times	277
6	The 4 prototiles of the chair tiling and the substitution of prototile p_1	283
7	The prototiles of the 2DTM tiling with an oriented dual graph G inscribed	306
8	The 8 prototiles of the fractal realization of the 2DTM tiling from (G, S) . Each oriented (single) edge is labelled by a pair of edges from G as pictured in the bounding squares (which are only present to label the edges)	306
9	The substitutions of the 8 fractal prototiles (not to scale)	306
10	A recurrent pair and resulting fractal for the dual graph of the Ammann–Beenker tiling	310
11	A patch of a fractal dual Ammann–Beenker tiling	310
12	The substitution rules of this fractal realization of the Ammann–Beenker octagonal tiling	310
13	A recurrent pair for the Penrose triangles; note that each triangle come in two reflections and so do their geometric graphs	311
14	A patch of the fractal pentaplexity tiling	311
15	A fractal version of the pentaplexity substitution	311
16	A recurrent pair and limiting fractal for the chairs as octagons	313
17	A patch of the octagonal chair tiling with a fractal realization overlaid	313
18	The three chair tiles and their fractal dual graphs	314
19	Substitution of the three chairs (not to scale)	314
20	A patch of the chair tiling with its fractal dual overlaid	314

1. Introduction

A tiling of the plane is a covering of \mathbb{R}^2 by closed subsets, called tiles, such that the interiors of two distinct tiles are disjoint. One method of producing tilings is by substitution: a rule which expands each tile by a fixed amount and then breaks each expanded tile into smaller pieces, each of which is an isometric copy of an original tile. The most famous substitution tiling is the Penrose tiling [25]; many other substitution tilings have been constructed and can be found online in the Tilings Encyclopedia [31].

1.1. Overview of the field. The study of substitution tilings is motivated by several disparate fields of mathematics and science. One thing they all have in common is that tilings are always constructed using a finite set of tiles called *prototiles* as their building blocks. The study of *aperiodic* prototile sets – finite prototile sets that can only form nonperiodic tilings – began with logician Hao Wang [32]. He tied them to the decidability of the Domino Problem, which asks if there is an algorithm that can determine whether any given set of (square) prototiles, with specified edge matching rules, can tile the plane. He conjectured that the problem was decidable, but his conjecture depended upon non-existence of aperiodic prototile sets. Wang’s student Robert Berger [11] proved the conjecture false by producing an aperiodic prototile set with 20,426 prototiles. It was the search for small aperiodic prototile sets that led Penrose to discover his famous tiles.

Research on aperiodic tilings was purely academic until Dan Shechtman’s Nobel Prize-winning discovery of quasicrystals [28]. In a sense, Shechtman’s discovery was predicted by Alan L. Mackay [22] when he computed the diffraction of a Penrose tiling. Its ‘impossible’ symmetry was found again in Shechtman’s quasicrystal and it became clear that Penrose tilings could be used as a mathematical model. Diffraction patterns of quasicrystals are now modeled by the spectral theory of tiling spaces, in particular the dynamical spectrum of the translation operator (see [3] for a beautiful exposition of the most recent advances in this theory). The dynamical spectrum of tilings is best understood in the case of substitution tilings.

An understanding of tiling spaces at this physical level led Bellissard to his gap-labelling conjecture [8], which connects the spectral gaps of the Schrödinger operator associated with a tiling space to the K-theory and cohomology of that space. Bellissard’s conjecture was proven independently in [8, 10, 18]. By the same token, tiling spaces provide interesting examples of C^* -algebras, and the geometry imposed by tilings makes them excellent candidates for study using Connes’ noncommutative geometry program [13].

1.2. Results and applications. The main goal of this paper is to establish a novel method for producing, from some given substitution tiling, infinitely many new substitution tilings. The method produces finite sets of prototiles which have fractal boundaries, so we call them *fractal realizations* of the original tiling. The construction depends on a synthesis of combinatorics and geometry in a way that has not previously been capitalized upon in the study of substitution tilings. Moreover, it allows for hands-on experimentation with an infinitude of choices, and for now it is not completely clear the significance of making one choice over another. In [15], we gave a somewhat ad-hoc method for making a fractal realization of the Pinwheel tiling. In this paper we show that a similar construction works for every primitive substitution tiling where tiles meet singly edge-to-edge and remove the ad-hoc nature of the construction in [15].

One advantage to our method is that it provides access to some of the information described above. For instance, our new substitutions are easily made to *force the border*, an essential property for cohomology computations. When a tiling does not force the border, our method requires less information than the original approach in [1] and uses the same amount of information, but is geometrically more elegant, than the most efficient known method [5].

The real power behind our approach comes from the geometric information encoded in the new substitution. One important application of our fractal dual tilings is to construct spectral triples on Kellendonk's C^* -algebra of a tiling [23]. Spectral triples are the fundamental object in Connes' noncommutative geometry program [13], and the existence of a spectral triple on a C^* -algebra is the noncommutative analogue of the existence of a Riemannian metric on a manifold. The spectral triples constructed in [23] are significant not only because of their existence, but because they provide the first class of spectral triples on Kellendonk's C^* -algebra [19] that respect the hierarchy of the substitution system. Moreover, each distinct fractal realization of a tiling provides a distinct noncommutative Riemannian metric on Kellendonk's algebra. We believe that this type of geometric information is an important aspect of tiling theory. After all, the space of tilings should not just be viewed as a topological object but also as a geometric object whose geometry must arise from the rigid structure of the prototiles and the way they fit together geometrically.

There are further interesting questions invoked by our fractal substitution tilings. For example, we give a formula for the fractal dimension of the boundary of the prototiles for any of our fractal tilings. But what information does the set of all possible fractal dimensions for a given tiling substitution carry? The answer to this question seems to rely on a deep connection between the combinatorial graph-

theoretic properties of the tiling and the geometric property of fractal dimension. An example of how fractal dimension is used appears in the spectral triples constructed in [23] where the dimension gives a scaling factor for the distance function between m - and n -supertiles. Further questions arise from a combinatorial standpoint; for example, which substitution tilings admit *combinatorial dual tilings*¹ that are also substitution tilings? We provide sufficient conditions but have not yet been able to provide a complete characterization.

Our method of obtaining fractal substitution tilings is different than those that have appeared in the literature. For example, fractal tilings arise in the seminal work of Kenyon [20] characterizing the possible expansion factors for substitution tilings. Since then fractal tilings have appeared many times, for instance in [2, 4, 7, 26, 15]. Our construction distinguishes itself from these in a few ways. The first is that we begin with a known substitution rule and construct an infinite family of substitutions whose tiling spaces are mutually locally derivable. In particular, we are able to construct substitution tilings that are mutually locally derivable to their original but with specified geometric properties that lend themselves more easily to computation. Our construction differs, also, because it is not simply a redrawing of existing tiles but rather a recombination that creates new tiles. We can understand, and indeed have some control over, the combinatorics of the new tiles within their tilings.

1.3. Methodology – an example. The main tool in our construction is a *recurrent pair* (G, S) , which is a combinatorial and geometric structure that is compatible with the original substitution. The best way to illustrate how to construct a recurrent pair is with a simple example. We begin with a two-dimensional version of the Thue–Morse substitution rule, which we call the “2DTM” substitution. It has two prototiles, labelled α and β , that are both unit squares. The substitution expands each prototile by a factor of four and replaces it with the patches shown in Figure 1.

To construct a recurrent pair (G, S) , we begin by embedding a planar graph G into each of the prototiles α and β , as shown on the left of Figure 2. Note that we have chosen G so that it has one vertex in the interior of each tile and connects to the interior of each edge. If G were placed in every tile in an infinite tiling of these tiles, it would construct a new tiling that is combinatorially dual to the original. In general it is natural to start with such a dual graph for several reasons that will be outlined in this paper.

¹ Two tilings are combinatorially dual if there is a one-to-one correspondence between their edge sets, between the tiles of one and the vertices of the other, and vice versa.




Figure 1. The two-dimensional Thue–Morse substitution rule.

Next we construct a new graph S on the prototiles by substituting the prototiles, this time without expanding, embedding the initial graph into each sub-tile, then selecting a subgraph S that is combinatorially equivalent to the initial graph if we ignore all vertices of degree two.² In Figure 2, the embedding stage is labelled $\mathcal{R}(G)$, and the selection stage is labelled S . We have made a choice of S that breaks the inherent symmetry of the 2DTM tilings, but this was not necessary.

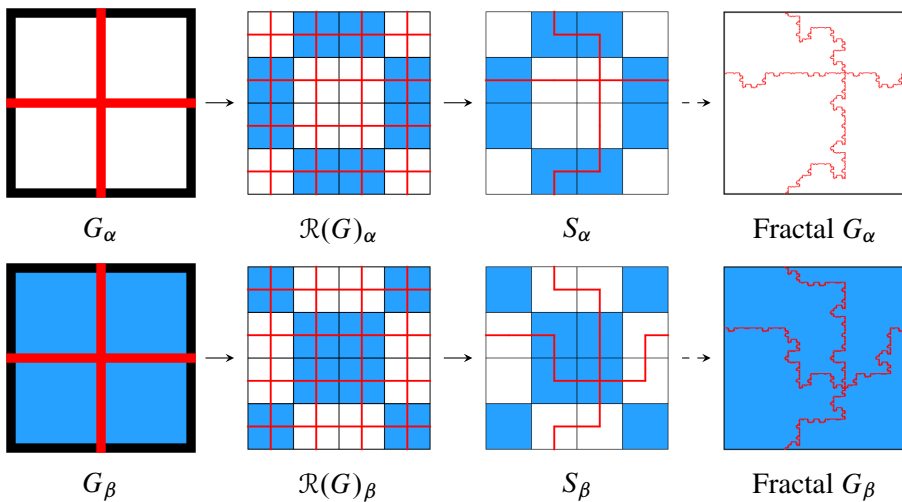


Figure 2. A recurrent pair on the 2-dimensional Thue–Morse substitution leading to a fractal graph.

² It is important to note that selecting such a graph is not necessarily possible in general and depends on the combinatorics of the tiling, the substitution, and G .

We call (G, S) a recurrent pair, and S can be thought of as of graph substitution of G . In fact, a recurrent pair forms a graph iterated function system (GIFS) and has an attractor (which is a fractal). Under the right conditions the attractor is a graph that is combinatorially equivalent to both G and S . We show the attractor for this example in the far right of Figure 2.

Figure 3 shows what happens when the attractor is placed into every tile of a 2DTM tiling. Everything works perfectly in this example and we obtain a combinatorially dual tiling that is itself a substitution tiling. The fact that it is a substitution tiling follows from the fact that the attractor of the recurrent pair (G, S) is invariant under the substitution rule of T .

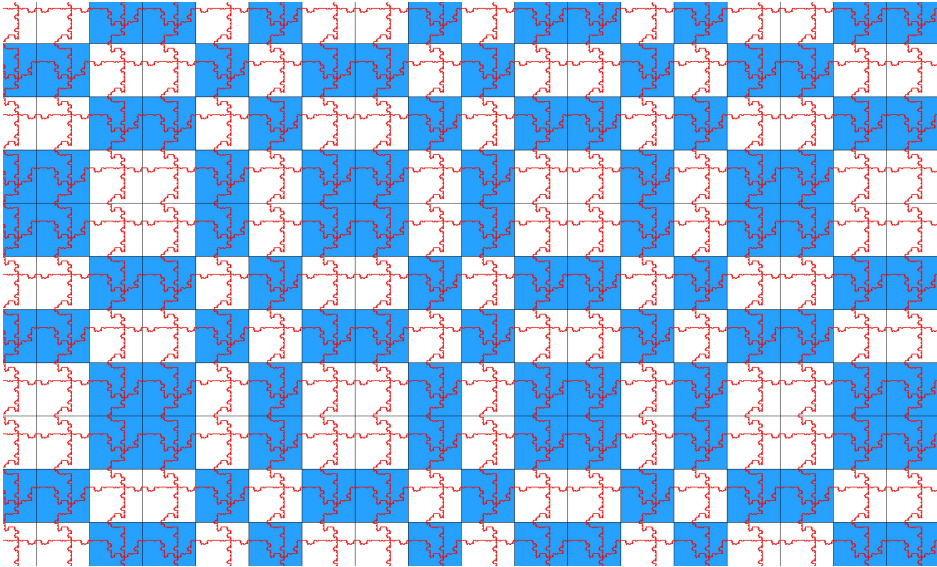


Figure 3. A patch of the 2-dimensional Thue–Morse tiling with a fractal dual tiling overlaid.

1.4. Organization of the paper. Given a substitution tiling, our main result is to construct new substitution tilings using a recurrent pair (G, S) . There is a natural map from the edges of G to the edges of the attractor associated with (G, S) . If this natural map is injective, then Theorem 5.7 says that the resulting tiling is mutually locally derivable to the original substitution tiling. In Theorem 6.2 we provide sufficient *injectivity conditions* on the recurrent pair for the map to be injective. Theorem 6.6 implies that every primitive substitution tiling whose tiles meet singly edge-to-edge has an infinite number of distinct recurrent pairs satisfying the injectivity conditions and whose fractal realization

is border-forcing. We conjecture that every substitution tiling of the plane with finite local complexity has a recurrent pair with the necessary injectivity conditions. We also believe that our techniques should extend to tilings with infinite rotational symmetry, such as the Pinwheel tiling, but have not addressed that in this paper.

We have organized the paper as follows. In Section 2 we introduce substitution tilings along with the definitions we require in the paper. Section 3 introduces the combinatorics of tilings. We give an alternative description of tiling substitutions in Section 4 using digit sets, and describe how these digit sets give rise to a contraction map on prototiles. Section 5 introduces the notion of a recurrent pair on a substitution tiling T and shows how a recurrent pair produces a fractal realization of T . In Section 6 we give injectivity conditions on a given recurrent pair and we prove they are what is needed to guarantee existence of a border-forcing fractal realization. We then show that every primitive substitution tiling with tiles meeting singly edge-to-edge has such a recurrent pair. In Section 7 we compute the fractal dimension of the new prototile boundaries. Section 8 shows how a recurrent pair can be used to compute the Čech cohomology of the original tiling space. Finally, in Appendix A we give several examples of fractal realizations of some famous substitution tilings.

Acknowledgements. We are extremely grateful to Michael Baake and Franz Gähler for helping to rectify an error in the cohomology computations from an early version of the paper. We are also indebted to Lorenzo Sadun for several helpful conversations and ideas and to Michael Mampusti for the Mathematica code used to produce some of the figures. Finally, we are indebted to the anonymous reviewer for several extremely helpful comments and suggestions.

2. Tiling definitions

In this section we introduce the basics of tilings. We provide the assumptions that will be used in the paper and introduce substitution rules. We establish the relationship between tiling substitution rules and self-similar tilings so that we may use these two notions interchangeably for the remainder of the paper.

The tilings in this paper are built out of a finite number of labelled shapes called prototiles. A *prototile* p consists of a closed subset of \mathbb{R}^2 that is homeomorphic to a topological disk, denoted $\text{supp}(p)$, and a label $\ell(p)$. The purpose of the labels is to distinguish between prototiles that have the same shape, and a common visualization is by color (for instance white vs. black unit squares). A *prototile set* \mathcal{P} is a finite set of prototiles.

A *tile* t is any translate of a prototile: for $p \in \mathcal{P}$ and $x \in \mathbb{R}^2$ we use the notation $t = p + x$ to mean the topological disk $\text{supp}(p) + x$ with label $\ell(t) = \ell(p)$. Two tiles that are translates of the same prototile are said to be *equivalent*; we note that equivalent tiles have congruent supports and carry the same label.

Definition 2.1. Given a set of prototiles \mathcal{P} , a *tiling* is a countable collection of tiles $\{t_1, t_2, \dots\}$ each of which is a translate of a prototile, and such that

- (1) $\bigcup_{i=1}^{\infty} \text{supp}(t_i) = \mathbb{R}^2$ and
- (2) $\text{Int}(\text{supp}(t_i)) \cap \text{Int}(\text{supp}(t_j)) = \emptyset$ for $i \neq j$.

The first condition implies that the tiles cover the plane and the second that they intersect on the boundaries only. These are sometimes referred to as the covering and packing conditions.

A *patch* is a finite set of tiles whose supports cover a connected set that intersect at most on their boundaries. Connected finite subsets of tilings are patches, and tilings are sometimes thought of as infinite patches. We denote the set of all patches from a prototile set \mathcal{P} by \mathcal{P}^* .

Given a patch or tiling Q and $x \in \mathbb{R}^2$, the set $Q + x := \{t + x \mid t \in Q\}$ is also a patch or tiling. Like tiles, patches and tilings are called *equivalent* if they are translations of the same patch or tiling.

Definition 2.2. A tiling T has *finite local complexity (FLC)* if the set of all two-tile patches appearing in T is finite up to equivalence. A tiling is *nonperiodic* if $T + x = T$ implies $x = 0$.

If Q is a patch of tiles or a tiling and S is a subset of \mathbb{R}^2 , we define the patch

$$[S]^Q := \{q \in Q \mid \text{supp}(q) \cap S \neq \emptyset\}.$$

We note that if Q is a tiling, the support of $[S]^Q$ contains S . Sometimes we may abuse notation and put a tile or patch in place of S , in which case it is understood to mean the Q -patch intersecting its support.

A very important form of equivalence between tilings with (potentially) different prototile sets is mutual local derivability.

Definition 2.3. We say a tiling T' is *locally derivable* from a tiling T if there is an $R > 0$ such that if $[B_R(x)]^T = [B_R(y)]^T + z$ then $[x]^{T'} = [y]^{T'} + z$. If both T and T' are locally derivable from each other then we say they are *mutually locally derivable (MLD)*.

Remark on notation and assumptions. Throughout this paper, unless otherwise noted, reference to a tiling T implies the presence of a finite prototile set \mathcal{P} that will only be mentioned explicitly if there is danger of confusion. Tilings will always be assumed to have finite local complexity.

2.1. Tiling substitutions and self-similar tilings

Definition 2.4. A function $\omega : \mathcal{P} \rightarrow \mathcal{P}^*$ is called a *tiling substitution* if there exists $\lambda > 1$ such that for every $p \in \mathcal{P}$,

$$\lambda \operatorname{supp}(p) = \operatorname{supp}(\omega(p)).$$

In this case λ is called the *expansion factor* of the substitution.

It is natural to extend the substitution ω to tiles, patches, and tilings. The substitution of a tile $t = p + x$, for $p \in \mathcal{P}$ and $x \in \mathbb{R}^2$, is the patch $\omega(t) := \omega(p) + \lambda x$. The substitution of a patch or tiling is the substitution applied to each of its tiles. Substitution rules naturally give rise to tilings by constructing a self-similar tiling as described in [27, p. 13].

An alternative approach to substitution tilings is to consider a self-similar tiling, and then find the substitution ω it determines. In order to define a self-similar tiling we require some notation. The boundary of a tiling T , denoted ∂T , is the subset of \mathbb{R}^2 given by the boundaries of all the supports of tiles in T , i.e. $\partial T := \bigcup_{t \in T} \partial(\operatorname{supp}(t))$.

Definition 2.5. A tiling T is *self-similar* if

- (1) there exists some $\lambda > 1$ such that $\lambda \partial T \subset \partial T$; and
- (2) if $t_1, t_2 \in T$ are translationally equivalent, then the patches enclosed by $\lambda \operatorname{supp}(t_1)$ and $\lambda \operatorname{supp}(t_2)$ are translationally equivalent.

The central patch of a self-similar tiling for the 2DTM substitution is shown in Figure 4 with the origin of \mathbb{R}^2 marked at the center. It is routine to check that the same patch sits at the center of a substituted version of this patch, and so on. Thus, we can extend this patch to a tiling by substituting an infinite number of times. Since the (ever expanding) central patch is always invariant under subsequent substitutions, we obtain a self-similar tiling of the plane. We call this particular tiling a self-similar version of the 2DTM tiling.




Figure 4. A portion of a self-similar 2DTM tiling, with the origin at the center.

To extract ω from a self-similar tiling T , for each prototile p find any equivalent tile $t = p + x \in T$, then define $\omega(p)$ to be the patch with support $\lambda \text{supp}(t)$ translated by $-\lambda x$. For the self-similar version of the 2DTM tiling, notice that the white prototile is sitting in the unit square in the first quadrant, its substitution is sitting at the 4×4 square in the first quadrant, its second substitution will be the 16×16 square in the first quadrant, and so on.

3. The combinatorics of tilings

Key to our methodology is the connection between the combinatorics of tilings and their geometry. In this section we collect the definitions necessary for our construction.

3.1. Combinatorial and geometric graphs. A *combinatorial graph* K consists of a set $V(K)$, whose elements we call *vertices*, and a set $E(K)$, whose elements we call *edges*. Each edge in $E(K)$ is defined to be an unordered pair of vertices; we call these vertices the *endpoints* of that edge. The *degree* of a vertex is the number of edges for which it is an endpoint, and a *dangling* vertex is a vertex of degree 1. In this paper we will never need to consider vertices with degree 0.

Following Gross and Tucker [16], we define the *topological realization* \tilde{K} of a combinatorial graph K as follows. Each edge of K is identified with a copy of $[0, 1]$, where 0 represents one endpoint vertex of the edge and 1 represents the other. Whenever two edges share a common vertex we identify the appropriate endpoints of their copies of $[0, 1]$. The end result is a topological space \tilde{K} that encodes the combinatorics of K .

It is essential for our purposes to visualize combinatorial graphs in the plane by their topological realizations. Intuitively, we to do this by first choosing a point set in the plane in one-to-one correspondence with $V(K)$. Then any pair of points whose associated vertices make up an edge are connected by a Jordan arc.

Definition 3.1. A *geometric graph* G is an embedding of the topological realization \tilde{K} of a combinatorial graph K into the plane. Let ι_G denote the embedding map. We say that two geometric graphs G, H are *equivalent* if they are embeddings of the same combinatorial graph, and write $G \sim H$.

Combinatorial graphs that can be embedded in the plane are often called *planar* or *plane graphs*, and we call their embeddings *geometric graphs*. The combinatorial graph that gives rise to a given connected geometric graph G is not unique because degree two vertices cannot be detected. However, there is always a unique combinatorial graph K with no vertices of degree two for which $G = \iota_G(\tilde{K})$. This makes the following definition of the edge set $E(G)$ quite natural.

Definition 3.2. Let G be a geometric graph and K a combinatorial graph with no vertices of degree two such that $\iota_G(\tilde{K}) = G$, where \tilde{K} is the topological realization of K . An *edge* (resp. *vertex*) of G is the image under ι_G of the topological realization of an edge (resp. vertex) in K .

3.2. Tilings as geometric graphs. A tiling T in the plane gives rise to a canonical geometric graph: each point at which three or more tiles meet represents a vertex, and any arc along which two tiles meet represents an edge. Although this combinatorial graph is quite natural, it can cause problems because ideally, prototiles would have well-defined edges and vertices that carry throughout T . However, there are the tilings, such as the chair tiling, whose prototiles will not have a well defined vertex set unless either the prototile set or the vertex set is enlarged to account for them. Thus we choose to make a more subtle definition for the graph of a tiling.

Suppose T is a tiling with finite local complexity and prototile set \mathcal{P} with boundary $\partial\mathcal{P}$. We form an equivalence relation on $\partial\mathcal{P}$ by $a \sim_T b$ if there exists $x, y \in \mathbb{R}^2$ and $p, q \in \mathcal{P}$ such that $p + x, q + y \in T$ and $a + x = b + y$. A set $F \subset \partial\mathcal{P}$ is *T -consistent* if it is a union of equivalence classes of \sim_T .

Definition 3.3. A *natural vertex* of T is a point at which three or more tiles meet; the natural vertices of \mathcal{P} are their representatives on the prototile set. The *vertex set* of \mathcal{P} is the T -consistent set of points generated by the natural vertices of \mathcal{P} . An *edge* in $p \in \mathcal{P}$ is a Jordan arc connecting pairs of neighboring vertices along the boundary of p .

In this way the boundary of each prototile is associated with a combinatorial and geometric graph. In Figure 5 we show how the vertex set for a chair tile is determined in this perspective.




Figure 5. A patch from the third substitution of the chair tiling (see Example 4.3). The natural vertices marked with a \bullet , and the vertices in the T -consistent set generated by the natural vertices are marked with an \times .

We use this definition of prototile edges and vertices to define a canonical combinatorial graph K_T associated with T . Since every tile in T is a translation of a prototile, the edges and vertices of the tile are inherited from the prototile by translation. Since the edges and vertices of tiles are embedded on the boundaries of the tiles, they appear in the supports of more than one tile. To define the edge and vertex sets of K_T we include only one vertex and one edge from any given pair of adjacent tiles. This combinatorial graph is conveniently equipped with the embedding ι_T that places the vertices and edges where they came from in the first place; the image of ι_T is ∂T , the boundary of T . We note that if T and T' are translates of one another, then K_T and $K_{T'}$ are graph isomorphic.

Definition 3.4. A tiling is defined to be *edge-to-edge* if any two tile edges intersect either completely, at a common vertex, or not at all. A tiling is defined to be *singly edge-to-edge* provided any two tiles intersect along at most one edge.

If T is a FLC tiling whose edges and vertices come from the edges and vertices of ∂T , then T is automatically edge-to-edge. The chair tiling is an instructive example; by adding vertices, as in Figure 5, we have ensured that tile edges always line up. However, doing so means that chair tilings are no longer singly edge-to-edge.

3.3. Geometric graphs on prototiles and the tilings they induce. A key construction used in this work is to make new tilings from a given tiling T by marking all its tiles in a specified way. This process of turning old tilings into new ones by carving up the prototile set is what is referred to as *recomposition* in [17]. Our method for doing this is by embedding geometric graphs into the supports of the prototiles, then extending to all the tiles of T .

Definition 3.5. A *geometric graph* G on a prototile set \mathcal{P} is the disjoint union $\bigsqcup_{p \in \mathcal{P}} G_p$ of finite geometric graphs such that $G_p \subset \text{supp}(p)$ for all $p \in \mathcal{P}$. We denote by K_p the underlying combinatorial graph of G_p that has no degree 2 vertices; we call its embedding map $(\iota_G)_p$. A vertex of G_p that is contained on the boundary of $\text{supp}(p)$ is called a *boundary vertex* of G and a vertex contained in the interior of $\text{supp}(p)$ is called an *interior vertex*.

We now introduce a condition ensuring that edges of geometric graphs always meet when we translate prototiles to make a tiling.

Definition 3.6. Suppose T is a FLC tiling and let G be a geometric graph on its prototile set \mathcal{P} . We say that G is *T -consistent* if

- (i) for all $p \in \mathcal{P}$, G_p only intersects the boundary of $\text{supp}(p)$ at boundary vertices and
- (ii) the boundary vertices of G form a T -consistent set.³

The next definition gives sufficient (but not necessary) conditions on G so that it induces a tiling T_G by recomposition.

Definition 3.7. Suppose T is a FLC tiling and let G be a T -consistent geometric graph on its prototile set \mathcal{P} . We say that G is a *quasi-dual graph* if

- (i) the graph G_p is a connected tree for all $p \in \mathcal{P}$,
- (ii) the interior of each prototile edge contains exactly one boundary vertex of G , situated on the interior of that prototile edge, and
- (iii) every interior vertex of G has degree at least 3.

If G is a quasi-dual graph such that every prototile contains exactly one interior vertex then we call G a *dual graph*.

An example of a dual graph for the 2DTM tiling is given on the left hand side of Figure 2 in the introduction.

A T -consistent graph G on \mathcal{P} extends to a graph in \mathbb{R}^2 as follows. For each $t \in T$, denote by x_t the translation vector of t for which $t = p + x_t$ for its $p \in \mathcal{P}$. Then we define

$$\Gamma(T, G) = \bigcup_{p \in \mathcal{P}} \bigcup_{\{t \in T : t \text{ is of type } p\}} (G_p + x_t) \tag{3.1}$$

³ For readers familiar with the Anderson–Putnam complex [1] of a tiling space, a T -consistent embedding is a geometric graph G on the Anderson–Putnam complex with no dangling vertices.

This is a union of Jordan arcs that can also be seen as a geometric graph in \mathbb{R}^2 . It has no dangling vertices since G is T -consistent; the only way it can fail to be a tiling is if some of the Jordan arcs do not close up to become Jordan curves.

An *empty* Jordan curve in $\Gamma(T, G)$ is a Jordan curve that contains no portion of $\Gamma(T, G)$ in its interior. In the following lemma we define labelled empty Jordan curves to be the tiles of a new tiling induced by a T -consistent graph.

Lemma 3.8. *Suppose G is a T -consistent graph such that each arc in $\Gamma(T, G)$ is part of an empty Jordan curve. Then $\Gamma(T, G)$ is the boundary of a tiling T_G such that every empty Jordan curve is the boundary of a tile in T_G .*

Proof. We begin by constructing a prototile set \mathcal{P}_G which we will use to define the tiling T_G . A label set \mathcal{L} is defined by

$$\mathcal{L} := \{[J + x]^T : J \text{ is an empty Jordan curve and } J + x \in \Gamma(T, G)\} / \sim,$$

where \sim denotes translational equivalence on patches of tiles in T . Since T has FLC the set \mathcal{L} is finite. A set of prototiles \mathcal{P}_G is defined by taking the closed set bounded by a representative from each translational equivalence class of empty Jordan curves for each label in \mathcal{L} . Then each empty Jordan curve in $\Gamma(T, G)$ uniquely defines a tile that is a translation of a prototile in \mathcal{P}_G . The union of all the tiles defined by empty Jordan curves in $\Gamma(T, G)$ defines a tiling T_G , as required. \square

If G is a quasi-dual graph on a substitution tiling T , then the hypotheses of Lemma 3.8 are satisfied and T_G is a tiling. We call tilings arising from quasi-dual graphs *quasi-dual tilings*. Moreover, if G is a dual graph, then T_G is a labelled combinatorial dual of T .

Proposition 3.9. *Suppose T is a tiling with FLC and G is a quasi-dual graph on \mathcal{P} , then the quasi-dual tiling T_G has FLC and is mutually locally derivable to T . Moreover, the vertex patterns in T are in one-to-one correspondence with the tiles in T_G .*

Proof. Let v be a vertex in T . Consider the edges emanating from v in clockwise order; there is a unique path through $\Gamma(T, G)$ connecting each edge to the next by conditions (i) and (ii) of Definition 3.7. These edges form the boundary of the tile in T_G that corresponds to v . Moreover, Definition 3.7(i) guarantees that there can be no additional tiles in T_G .

Mutual local derivability follows from the fact that vertex patches in T give rise to tiles in T_G while tiles in T_G specify their corresponding patches in T by their label. \square

4. Alternate views of tiling substitutions

In this section we describe tiling substitutions using digit sets and digit matrices. We then use digit sets to define a contraction map on substitution tilings, which turns out to be essential for our construction. At the end of the section we calculate the digit matrix and contraction map for the chair substitution.

4.1. Matrices associated with tiling substitutions. We follow the method of Moody and Lee [21, Section 2]. Consider a fixed tiling substitution ω on prototile set \mathcal{P} with expansion factor λ as in Definition 2.4. Suppose $|\mathcal{P}| = m$ and that when \mathcal{P} is used as an index set for the rows and columns of a matrix we keep some fixed order.

We now introduce the substitution matrix of a substitution tiling. We note that our definition is the transpose of what is usually called the substitution matrix, but is better suited to our subsequent definitions.

Definition 4.1. The *substitution matrix* of ω is the $m \times m$ matrix M indexed by \mathcal{P} whose (p, q) -entry is equal to the number of copies of the prototile q occurring in $\omega(p)$.

The substitution matrix contains basic information about ω . A matrix that keeps more geometric information about the substitution is the *digit matrix*

$$D = (D_{pq}),$$

which contains the translation vectors required to implement the tiling substitution. The patch given by $\omega(p)$ is a collection of translates of the prototiles $q \in \mathcal{P}$, where the number of copies of each q is M_{pq} . The translation vectors for q in $\omega(p)$ are denoted $d_{pq}^1, d_{pq}^2, \dots, d_{pq}^{M_{pq}}$ and we define the entries of the digit matrix D to be these sets of translation vectors:

$$D_{pq} = \begin{cases} \{d_{pq}^k \in \mathbb{R}^2 : k = 1, \dots, M_{pq}\} & \text{if } M_{pq} \neq 0, \\ \emptyset & \text{if } M_{pq} = 0. \end{cases}$$

We will use D to translate both subsets of \mathbb{R}^2 (especially geometric graphs on \mathcal{P}) and prototiles q as follows. For $S \subset \mathbb{R}^2$ and $D_{pq} \neq \emptyset$ define

$$S + D_{pq} := \bigcup_{k=1}^{M_{pq}} (S + d_{pq}^k).$$

If $D_{pq} = \emptyset$ we take the convention that $S + \emptyset = \emptyset$. For a tile q and $D_{pq} \neq \emptyset$, we define $q + D_{pq}$ to be the set of tiles $\{q + d_{pq}^k : k \in 1, 2, \dots, M_{pq}\}$, and again use the convention that $q + \emptyset = \emptyset$. With this notation we can write

$$\omega(p) = \bigcup_{q \in \mathcal{P}} (q + D_{pq}). \tag{4.1}$$

4.2. Tiling contraction maps. In this section we will define a tiling contraction map \mathcal{R} . When restricted to prototiles we want \mathcal{R} to be equal to $\lambda^{-1}\omega$, which subdivides each prototile into subtiles rather than inflating and then subdividing. However, we want \mathcal{R} to be a more general map and we will see that the domain of \mathcal{R} plays an important role in its definition. In particular, we would like \mathcal{R} to be a ‘shrink-and-replace’ rule on patches of tiles and on disjoint unions of compact subsets of \mathbb{R}^2 , though our prototypical application is to geometric graphs embedded in tiles. We will use the notation \mathcal{R} regardless of what type of object is being acted on, with the understanding that the output is always the same type as the input.

Before defining \mathcal{R} we will define its domain. In order to define the domain of \mathcal{R} we need to decide whether we want it to act on disjoint unions of compact subsets of \mathbb{R}^2 or on patches of scaled tiles. For the former, let $H(\mathbb{R}^2)$ denote the set of nonempty compact subsets of \mathbb{R}^2 and define the domain of \mathcal{R} to be

$$X := \bigsqcup_{p \in \mathcal{P}} H(\mathbb{R}^2).$$

The latter is slightly trickier. Given a prototile set \mathcal{P} , consider the set of all nonempty patches of translated prototiles, denoted \mathcal{P}^* . For any $0 < \kappa \leq 1$ the set $\kappa\mathcal{P}^*$ is the set of all nonempty patches of prototiles scaled by κ . When we want \mathcal{R} to act on scaled patches of tiles the domain is then

$$X_{\mathcal{P}} := \bigcup_{0 < \kappa \leq 1} \bigsqcup_{p \in \mathcal{P}} \kappa\mathcal{P}^*.$$

Definition 4.2. We define *tiling contraction map* \mathcal{R} as follows. Let $B = \bigsqcup_{p \in \mathcal{P}} B_p$ be an element of X or $X_{\mathcal{P}}$. For each $p \in \mathcal{P}$ let

$$\mathcal{R}(B)_p = \bigcup_{q \in \mathcal{P}} \lambda^{-1} (B_q + D_{pq}), \tag{4.2}$$

and then $\mathcal{R}(B) = \bigsqcup_{p \in \mathcal{P}} \mathcal{R}(B)_p$.

Since \mathcal{R} has the same definition regardless of the domain we now illustrate the difference. The first case is if $B \in X$, and then $\mathcal{R}(B) \in X$ as well. By comparing equation (4.2) to equation (4.1) we see that if $B = \bigsqcup_{p \in \mathcal{P}} \text{supp}(p)$ then we have

$$\mathcal{R}(B) = \bigsqcup_{p \in \mathcal{P}} \bigcup_{q \in \mathcal{P}} \lambda^{-1}(\text{supp}(q) + D_{pq}) = \bigsqcup_{p \in \mathcal{P}} \text{supp}(p) = B.$$

Thus the disjoint union of prototile supports is fixed under the action of \mathcal{R} and is therefore the attractor of the dynamical system (X, \mathcal{R}) . In the special case that G is a geometric graph on \mathcal{P} , then $\mathcal{R}(G)$ is also a geometric graph on \mathcal{P} . If G is T -consistent then $\mathcal{R}(G)$ will be T -consistent with no dangling vertices in its interior. Moreover, if G is quasi-dual then $\mathcal{R}(G)$ will be connected but will have too many boundary vertices to be quasi-dual itself. In the next section we will be selecting subgraphs of $\mathcal{R}(G)$ to define recurrent pairs.

On the other hand, if $B \in X_{\mathcal{P}}$, then $\mathcal{R}(B)$ may fail to be in $X_{\mathcal{P}}$. This can happen when the patches that make up $\mathcal{R}(B)_p$ have overlapping interiors. To deal with this annoyance we will restrict our attention to those B for which $\text{supp}(B_p) = \text{supp}(p)$ for all $p \in \mathcal{P}$, in which case both $\mathcal{R}(B)$ and $\mathcal{R}^n(B)$ are elements of $X_{\mathcal{P}}$ for all $n \in \mathbb{N}$. With a slight abuse of terminology we write $\mathcal{R}^n(\mathcal{P})$ for $\mathcal{R}^n(\bigsqcup_{p \in \mathcal{P}} p)$. Then $\mathcal{R}(\mathcal{P})_p = \lambda^{-1}\omega(p)$ for all $p \in \mathcal{P}$. We call a tile in $\mathcal{R}(\mathcal{P})_p$ a *subtile*; and \mathcal{R} can be applied repeatedly so that $\mathcal{R}^n(\mathcal{P})$ consists of patches of \mathcal{P} -tiles scaled by λ^{-n} that we call *n-subtiles*.

Example 4.3 (the chair tiling). Let \mathcal{P} be the set of four prototiles with long side length 1 and the origin in the corner marked by the dot, as depicted in Figure 6. These are the prototiles of the chair tiling.

The expansion factor for the chair substitution is $\lambda = 2$, and the substitution of prototile p_1 also appears in Figure 6; the substitution of each of its three rotations are just the rotation of this substitution about the origin. The digit matrix for the chair substitution is

$$D = \begin{pmatrix} \{(0, 0), (\frac{1}{2}, \frac{1}{2})\} & \{(0, 2)\} & \emptyset & \{(2, 0)\} \\ \{(0, -2)\} & \{(0, 0), (\frac{1}{2}, -\frac{1}{2})\} & \{(2, 0)\} & \emptyset \\ \emptyset & \{(-2, 0)\} & \{(0, 0), (-\frac{1}{2}, -\frac{1}{2})\} & \{(0, -2)\} \\ \{(-2, 0)\} & \emptyset & \{(0, 2)\} & \{(0, 0), (-\frac{1}{2}, \frac{1}{2})\} \end{pmatrix}.$$




Figure 6. The 4 prototiles of the chair tiling and the substitution of prototile p_1 .

Thus, for $B = \bigsqcup_{i=1}^4 B_i$ in X or $X_{\mathcal{P}}$, the tiling contraction map \mathcal{R} is given by

$$\mathcal{R} \begin{pmatrix} B_1 \\ B_2 \\ B_3 \\ B_4 \end{pmatrix} = \frac{1}{2} \begin{pmatrix} B_1 \cup (B_1 + (\frac{1}{2}, \frac{1}{2})) \cup (B_2 + (0, 2)) \cup (B_4 + (2, 0)) \\ (B_1 + (0, -2)) \cup B_2 \cup (B_2 + (\frac{1}{2}, -\frac{1}{2})) \cup (B_3 + (2, 0)) \\ (B_2 + (-2, 0)) \cup B_3 \cup (B_3 + (-\frac{1}{2}, \frac{1}{2})) \cup (B_4 + (0, -2)) \\ (B_1 + (-2, 0)) \cup (B_3 + (0, 2)) \cup B_4 \cup (B_4 + (-\frac{1}{2}, \frac{1}{2})) \end{pmatrix}.$$

5. Constructing fractal tilings from self-similar tilings

In this section we construct fractal substitution tilings. Our method uses a self-similar tiling to construct the fractal tiling explicitly. We use this method rather than standard techniques from fractal geometry because we obtain more information about the relationship between the the original tiling and a fractal realization. In particular, our techniques allow for the possibility of experimentation and we are able to pinpoint a necessary condition for a recurrent pair on a substitution tiling to give rise to a fractal substitution tiling.

Here is a brief overview of the method. Suppose T is a fixed self-similar tiling with prototile set \mathcal{P} and substitution map ω . Given a geometric graph G on \mathcal{P} satisfying the consistency conditions of Definition 3.6, we let $T^{(0)} := T_G$ be the tiling defined in Definition 3.8 obtained by embedding G in all the tiles of T . We then select a geometric graph $S \subset \mathcal{R}(G)$ with the same combinatorics as G , forming a recurrent pair (G, S) . The map taking G to S acts as a substitution rule on edges, and we define the map ψ to be the extension of the substitution to the boundary of $T^{(0)}$ and obtain a tiling $T^{(1)} := \psi(T^{(0)})$. Assuming certain injectivity conditions on the recurrent pair (G, S) and iterating this process leads to a limiting tiling $T^{(\infty)}$, which is also a self-similar tiling.

5.1. Recurrent pairs and edge refinement maps. Recall that two geometric graphs B and C are said to be equivalent, and we write $B \sim C$, provided there exists a combinatorial graph K and embeddings ι_B and ι_C of the topological realization of K such that $\text{Image}(\iota_B) = B$ and $\text{Image}(\iota_C) = C$. In this case there is a homeomorphism between B and C given by $\iota_C \circ \iota_B^{-1}$.

Definition 5.1. A pair of T -consistent geometric graphs (G, S) on \mathcal{P} is said to be a *recurrent pair* for (T, ω) if the following conditions hold:

- (i) $S \subset \mathcal{R}^N(G)$ for some $N \in \mathbb{N}$,
- (ii) for all $p \in \mathcal{P}$, $G_p \sim S_p$ with homeomorphism denoted by $\tilde{\psi}_p : G_p \rightarrow S_p$.
- (iii) v is a boundary vertex of G_p if and only if $\tilde{\psi}_p(v)$ is a boundary vertex of S_p , in which case both lie in the same edge of p .

Notice also that we can always take $N = 1$ in the definition of a recurrent pair by replacing \mathcal{R} with \mathcal{R}^N . So after we find a recurrent pair (G, S) we will typically assume $S \subset \mathcal{R}(G)$.

Figure 2 of the introduction shows how a recurrent pair for the 2DTM substitution is constructed. The leftmost images show the prototiles α and β with dual graphs G_α and G_β inscribed. The disjoint union of G_α and G_β is the graph G . The images labelled $\mathcal{R}(G)_\alpha$ and $\mathcal{R}(G)_\beta$ show the graph G inscribed in each tile of the substitution. The images labelled S_α and S_β show the selection of a graph S that makes (G, S) a recurrent pair for the 2DTM substitution.

Notice that in a recurrent pair, since $S \subset \mathcal{R}(G)$, an edge of S is made up of a union of edges from G rescaled by λ^{-1} , and S passes through many vertices of $\mathcal{R}(G)$ that do not affect the mutual underlying combinatorial graph K since they are of degree 2. There is a very meaningful sense in which we think of S as being the substitution of the graph G ; describing how to iterate the substitution to produce fractals is our next task.⁴

In all that follows, we suppose that when G is drawn in all the tiles of T , it forms the boundary of a tiling called T_G , defined by Lemma 3.8, which we take to be our initial tiling $T^{(0)}$. Since the boundary graph of $T^{(0)}$ consists of translated elements of G we can extend the $\tilde{\psi}_p$ maps to a map $\psi : \partial T^{(0)} \rightarrow \mathbb{R}^2$ that redraws an edge of $T^{(0)}$ the same way $\tilde{\psi}_p$ redraw it in its corresponding prototile. That is, if $z \in \partial T^{(0)}$ is in the support of the tile $t = p + x \in T$ for $p \in \mathcal{P}$ and $x \in \mathbb{R}^2$, then $\psi(z) = \tilde{\psi}_p(z - x) + x$.

⁴ We will make this definition with the aid of our self-similar tiling T ; a different but equivalent formulation can be made using the prototile set \mathcal{P} only.

Lemma 5.2. *The map $\psi : \partial T^{(0)} \rightarrow \mathbb{R}^2$ is a homeomorphism onto its image.*

Proof. Since ψ restricted to the interior of any T -tile is the translation of a homeomorphism, the only question is what happens on $\partial T^{(0)} \cap \partial T$. Part (iii) in the definition of a recurrent pair along with the fact that S is T -consistent imply that ψ is well-defined and continuous. The fact that G is T -consistent implies that ψ is injective. Thus ψ is a homeomorphism. \square

We define the tiling $T^{(1)}$ to be the tiling with boundary $\psi(\partial T^{(0)})$; its tiles are given by the empty Jordan curves in this graph and we allow a tile in $T^{(1)}$ to inherit the label of its corresponding tile in $T^{(0)}$. It should be noted that $\partial T^{(1)} = \psi(\partial T^{(0)})$ is the boundary of the tiling T_S induced by S on T . Like $T^{(0)}$, the tiling $T^{(1)}$ is typically not self similar but is pseudo-self-similar [26] and the two tilings are combinatorially equivalent and mutually locally derivable.

The set $\psi(\partial T^{(0)}) = \partial T^{(1)}$ is not contained in $\partial T^{(0)}$, but because it is induced from the self-similar tiling T it is contained in $\lambda^{-1}(\partial T^{(0)})$. Thus if we wish to apply ψ again, we must do so on $\lambda\psi(\partial T^{(0)})$. Taking this into account we define $\psi^{(n)} : \partial T^{(0)} \rightarrow \mathbb{R}^2$ by

$$\psi^{(n)}(z) = \lambda^{-n}(\lambda\psi)^n(z) \tag{5.1}$$

That is, the edges in $\partial T^{(0)}$ are redrawn and then inflated by λ so that the result is again in $\partial T^{(0)}$, after which they are redrawn and inflated again until this has happened n times. Rescaling the result by λ^{-n} brings it back to the original scale. The map $(\lambda\psi)^n$ can be thought of as a kind of ‘inflate-and-subdivide’ rule for edges of G , taking an edge in $\partial T^{(0)}$ to a sequence of edges in $\partial T^{(0)}$ of length approximately λ^n times as long.

The map $\psi^{(n)}$ is a homeomorphism for any finite n and we define $T^{(n)}$ to be the tiling with boundary graph $\psi^{(n)}(\partial T^{(0)})$ whose tiles inherit the labels of their corresponding tiles in $T^{(0)}$:

Definition 5.3. For each tile $t \in T^{(0)}$ we define the tile $t^{(n)}$ in $T^{(n)}$ as being supported by the set enclosed by $\psi^{(n)}(\partial t)$ and carrying the label of t .

We define $\psi^{(\infty)}$ to be the pointwise limit of $\psi^{(n)}$, whose existence is proved in the following lemma.

Lemma 5.4. *The sequence $(\psi^{(n)})_{n=1}^{\infty}$ is uniformly Cauchy. Hence the pointwise limit, denoted $\psi^{(\infty)}$, is continuous.*

Proof. Fix $\varepsilon > 0$. Let $C = \max_{t \in T} \text{diam } t$, and choose $N \in \mathbb{N}$ such that $C\lambda^{-N} < \varepsilon$. Fix $m, n > N$, and assume without loss of generality that $m < n$. For $z \in \partial T$ we have

$$\begin{aligned} & |\psi^{(m)}(z) - \psi^{(n)}(z)| \\ &= |\lambda^{-m}(\lambda\psi)^m(z) - \lambda^{-n}(\lambda\psi)^n(z)| \\ &= \lambda^{-m}|(\lambda\psi)^m(z) - \lambda^{-(n-m)}(\lambda\psi)^{n-m}(\lambda\psi)^m(z)| \\ &= \lambda^{-m}|y - \psi^{(n-m)}(y)| \quad (\text{letting } y = (\lambda\psi)^m(z)) \\ &\leq \lambda^{-m}C \quad (\text{since } y \text{ and } \psi^{(n-m)}(y) \text{ are in the same tile of } T) \\ &< \varepsilon. \end{aligned}$$

So $(\psi^{(n)})_{n=1}^\infty$ is uniformly Cauchy, and $\psi^{(\infty)}$ is continuous. \square

5.2. The limiting fractal tiling. When $\psi^{(\infty)}$ is injective its image forms the boundary of a tiling which we denote $T^{(\infty)}$ and which we will show is self-similar. Under the condition of injectivity we make the formal definition of a fractal realization.

Definition 5.5. For each tile $t \in T^{(0)}$ we define the tile $t^{(\infty)}$ in $T^{(\infty)}$ as being supported by the set enclosed by $\psi^{(\infty)}(\partial t)$ and carrying the label of t . We call $T^{(\infty)}$ a *fractal realization* of $T^{(0)}$ (or of G).

We collect a few observations about $\psi^{(n)}$, $\psi^{(\infty)}$, $T^{(n)}$ and $T^{(\infty)}$ in the following lemma.

Lemma 5.6. *When $\psi^{(\infty)}$ is injective,*

- (1) $\psi^{(1)} = \psi$;
- (2) if $z \in \partial T^{(0)}$ and $z \in \text{supp}(t)$ for a T -tile t , then $\psi^{(n)}(z) \in \text{supp}(t)$ for all n ;
- (3) T , $T^{(0)}$, $T^{(n)}$, and $T^{(\infty)}$ are MLD for all $n \in \mathbb{N}$, and
- (4) For any $z \in \partial T^{(0)}$ and $n \in \mathbb{N}$;

$$\lambda^n \psi^{(\infty)}(z) = \psi^{(\infty)}((\lambda\psi)^n(z)). \quad (5.2)$$

Proof. The first part follows directly from the definition of $\psi^{(k)}$. The second follows from the construction of $\psi^{(n)}$ by induction, since it is true for $\psi = \psi^{(1)}$. The fact that $T^{(0)}, T^{(n)}$, and $T^{(\infty)}$ are MLD is immediate from the definitions. Since every tile in $T^{(0)}$ corresponds to a patch in T we see that $T^{(0)}$ and T are MLD. Part (3) can be seen by considering only $k > n$ in the following:

$$\begin{aligned} \lambda^n \psi^{(\infty)}(z) &= \lim_{k \rightarrow \infty} \lambda^{n-k} (\lambda \psi)^k(z) \\ &= \lim_{k \rightarrow \infty} \lambda^{n-k} (\lambda \psi)^{k-n} ((\lambda \psi)^n(z)) \\ &= \psi^{(\infty)}((\lambda \psi)^n(z)). \end{aligned} \quad \square$$

Theorem 5.7. *Let T be self-similar tiling with expansion factor λ , and let (G, S) be a recurrent pair for T . If $\psi^{(\infty)}$ is injective then $T^{(\infty)}$ is a self-similar tiling with expansion factor λ and finite local complexity.*

Proof. Since finite local complexity is preserved by MLD equivalence, by Lemma 5.6(3), $T^{(\infty)}$ has finite local complexity.

To prove $T^{(\infty)}$ is self-similar we show that its boundary is invariant under scaling by λ and that equivalent tiles inflate to equivalent patches. The former follows directly from Lemma 5.6, since if $x = \psi^{(\infty)}(z) \in \partial T^{(\infty)}$,

$$\lambda x = \psi^{(\infty)}((\lambda \psi)(z)) \in \partial T^{(\infty)},$$

as desired. To prove the latter suppose t_1 and t_2 are equivalent tiles in $T^{(0)}$ so that $t_1^{(\infty)}$ and $t_2^{(\infty)}$ are equivalent tiles in $T^{(\infty)}$. Denote by $[t_1]^T$ and $[t_2]^T$ the patches in T that intersect the supports of t_1 and t_2 , respectively; these are also equivalent, by definition, since t_1 and t_2 carry the same label. Since T is self-similar the T -patches in $\lambda \text{supp}[t_1]^T$ and $\lambda \text{supp}[t_2]^T$ are equivalent. This means that the $T^{(0)}$ - and $T^{(\infty)}$ -patches contained completely inside them are equivalent as well. Since $\text{supp}(t_1^{(\infty)})$ is contained in $[t_1]^T$ the same way $\text{supp}(t_2^{(\infty)})$ is contained in $[t_2]^T$, $\lambda \text{supp}(t_1^{(\infty)})$ and $\lambda \text{supp}(t_2^{(\infty)})$ support equivalent $T^{(\infty)}$ -patches, as required. \square

Under these circumstances we will let $\mathcal{P}^{(\infty)}$ denote the set of prototiles for the self-similar tiling $T^{(\infty)}$, and let $\omega_{(\infty)}$ denote its substitution map.

6. The existence of fractal dual and quasi-dual tilings

The process described in the previous section can be used to find recurrent pairs with injective $\psi^{(\infty)}$ maps by experimentation. This section contains technical results related to the existence of such pairs in general.

In the first part of this section we provide sufficient conditions on a recurrent pair (G, S) for the map $\psi^{(\infty)}$ to be injective. These conditions result in fractal tilings that are quasi-dual to their original self-similar tilings. In the second part of the section we show that every singly edge-to-edge self-similar tiling has a recurrent pair satisfying the injectivity conditions. In the third part of this section we prove that if a singly edge-to-edge self-similar tiling has convex prototiles, then its labelled combinatorial dual has a fractal realization.

6.1. Conditions under which $\psi^{(\infty)}$ is injective. We introduce conditions that ensure that the combinatorics of $\psi^{(\infty)}(\partial T^{(0)})$ are identical to the combinatorics of $T^{(0)}$. One way this can fail if the interior of an edge in $\psi^{(n)}(\partial T^{(0)})$ approaches the boundary of a tile in the original self-similar tiling T as $n \rightarrow \infty$. If $\psi^{(\infty)}(z) \in \partial T$ and $z \notin \partial T$, it can happen that an additional tile (or tiles) can arise in $T^{(\infty)}$ either by two edges coming together that were apart in $T^{(0)}$ or by an edge doubling back on itself. Infinitely many arbitrarily small tiles are also a danger in this situation. We can avoid these troubles by keeping $\psi^{(n)}(\partial T^{(0)})$ away from the boundary of T except at T -consistent boundary vertices.

Another way the combinatorics of $\psi^{(\infty)}(\partial T^{(0)})$ can fail to be those of $\partial T^{(0)}$ is if $\psi^{(n)}(z)$ approaches the vertex of a tile in T as $n \rightarrow \infty$. In this situation there may still be a perfectly good self-similar tiling $T^{(\infty)}$, but we've lost control of its combinatorics and its substitution rule is not guaranteed to be border forcing.

Before we give the injectivity conditions we set some notation. In what follows, for a real number $k > 0$ and any tiling T , the notation kT represents the tiling obtained by rescaling all the tiles of T by a factor of k and keeping the labels the same. Recall that if B is any subset of \mathbb{R}^2 or patch of tiles and Q is a tiling or patch the notation $[B]^Q$ represents the patch in Q whose support intersects B . In particular, for $p \in \mathcal{P}$ note that $[B]^{\mathcal{R}^n(\mathcal{P})}_p$ is the patch of n -subtiles of $\mathcal{R}^n(\mathcal{P})$ in p that intersect B .

Definition 6.1 (injectivity conditions). Let S be a geometric graph on \mathcal{P} . If there is an $N \in \mathbb{N}$ for which the following four conditions hold, we say that S *satisfies the injectivity conditions for N -subtiles*.

- (I1) For any $p \in \mathcal{P}$ and edges $e \neq f$ in S_p , the patches $[e]^{\mathcal{R}^N(\mathcal{P})}_p$ and $[f]^{\mathcal{R}^N(\mathcal{P})}_p$ have intersecting interiors if and only if e and f share a common vertex v .
- (I2) If $e \neq f \in S_p$ share a common vertex v then

$$[e]^{\mathcal{R}^N(\mathcal{P})}_p \cap [f]^{\mathcal{R}^N(\mathcal{P})}_p = [v]^{\mathcal{R}^N(\mathcal{P})}_p,$$

the single subtile containing v , and this subtile is contained in the interior of $\text{supp}(p)$.

- (I3) For each $p \in \mathcal{P}$, $[S_p]^{\mathcal{R}^N(\mathcal{P})_p}$ does not contain any vertex of p in its support.
- (I4) For e in S_p , the support of $[e]^{\mathcal{R}^N(\mathcal{P})_p}$ intersects the boundary of p if and only if e does, in which case $\text{supp}[e]^{\mathcal{R}^N(\mathcal{P})_p} \cap \partial p$ is connected.

Notice that if (G, S) is a recurrent pair, G is a quasi-dual graph, and S satisfies the injectivity conditions for N -subtiles, then condition (I3) along with the definition of a recurrent pair implies that S is also a quasi-dual graph. In general these conditions imply that the interior of the patch of subtiles that S runs through retracts to a geometric graph that has nearly the same combinatorics as S . If the combinatorics differ it is at interior vertices of S , several of which could be collapsed into a single vertex in the retraction.

Given a recurrent pair (G, S) we make two assumptions on the embeddings ι_G and ι_S , which can be considered “without loss of generality.” The first is that G is a piecewise linear graph. If for some reason it is not, it can be redrawn as one: the topological realization of the combinatorial graph K corresponding to the geometric graph G can always be embedded into $\text{supp}(\mathcal{P})$ in a piecewise linear fashion. We assume, then, that such an embedding has been chosen for G . The second is on the embedding ι_S of S , which we can control since S , being a subgraph of $\mathcal{R}(G)$, is also a piecewise linear graph. Letting $B \subset e \in E(G)$ or $E(S)$, denote by $|B|$ the arc length of the smallest sub-arc of e containing B . In this case we assume that ι_S has been chosen so that for any $e \in E(G)$ there is a K_e such that for all $B \subset e$ we have $|\iota_S \circ \iota_G^{-1}(B)| = K_e|B|$. In other words, $\tilde{\psi}$ is a *piecewise constant-speed parameterization* taking G to S .

One further technical note is on the edge set of $\partial(T^{(0)})$, which can be thought of in two different ways. One way is to consider the vertex set endowed by the tiling $T^{(0)}$, and another is to allow $\partial(T^{(0)})$ to inherit the edge and vertex sets of G . The latter has extra vertices wherever the boundary of $T^{(0)}$ intersects the boundary of T . In all the proofs in this section we consider an edge of $T^{(0)}$ to mean a copy of an edge in G .

Proposition 6.2. *Suppose (G, S) is a recurrent pair such that G is quasi-dual and the parameterization $\tilde{\psi}: G \rightarrow S$ is piecewise constant-speed. If S satisfies the injectivity conditions on N -subtiles for some N , then $\psi^{(\infty)}: \partial T^{(0)} \rightarrow \partial T^{(\infty)}$ is injective.*

We prove Proposition 6.2 using a sequence of lemmas, each of which implicitly has the assumptions of Proposition 6.2. For convenience of notation we will assume $N = 1$; the general case can be surmised by replacing the substitution rule ω by ω^N , which changes the expansion factor to λ^N and the tiling contraction map to \mathcal{R}^N .

The injectivity conditions (II)–(I4) are designed to keep $\psi^{(n)}(e)$ away from the boundaries of both of $T^{(0)}$ and T except where the combinatorics of ∂T explicitly require it, while the constant-speed assumption keeps edges from collapsing onto themselves in the limit.

A key observation in everything that follows is that if e is an edge in $\partial T^{(0)}$ for which $e = e_p + z$ for some edge $e_p \in G_p$ and $z \in \mathbb{R}^2$, then

$$[\psi(e)]^{\lambda^{-1}T} = [\tilde{\psi}(e_p)]^{\mathcal{R}(P)_p} + z$$

by the self-similarity of T . Thus edges in $\psi(\partial T^{(0)})$ follow the same injectivity conditions, when intersected with $\lambda^{-1}T$, as edges in S do when intersected with $\mathcal{R}(P)$.

Lemma 6.3. *Suppose that $e \in E(\partial T^{(0)})$ and B is an edge in ∂T such that $B \cap e = \emptyset$. Then $[\lambda\psi(e)]^T \cap \lambda B = \emptyset$.*

Proof. If e has no vertex in ∂T then by (I4), $[\psi(e)]^{\lambda^{-1}T} \cap B = \emptyset$. If e has a vertex in ∂T then by (I3) and (I4), $[\psi(e)]^{\lambda^{-1}T}$ must intersect ∂T only on the edge that e does. However, $B \cap e = \emptyset$, and $[\lambda\psi(e)]^T \cap \lambda B = \emptyset$. □

Lemma 6.4. *Suppose that $x, y \in \partial T^{(0)}$ and that there exists $n \in \mathbb{N}$ such that $(\lambda\psi)^n(x)$ and $(\lambda\psi)^n(y)$ are on disjoint edges in $\partial T^{(0)}$. Then $\psi^{(\infty)}(x) \neq \psi^{(\infty)}(y)$.*

Proof. Denote by e_x and e_y the edges in $\partial T^{(0)}$ containing x and y , respectively. If x and y are in distinct tiles in T , then even if those tiles share a common boundary segment, that segment cannot intersect both e_x and e_y without violating either the fact that G is quasi-dual or that $e_x \cap e_y = \emptyset$. Thus, in this case, it follows that $\psi^{(\infty)}(x) \neq \psi^{(\infty)}(y)$.

Next we consider the case where e_x and e_y are in the same T -tile. Since $\lambda\psi$ is injective, we have that $(\lambda\psi(e_x)) \cap (\lambda\psi(e_y)) = \emptyset$. Then (II) implies that

$$\text{int}(\text{supp}[\lambda\psi(e_x)]^T) \cap \text{int}(\text{supp}[\lambda\psi(e_y)]^T) = \emptyset.$$

If we have $\text{supp}[\lambda\psi(e_x)]^T \cap \text{supp}[\lambda\psi(e_y)]^T = \emptyset$, then we are done since $\psi^{(\infty)}(x) \in \text{supp} \lambda^{-1}[\lambda\psi(e_x)]^T$ and $\psi^{(\infty)}(y) \in \text{supp} \lambda^{-1}[\lambda\psi(e_y)]^T$. If not, then $[\lambda\psi(e_x)]^T$ and $[\lambda\psi(e_y)]^T$ share some part B of their boundary. Since G is quasi-dual, neither $\lambda\psi(e_x)$ nor $\lambda\psi(e_y)$ intersect B , for otherwise $\lambda\psi(e_x)$ would pass through B into the interior of $[\lambda\psi(e_y)]^T$ or vice versa. By Lemma 6.3, we have that $[(\lambda\psi)^2(e_x)]^T \cap \lambda B = \emptyset$. So there exist $t_x, t_y \in T$ satisfying $(\lambda\psi)^2(x) \in t_x$, $(\lambda\psi)^2(y) \in t_y$ and $t_x \cap t_y = \emptyset$. Since $\lambda^2\psi^{(\infty)}(x) = \psi^{(\infty)}((\lambda\psi)^2(x)) \in \text{supp}(t_x)$ and $\lambda^2\psi^{(\infty)}(y) = \psi^{(\infty)}((\lambda\psi)^2(y)) \in \text{supp}(t_y)$, we have $\psi^{(\infty)}(x) \neq \psi^{(\infty)}(y)$.

Now suppose that there exists $n \in \mathbb{N}$ such that $(\lambda\psi)^n(x)$ and $(\lambda\psi)^n(y)$ are on disjoint edges in $E(\partial T^{(0)})$. Then by the previous argument,

$$\psi^{(\infty)}((\lambda\psi)^n(x)) \neq \psi^{(\infty)}((\lambda\psi)^n(y)),$$

making

$$\lambda^n \psi^{(\infty)}(x) \neq \lambda^n \psi^{(\infty)}(y),$$

and in turn

$$\psi^{(\infty)}(x) \neq \psi^{(\infty)}(y),$$

as desired. \square

Lemma 6.4 implies that edges that are disjoint in $T^{(0)}$ remain disjoint in $T^{(\infty)}$. The next lemma shows that $\psi^{(\infty)}$ restricted to edges of $T^{(0)}$ is injective, meaning that an edge in $T^{(0)}$ cannot loop back onto itself in $T^{(\infty)}$. This is the part of the proof that uses the assumption that the parameterizations are piecewise constant-speed.

Lemma 6.5. *For each $e \in E(\partial T^{(0)})$, $\psi^{(\infty)}|_e$ is injective.*

Proof. Fix $x, y \in e \in E(\partial T^{(0)})$, and suppose that $\psi^{(\infty)}(x) = \psi^{(\infty)}(y)$. We will show $x = y$. Write $x_n := (\lambda\psi)^n(x)$ and $y_n := (\lambda\psi)^n(y)$. In order for $\psi^{(\infty)}(x)$ and $\psi^{(\infty)}(y)$ to be equal, Lemma 6.4 implies that for each $n \in \mathbb{N}$, x_n and y_n are either on the same or adjacent edges in $E(\partial T^{(0)})$. This breaks into two cases: either they are on the same edge for all n or there is an N such that if $n \geq N$, x_n and y_n are on adjacent edges.

We first suppose that x_n and y_n are on the same edge e_n in $\partial T^{(0)}$ for all n . For each n , we have $e_{n+1} \subset \lambda\psi(e_n)$, and we let l_n be the arc in e_n for which $\lambda\psi(l_n) = e_{n+1}$. Set $\rho = \max_{n \in \mathbb{N}}(|l_n|/|e_n|)$. We can deduce that $\rho < 1$ by injectivity condition (I2) because for every $e \in E(\partial T)$, $\lambda\psi(e)$ is a union of at least 2 edges in $E(\partial T)$. The constant-speed assumption on $\tilde{\psi}$ extends to guarantee that $\lambda\psi$ is piecewise constant-speed as well.

For each $n \in \mathbb{N}$ and $k \leq n$ we know that $(\lambda\psi)^{-(n-k)}(e_n)$ and $(\lambda\psi)^{-(n-k)}(l_n)$ are in the same edge in $\partial T^{(0)}$ and the constant-speed assumption implies that

$$\frac{|(\lambda\psi)^{-n}(l_n)|}{|(\lambda\psi)^{-n}(e_n)|} = \frac{|(\lambda\psi)^{-(n-1)}(l_n)|}{|(\lambda\psi)^{-(n-1)}(e_n)|} = \dots = \frac{|l_n|}{|e_n|} \leq \rho.$$

Since $x, y \in (\lambda\psi)^{-n}(I_n)$ for all n , we have that

$$\begin{aligned} |x - y| &\leq |(\lambda\psi)^{-n}(I_n)| \leq \rho |(\lambda\psi)^{-n}(e_n)| \\ &= \rho |(\lambda\psi)^{-(n-1)}(I_{n-1})| \\ &\leq \rho^2 |(\lambda\psi)^{-(n-1)}(e_{n-1})| \\ &= \rho^2 |(\lambda\psi)^{-(n-2)}(I_{n-2})| \\ &\leq \\ &\vdots \\ &\leq \rho^n |e_0|. \end{aligned}$$

Since this is true for all n and ρ is strictly less than one we have shown $x = y$ when x_n and y_n are on the same edge of $\partial T^{(0)}$ for all n .

Now suppose there is an N such that if $n \geq N$ then x_n and y_n are on adjacent edges that share a vertex v_n . Since ψ is continuous, the edges containing $(\lambda\psi)^n(x_N)$ and $(\lambda\psi)^n(y_N)$ share the vertex $(\lambda\psi)^n(v_N)$ for all $n \in \mathbb{N}$. Then

$$\|(\lambda\psi)^n(x_N) - (\lambda\psi)^n(v_N)\| \leq \max_{t \in T} \text{diam}(t)$$

for all $n \in \mathbb{N}$. Multiplying both sides by λ^{-n} then gives

$$\|\psi^{(n)}((\lambda\psi)^N(x)) - \psi^{(n)}(v_N)\| \leq \lambda^{-n} \max_{t \in T} \text{diam}(t)$$

for all $n \in \mathbb{N}$. This implies that $\psi^{(\infty)}(x_N) = \psi^{(\infty)}(v_N)$, and since x_N and v_N are always on the same edge the first case in this proof shows that $x_N = v_N$. This argument shows that also $y_N = v_N$. Since $(\lambda\psi)^N$ is injective this implies that $x = y$ and we have shown that $\psi^{(\infty)}$ is injective on edges of $\partial T^{(0)}$. \square

Proof of Proposition 6.2. Fix $x, y \in \partial T$ and suppose that $x \neq y$. If $\psi^{(n)}(x)$ and $\psi^{(n)}(y)$ are on the same or on disjoint edges for some n , then Lemmas 6.4 and 6.5 imply that $\psi^{(\infty)}(x) \neq \psi^{(\infty)}(y)$. The only other option is that $\psi^{(n)}(x)$ and $\psi^{(n)}(y)$ are on adjacent edges for all n , but the proof of Lemma 6.5 shows that this cannot happen if $x \neq y$. Thus $\psi^{(\infty)}$ is injective. \square

6.2. Existence of recurrent pairs with injective edge refinements. In this section we prove that every singly edge-to-edge, primitive self-similar tiling T with finite local complexity has a recurrent pair (G, S) satisfying the conditions of Proposition 6.2. In this case the map $\psi^{(\infty)}$ is injective and the tiling T gives

rise to a fractal substitution tiling $T^{(\infty)}$. The construction of the recurrent pair often requires multiple iterations of the substitution ω , so the expansion factor of $T^{(\infty)}$ is always a power of the expansion factor of T .

Theorem 6.6. *Suppose T is a singly edge-to-edge, primitive self-similar tiling with FLC. Let T have substitution ω and prototile set \mathcal{P} . Then T has a recurrent pair (G, S) satisfying the conditions of Proposition 6.2, and hence T has a fractal realization.*

The proof of the Theorem 6.6 is an algorithm for constructing the recurrent pair (G, S) . The algorithm begins with Lemma 6.7 followed by a finite number of iterations of Lemma 6.8. Lemmas 6.7 and 6.8 are implicitly assumed to have the assumptions of Theorem 6.6.

Recall that we say that a graph G satisfies the injectivity conditions for N -subtiles if conditions (II)–(IV) of Definition 6.1 are satisfied with respect to the subtiles $\mathcal{R}^N(\mathcal{P})$. Also recall that we say a graph G_0 is T -consistent if it satisfies Definition 3.6.

Lemma 6.7. *Let G_0 be a T -consistent dual graph in \mathcal{P} . Then there exists $N \in \mathbb{N}$ and a T -consistent quasi-dual graph $G_1 \subset \mathcal{R}^N(G_0)$ such that G_1 satisfies the injectivity conditions for N -subtiles.*

Proof. By primitivity, there exists $n \in \mathbb{N}$ such that for each $p \in \mathcal{P}$, there is a copy of p interior to $\mathcal{R}^n(\mathcal{P})_p$. More precisely, there exists $x \in \mathbb{R}^2$ such that the n -subtile $p_0 = \lambda^{-n}p + x$ has support contained in the interior of $\mathcal{R}^n(\mathcal{P})_p$.

Now fix $p \in \mathcal{P}$ and let V_p denote the vertices of p . Since $p_0 \subset \text{int}(p)$, there exists a geometric graph $H_p \subset p \setminus \text{int}(p_0)$ consisting of $|V_p|$ disjoint edges, where the edge associated to $v \in V_p$ has endpoints $v \in p$ and $v^0 = \lambda^{-n}v + x \in p_0$. Let V_p^0 denote the collection of such vertices. Note that the geometric graph H_p connects the vertices of p with their corresponding vertices p_0 , and the existence of H_p is guaranteed since $p \setminus p_0$ is an annulus. Let H be the disjoint union

$$H := \bigsqcup_{p \in \mathcal{P}} H_p.$$

Let $M \in \mathbb{N}$ be large enough so that for each $p \in \mathcal{P}$ there are at least three M -subtile edges along each edge of p and three M -subtile edges along each edge of p_0 . Now let $E(H_p)$ denote the set of edges of H_p . Set ε to be less than the minimum of

$$\varepsilon_1 = \min_{p \in \mathcal{P}} \{d(e, f) \mid e \neq f \in E(H_p)\},$$

and

$$\varepsilon_2 = \min_{p \in \mathcal{P}} \{d(e, f) \mid e \in E(H_p), f \in \partial p \setminus \partial([V_p \cup V_p^0]^{\mathcal{R}^M(\mathcal{P})_p})\}$$

Then since $\lambda > 1$, we can choose $N \geq M$ such that $\lambda^{-N} \max_{p \in \mathcal{P}} \text{diam } p < \varepsilon/3$.

Let

$$O_B = \bigcup_{p \in \mathcal{P}} \{\partial p \setminus \partial([V_p]^{\mathcal{R}^M(\mathcal{P})_p})\} \quad \text{and} \quad I_B = \partial p_0 \setminus \partial([V_p^0]^{\mathcal{R}^M(\mathcal{P})_p}).$$

Choose a T -consistent set of vertices contained in the set $O_B \cap \mathcal{R}^N(G_0)$ with exactly one vertex for each edge $e \in \mathcal{P}$, and denote the set of vertices by O_V . By our choice of N with respect to $\varepsilon > 0$, for every edge $e \in \mathcal{P}$ there is at least one path in $\mathcal{R}^N(G_0)$ connecting O_V with any vertex in $I_B \cap \mathcal{R}^N(G_0)$ that does not cross any edge of H_p . We choose one such path, which we denote x_e , with the additional restriction that once it enters an interior N -subtile it continues through interior subtiles until it reaches its destination on the boundary of p_0 . Note that x_e must connect the edge $e \in E(p)$ to its counterpart $\lambda^{-n}e + x$ in p_0 , and $[x_e]^{\mathcal{R}^N(\mathcal{P})_p}$ must intersect the boundary of p in a connected set. Moreover, for $e \neq f \in p$, we have $[x_e]^{\mathcal{R}^N(\mathcal{P})_p} \cap [x_f]^{\mathcal{R}^N(\mathcal{P})_p} = \emptyset$. Let $X_p = \bigcup_{e \in p} x_e$. Although the component graphs of $X = \bigsqcup_{p \in \mathcal{P}} X_p$ are disconnected, X satisfies the injectivity conditions for N -subtiles.

In order to complete the proof, for each $p \in \mathcal{P}$, we need to connect the points $X_p \cap I_B \subset \partial p_0$ by a tree in p_0 . To that end, consider the connected graph $\mathcal{R}^N(G_0)_{p_0}$, which we can prune until we have a tree connecting the points $X_p \cap I_B$ with no dangling vertices. Let Y_p denote this tree.

Set $G_1 := \bigsqcup_{p \in \mathcal{P}} X_p \cup Y_p$. Then, by construction, the graph G_1 is quasi-dual and satisfies the injectivity conditions for N -subtiles. □

When $G' \subset \mathcal{R}^N(G)$ is a geometric graph we have been interested in the patches of N -subtiles intersecting the edges of G' . The next lemma states that the paths in $\mathcal{R}^N(G')$ running through those same patches are unique when G and G' are both quasi-dual.

Lemma 6.8. *Suppose that G and $G' \subset \mathcal{R}^N(G)$ are quasi-dual graphs in \mathcal{P} such that, for some $M \in \mathbb{N}$, G' satisfies the injectivity conditions for M -subtiles. Then there is a unique quasi-dual graph $H \subset \mathcal{R}^N(G')$ such that*

$$[G'_p]^{\mathcal{R}^N(\mathcal{P})_p} = [H_p]^{\mathcal{R}^N(\mathcal{P})_p} \quad \text{for all } p \in \mathcal{P},$$

which satisfies the injectivity conditions for $(M + N)$ -subtiles.

Proof. For each $p \in \mathcal{P}$, consider the patch of N -subtiles $Q_p := [G'_p]^{\mathcal{R}^N(\mathcal{P})_p}$ and boundary N -subedges $E_p := \{e \in \mathcal{R}^N(\mathcal{P})_p \cap \partial p \mid e \cap G' \neq \emptyset\}$. Since G' is quasi-dual and T is singly edge-to-edge there is a unique quasi-dual graph H_p with edges in $\mathcal{R}^N(G')_p$ through the patch Q_p with dangling vertices exactly on the edges E_p . Let $H := \bigsqcup_{p \in \mathcal{P}} H_p$. By construction, H is a quasi-dual graph satisfying $[G'_p]^{\mathcal{R}^N(\mathcal{P})_p} = [H_p]^{\mathcal{R}^N(\mathcal{P})_p}$ for all $p \in \mathcal{P}$.

Since G' satisfies the injectivity conditions for M -subtiles and $H \subset \mathcal{R}^N(G')$, it follows that H satisfies the injectivity conditions for $(M + N)$ -subtiles. \square

We can now prove Theorem 6.6.

Proof of Theorem 6.6. Let G_0 and G_1 be the quasi-dual graphs from Lemma 6.7. If $G_0 \sim G_1$, the proof is complete with $G = G_0$ and $S = G_1$. If not, apply Lemma 6.8 with $G_0 = G$ and $G_1 = G'$ to obtain $G_2 := H$.

If $G_1 \sim G_2$ then we are done with $G = G_1$ and $S = G_2$. If not, we continue iterating this process of applying Lemma 6.8 to the last two graphs in the sequence. We claim that there exists $n \geq 2$ such that $G_n \sim G_{n+1}$.

Since each prototile has only a finite number of edges and each graph G_j is quasi-dual, the number of interior vertices (with degree greater than two) is bounded by half the number of edges of the prototile. Suppose $G_i \not\sim G_{i+1}$, for $i \geq 2$, then there exists $p \in \mathcal{P}$ such that $(G_i)_p \not\sim (G_{i+1})_p$. Then the construction in Lemma 6.8 implies that the number of interior vertices of $(G_{i+1})_p$ is greater than the number of interior vertices of $(G_i)_p$. Since the number of interior vertices of each prototile is bounded, this process must end. So there exists $n \in \mathbb{N}$ such that $G_n \sim G_{n+1}$, and we let $G := G_n$ and $S := G_{n+1}$.

Lemma 6.7 shows that G_1 satisfies the injectivity conditions for N -subtiles. Lemma 6.8 implies that each graph G_i in the algorithm satisfies the injectivity conditions for iN -subtiles. Thus, every graph in the sequence continues to satisfy the injectivity conditions. In particular, $S \subset \mathcal{R}^N(G)$ satisfies the injectivity conditions for $(n + 1)N$ -subtiles. Moreover, since G_0 is a dual tiling and $G \subset \mathcal{R}^{nN}(G_0)$ and $S \subset \mathcal{R}^{(n+1)N}(G_0)$ the map $\tilde{\psi}$ can be chosen to be piecewise constant-speed. Thus, (G, S) satisfies the conditions of Proposition 6.2 and Theorem 6.6 is proved. \square

6.3. Self-similar combinatorial dual tilings. The construction in Section 6.2 invites the question of when the (G, S) recurrent pair are both dual graphs. We give a sufficient condition in the following result.

Theorem 6.9. *Let T be a singly edge-to-edge, primitive self-similar tiling whose tiles are all convex. Then the combinatorial dual tiling of T has a self-similar realization.*

Proof. Let G be a T -consistent dual graph \mathcal{P} , embedded such that G_p has exactly one boundary vertex in the interior of each edge of p and there is one interior vertex that is connected by a straight-line edge to each boundary vertex. By primitivity we can choose N such that

- (1) $\mathcal{R}^N(\mathcal{P})_p$ contains a tile of type p in its interior for each $p \in \mathcal{P}$ and
- (2) there is a T -consistent set of boundary vertices in $\mathcal{R}^N(G)$ such that if v is in this set and in $p \in \mathcal{P}$, then $\text{supp}[v]^{\mathcal{R}^N(\mathcal{P})_p}$ does not contain any vertex of p . Denote by V_B such a T -consistent set that has exactly one vertex per prototile edge.

We let the vertex set of S_p be $(V_B)_p$ along with a single interior vertex $v_{p,\text{int}}$ chosen from the interior of an N -subtile p_0 of type p lying in the interior of $\mathcal{R}^N(\mathcal{P})_p$. We construct the edges of S_p as follows. Let $l(v)$ be the line connecting the boundary vertex $v \in S_p$ to its counterpart in p_0 . We take as the edge in S_p a path in $\mathcal{R}^N(G)$ through $[l(v)]^{\mathcal{R}^N(\mathcal{P})_p}$ connecting v to $v_{p,\text{int}}$.

Notice that if $v \neq w$ are boundary vertices in p , by convexity the $\mathcal{R}^N(\mathcal{P})_p$ -patches intersecting $l(v)$ and $l(w)$ cannot intersect except at p_0 and perhaps on their boundaries, and thus (II) and (I2) are satisfied. G was already dual and thus quasi-dual and (I3) and (I4) follow from the construction and by convexity. The map $\tilde{\psi}$ can be chosen to be piecewise constant-speed so that all the conditions of Proposition 6.2 hold.

Thus we have a recurrent pair (G, S) for which $T_G = T^{(0)}$ is a labelled dual tiling of T . Since $T^{(\infty)}$ is combinatorially equivalent to T_G the result follows. \square

6.4. Applications and comments on the main result. In Theorem 6.6 we showed the existence of fractal realizations for a large class of tilings. We believe the restriction of singly edge-to-edge tilings can be lifted and that every tiling has a fractal dual tiling rather than a fractal quasi-dual tiling. In particular, we are fairly convinced that every two-dimensional tiling that appears in the Tiling Encyclopedia [31] has a fractal dual tiling. We also note that the proof of Theorem 6.6

implies that there are an infinite number of distinct fractal realizations for each tiling satisfying the hypotheses.

In [17, Section 4.2], Grunbaum and Shephard show the importance of dual tilings to the study of tilings. For instance, we can use [17, Theorem 4.2.1] to deduce that all fractal dual tilings $T^{(\infty)}$ obtained in Theorem 6.9 are homeomorphic (note that this is a very different statement than that their tiling spaces are mutually locally derivable). However, the usual combinatorial dual of a substitution tiling is not typically a substitution tiling itself. The combinatorial (quasi-) dual tilings we construct are substitution tilings, and hence we benefit from the dual construction while maintaining a substitution system. As a particular application, in Section 8 we show how this allows us to compute the cohomology of a tiling that doesn't force its border.

As mentioned in the introduction, in [23] spectral triples are constructed on Kellendonk's C^* -algebra associated with a substitution tiling. In the construction, the existence of a fractal realization corresponds to the existence of a fractal tree overlaying the original tiling that uniquely connects every pair of tiles in each self-similar tiling. Perron–Frobenius theory allows us to define a fractal distance function that scales appropriately between supertiles, which in turn defines the Dirac operator required in a spectral triple. Thus, each fractal realization provides a distinct noncommutative geometry on Kellendonk's algebra. In future work we hope to explore the implications of making different choices of recurrent pairs in a noncommutative framework.

7. Edge contraction map and Hausdorff dimension of $\partial T^{(\infty)}$

We mentioned earlier that it is possible to see the edge refinement map $\tilde{\psi}$ as a substitution rule on prototile supports. We make this precise here, using the formalism of digit sets and the tiling contraction map used to describe a substitution rule. For simplicity we assume that $S \subset \mathcal{R}(G)$ but all results extend to the general case of $S \subset \mathcal{R}^N(G)$.

Once this is done it is possible to see the edges of $T^{(\infty)}$, or more precisely their counterparts in \mathcal{P} , as the invariant set of a graph iterated function system. This system satisfies the strong open set condition, so [24] implies that the Hausdorff dimension of its attractor (our edges) is the Mauldin-Williams dimension. In this way we obtain a geometric invariant on the set of fractal realizations of a particular tiling.

7.1. Edge refinement as a contraction map \mathcal{R}^E . Suppose that (G, S) is a recurrent pair for a substitution ω with expansion λ .

Let the edge set of G , denoted $E(G)$, be the collection of all edges in G : $E(G) = \bigcup_{p \in \mathcal{P}} E(G_p)$. Analogous to the tiling contraction case, we assign a copy of $H(\mathbb{R}^2)$ for each $e \in E(G)$ and define $X^E = \bigsqcup_{e \in E(G)} H(\mathbb{R}^2)$. Consider a fixed $e \in E(G_p)$ and consider its counterpart $\tilde{\psi}(e) \subset S$. By definition we know that $\tilde{\psi}(e)$ is a subset of $\mathcal{R}(G)_p$ and thus is comprised of subedges of the form $\lambda^{-1}(f + d)$, where f is an edge in a prototile q and d is some element of the digit set D_{pq} . We define the matrix M^E by letting M_{ef}^E be the number of copies of f appearing in $\tilde{\psi}(e)$ for all $f \in E(G)$. This makes M^E a nonnegative integer matrix whose rows and columns are indexed by $E(G)$.

Suppose that f is an edge in the prototile q . The digit set D_{ef}^E is obtained by taking all of the digits in D_{pq} giving copies of f in $\tilde{\psi}(e)$:

$$D_{ef}^E = \{d \in D_{pq} \mid \tilde{\psi}(e) \cap \lambda^{-1}(f + d) = \lambda^{-1}(f + d)\} \tag{7.1}$$

For $B \in X^E$ and $e \in E(G)$ we define

$$\mathcal{R}^E(B)_e = \bigcup_{q \in \mathcal{P}} \bigcup_{f \in E(G_q)} \lambda^{-1}(B_f + D_{ef}^E) \tag{7.2}$$

With this definition we see that for $e \in E(G)$,

$$\tilde{\psi}(e) = \bigcup_{q \in \mathcal{P}} \bigcup_{f \in E(G_q)} \lambda^{-1}(f + D_{ef}^E) = \left(\mathcal{R}^E \left(\bigsqcup_{f \in E(G)} f \right) \right)_e.$$

We can iterate \mathcal{R}^E but not $\tilde{\psi}$, since the latter is only defined for edges $e \in E(G)$. However we can define $\tilde{\psi}^{(n)}(e)$ for $n \in \mathbb{N} \cup \{\infty\}$ using $\psi^{(n)}$: take any edge $e' \in T^{(0)}$ for which $e' = e + x$, then $\tilde{\psi}^{(n)}(e) = \psi^{(n)}(e') - x$. It is a tedious, but not difficult, chase through notation to verify that if $B = \bigsqcup_{f \in E(G)} f$, then

$$\tilde{\psi}^{(n)}(e) = ((\mathcal{R}^E)^n(B))_e$$

Let $A^{(\infty)} \subset X^E$ represent the attractor of \mathcal{R}^E , and let $S^{(\infty)} \subset X$ represent the canonical projection of $A^{(\infty)}$ onto the support of the prototiles:

$$S_p^{(\infty)} = \bigcup_{e \in E(G_p)} A_e^{(\infty)}.$$

Proposition 7.1. *Let (G, S) be a recurrent pair and T a self-similar tiling for the substitution ω . Then $\psi^{(\infty)}(\partial T^{(0)}) = \Gamma(T, S^{(\infty)})$, where Γ is defined in (3.1). If $\psi^{(\infty)}$ is injective, then $\partial T^{(\infty)} = \partial T_S^{(\infty)}$.*

The first part says that even if $\Gamma(T, S^{(\infty)})$ is not the boundary of a tiling, it still corresponds to $\psi^{(\infty)}(\partial T^{(0)})$. The second part does not quite extend to imply that $T_{S^{(\infty)}}$ and $T^{(\infty)}$ are actually equal, because there is the possibility that the tile labels differ slightly if the tiles of $T_{S^{(\infty)}}$ intersect larger T -patches than those that $T^{(0)}$ did, since, by definition, those endow the tiles of $T^{(\infty)}$ with their labels.

Proof. We must show that if $e \in E(G)$, then $\tilde{\psi}^{(\infty)}(e) = A_e^{(\infty)}$. We have

$$\begin{aligned} \bigsqcup_{e \in E(G)} \tilde{\psi}^{(\infty)}(e) &= \bigsqcup_{e \in E(G)} \left(\lim_{n \rightarrow \infty} \tilde{\psi}^{(n)}(e) \right) \\ &= \lim_{n \rightarrow \infty} \bigsqcup_{e \in E(G)} \tilde{\psi}^{(n)}(e) \\ &= \lim_{n \rightarrow \infty} \mathcal{R}_E^n \left(\bigsqcup_{e \in E(G)} e \right) \\ &= A^{(\infty)}, \end{aligned}$$

since $A^{(\infty)}$ is the attractor of \mathcal{R}_E . □

7.2. Hausdorff dimension. We will prove the following theorem.

Theorem 7.2. *Suppose that (G, S) is a recurrent pair for a tiling T with substitution ω having expansion factor λ which satisfies the injectivity conditions of Definition 6.1. Let M^E be the edge substitution matrix with largest eigenvalue λ_E . Then the Hausdorff dimension of $\partial T^{(\infty)}$ is $\frac{\ln \lambda_E}{\ln \lambda}$.*

It is important to notice that while M is always a primitive matrix, by assumption, it often happens that M^E is not. We note that if M^E is not primitive, then a power of M^E must be block triangular. Each block will contain classes of edges, and the Hausdorff dimension of each block is related to the leading eigenvalue of that block. Theorem 7.2 can be deduced almost directly from the definition of the Hausdorff dimension when M^E is primitive. In the interest of completeness we prove Theorem 7.2 by describing a fractal realization as the fixed point of a graph-directed iterated function system (GIFS), as defined in [14, 24]. In particular, [24] allows us to compute the Hausdorff dimension of the boundary of $T^{(\infty)}$ by computing the Mauldin–Williams dimension of the GIFS.

We begin by defining a directed combinatorial graph Δ that has one vertex, denoted v_e , for each edge $e \in E(G)$. The edge set $E(\Delta)$ contains an edge ϵ from v_e to v_f for each copy of f that appears as a subedge in $\tilde{\psi}(e)$. That is, there are M_{ef}^E edges pointing from v_e to v_f . To each vertex v_e we associate a copy of \mathbb{R}^2 .

To each edge $\epsilon \in E(\Delta)$ from e to f we assign a digit $d \in D_{ef}^E$ and construct the map $x \mapsto \lambda^{-1}(x + d)$ and call this map h_ϵ . These are the components necessary to define a GIFS; since the maps h_ϵ are taken directly from the definition of \mathcal{R}^E we refer to this as the GIFS given by \mathcal{R}^E .

Our fractal graph $A^{(\infty)}$ (as a vector of edges) is invariant for this GIFS in the sense that for each vertex $v_e \in V(\Delta)$, $A_e^{(\infty)} = \bigcup_{v_f \in V(\Delta)} \bigcup_{\epsilon \in E_{ef}} h_\epsilon(A_e^{(\infty)})$, where E_{ef} denotes edges from v_e to v_f .

The Hausdorff dimension of the invariant set of a GIFS has bounds that depend on the contraction factors of the maps h_ϵ . In our case, each contraction factor is exactly λ^{-1} and this leads to the lower and upper bounds being equal, and in fact $\ln \lambda_E / \ln \lambda$ as soon as we can establish that the GIFS satisfies the strong open set condition given here (and adapted to our notation.)

Definition 7.3 ([14], Definition 3.11). A GIFS with attractor set $A^{(\infty)}$ satisfies the *strong open set condition* if, for each $v_e \in V(\Delta)$, there exists an open set $U_e \subset \mathbb{R}^2$ satisfying:

- (1) for all vertices v_e, v_f and $\epsilon \in E_{ef}$, $h_\epsilon(U_f) \subset U_e$;
- (2) for all vertices v_e, v_f and $v_{f'}, \epsilon \in E_{ef}$ and $\epsilon' \in E_{ef'}$ with $\epsilon \neq \epsilon'$, we have $h_\epsilon(U_f) \cap h_{\epsilon'}(U_{f'}) = \emptyset$;
- (3) for each vertex v_e , $U_e \cap A_e^{(\infty)} \neq \emptyset$.

Proposition 7.4. *Suppose (G, S) is a recurrent pair satisfying the injectivity conditions (II)–(IV) of Definition 6.1. Then the GIFS given by \mathcal{R}^E satisfies the strong open set condition.*

Proof. Fix an edge $e \in E(G)$ in the prototile p ; e is associated to a unique vertex of our GIFS. Denote by $e^{(\infty)} = A_e^{(\infty)}$ its counterpart in the attractor and define U_e as follows. Let a_e be the set of interior vertices of $e^{(\infty)}$ that are endpoints of $e^{(\infty)}$ (i.e., endpoints of $e^{(\infty)}$ interior to the support of p ; by the injectivity conditions this has either one or two elements). Let $U_{e,n} = \text{int}([e^{(\infty)}]_{\mathcal{R}^n(\mathcal{P})_p} \setminus [a_e]_{\mathcal{R}^n(\mathcal{P})_p})$. Then set $U_e = \bigcup_{n \in \mathbb{N}} U_{e,n}$. Then the injectivity conditions (II)–(IV) then guarantee that these U_e satisfy conditions (1)–(3) of Definition 7.3. □

The *Mauldin–Williams dimension* [24] of the invariant set F of a GIFS is defined as follows. Suppose each h_ϵ has similarity ratio $r_\epsilon < 1$. Define a matrix $M^E(s)$ by $M_{ef}^E(s) = \sum_{\epsilon \in E_{ef}} r_\epsilon^s$, where e and f represent edges in G and therefore vertices in Δ . Let $\Phi(s)$ denote the spectral radius of $M^E(s)$. The Mauldin–Williams dimension $\dim_{MW}(F)$ of F is the unique s_1 for which $\Phi(s_1) = 1$.

Proof of Theorem 7.2. In [14, Section 3] it is shown that if $\|h_\epsilon(x) - h_\epsilon(y)\| = r_\epsilon \|x - y\|$ for each $\epsilon \in E_{ef}$ and x, y in the copy of \mathbb{R}^2 associated with v_f , and if the GIFS satisfies the strong open set condition, then the Hausdorff dimension $\dim_H A_e^{(\infty)} = \dim_{MW} (A_e^{(\infty)})$ for each vertex v_e . Our GIFS satisfies the strong open set condition by Proposition 7.4 and has $r_\epsilon = \lambda^{-1}$ for all ϵ . This means that $M^E(s) = \lambda^{-s} M^E$, where M^E is the edge substitution matrix given by (G, S) . This implies that $\Phi(s) = \lambda^{-s} \lambda_E$, where λ_E is the spectral radius of M^E . Solving $\lambda^{-s} \lambda_E = 1$ gives $s = \ln \lambda_E / \ln \lambda$. \square

8. Čech cohomology of a tiling space

In this section we relate our fractal tiling construction to the Anderson-Putnam complex and Čech cohomology. Given an nonperiodic and primitive substitution tiling along with a recurrent pair satisfying the injectivity conditions of Definition 6.1, we show that defining an orientation on the recurrent pair (G, S) gives rise to the substitution maps and boundary maps for the associated fractal realization. If the recurrent pair is quasi-dual then the fractal realization forces its border and is mutually locally derivable to the original substitution tiling. Putting this together, a recurrent pair (G, S) can be used to compute the cohomology of the original tiling. The upshot of our construction is that Čech cohomology is readily computable for tilings that do not force the border. We illustrate our method by computing the cohomology for the 2-dimensional Thue–Morse (2DTM) tiling. We begin with some background on substitution tiling spaces and their cohomology.

8.1. A brief description of the tiling space of a substitution. A substitution ω with expansion λ generates a topological space Ω of tilings often called the *hull*. This consists of all tilings T such that for every finite patch of tiles in T there exists a prototile $p \in \mathcal{P}$ and an N for which a copy of the patch appears inside $\omega^N(p)$. That is, every patch in T is admissible by the substitution.

The topology on Ω is generated by the ‘big ball’ metric, which measures how close two tilings are by how little they differ on big balls around the origin.⁵ We will not define it precisely here but refer readers to [29, p. 5] for details. It is possible to define Ω as the orbit closure of a fixed self-similar tiling T under the action of translation, and in this view Ω is often called the hull of T .

⁵ This is a continuous analogue to the standard metric used on shift spaces in symbolic dynamics.

The substitution ω extends to all tilings in Ω in a natural way. Given a tiling $T \in \Omega$ and a tile $t \in T$, we place the patch $\omega(t)$ atop $\lambda \operatorname{supp}(t)$. Doing this for all the tiles of T yields a tiling we call $\omega(T)$ which also lives in Ω . When the substitution on Ω is invertible we call the substitution *recognizable*. This means that in any location of any tiling T it is possible to determine the exact type and location of the supertile $\omega(t)$ at that spot using only local information. Since all the tilings considered in this paper are nonperiodic, [30] implies that they are recognizable.

8.2. Border-forcing. For topological analysis of tiling spaces, it is crucial that the substitution have the property of forcing its border. This means that whenever a supertile $\omega^N(t)$ appears in $T \in \Omega$, the patch $[\omega^N(t)]^T$ of tiles (which includes those immediately adjacent to $\omega^N(t)$ in T) are independent of T and appear in identical fashion anywhere else $\omega^N(t)$ appears in any tiling in Ω . Formally, the substitution forces the border if whenever t and t' are equivalent tiles in T , then $[\omega^N(t)]^T$ and $[\omega^N(t')]^T$ are equivalent patches.

When ω is border-forcing its tiling space is homeomorphic to the inverse limit of its approximants as discussed below, which is key to cohomology computations on Ω . The following theorem gives conditions for border-forcing that are sufficient but not necessary; it is often possible to tell by inspection whether a given recurrent pair will produce a substitution with the desired result.

Proposition 8.1. *Let (G, S) be a recurrent pair for a self-similar tiling T satisfying the conditions of Proposition 6.2. Then the substitution map $\omega_{(\infty)}$ for the tiling $T^{(\infty)}$ is border-forcing.*

Proof. The fact that G and S are quasi-dual and form a recurrent pair means that for any $t \in T^{(0)}$ and its counterpart $t^{(\infty)} \in T^{(\infty)}$ we have $[t]^T = [t^{(\infty)}]^T$, and that both t and $t^{(\infty)}$ are on the interior of the same patch of T -tiles. By definition, the equivalence class of this patch forms the label for both t and $t^{(\infty)}$. Let $k_1 = \max_{t \in T^{(0)}} \operatorname{diam}[t]^T$. Then every ball of radius $2k_1$ contains a patch of tiles in T which determine at least one tile in $T^{(\infty)}$.

There is a strictly positive minimum distance k_2 between the boundaries $\partial t^{(\infty)}$ and $\partial \operatorname{supp}[t^{(\infty)}]^T$, where the minimum is taken over all $t^{(\infty)} \in T^{(\infty)}$. Choose N such that $\lambda^N k_2 > 2k_1$. For this N , we know that $\omega^N([t^{(\infty)}]^T)$ determines all of the T -tiles on its interior, and this includes all of the T -tiles a distance of $2k_1$ or less from $\lambda^N(\operatorname{supp} t^{(\infty)})$. By choice of k_1 this patch determines $[\omega_{(\infty)}^N(t^{(\infty)})]^T$, as desired. \square

Corollary 8.2. *Suppose T is a self-similar tiling satisfying the conditions of Theorem 6.6 and let Ω denote its hull. Then there are infinitely many border-forcing substitutions $\omega_{(\infty)}$ such that $(\Omega^{(\infty)}, \omega_{(\infty)})$ is topologically conjugate to (Ω, ω) .*

Proof. The proof of Theorem 6.6 can be adapted to produce an infinite number of distinct recurrent pairs. In particular, Lemma 6.7 can be adapted by

- (1) using a quasi-dual graph in place of G_0 or
- (2) allowing N to increase, giving several distinct G_1 graphs for each N . \square

8.3. Čech cohomology of a substitution tiling. We begin with a very brief description of the Anderson-Putnam complex and Čech cohomology of a border forcing tiling space in two dimensions. We refer the reader to [1] or [27] for further details. Suppose Ω is the tiling space of a border forcing nonperiodic substitution tiling with finite local complexity and prototile set \mathcal{P} . A finite CW complex Γ , called the approximant, is defined by identifying edges and vertices in two prototiles whenever those edges or vertices are common in any two translates that occur in a patch in the tiling space. Anderson and Putnam [1] have shown that the tiling space Ω is the inverse limit of the approximant for tilings that force the border

$$\Omega = \varprojlim (\Gamma, f), \tag{8.1}$$

where f is the forgetful map described in [27, Section 2.5].

Since the approximant is a finite CW complex, we obtain a chain complex in simplicial homology. However, homology is not well-behaved with respect to inverse limits so we use cohomology instead, and consider the dual cochain complex. Let Γ^0 , Γ^1 and Γ^2 denote the functions from the vertices, edges, and tiles into the group of integers, respectively. There are coboundary maps $\delta_0: \Gamma^0 \rightarrow \Gamma^1$ and $\delta_1: \Gamma^1 \rightarrow \Gamma^2$ defined by taking the transpose of the matrix defining the homology boundary maps. We then have the following cochain complex:

$$0 \longrightarrow \Gamma^0 \xrightarrow{\delta_0} \Gamma^1 \xrightarrow{\delta_1} \Gamma^2 \longrightarrow 0.$$

From [27, Theorem 3.4], Čech cohomology and singular cohomology are equal on finite CW complexes. Thus we obtain the cohomology groups of the approximant Γ by

$$\check{H}^0(\Gamma) = \ker(\delta_0), \quad \check{H}^1(\Gamma) = \ker(\delta_1) / \text{Im}(\delta_0), \quad \check{H}^2(\Gamma) = \Gamma^2 / \text{Im}(\delta_1).$$

Since we have a substitution tiling there are substitution maps on vertices, edges, and tiles, which we denote by A_0 , A_1 , and A_2 respectively. Since cohomology is contravariant, the inverse limit appearing in (8.1) turns into direct limits on the cohomology groups of the approximant. Thus, the Čech cohomology of Ω is given by

$$\check{H}^i(\Omega) = \varinjlim(H^i(\Gamma), A_i^*) \quad (8.2)$$

where A_i^* denotes the map induced by the substitution on the singular cohomology group $H^i(\Gamma)$.

8.4. Cohomology from a recurrent pair (G, S) . We now relate the Anderson–Putnam complex to our situation in order to compute the cohomology of a tiling T using a fractal realization of T . In fact, it is not necessary to see the fractal at all: once a recurrent pair (G, S) satisfying the injectivity conditions is identified, it is possible to do all computations using the graphs G and S . We illustrate the construction in Section 8.5 by computing the Čech cohomology of the two dimensional Thue–Morse tiling using a recurrent pair (G, S) . We note that it may be useful to have the example in Section 8.5 at hand while reading this section.

Suppose T is a strongly nonperiodic tiling with finite local complexity and prototile set \mathcal{P} admitting a recurrent pair (G, S) satisfying the injectivity conditions of Definition 6.1. Let Ω be the tiling space associated with T . Using the construction in Section 5 we obtain a fractal tiling $T^{(\infty)}$ with tiling space $\Omega^{(\infty)}$ and prototile set $\mathcal{P}^{(\infty)}$. Proposition 5.7 implies that $T^{(\infty)}$ is a self-similar nonperiodic tiling and Proposition 8.1 implies that $T^{(\infty)}$ forces the border.

We will show that the Anderson–Putnam complex is completely determined by the recurrent pair (G, S) . First the approximant Γ of $\Omega^{(\infty)}$.

Let V_I denote the set of interior vertices of G and for each $v \in V_I$ let f_v denote the function that takes the value one on v and zero on $V_I \setminus v$. Then Γ^0 is generated by $\{f_v \mid v \in V_I\}$ and $\Gamma^0 \cong \mathbb{Z}^{|V_I|}$.

We now look at Γ^1 . Since we have assumed that (G, S) satisfies the conditions of Proposition 6.2, $\psi^{(\infty)}$ is a bijective map from $T^{(0)}$ to $T^{(\infty)}$, which induces a bijection taking edges in G to fractal edges whose translations make up $T^{(\infty)}$. More specifically, there are two types of edges involved in computing the cohomology of the approximant. The first are pairs of edges in G that form a single edge when we take the quotient by the equivalence relation \sim_T described in Section 3.2 (two boundary vertices are identified if they ever meet in any patch in T). We call pairs of these edges a (single) boundary edge. The second type are edges that connect two interior vertices of G within a prototile, and we call these edges interior edges. Let E be the union of the boundary edges and the interior edges.

For $e \in E$ let f_e denote the function that takes the value one on e and zero on $E \setminus e$. Then Γ^1 is generated by $\{f_e \mid e \in E\}$ and $\Gamma^1 \cong \mathbb{Z}^{|E|}$.

For Γ^2 , recall that the prototile set $\mathcal{P}^{(\infty)}$ is in bijective correspondence with distinct vertex patterns in T . For $p \in \mathcal{P}^{(\infty)}$, let f_p denote the function that takes the value one on p and zero on $\mathcal{P}^{(\infty)} \setminus p$. Then Γ^2 is generated by $\{f_p \mid p \in \mathcal{P}^{(\infty)}\}$ and $\Gamma^2 \cong \mathbb{Z}^{|\mathcal{P}^{(\infty)}|}$.

To compute the cohomology of the approximant it remains to determine the coboundary maps. We accomplish this by finding the boundary maps in homology and then taking adjoints. The first step is to assign an orientation to each edge in E (forming the dual space of Γ^1). As in [27, Section 3.2], let ∂_2 be the matrix with a row for each oriented edge in E and a column for each prototile in $\mathcal{P}^{(\infty)}$. For $e \in E$ and $p \in \mathcal{P}^{(\infty)}$, a $+1$ is added to the (e, p) entry if edge e appears as an edge in p in the same orientation and a -1 is added if it appears in the opposite orientation. The coboundary map $\delta_1: \Gamma^1 \rightarrow \Gamma^2$ is the transpose matrix $\delta_1 = \partial_2^t$. Let ∂_1 be the matrix with a row for each interior vertex V_I , and a column for each oriented edge in E . For each $e \in E$, place a $+1$ in the (v, e) entry if v is the range vertex of e and place a -1 in the (w, e) entry if w is the source vertex of e . The coboundary map $\delta_0: \Gamma^0 \rightarrow \Gamma^1$ is the transpose matrix $\delta_0 = \partial_1^t$. Putting all of this together leads to the cohomology groups of the approximants $\check{H}(\Gamma^i)$ for $i = 0, 1, 2$.

In order to compute the Čech cohomology of the tiling space $\Omega^{(\infty)}$ the final ingredients are the substitution maps. The matrix A_0 is the $|V_I| \times |V_I|$ substitution matrix on interior vertices, described entirely by the substitution of interior vertices from G to S . The matrix A_1 is an $|E| \times |E|$ matrix with a $+1$ added to the (e, f) entry if f appears in the substitution of e in the same orientation as e and a -1 if f appears in the opposite orientation, where the substitution is determined by the graph S . The matrix A_2 is a $|\mathcal{P}^{(\infty)}| \times |\mathcal{P}^{(\infty)}|$ matrix where entry (p, q) is the number of translations of prototile q appearing in the substitution of prototile p .

Computing the Čech cohomology of $\Omega^{(\infty)}$ is now an exercise in linear algebra. Lemma 5.6(3) implies that $\Omega^{(\infty)}$ is mutually locally derivable to Ω , and hence the Čech cohomology groups of the two tiling spaces are equal.

8.5. Cohomology of the two dimensional Thue–Morse tiling. We compute the cohomology of the 2-dimensional Thue–Morse (2DTM) tiling using the recurrent pair shown in Figure 2 of the introduction.

We label the edges of the graph G and define an orientation on these edges as in Figure 8 and extend the orientation to S . It is essential that the orientation be consistent across pairs of edges connected by boundary vertices, since they become edges in the AP complex of $T^{(\infty)}$.

Applying the machinery from Section 5 we obtain the tiles appearing in Figure 8 along with their substitution appearing in Figure 9.




Figure 7. The prototiles of the 2DTM tiling with an oriented dual graph G inscribed.

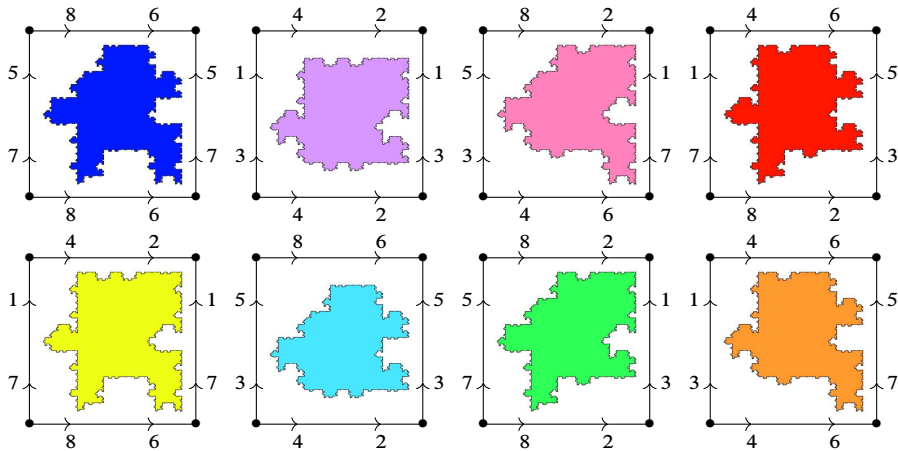


Figure 8. The 8 prototiles of the fractal realization of the 2DTM tiling from (G, S) . Each oriented (single) edge is labelled by a pair of edges from G as pictured in the bounding squares (which are only present to label the edges).

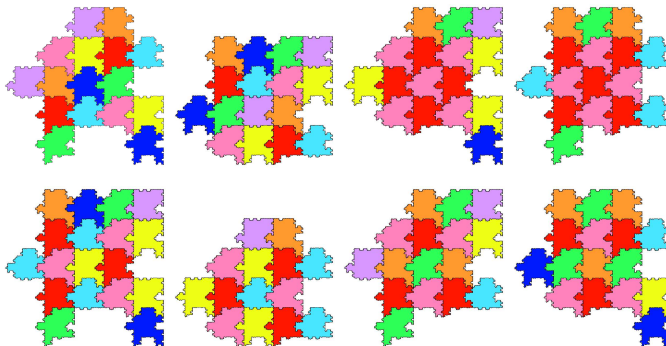


Figure 9. The substitutions of the 8 fractal prototiles (not to scale).

We begin by defining the matrices used to compute the cohomology $\check{H}^i(\Omega^{(\infty)})$. The fractal edges of the tiles are in bijective correspondence with pairs of edges in G whose endpoints are identified in S by boundary vertices. Since we will be interested in the substitution matrices of these edges, we label them in the ordered set

$$\{42, 46, 82, 86, 31, 35, 71, 75\}.$$

The pictures above give rise to the matrices:

$$A_0 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix},$$

$$\delta_1 = \begin{pmatrix} 0 & 0 \\ -1 & 1 \\ 1 & -1 \\ 0 & 0 \\ 0 & 0 \\ -1 & 1 \\ 1 & -1 \\ 0 & 0 \end{pmatrix}, \quad \delta_2 = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & -1 & 1 & 0 & 0 & 1 & -1 & 0 \\ 0 & 1 & -1 & 0 & 0 & -1 & 1 & 0 \\ 1 & 0 & 0 & -1 & 0 & 0 & 0 & 0 \\ -1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & -1 \end{pmatrix},$$

$$A_1 = \begin{pmatrix} 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 2 & 1 & 1 & 0 & 0 & 0 & -1 \\ 1 & 1 & 2 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 & 0 & -1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & -1 & 0 & 2 & 1 & 0 \\ -1 & 0 & 0 & 1 & 1 & 1 & 2 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \end{pmatrix}, \quad A_2 = \begin{pmatrix} 2 & 2 & 2 & 2 & 2 & 2 & 2 & 2 \\ 2 & 2 & 2 & 2 & 2 & 2 & 2 & 2 \\ 1 & 1 & 5 & 4 & 3 & 0 & 1 & 1 \\ 0 & 0 & 4 & 5 & 0 & 3 & 2 & 2 \\ 2 & 1 & 3 & 3 & 3 & 3 & 2 & 1 \\ 0 & 1 & 3 & 3 & 3 & 3 & 0 & 1 \\ 0 & 2 & 3 & 3 & 1 & 1 & 3 & 3 \\ 2 & 0 & 3 & 3 & 1 & 1 & 3 & 3 \end{pmatrix}.$$

The zeroth cohomology of the approximant Γ is generated by $\ker(\delta_0) = (11)^t$, which is viewed as the function assigning the value 1 to each vertex in Γ . So

$$\check{H}^0(\Gamma) \cong \mathbb{Z}.$$

Since A_0 is the identity matrix, it follows that

$$\check{H}^0(\Omega^{(\infty)}) \cong \mathbb{Z}$$

as well.

The first cohomology of the approximant Γ is given by

$$\check{H}^1(\Gamma) = \ker(\delta_1) / \text{Im}(\delta_0),$$

and routine linear algebra shows that

$$\check{H}^1(\Gamma) = \text{span} \left\{ \begin{pmatrix} 1 \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \end{pmatrix} \right\} \cong \mathbb{Z}^4.$$

The induced matrix A_1^* on $\check{H}^1(\Gamma)$ has eigenvalues $\{4, 4, 1, 1\}$ and eigenvectors $\text{EV}(A_1^*)$:

$$A_1^* = \begin{pmatrix} 2 & 1 & 0 & 0 \\ 2 & 3 & 0 & 0 \\ 0 & -1 & 2 & 2 \\ 1 & 0 & 1 & 3 \end{pmatrix}, \quad \text{EV}(A_1^*) = \left\{ \begin{pmatrix} 1 \\ 2 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ -2 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} -2 \\ 2 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ 1 \\ 1 \\ 0 \end{pmatrix} \right\}.$$

Thus, we have $\check{H}^1(\Omega^{(\infty)}) \cong \mathbb{Z}[1/4]^2 \oplus \mathbb{Z}^2$.

Finally, $\check{H}^2(\Gamma) = \mathbb{Z}^8 / \text{Im}(\delta_1)$ and we see that

$$\check{H}^2(\Gamma) = \text{span} \left\{ \begin{pmatrix} 1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \\ 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} \right\}.$$

The induced matrix A_2^* on $\check{H}^2(\Gamma)$ has nonzero eigenvalues $\{16, 4, 4, 1\}$ with eigenvectors $\text{EV}(A_2^*)$:

$$A_2^* = \begin{pmatrix} 4 & 1 & 2 & 2 & 0 \\ 4 & 9 & 6 & 6 & 0 \\ 4 & 3 & 6 & 2 & 0 \\ 4 & 3 & 2 & 6 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}, \quad \text{EV}(A_2^*) = \left\{ \begin{pmatrix} 1 \\ 4 \\ 2 \\ 2 \\ 0 \end{pmatrix}, \begin{pmatrix} 4 \\ 9 \\ 6 \\ 6 \\ 0 \end{pmatrix}, \begin{pmatrix} 4 \\ 3 \\ 6 \\ 2 \\ 0 \end{pmatrix}, \begin{pmatrix} 4 \\ 3 \\ 2 \\ 6 \\ 0 \end{pmatrix} \right\}.$$

Thus, we have $\check{H}^2(\Omega^{(\infty)}) \cong \mathbb{Z}[1/16] \oplus \mathbb{Z}[1/4]^2 \oplus \mathbb{Z}$.

Appendix A. Examples

A.1. A fractal dual for the Ammann–Beenker (“octagonal”) tiling. We consider the version of this tiling using two labelled right triangles and a parallelogram, where the labels keep track of the handedness of the tiles. These tiles come in rotations by $\pi/4$, all of which we treat the same way. In Figure 10 we show the parallelogram and one copy of the triangle; all rotations and reflections will carry along geometric graphs. Denoting $\mathcal{P} = \{\alpha, \beta\}$ we define G to be a T -consistent dual graph in both α and β . We show $G, \mathcal{R}(\mathcal{P}), \mathcal{R}(G)$, and select S as in Figure 10. The fractal realization of the Ammann–Beenker Tiling induced by (G, S) is shown in Figure 11.

In Figure 12 we show the substitution rule for the prototiles of the fractal realization. We note that the 4 different colours (labels) of the tiles with 8-fold symmetry come from the distinct vertex patterns in the original octagonal tiling making up each fractile. As in the proof of Lemma 3.8 these vertex patterns form the label set \mathcal{L} for the prototiles of the fractal realization.

A.2. Penrose’s “pentaplexity” tiling [25]. The simplest self-similar version of the Penrose tilings has a prototile set with forty triangles, the two tiles on the left of Figure 13, their labelled reflections, and all rotation by $\pi/5$ of these four prototiles. These triangles can be combined to produce either the kite and dart or the rhombus tilings, and all three prototile sets have matching rules that make them aperiodic tile sets. The kite/dart and rhombus versions are only pseudo-self-similar and not actually self-similar, so we’ve chosen to work with the triangle version instead. One interesting note about the fractal tiling we obtain is that it is closely related to the pentaplexity version of the Penrose tiling that arises in the first six images of [25].

We choose an initial graph G on the two prototiles as pictured in Figure 13 along with a choice of S that produces a recurrent pair. We then extend these graphs to all forty prototiles by reflection and rotation. What we gain in simplicity by working with one iteration of the substitution we pay for with an inability to construct S in a way that satisfies the injectivity conditions for any level of N -subtiles. The recurrent pair (G, S) does not give rise to an injective $\psi^{(\infty)}$ because one edge in each prototile is collapsed into a vertex, so Theorem 5.7 does not apply. Nonetheless $\psi^{(\infty)}(\partial T^{(0)})$ forms the boundary of a self-similar tiling $T^{(\infty)}$, a patch of which is shown in Figure 14.

There are five tile types in this fractal Penrose prototile set (ignoring rotations), some of which are shown along with their substitutions in Figure 15. The connection to the tiles developed by Penrose at the beginning of [25] were discussed in a 2013 talk [12].




Figure 10. A recurrent pair and resulting fractal for the dual graph of the Ammann–Beenker tiling.

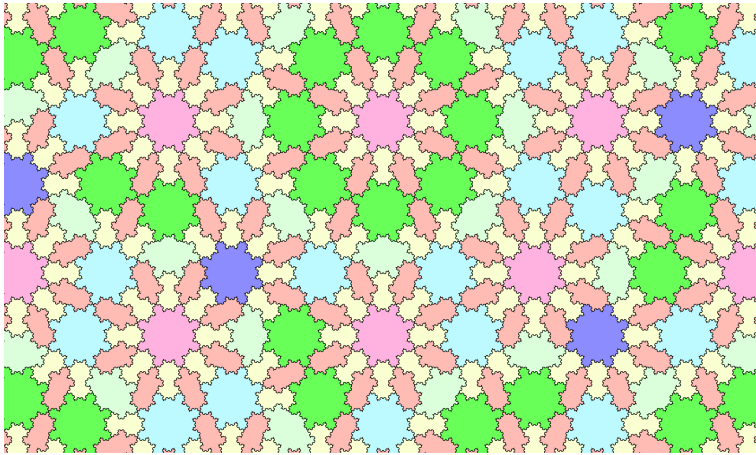


Figure 11. A patch of a fractal dual Ammann–Beenker tiling.




Figure 12. The substitution rules of this fractal realization of the Ammann–Beenker octagonal tiling.

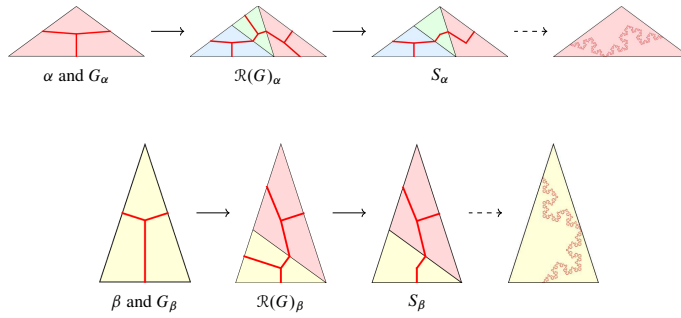


Figure 13. A recurrent pair for the Penrose triangles; note that each triangle come in two reflections and so do their geometric graphs.

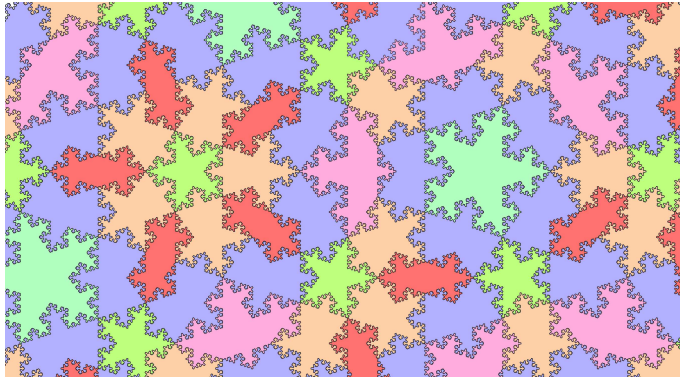


Figure 14. A patch of the fractal pentaplexity tiling.

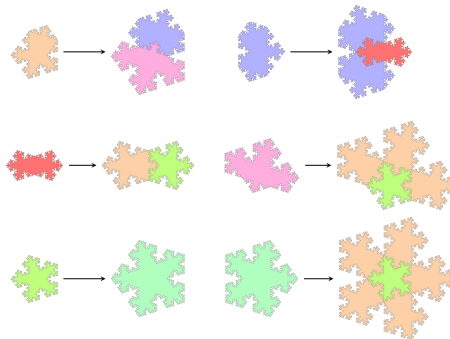


Figure 15. A fractal version of the pentaplexity substitution.

A.3. Two ways to construct fractal realizations of the chair tiling. If we endow the chair tiles with vertices from a tiling T as described in section 3.2, then each tile is an octagon and the tiling is edge-to-edge but not singly edge-to-edge. Thus the existence theorem does not directly apply. Nonetheless it is possible to use the existence algorithm in Section 6.2 to establish numerous fractal realizations by quasi-dual graphs. Interestingly, none of these quasi-dual graphs can actually be duals.

Another approach is to label the tiles of T by the number of natural edges they have. It turns out that there are three types of tiles in this case and they have 4, 5, and 6 edges. We can expand our prototile set to have twelve labels rather than four (once rotations are accounted for) and it is routine to write down the tiling substitution for each of the three tile types. We exhibit a self-similar fractal realization for the dual tiling in this situation.

There are many other possible ways to manage the chair tiling, for instance by passing to the “square chair” version [27, p. 16], which subdivides each chair tile into three labelled squares. The square chair is singly edge-to-edge so that fractal realizations are guaranteed by Theorem 6.6. This is an MLD operation and, as above, yields equivalent hulls, but the geometry and combinatorics are not as well preserved.

Example A.1 (chairs as octagons). We consider there to be four chair tiles, each octagons, with natural and inherited vertices as shown in Figure 5. As usual, we treat all rotations the same so our figures contain only one tile.

We begin by noticing the simple reason why the dual tiling can never have a self-similar realization. The dual graph G_0 has, on each prototile, a single interior vertex of degree 8. Trying to find a subgraph S of $\mathcal{R}^N(G_0)$ that is equivalent to G_0 is doomed to failure: no matter which vertex of degree 8 is selected from the interior, there will be two edges emanating from that vertex that come back together in an adjacent subtile and cannot therefore both be part of a tree. This illustrates how essential the singly edge-to-edge condition is to our construction.

It is possible to run the algorithm guaranteeing existence in Theorem 6.6 to obtain a quasi-dual recurrent pair for the octagonal chair tiling even though the tiling is not singly edge-to-edge. The ultimate choice of G has a minimum of two vertices on the interior of the prototile and the minimum number of iterations required to satisfy the injectivity conditions is 3. Instead, we choose to exhibit a recurrent pair that is not quasi-dual and whose limiting fractal intersects prototile vertices but for which $\psi^{(\infty)}$ is injective nonetheless.

The initial graph G has 9 edges and two interior vertices, and two of the boundary vertices have degree 2; we show this graph on the left of Figure 16. Conve-

niently, there is a subgraph of $\mathcal{R}(G)_\alpha$ that is isomorphic to G which we extract as shown in the same figure.




Figure 16. A recurrent pair and limiting fractal for the chairs as octagons.

On the far right of Figure 16 we see the result of iterating the graph substitution induced by this recurrent pair. The central edge has migrated to the boundary of the prototile and touches one of its vertices. The fact that there is no problem with the injectivity of $\psi^{(\infty)}$ can be seen by looking at the patch of fractiles superimposed on chairs shown in Figure 17.

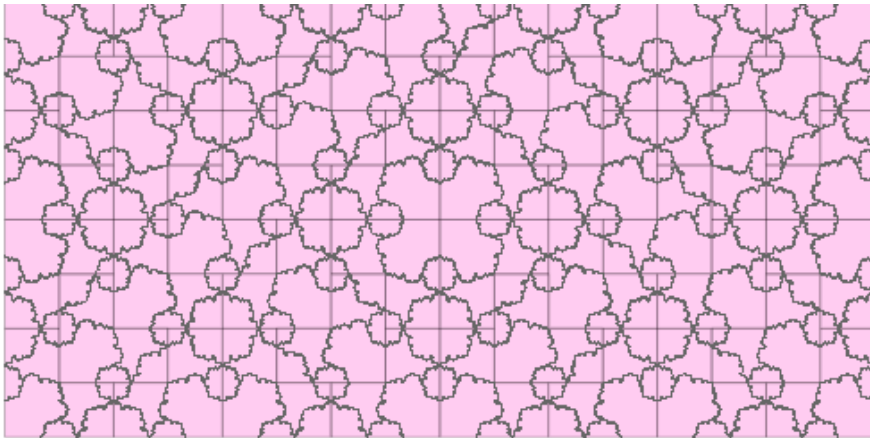


Figure 17. A patch of the octagonal chair tiling with a fractal realization overlaid.

Example A.2 (Colored chairs). Figure 18 shows the three types of chairs along with their dual graphs, embedded in their final fractal form. Although the graph edges appear to intersect prototile vertices, they do not: the natural vertices of these tile types do not include the graph vertices as they appear in the tiling.




Figure 18. The three chair tiles and their fractal dual graphs.

The substitutions of the three tile types appear in Figure 19. It should be noted that one substitution is not enough to find a recurrent pair for the dual substitution, however: none of the three graphs in $\mathcal{R}(G)$ contain combinatorially equivalent subgraphs. Thus, we used $\mathcal{R}^2(G)$ to produce the graph S for our recurrent pair.




Figure 19. Substitution of the three chairs (not to scale).

In Figure 20 we show a patch of the chair with the self-similar dual tiling overlaid atop it.

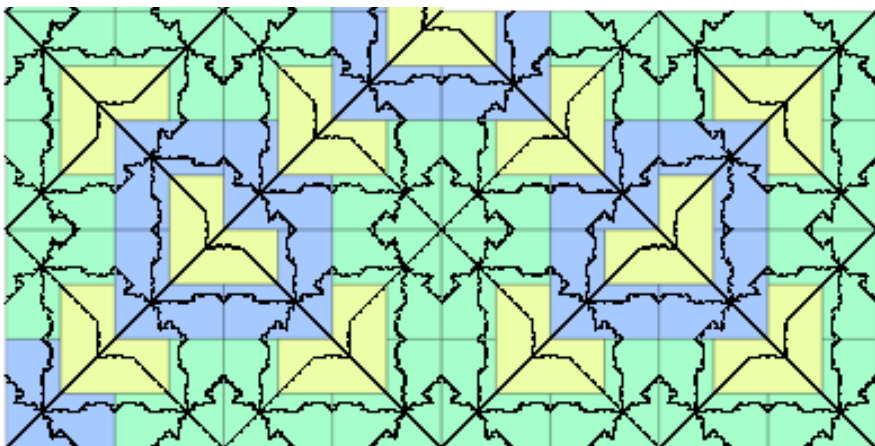


Figure 20. A patch of the chair tiling with its fractal dual overlaid.

References

- [1] J. Anderson and I. F. Putnam, Topological invariants for substitution tilings and their C^* -algebras. *Ergodic Theory Dynam. Systems* **18** (1998), no. 3, 509–537. [MR 1631708](#) [Zbl 1053.46520](#)
- [2] P. Arnoux, M. Furukado, E. Harriss, and S. Ito, Algebraic numbers, free group automorphisms and substitutions on the plane. *Trans. Amer. Math. Soc.* **363** (2011), no. 9, 4651–4699. [MR 2806687](#) [Zbl 1254.37015](#)
- [3] M. Baake and U. Grimm, *Aperiodic order*. A mathematical invitation. With a foreword by Roger Penrose. Encyclopedia of Mathematics and its Applications, 149. Cambridge University Press, Cambridge, 2013. [MR 3136260](#) [Zbl 1295.37001](#)
- [4] C. Bandt and P. Gummelt, Fractal Penrose tilings I. Construction and matching rules. *Aequationes Math.* **53** (1997), no. 3, 295–307. [MR 1444180](#) [Zbl 0882.52008](#)
- [5] M. Barge, B. Diamond, J. Hunton, and L. Sadun, Cohomology of substitution tiling spaces. *Ergodic Theory Dynam. Systems* **30** (2010), no. 6, 1607–1627. [MR 2736888](#) [Zbl 1225.37021](#)
- [6] M. F. Barnsley, *Fractals everywhere*. Second edition. Revised with the assistance of and with a foreword by H. Rising, III. Academic Press Professional, Boston, MA, 1993. [MR 1231795](#) [Zbl 0784.58002](#)
- [7] M. F. Barnsley and A. Vince, Fractal tilings from iterated function systems. *Discrete Comput. Geom.* **51** (2014), no. 3, 729–752. [MR 3201253](#) [Zbl 06312360](#)
- [8] J. Bellissard, Gap labelling theorems for Schrödinger operators. In M. Waldschmidt, P. Moussa, J. M. Luck and C. Itzykson (eds.), *From number theory to physics*. Papers from the Meeting on Number Theory and Physics held in Les Houches, March 7–16, 1989. Springer, Berlin, 1992, 538–630. [MR 1221111](#) [MR 1221099](#) (collection) [Zbl 0833.47056](#) [Zbl 0784.00021](#) (collection)
- [9] J. Bellissard, R. Benedetti and J.-M. Gambaudo, Spaces of tilings, finite telescopic approximations, and gap-labeling. *Comm. Math. Phys.* **261** (2006), no. 1, 1–41. [MR 2193205](#) [Zbl 1116.46063](#)
- [10] M.-T. Benameur and H. Oyono-Oyono, Gap-labelling for quasi-crystals (proving a conjecture by J. Bellissard). In J.-M. Combes, J. Cuntz, G. A. Elliott, Gh. Nenciu, H. Siedentop and Ş. Strătilă (eds.), *Operator algebras and mathematical physics*. Proceedings of the conference held in Constanța, July 2–7, 2001. Theta, Bucharest, 2003, 11–22. [MR 2018220](#) [MR 2018218](#) (collection) [Zbl 1284.46057](#) [Zbl 1095.46001](#) (collection)
- [11] R. Berger, The undecidability of the domino problem. *Mem. Amer. Math. Soc.* **66**, Providence, 72 pp. [MR 0216954](#) [Zbl 0199.30802](#)
- [12] A. Brunell, N. P. Frank, and D. Sherwood, Pentaplexity: Comparing fractal tilings and Penrose tilings. Paper presented at the Hudson River Undergraduate Mathematics Conference, Williams College, April 2013.
- [13] A. Connes, *Noncommutative geometry*. Transl. from the French by S. Berberian academic Press, San Diego, CA, 1994. [MR 1303779](#) [Zbl 0818.46076](#)

- [14] G. A. Edgar and J. Golds, A fractal dimension estimate for a graph-directed iterated function system of non-similarities. *Indiana Univ. Math. J.* 48 (1999), no. 2, 429–447. [MR 1722803](#) [Zbl 0944.28008](#)
- [15] N. P. Frank and M.F. Whittaker, A Fractal version of the Pinwheel tiling. *Math. Intelligencer* 33 (2011), no. 2, 7–17. [MR 2813257](#) [Zbl 1235.52030](#)
- [16] J. L. Gross and T. W. Tucker, *Topological graph theory*. Reprint of the 1987 original (John Wiley & Sons, New York), with a new preface and supplementary bibliography. Dover Publications, Mineola, N.Y., 2001 [MR 1855951](#) [Zbl 0991.05001](#)
- [17] B. Grunbaum and G.C. Shephard, *Tilings and Patterns*. W. H. Freeman and Company, New York, N.Y., 1987. [MR 0857454](#) [Zbl 0601.05001](#)
- [18] J. Kaminker and I. Putnam, A proof of the gap labeling conjecture. *Michigan Math. J.* 51 (2003), no. 3, 537–546. [MR 2018220](#) [Zbl 1054.46047](#)
- [19] J. Kellendonk, Noncommutative geometry of tilings and gap labelling. *Rev. Math. Phys.* 7 (1995), no. 7, 1133–1180. [MR 1359991](#) [Zbl 0847.52022](#)
- [20] R. Kenyon, The construction of self-similar tilings. *Geom. Funct. Anal.* 6 (1996), no. 3, 471–488. [MR 1392326](#) [Zbl 0866.52014](#)
- [21] J-Y. Lee and R.V. Moody, Lattice substitution systems and model sets. *Discrete Comput. Geom.* 25 (2001), no. 2, 173–201. [MR 1811757](#) [Zbl 0997.52013](#)
- [22] A. L. Mackay, Crystallography and the Penrose pattern. In L. L. Boyle and A. P. Cracknell. *Proceedings of the Tenth International Colloquium on Group-Theoretical Methods in Physics*. Held at the University of Kent, Canterbury, August 31–September 4, 1981. *Phys. A* 114 (1982), no. 1-3, 609–613. [MR 0678468](#) [MR 0678378](#) (collection) [Zbl 0509.00036](#) (collection)
- [23] M. Mampusti and M.F. Whittaker, Fractal spectral triples on Kellendonk’s C^* -algebra of a substitution tiling. Preprint 2015. [arXiv:1504.05296](#) [math.OA]
- [24] R. D. Mauldin and S. C. Williams, Hausdorff dimension in graph directed constructions. *Trans. Amer. Math. Soc.* 309 (1988), no. 2, 811–829. [MR 0961615](#) [Zbl 0706.28007](#)
- [25] R. Penrose, Pentaplexity: A class of non-periodic tilings of the plane. *Math. Intelligencer* 2 (1979/80), no. 1, 32–37. [MR 0558670](#) [Zbl 0426.52005](#)
- [26] N. Priebe and B. Solomyak, Characterization of planar pseudo-self-similar tilings. *Discrete Comput. Geom.* 26 (2001), no. 3, 289–306. [MR 1854103](#) [Zbl 0997.52012](#)
- [27] L. Sadun, *Topology of tiling spaces*. University Lecture Series, 46. American Mathematical Society, Providence, R.I., 2008. [MR 2446623](#) [Zbl 1166.52001](#)
- [28] D. Shechtman, Metallic phase with long-range orientational order and no translational symmetry. *Phys. Rev. Lett.* 53 (1984), 1951–1953.
- [29] B. Solomyak, Dynamics of self-similar tilings. *Ergodic Theory Dynam. Systems* 17 (1997), no. 3, 695–738; corrections, *Ergodic Theory Dynam. Systems* 19 (1999), no. 6, 1685. [MR 1452190](#) [MR 1738956](#) (corrections) [Zbl 0884.58062](#)

- [30] B. Solomyak, Nonperiodicity implies unique composition for self-similar translationally finite tilings. *Discrete Comput. Geom.* **20** (1998), no. 2, 265–279. [MR 1637896](#)
[Zbl 0919.52017](#)
- [31] The Tilings Encyclopedia, <http://tilings.math.uni-bielefeld.de/>
- [32] H. Wang, Proving theorems by pattern recognition, II. *Bell Sys. Tech. J.* **40** (1961), 1–41.

Received February 27, 2015; revised August 21, 2015

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

e-mail: nafrank@vassar.edu

Samuel B.G. Webster, School of Mathematics and Applied Statistics,
The University of Wollongong, NSW 2522, Australia

e-mail: sbgwebster@gmail.com

Michael F. Whittaker, School of Mathematics and Statistics, University of Glasgow,
University Gardens, Glasgow Q12 8QW, Scotland

e-mail: Mike.Whittaker@glasgow.ac.uk