Scaled pier fractals do not strictly self-assemble*

David Furcy† Scott M. Summers‡

Abstract

A *pier fractal* is a discrete self-similar fractal whose generator contains at least one *pier*, that is, a member of the generator with exactly one adjacent point. Tree fractals and pinch-point fractals are special cases of pier fractals. In this paper, we study *scaled pier fractals*, where a *scaled fractal* is the shape obtained by replacing each point in the original fractal by a $c \times c$ block of points, for some $c \in \mathbb{Z}^+$. We prove that no scaled discrete self-similar pier fractal strictly self-assembles, at any temperature, in Winfree's abstract Tile Assembly Model.

1 Introduction

The stunning, often mysterious complexities of the natural world, from nanoscale crystalline structures to unthinkably massive galaxies, all arise from the same elemental process known as *self-assembly*. In the absence of a mathematically rigorous definition, self-assembly is colloquially thought of as the process through which simple, unorganized components spontaneously combine, according to local interaction rules, to form some kind of organized final structure. A major objective of nanotechnology is to harness the power of self-assembly, perhaps for the purpose of engineering atomically precise medical, digital and mechanical components at the nanoscale. One strategy for doing so, developed by Nadrian Seeman, is *DNA tile self-assembly* [8, 9].

In DNA tile self-assembly, the fundamental components are "tiles", which are comprised of interconnected DNA strands. Remarkably, these DNA tiles can be "programmed", via the careful configuration of their constituent DNA strands, to automatically coalesce into a desired target structure, the characteristics of which are completely determined by the "programming" of the DNA tiles. In order to fully realize the power of DNA tile self-assembly, we must study the algorithmic and mathematical underpinnings of tile self-assembly.

Perhaps the simplest mathematical model of algorithmic tile self-assembly is Erik Winfree's abstract Tile Assembly Model (aTAM) [11]. The aTAM is a deliberately over-simplified, combinatorial model of nanoscale (DNA) tile self-assembly that "effectivizes" classical Wang tiling [10] in the sense that the former augments the latter with a mechanism for sequential "growth" of a tile assembly. Very briefly, in the aTAM, the fundamental components are un-rotatable, translatable square "tile types" whose sides are labeled with (alpha-numeric) glue "colors" and (integer) "strengths". Two tiles that are placed next to each other *bind* if the glues on their abutting sides match in both color and strength, and the common strength is at least a certain (integer) "temperature". Self-assembly starts from a "seed" tile type, typically assumed to be placed at the origin of the coordinate system, and proceeds nondeterministically and asynchronously as tiles bind to the seed-containing assembly one at a time.

Despite its deliberate over-simplification, the aTAM is a computationally expressive model. For example, Winfree [11] proved that the model is Turing universal, which means that, in principle, the process of self-assembly can be directed by any algorithm. In this paper, we study the extent to which tile sets in the aTAM can be algorithmically directed to "strictly" self-assemble (i.e., place tiles at and only at locations that belong to) shapes that are self-similar "pier fractals".

Intuitively, a "pier fractal" is a just-barely connected, self-similar fractal that contains the origin, as well as infinitely many, arbitrarily-large subsets of specially-positioned points that either lie at or "on the far side" from the

*This work is an extension of the conference paper titled, "Scaled tree fractals do not strictly self-assemble" [1].

†Department of Computer Science University of Wisconsin–Oshkosh, Oshkosh, WI 54901, USA. furcyd@uwosh.edu
‡Department of Computer Science University of Wisconsin–Oshkosh, Oshkosh, WI 54901, USA. summerss@uwosh.edu

origin of a pinch-point location (we make this notion precise in Section 2.2). Note that "tree" and "pinch-point" fractals constitute notable, previously-studied sub-classes of pier fractals (e.g., see [3, 4, 6] for definitions).

There are examples of prior results related to the strict self-assembly of fractals in the aTAM. For example, Theorem 3.2 of [3] bounds from below the size of the smallest tile set in which an arbitrary shape X strictly self-assembles, by the depth of X's largest finite sub-tree. Although not stated explicitly, an immediate corollary of this result is that no tree-fractal strictly self-assembles in the aTAM. In [4], Lutz and Shutters prove that a notable example of a tree fractal, the Sierpinski triangle, does not even "approximately" strictly self-assemble, in the sense that the discrete fractal dimension (see [2]) of the symmetric difference of any set that strictly self-assembles and the Sierpinski triangle is at least that of the latter, which is approximately $\log_2 3$. Theorem 3.12 of [6], the only prior result related to (the impossibility of) the strict self-assembly of pinch-point fractals, is essentially a qualitative generalization of Theorem 3.2 of [3].

While the strict self-assembly of certain classes of fractals in the aTAM has been studied previously, *nothing* is known about the strict self-assembly in the aTAM of scaled-up versions of fractals, where "scaled-up" means that each point in the original shape is replaced by a $c \times c$ block of points, for some $c \in \mathbb{Z}^+$. After all, certain classes of fractals defined by intricate geometric properties, such as the existence of "pinch-points" or "tree-ness", are not closed under the scaling operation. To see this, consider the full connectivity graph of any shape in which each point in the shape is represented by one vertex and edges exist between vertices that represent adjacent points in the shape. If this graph is a tree and/or contains one or more pinch-points, then the scaled-up version of the original shape (with $c > 1$) is not a tree and does not contain any pinch points. This means that prior proof techniques that exploit similar subtle geometric sub-structures of fractals (e.g., [3, 4, 6]) simply cannot be applied to scaled-up versions of fractals. Thus, in this paper, we ask if it is possible for a scaled-up version of a pier fractal to strictly self-assemble in the aTAM.

The main contribution of this paper provides an answer to the previous question, perhaps not too surprisingly to readers familiar with the aTAM, in the negative: we prove that there is no pier fractal that strictly self-assembles in the aTAM at any positive scale factor. Furthermore, our definition of pier fractal includes, as a strict subset, the set of all pinch-point fractals from [6]. Our proof makes crucial use of a (modified version of a) recent technical lemma developed by Meunier, Patitz, Summers, Theyssier, Winslow and Woods [5], known as the "Window Movie Lemma" (WML). This (standard) WML is a kind of pumping lemma for self-assembly since it gives a sufficient condition for taking any pair of tile assemblies, at any temperature, and "splicing" them together to create a new valid tile assembly. Our modified version of the WML, which we call the "Closed Window Movie Lemma" (see Section 2.3 for a formal statement and proof), allows one to replace a portion of a tile assembly with another portion of the same assembly, assuming a certain extra "containment" condition is met. Moreover, unlike in the standard WML that lacks the extra containment assumptions, the replacement of one part of the tile assembly with another in our Closed WML only goes "one way", i.e., the part of the tile assembly being used to replace another part cannot itself be replaced by the part of the tile assembly it is replacing.

2 Definitions

In this section, we give a formal definition of Erik Winfree's abstract Tile Assembly Model (aTAM), define "pier fractals" and develop a "Closed" Window Movie Lemma.

2.1 Formal description of the abstract Tile Assembly Model

This section gives a formal definition of the abstract Tile Assembly Model (aTAM) [11]. For readers unfamiliar with the aTAM, [7] gives an excellent introduction to the model.

Fix an alphabet Σ. Σ^* is the set of finite strings over Σ. Let $\mathbb{Z}$, $\mathbb{Z}^+$, and $\mathbb{N}$ denote the set of integers, positive integers, and nonnegative integers, respectively. Given $V \subseteq \mathbb{Z}^2$, the *full grid graph* of V is the undirected graph $G^{\mathrm{f}}_V = (V, E)$, such that, for all $\vec{x}, \vec{y} \in V$, $\{\vec{x}, \vec{y}\} \in E \iff \|\vec{x} - \vec{y}\| = 1$, i.e., if and only if $\vec{x}$ and $\vec{y}$ are adjacent in the 2-dimensional integer Cartesian space.

A *tile type* is a tuple $t \in (\Sigma^* \times \mathbb{N})^4$, e.g., a unit square, with four sides, listed in some standardized order, and each side having a *glue* $g \in \Sigma^* \times \mathbb{N}$ consisting of a finite string *label* and a nonnegative integer *strength*.

We assume a finite set of tile types, but an infinite number of copies of each tile type, each copy referred to as a *tile*. A tile set is a set of tile types and is usually denoted as T.

A *configuration* is a (possibly empty) arrangement of tiles on the integer lattice $\mathbb{Z}^2$, i.e., a partial function $\alpha : \mathbb{Z}^2 \dashrightarrow T$. Two adjacent tiles in a configuration *interact*, or are *attached*, if the glues on their abutting sides are equal (in both label and strength) and have positive strength. Each configuration α induces a *binding graph* G^{b}_{α}, a grid graph whose vertices are positions occupied by tiles, according to α, with an edge between two vertices if the tiles at those vertices bind. An *assembly* is a connected, non-empty configuration, i.e., a partial function $\alpha : \mathbb{Z}^2 \dashrightarrow T$ such that $G^{\mathrm{f}}_{\mathrm{dom}\,\alpha}$ is connected and $\mathrm{dom}\,\alpha \neq \varnothing$.

Given $\tau \in \mathbb{Z}^+$, α is *τ-stable* if every cut-set of G^{b}_{α} has weight at least τ, where the weight of an edge is the strength of the glue it represents.[1] When τ is clear from context, we say α is *stable*. Given two assemblies α, β, we say α is a *subassembly* of β, and we write $\alpha \sqsubseteq \beta$, if $\mathrm{dom}\,\alpha \subseteq \mathrm{dom}\,\beta$ and, for all points $\vec{p} \in \mathrm{dom}\,\alpha$, $\alpha(\vec{p}) = \beta(\vec{p})$. For two non-overlapping assemblies α and β, $\alpha \cup \beta$ is defined as the unique assembly γ satisfying, for all $\vec{x} \in \mathrm{dom}\,\alpha$, $\gamma(\vec{x}) = \alpha(\vec{x})$, for all $\vec{x} \in \mathrm{dom}\,\beta$, $\gamma(\vec{x}) = \beta(\vec{x})$, and $\gamma(\vec{x})$ is undefined at any point $\vec{x} \in \mathbb{Z}^2 \setminus (\mathrm{dom}\,\alpha \cup \mathrm{dom}\,\beta)$.

A *tile assembly system* (TAS) is a triple $\mathcal{T} = (T, \sigma, \tau)$, where T is a tile set, $\sigma : \mathbb{Z}^2 \dashrightarrow T$ is the finite, τ-stable, *seed assembly*, and $\tau \in \mathbb{Z}^+$ is the *temperature*.

Given two τ-stable assemblies α, β, we write $\alpha \to_1^{\mathcal{T}} \beta$ if $\alpha \sqsubseteq \beta$ and $|\mathrm{dom}\,\beta \setminus \mathrm{dom}\,\alpha| = 1$. In this case we say α *$\mathcal{T}$-produces β in one step*. If $\alpha \to_1^{\mathcal{T}} \beta$, $\mathrm{dom}\,\beta \setminus \mathrm{dom}\,\alpha = \{\vec{p}\}$, and $t = \beta(\vec{p})$, we write $\beta = \alpha + (\vec{p} \mapsto t)$. The *$\mathcal{T}$-frontier* of α is the set $\partial^{\mathcal{T}} \alpha = \bigcup_{\alpha \to_1^{\mathcal{T}} \beta} (\mathrm{dom}\,\beta \setminus \mathrm{dom}\,\alpha)$, i.e., the set of empty locations at which a tile could stably attach to α. The *t-frontier* of α, denoted $\partial^{\mathcal{T}}_t \alpha$, is the subset of $\partial^{\mathcal{T}} \alpha$ defined as $\left\{\, \vec{p} \in \partial^{\mathcal{T}} \alpha \;\middle|\; \alpha \to_1^{\mathcal{T}} \beta \text{ and } \beta(\vec{p}) = t \,\right\}$.

Let $\mathcal{A}^T$ denote the set of all assemblies of tiles from T, and let $\mathcal{A}^T_{<\infty}$ denote the set of finite assemblies of tiles from T. A sequence of $k \in \mathbb{Z}^+ \cup \{\infty\}$ assemblies $\alpha_0, \alpha_1, \ldots$ over $\mathcal{A}^T$ is a *$\mathcal{T}$-assembly sequence* if, for all $1 \leq i < k$, $\alpha_{i-1} \to_1^{\mathcal{T}} \alpha_i$. The *result* of an assembly sequence $\vec{\alpha}$, denoted as $\mathrm{res}(\vec{\alpha})$, is the unique limiting assembly (for a finite sequence, this is the final assembly in the sequence).

We write $\alpha \to^{\mathcal{T}} \beta$, and we say α *$\mathcal{T}$-produces* β (in 0 or more steps), if there is a $\mathcal{T}$-assembly sequence $\alpha_0, \alpha_1, \ldots$ of length $k = |\mathrm{dom}\,\beta \setminus \mathrm{dom}\,\alpha| + 1$ such that (1) $\alpha = \alpha_0$, (2) $\mathrm{dom}\,\beta = \bigcup_{0 \leq i < k} \mathrm{dom}\,\alpha_i$, and (3) for all $0 \leq i < k$, $\alpha_i \sqsubseteq \beta$. If k is finite then it is routine to verify that $\beta = \alpha_{k-1}$.

We say α is *$\mathcal{T}$-producible* if $\sigma \to^{\mathcal{T}} \alpha$, and we write $\mathcal{A}[\mathcal{T}]$ to denote the set of $\mathcal{T}$-producible assemblies. The relation $\to^{\mathcal{T}}$ is a partial order on $\mathcal{A}[\mathcal{T}]$ [3,7].

An assembly α is *$\mathcal{T}$-terminal* if α is τ-stable and $\partial^{\mathcal{T}} \alpha = \varnothing$. We write $\mathcal{A}_{\Box}[\mathcal{T}] \subseteq \mathcal{A}[\mathcal{T}]$ to denote the set of $\mathcal{T}$-producible, $\mathcal{T}$-terminal assemblies. If $|\mathcal{A}_{\Box}[\mathcal{T}]| = 1$ then $\mathcal{T}$ is said to be *directed*.

We say that a TAS $\mathcal{T}$ *strictly (or uniquely) self-assembles* $X \subseteq \mathbb{Z}^2$ if, for all $\alpha \in \mathcal{A}_{\Box}[\mathcal{T}]$, $\mathrm{dom}\,\alpha = X$, i.e., if every terminal assembly produced by $\mathcal{T}$ places a tile on every point in X and does not place any tiles on points in $\mathbb{Z}^2 \setminus X$.

In this paper, we consider scaled-up versions of subsets of $\mathbb{Z}^2$. Formally, if X is a subset of $\mathbb{Z}^2$ and $c \in \mathbb{Z}^+$, then a *c-scaling* of X is defined as the set $X^c = \left\{ (x,y) \in \mathbb{Z}^2 \;\middle|\; \left(\left\lfloor \frac{x}{c} \right\rfloor, \left\lfloor \frac{y}{c} \right\rfloor \right) \in X \right\}$. Intuitively, X^c is the subset of $\mathbb{Z}^2$ obtained by replacing each point in X with a $c \times c$ block of points. We refer to the natural number c as the *scaling factor* or *resolution loss*.

2.2 Pier fractals

In this section, we first introduce some terminology and then define a class of fractals called "pier fractals" that is the focus of this paper.

Notation. We use $\mathbb{N}_g$ to denote the subset $\{0, \ldots, g-1\}$ of $\mathbb{N}$.

Notation. If A and B are subsets of $\mathbb{N}^2$ and $k \in \mathbb{N}$, then $A + kB = \{\vec{m} + k\vec{n} \mid \vec{m} \in A \text{ and } \vec{n} \in B\}$.

The following definition is a modification of Definition 2.11 in [6].

[1] A *cut-set* is a subset of edges in a graph which, when removed from the graph, produces two or more disconnected subgraphs. The *weight* of a cut-set is the sum of the weights of all of the edges in the cut-set.

Definition 2.1. Let $1 < g \in \mathbb{N}$ and $\mathbf{X} \subset \mathbb{N}^2$. We say that $\mathbf{X}$ is a *g-discrete self-similar fractal* (or *g-dssf* for short), if there is a set $\{(0,0)\} \subset G \subset \mathbb{N}_g^2$ with at least one point in every row and column, such that $\mathbf{X} = \bigcup_{i=1}^{\infty} X_i$, where X_i, the i^{th} *stage* of $\mathbf{X}$, is defined by $X_1 = G$ and $X_{i+1} = X_i \, + \, g^i G$. We say that G is the *generator* of $\mathbf{X}$.

Intuitively, a g-dssf is built as follows. Start by selecting points in $\mathbb{N}_g^2$ satisfying the constraints listed in Definition 2.1. This first stage of the fractal is the generator. Then, each subsequent stage of the fractal is obtained by adding a full copy of the previous stage for every point in the generator and translating these copies so that their relative positions are identical to the relative positions of the individual points in the gnerator.

Definition 2.2. Let S be any finite subset of $\mathbb{Z}^2$. Let l, r, b, and t denote the following integers:

$$l_S = \min_{(x,y)\in S} x \qquad r_S = \max_{(x,y)\in S} x \qquad b_S = \min_{(x,y)\in S} y \qquad t_S = \max_{(x,y)\in S} y$$

An *h-bridge* of S is any subset of S of the form $hb_S(y) = \{(l_S, y), (r_S, y)\}$. Similarly, a *v-bridge* of S is any subset of S of the form $vb_S(x) = \{(x, b_S), (x, t_S)\}$. We say that a bridge is *connected* if there is a simple path in S connecting the two bridge points.

Notation. Let S be any finite subset of $\mathbb{Z}^2$. We will denote by nhb_S and nvb_S, respectively, the number of h-bridges and the number of v-bridges of S.

Notation. The directions $\mathscr{D} = \{N, E, S, W\}$ will be used as functions from $\mathbb{Z}^2$ to $\mathbb{Z}^2$ such that $N(x,y) = (x, y+1)$, $E(x,y) = (x+1, y)$, $S(x,y) = (x, y-1)$ and $W(x,y) = (x-1, y)$. Note that $N^{-1} = S$ and $W^{-1} = E$.

Notation. Let $X \subseteq \mathbb{Z}^2$. We say that a point $(x,y) \in X$ is *D-free* in X, for some direction $D \in \mathscr{D}$, if $D(x,y) \notin X$.

Definition 2.3. Let G be the generator of any g-discrete self-similar fractal. A *pier* is a point in G that is D-free in G for exactly three of the four directions in $\mathscr{D}$. We say that a pier (x,y) is *D-pointing* (or *points D*) if $D^{-1}(x,y) \in G$. Note that a pier always points in exactly one direction.

Definition 2.4. Let G be the generator of any g-discrete self-similar fractal with exactly one h-bridge and one v-bridge. G may contain up to four distinct types of piers characterized by the number of bridges they belong to. Each pier may belong to no more than two bridges. A *real pier* is a pier that does not belong to any bridge in G. A *single-bridge pier* belongs to exactly one bridge. A *double-bridge pier* belongs to exactly two bridges. Finally, we will distinguish between two sub-types of single-bridge piers. If the pier is pointing in a direction that is parallel to the direction of the bridge (i.e., if the pier points north or south and belongs to a v-bridge, or the pier points east or west and belongs to an h-bridge), the pier is a *parallel single-bridge pier*. If the pier is pointing in a direction that is orthogonal to the direction of the bridge (i.e., if the pier points north or south and belongs to an h-bridge, or the pier points east or west and belongs to a v-bridge), the pier is an *orthogonal single-bridge pier*.

For example, the generator in Figure 1 below contains the h-bridge $\{(0,0),(4,0)\}$ and the v-bridge $\{(4,0),(4,4)\}$. The point $(1,4)$ is a real pier. The point $(0,0)$ is an orthogonal single-bridge pier. The point $(4,4)$ is a parallel single-bridge pier. The point $(4,0)$ is a double-bridge pier.

We are now ready to define the class of fractals that is the main focus of this paper.

Definition 2.5. $\mathbf{P}$ is a *pier fractal* if and only if $\mathbf{P}$ is a discrete self-similar fractal with generator G such that:

a. The full grid graph of G is connected, and
b. $nhb_G = nvb_G = 1$, and
c. G contains at least one pier.

2.3 The Closed Window Movie Lemma

In this subsection, we develop a more accommodating (modified) version of the standard Window Movie Lemma (WML) [5]. Our version of the WML, which we call the "Closed Window Movie Lemma", allows us to replace one portion of a tile assembly with another, assuming certain extra "containment" conditions are met. Moreover, unlike in the standard WML that lacks the extra containment assumptions, the replacement of a portion of one tile assembly

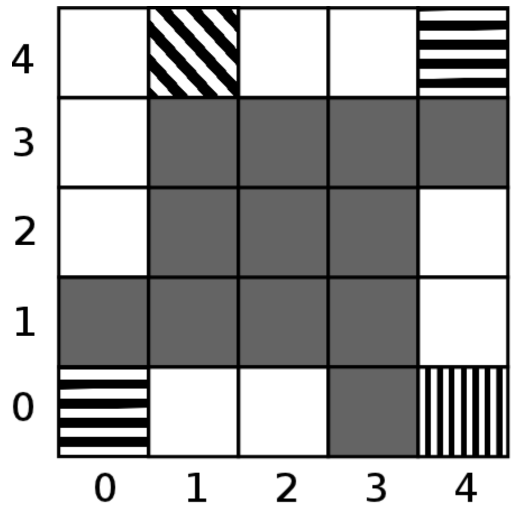

Figure 1: A 5×5 generator containing one h-bridge, one v-bridge and four piers.

with another portion of the same assembly in our Closed WML only goes "one way", i.e., the part of the tile assembly being used to replace another part cannot itself be replaced by the part of the tile assembly it is replacing. We must first define some notation that we will use in our Closed Window Movie Lemma.

A *window* w is a set of edges forming a cut-set of the full grid graph of $\mathbb{Z}^2$. For the purposes of this paper, we say that a *closed window* w induces a cut[2] of the full grid graph of $\mathbb{Z}^2$, written as $C_w = (C_{<\infty}, C_\infty)$, where C_∞ is infinite, $C_{<\infty}$ is finite and for all pairs of points $\vec{x}, \vec{y} \in C_{<\infty}$, no simple path connecting $\vec{x}$ and $\vec{y}$ in the full grid graph of $C_{<\infty}$ crosses the cut C_w. We call the set of vertices that make up $C_{<\infty}$ the *inside* of the window w, and write $inside(w) = C_{<\infty}$ and $outside(w) = \mathbb{Z}^2 \setminus inside(w) = C_\infty$. We say that a window w is *enclosed* in another window w' if $inside(w) \subseteq inside(w')$.

Given a window w and an assembly α, a window that *intersects* α is a partitioning of α into two configurations (i.e., after being split into two parts, each part may or may not be disconnected). In this case we say that the window w cuts the assembly α into two configurations α_L and α_R, where $\alpha = \alpha_L \cup \alpha_R$. If w is a closed window, for notational convenience, we write α_I for the configuration inside w and α_O for the configuration outside w. Given a window w, its translation by a vector $\vec{c}$, written $w + \vec{c}$ is simply the translation of each of w's elements (edges) by $\vec{c}$.

For a window w and an assembly sequence $\vec{\alpha}$, we define a window movie M to be the order of placement, position and glue type for each glue that appears along the window w in $\vec{\alpha}$. Given an assembly sequence $\vec{\alpha}$ and a window w, the associated *window movie* is the maximal sequence $M_{\vec{\alpha},w} = (v_0, g_0), (v_1, g_1), (v_2, g_2), \ldots$ of pairs of grid graph vertices v_i and glues g_i, given by the order of the appearance of the glues along window w in the assembly sequence $\vec{\alpha}$. Furthermore, if k glues appear along w at the same instant (this happens upon placement of a tile which has multiple sides touching w) then these k glues appear contiguously and are listed in lexicographical order of the unit vectors describing their orientation in $M_{\vec{\alpha},w}$.

Let w be a window and $\vec{\alpha}$ be an assembly sequence and $M = M_{\vec{\alpha},w}$. We use the notation $\mathscr{B}(M)$ to denote the *bond-forming submovie* of M, i.e., a restricted form of M, which consists of only those steps of M that place glues that eventually form positive-strength bonds in the assembly $\alpha = \text{res}(\vec{\alpha})$. Note that every window movie has a unique bond-forming submovie.

Lemma 2.6 (Closed Window Movie Lemma). *Let $\vec{\alpha} = (\alpha_i \mid 0 \leq i < l)$, with $l \in \mathbb{Z}^+ \cup \{\infty\}$, be an assembly sequence in some TAS $\mathscr{T}$ with result α. Let w be a closed window that partitions α into α_I and α_O, and w' be a closed window that partitions α into α'_I and α'_O. If $\mathscr{B}(M_{\vec{\alpha},w}) + \vec{c} = \mathscr{B}(M_{\vec{\alpha},w'})$ for some $\vec{c} \neq (0,0)$ and the window $w + \vec{c}$ is enclosed in w', then the assembly $\alpha'_O \cup (\alpha_I + \vec{c})$ is in $\mathscr{A}[\mathscr{T}]$.*

Proof. Before we proceed with the proof, the next paragraph introduces some notation taken directly from [5].

[2] A *cut* is a partition of the vertices of a graph into two disjoint subsets that are joined by at least one edge.

For an assembly sequence $\vec{\alpha} = (\alpha_i \mid 0 \leq i < l)$, we write $|\vec{\alpha}| = l$ (note that if $\vec{\alpha}$ is infinite, then $l = \infty$). We write $\vec{\alpha}[i]$ to denote $\vec{x} \mapsto t$, where $\vec{x}$ and t are such that $\alpha_{i+1} = \alpha_i + (\vec{x} \mapsto t)$, i.e., $\vec{\alpha}[i]$ is the placement of tile type t at position $\vec{x}$, assuming that $\vec{x} \in \partial_t \alpha_i$. We write $\vec{\alpha}[i] + \vec{c}$, for some vector $\vec{c}$, to denote $(\vec{x} + \vec{c}) \mapsto t$. We define $\vec{\alpha} = \vec{\alpha} + (\vec{x} \mapsto t) = (\alpha_i \mid 0 \leq i < k+1)$, where $\alpha_k = \alpha_{k-1} + (\vec{x} \mapsto t)$ if $\vec{x} \in \partial_t \alpha_{k-1}$ and undefined otherwise, assuming $|\vec{\alpha}| > 0$. Otherwise, if $|\vec{\alpha}| = 0$, then $\vec{\alpha} = \vec{\alpha} + (\vec{x} \mapsto t) = (\alpha_0)$, where α_0 is the assembly such that $\alpha_0(\vec{x}) = t$ and is undefined at all other positions. This is our notation for appending steps to the assembly sequence $\vec{\alpha}$: to do so, we must specify a tile type t to be placed at a given location $\vec{x} \in \partial_t \alpha_i$. If $\alpha_{i+1} = \alpha_i + (\vec{x} \mapsto t)$, then we write $Pos(\vec{\alpha}[i]) = \vec{x}$ and $Tile(\vec{\alpha}[i]) = t$. For a window movie $M = (v_0, g_0), (v_1, g_1), \ldots$, we write $M[k]$ to be the pair (v_k, g_k) in the enumeration of M and $Pos(M[k]) = v_k$, where v_k is a vertex of a grid graph.

We now proceed with the proof, throughout which we will assume that $M = \mathscr{B}\left(M_{\vec{\alpha},w}\right)$ and $M' = \mathscr{B}\left(M_{\vec{\alpha},w'}\right)$. Since $M + \vec{c} = M'$ for some $\vec{c} \neq (0,0)$ and w and w' are both closed windows, it must be the case that the seed tile of α is in dom $\alpha_O \cap$ dom α'_O or in dom $\alpha_I \cap$ dom α'_I. In other words, the seed tile cannot be in dom $\alpha_I \setminus$ dom α'_I nor in dom $\alpha'_I \setminus$ dom α_I. Therefore, assume without loss of generality that the seed tile is in dom $\alpha_O \cap$ dom α'_O.

The algorithm in Figure 2 describes how to produce a new valid assembly sequence $\vec{\gamma}$.

```
Initialize i, j = 0 and γ⃗ to be empty
for k = 0 to |M| − 1 do
    if Pos(M′[k]) ∈ dom α′_O then
        while Pos(α⃗[i]) ≠ Pos(M′[k]) do
            if Pos(α⃗[i]) ∈ dom α′_O then
                γ⃗ = γ⃗ + α⃗[i]
            i = i + 1
        γ⃗ = γ⃗ + α⃗[i]
        i = i + 1
    else
        while Pos(α⃗[j]) ≠ Pos(M[k]) do
            if Pos(α⃗[j]) ∈ dom α_I then
                γ⃗ = γ⃗ + (α⃗[j] + c⃗)
            j = j + 1
        γ⃗ = γ⃗ + (α⃗[j] + c⃗)
        j = j + 1
while inside(w) ∩ ∂res(γ⃗) ≠ ∅ do
    if Pos(α⃗[j]) ∈ dom α_I then
        γ⃗ = γ⃗ + (α⃗[j] + c⃗)
    j = j + 1
while i < |α⃗| do
    if Pos(α⃗[i]) ∈ dom α′_O then
        γ⃗ = γ⃗ + α⃗[i]
    i = i + 1
return γ⃗
```

Figure 2: The algorithm to produce a valid assembly sequence $\vec{\gamma}$.

If we assume that the assembly sequence $\vec{\gamma}$ ultimately produced by the algorithm is valid, then the result of $\vec{\gamma}$ is indeed $\alpha'_O \cup (\alpha_I + \vec{c})$. Observe that α_I must be finite, which implies that M is finite. If $|\vec{\alpha}| < \infty$, then all loops will terminate. If $|\vec{\alpha}| = \infty$, then $|\alpha'_O| = \infty$ and the first two loops will terminate and the last loop will run forever. In either case, for every tile in α'_O and $\alpha_I + \vec{c}$, the algorithm adds a step to the sequence $\vec{\gamma}$ involving the addition of this tile to the assembly. However, we need to prove that the assembly sequence $\vec{\gamma}$ is valid. It may be the case that either: 1. there is insufficient bond strength between the tile to be placed and the existing neighboring tiles, or 2. a tile is already present at this location.

Case 1: In this case, we claim the following: at each step of the algorithm, the current version of $\vec{\gamma}$ is a valid assembly sequence whose result is a producible subassembly of $\alpha'_O \cup (\alpha_I + \vec{c})$. Note that the three loops in the algorithm iterate through all steps of $\vec{\alpha}$, such that, when adding $\vec{\alpha}[i]$ (or $\vec{\alpha}[j] + \vec{c}$) to $\vec{\gamma}$, all steps of the window movie corresponding to the positions/glues of tiles to which $\vec{\alpha}[i]$ (or $\vec{\alpha}[j]+\vec{c}$) initially bind in $\vec{\alpha}$ have occurred. In other words, when adding $\vec{\alpha}[i]$ (or $\vec{\alpha}[j]+\vec{c}$) to $\vec{\gamma}$, the tiles to which $\vec{\alpha}[i]$ (or $\vec{\alpha}[j]+\vec{c}$) initially bind have already been added to $\vec{\gamma}$ by the algorithm. Similarly, all tiles in α'_O (or $\alpha_I+\vec{c}$) added to α before step i (or j) in the assembly sequence $\vec{\alpha}$ have already been added to $\vec{\gamma}$.

So, if the tile $Tile(\vec{\alpha}[i])$ that is added to the subassembly of α produced after $i-1$ steps can bond at a location in α'_O to form a τ-stable assembly, then the same tile added to the result of $\vec{\gamma}$, which is producible, must also bond to the same location in the result of $\vec{\gamma}$, as the neighboring glues consist of (i) an identical set of glues from tiles in the subassembly of α'_O and (ii) glues on the side of the window movie containing $\alpha_I+\vec{c}$. Similarly, the tiles of $\alpha_I+\vec{c}$ must also be able to bind.

Case 2: Since we only assume that $\mathscr{B}\left(M_{\vec{\alpha},w}\right)+\vec{c}=\mathscr{B}\left(M_{\vec{\alpha},w'}\right)$, as opposed to the stronger condition $\mathscr{B}\left(M_{\vec{\alpha},w+\vec{c}}\right)=\mathscr{B}\left(M_{\vec{\alpha},w'}\right)$, which is assumed in the standard WML, we must show that $\text{dom}\ (\alpha_I+\vec{c})\cap\text{dom}\ \alpha'_O=\varnothing$. To see this, observe that, by assumption, $w+\vec{c}$ is enclosed in w', which, by definition, means that $inside(w+\vec{c})\subseteq inside(w')$. Then we have $\vec{x}\in\text{dom}\ \alpha'_O\Rightarrow\vec{x}\in outside(w')\Rightarrow\vec{x}\notin inside(w')\Rightarrow\vec{x}\notin inside(w+\vec{c})\Rightarrow\vec{x}\notin\text{dom}\ (\alpha_I+\vec{c})$. Thus, locations in $\alpha_I+\vec{c}$ only have tiles from α_I placed in them, and locations in α'_O only have tiles from α'_O placed in them.

So the assembly sequence of $\vec{\gamma}$ is valid, i.e., every single-tile addition in $\vec{\gamma}$ adds a tile to the assembly to form a new producible assembly. Since we have a valid assembly sequence, as argued above, the resulting producible assembly is $\alpha'_O\cup(\alpha_I+\vec{c})$.

□

3 Scaled pier fractals do not strictly self-assemble in the aTAM

In this section, we first define some notation and establish preliminary results. Then we prove our main result, namely that no scaled pier fractal self-assembles in the aTAM. Finally, we prove corollaries of our main result, including the fact that no scaled tree fractal self-assembles in the aTAM.

3.1 Preliminaries

Recall that each stage X_s ($s>1$) of a g-dssf (scaled by a factor c) is made up of copies of the previous stage X_{s-1}, each of which is a square of size cg^{s-1}.

In the proof of our main result, we will need to refer to one of the squares of size cg^{s-2} inside the copies of stage X_{s-1}, leading to the following notation.

Notation. Let $c\in\mathbb{Z}^+$, $1<s\in\mathbb{N}$ and $1<g\in\mathbb{N}$. Let $e,f,p,q\in\mathbb{N}_g$. We use $S_s^c(e,f,p,q)$ to denote $\{0,1,\ldots,cg^{s-2}-1\}^2+cg^{s-1}(e,f)+cg^{s-2}(p,q)$ and $W_s^c(e,f,p,q)$ to denote the square-shaped, closed window whose inside is $S_s^c(e,f,p,q)$.

In Figure 3 below, the bottom and top (circular) magnifications show the windows $W_2^1(0,1,3,2)$ and $W_3^1(0,1,3,2)$, respectively.

Next, we will need to translate a small window to a position inside a larger window. These two windows will correspond to squares at the same relative position in different stages i and j of a g-dssf.

Notation. Let $c\in\mathbb{Z}^+$, $1<i\in\mathbb{N}$, $1<j\in\mathbb{N}$, with $i<j$, and $e,f,xsp,q\in\mathbb{N}_g$. We use $\vec{t}^{\,c}_{i\to j}(e,f,p,q)$ to denote the vector joining the southwest corner of $W_i^c(e,f,p,q)$ to the southwest corner of $W_j^c(e,f,p,q)$. In other words, $\vec{t}^{\,c}_{i\to j}(e,f,p,q)=\left(c\left(g^{j-1}-g^{i-1}\right)e+c\left(g^{j-2}-g^{i-2}\right)p,c\left(g^{j-1}-g^{i-1}\right)f+c\left(g^{j-2}-g^{i-2}\right)q\right)$.

For example, in Figure 3 below, $\vec{t}^{\,1}_{2\to3}(0,1,3,2)=(9,18)$.

To apply Lemma 2.6, we will need the bond-forming submovies to line up. Therefore, once the two square windows share their southwest corner after using the translation defined above, we will need to further translate the smallest one either up or to the right, or both, depending on which side of the windows contains the bond-forming glues, which, in the case of scaled pier fractals, always form a straight (vertical or horizontal) line of length c. We will compute the coordinates of this second translation in our main proof. For now, we establish an upper bound on these coordinates that will ensure that the translated window will remain enclosed in the larger window.

Lemma 3.1. *Let $c \in \mathbb{Z}^+$, $1 < i \in \mathbb{N}$, $1 < j \in \mathbb{N}$, with $i < j$, $e, f, p, q \in \mathbb{N}_g$, and $x, y \in \mathbb{N}$. Let $m = c(g^{j-2} - g^{i-2})$. If $x \leq m$ and $y \leq m$, then the window $W_i^c(e, f, p, q) + \vec{t}^c_{i \to j}(e, f, p, q) + (x, y)$ is enclosed in the window $W_j^c(e, f, p, q)$.*

Proof. Let W and w denote $W_j^c(e, f, p, q)$ and $W_i^c(e, f, p, q) + \vec{t}^c_{i \to j}(e, f, p, q)$, respectively. Since W and w are square windows that have the same southwest corner and whose respective sizes are cg^{j-2} and cg^{i-2}, W encloses w. The eastern side of $w + (x, 0)$ still lies within W, because the maximum value of x is equal to the difference between the size of W and the size of w. The same reasoning applies to a northward translation of w by $(0, y)$. In conclusion, $w + (x, y)$ must be enclosed in W, as long as neither x nor y exceeds m. $\square$

Finally, in our main result, we will use the fact that, for any scaled pier fractal $\mathbf{P}^c$, we can find an infinite number of closed windows that all cut the fractal along a single line of glues (see Lemma 3.5 below), the proof of which uses the following three intermediate lemmas.

Lemma 3.2. *If $\mathbf{P}$ is any pier fractal with generator G, then G contains at least one pier that is not a double-bridge pier.*

Proof. For the sake of obtaining a contradiction, assume that G contains exactly one pier, say (p, q), and that (p, q) is a double-bridge pier. Note that any double-bridge pier must be positioned at one of the corners of G, that is, $(p, q) \in \{(0, 0), (g-1, 0), (0, g-1), (g-1, g-1)\}$. Without loss of generality, assume that $(p, q) = (g-1, 0)$, as in Figure 1 above. Since (p, q) is a double-bridge pier, $(0, 0)$ must be the other point in the h-bridge and $(g-1, g-1)$ must be the other point in the v-bridge. Thus, $(0, 0) \in G$ (this is also true by definition of G) and $(g-1, g-1) \in G$. Since (p, q) is the only pier in G, $(0, 0)$ cannot be north-free (nor east-free), which implies that $(0, 1) \in G$. Therefore, $(g-1, 1) \notin G$ (otherwise, G would contain a second h-bridge). Similarly, since (p, q) is the only pier in G, $(g-1, g-1)$ cannot be west-free (nor south-free), which implies that $(g-2, g-1) \in G$. Therefore, $(g-2, 0) \notin G$ (otherwise, G would contain a second v-bridge). In conclusion, the point $(p, q) = (g-1, 0)$ is in G but it is not connected to the rest of G, which contradicts the definition of $\mathbf{P}$, whose generator must be connected. $\square$

Lemma 3.3. *Let $\mathbf{P}$ be any pier fractal with generator G such that $(p, q) \in G$ is a parallel single-bridge pier. If $c \in \mathbb{Z}^+$, then it is always possible to pick one point (e, f) in G such that, for $1 < s \in \mathbb{N}$, $W_s^c(e, f, p, q)$ encloses a configuration that is connected to $\mathbf{P}^c$ via a single connected line of glues of length c.*

Proof. Without loss of generality, assume that the pier (p, q) is pointing north, that it belongs to a v-bridge and that $q = g-1$ (a similar reasoning holds if $q = 0$ and the pier points south, or if the pier belongs to an h-bridge and points either west or east). Now, we must pick a point (e, f) such that any window of the form $W_s^c(e, f, p, q)$ has exactly three free sides. We distinguish two cases.

1. If $p = 0$, that is, the pier is in the leftmost column of G, then $(1, g-1) \notin G$, since $(0, g-1)$ is a north-pointing pier. Therefore, there must exist at least one point in $G \cap (\{1\} \times \mathbb{N}_{g-1})$, say $(1, y)$, with $0 \leq y < g-1$, that is north-free. In this case, we pick (e, f) to be equal to $(1, y)$. Now, consider any window w of the form $W_s^c(e, f, p, q)$. The north side of w is free (since $q = g-1$, (e, f) is north-free in G and $f = y < g-1$), the east side of w is free (since $(1, g-1) \notin G$), and the west side of w is free (since the facts that $(0, 0) \in G$, $(0, g-1) \in G$ and $(0, g-1)$ is a single-bridge pier together imply that $(g-1, g-1) \notin G$). Furthermore, since $(0, g-1)$ is a north-pointing pier, $S(0, g-1) \in G$.

2. If $p > 0$, then $(p-1, g-1) \notin G$. Therefore, there must exist at least one point in $G \cap (\{p-1\} \times \mathbb{N}_{g-1})$, say $(p-1, y)$, with $0 \leq y < g-1$, that is north-free. In this case, we pick (e, f) to be equal to $(p-1, y)$. Now, consider any window w of the form $W_s^c(e, f, p, q)$. The north side of w is free (since $q = g-1$, (e, f) is north-free in G and $f = y < g-1$), the west side of w is free (because $(p-1, g-1) \notin G$), and the east side of w is free (since, either $p < g-1$ and $(p+1, g-1) \notin G$, or $p = g-1$, in which case the facts that $(g-1, g-1) \in G$, $(g-1, 0) \in G$ and $(g-1, g-1)$ is a single-bridge pier together imply that $(0, g-1) \notin G$). Furthermore, since $(p, g-1)$ is a north-pointing pier, $S(p, g-1) \in G$.

Therefore, in both cases, $W_s^c(e, f, p, q)$ has exactly three free sides and encloses a configuration that is connected to $\mathbf{P}^c$ via a single connected horizontal line of glues of length c positioned on the south side of the window.

$\square$

Lemma 3.4. *Let* $\mathbf{P}$ *be any pier fractal with generator* G *such that* $(p,q) \in G$ *is an orthogonal single-bridge pier. If* $c \in \mathbb{Z}^+$*, then it is always possible to pick one point* (e, f) *in* G *such that, for* $1 < s \in \mathbb{N}$, $W_s^c(e, f, p,q)$ *encloses a configuration that is connected to* $\mathbf{P}^c$ *via a single connected line of glues of length* c*.*

Proof. Without loss of generality, assume that the pier (p,q) is pointing east, that it belongs to a v-bridge and that $q = g-1$ (a similar reasoning holds if $q = 0$, or if the pier points west, or if the pier belongs to an h-bridge and points either north or south). Note that, in this case, g must be strictly greater than 2, since $(p,g-1) \in G$, $(p,0) \in G$ but $(p,g-2) \notin G$. Now, we must pick a point (e, f) such that any window of the form $W_s^c(e, f, p,q)$ has exactly three free sides. We distinguish two cases.

1. If $p < g-1$, then $(p,g-2) \notin G$, since $(p,g-1)$ is an east-pointing pier. Therefore, there must exist at least one point in $G \cap (\{p\} \times \mathbb{N}_{g-2})$, say (p,y), with $0 \le y < g-2$, that is north-free. In this case, we pick (e, f) to be equal to (p,y). Now, consider any window w of the form $W_s^c(e, f, p,q)$. The north side of w is free (since $q = g-1$, (e, f) is north-free in G and $f = y < g-2 < g-1$), the east side of w is free (since $p < g-1$ and $(p+1,g-1) \notin G$), and the south side of w is free (since $(p,g-2) \notin G$). Furthermore, since $(p,g-1)$ is an east-pointing pier, $W(p,g-1) \in G$.

2. If $p = g-1$, that is, the pier is in the rightmost column of G, then the facts that $(g-1,0) \in G$, $(g-1,g-1) \in G$ and $(g-1,g-1)$ is a single-bridge pier together imply that $(0,g-1) \notin G$. This, together with the fact that $(0,0) \in G$, implies that there must exist at least one point in $G \cap (\{0\} \times \mathbb{N}_{g-1})$, say $(0,y)$, with $0 \le y < g-1$, that is north-free. In this case, we pick (e, f) to be equal to $(0,y)$. Now, consider any window w of the form $W_s^c(e, f, p,q)$. The north side of w is free (since $q = g-1$, (e, f) is north-free in G and $f = y < g-1$), the east side of w is free (since $(0,g-1) \notin G$), and the south side of w is free (since $(p,g-2) \notin G$). Furthermore, since $(g-1,g-1)$ is an east-pointing pier, $W(g-1,g-1) \in G$.

Therefore, in both cases, $W_s^c(e, f, p,q)$ has exactly three free sides and encloses a configuration that is connected to $\mathbf{P}^c$ via a single connected horizontal line of glues of length c positioned on the west side of the window.

□

Lemma 3.5. *Let* $\mathbf{P}$ *be any pier fractal with generator* G*. If* $c \in \mathbb{Z}^+$*, then it is always possible to pick one pier* (p,q) *and one point* (e, f)*, both in* G*, such that, for* $1 < s \in \mathbb{N}$, $W_s^c(e, f, p,q)$ *encloses a configuration that is connected to* $\mathbf{P}^c$ *via a single connected (horizontal or vertical) line of glues of length* c*.*

Proof. Let $\mathbf{P}$ be any pier fractal with generator G. Let $c \in \mathbb{Z}^+$ and $1 < s \in \mathbb{N}$. By definition of a pier fractal, G contains at least one pier. We will pick one of these piers carefully.

According to Lemma 3.2, it is always possible to choose a pier in G that is not a double-bridge pier. Therefore, we can always choose either a real pier or a single-bridge pier. We now consider the three possible cases.

First, if G contains one or more real piers, we can simply choose one of them as (p,q). In this case, we pick $(e, f) = (p,q)$, since any window of the form $W_s^c(p,q, p,q)$, where (p,q) is a real pier in G, must have exactly three free sides. Therefore, $W_s^c(p,q, p,q)$ must enclose a configuration that is connected to $\mathbf{P}^c$ via a single line of glues of length c, namely on its non-free side.

Second, if G does not contain any real piers, G must contain at least one single-bridge pier. So we wrap up this proof by considering the two types of single-bridge piers.

If G contains at least one parallel single-bridge pier, according to Lemma 3.3, it is always possible to choose one pier (p,q) and one point (e, f), both in G, such that, for $1 < s \in \mathbb{N}$, $W_s^c(e, f, p,q)$ encloses a configuration that is connected to $\mathbf{P}^c$ via a single connected line of glues of length c.

Finally, if G contains at least one orthogonal single-bridge pier, according to Lemma 3.4, it is always possible to choose one pier (p,q) and one point (e, f), both in G, such that, for $1 < s \in \mathbb{N}$, $W_s^c(e, f, p,q)$ encloses a configuration that is connected to $\mathbf{P}^c$ via a single connected line of glues of length c. □

3.2 Main result

We are now ready to prove our main result.

Theorem 3.6. *Let* $\mathbf{P}$ *be any pier fractal. If* $c \in \mathbb{Z}^+$*, then* $\mathbf{P}^c$ *does not strictly self-assemble in the aTAM.*

Proof. Let $\mathbf{P}$ be any pier fractal with a $g \times g$ generator G, where $1 < g \in \mathbb{N}$. Let c be any positive integer. For the sake of obtaining a contradiction, assume that $\mathbf{P}^c$ does strictly self-assemble in some TAS $\mathscr{T} = (T, \sigma, \tau)$. Further assume that $\vec{\alpha}$ is some assembly sequence in $\mathscr{T}$ whose result is α, such that dom $\alpha = \mathbf{P}^c$.

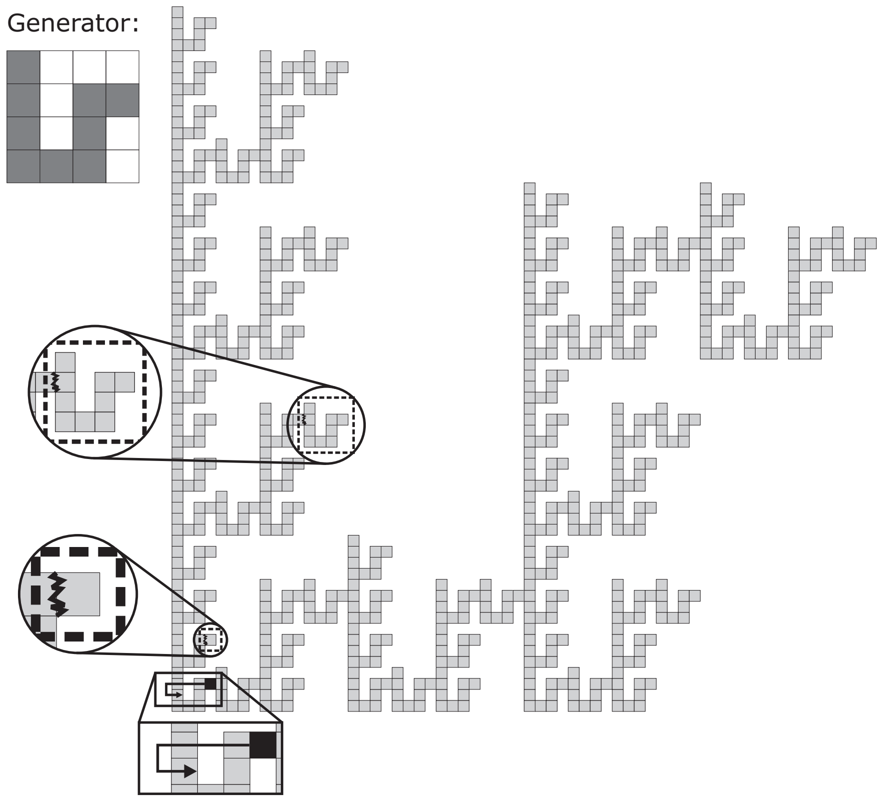

Figure 3: First three stages ($s = 1, 2, 3$) of an unscaled ($c = 1$) pier fractal with an east-pointing pier at position $(3, 2)$. The east-free point $(0, 1)$ is at the tip of the arrow (see the rectangular magnification box). In other words, $g = 4$, $(p, q) = (3, 2)$, and $(e, f) = (0, 1)$.

According to Lemma 3.5, we can always pick one pier (p, q) and a point (e, f), both in G, such that, for $1 < s \in \mathbb{N}$, the window $W_s^c(e, f, p, q)$, which we will abbreviate w_s, encloses a configuration that is connected to $\mathbf{P}^c$ via a single line of glues of length c.[3] The maximum number of distinct combinations and orderings of glue positionings along this line of glues is finite.[4] By the generalized pigeonhole principle, since $|\{w_s \mid 1 < s \in \mathbb{N}\}|$ is infinite, there must be at least one bond-forming submovie such that an infinite number of windows generate this submovie (up to translation). Let us pick two such windows, say, w_i and w_j with $i < j$, such that $\mathscr{B}(M_{\vec{\alpha}, w_i})$ and $\mathscr{B}(M_{\vec{\alpha}, w_j})$ are equal (up to translation). We must pick these windows carefully, since as stated in the proof of Lemma 2.6, the seed of α must be either in both windows or in neither. This condition can always be satisfied. The only case where the seed is in more than one window is when it is at position $(0, 0)$ and $e = f = p = q = 0$, which implies that all windows include the origin. So, in this case, any choice of i and $j > i$ will do. In all other cases, none of the windows overlap. So, if the seed belongs to one of them, say w_k, then we can pick any i greater than k (and $j > i$). Finally, if the seed does not belong to any windows, then any choice of i and $j > i$ will do.

[3] Without loss of generality, we will assume that this line of glues is positioned on the western side of the windows and is thus vertical (see the jagged lines in Figure 3, where $s = 2$ and $s = 3$ for the small and large windows, respectively, and $(p, q) = (3, 2)$ and $(e, f) = (0, 1)$), because the chosen pier in our example points east. A similar reasoning holds for piers pointing north, south or west.

[4] This number is bounded above by $T_{glue}^{2c} \cdot (2c)!$, where T_{glue} is the total number of distinct glue types in T.

	Translation Formulas for (x,y)	
North-pointing Pier Vertical Bridge with Southern Bridge Point at $(a,0)$	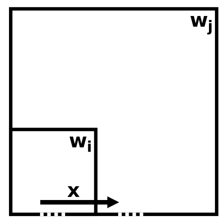	$x = ac\Sigma_{k=i-2}^{j-3} g^k$ $y = 0$
East-pointing Pier Horizontal Bridge with Western Bridge Point at $(0,b)$	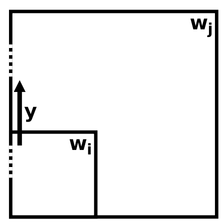	$x = 0$ $y = bc\Sigma_{k=i-2}^{j-3} g^k$
South-pointing Pier Vertical Bridge with Northern Bridge Point at $(a,g-1)$	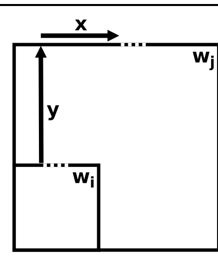	$x = ac\Sigma_{k=i-2}^{j-3} g^k$ $y = c(g^{j-2} - g^{i-2})$
West-pointing Pier Horizontal Bridge with Eastern Bridge Point at $(g-1,b)$	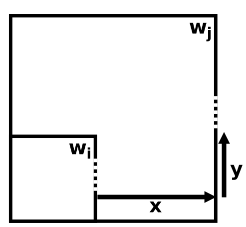	$x = c(g^{j-2} - g^{i-2})$ $y = bc\Sigma_{k=i-2}^{j-3} g^k$

Figure 4: Computing the coordinates (x,y) of the translation that aligns the bond-forming glues (shown as a dotted line) of the windows w_i and w_j. Note that (a,b) with $a \in \mathbb{N}_g$ and $b \in \mathbb{N}_g$ are the coordinates of the point in G within the (horizontal or vertical) bridge that determines the position of the bond-forming glues.

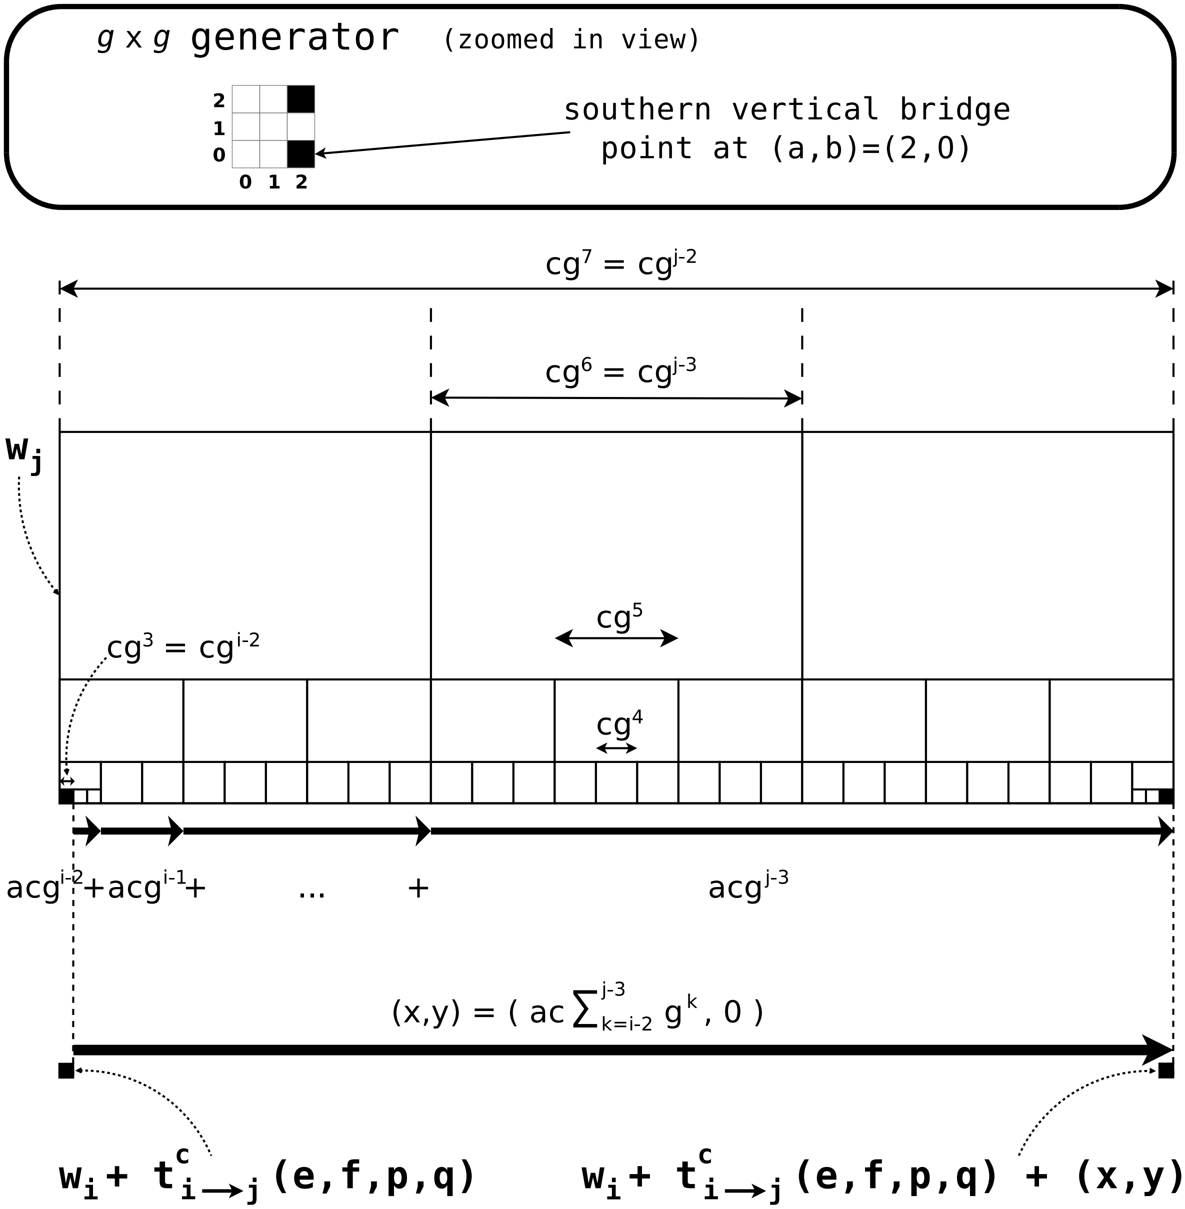

Figure 5: (x,y) translation needed to align w_i and w_j on their east side once their southwest corners already match. Example with a north-pointing pier (not shown) and $g = 3$, $i = 5$, $j = 9$, and southern vertical bridge point at location $(a,b) = (2,0)$.

We will now prove that w_i and w_j satisfy the two conditions of Lemma 2.6.

First, we compute $\vec{c}$ such that $\mathscr{B}(M_{\vec{\alpha},w_i}) + \vec{c} = \mathscr{B}(M_{\vec{\alpha},w_j})$. We know that $w_i + \vec{t}^c_{i\to j}(e,f,p,q)$ and w_j share their southwest corner. We need to perform one more translation (x,y) to align the bond-forming glues of w_i and w_j. The values of x and y depend on the direction in which the chosen pier is pointing. The formulas corresponding to all four directions are given in Figure 4. Furthermore, a justification for the recurring summation in the formulas of Figure 4 is provided in Figure 5. In that figure, the chosen case is a north-pointing pier. However, our example in Figure 3 above uses an east-pointing pier. We now complete the proof for this case. To align the bond-forming glues of w_i and w_j, we must translate $w_i + \vec{t}^c_{i\to j}(e,f,p,q)$ by $(x,y) = \left(0, bc\sum_{k=i-2}^{j-3} g^k\right)$, with $b \leq g-1$. Since $x = 0 \leq m$ (as defined in Lemma 3.1) and $y = bc\sum_{k=i-2}^{j-3} g^k \leq (g-1)c\sum_{k=i-2}^{j-3} g^k = c\left(\sum_{k=i-1}^{j-2} g^k - \sum_{k=i-2}^{j-3} g^k\right) = c\left(g^{j-2} - g^{i-2}\right) = m$, we can apply Lemma 3.1 to infer that, with $\vec{c} = \vec{t}^c_{i\to j}(e,f,p,q) + (x,y)$, $w_i + \vec{c}$ is enclosed in w_j. Therefore, the second condition of Lemma 2.6 holds.

Second, by construction, $\mathscr{B}(M_{\vec{\alpha},w_i}) + \vec{c} = \mathscr{B}(M_{\vec{\alpha},w_j})$. Therefore, the first condition of Lemma 2.6 holds.

In conclusion, the two conditions of Lemma 2.6 are satisfied, with α_I and α'_O defined as the intersection of $\mathbf{P}^c$ with the inside of w_i and the outside of w_j, respectively. We can thus conclude that the assembly $\alpha'_O \cup (\alpha_I + \vec{c})$ is producible in $\mathscr{T}$. Note that this assembly is identical (up to translation) to $\mathbf{P}^c$, except that the interior of the large window w_j is replaced by the interior of the small window w_i. Since the configurations in these two windows cannot be identical, we have proved that $\mathscr{T}$ does not strictly self-assemble $\mathbf{P}^c$, which is a contradiction.

□

3.3 Corollaries of our main result

In this section, we discuss both special cases and generalizations of our main result.

3.3.1 Specializations of our main result

In [1], we proved that no scaled tree fractal strictly self-assembles in the aTAM, where a *tree fractal* is a discrete self-similar fractal whose underlying graph is a tree. In this section, we start by proving a new characterization of tree fractals in terms of simple connectivity properties of their generator.

Theorem 3.7. $\mathbf{T} = \bigcup_{i=1}^{\infty} T_i$ *is a g-discrete self-similar tree fractal, for some* $g > 1$, *with generator G if and only if*

a. G is a tree and

b. $nhb_G = nvb_G = 1$.

The proof of this theorem is in the appendix. Next, the following observation follows from the fact that a tree with more than one vertex must contain at least two leaf nodes.

Observation 3.8. *If G is the generator of any discrete self-similar fractal and G is a tree, then it must contain at least two piers.*

Finally, we can recast the main result in [1] as a special case of our main result.

Corollary 3.9. *[From [1]] Let* $\mathbf{T}$ *be any tree fractal. If* $c \in \mathbb{Z}^+$, *then* $\mathbf{T}^c$ *does not strictly self-assemble in the aTAM.*

Proof. Let $\mathbf{T}$ be any tree fractal with generator G. According to Theorem 3.7, the full grid graph of G is a tree and is thus connected, and $nhb_G = nvb_G = 1$. Furthermore, according to Observation 3.8, G must contain at least one pier. Therefore, $\mathbf{T}$ is a pier fractal and $\mathbf{T}^c$ does not strictly self-assemble in the aTAM. □

We now turn our attention to a second specialization of our main result by considering "pinch-point fractals," which are defined in [6] as follows.

Definition 3.10. Let $\mathbf{X} \subset \mathbb{N}^2$ be a g-discrete self-similar fractal with generator G. We say that $\mathbf{X}$ is a *pinch-point discrete self-similar fractal* if G satisfies the following four conditions:

1. $\{(0,0),(0,g-1),(g-1,0)\} \subseteq G$.
2. $G \cap (\{1,\ldots,g-1\} \times \{g-1\}) = \varnothing$.
3. $G \cap (\{g-1\} \times \{1,\ldots,g-1\}) = \varnothing$.
4. The full grid graph of G is connected.

Theorem 3.12 in [6] establishes that no pinch-point fractal strictly self-assembles in the aTAM. We can now generalize this result as follows.

Corollary 3.11. *Let* $\mathbf{X}$ *be any pinch-point discrete self-similar fractal. If* $c \in \mathbb{Z}^+$*, then* $\mathbf{X}^c$ *does not strictly self-assemble in the aTAM.*

Proof. Let $\mathbf{X}$ be any pinch-point discrete self-similar fractal with generator $G \subset \mathbb{N}_g^2$, for some $g > 1$. First, by definition of a pinch-point fractal, the full grid graph of G is connected. Second, since the point $(g-1,0)$ is the only point of G that belongs to $\{g-1\} \times \mathbb{N}$ and the point $(0,0)$ also belongs to G, $nhb_G = 1$. Similarly, since the point $(0,g-1)$ is the only point of G that belongs to $\mathbb{N} \times \{g-1\}$ and the point $(0,0)$ also belongs to G, $nvb_G = 1$. Third, since the point $(0,g-1)$ is a pier in G, G contains at least one pier. Therefore, $\mathbf{X}$ is a pier fractal. In conclusion, if $c \in \mathbb{Z}^+$, then $\mathbf{X}^c$ does not strictly self-assemble in the aTAM.

□

3.3.2 Generalizations of our main result

We now discuss how to extend our main result to different classes of fractals. More specifically, we will relax the last two conditions in the definition of pier fractals and still be able to use the same reasoning as we did in the proof of our main result.

First, our proof of Theorem 3.6 uses the fact that there exist an infinite collection of square windows, each of which encloses a sub-configuration of the fractal that is attached to the rest of the fractal at a single point (or single line of points). In other words, each window in the collection has three free sides. If, for example, the east and west sides of each window are free, then the number of horizontal bridges in the generator G does not matter. Even if $nhb_G > 1$, our construction for the windows still works. Figure 6 is one example of such a fractal to which our main result generalizes, with the first three windows shown as thick, black squares. In this case, our proof technique still works, even though the generator contains three horizontal bridges. Here is a precise statement of the corollary.

Corollary 3.12. *Let* $\mathbf{F}$ *be a discrete self-similar fractal with generator* G *such that the full grid graph of* G *is connected,* $nvb_G = 1$*, and* G *contains at least one north-pointing pier or one south-pointing pier. If* $c \in \mathbb{Z}^+$*, then* $\mathbf{F}^c$ *does not strictly self-assemble in the aTAM.*

Symmetrically, a similar result holds for fractals whose generator G contains either at least one west-pointing pier or at least one east-pointing pier, and such that $nhb_G = 1$.

Second, having relaxed the second condition (part *b*) of the definition of pier fractals, we can now relax the third condition (part *c*) as well. To apply our Closed Window Movie Lemma, a pier is not strictly needed. Instead, the generator only need contain a *pier-like sub-configuration*, that is, a sub-configuration of one or more tiles that is attached to the rest of the fractal at a single point. Figure 7 gives one example of such a fractal with the first two windows shown as thick, solid, black squares. In this case, our proof technique still works, even though the generator contains five horizontal bridges and no pier. Here is a precise statement of the corollary.

Corollary 3.13. *Let* $\mathbf{F}$ *be a discrete self-similar fractal with generator* G *such that the full grid graph of* G *is connected,* $nvb_G = 1$*, and* G *contains at least one north-pointing pier-like sub-configuration or at least one south-pointing pier-like sub-configuration. If* $c \in \mathbb{Z}^+$*, then* $\mathbf{F}^c$ *does not strictly self-assemble in the aTAM.*

Symmetrically, a similar result holds for fractals whose generator G contains either at least one west-pointing pier-like sub-configuration or at least one east-pointing pier-like sub-configuration, and such that $nhb_G = 1$.

Finally, the Closed Window Movie Lemma may be applicable even when the generator does not contain any pier-like sub-configuration. The key requirement in the proof of our main result is to be able to find at least two

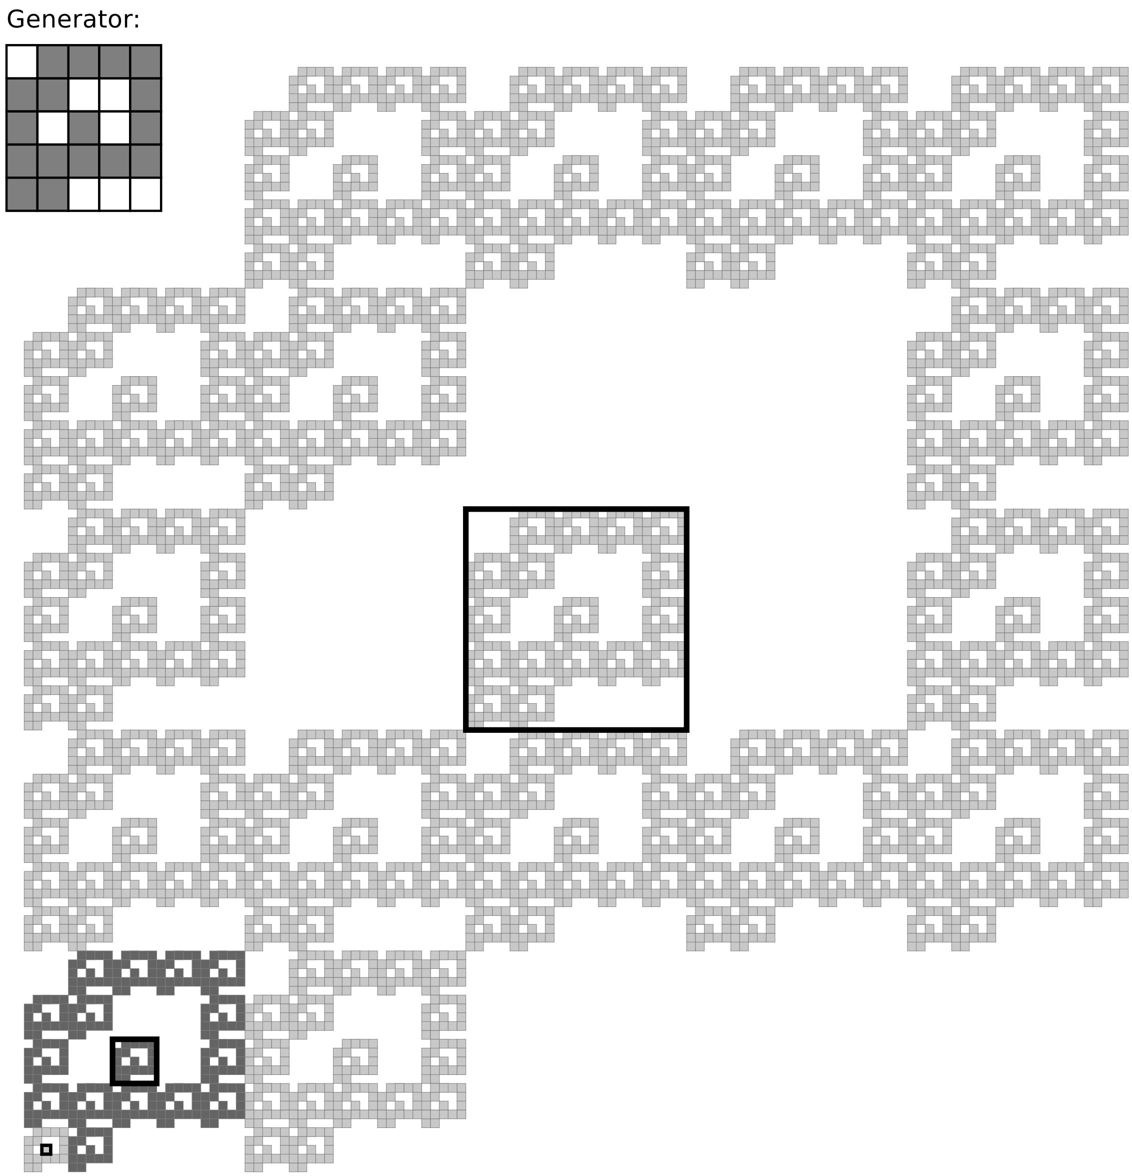

Figure 6: First three stages ($s = 1,2,3$) of an unscaled ($c = 1$) 5-discrete self-similar fractal with a north-pointing pier, $nhb_G = 3$, $nvb_G = 1$, $(p,q) = (2,2)$, and $(e,f) = (0,0)$.

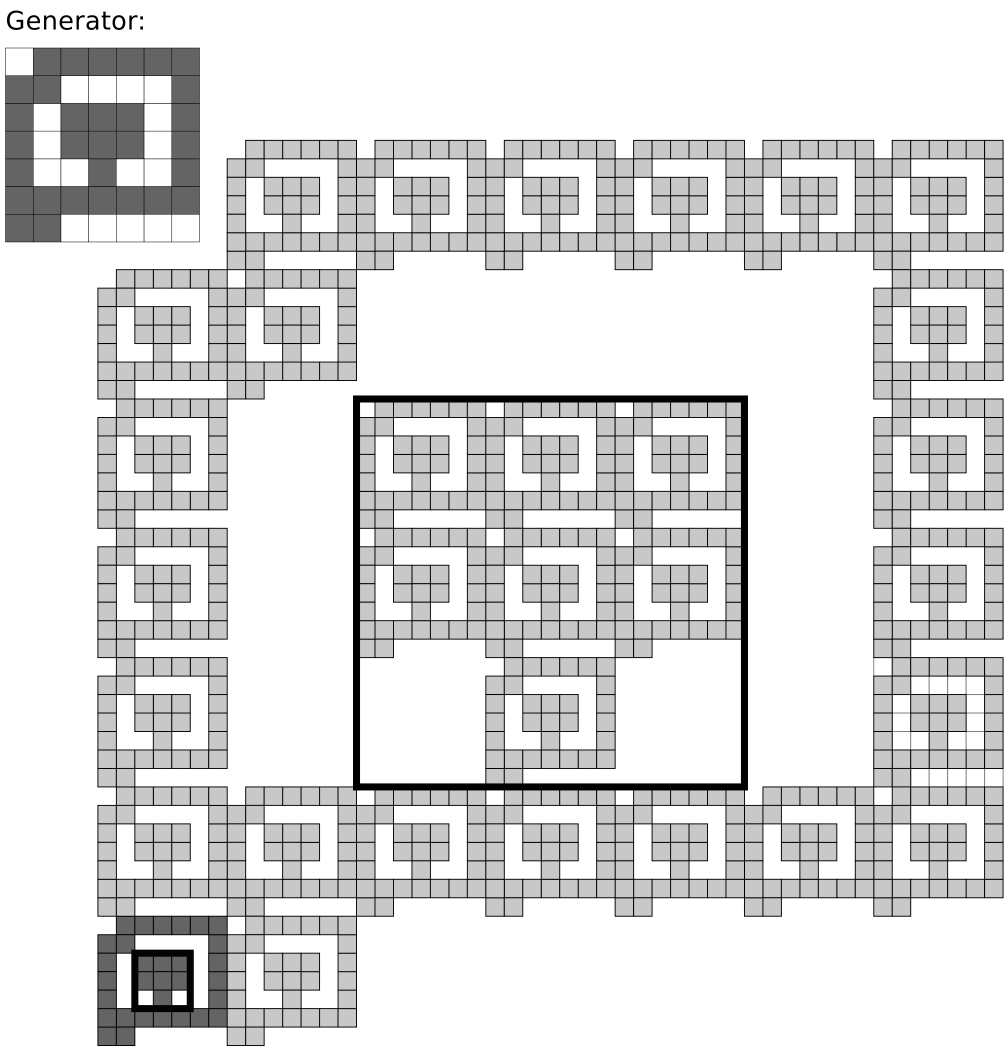

Figure 7: First two stages ($s = 1, 2$) of an unscaled ($c = 1$) 7-discrete self-similar fractal with a north-pointing pier-like sub-configuration, $nhb_G = 5$ and $nvb_G = 1$.

windows that share a common bond-forming window movie but whose insides contain different sub-configurations. This requirement can be met even when the sub-configuration contained in each window is attached to the rest of the fractal at more than one point. Figure 8 illustrates such a situation. For a general characterization, we need some definitions.

If G is a $g \times g$ generator, then a *column* is any set $G \cap (\{x\} \times \mathbb{N}_g)$, where $x \in \mathbb{N}_g$ is the index of the column. Therefore, columns are indexed from left to right starting at 0. Two columns are *equivalent* if they contain the same number of points and, for each point in one column, there is a point in the other column with the same y coordinate. A *vertical cut* is any set of edges connecting two adjacent columns of G. Two vertical cuts are *equivalent* if they contain the same number of edges and, for each edge in one cut, there is an edge in the other cut with the same y coordinate. Figure 8 depicts a 5×5 generator G in which columns 2 and 3 are equivalent. Furthermore, in this example, cuts 2 and 3 are also equivalent.[5] Note that the fact that two columns i and j are equivalent does not imply that the vertical cuts i and j are also equivalent. That both facts hold is just a coincidence in this example. If, for example, the point $(4,1)$ were removed from the generator in Figure 8, then columns 2 and 3 would still be equivalent, but vertical cuts 2 and 3 would no longer be equivalent.[6] In general, there is no correlation between the indices of equivalent columns and the indices of equivalent vertical cuts. However, the co-existence of equivalent columns and equivalent vertical cuts in the same generator may render the Window Movie Lemma applicable.

In the example of Figure 8, vertical cut 2 is to the east of the vertical bridge (which is a subset of column 1). Therefore, if we can find an east-free point in G, e.g., the point $(1,0)$ in our running example, we will be able to position a closed window that only cuts the fractal on one side (here, its western side), e.g., the smallest of the two solid windows in Figure 8. Similarly, we can position another closed window of the same size that cuts the generator through vertical cut 3, e.g., the dotted window that overlaps the small solid window. By construction, the window movies corresponding to these two windows have the same length and contain exactly the same positions (up to translation). Of course, these window movies may not be equal up to translation because the glues in their respective positions may not match. But this is where we can take advantage of the existence of two equivalent columns. By self-similarity, these two columns will, in the next stage of the fractal, become two 5-wide sets of columns of height 20 that are pairwise equivalent, that is, columns c_1 and c_1' are equivalent, columns c_2 and c_2' are equivalent, $\ldots$, and columns c_5 and c_5' are equivalent. More importantly, the 20-high cuts labeled a, b, c, and d in Figure 8 are all pairwise equivalent. Therefore, at this stage of the fractal, we can build four larger square windows, as shown in Figure 8. Furthermore, at each successive stage of the fractal, we will be able to build twicef[7] as many square windows that all generate window movies of the same length and with positions that are equal up to translation. Since the number of window movies grows without bound as the stage number increases, but the number of distinct combinations and orderings of glue positionings is finite (following a reasoning similar to the one in Footnote 4), there is always a stage (in fact, an infinite number of them) that contains two bond-forming window movies that are identical up to translation. The sub-configurations inside the two corresponding windows cannot be equivalent because of the way the windows overlap. Additionally, since the two windows have exactly the same shape and size, the translation of the eastmost one is enclosed in (in fact, equal to) the other one. Therefore, we can apply the Closed Window Movie Lemma and conclude the proof by contradiction. Here is a precise statement of the corollary that covers the class of similar situations.

Corollary 3.14. *Let* $\mathbf{F}$ *be a discrete self-similar fractal with generator G such that the full grid graph of G is connected, G contains two equivalent columns, and G contains two equivalent vertical cuts that are positioned on the same side of all vertical bridges. If* $c \in \mathbb{Z}^+$, *then* $\mathbf{F}^c$ *does not strictly self-assemble in the aTAM.*

Symmetrically, a similar result holds for fractals with equivalent rows and equivalent horizontal cuts.

To conclude this section, we note that Corollary 3.14 could have been proved using the standard Window Movie Lemma introduced in [5], since the windows used in the proof have exactly the same shape and size. In the next section, we motivate our introduction of the Closed Window Movie Lemma as a more convenient tool in the study of scaled pier fractals.

[5]The index of a vertical cut is given by the index of the leftmost of the two columns that its edges connect.

[6]We included point $(4,1)$ in the generator to exclude piers from the generator in this example.

[7]Note that, in this example, cut 1 is also equivalent to cuts 2 and 3. So we can actually build three windows for each one of the equivalent columns in the generator. Therefore, we could could have drawn 3, 6, 12, etc. windows for stages 2, 3, 4, etc., respectively. However, we chose to use only two of the three equivalent cuts in our discussion in order to keep the figure as legible as possible.

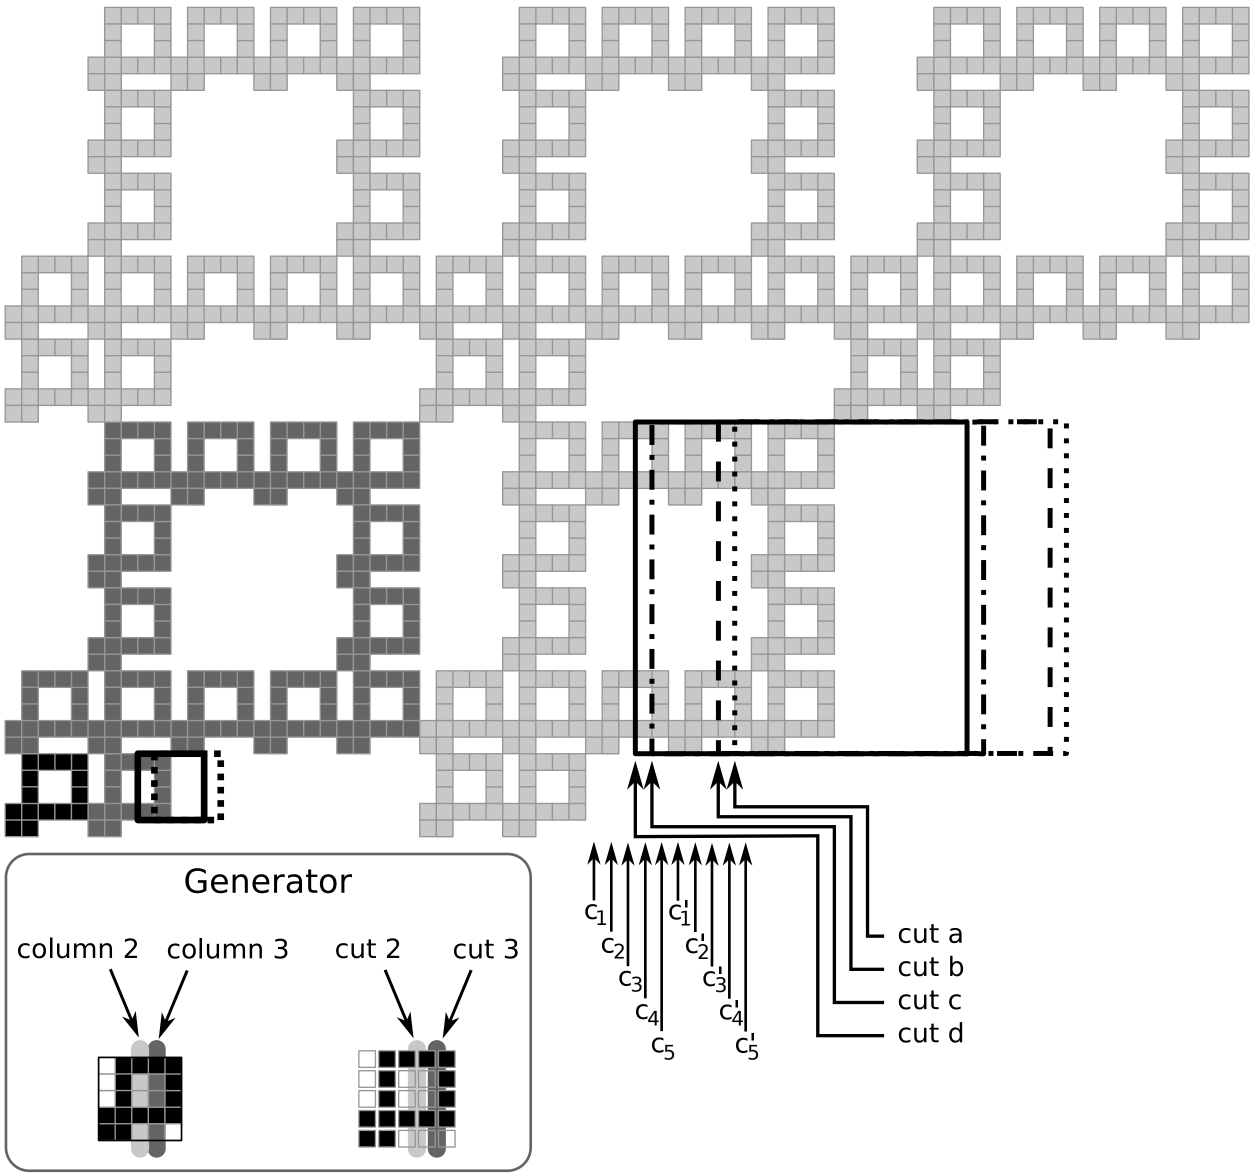

Figure 8: First two stages (and part of the third stage) of an unscaled ($c = 1$) 5-discrete self-similar fractal with two equivalent columns and two equivalent vertical cuts.

4 Discussion

A fair question for one to ask is: why not simply prove Theorem 3.6 using the standard Window Movie Lemma from [5]? Our response is that we currently do not know that we cannot.

For the sake of discussion, the statement of the standard WML, restricted to bond-forming submovies, is as follows.

Lemma 4.1 (Standard Window Movie Lemma [5])**.** *Let $\vec{\alpha} = (\alpha_i \mid 0 \leq i < l)$ and $\vec{\beta} = (\beta_i \mid 0 \leq i < m)$, with $l, m \in \mathbb{Z}^+ \cup \{\infty\}$, be assembly sequences in some TAS $\mathcal{T}$ with results α and β, respectively. Let w be a window that partitions α into two configurations α_L and α_R, and, for some $\vec{c} \neq (0,0)$, $w' = w + \vec{c}$ be a translation of w that partitions β into two configurations β_L and β_R. Furthermore, define $M_{\vec{\alpha},w}$ and $M_{\vec{\beta},w'}$ to be the respective window movies for $\vec{\alpha}, w$ and $\vec{\beta}, w'$ and define α_L, β_L to be the sub-configurations of α and β containing the seed tiles of α and β, respectively. Then, if $\mathscr{B}\left(M_{\vec{\alpha},w}\right) + \vec{c} = \mathscr{B}\left(M_{\vec{\beta},w'}\right)$, it is the case that the following two assemblies are also producible: (1) the assembly $\alpha_L\beta'_R = \alpha_L \cup \beta'_R$ and (2) the assembly $\beta'_L\alpha_R = \beta'_L \cup \alpha_R$, where $\beta'_L = \beta_L - \vec{c}$ and $\beta'_R = \beta_R - \vec{c}$.*

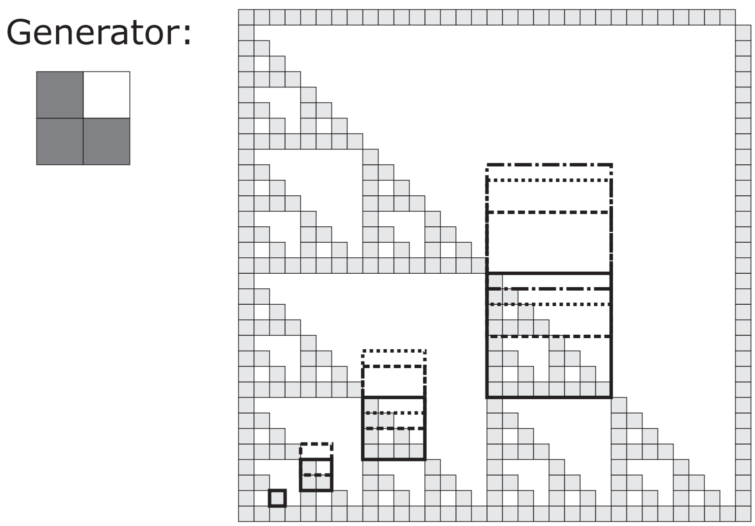

Figure 9: In each stage of the Sierpinski triangle, it is possible to define a sequence of closed-rectangular window movies, with the following properties: the number of window movies in the sequence is proportional to the stage number and the set of points contained in each window is unique.

Basically, the reason we do not use the standard WML to prove Theorem 3.6 is because we simply are not able to devise a unified strategy for finding two closed-rectangular window movies in a pier-fractal-shaped assembly that (1) have equivalent (up to translation) bond-forming submovies and (2) contain different sub-shapes of the assembly. On the one hand, it is trivial to find two such closed-rectangular window movies in a pier-fractal-shaped assembly whose sub-shapes are equal. But this does not help us derive the contradiction that we need to prove Theorem 3.6. On the other hand, it is also trivial to find two closed-rectangular window movies that contain different sub-shapes of the assembly, but, as a result of the self-similarity of pier fractals, do not have equivalent (up to translation) bond-forming submovies, at which point the conditions of the hypothesis of the standard WML are no longer satisfied.

In our attempts to resolve this dilemma, we investigated the use of an infinite-open window strategy, as opposed to a closed-rectangular window strategy. But this approach has its own set of technical challenges. Fortunately, these challenges can be dismissed! One must simply observe that, in order to prove Theorem 3.6, one does not need the "two-way-assembly-replacement" power offered by the conclusion of the standard WML. In fact, in order to derive a contradiction to prove Theorem 3.6, one merely needs to be able to replace one portion of a tile assembly with another portion in a strictly "one-way" fashion, i.e., the part of the tile assembly being used to replace another part does not need to be able to be replaced by the part of the tile assembly it is replacing. Thus, we weaken the conclusion

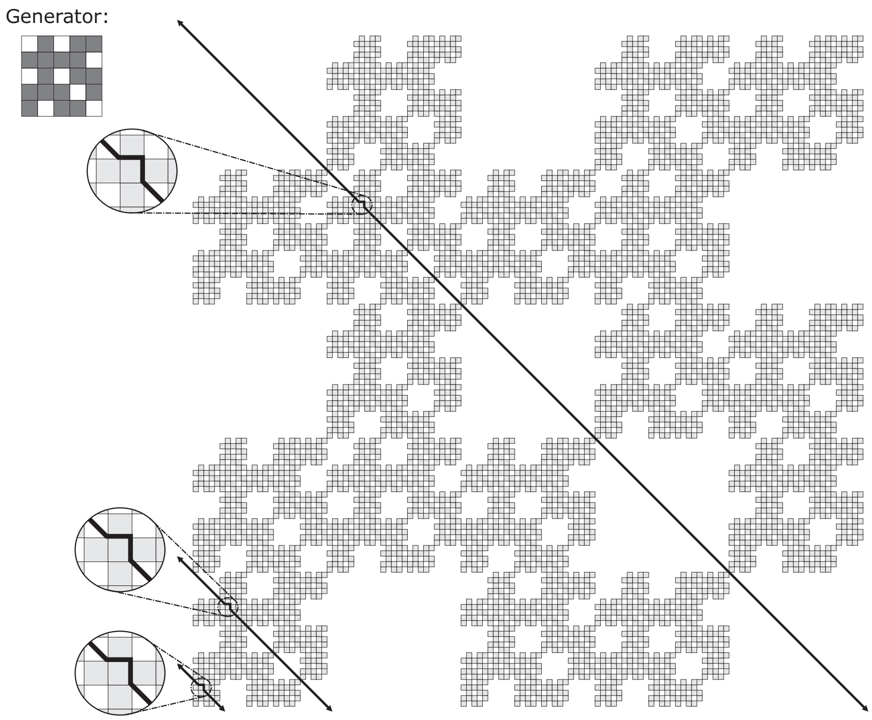

Figure 10: A generator for a pier fractal and the first three stages of an unscaled version of it. Note that it is possible to apply the standard WML to this pier fractal using infinite-open windows.

and strengthen the hypothesis of the standard WML to get the Closed WML, which turns out to be much more accommodating to a unified, closed-rectangular window proof technique for pier fractals.

It appears that the hypothesis of the standard WML, unlike that of the Closed WML, is too strong to be able to "handle" all pier fractals under a unified closed-rectangular window proof technique. However, it is worthy of note that in some special cases, it is possible to use the standard WML to prove that certain pier fractals do not strictly self-assemble. For example, it is possible to prove that the Sierpinski triangle does not strictly self-assemble at any positive scale factor (see Figure 9 for the proof idea). Next, consider the tree fractal defined by the generator given in Figure 10. In this case, it is possible to apply the standard WML using an open-infinite window proof technique (informally depicted in Figure 10). Unfortunately, depending on the geometry of the particular fractal, neither of the previous two applications of the standard WML, either with closed-rectangular or open-infinite windows, immediately generalizes to even the set of all tree fractals, which is a strict sub-class of pier fractals. Even more troubling, we suspect that, for the pier fractal whose generator is shown in Figure 11, it is not possible to apply the standard WML, with windows of any shape, to prove that it does not strictly self-assemble at any positive scale factor.

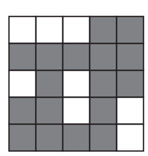

Figure 11: How can one apply the standard WML to prove that any scaled version of the pier fractal with this generator does not strictly self-assemble?

5 Conclusion

In this paper, we made three contributions. First, we gave a new characterization of tree fractals in terms of simple geometric properties of their generator (see Section 6). Second, we proved a new variant of the Window Movie Lemma in [5], which we call the "Closed Window Movie Lemma" (see Section 2.3). Third, we proved that no scaled-up version of any discrete self-similar pier fractal strictly self-assembles in the aTAM (see Section 3.2).

As we pointed out in Section 3.3.2, the scope of applicability of the Closed Window Movie Lemma is much wider than the class of pier fractals. Recall that Corollary 3.14 applies the Closed Window Movie Lemma to discrete self-similar fractals with no pier-like sub-configurations and an arbitrary number of vertical and horizontal bridges. In future work, we would like to provide a characterization of the class of all fractals to which the Closed Window Movie Lemma applies, that is, a strict super-class of the class of pier fractals. In addition, it would be satisfying to find a crisp characterization of the differences (if any) between the scope of applicability of the standard WML and that of the Closed WML. For instance, we would like to prove our conjecture that it is not possible to use the standard WML to prove that any scaled version of the pier fractal whose generator is shown in Figure 11 does not strictly self-assemble in the aTAM.

Acknowledgments

We would like to thank Kimberly Barth and Paul Totzke for contributing to an earlier proof of Corollary 3.9 that appeared in [1] and that we generalized to obtain the proof of our main result in this work.

References

[1] Kimberly Barth, David Furcy, Scott M. Summers, and Paul Totzke, *Scaled tree fractals do not strictly self-assemble*, Proceedings of the 13th International Conference on Unconventional Computation and Natural Computation (UCNC 2014), to appear.

[2] David Doty, Xiaoyang Gu, Jack H. Lutz, Elvira Mayordomo, and Philippe Moser, *Zeta-Dimension*, Proceedings of the Thirtieth International Symposium on Mathematical Foundations of Computer Science, Springer-Verlag, 2005, pp. 283–294.

[3] James I. Lathrop, Jack H. Lutz, and Scott M. Summers, *Strict self-assembly of discrete Sierpinski triangles*, Theoretical Computer Science **410** (2009), 384–405.

[4] Jack H. Lutz and Brad Shutters, *Approximate self-assembly of the sierpinski triangle*, Theory Comput. Syst. **51** (2012), no. 3, 372–400.

[5] P.-E. Meunier, M. J. Patitz, S. M. Summers, G. Theyssier, A. Winslow, and D. Woods, *Intrinsic universality in tile self-assembly requires cooperation*, Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2014), 2014, pp. 752–771.

[6] Matthew J. Patitz and Scott M. Summers, *Self-assembly of discrete self-similar fractals*, Natural Computing **1** (2010), 135–172.

[7] Paul W. K. Rothemund, *Theory and experiments in algorithmic self-assembly*, Ph.D. thesis, University of Southern California, December 2001.

[8] Nadrian C. Seeman, *Nucleic-acid junctions and lattices*, Journal of Theoretical Biology **99** (1982), 237–247.

[9] ______, *De novo design of sequences for nucleic acid structural engineering*, Journal of Biomolecular Structural Dynamics **8** (1990), 573–581.

[10] Hao Wang, *Proving theorems by pattern recognition – II*, The Bell System Technical Journal **XL** (1961), no. 1, 1–41.

[11] Erik Winfree, *Algorithmic self-assembly of DNA*, Ph.D. thesis, California Institute of Technology, June 1998.

6 Appendix

Definition 6.1. If G is the generator of any g-discrete self-similar fractal, then the *interior* of G is $G \cap (\mathbb{N}_{g-1} \times \mathbb{N}_{g-1})$.

Lemma 6.2. *Let G be any finite subset of $\mathbb{N}^2$ that has at least one connected h-bridge. If G contains a connected component $C \subset G$ such that $C \cap (\mathbb{N} \times \{t_G\}) \neq \varnothing$ and $C \cap (\{l_G\} \times \mathbb{N}) = \varnothing$, then there exists a point $\vec{x}_N \in G \backslash C$ such that $N(\vec{x}_N) \notin G$ and $\vec{x}_N \notin \mathbb{N} \times \{t_G\}$.*

Proof. Let h be a connected h-bridge in G and let π be a connected component in G that contains a path connecting the two points in h. Since π connects the leftmost and rightmost columns of G and C does not contain any point in the leftmost column of G, $C \cap \pi = \varnothing$. Since C is a connected component that extends vertically from row $t_C = t_G$ down to row b_C and $C \cap \pi = \varnothing$, π must go around (and below) C. Furthermore, no point in C is adjacent to any point in π. Let $\vec{p}$ denote a bottommost point (x, b_C) in C, with $l_G < x \leq r_G$. Let $\vec{q}$ denote the topmost point (x,y) in $\pi \cap (\{x\} \times \mathbb{N}_{b_C})$. Note that $\vec{p}$ and $\vec{q}$ are in the same column and that $\vec{p}$ is above (but not adjacent to) $\vec{q}$, that is, $y < b_C - 1$. Furthermore, $N(\vec{q}) \notin G$. Since $\vec{q} \in \pi \subset G$ and $\vec{q} \notin C$, $\vec{q} \in G \backslash C$. Furthermore, since $\vec{q} = (x,y)$ and $y < b_C - 1 < b_C \leq t_C = t_G$, $\vec{q} \notin \mathbb{N} \times \{t_G\}$. In conclusion, $\vec{q}$ exists and is a candidate for the role of x_N. □

Lemma 6.3. *Let G be any finite subset of $\mathbb{N}^2$ that has at least one connected v-bridge. If G contains a connected component $C \subset G$ such that $C \cap (\{r_G\} \times \mathbb{N}) \neq \varnothing$, $C \cap (\mathbb{N} \times \{t_G\}) \neq \varnothing$ and $C \cap (\mathbb{N} \times \{b_G\}) = \varnothing$, then there exists a point $\vec{x}_{NE} \in G \backslash C$ such that $E(\vec{x}_{NE}) \notin G$, $\vec{x}_{NE} \in \mathbb{N} \times \{t_G\}$ and $\vec{x}_{NE} \notin \{r_G\} \times \mathbb{N}$.*

Proof. Let v be a connected v-bridge in G and let π be a connected component in G that contains a path connecting the two points in v. Let π_t denote the set $\pi \cap (\mathbb{N} \times \{t_G\})$. Since this set cannot be empty, let us call its rightmost point $\vec{p} = (x_\pi, t_G)$. Similarly, let C_t denote $C \cap (\mathbb{N} \times \{t_G\})$. Since this set cannot be empty, let us call its leftmost point $\vec{q} = (x_C, t_G)$.

Since π connects the topmost and bottommost rows of G and C does not contain any point in the bottommost row of G, $C \cap \pi = \varnothing$. This, together with the fact that C contains a path from the topmost row to the rightmost column of G (that is, C "cuts off" the subset of G that lies to the north-east of C from the rest of G), implies that each point in π_t must appear to the left of all the points in C_t, namely $x_\pi < x_C$. In fact, since π and C cannot be connected, $\vec{p}$ and $\vec{q}$ cannot be adjacent, i.e., $x_\pi < x_C - 1$. Therefore, $\vec{p}$ and $\vec{q}$ are both in the topmost row of G (thus $\vec{p} \in \mathbb{N} \times \{t_G\}$) and $\vec{p}$ is to the left of $\vec{q}$ (thus $\vec{p} \notin \{r_G\} \times \mathbb{N}$). Finally, by construction, $\vec{p} \in G \backslash C$ and $E(\vec{p}) \notin G$. In conclusion, $\vec{p}$ is a candidate for the role of $\vec{x}_{NE}$. □

Lemma 6.4. *Let G be any finite subset of $\mathbb{N}^2$ that has at least one connected v-bridge. If G contains a connected component $C \subset G$ such that $C \cap (\{r_G\} \times \mathbb{N}) \neq \varnothing$ and $C \cap (\mathbb{N} \times \{b_G\}) = \varnothing$, then there exists a point $\vec{x}_E \in G \backslash C$ such that $E(\vec{x}_E) \notin G$ and $\vec{x}_E \notin \{r_G\} \times \mathbb{N}$.*

Proof. Let v be a connected v-bridge in G and let π be a connected component in G that contains a path connecting the two points in v. Since π connects the topmost and bottommost rows of G and C does not contain any point in the bottommost row of G, $C \cap \pi = \varnothing$. Since C is a connected component that extends horizontally from column l_C to column $r_C = r_G$ and $C \cap \pi = \varnothing$, π must go around (and to the left of) C. Furthermore, no point in C is adjacent to any point in π. Let $\vec{p}$ denote a leftmost point (l_C, y) in C, with $b_G < y \leq t_G$. Let $\vec{q}$ denote the rightmost point (x,y) in $\pi \cap (\mathbb{N}_{l_C} \times \{y\})$. Note that $\vec{p}$ and $\vec{q}$ are in the same row and that $\vec{q}$ is to the left of (but not adjacent to) $\vec{p}$, that is, $x < l_C - 1$. Furthermore, $E(\vec{q}) \notin G$. Since $\vec{q} \in \pi \subset G$ and $\vec{q} \notin C$, $\vec{q} \in G \backslash C$. Furthermore, since $\vec{q} = (x,y)$ and $x < l_C - 1 < l_C \leq r_C = r_G$, $\vec{q} \notin \{r_G\} \times \mathbb{N}$. In conclusion, $\vec{q}$ exists and is a candidate for the role of x_E. □

Lemma 6.5. *Let $\mathbf{X} = \bigcup_{i=1}^{\infty} X_i$ be a g-discrete self-similar fractal with generator G. If $\mathbf{X}$ is a tree, then G must have at least one connected h-bridge and at least one connected v-bridge.*

Proof. In this proof, we assume only that G does not have a connected h-bridge and reach a contradiction. We omit the symmetric reasoning that would allow us to prove that G must contain at least one connected v-bridge. Together, these two subproofs establish the fact that G must ontain at least one connected h-bridge and at least one connected v-bridge.

Assume that G does not have a connected h-bridge. We consider two cases characterized by the number of points in the leftmost column of G.

Case 1: $|G \cap (\{0\} \times \mathbb{N})| = g$. Then, the following three propositions hold:

(a) For every point $(1,y) \in G$, $N(1,y) \notin G$. Indeed, if $(1,y) \in G$ and $N(1,y) \in G$, then $\{(0,y), N(0,y), N(1,y), (1,y)\} \subset$ **X** would constitute a cycle in **X**, which contradicts the fact that **X** is a tree.

(b) For every point $(1,y) \in G$, $S(1,y) \notin G$. The justification is similar to the one for (a) above.

(c) $(1,g-1) \in G \Rightarrow (1,0) \notin G$. Indeed, if $(1,0) \in G$ and $(1,g-1) \in G$, then $\{(0,g-1), N(0,g-1), N(1,g-1), (1,g-1)\} \subset$ **X** would constitute a cycle in **X**, which contradicts the fact that **X** is a tree.

We will now prove that there is no path in **X** from the origin to any point $(x,y) \in$ **X** with $x \geq 2g$. If there were such a path π, it would include at least one pair of consecutive points $(2g-1,y')$ and $(2g,y')$. Let us consider the first such pair in π and let $\lfloor \frac{y'}{g} \rfloor = a$. Then $(2g-1,y') \in G + (g,ag)$. Since this copy of G belongs to the second column of copies of G in X_2, we can use the conjunction of propositions (a), (b) and (c) above to infer that $\mathbf{X} \cap (G + (g,(a-1)g) = \varnothing$ and $\mathbf{X} \cap (G + (g,(a+1)g) = \varnothing$. Therefore, π must contain a sub-path π' from the leftmost column of $G + (g,ag)$ to $(2g-1,y')$, that is, a path from (g,y'') to $(2g-1,y')$, for $ag \leq y'' < (a+1)g$. But since the leftmost column of $G + (g,ag)$ contains g points, there must be a (vertical) path from (g,y') to (g,y'') fully contained in the leftmost column of $G + (g,ag)$. Therefore, by concatenation of this path to π', $G + (g,ag)$ must contain a path from (g,y') to $(2g-1,y')$. But this path would be a connected h-bridge of $G + (g,ag)$, which would imply that G contains a connected h-bridge. So we can conclude that there is no path in **X** from the origin to any point east of the line $x = 2g-1$. Since **X** contains an infinite number of points in this region of $\mathbb{N}^2$, **X** cannot be connected, which is impossible since **X** is a tree. This contradiction implies that G must contain at least one connected h-bridge.

Case 2: $|G \cap (\{0\} \times \mathbb{N})| < g$. Since G does not have a connected h-bridge, one can show via a case analysis that either **X** is disconnected or contains a cycle. However, both of these scenarios are impossible since **X** is a tree.

□

Notation. Let $c, s \in \mathbb{Z}^+$ and $1 < g \in \mathbb{N}$. Let $e, f \in \mathbb{N}_g$. We use $S_s^c(e,f)$ to denote $\{0,1,\ldots,cg^{s-1}-1\}^2 + cg^{s-1}(e,f)$.

Notation. Let $1 < g \in \mathbb{N}$. Let $\mathbf{X} = \bigcup_{i=1}^{\infty} X_i$ be a g-discrete self-similar fractal. If $s \in \mathbb{Z}^+$, we use $P_{\mathbf{X}}(s)$ to denote the property: " X_s is a tree and $nhb_{X_s} = nvb_{X_s} = 1$".

Lemma 6.6. *Let $1 < g \in \mathbb{N}$. If **X** is a g-discrete self-similar fractal, then $P_{\mathbf{X}}(i) \Rightarrow P_{\mathbf{X}}(i+1)$ for $i \in \mathbb{Z}^+$.*

Proof. Let **X** be any g-discrete self-similar fractal. Let $i \in \mathbb{Z}^+$. We will abbreviate $X_i \cap S_i^1(x,y)$ and $X_{i+1} \cap S_{i+1}^1(x,y)$ to $U(x,y)$ and $V(x,y)$, respectively, where $x,y \in \mathbb{N}_g$. The definition of **X** implies that the following proposition, which we refer to as $(*)$, is true: "Every non-empty V subset of X_{i+1} is a translated copy of X_i".

Assume that $P_{\mathbf{X}}(i)$ holds.

First, we prove that X_{i+1} is connected. Pick any two distinct points $\vec{p}$ and $\vec{q}$ in X_{i+1}. If $\vec{p}$ and $\vec{q}$ belong to the same V subset of X_{i+1}, then there is a simple path from $\vec{p}$ to $\vec{q}$ (because of $(*)$ and the fact that X_i is connected, by $P_{\mathbf{X}}(i)$). If $\vec{p}$ and $\vec{q}$ belong to two distinct V subsets of X_{i+1}, say, $V(x_0,y_0)$ and $V(x_k,y_k)$, then consider the corresponding two U subsets $U(x_0,y_0)$ and $U(x_k,y_k)$ of X_i, neither of which can be empty. $P_{\mathbf{X}}(i)$ implies that there exists a simple path from any point in $U(x_0,y_0)$ to any point in $U(x_k,y_k)$. Assume that this path goes through the following sequence P_i of U subsets of X_i: $U(x_0,y_0), U(x_1,y_1), \ldots, U(x_{k-1},y_{k-1}), U(x_k,y_k)$. $P_{\mathbf{X}}(i)$ and $(*)$ together imply that each one of the corresponding V subsets of X_{i+1}, i.e., $V(x_0,y_0)$, $\ldots$, $V(x_k,y_k)$, is connected and contains a connected h-bridge and a connected v-bridge. Furthermore, since any pair of consecutive U subsets in P_i are adjacent in X_i, the same is true of the V subsets of X_{i+1} in the sequence P_{i+1}: $V(x_0,y_0), V(x_1,y_1), \ldots, V(x_{k-1},y_{k-1}), V(x_k,y_k)$. Since, for $i \in \mathbb{N}_k$, $V(x_i,y_i)$ is adjacent to $V(x_{i+1},y_{i+1})$ and each one of these subsets is connected and has at least one horizontal bridge and one vertical bridge, there must be at least one simple path from any point in $V(x_0,y_0)$ to any point in $V(x_k,y_k)$. Therefore, there exists a simple path between $\vec{p} \in V(x_0,y_0)$ and $\vec{q} \in V(x_k,y_k)$. Since this is true for any two distinct points $\vec{p}$ and $\vec{q}$ in X_{i+1}, X_{i+1} is connected.

Second, we prove that $nhb_{X_{i+1}} = nvb_{X_{i+1}} = 1$. Since the reasoning is similar for both horizontal and vertical bridges, we only deal with $nhb_{X_{i+1}}$ here. By $P_{\mathbf{X}}(i)$, X_i contains exactly one horizontal bridge. Therefore, there are exactly two subsets of X_i of the form $U(0,y)$ and $U(g-1,y)$, for some y in $\mathbb{N}_g$, such that there exist exactly two points $\vec{p} = (x_p,y_p)$ in $U(0,y)$ and $\vec{q} = (x_q,y_q)$ in $U(g-1,y)$ with $y_p = y_q$. Now consider $V(0,y)$ and $V(g-1,y)$. Since each one of these subsets of X_{i+1} is a translated copy of X_i, the westmost column of $V(0,y)$ is identical to the westmost column of X_i and the eastmost column of $V(g-1,y)$ is identical to the eastmost column of X_i. Therefore, the number of horizontal bridges in X_{i+1} that belong to $V(0,y) \cup V(g-1,y)$ is equal to $nhb_{X_i} = 1$. In other words, $nhb_{X_{i+1}} \geq 1$. Since both X_i

and X_{i+1} are built out of copies of their preceding stage according to the same pattern (namely the generator of $\mathbf{X}$) and we argued above that the only horizontal bridges in X_i belong to $U(0,y) \cup U(g-1,y)$, the horizontal bridges in X_{i+1} can only belong to the subsets $V(0,y)$ and $V(g-1,y)$. In other words, $nhb_{X_{i+1}} \leq 1$. Finally, $nhb_{X_{i+1}} = 1$.

Third, we prove that X_{i+1} is acyclic. For the sake of obtaining a contradiction, assume that there exists a simple cycle C in X_{i+1}. Let the sequence P_{i+1} of adjacent V subsets that C traverses be $V(x_0,y_0), \ldots, V(x_k,y_k)$. If P_{i+1} has length one, then C is contained in a single (translated) copy of X_i (by $(*)$), which contradicts the fact that X_i is acyclic (by $P_{\mathbf{X}}(i)$). Otherwise, C traverses all of the V subsets in P_{i+1}, whose length is at least two. Following the same reasoning as above, there must exist a corresponding sequence P_i, namely $U(x_0,y_0), \ldots, U(x_k,y_k)$, of U subsets in X_i. Since each subset in this sequence is connected, contains one horizontal bridge and one vertical bridge (by $P_{\mathbf{X}}(i)$), and is adjacent to its neighbors in the sequence, the union of these subsets forms a connected component that must contain at least one simple cycle, which contradicts the fact that X_i is a tree (by $P_{\mathbf{X}}(i)$). In all cases, we reached a contradiction. Therefore, X_{i+1} is acyclic.

Finally, since X_{i+1} is a tree and $nhb_{X_{i+1}} = nvb_{X_{i+1}} = 1$, $P_{\mathbf{X}}(i+1)$ holds.

□

Theorem 6.2. $\mathbf{T} = \bigcup_{i=1}^{\infty} T_i$ *is a g-discrete self-similar tree fractal, for some* $g > 1$*, with generator G if and only if*

a. G is a tree, and

b. $nhb_G = nvb_G = 1$

Proof. Assume that $\mathbf{T}$ is a g-discrete self-similar tree fractal with generator G. Thus, $\mathbf{T}$ is acyclic and connected. If $nhb_G < 1$ or $nvb_G < 1$, then $\mathbf{T}$ is trivially disconnected. Thus, $nhb_G \geq 1$, $nvb_G \geq 1$.

Since $\mathbf{T}$ is acyclic, G must be acyclic as well, for if G were not acyclic, then $\mathbf{T}$ would not be, as $G \subset \mathbf{T}$.

We will now show that G is connected. To see this, assume that G is disconnected. First, note that, if G has a connected component contained strictly within its interior, then $\mathbf{T}$ is trivially disconnected.

Second, if G is disconnected, then G contains a connected component that touches at most two sides of G. To see this, note that Lemma 6.5 says that G has at least one connected h-bridge and at least one connected v-bridge. If G had a connected component, say C, that touched three or more sides of G, then due to the existence of at least one connected h-bridge and at least one connected v-bridge, G would necessarily have another connected component, say C', that could only touch at most two sides of G.

We now proceed with a case analysis based on the number of sides of G that the connected component touches (one or two sides) and the relative positions of these sides (adjacent or opposite sides).

Case 1: Assume that G has a connected component, say C, that does not contain the origin but does contain points in the northmost row and eastmost column of G (and there is no path in G from the origin to any point in C). We will call this case "NE". Lemma 6.5 says that G has at least one connected h-bridge and at least one connected v-bridge. Therefore, Lemma 6.2 says that G has a north-free point not in the northmost row, say $\vec{x}_N$, and Lemma 6.3 says that G has an east-free point in the northmost row but not in the eastmost column, say $\vec{x}_{NE}$. Let $C' = C + g^2\vec{x}_N + g\vec{x}_{NE}$. Since $\vec{x}_N$ is north-free and not in the northmost row of G, $N(\vec{x}_N) \notin G$, whence $\mathbf{T} \cap \left(\left\{0,\ldots,g^2-1\right\}^2 + g^2 N(\vec{x}_N) \right) = \varnothing$. Since $\vec{x}_{NE}$ is in the northmost row, this means the northmost point in every column of C' is north-free in $\mathbf{T}$. Since $\vec{x}_{NE}$ is not in the eastmost column of G, $E(\vec{x}_{NE}) \notin G$, whence $\mathbf{T} \cap \left(g^2\vec{x}_N + \left(\{0,\ldots,g-1\}^2 + gE(\vec{x}_{NE}) \right) \right) = \varnothing$. This means that the eastmost point in every row of C' is east-free in $\mathbf{T}$. We also know that the westmost point in every row of C is west-free in G and the southmost point in every column of C is south-free in G, therefore the westmost point in every row of C' is west-free in $\mathbf{T}$ and the southmost point in every column of C' is south-free in $\mathbf{T}$. Thus, there is no path in $\mathbf{T}$ from any point in C' to the origin, which contradicts the assumption that $\mathbf{T}$ is connected. The "NW" and "SE" cases can be handled with a similar argument. Note that, in the "SW" case, the connected component is contained strictly within the interior of the generator. Such situations were handled above.

Case 2: Assume that G has a connected component, say C, that contains points in the eastmost column but does not contain the origin nor points in the westmost column of G nor the northmost or southmost rows of G. This is the "E" case. In this case, Lemma 6.4 says that there is an east-free point in G that is not in the eastmost column of G. Call this point $\vec{x}_E$ and define $C' = C + g\vec{x}_E$. Following directly from the definition of the "E" case, the northmost point in every column of C is north-free in G, the southmost point in every column of C is south-free in G and the westmost point in every row of C is west-free in G. From the definition of C' and the fact that $\vec{x_E}$ is east-free, it follows that the eastmost (respectively, westmost) point in every row of C' is east-free (respectively, west-free) in $\mathbf{T}$. Similarly, the northmost (respectively, southmost) point in every column of C' is north-free (respectively, south-free) in $\mathbf{T}$. Therefore, there is no path in $\mathbf{T}$ from any point in C' to the origin, which contradicts the assumption that $\mathbf{T}$ is connected. The "N" case can be handled with a similar argument. Note that, in the "W" and "S" cases, the connected component is contained strictly within the interior of the generator. Such situations were handled above.

Case 3: Assume that G has a connected component, say C, that contains points in both the eastmost and westmost columns of G. This is the "EW" case. In this case, since G contains at least one connected v-bridge, C must contain all connected v-bridges of G (since C must have a non-empty intersection with each connected v-bridge in G). Therefore, C touches all four sides of G. If C contains the origin, then there must exist another disconnected component, say C', that does not contain the origin and C' must belong to one of the previous cases. If C does not contain the origin, then the origin itself must be part of a connected component that is not connected to C nor to any other point in $\mathbf{T}$, which contradicts the assumption that $\mathbf{T}$ is connected. The "NS" case can be handled with a similar argument.

Therefore, in all cases, G is connected. Since we argued above that G is acyclic, we may conclude that G is a tree.

Finally, since G is a tree, it must be the case that $nvb_G \leq 1$ and $nhb_G \leq 1$, otherwise $\mathbf{T}$ would contain a cycle, whence $nvb_G = nhb_G = 1$.

Now we prove that if G is a tree and $nhb_G = nvb_G = 1$, then $\mathbf{T}$ is a tree.

Assume that G is a tree and $nhb_G = nvb_G = 1$. Then $P_{\mathbf{T}}(1)$ holds (since $G = T_1$). Furthermore, by Lemma 6.6, $P_{\mathbf{T}}(i) \Rightarrow P_{\mathbf{T}}(i+1)$ for $i \in \mathbb{Z}^+$. Thus, by induction, $P_{\mathbf{T}}(i)$ holds for $i \in \mathbb{Z}^+$, which implies that each stage in $\mathbf{T}$ is a tree. We now prove that $\mathbf{T}$ is a tree.

First, $\mathbf{T}$ is connected, since each stage of $\mathbf{T}$ is connected.

Second, we prove that $\mathbf{T}$ cannot contain a cycle. Assume, for the sake of obtaining a contradiction, that there exist two distinct points $\vec{p}$ and $\vec{q}$ in $\mathbf{T}$ such that there exist two distinct simple paths from $\vec{p}$ to $\vec{q}$. Since both of these paths must be finite, the cycle that they form must also be finite. Therefore, this cycle must be fully contained in some stage of $\mathbf{T}$, which contradicts the fact that all stages of $\mathbf{T}$ are trees.

In conclusion, $\mathbf{T}$ is connected and acyclic, and is thus a tree. $\square$