Electronic Colloquium on Computational Complexity, Report No. 35 (2008)

Strict Self-Assembly of Discrete Sierpinski Triangles

James I. Lathrop[*], Jack H. Lutz[†], and Scott M. Summers.[‡]

Abstract

Winfree (1998) showed that discrete Sierpinski triangles can self-assemble in the Tile Assembly Model. A striking molecular realization of this self-assembly, using DNA tiles a few nanometers long and verifying the results by atomic-force microscopy, was achieved by Rothemund, Papadakis, and Winfree (2004).

Precisely speaking, the above self-assemblies tile completely filled-in, two-dimensional regions of the plane, with labeled subsets of these tiles representing discrete Sierpinski triangles. This paper addresses the more challenging problem of the *strict self-assembly* of discrete Sierpinski triangles, i.e., the task of tiling a discrete Sierpinski triangle and nothing else.

We first prove that the standard discrete Sierpinski triangle *cannot* strictly self-assemble in the Tile Assembly Model. We then define the *fibered Sierpinski triangle*, a discrete Sierpinski triangle with the same fractal dimension as the standard one but with thin fibers that can carry data, and show that the fibered Sierpinski triangle strictly self-assembles in the Tile Assembly Model. In contrast with the simple XOR algorithm of the earlier, non-strict self-assemblies, our strict self-assembly algorithm makes extensive, recursive use of optimal counters, coupled with measured delay and corner-turning operations. We verify our strict self-assembly using the local determinism method of Soloveichik and Winfree (2007).

1 Introduction

Structures that self-assemble in naturally occurring biological systems are often fractals of low dimension, by which we mean that they are usefully modeled as fractals and that their fractal dimensions are less than the dimension of the space or surface that they occupy. The advantages of such fractal geometries for materials transport, heat exchange, information processing, and robustness imply that structures engineered by nanoscale self-assembly in the near future will also often be fractals of low dimension.

The simplest mathematical model of nanoscale self-assembly is the Tile Assembly Model (TAM), an extension of Wang tiling [17, 18] that was introduced by Winfree [20] and refined by Rothemund and Winfree [13, 12]. (See also [1, 11, 16].) This elegant model, which is described in section 2, uses tiles with various types and strengths of "glue" on their edges as abstractions of molecules adsorbing to a growing structure. (The tiles are squares in the two-dimensional TAM, which is most widely used, cubes in the three-dimensional TAM, etc.) Despite the model's deliberate oversimplification of molecular geometry and binding, Winfree [20] proved that the TAM is computationally universal in two or more dimensions. Self-assembly in the TAM can thus be directed algorithmically.

This paper concerns the self-assembly of fractal structures in the Tile Assembly Model. The typical test bed for a new research topic involving fractals is the Sierpinski triangle, and this is certainly the case for fractal self-assembly. Specifically, Winfree [20] showed that the *standard discrete Sierpinski triangle* **S**, which is illustrated in Figure 1, self-assembles from a set of seven tile types in the Tile Assembly Model.

*Department of Computer Science, Iowa State University, Ames, IA 50014, USA. jil@cs.iastate.edu.

†Department of Computer Science, Iowa State University, Ames, IA 50014, USA. lutz@cs.iastate.edu. This author's research was supported in part by National Science Foundation Grants 0344187, 0652569, and 0728806 and in part by Spanish Government MEC Project TIN 2005-08832-C03-02.

‡Department of Computer Science, Iowa State University, Ames, IA 50014, USA. summers@cs.iastate.edu. This author's research was supported in part by NSF-IGERT Training Project in Computational Molecular Biology Grant number DGE-0504304

ISSN 1433-8092

Formally, **S** is a set of points in the discrete Euclidean plane $\mathbb{Z}^2$. The obvious and well-known resemblance between **S** and the Sierpinski triangle in $\mathbb{R}^2$ that is studied in fractal geometry [8] is a special case of a general correspondence between "discrete fractals" and "continuous fractals" [19]. Continuous fractals are typically bounded (in fact, compact) and have intricate structure at arbitrarily small scales, while discrete fractals like **S** are unbounded and have intricate structure at arbitrarily large scales.

A striking molecular realization of Winfree's self-assembly of **S** was reported in 2004. Using DNA double-crossover molecules (which were first synthesized in pioneering work of Seeman and his co-workers [15]) to construct tiles only a few nanometers long, Rothemund, Papadakis and Winfree [14] implemented the molecular self-assembly of **S** with low enough error rates to achieve correct placement of 100 to 200 tiles, confirmed by atomic force microscopy (AFM). This gives strong evidence that self-assembly can be algorithmically directed at the nanoscale.

The abstract and laboratory self-assemblies of **S** described above are impressive, but they are not (nor were they intended or claimed to be) true fractal self-assemblies. Winfree's abstract self-assembly of **S** actually tiles an *entire quadrant* of the plane in such a way that five of the seven tile types occupy positions corresponding to points in **S**. Similarly, the laboratory self-assemblies tile completely filled-in, two-dimensional regions, with DNA tiles at positions corresponding to points of **S** marked by inserting hairpin sequences for AFM contrast. To put the matter figuratively, what self-assembles in these assemblies is not the fractal **S** but rather a two-dimensional canvas on which **S** has been painted.

In order to achieve the advantages of fractal geometries mentioned in the first paragraph of this paper, we need self-assemblies that construct fractal shapes *and nothing more*. Accordingly, we say that a set $F \subseteq \mathbb{Z}^2$ *strictly self-assembles* in the Tile Assembly Model if there is a (finite) tile system that eventually places a tile on each point of F and never places a tile on any point of the complement, $\mathbb{Z}^2 - F$. (This condition is defined precisely in section 2.)

The specific topic of this paper is the strict self-assembly of discrete Sierpinski triangles in the Tile Assembly Model. We present two main results on this topic, one negative and one positive.

Our negative result is that the standard discrete Sierpinski triangle **S** *cannot* strictly self-assemble in the Tile Assembly Model. That is, there is no tile assembly system that places tiles on all the points of **S** and on none of the points of $\mathbb{Z}^2 - \mathbf{S}$. This theorem appears in section 3. The key to its proof is an extension of the theorem of Adleman, Cheng, Goel, Huang, Kempe, Moisset de Espanés, and Rothemund [2] on the number of tile types required for a finite tree to self-assemble from a single seed tile at its root.

Our positive result is that a slight modification of **S**, the *fibered Sierpinski triangle* **T** illustrated in Figure 2, strictly self-assembles in the Tile Assembly Model. Intuitively, the fibered Sierpinski triangle **T** (defined precisely in section 4) is constructed by following the recursive construction of **S** but also adding a thin *fiber* to the left and bottom edges of each stage in the construction. These fibers, which carry data in an algorithmically directed self-assembly of **T**, have thicknesses that are logarithmic in the sizes of the corresponding stages of **T**. This means that **T** is visually indistinguishable from **S** at sufficiently large scales. Mathematically, it implies that **T** has the same fractal dimension as **S**.

Since our strict self-assembly must tile the set **T** "from within," the algorithm that directs it is perforce more involved than the simple XOR algorithm that directs Winfree's seven-tile-type, non-strict self-assembly of **S**. Our algorithm, which is described in section 5, makes extensive, recursive use of optimal counters [5], coupled with measured delay and corner-turning operations. It uses 51 tile types, but these are naturally partitioned into small functional groups, so that we can use Soloveichik and Winfree's local determinism method [16] to prove that **T** strictly self-assembles.

2 Preliminaries

2.1 Notation and Terminology

We work in the discrete Euclidean plane $\mathbb{Z}^2 = \mathbb{Z} \times \mathbb{Z}$. We write U_2 for the set of all *unit vectors*, i.e., vectors of length 1, in $\mathbb{Z}^2$. We regard the four elements of U_2 as (names of the cardinal) *directions* in $\mathbb{Z}^2$.

We write $[X]^2$ for the set of all 2-element subsets of a set X. All *graphs* here are undirected graphs, i.e., ordered pairs $G = (V, E)$, where V is the set of *vertices* and $E \subseteq [V]^2$ is the set of *edges*. A *cut* of a graph $G = (V, E)$ is a partition $C = (C_0, C_1)$ of V into two nonempty, disjoint subsets C_0 and C_1.

A *binding function* on a graph $G = (V, E)$ is a function $\beta : E \to \mathbb{N}$. (Intuitively, if $\{u, v\} \in E$, then $\beta(\{u, v\})$ is the strength with which u is bound to v by $\{u, v\}$ according to β. If β is a binding function on a graph $G = (V, E)$ and $C = (C_0, C_1)$ is a cut of G, then the *binding strength* of β on C is

$$\beta_C = \{\beta(e) \,|\, e \in E, e \cap C_0 \neq \emptyset, \text{ and } e \cap C_1 \neq \emptyset\}.$$

The *binding strength* of β on the graph G is then

$$\beta(G) = \min\{\beta_C \,|\, C \text{ is a cut of } G\}.$$

A *binding graph* is an ordered triple $G = (V, E, \beta)$, where (V, E) is a graph and β is a binding function on (V, E). If $\tau \in \mathbb{N}$, then a binding graph $G = (V, E, \beta)$ is τ-*stable* if $\beta(V, E) \geq \tau$.

A *grid graph* is a graph $G = (V, E)$ in which $V \subseteq \mathbb{Z}^2$ and every edge $\{\vec{m}, \vec{n}\} \in E$ has the property that $\vec{m} - \vec{n} \in U_2$. The *full grid graph* on a set $V \subseteq \mathbb{Z}^2$ is the graph $G^{\#}_V = (V, E)$ in which E contains *every* $\{\vec{m}, \vec{n}\} \in [V]^2$ such that $\vec{m} - \vec{n} \in U_2$.

We say that f is a *partial function* from a set X to a set Y, and we write $f : X \dashrightarrow Y$, if $f : D \to Y$ for some set $D \subseteq X$. In this case, D is the *domain* of f, and we write $D = \text{dom } f$.

All logarithms here are base-2.

2.2 The Tile Assembly Model

We review the basic ideas of the Tile Assembly Model. Our development largely follows that of [13, 12], but some of our terminology and notation are specifically tailored to our objectives. In particular, our version of the model only uses nonnegative "glue strengths", and it bestows equal status on finite and infinite assemblies. We emphasize that the results in this section have been known for years, e.g., they appear, with proofs, in [12].

Definition. A *tile type* over an alphabet Σ is a function $t : U_2 \to \Sigma^* \times \mathbb{N}$. We write $t = (\text{col}_t, \text{str}_t)$, where $\text{col}_t : U_2 \to \Sigma^*$, and $\text{str}_t : U_2 \to \mathbb{N}$ are defined by $t(\vec{u}) = (\text{col}_t(\vec{u}), \text{str}_t(\vec{u}))$ for all $\vec{u} \in U_2$.

Intuitively, a tile of type t is a unit square. It can be translated but not rotated, so it has a well-defined "side $\vec{u}$ " for each $\vec{u} \in U_2$. Each side $\vec{u}$ of the tile is covered with a "glue" of *color* $\text{col}_t(\vec{u})$ and *strength* $\text{str}_t(\vec{u})$. If tiles of types t and t' are placed with their centers at $\vec{m}$ and $\vec{m} + \vec{u}$, respectively, where $\vec{m} \in \mathbb{Z}^2$ and $\vec{u} \in U_2$, then they will *bind* with strength $\text{str}_t(\vec{u}) \cdot [\![t(\vec{u}) = t'(-\vec{u})]\!]$ where $[\![\phi]\!]$ is the *Boolean* value of the statement ϕ. Note that this binding strength is 0 unless the adjoining sides have glues of both the same color and the same strength.

For the remainder of this section, unless otherwise specified, T is an arbitrary set of tile types, and $\tau \in \mathbb{N}$ is the "temperature."

Definition. A T-*configuration* is a partial function $\alpha : \mathbb{Z}^2 \dashrightarrow T$.

Intuitively, a configuration is an assignment α in which a tile of type $\alpha(\vec{m})$ has been placed (with its center) at each point $\vec{m} \in \text{dom } \alpha$. The following data structure characterizes how these tiles are bound to one another.

Definition. The *binding graph of* a T-configuration $\alpha : \mathbb{Z}^2 \dashrightarrow T$ is the binding graph $G_\alpha = (V, E, \beta)$, where (V, E) is the grid graph given by

$$V = \text{dom } \alpha,$$

$$E = \left\{ \{\vec{m}, \vec{n}\} \in [V]^2 \middle| \vec{m} - \vec{n} \in U_n, \text{col}_{\alpha(\vec{m})}(\vec{n} - \vec{m}) = \text{col}_{\alpha(\vec{n})}(\vec{m} - \vec{n}), \text{and } \text{str}_{\alpha(\vec{m})}(\vec{n} - \vec{m}) > 0 \right\},$$

and the binding function $\beta : E \to \mathbb{Z}^+$ is given by

$$\beta\left(\{\vec{m}, \vec{n}\}\right) = \mathrm{str}_{\alpha(\vec{m})}\left(\vec{n} - \vec{m}\right)$$

for all $\{\vec{m}, \vec{n}\} \in E$.

Definition.

1. A T-configuration α is *τ-stable* if its binding graph G_α is τ-stable.
2. A *τ-T-assembly* is a T-configuration that is τ-stable. We write $\mathcal{A}_T^\tau$ for the set of all τ-T-assemblies.

Definition. Let α and α' be T-configurations.

1. α is a *subconfiguration* of α', and we write $\alpha \sqsubseteq \alpha'$, if $\mathrm{dom}\ \alpha \subseteq \mathrm{dom}\ \alpha'$ and, for all $\vec{m} \in \mathrm{dom}\ \alpha$, $\alpha(\vec{m}) = \alpha'(\vec{m})$.
2. α' is a *single-tile extension* of α if $\alpha \sqsubseteq \alpha'$ and $\mathrm{dom}\ \alpha' - \mathrm{dom}\ \alpha$ is a singleton set. In this case, we write $\alpha' = \alpha + (\vec{m} \mapsto t)$, where $\{\vec{m}\} = \mathrm{dom}\ \alpha' - \mathrm{dom}\ \alpha$ and $t = \alpha'(\vec{m})$.

Note that the expression $\alpha + (\vec{m} \mapsto t)$ is only defined when $\vec{m} \in \mathbb{Z}^2 - \mathrm{dom}\ \alpha$.

We next define the "τ-t-frontier" of a τ-T-assembly α to be the set of all positions at which a tile of type t can be "τ-stably added" to the assembly α.

Definition. Let $\alpha \in \mathcal{A}_T^\tau$.

1. For each $t \in T$, the *τ-t-frontier* of α is the set

$$\partial_t^\tau \alpha = \left\{ \vec{m} \in \mathbb{Z}^2 - \mathrm{dom}\ \alpha \;\middle|\; \sum_{\vec{u} \in U_2} \mathrm{str}_t(\vec{u}) \cdot [\![\alpha(\vec{m} + \vec{u})(-\vec{u}) = t(\vec{u})]\!] \geq \tau \right\}.$$

2. The *τ-frontier* of α is the set

$$\partial^\tau \alpha = \bigcup_{t \in T} \partial_t^\tau \alpha.$$

The following lemma shows that the definition of $\partial_t^\tau \alpha$ achieves the desired effect.

Lemma 2.1. Let $\alpha \in \mathcal{A}_T^\tau$, $\vec{m} \in \mathbb{Z}^2 - \mathrm{dom}\ \alpha$, and $t \in T$. Then $\alpha + (\vec{m} \mapsto t) \in \mathcal{A}_T^\tau$ if and only if $\vec{m} \in \partial_t^\tau \alpha$.

Notation. We write $\alpha \xrightarrow[\tau,T]{1} \alpha'$ (or, when τ and T are clear from context, $\alpha \xrightarrow{1} \alpha'$) to indicate that $\alpha, \alpha' \in \mathcal{A}_T^\tau$ and α' is a single-tile extension of α.

In general, self-assembly occurs with tiles adsorbing nondeterministically and asynchronously to a growing assembly. We now define assembly sequences, which are particular "execution traces" of how this might occur.

Definition. A *τ-T-assembly sequence* is a sequence $\vec{\alpha} = (\alpha_i \mid 0 \leq i < k)$ in $\mathcal{A}_T^\tau$, where $k \in \mathbb{Z}^+ \cup \{\infty\}$ and, for each i with $1 \leq i + 1 < k$, $\alpha_i \xrightarrow[\tau,T]{1} \alpha_{i+1}$.

Note that assembly sequences may be finite or infinite in length. Note also that, in any τ-T-assembly sequence $\vec{\alpha} = (\alpha_i \mid 0 \leq i < k)$, we have $\alpha_i \sqsubseteq \alpha_j$ for all $0 \leq i \leq j < k$.

Definition. The *result* of a τ-T-assembly sequence $\vec{\alpha} = (\alpha_i \mid 0 \leq i < k)$ is the unique T-configuration $\alpha = \mathrm{res}(\vec{\alpha})$ satisfying $\mathrm{dom}\ \alpha = \bigcup_{0 \leq i < k} \mathrm{dom}\ \alpha_i$ and $\alpha_i \sqsubseteq \alpha$ for each $0 \leq i < k$.

It is clear that $\mathrm{res}(\vec{\alpha}) \in \mathcal{A}_T^\tau$ for every τ-T-assembly sequence $\vec{\alpha}$.

Definition. Let $\alpha, \alpha' \in \mathcal{A}^{\tau}_{T}$.

1. A *τ-T-assembly sequence from α to α'* is a τ-T-assembly sequence $\vec{\alpha} = (\alpha_i \mid 0 \leq i < k)$ such that $\alpha_0 = \alpha$ and $\text{res}(\vec{\alpha}) = \alpha'$.
2. We write $\alpha \xrightarrow[\tau,T]{} \alpha'$ (or, when τ and T are clear from context, $\alpha \longrightarrow \alpha'$) to indicate that there exists a τ-T-assembly sequence from α to α'.

A routine dovetailing argument extends the following observation of [12] to assembly sequences that may have infinite length.

Theorem 2.2. The binary relation $\xrightarrow[\tau,T]{}$ is a partial ordering of $\mathcal{A}^{\tau}_{T}$.

Definition. An assembly $\alpha \in \mathcal{A}^{\tau}_{T}$ is *terminal* if it is a $\xrightarrow[\tau,T]{}$-maximal element of $\mathcal{A}^{\tau}_{T}$.

It is clear that an assembly α is terminal if and only if $\partial^{\tau}\alpha = \emptyset$.
We now note that every assembly is $\xrightarrow[\tau,T]{}$-bounded by (i.e., can lead to) a terminal assembly.

Lemma 2.3. For each $\alpha \in \mathcal{A}^{\tau}_{T}$, there exists $\alpha' \in \mathcal{A}^{\tau}_{T}$ such that $\alpha \xrightarrow[\tau,T]{} \alpha'$ and α' is terminal.

We now define tile assembly systems.

Definition.

1. A *generalized tile assembly system* (*GTAS*) is an ordered triple
$$\mathcal{T} = (T, \sigma, \tau),$$
where T is a set of tile types, $\sigma \in \mathcal{A}^{\tau}_{T}$ is the *seed assembly*, and $\tau \in \mathbb{N}$ is the *temperature*.
2. A *tile assembly system* (*TAS*) is a GTAS $\mathcal{T} = (T, \sigma, \tau)$ in which the sets T and dom σ are finite.

Intuitively, a "run" of a GTAS $\mathcal{T} = (T, \sigma, \tau)$ is any τ-T-assembly sequence $\vec{\alpha} = (\alpha_i \mid 0 \leq i < k)$ that begins with $\alpha_0 = \sigma$. Accordingly, we define the following sets.

Definition. Let $\mathcal{T} = (T, \sigma, \tau)$ be a GTAS.

1. The *set of assemblies produced by $\mathcal{T}$* is
$$\mathcal{A}[\mathcal{T}] = \left\{ \alpha \in \mathcal{A}^{\tau}_{T} \,\middle|\, \sigma \xrightarrow[\tau,T]{} \alpha \right\}.$$
2. The *set of terminal assemblies produced by $\mathcal{T}$* is
$$\mathcal{A}_{\Box}[\mathcal{T}] = \{ \alpha \in \mathcal{A}[\mathcal{T}] | \, \alpha \text{ is terminal} \}.$$

Definition. A GTAS $\mathcal{T} = (T, \sigma, \tau)$ is *directed* if the partial ordering $\xrightarrow[\tau,T]{}$ directs the set $\mathcal{A}[\mathcal{T}]$, i.e., if for each $\alpha, \alpha' \in \mathcal{A}[\mathcal{T}]$ there exists $\hat{\alpha} \in \mathcal{A}[\mathcal{T}]$ such that $\alpha \xrightarrow[\tau,T]{} \hat{\alpha}$ and $\alpha' \xrightarrow[\tau,T]{} \hat{\alpha}$.

We are using the terminology of the mathematical theory of relations here. The reader is cautioned that the term "directed" has also been used for a different, more specialized notion in self-assembly [3].
Directed tile assembly systems are interesting because they are precisely those tile assembly systems that produce unique terminal assemblies.

Theorem 2.4. A GTAS $\mathcal{T}$ is directed if and only if $|\mathcal{A}_{\Box}[\mathcal{T}]| = 1$.

In the present paper, we are primarily interested in the self-assembly of sets.

Definition. Let $\mathcal{T} = (T, \sigma, \tau)$ be a GTAS, and let $X \subseteq \mathbb{Z}^2$.

1. The set X *weakly self-assembles* in $\mathcal{T}$ if there is a set $B \subseteq T$ such that, for all $\alpha \in \mathcal{A}_\Box[\mathcal{T}]$, $\alpha^{-1}(B) = X$.
2. The set X *strictly self-assembles* in $\mathcal{T}$ if, for all $\alpha \in \mathcal{A}_\Box[\mathcal{T}]$, dom $\alpha = X$.

Intuitively, a set X weakly self-assembles in $\mathcal{T}$ if there is a designated set B of "black" tile types such that every terminal assembly of $\mathcal{T}$ "paints the set X - and only the set X - black". In contrast, a set X strictly self-assembles in $\mathcal{T}$ if every terminal assembly of $\mathcal{T}$ has tiles on the set X and only on the set X. Clearly, every set that strictly self-assembles in a GTAS $\mathcal{T}$ also weakly self-assembles in $\mathcal{T}$.

We now have the machinery to say what it means for a set in the discrete Euclidean plane to self-assemble in either the weak or the strict sense.

Definition. Let $X \subseteq \mathbb{Z}^2$.

1. The set X *weakly self-assembles* if there is a TAS $\mathcal{T}$ such that X weakly self-assembles in $\mathcal{T}$.
2. The set X *strictly self-assembles* if there is a TAS $\mathcal{T}$ such that X strictly self-assembles in $\mathcal{T}$.

Note that $\mathcal{T}$ is required to be a TAS, i.e., finite, in both parts of the above definition.

2.3 Local Determinism

The proof of our second main theorem uses the local determinism method of Soloveichik and Winfree [16], which we now review.

Notation. For each T-configuration α, each $\vec{m} \in \mathbb{Z}^2$, and each $\vec{u} \in U_2$,

$$\mathrm{str}_\alpha(\vec{m}, \vec{u}) = \mathrm{str}_{\alpha(\vec{m})}(\vec{u}) \cdot [\![\alpha(\vec{m})(\vec{u}) = \alpha(\vec{m} + \vec{u})(-\vec{u})]\!].$$

(The Boolean value on the right is 0 if $\{\vec{m}, \vec{m} + \vec{u}\} \nsubseteq$ dom α.)

Notation. If $\vec{\alpha} = (\alpha_i | 0 \le i < k)$ is a τ-T-assembly sequence and $\vec{m} \in \mathbb{Z}^2$, then the $\vec{\alpha}$-*index* of $\vec{m}$ is

$$i_{\vec{\alpha}}(\vec{m}) = \min\{ i \in \mathbb{N} \,| \vec{m} \in \mathrm{dom}\, \alpha_i \}.$$

Observation 2.5. $\vec{m} \in \mathrm{dom}\, \mathrm{res}(\vec{\alpha}) \Leftrightarrow i_{\vec{\alpha}}(\vec{m}) < \infty$.

Notation. If $\vec{\alpha} = (\alpha_i | 0 \le i < k)$ is a τ-T-assembly sequence, then, for $\vec{m}, \vec{m}' \in \mathbb{Z}^2$,

$$\vec{m} \prec_{\vec{\alpha}} \vec{m}' \Leftrightarrow i_{\vec{\alpha}}(\vec{m}) < i_{\vec{\alpha}}(\vec{m}').$$

Definition. (Soloveichik and Winfree [16]) Let $\vec{\alpha} = (\alpha_i | 0 \le i < k)$ be a τ-T-assembly sequence, and let $\alpha = \mathrm{res}(\vec{\alpha})$. For each location $\vec{m} \in \mathrm{dom}\, \alpha$, define the following sets of directions.

1. $\mathrm{IN}^{\vec{\alpha}}(\vec{m}) = \left\{ \vec{u} \in U_2 \,\middle|\, \vec{m} + \vec{u} \prec_{\vec{\alpha}} \vec{m} \text{ and } \mathrm{str}_{\alpha_{i_{\vec{\alpha}}(\vec{m})}}(\vec{m}, \vec{u}) > 0 \right\}$.
2. $\mathrm{OUT}^{\vec{\alpha}}(\vec{m}) = \left\{ \vec{u} \in U_2 \,\middle|\, -\vec{u} \in \mathrm{IN}^{\vec{\alpha}}(\vec{m} + \vec{u}) \right\}$.

Intuitively, $\mathrm{IN}^{\vec{\alpha}}(\vec{m})$ is the set of sides on which the tile at $\vec{m}$ initially binds in the assembly sequence $\vec{\alpha}$, and $\mathrm{OUT}^{\vec{\alpha}}(\vec{m})$ is the set of sides on which this tile propagates information to future tiles.

Note that $\mathrm{IN}^{\vec{\alpha}}(\vec{m}) = \emptyset$ for all $\vec{m} \in \alpha_0$.

Notation. If $\vec{\alpha} = (\alpha_i | 0 \le i < k)$ is a τ-T-assembly sequence, $\alpha = \mathrm{res}(\vec{\alpha})$, and $\vec{m} \in \mathrm{dom}\, \alpha - \mathrm{dom}\, \alpha_0$, then

$$\vec{\alpha} \setminus \vec{m} = \alpha \restriction \left(\mathrm{dom}\, \alpha - \{\vec{m}\} - \left(\vec{m} + \mathrm{OUT}^{\vec{\alpha}}(\vec{m}) \right) \right).$$

(Note that $\vec{\alpha} \setminus \vec{m}$ is a T-configuration that may or may not be a τ-T-assembly.

Definition. (Soloveichik and Winfree [16]). A τ-T-assembly sequence $\vec{\alpha} = (\alpha_i | 0 \leq i < k)$ with result α is *locally deterministic* if it has the following three properties.

1. For all $\vec{m} \in \text{dom } \alpha - \text{dom } \alpha_0$,
$$\sum_{\vec{u} \in \text{IN}^{\vec{\alpha}}(\vec{m})} \text{str}_{\alpha_{i_{\vec{\alpha}}(\vec{m})}}(\vec{m}, \vec{u}) = \tau.$$
2. For all $\vec{m} \in \text{dom } \alpha - \text{dom } \alpha_0$ and all $t \in T - \{\alpha(\vec{m})\}$, $\vec{m} \notin \partial_t^\tau(\vec{\alpha} \setminus \vec{m})$.
3. $\partial^\tau \alpha = \emptyset$.

That is, $\vec{\alpha}$ is locally deterministic if (1) each tile added in $\vec{\alpha}$ "just barely" binds to the assembly; (2) if a tile of type t_0 at a location $\vec{m}$ and its immediate "OUT-neighbors" are deleted from the *result* of $\vec{\alpha}$, then no tile of type $t \neq t_0$ can attach itself to the thus-obtained configuration at location $\vec{m}$; and (3) the result of $\vec{\alpha}$ is terminal.

Definition. A GTAS $\mathcal{T} = (T, \sigma, \tau)$ is *locally deterministic* if there exists a locally determinstic τ-T-assembly sequence $\vec{\alpha} = (\alpha_i | 0 \leq i < k)$ with $\alpha_0 = \sigma$.

Theorem 2.6. (Soloveichik and Winfree [16]) Every locally deterministic GTAS is directed.

2.4 Zeta-Dimension

The most commonly used dimension for discrete fractals is zeta-dimension, which we use in this paper. The discrete-continuous correspondence mentioned in the introduction preserves dimension somewhat generally. Thus, for example, the zeta-dimension of the discrete Sierpinski triangle is the same as the Hausdorff dimension of the continuous Sierpinski triangle.

Zeta-dimension has been re-discovered several times by researchers in various fields over the past few decades, but its origins actually lie in Euler's (real-valued predecessor of the Riemann) zeta-function [7] and Dirichlet series. For each set $A \subseteq \mathbb{Z}^2$, define the *A-zeta-function* $\zeta_A : [0, \infty) \to [0, \infty]$ by $\zeta_A(s) = \sum_{(0,0) \neq (m,n) \in A} (|m| + |n|)^{-s}$ for all $s \in [0, \infty)$. Then the *zeta-dimension* of A is

$$\text{Dim}_\zeta(A) = \inf\{s | \zeta_A(s) < \infty\}.$$

It is clear that $0 \leq Dim_\zeta(A) \leq 2$ for all $A \subseteq \mathbb{Z}^2$. It is also easy to see (and was proven by Cahen in 1894; see also [4, 10]) that zeta-dimension admits the "entropy characterization"

$$\text{Dim}_\zeta(A) = \limsup_{n \to \infty} \frac{\log |A_{\leq n}|}{\log n}, \tag{2.1}$$

where $A_{\leq n} = \{(m, n) \in A \mid |m| + |n| \leq n\}$. Various properties of zeta-dimension, along with extensive historical citations, appear in the recent paper [6], but our technical arguments here can be followed without reference to this material. We use the fact, verifiable by routine calculation, that (2.1) can be transformed by changes of variable up to exponential, e.g.,

$$\text{Dim}_\zeta(A) = \limsup_{n \to \infty} \frac{\log |A_{[0,2^n]}|}{n}$$

also holds.

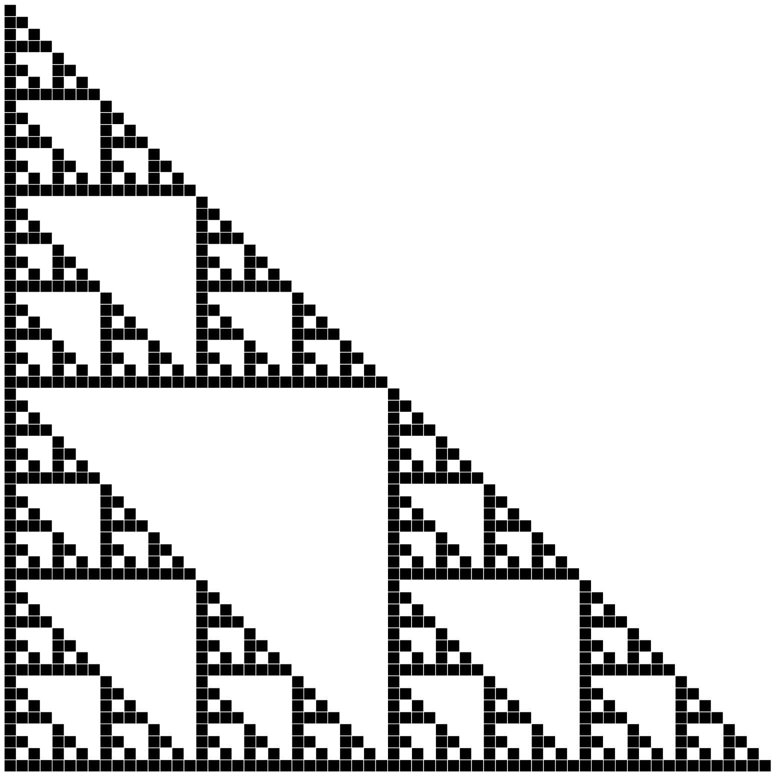

Figure 1: The standard discrete Sierpinski triangle $\mathbf{S}$.

2.5 The Standard Discrete Sierpinski Triangle S

We briefly review the standard discrete Sierpinski triangle and the calculation of its zeta-dimension.

Let $V = \{(1,0),(0,1)\}$. Define the sets $S_0, S_1, S_2, \cdots \subseteq \mathbb{Z}^2$ by the recursion

$$\begin{aligned} S_0 &= \{(0,0)\}, \\ S_{i+1} &= S_i \cup \left(S_i + 2^i V\right), \end{aligned} \tag{2.2}$$

where $A + cB = \{\vec{m} + c\vec{n} | \vec{m} \in A \text{ and } \vec{n} \in B\}$. Then the *standard discrete Sierpinski triangle* is the set

$$\mathbf{S} = \bigcup_{i=0}^{\infty} S_i,$$

which is illustrated in Figure 1. It is well known that $\mathbf{S}$ is the set of all $(k,l) \in \mathbb{N}^2$ such that the binomial coefficient $\binom{k+l}{k}$ is odd. For this reason, the set $\mathbf{S}$ is also called *Pascal's triangle modulo 2.* It is clear from the recursion (2.2) that $|S_i| = 3^i$ for all $i \in \mathbb{N}$. The zeta-dimension of $\mathbf{S}$ is thus

$$\begin{aligned}
\mathrm{Dim}_\zeta(\mathbf{S}) &= \limsup_{n\to\infty} \frac{\log\left|\mathbf{S}_{[0,2^n]}\right|}{n} \\
&= \limsup_{n\to\infty} \frac{\log|S_n|}{n} \\
&= \log 3 \\
&\approx 1.585.
\end{aligned}$$

3 Impossibility of Strict Self-Assembly of S

This section presents our first main theorem, which says that the standard discrete Sierpinski triangle $\mathbf{S}$ does not strictly self-assemble in the Tile Assembly Model. In order to prove this theorem, we first develop a lower bound on the number of tile types required for the self-assembly of a set X in terms of the depths of finite trees that occur in a certain way as subtrees of the full grid graph $G^{\#}_X$ of X.

Intuitively, given a set D of vertices of $G^{\#}_X$ (which is in practice the domain of the seed assembly), we now define a D-subtree of $G^{\#}_X$ to be any rooted tree in $G^{\#}_X$ that consists of *all* vertices of $G^{\#}_X$ that lie at or on the far side of the root from D. For simplicity, we state the definition in an arbitrary graph G.

Definition. Let $G = (V, E)$ be a graph, and let $D \subseteq V$.

1. For each $r \in V$, the *D-r-rooted subgraph* of G is the graph $G_{D,r} = (V_{D,r}, E_{D,r})$, where

   $$V_{D,r} = \{v \in V \mid \text{every path from } v \text{ to (any vertex in) } D \text{ in } G \text{ goes through } r\}$$

   and

   $$E_{D,r} = E \cap [V_{D,r}]^2 .$$

   (Note that $r \in V_{D,r}$ in any case.)

2. A *D-subtree* of G is a rooted tree B with root $r \in V$ such that $B = G_{D,r}$.

3. A *branch* of a D-subtree B of G is a simple path $\pi = (v_0, v_1, \ldots)$ in B that starts at the root of B and either ends at a leaf of B or is infinitely long.

We use the following quantity in our lower bound theorem.

Definition. Let $G = (V, E)$ be a graph, and let $D \subseteq V$. The *finite-tree depth* of G *relative to* D is

$$\text{ft-depth}_D(G) = \sup\{\text{depth(B)} | B \text{ is a finite } D\text{-subtree of } G\}.$$

We emphasize that the above supremum is only taken over *finite* D-subtrees. It is easy to construct an example in which G has a D-subtree of infinite depth, but $\text{ft-depth}_D(G) < \infty$.

To prove our lower bound result, we use the following theorem from [2].

Theorem 3.1. (Adleman, Cheng, Goel, Huang, Kempe, Moisset de Espanés, and Rothemund [2]) Let $X \subseteq \mathbb{Z}^2$ with $|X| < \infty$ be such that $G^{\#}_X$ is a tree rooted at the origin. If X strictly self-assembles in a GTAS $\mathcal{T} = (T, \sigma, 2)$ whose seed σ consists of a single tile at the origin, then $|T| \geq \text{depth}\left(G^{\#}_X\right)$.

Our lower bound result is the following.

Theorem 3.2. Let $X \subseteq \mathbb{Z}^2$. If X strictly self-assembles in a GTAS $\mathcal{T} = (T, \sigma, \tau)$, then $|T| \geq \text{ft-depth}_{\text{dom}\,\sigma}\left(G^{\#}_X\right)$.

Proof. Assume the hypothesis, and let B be a finite dom σ-subtree of $G^{\#}_{X}$. If suffices to prove that $|T| \geq \text{depth}(B)$.

Let $\alpha \in \mathcal{A}_{\Box}[\mathcal{T}]$, and let $\vec{r}$ be the root of B. Let σ' be the assembly with dom $\sigma' = \{\vec{r}\}$ and $\vec{u} \in U_2$. We define $\sigma'(\vec{r})$ as follows.

$$\sigma'(\vec{r})(\vec{u}) = \begin{cases} (\text{col}_{\alpha(\vec{r})}(\vec{u}), \text{str}_{\alpha(\vec{r})}(\vec{u})) & \text{if } \vec{r}+\vec{u} \in B \\ (\text{col}_{\alpha(\vec{r})}(\vec{u}), 0) & \text{otherwise.} \end{cases}$$

Then $\mathcal{T}' = (T, \sigma', \tau)$ is a GTAS in which B self-assembles. By Theorem 3.1, this implies that $|T| \geq \text{depth}(B)$. □

We next show that the standard discrete Sierpinski triangle $\mathbf{S}$ has infinite finite-tree depth.

Lemma 3.3. For every finite set $D \subseteq S$, $\text{ft-depth}_D\left(G^{\#}_{\mathbf{S}}\right) = \infty$.

Proof. Let $D \subseteq \mathbf{S}$ be finite, and let m be a positive integer. It suffices to show that $\text{ft-depth}_D\left(G^{\#}_{\mathbf{S}}\right) > m$. Choose $k \in \mathbb{N}$ large enough to satisfy the following two conditions.

(i) $2^k > \max\{a \in \mathbb{N} | (\exists b \in \mathbb{N})(a, b) \in D\}$.

(ii) $2^k > m$.

Let $\vec{r}_k = (2^{k+1}, 2^k)$, and let

$$B_k = \left\{(a,b) \in \mathbf{S} \,\middle|\, a \geq 2^{k+1}, b \geq 2^k \text{ and } a+b \leq 2^{k+2}-1\right\}.$$

It is routine to verify that $G^{\#}_{B_k}$ is a finite D-subtree of $G^{\#}_{\mathbf{S}}$ with root at $\vec{r}$ and depth 2^k. It follows that

$$\text{ft-depth}_D\left(G^{\#}_{\mathbf{S}}\right) \geq \text{depth}\left(G^{\#}_{B_k}\right) = 2^k > m.$$

□

We now have the machinery to prove our first main theorem.

Theorem 3.4. $\mathbf{S}$ does not strictly self-assemble in the Tile Assembly Model.

Proof. Let $\mathcal{T} = (T, \sigma, \tau)$ be a GTAS in which $\mathbf{S}$ strictly self-assembles. It suffices to show that $\mathcal{T}$ is not a TAS. If dom σ is infinite, this is clear, so assume that dom σ is finite. Then Theorem 3.2 and Lemma 3.3 tell us that $|T| = \infty$, whence $\mathcal{T}$ is not a TAS. □

Before moving on, we note that Theorem 3.4 implies the following lower bound on the number of tile types needed to strictly assemble any *finite* stage S_n of $\mathbf{S}$.

Corollary 3.5. If a stage S_n of $\mathbf{S}$ strictly self-assembles in a TAS $\mathcal{T} = (T, \sigma, \tau)$ in which σ consists of a single tile at the origin, then $|T| \geq 2^n$.

If we let $N = |S_n| = 3^n$, then the above lower bound exceeds $N^{0.63}$. As Rothemund [12] has noted, a structure of N tiles that requires $\sqrt{N}$ or more tile types for its self-assembly cannot be said to feasibly self-assemble.

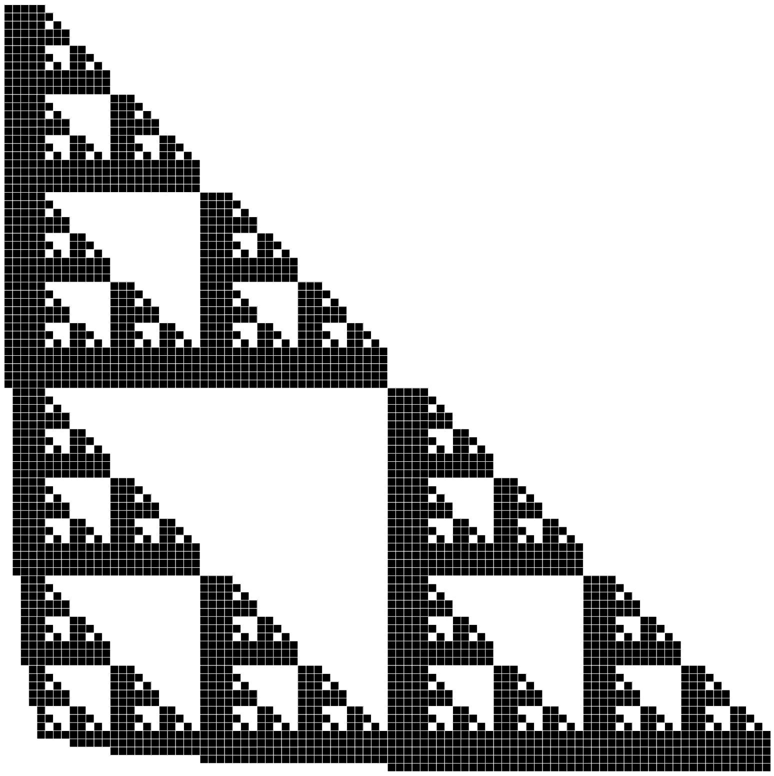

Figure 2: The fibered Sierpinski triangle **T**.

4 The Fibered Sierpinski Triangle T

We now define the fibered Sierpinski triangle and show that it has the same zeta-dimension as the standard discrete Sierpinski triangle.

As in Section 2, let $V = \{(1,0),(0,1)\}$. Our objective is to define sets $T_0, T_1, T_2, \cdots \subseteq \mathbb{Z}^2$, sets $F_0, F_1, F_2, \cdots \subseteq \mathbb{Z}^2$, and functions $l, f, t : \mathbb{N} \to \mathbb{N}$ with the following intuitive meanings.

1. T_i is the i^{th} stage of our construction of the fibered Sierpinski triangle.

2. F_i is the *fiber* associated with T_i, a thin strip of tiles along which data moves in the self-assembly process of Section 5. It is the smallest set whose union with T_i has a vertical left edge and a horizontal bottom edge, together with one additional layer added to these two now-straight edges.

3. $l(i)$ is the length of (number of tiles in) the left (or bottom) edge of $T_i \cup F_i$.

4. $f(i) = |F_i|$.

5. $t(i) = |T_i|$.

These five entities are defined recursively by the equations

$$
\begin{aligned}
&T_0 = S_2 \text{ (stage 2 in the construction of } S\text{)},\\
&F_0 = (\{-1\} \times \{-1,0,1,2,3\}) \cup (\{-1,0,1,2,3\} \times \{-1\}),\\
&l(0) = 5,\\
&f(0) = 9,\\
&t(0) = 9,\\
&T_{i+1} = T_i \cup ((T_i \cup F_i) + l(i)V), && (4.1)\\
&F_{i+1} = F_i \cup (\{-i-2\} \times \{-i-2, -i-1, \cdots, l(i+1)-i-3\})\\
&\cup (\{-i-2, -i-1, \cdots, l(i+1)-i-3\} \times \{-i-2\}), && (4.2)\\
&l(i+1) = 2l(i) + 1,\\
&f(i+1) = f(i) + 2l(i+1) - 1,\\
&t(i+1) = 3t(i) + 2f(i).
\end{aligned}
$$

Comparing the recursions (2.1) and (4.1) shows that the sets $T_0, T_1, T_2, \cdots$ are constructed exactly like the sets $S_0, S_1, S_2, \cdots$, except that the fibers F_i are inserted into the construction of the sets T_i. A routine induction verifies that this recursion achieves conditions 2, 3, 4, and 5 above. The *fibered Sierpinski triangle* is the set

$$\mathbf{T} = \bigcup_{i=0}^{\infty} T_i \tag{4.3}$$

which is illustrated in Figure 2. The resemblance between $\mathbf{S}$ and $\mathbf{T}$ is clear from the illustrations. We now verify that $\mathbf{S}$ and $\mathbf{T}$ have the same zeta-dimension.

Lemma 4.1. $\mathrm{Dim}_\zeta(\mathbf{T}) = \mathrm{Dim}_\zeta(\mathbf{S})$.

Proof. Solving the recurrences for l, f, and t, in that order, gives the formulas

$$
\begin{aligned}
l(i) &= 3 \cdot 2^{i+1} - 1,\\
f(i) &= 3\left(2^{i+3} - i - 5\right),\\
t(i) &= \frac{3}{2}\left(3^{i+3} - 2^{i+5} + 2i + 11\right),
\end{aligned}
$$

which can be routinely verified by induction. It follows readily that

$$\mathrm{Dim}_\zeta(\mathbf{T}) = \limsup_{n\to\infty} \frac{\log t(n)}{\log l(n)} = \log 3 = \mathrm{Dim}_\zeta(\mathbf{S}).$$

□

We note that the thickness $i+1$ of a fiber F_i is $O(\log l(i))$, i.e., logarithmic in the side length of T_i. Hence the difference between S_i and T_i is asymptotically negligible as $i \to \infty$. Nevertheless, we show in the next section that $\mathbf{T}$, unlike $\mathbf{S}$, strictly self-assembles in the Tile Assembly Model.

5 Strict Self-Assembly of T

This section is devoted to proving our second main theorem, which is the fact that the fibered Sierpinski triangle $\mathbf{T}$ strictly self-assembles in the Tile Assembly Model. Our proof is constructive, i.e., we exhibit a specific tile assembly system in which $\mathbf{T}$ strictly self-assembles.

Our strict self-assembly of $\mathbf{T}$ is not based directly upon the recursive definition (4.1). A casual inspection of Figure 2 suggests that $\mathbf{T}$ can also be regarded as a structure consisting of many horizontal and vertical bars, with each large bar having many smaller bars perpendicular to it. In subsection 5.1 we give a precise statement and proof of this "bar characterization" of $\mathbf{T}$, which is the basis of our strict self-assembly. In subsections 5.2 and 5.3 we present the main functional subsystems of our construction. This gives us a tile assembly system $\mathcal{T}_{\mathbf{T}} = (T_{\mathbf{T}}, \sigma_{\mathbf{T}}, \tau)$, where

(i) the tile set $T_{\mathbf{T}}$ consists of 51 tile types;

(ii) the seed assembly $\sigma_{\mathbf{T}}$ consists of a single 'S' tile at the origin; and

(iii) the temperature τ is 2.

Subsection 5.4 proves that the fibered Sierpinski triangle $\mathbf{T}$ strictly self-assembles in $\mathcal{T}_{\mathbf{T}}$.

Throughout this section, the temperature τ is 2. Tiles are depicted as squares whose various sides are dotted lines, solid lines, or doubled lines, indicating whether the glue strengths on these sides are 0, 1, or 2, respectively. Thus, for example, a tile of the type shown in Figure 3 has glue of strength 0 on the left and

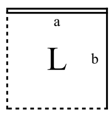

Figure 3: An example tile type.

bottom, glue of color 'a' and strength 2 on the top, and glue of color 'b' and strength 1 on the right. This tile also has a label 'L', which plays no formal role but may aid our understanding and discussion of the construction.

5.1 Bar Characterization of T

We now formulate the characterization of $\mathbf{T}$ that guides its strict self-assembly. At the outset, in the notation of section 4, we focus on the manner in which the sets $T_i \cup F_i$ can be constructed from horizontal and vertical bars. Recall that

$$l(i) = 3 \cdot 2^{i+1} - 1$$

is the length of (number of tiles in) the left or bottom edge of $T_i \cup F_i$.

Definition. Let $-1 \leq i \in \mathbb{Z}$.

1. The S_i-*square* is the set
$$S_i = \{-i-1, \ldots, 0\} \times \{-i-1, \ldots, 0\}.$$

2. The X_i-*bar* is the set
$$X_i = \{1, \ldots, l(i) - i - 2\} \times \{-i-1, \ldots, 0\}.$$

3. The Y_i-*bar* is the set
$$Y_i = \{-i-1, \ldots, 0\} \times \{1, \ldots, l(i) - i - 2\}.$$

It is clear that the set

$$S_i \cup X_i \cup Y_i$$

is the "outer framework" of $T_i \cup F_i$. Our attention thus turns to the manner in which smaller and smaller bars are recursively attached to this framework.

We use the *ruler function*

$$\rho : \mathbb{Z}^+ \to \mathbb{N}$$

defined by the recurrence

$$\begin{aligned} \rho(2k+1) &= 0, \\ \rho(2k) &= \rho(k) + 1 \end{aligned}$$

for all $k \in \mathbb{N}$. It is easy to see that $\rho(n)$ is the (exponent of the) largest power of 2 that divides n. Equivalently, $\rho(n)$ is the number of 0's lying to the right of the rightmost 1 in the binary expansion of n [9]. An easy induction can be used to establish the following observation.

Observation 5.1. For all $n \in \mathbb{N}$,

$$\sum_{j=1}^{2^{n+1}-1} \rho(j) = 2^{n+1} - n - 2.$$

Using the ruler function, we define the function

$$\theta : \mathbb{Z}^+ \to \mathbb{Z}^+$$

by the recurrence

$$\begin{aligned} \theta(1) &= 2, \\ \theta(j+1) &= \theta(j) + \rho(j+1) + 2 \end{aligned}$$

for all $j \in \mathbb{Z}^+$.

We now use the function θ to define the points at which smaller bars are attached to the X_i- and Y_i-bars.

Definition.

1. The j^{th} *θ-point* of X_i is the point
$$\theta_j(X_i) = (\theta(j), 1),$$
lying just above the X_i-bar.
2. The j^{th} *θ-point* of Y_i is the point
$$\theta_j(Y_i) = (1, \theta(j)),$$
lying just to the right of the Y_i-bar.

The following recursion attaches smaller bars to larger bars in a recursive fashion.

Definition. The *θ-closures* of the bars X_i and Y_i are the sets $\theta(X_i)$ and $\theta(Y_i)$ defined by the mutual recursion

$$\begin{aligned} \theta\left(X_{-1}\right) &= X_{-1} \\ \theta\left(Y_{-1}\right) &= Y_{-1} \\ \theta(X_i) &= X_i \cup \bigcup_{j=1}^{2^{i+1}-1} \left(\theta_j(X_i) + \theta\left(Y_{\rho(j)-1}\right)\right), \\ \theta(Y_i) &= Y_i \cup \bigcup_{j=1}^{2^{i+1}-1} \left(\theta_j(Y_i) + \theta\left(X_{\rho(j)-1}\right)\right), \end{aligned}$$

for all $i \in \mathbb{N}$.

This definition, along with the symmetry of ρ, admit the following characterizations of $\theta(X_i)$ and $\theta(Y_i)$.

Observation 5.2. Let $0 \le i \in \mathbb{N}$.

1. $\theta(X_{i+1}) = \theta(X_i) \cup ((S_i \cup \theta(X_i) \cup \theta(Y_i)) + (l(i), 0)) \cup (\{1, \ldots l(i+1) - i - 3\} \times \{-i-2\})$.
2. $\theta(Y_{i+1}) = \theta(Y_i) \cup ((S_i \cup \theta(X_i) \cup \theta(Y_i)) + (0, l(i))) \cup (\{-i-2\} \times \{1, \ldots l(i+1) - i - 3\})$.

We have the following characterization of the sets $T_i \cup F_i$.

Lemma 5.3. For all $i \in \mathbb{N}$,

$$T_i \cup F_i = S_i \cup \theta(X_i) \cup \theta(Y_i).$$

Proof. We proceed by induction on i, and note that the case when $i = 0$ is trivial. Assume that, for all $i \in \mathbb{N}$, the lemma holds. Then we have

$$\begin{aligned}
T_{i+1} \cup F_{i+1} &\overset{(4.1)}{=} T_i \cup ((T_i \cup F_i) + l(i)V) \cup F_{i+1} \\
&\overset{(4.2)}{=} (T_i \cup F_i) \cup ((T_i \cup F_i) + l(i)V) \cup \\
&\quad (\{-i-2\} \times \{-i-2, \ldots, l(i+1) - i - 3\}) \cup \\
&\quad\quad (\{-i-2, \ldots, l(i+1) - i - 3\} \times \{-i-2\}) \\
&\overset{\text{Ind Hyp}}{=} (S_i \cup \theta(X_i) \cup \theta(Y_i)) \cup ((S_i \cup \theta(X_i) \cup \theta(Y_i)) + l(i)V) \cup \\
&\quad (\{-i-2\} \times \{-i-2, \ldots, l(i+1) - i - 3\}) \cup \\
&\quad\quad (\{-i-2, \ldots, l(i+1) - i - 3\} \times \{-i-2\}) \\
&= (S_i \cup (\{-i-2, \ldots, 0\} \times \{-i-2\}) \cup (\{-i-2\} \times \{-i-2, \ldots, 0\})) \cup \\
&\quad (\theta(X_i) \cup ((S_i \cup \theta(X_i) \cup \theta(Y_i)) + (l(i), 0))) \cup \\
&\quad\quad (\{1, \ldots, l(i+1) - i - 3\} \times \{-i-2\}) \cup \\
&\quad\quad\quad (\theta(Y_i) \cup ((S_i \cup \theta(X_i) \cup \theta(Y_i)) + (0, l(i)))) \cup \\
&\quad\quad\quad\quad (\{-i-2\} \times \{1, \ldots, l(i+1) - i - 3\}) \\
&\overset{\text{Obs 5.1}}{=} S_{i+1} \cup \theta(X_{i+1}) \cup \theta(Y_{i+1}).
\end{aligned}$$

□

We now shift our attention to the global structure of the set $\mathbf{T}$.

Definition.

1. The *x-axis* of $\mathbf{T}$ is the set
$$X = \{(m, n) \in \mathbf{T} \mid m > 0, \text{ and } n \le 0\}.$$
2. The *y-axis* of $\mathbf{T}$ is the set
$$Y = \{(m, n) \in \mathbf{T} \mid m \le 0, \text{ and } n > 0\}.$$

Intuitively, the x-axis of $\mathbf{T}$ is the part of $\mathbf{T}$ that is a "gradually thickening bar" lying on and below the (actual) x-axis in $\mathbb{Z}^2$. (see Figure 2.) For technical convenience, we have omitted the origin from this set. Similar remarks apply to the y-axis of $\mathbf{T}$.

Define the sets

$$\begin{aligned}
\widetilde{X}_{-1} &= \{(1, 0), (2, 0), (3, 0)\}, \\
\widetilde{Y}_{-1} &= \{(0, 1), (0, 2), (0, 3)\}.
\end{aligned}$$

For each $i \in \mathbb{N}$, define the translations

$$\begin{aligned} S_i^{\rightarrow} &= (l(i), 0) + S_i, \\ S_i^{\uparrow} &= (0, l(i)) + S_i, \\ X_i^{\rightarrow} &= (l(i), 0) + X_i, \\ Y_i^{\uparrow} &= (0, l(i)) + Y_i \end{aligned}$$

of S_i, X_i, and Y_i. It is clear by inspection that X is the disjoint union of the sets

$$\widetilde{X}_{-1}, S_0^{\rightarrow}, X_0^{\rightarrow}, S_1^{\rightarrow}, X_1^{\rightarrow}, S_2^{\rightarrow}, X_2^{\rightarrow}, \ldots,$$

which are written in their left-to-right order of position in X. More succinctly, we have the following.

Observation 5.4.

1. $X = \widetilde{X}_{-1} \cup \bigcup_{i=0}^{\infty} (S_i^{\rightarrow} \cup X_i^{\rightarrow})$.
2. $Y = \widetilde{Y}_{-1} \cup \bigcup_{i=0}^{\infty} \left(S_i^{\uparrow} \cup Y_i^{\uparrow}\right)$.

Moreover, both of these are disjoint unions.

In light of Observation 5.4, it is convenient to define, for each $-1 \le n \in \mathbb{Z}$, the initial segment

$$\widetilde{X}_n = \widetilde{X}_{-1} \cup \bigcup_{i=0}^{n} (S_i^{\rightarrow} \cup X_i^{\rightarrow})$$

of X and the initial segment

$$\widetilde{Y}_n = \widetilde{Y}_{-1} \cup \bigcup_{i=0}^{n} \left(S_i^{\uparrow} \cup Y_i^{\uparrow}\right)$$

of Y. (Note that this is consistent with earlier usage when $n = -1$.)

The following definition specifies the manner in which bars are recursively attached to the x- and y-axes of $\mathbf{T}$.

Definition. Let $j \in \mathbb{Z}^+$.

1. The j^{th} *θ-point* of X is the point
$$\theta_j(X) = (\theta(j), 1)$$
lying just above X.
2. The j^{th} *θ-point* of Y is the point
$$\theta_j(Y) = (1, \theta(j))$$
lying just to the right of Y.

Definition. For all $-1 \le n \in \mathbb{Z}$, the *θ-closures* of the initial segment of the axes $\widetilde{X}_n$ and $\widetilde{Y}_n$ are the sets

$$\theta\left(\widetilde{X}_n\right) = \widetilde{X}_n \cup \bigcup_{j=1}^{2^{n+2}-1} \left(\theta_j(X) + \theta\left(Y_{\rho(j)-1}\right)\right)$$

and

$$\theta\left(\widetilde{Y}_n\right) = \widetilde{Y}_n \cup \bigcup_{j=1}^{2^{n+2}-1} \left(\theta_j(Y) + \theta\left(X_{\rho(j)-1}\right)\right),$$

respectively.

The following observation is an immediate consequence of the previous definition.

Observation 5.5. Let $0 \leq n \in \mathbb{N}$.

1. $\theta\left(\widetilde{X}_n\right) = \theta\left(\widetilde{X}_{n-1}\right) \cup \left(\left(S_n \cup \theta\left(X_n\right) \cup \theta\left(Y_n\right)\right) + (l(n), 0)\right)$.
2. $\theta\left(\widetilde{Y}_n\right) = \theta\left(\widetilde{Y}_{n-1}\right) \cup \left(\left(S_n \cup \theta\left(X_n\right) \cup \theta\left(Y_n\right)\right) + (0, l(n))\right)$.

We have the following characterization of T_n.

Lemma 5.6. For all $-1 \leq n \in \mathbb{Z}$,

$$T_{n+1} = \{(0,0)\} \cup \theta\left(\widetilde{X}_n\right) \cup \theta\left(\widetilde{Y}_n\right).$$

Proof. We proceed by induction on n. When $n = -1$, it is easy to see that

$$\{(0,0)\} \cup \theta\left(\widetilde{X}_{-1}\right) \cup \theta\left(\widetilde{Y}_{-1}\right) = T_0.$$

Now assume that, for all $-1 \leq n \in \mathbb{N}$, the lemma holds. Then we have

$$\begin{aligned}
T_{n+2} &\overset{(4.1)}{=} T_{n+1} \cup \left(\left(T_{n+1} \cup F_{n+1}\right) + l(n+1)V\right) \\
&\overset{\text{Ind Hyp}}{=} \left(\{(0,0)\} \cup \theta\left(\widetilde{X}_n\right) \cup \theta\left(\widetilde{Y}_n\right)\right) \cup \\
&\qquad \left(\left(T_{n+1} \cup F_{n+1}\right) + l(n+1)V\right) \\
&\overset{\text{Lemma 5.2}}{=} \left(\{(0,0)\} \cup \theta\left(\widetilde{X}_n\right) \cup \theta\left(\widetilde{Y}_n\right)\right) \cup \\
&\qquad \left(\left(S_{n+1} \cup \theta\left(X_{n+1}\right) \cup \theta\left(Y_{n+1}\right)\right) + l(n+1)V\right) \\
&= \{(0,0)\} \cup \left(\theta\left(\widetilde{X}_n\right) \cup \left(\left(S_{n+1} \cup \theta\left(X_{n+1}\right) \cup \theta\left(Y_{n+1}\right)\right) + (l(n+1), 0)\right)\right) \\
&\qquad \left(\theta\left(\widetilde{Y}_n\right) \cup \left(\left(S_{n+1} \cup \theta\left(X_{n+1}\right) \cup \theta\left(Y_{n+1}\right)\right) + (0, l(n+1))\right)\right) \\
&\overset{\text{Obs 5.4}}{=} \{(0,0)\} \cup \theta\left(\widetilde{X}_{n+1}\right) \cup \theta\left(\widetilde{Y}_{n+1}\right).
\end{aligned}$$

□

Definition. The θ-*closures* of the axes X and Y are the sets

$$\theta(X) = X \cup \bigcup_{j=1}^{\infty} \left(\theta_j(X) + \theta\left(Y_{\rho(j)-1}\right)\right)$$

and

$$\theta(Y) = Y \cup \bigcup_{j=1}^{\infty} \left(\theta_j(Y) + \theta\left(X_{\rho(j)-1}\right)\right),$$

respectively.

Figure 3 shows the structure of the Y-axis.

Lemma 5.7. Let $1 \leq i \in \mathbb{Z}$.

1. $\theta(X) = \bigcup_{i=-1}^{\infty} \theta\left(\widetilde{X}_i\right)$.