

Limitations of Self-Assembly at Temperature 1*

David Doty[†], Matthew J. Patitz, and Scott M. Summers[‡]

Department of Computer Science
Iowa State University
Ames, IA 50011, USA
{ddoty,mpatitz,summers}@cs.iastate.edu

Abstract

We prove that if a set $X \subseteq \mathbb{Z}^2$ weakly self-assembles at temperature 1 in a deterministic (Winfree) tile assembly system satisfying a natural condition known as *pumpability*, then X is a finite union of semi-doubly periodic sets. This shows that only the most simple of infinite shapes and patterns can be constructed using pumpable temperature 1 tile assembly systems, and gives evidence for the thesis that temperature 2 or higher is required to carry out general-purpose computation in a tile assembly system. Finally, we show that general-purpose computation *is* possible at temperature 1 if negative glue strengths are allowed in the tile assembly model.

1 Introduction

Self-assembly is a bottom-up process by which a small number of fundamental components automatically coalesce to form a target structure. In 1998, Winfree [10] introduced the (abstract) Tile Assembly Model (TAM) – an “effectivization” of Wang tiling [8, 9] – as an over-simplified mathematical model of the DNA self-assembly pioneered by Seeman [7]. In the TAM, the fundamental components are un-rotatable, but translatable square “tile types” whose sides are labeled with glue “colors” and “strengths.” Two tiles that are placed next to each other *interact* if the glue colors on their abutting sides match, and they *bind* if the strength on their abutting sides matches with total strength at least a certain ambient “temperature,” usually taken to be 1 or 2.

Despite its deliberate over-simplification, the TAM is a computationally and geometrically expressive model at temperature 2. The reason is that, at temperature 2, certain tiles are not permitted to bond until *two* tiles are already present to match the glues on the bonding sides, which enables cooperation between different tile types to control the placement of new tiles. Winfree [10] proved that at temperature 2 the TAM is computationally universal and thus can be directed algorithmically.

In contrast, we aim to prove that the lack of cooperation at temperature 1 inhibits the sort of complex behavior possible at temperature 2. Our main theorem concerns the *weak self-assembly*

*This research was supported in part by National Science Foundation grants 0652569 and 0728806.

[†]This author’s research was partially supported by NSF grant CCF:0430807.

[‡]This author’s research was supported in part by NSF-IGERT Training Project in Computational Molecular Biology Grant number DGE-0504304.

of subsets of \mathbb{Z}^2 using temperature 1 tile assembly systems. Informally, a set $X \subseteq \mathbb{Z}^2$ weakly self-assembles in a tile assembly system \mathcal{T} if some of the tile types of \mathcal{T} are painted black, and \mathcal{T} is guaranteed to self-assemble into an assembly α of tiles such that X is precisely the set of integer lattice points at which α contains black tile types. As an example, Winfree [10] exhibited a temperature 2 tile assembly system, directed by a clever XOR-like algorithm, that weakly self-assembles a well-known set, the discrete Sierpinski triangle, onto the first quadrant. Note that the underlying *shape* (set of all points containing a tile, whether black or not) of Winfree’s construction is an infinite canvas that covers the entire first quadrant, onto which a more sophisticated set, the discrete Sierpinski triangle, is painted.

We show that under a plausible assumption, temperature 1 tile systems weakly self-assemble only a limited class of sets. To prove our main result, we require the hypothesis that the tile system is *pumpable*. Informally, this means that every sufficiently long path of tiles in an assembly of this system contains a segment in which the same tile type repeats (a condition clearly implied by the pigeonhole principle), and that furthermore, the subpath between these two occurrences can be repeated indefinitely (“pumped”) along the same direction as the first occurrence of the segment, without “colliding” with a previous portion of the path. We give a counterexample in Section 3 (Figure 1) of a path in which the same tile type appears twice, yet the segment between the appearances cannot be pumped without eventually resulting in a collision that prevents additional pumping. The hypothesis of pumpability states (roughly) that in every sufficiently long path, despite the presence of some repeating tiles that cannot be pumped, *there exists* a segment in which the same tile type repeats that *can* be pumped. In the above-mentioned counterexample, the paths constructed to create a blocked segment always contain some previous segment that *is* pumpable. We conjecture that this phenomenon, pumpability, occurs in every temperature 1 tile assembly system that produces a unique infinite structure. We discuss this conjecture in greater detail in Section 5.

A *semi-doubly periodic* set $X \subseteq \mathbb{Z}^2$ is a set of integer lattice points with the property that there are three vectors \vec{b} (the “base point” of the set), \vec{u} , and \vec{v} (the two periods of the set), such that $X = \left\{ \vec{b} + n \cdot \vec{u} + m \cdot \vec{v} \mid n, m \in \mathbb{N} \right\}$. That is, a semi-doubly periodic set is a set that repeats infinitely along two vectors (linearly independent vectors in the non-degenerate case), starting at some base point \vec{b} . We show that any directed, pumpable, temperature 1 tile assembly system weakly self-assembles a set $X \subseteq \mathbb{Z}^2$ that is a finite union of semi-doubly periodic sets.

It is our contention that weak self-assembly captures the intuitive notion of what it means to “compute” with a tile assembly system. For example, the use of tile assembly systems to build shapes is captured by requiring all tile types to be black, in order to ask what set of integer lattice points contain any tile at all (so-called *strict self-assembly*). However, weak self-assembly is a more general notion. For example, Winfree’s above mentioned result shows that the discrete Sierpinski triangle weakly self-assembles at temperature 2 [6], yet this shape does not strictly self-assemble at *any* temperature [2]. Hence weak self-assembly allows for a more relaxed notion of set building, in which intermediate space can be used for computation, without requiring that the space filled to carry out the computation also represent the final result of the computation.

As another example, there is a standard construction [10] by which a single-tape Turing machine may be simulated by a temperature 2 tile assembly system. Regardless of the semantics of the Turing machine (whether it decides a language, enumerates a language, computes a function, etc.), it is routine to represent the result of the computation by the weak self-assembly of some set. For example, Patitz and Summers [3] showed that for any decidable language $A \subseteq \mathbb{N}$, A ’s projection

along the X -axis (the set $\{ (x, 0) \in \mathbb{N}^2 \mid x \in A \}$) weakly self-assembles in a temperature 2 tile assembly system. As another example, if a Turing machine computes a function $f : \mathbb{N} \rightarrow \mathbb{N}$, it is routine to design a tile assembly system based on Winfree’s construction such that, if the seed assembly is used to encode the binary representation of a number $n \in \mathbb{N}$, then the tile assembly system weakly self-assembles the set

$$\left\{ (k, 0) \in \mathbb{N}^2 \mid \begin{array}{l} \text{the } k^{\text{th}} \text{ least significant bit of the} \\ \text{binary representation of } f(n) \text{ is 1} \end{array} \right\}.$$

Our result is motivated by the thesis that if a tile assembly system can reasonably be said to “compute”, then the result of this computation can be represented in a straightforward manner as a set $X \subseteq \mathbb{Z}^2$ that weak self-assembles in the tile assembly system, or a closely related tile assembly system. Our examples above provide evidence for this thesis, although it is as informal and unprovable as the Church-Turing thesis. On the basis of this claim, and the triviality of semi-doubly periodic sets (shown more formally in Observation 4.2), we conclude that our main result implies that directed, pumpable, temperature 1 tile assembly systems are incapable of general-purpose deterministic computation, without further relaxing the model.

This paper is organized as follows. Section 2 provides background definitions and notation for the abstract TAM. Section 3 introduces new definitions and concepts that are required in proving our main theorem. Section 4 states our main theorem, and contains an observation justifying the suggestion that semi-doubly periodic sets are computationally very simple, based on their relationship to regular languages. Section 5 concludes the paper and discusses open questions. Section 5 also observes that a relaxation of the tile assembly model, allowing negative glue strengths and allowing glues with different colors to interact, *is* capable of general-purpose computation. Due to space constraints, we omit full proofs from this version of the paper. A full version of the paper with proofs can be found at <http://www.cs.iastate.edu/~lnsa>.

2 The Abstract Tile Assembly Model

We work in the 2-dimensional discrete space \mathbb{Z}^2 . Define the set $U_2 = \{(0, 1), (1, 0), (0, -1), (-1, 0)\}$ to be the set of all *unit vectors*, i.e., vectors of length 1 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*.

We now give a brief and intuitive sketch of the Tile Assembly Model that is adequate for reading this paper. More formal details and discussion may be found in [2, 4, 5, 10]. Our notation is that of [2], which contains a self-contained introduction to the Tile Assembly Model for the reader unfamiliar with the model.

Intuitively, a tile type t is a unit square that can be translated, but not rotated, having a well-defined “side \vec{u} ” for each $\vec{u} \in U_2$. Each side \vec{u} of t has a “glue” of “color” $\text{col}_t(\vec{u})$ – a string over some fixed alphabet Σ – and “strength” $\text{str}_t(\vec{u})$ – a nonnegative integer – specified by its type t . Two tiles t and t' that are placed at the points \vec{a} and $\vec{a} + \vec{u}$ respectively, *bind* with *strength* $\text{str}_t(\vec{u})$ if and only if $(\text{col}_t(\vec{u}), \text{str}_t(\vec{u})) = (\text{col}_{t'}(-\vec{u}), \text{str}_{t'}(-\vec{u}))$.

Given a set T of tile types, an *assembly* is a partial function $\alpha : \mathbb{Z}^2 \dashrightarrow T$, with points $\vec{x} \in \mathbb{Z}^2$ at which $\alpha(\vec{x})$ is undefined interpreted to be empty space, so that $\text{dom } \alpha$ is the set of points with tiles. α is *finite* if $|\text{dom } \alpha|$ is finite. For assemblies α and α' , we say that α is a *subconfiguration* of α' , and write $\alpha \sqsubseteq \alpha'$, if $\text{dom } \alpha \subseteq \text{dom } \alpha'$ and $\alpha(\vec{x}) = \alpha'(\vec{x})$ for all $x \in \text{dom } \alpha$.

Let α be an assembly and $B \subseteq \mathbb{Z}^2$. α *restricted to* B , written as $\alpha \upharpoonright B$, is the unique assembly satisfying $(\alpha \upharpoonright B) \sqsubseteq \alpha$, and $\text{dom}(\alpha \upharpoonright B) = B$. If π is a sequence over \mathbb{Z}^2 (such as a path), then we write $\alpha \upharpoonright \pi$ to mean α restricted to the set of points in π . If $A \subseteq \text{dom} \alpha$, we write $\alpha \setminus A = \alpha \upharpoonright (\text{dom} \alpha - A)$. If $\vec{0} \neq \vec{v} \in \mathbb{Z}^2$, then the *translation of α by \vec{v}* is defined as the assembly $(\alpha + \vec{v})$ satisfying, for all $\vec{a} \in \mathbb{Z}^2$,

$$(\alpha + \vec{v})(\vec{a}) = \begin{cases} \alpha(\vec{a}) & \text{if } \vec{a} - \vec{v} \in \text{dom } \alpha \\ \uparrow & \text{otherwise.} \end{cases}$$

A *grid graph* is a graph $G = (V, E)$ in which $V \subseteq \mathbb{Z}^2$ and every edge $\{\vec{a}, \vec{b}\} \in E$ has the property that $\vec{a} - \vec{b} \in U_2$. The *binding graph* of an assembly α is the grid graph $G_\alpha = (V, E)$, where $V = \text{dom } \alpha$, and $\{\vec{m}, \vec{n}\} \in E$ if and only if (1) $\vec{m} - \vec{n} \in U_2$, (2) $\text{col}_{\alpha(\vec{m})}(\vec{n} - \vec{m}) = \text{col}_{\alpha(\vec{n})}(\vec{m} - \vec{n})$, and (3) $\text{str}_{\alpha(\vec{m})}(\vec{n} - \vec{m}) > 0$. An assembly is τ -*stable*, where $\tau \in \mathbb{N}$, if it cannot be broken up into smaller assemblies without breaking bonds of total strength at least τ (i.e., if every cut of G_α cuts edges, the sum of whose strengths is at least τ).

Self-assembly begins with a *seed assembly* σ (typically assumed to be finite and τ -stable) and proceeds asynchronously and nondeterministically, with tiles adsorbing one at a time to the existing assembly in any manner that preserves stability at all times.

A *tile assembly system (TAS)* is an ordered triple $\mathcal{T} = (T, \sigma, \tau)$, where T is a finite set of tile types, σ is a seed assembly with finite domain, and τ is the temperature. In subsequent sections of this paper, we assume that $\tau = 1$ unless explicitly stated otherwise. An *assembly sequence* in a TAS $\mathcal{T} = (T, \sigma, 1)$ is a (possibly infinite) sequence $\vec{\alpha} = (\alpha_i \mid 0 \leq i < k)$ of assemblies in which $\alpha_0 = \sigma$ and each α_{i+1} is obtained from α_i by the “ τ -stable” addition of a single tile. The *result* of an assembly sequence $\vec{\alpha}$ is the unique assembly $\text{res}(\vec{\alpha})$ satisfying $\text{dom } \text{res}(\vec{\alpha}) = \bigcup_{0 \leq i < k} \text{dom } \alpha_i$ and, for each $0 \leq i < k$, $\alpha_i \sqsubseteq \text{res}(\vec{\alpha})$.

Note that, if $\mathcal{T} = (T, \sigma, 1)$, then there is a finite path from the seed to a point $\vec{x} \in \text{dom } \alpha$, denoted as $\pi_{\vec{0}, \vec{x}}$ if and only if there is an assembly sequence $\vec{\alpha}$ satisfying $\text{res}(\vec{\alpha}) = \alpha \upharpoonright \pi_{\vec{0}, \vec{x}}$.

We write $\mathcal{A}[\mathcal{T}]$ for the *set of all producible assemblies of \mathcal{T}* . An assembly α is *terminal*, and we write $\alpha \in \mathcal{A}_\square[\mathcal{T}]$, if no tile can be stably added to it. We write $\mathcal{A}_\square[\mathcal{T}]$ for the *set of all terminal assemblies of \mathcal{T}* . A TAS \mathcal{T} is *directed*, or *produces a unique assembly*, if it has exactly one terminal assembly i.e., $|\mathcal{A}_\square[\mathcal{T}]| = 1$. The reader is cautioned that the term “directed” has also been used for a different, more specialized notion in self-assembly [1]. We interpret “directed” to mean “deterministic”, though there are multiple senses in which a TAS may be deterministic or nondeterministic.

A set $X \subseteq \mathbb{Z}^2$ *weakly self-assembles* if there exists a TAS $\mathcal{T} = (T, \sigma, 1)$ and a set $B \subseteq T$ (B constitutes the “black” tiles) such that $\alpha^{-1}(B) = X$ holds for every assembly $\alpha \in \mathcal{A}_\square[\mathcal{T}]$. A set X *strictly self-assembles* if there is a TAS \mathcal{T} for which every assembly $\alpha \in \mathcal{A}_\square[\mathcal{T}]$ satisfies $\text{dom } \alpha = X$. Note that if X strictly self-assembles, then X weakly self-assembles. (Let all tiles be black.)

3 Pumpability, Finite Closures, and Combs

Throughout this section, let $\mathcal{T} = (T, \sigma, 1)$ be a directed TAS, and α be the unique assembly satisfying $\alpha \in \mathcal{A}_\square[\mathcal{T}]$. Further, we assume without loss of generality that σ consists of a single “seed” tile type placed at the origin.

Given, $\vec{0} \neq \vec{v} \in \mathbb{Z}^2$, a \vec{v} -*semi-periodic path in α originating at $\vec{a}_0 \in \text{dom } \alpha$* is an infinite, simple path $\pi = (\vec{a}_0, \vec{a}_1, \dots)$ in the binding graph G_α such that there is a constant $k \in \mathbb{N}$ such that, for all

$j \in \mathbb{N}$, $\pi[j+k] = \pi[j] + \vec{v}$, and $\alpha(\pi[j+k]) = \alpha(\pi[j])$. Intuitively, \vec{v} is the “geometric” period of the path – the straight-line vector between two repeating tile types – while k is the “linear” period – the number of tiles that must be traversed along the path before the tile types repeat, which is at least $\|\vec{v}\|_1$, but possibly larger if the segment from $\pi[j]$ to $\pi[j] + \vec{v}$ is “winding”.

An *eventually \vec{v} -semi-periodic path in α originating at $\vec{a}_0 \in \text{dom } \alpha$* is an infinite, simple path $\pi = (\vec{a}_0, \vec{a}_1, \dots)$ in the binding graph G_α for which there exists $s \in \mathbb{N}$ such that the path $\pi' = (\pi[s], \pi[s+1], \dots)$ is a \vec{v} -semi-periodic path in α originating at $\pi[s]$. Let the *initial segment length* be the smallest index s^* such that $\pi' = (\pi[s^* - k], \pi[s^* - k + 1], \dots)$ is a \vec{v} -semi-periodic path originating at the point $\pi[s^* - k]$. The *initial segment of π* is the path $\pi[0 \dots s^* - 1]$ (for technical reasons, we enforce the initial segment of π to contain the simple path $\pi[0 \dots s]$ along with one period of π'). The *tail of π* is $\pi[s^* \dots]$. Note that the tail of an eventually \vec{v} -semi-periodic path is simply a \vec{v} -semi-periodic path originating at $\pi[s^*]$. An *eventually \vec{v} -periodic path in α* is an eventually \vec{v} -periodic path in α originating at \vec{a}_0 for some $\vec{a}_0 \in \text{dom } \alpha$. We say that π is a *\vec{v} -periodic path in α* if $\pi = (\dots, \vec{a}_{-1}, \vec{a}_0, \vec{a}_1, \dots)$ is a two-way infinite simple path such that, for all $j \in \mathbb{Z}$, $\alpha(\pi[j] + \vec{v}) = \alpha(\pi[j])$.

Let $\vec{w}, \vec{x} \in \text{dom } \alpha$, π be a simple path from \vec{w} to \vec{x} in the binding graph G_α , and $i, k \in \mathbb{N}$ with $0 \leq i < k \leq |\pi|$. We say that π has a *pumpable segment $\pi[i \dots k]$* (with respect to $\vec{v} = \pi[k] - \pi[i]$) if there exists a \vec{v} -semi-periodic path π' in α originating at $\pi[i]$ and $\pi'[0 \dots k - i] = \pi[i \dots k]$.

Intuitively, the path π has a pumpable segment if, after some initial sequence of tile types, it consists of a subsequence of tile types which is repeated in the same direction an infinite number of times, one after another. Figure 1 shows an assembly in which the same tile type repeats along




Figure 1: An assembly containing a path with repeating tiles A-A that do *not* form a pumpable segment, because they are blocked from infinite growth by the existing assembly. Note, however, that any sufficiently long path from the origin (at the left) contains a pumpable segment, namely the repeating segment 1-2-3-4-1 along the bottom row, which can be pumped infinitely to the right.

a path, but the segment between the occurrences is not pumpable.

Let $c \in \mathbb{N}$ and $\vec{v} \in \mathbb{Z}^2$. The *diamond of radius c centered about the point \vec{v}* is the set of points defined as $D(c, \vec{v}) = \{ (x, y) + \vec{v} \mid |x| + |y| \leq c \}$. Let $\vec{w}, \vec{x} \in \text{dom } \alpha$, $c \in \mathbb{N}$, and π be a simple path from \vec{w} to \vec{x} in the binding graph G_α . We say that π is a *pumpable path from \vec{w} to \vec{x} in α* if it contains a pumpable segment $\pi[i \dots k]$ for some $i, k \in \mathbb{N}$ such that $0 \leq i < k \leq |\pi|$. We say that \mathcal{T} is *c -pumpable* if there exists $c \in \mathbb{N}$ such that for every $\vec{w}, \vec{x} \in \text{dom } \alpha$ with $\vec{x} \notin D(c, \vec{w})$, there exists a pumpable path π from \vec{x} to \vec{w} in α .

Figure 2 shows, from left to right, (1) a partially complete assembly beginning from the (dark grey) seed tile, where the dark notches between adjacent tiles represent strength 1 bonds, and a tile selected for the example, (2) the full path leading from the seed to the selected tile, (3) the tile types for a segment of the path, showing the repeating pattern of tile types ‘1-2-3-4’, and (4) an extended version of the path which shows its ability to be pumped.

In our proof, it is helpful to consider extending an assembly in such a way that no individual




Figure 2: A partial assembly, selected path, pumpable segment, and pumpable path

tile in the existing assembly is extended by more than a finite amount (though an infinite assembly may have an infinite number of tiles that can each be extended by a finite amount). We call such an extension the *finite closure* of the assembly, and define it formally as follows. Let $\alpha' \in \mathcal{A}[\mathcal{T}]$. We say that the *finite closure* of α' is the unique assembly $\mathcal{F}(\alpha')$ satisfying

1. $\alpha' \sqsubseteq \mathcal{F}(\alpha')$, and
2. $\text{dom } \mathcal{F}(\alpha')$ is the set of all points $\vec{x} \in \mathbb{Z}^2$ such that every infinite simple path in the binding graph G_α containing \vec{x} intersects $\text{dom } \alpha'$.

Intuitively, this means that if we extend α' by only those “portions” that will eventually stop growing, the finite closure is the super-assembly that will be produced. That is, any attempt to “leave” α' through the finite closure and go infinitely far will eventually run into α' again. If α' is terminal, then α' is its own finite closure. Note that in general, the finite closure of an assembly α' is not the result of adding finitely many tiles to α' . For instance, if infinitely many points of α' allow exactly one tile to be added, the finite closure adds infinitely many points to α' . However, the finite closure of a finite assembly is always a finite assembly.




Figure 3: Example of a finite closure. The dark gray points represent locations at which tiles can attach.

For an example of a finite closure of an assembly, see Figure 3. Figure 3c shows the terminal assembly which consists of three rows of tiles that continue infinitely to the right (denoted by the arrow), with a 10 tile upward projection occurring at every fourth column. Figure 3a shows an assembly α' which consists of two rows of tiles continuing infinitely to the right but with an incomplete bottom row and none of the full upward projections. Figure 3b shows $\mathcal{F}(\alpha')$, which is the finite closure of α' . Notice that an infinite number of tiles were added to α' to create $\mathcal{F}(\alpha')$ since an infinite number of upward projections were added, each consisting of only 10 tiles. However, also

note that the bottom row is not grown because that row consists of an infinite path and therefore cannot be part of the finite closure.




Figure 4: Example of a comb connected to a hard-coded assembly (dark tiles that spiral). The seed is the tile at the center of the spiral. The comb originates at the bottom most black tile. Note that a finite number of “attempts at teeth” (one in the above example) may be blocked before the infinite teeth are allowed to grow.

Let π^{\rightarrow} be an eventually \vec{v} -semi-periodic path in α originating at $\vec{0}$ where $\vec{0} \neq \vec{v} \in \mathbb{Z}^2$. Let $\vec{u} \in \mathbb{Z}^2$ such that $\vec{u} \neq z \cdot \vec{v}$ for all $z \in \mathbb{R}$. Suppose that there is a point \vec{b} on the tail of π^{\rightarrow} such that there is an eventually \vec{u} -semi-periodic path π^{\uparrow} in α originating at \vec{b} such that $\pi^{\rightarrow} \cap \pi^{\uparrow} = \{\vec{b}\}$.

Define the assembly $\alpha^* = \alpha \upharpoonright \left(\pi^{\rightarrow} \cup \bigcup_{n \in \mathbb{N}} (\pi^{\uparrow} + n \cdot \vec{v}) \right)$. It is easy to see that α^* need not be a

producible assembly. We say that α^* is a *comb in* (with respect to π^{\rightarrow} and π^{\uparrow}) if, for every $n \in \mathbb{N}$, $\alpha^* \upharpoonright \pi^{\uparrow} + n \cdot \vec{v} = \alpha^* \upharpoonright (\pi^{\uparrow} + n \cdot \vec{v})$. We refer to the assembly $\alpha^* \upharpoonright \pi^{\rightarrow}$ as the *base* of α^* . For any $n \in \mathbb{N}$, we define the n^{th} *tooth* of α^* to be the assembly $\alpha^* \upharpoonright (\pi^{\uparrow} + n \cdot \vec{v})$. We say that the comb α^* *starts at the point* $\pi[s]$ (the point at which the eventually-periodic path π^{\rightarrow} becomes periodic). It follows from the definition that, if α^* is a comb in α , then $\alpha^* \in \mathcal{A}[\mathcal{T}]$.

See Figure 4 for an example of a comb. A comb, intuitively, is a generalization of the assembly in which an infinite periodic one-way path (the base) grows along the positive x -axis, and once per period, an infinite periodic path (a tooth) grows in the positive y direction, creating an infinite number of “teeth”. The generalizations are that (1) the base and teeth need not run parallel to either axis, and (2), the teeth may have some initial hard-coded tiles before the repeating periodic segment begins. Also, it is possible at temperature 1 to build multiple combs with the same base, but with different teeth, growing in either direction.

4 Main Result

We show in this section that only “simple” sets weakly self-assemble in directed, pumpable tile assembly systems at temperature 1. We now formally define “simple.”

Definition 4.1. A set $X \subseteq \mathbb{Z}^2$ is *semi-doubly periodic* if there exist three vectors \vec{b} , \vec{u} , and \vec{v} such that

$$X = \left\{ \vec{b} + n \cdot \vec{u} + m \cdot \vec{v} \mid n, m \in \mathbb{N} \right\}.$$

Note that a set that is periodic along only one dimension is also semi-doubly periodic, by our definition, since this corresponds to the condition that exactly one of the vectors \vec{u} or \vec{v} is equal to $\vec{0}$ (or if $\vec{u} = \vec{v}$). Similarly, if $\vec{u} = \vec{v} = \vec{0}$, then the definition of semi-doubly periodic is equivalent to A being a singleton set. The following observation justifies the intuition that finite unions of semi-doubly periodic sets constitute only the computationally simplest subsets of \mathbb{Z}^2 .

Observation 4.2. *Let $A \subseteq \mathbb{Z}^2$ be a finite union of semi-doubly periodic sets. Then the unary languages $L_{A,x} = \{ 0^{|x|} \mid (x, y) \in A \text{ for some } y \in \mathbb{Z} \}$ and $L_{A,y} = \{ 0^{|y|} \mid (x, y) \in A \text{ for some } x \in \mathbb{Z} \}$ consisting of the unary representations of the projections of A onto the x -axis and y -axis, respectively, are regular languages.*

The following theorem is the main result of this paper.

Theorem 4.3. *Let $\mathcal{T} = (T, \sigma, 1)$ be a directed, pumpable TAS. If a set $X \subseteq \mathbb{Z}^2$ weakly self-assembles in \mathcal{T} , then X is a finite union of semi-doubly periodic sets.*

The proof idea of Theorem 4.3 is as follows. Suppose that $\alpha \in \mathcal{A}_{\square}[\mathcal{T}]$. Note that α is unique since \mathcal{T} is directed. Either α is an infinite “grid” that fills the plane, or there exists finitely many semi-doubly periodic combs and semi-periodic paths that, taken together, “cover” every point in X (in the sense that each such point is in the finite closure of one of these combs or paths).

The reason for this is that each comb is defined by two vectors \vec{u} (the base) and \vec{v} (the teeth), and these vectors form a “basis” for the space of points located within the cone formed by the base and the first tooth of the comb. While the vectors do not reach every point in this cone, they reach within a constant distance of every point in the cone, and the doubly periodic regularity of the teeth and base enforces doubly periodic regularity in between the teeth as well. Of course, not all combs have teeth, in which case the comb is just a periodic path.

We associate each black tile with some periodic path or comb that begins in a fixed radius about the origin (utilizing the fact that a path cannot go far from the origin without having pumpable segments that can be used to construct a periodic path or comb). The finite number of combs and periodic paths originating within this radius tells us that the number of semi-doubly periodic sets of (locations tiled by) black tiles that they each define is finite.

A simple application of Theorem 4.3 is that no discrete self-similar fractal weakly (and hence strictly) self-assembles in any temperature 1 tile assembly system that is pumpable and directed. Since Winfree [10] showed that one particular discrete self-similar fractal, the discrete Sierpinski triangle, weakly self-assembles at temperature 2, this provides a concrete example of computation that is possible (and simple) at temperature 2, but impossible at temperature 1, assuming directedness and pumpability.

5 Conclusion

We have studied the class of shapes that self-assemble in Winfree’s abstract tile assembly model at temperature 1. We introduced the notion of a pumpable temperature 1 tile assembly system and then proved that, if X weakly self-assembles in a pumpable, directed tile assembly system, then

X is necessarily “simple” in the sense that X is merely a finite union of semi-doubly periodic sets. Finally, we conjecture that our results hold in the absence of the pumpability hypothesis.

Conjecture 5.1. *Let $\mathcal{T} = (T, \sigma, 1)$ be a directed tile assembly system and $\alpha \in \mathcal{A}_{\square}[\mathcal{T}]$. If $\text{dom } \alpha$ is infinite, then \mathcal{T} is pumpable.*

It is always possible to produce long paths in which the presence of a segment with two repetitions of a tile type does not imply that the segment is pumpable. However, in every case we consider, there is always a previous segment of the path that is pumpable. Proving Conjecture 5.1 would imply that every directed tile set weakly self-assembles a finite union of doubly periodic sets.

We also leave open the question of whether the hypothesis of directedness may be removed. We use the property of directedness at many points in our proof, but in some cases, a more careful and technically convoluted argument could be used to show that the tile set need not be directed. Intuitively, an undirected tile set \mathcal{T} that weakly self-assembles a set $X \subseteq \mathbb{Z}^2$ is deterministic in that all terminal assemblies of \mathcal{T} “paint” exactly the points in X black, but is nondeterministic in the sense that different terminal assemblies of \mathcal{T} may place different tiles in the same location (including different black tiles at locations in X), and may even place non-black tiles at locations in one terminal assembly that are left empty in other terminal assemblies. Undirected tile assembly systems that weakly self-assemble a unique set X exist, but in every case that we know of, the undirected tile set may be replaced by a directed tile set self-assembling the same set.

If both hypotheses of directedness and pumpability could be removed from the entire proof, then our main result would settle the case of computation via self-assembly at temperature 1, by showing that every temperature 1 tile assembly system weakly self-assembles a finite union of semi-doubly periodic sets if it weakly self-assembles any set at all. As indicated in the introduction, we interpret this statement to imply that general-purpose deterministic computation is not possible with temperature 1 tile assembly systems.

Finally, we make the observation that universal computation at temperature 1 is possible by changing the underlying model of self-assembly. For instance, a construction that was personally communicated by Matt Cook (also mentioned briefly in [4]) establishes that the abstract TAM is computationally universal with respect to directed, temperature 1 tile assembly systems that place tiles in three spatial dimensions. In fact, Cook’s construction only uses two integer planes, “stacked” one on top of the other, to simulate an arbitrary cellular automaton (no such directed two-dimensional universality result is known at the time of this writing). Moreover, universality can also be achieved if negative glue strengths and interaction between differently colored glues (a so-called *non-diagonal strength function*) are allowed.

Theorem 5.2. *For every single-tape Turing machine M , there is a tile set T with negative, non-diagonal glue strengths, which simulates M in the following way. Given an input string x , define $\mathcal{T}_x = (T, \sigma_x, 1)$ to be the temperature 1 TAS where σ_x is the seed assembly satisfying $\text{dom } \sigma_x = \{0, \dots, |x|-1\} \times \{0\}$ that encodes the initial configuration of M . Then \mathcal{T}_x simulates the computation of $M(x)$, with the configuration of $M(x)$ after n steps represented by the line $y = n$ in the terminal assembly of \mathcal{T}_x .*

Intuitively, by introducing negative glue strengths, we allow for the cooperation that is impossible with only nonnegative glue strengths. The difference with temperature 2 is that, at temperature 2, one may enforce that no tile appears at a certain position until two neighbors are present. At temperature 1 and with negative glue strengths, on the contrary, we do not enforce that two tiles

are present before a position can be given a tile. However, we enforce that if the wrong tile binds in this position, eventually the error is corrected by the presence of a neighboring tile which forces the removal of the incorrect tile using negative glue strengths.

Acknowledgment

We wish to thank Maria Axenovich, Matt Cook, and Jack Lutz for useful discussions. We would especially like to thank Niall Murphy, Turlough Neary and Damien Woods for inviting us to present a preliminary version of this research at the International Workshop on The Complexity of Simple Programs, University College Cork, Ireland on December 6th and 7th, 2008.

References

- [1] Leonard M. Adleman, Jarkko Kari, Lila Kari, and Dustin Reishus, *On the decidability of self-assembly of infinite ribbons*, Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002, pp. 530–537.
- [2] James I. Lathrop, Jack H. Lutz, and Scott M. Summers, *Strict self-assembly of discrete Sierpinski triangles*, Theoretical Computer Science **410** (2009), 384–405.
- [3] Matthew J. Patitz and Scott M. Summers, *Self-assembly of decidable sets*, Proceedings of The Seventh International Conference on Unconventional Computation (Vienna, Austria, August 25-28, 2008), 2008.
- [4] Paul W. K. Rothmund, *Theory and experiments in algorithmic self-assembly*, Ph.D. thesis, University of Southern California, December 2001.
- [5] Paul W. K. Rothmund and Erik Winfree, *The program-size complexity of self-assembled squares (extended abstract)*, STOC '00: Proceedings of the thirty-second annual ACM Symposium on Theory of Computing (New York, NY, USA), ACM, 2000, pp. 459–468.
- [6] Paul W.K. Rothmund, Nick Papadakis, and Erik Winfree, *Algorithmic self-assembly of dna sierpinski triangles*, PLoS Biology **2** (2004), no. 12.
- [7] Nadrian C. Seeman, *Nucleic-acid junctions and lattices*, Journal of Theoretical Biology **99** (1982), 237–247.
- [8] Hao Wang, *Proving theorems by pattern recognition – II*, The Bell System Technical Journal **XL** (1961), no. 1, 1–41.
- [9] ———, *Dominoes and the AEA case of the decision problem*, Proceedings of the Symposium on Mathematical Theory of Automata (New York, 1962), Polytechnic Press of Polytechnic Inst. of Brooklyn, Brooklyn, N.Y., 1963, pp. 23–55.
- [10] Erik Winfree, *Algorithmic self-assembly of DNA*, Ph.D. thesis, California Institute of Technology, June 1998.