One Tile to Rule Them All: Simulating Any Tile Assembly System with a Single Universal Tile * **

Erik D. Demaine[1], Martin L. Demaine[1], Sándor P. Fekete[2],
Matthew J. Patitz[3], Robert T. Schweller[4],
Andrew Winslow[5], and Damien Woods[6]

[1] Massachussetts Institute of Technology, Cambridge, MA 02139, USA.
{edemaine,mdemaine}@mit.edu
[2] TU Braunschweig, Germany. s.fekete@tu-bs.de
[3] University of Arkansas, Fayetteville, AR 72701, USA. patitz@uark.edu
[4] University of Texas–Pan American, Edinburg, TX 78539, USA.
rtschweller@utpa.edu
[5] Tufts University, Medford, MA 02155, USA. awinslow@cs.tufts.edu
[6] California Institute of Technology, Pasadena, CA 91125, USA. woods@caltech.edu

Abstract. In the classical model of tile self-assembly, unit square *tiles* translate in the plane and attach edgewise to form large crystalline structures. This model of self-assembly has been shown to be capable of asymptotically optimal assembly of arbitrary shapes and, via information-theoretic arguments, increasingly complex shapes necessarily require increasing numbers of distinct types of tiles.

We explore the possibility of complex and efficient assembly using systems consisting of a single tile. Our main result shows that any system of square tiles can be simulated using a system with a single tile that is permitted to flip and rotate. We also show that systems of single tiles restricted to translation only can simulate cellular automata for a limited number of steps given an appropriate seed assembly, and that any longer-running simulation must induce infinite assembly.

Keywords: DNA computing, algorithmic self-assembly, hexagonal tiles

1 Introduction

This paper shows that many copies of a single rotatable polygonal tile type suffices to simulate any square tile assembly system.

---

* A full version of this paper can be found at http://arxiv.org/abs/1212.4756

** Research of Matthew J. Patitz supported in part by NSF grant CCF-1117672. Research of Robert Schweller supported in part by NSF grant CCF-1117672. Research of Andrew Winslow supported in part by NSF grant CDI-0941538. Research of Damien Woods Supported by NSF grants 0832824 & 1317694 (the Molecular Programming Project), CCF-1219274, and CCF-1162589.

Winfree [15] introduced the *abstract Tile Assembly Model* (aTAM) as a clean theoretical model for nanoscale self-assembling systems. In several experiments of increasing complexity and reliability [16,2,5,13,3], this model has been shown to be physically practical, with tiles composed of DNA strands. As a result, the aTAM has become the standard in theoretical work on self-assembly, with previous work exploring its abilities and limitations in terms of its ability to use computation to assemble shapes and patterns [12,1,14,6], as well comparing its computational power and simulation abilities [7,10,18].

In the aTAM, we start with a single specific tile, or a connected assembly of tiles, (called the *seed*) and repeatedly add any tile to the assembly that has enough matching glues (colored edges) to "stick" to the rest of the assembly. Each glue type (color class) has a natural number *strength*, which represents the affinity for matching glues of that type, and a global *temperature* (typically 2) of the system specifies the total required strength for a tile to attach to the assembly. Unlike Wang tiling, in the aTAM we can never throw away partially formed assemblies; in fact, the aTAM can be seen as a special kind of asynchronous, and nondeterministic, cellular automaton. See Section 2 for more details.

1.1 Our results I: Universal self-assembly with one tile

We prove in Section 5 that any aTAM system can be simulated by just a single tile, in a generalization of the aTAM model called the polygonal free-body Tile Assembly Model (pfbTAM). More precisely, we show that any temperature-τ aTAM system can be converted into a temperature-τ pfbTAM system with a single tile type t_U such that the two systems have exactly the same producible assemblies, modulo isometry. This construction is *self-seeding* in the sense that it starts from a single copy of the very same tile t; it is even a challenge to get the next copy of t to attach without uncontrolled infinite growth.

Another contribution of this paper is in our proof technique: we use a chain of four simulations. Such long chains of simulations, or reductions, are commonplace throughout the theory of computation, but not seen (so far) in self-assembly. Our single tile simulates an arbitrary square tile assembly system $\mathcal{T}$ as follows. We begin with the fact that the aTAM is intrinsically universal [7], which means that there is a single set of square tiles U that can simulate any square tile assembly system $\mathcal{T}$. Via the construction in [7], the tile assembly system $\mathcal{T}$ that we wish to simulate is first encoded as a seed assembly using tiles from the intrinsically universal tile set U to give a tile assembly system $\mathcal{U}_\mathcal{T}$. Next, the system $\mathcal{U}_\mathcal{T}$ is simulated by a "low-strength" hexagonal tile assembly system, as described below. The main result of this paper is that the resulting hexagonal assembly system can be simulated by a pfbTAM system consisting of only a single tile t_U. And, of course, our single tile t_U works for any such system $\mathcal{T}$ we wish to simulate. In particular, both the geometry and the dynamics of the simulated system are, modulo rescaling, precisely simulated: for every sequence of tile placement in the simulated system there is a sequence of blocks of tile placements of t_U in the simulator system, and vice-versa. Hence, when appropriately seeded, t_U

assembles scaled-up versions of what is assembled by $\mathcal{T}$, and *in the same way* that $\mathcal{T}$ does it.

It is worth noting that the notion of intrinsic universality, with its strict notion of simulation, gives a framework to compare the power of tile assembly models that at first sight seem very different and perhaps difficult to compare. For example, in this paper we study rotatable and fliapable polygons, and translatable polyominos, yet simulation gives a way to directly compare the power of this model with the well-known square (aTAM) model. Intrinsic universality is giving rise to a kind of complexity theory for self-assembly systems allowing us to tease apart the power of different models [18].

Our universal tile t_U is a kind of geometric analog to a universal Turing machine, simultaneously simulating the shape construction and computational ability of an arbitrary tile assembly system (although our definition of simulation is in fact stronger—we care about dynamics, not merely input to output mappings). The existence of a system with just a single tile demonstrates that geometry alone (as opposed to, say, a large, although constant, number of square tile types [7]) suffices to bootstrap a system of self-assembly in even the most restrictive case where the system may only utilize copies of a single shape.

The pfbTAM model differs from aTAM in two ways: tiles consist of general simple polygons rather than squares and tiles are (possibly) permitted to rotate. Both are physically realistic aspects of self-assembling systems. For example, DNA origami [11] is a rapidly evolving technology that has been used to successfully build numerous complex shapes using strands of DNA. The technology has evolved to the point where free software automatically designs DNA to fold into essentially arbitrary desired shapes. Polyomino generalizations of square aTAM tiles have already been developed in practice [17] and studied in theory [8]. Rotation is clearly a natural attribute of all physical systems – prior work in the aTAM used a simple trick to eliminate rotation. Although our single polygonal tile t_U has such a large number of sides that its fabrication would be extremely challenging, our work here demonstrates that rotation can be used as an encoding mechanism to design systems that reuse a single tile at various rotations to achieve universal tile assembly systems. It would not be inconceivable to build a single tile with a more modest number of sides that simulates a simple square tile system.

The full version of this paper also shows the existence of a single tile that simulates any square tile system in the Wang model of plane tiling.

1.2 Our results II: Hexagonal tile assembly systems

As noted above, part of our proof involves the use hexagonal tiles: Section 4 describes aTAM systems with unit-sized hexagonal tiles on a hexagonal grid. The only previous paper considering this model [9] simply showed differences between squares and hexagons with respect to infinite constructions. Here we show that any temperature-2 square aTAM system can be simulated by a temperature-2 hexagonal aTAM system in which all glues have strength at most 1. The construction works at a scale factor of only 3: each square tile is simulated by a

3×3 block of hexagonal tiles. The main reason we use hexagons is that no such system is possible for square systems: any temperature-2 square aTAM system in which all glues have strength at most 1 cannot grow outside its bounding box (and so it cannot simulate arbitrary square-tile systems).

This result is a key step to proving our main positive result (aTAM simulation allowing translation and rotation). Specifically, we show in Section 5 how to simulate any temperature-2 hexagonal aTAM system that uses exclusively strength-1 glues with a rotatable polygon t_U that encodes different tile types by attaching at different rotation angles. Independent of their use in our constructions, hexagonal systems without strength-2 glues could be significantly easier to implement in the laboratory than square systems using both strength-1 and strength-2 glues in arbitrary arrangements on the tiles.

1.3 Our results III: Linear-time computation with a single translation-only tile

In Section 6, we prove both a positive and a negative result on single-tile systems where the tile is forbidden from rotating. On the positive side, we prove that single-tile translation-only systems have non-trivial power: they are capable of time-bounded simulation of computationally universal 1D cellular automaton. Any 1D cellular automata that runs for n steps can be simulated starting with a seed assembly of $O(n)$ tiles. This is proved by first showing that single-tile translation-only systems can simulate a restricted class of multi-tile systems, which have previously been shown to simulate computationally universal 1D cellular automata.

On the negative side, we prove that translation-only single-tile systems need a seed assembly consisting of at least four tiles to carry out any non-trivial assembly. More formally, we prove that any single-tile translation-only system with a seed tile consisting of one, two, or three tiles either produces an infinite assembly, or only the seed assembly. More generally, we conjecture that no finite seed suffices for unbounded computational power with single-tile, translation-only systems, in stark contrast to our result that general single-tile pfbTAM systems have this power starting with a single-tile seed.

2 Model Definitions

The *polygonal free-body Tile Assembly Model* or *pfbTAM* generalizes self-assembly models such as the aTAM by using tiles with arbitrary polygonal shapes that may be translated and rotated. In our positive results, we only utilize tiles whose shapes are convex regular n-gons with small surface geometries. Our negative results are valid for arbitrary polygons, as discussed in Section 6.

A pfbTAM system Γ is defined as $\Gamma = (T, \Sigma, \tau, \sigma)$, where T is a *tile set* of polygonal *tiles*, Σ is a collection of *glue types*, $\tau \in \mathbb{N}$ is the *temperature* of the system, and σ is a *seed assembly* consisting of an arrangement of tiles from T and their locations. Each tile in T has a shape defined by a simple polygon (a polygon without holes), and each side of the tile has is assigned a *glue* from the collection of glue types Σ of Γ. Each glue type $g \in \Sigma$ is assigned a positive integer value called a *strength* with the exception of special *null glue* whose strength is 0.

If a pair of tiles of Γ are arranged in the plane such that their interiors do not overlap, and a pair of their edges are coincident, then these edges are said to form a *bond*. The *strength* of this bond is determined by the glues of both sides, with the strength equal to the strength of the glue if both sides have the same glue type, and zero otherwise.

A collection of tiles arranged in the plane whose interiors are pairwise disjoint is an *assembly*. The *bond graph* of an assembly is the (planar) multigraph consisting of labeled nodes for each tile, and an edge between a pair of nodes for each positive-strength bond the tiles share. An assembly is *τ-stable* if any edge-cut of the bond graph of the assembly has cut edges whose total corresponding bond strength meets or exceeds τ, the temperature of the system.

The seed assembly σ of Γ is a τ-stable assembly consisting of the tiles in T. The assembly process consists of attaching single tiles of T to a growing τ-stable assembly, beginning with σ, the seed assembly of the system. Because each single-tile attachment must yield a τ-stable assembly, a tile can attach to the growing assembly if and only if it is able to form bonds with assembly whose total strength is at least τ. Any assembly A that can be formed by this process is a *producible assembly of Γ* and is said to be *produced by Γ*. If no tile can attach to A, then A is also a *terminal assembly of Γ*. In some cases we consider pfbTAM systems in which tiles are not permitted to rotate, but instead merely translate. We call these system *translation-only pfbTAM systems*.

The well-studied aTAM and hTAM models (reviewed in Section 1) are both special cases of translation-only pfbTAM systems. An *abstract Tile Assembly Model (aTAM) system* is a translation-only pfbTAM system $\Gamma = (T, \Sigma, \tau, \sigma)$ where the tiles in T are unit squares, while a *hexagonal Tile Assembly Model (hTAM) system* is a system where the tiles in T are unit hexagons.

In Section 6 we also consider a restricted class of aTAM systems (rotated by 45°) called *pyramid aTAM systems*, proving that single-tile, translation-only pfbTAM systems are capable of simulating them. An aTAM system $\Gamma = (T, \Sigma, 2, \sigma)$ is said to be a *pyramid aTAM system* if three conditions hold. First, σ contains n tiles configured in the format described in Figure 1, with the property that all coincident tile edges have matching glues. Second, all glues in T have strength 1. Third, the tile set T and seed σ are such that no tiles can attach to the southern face of the seed. In addition, a pyramid aTAM system is said to be *double-checkerboarded* if every tile attaching to the seed assembly has distinct tile types at the locations to the southwest, south, and southeast of the tile's attachment location. An example of a coloring scheme that denotes which tiles must be of differing type is shown in Figure 1.

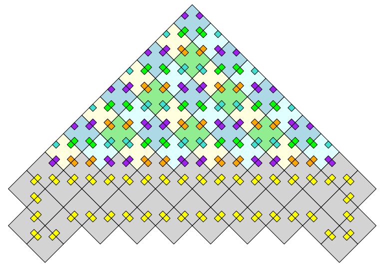

Fig. 1. Pyramid aTAM systems start with a seed assembly (gray) and grow upwards with cooperative temperature-2 bonding, yielding an assembly that is maximally pyramid-shaped. They can be used to simulate the light cone of a cellular automaton.

3 Simulation Definitions

En route to proving that pfbTAM systems with a single tile type are powerful, we prove that pfbTAM and hTAM systems can capture the behavior of, or *simulate*, aTAM systems. Below we define what it means for a pfbTAM system to simulate hTAM and aTAM systems, and for an hTAM system to simulate an aTAM system.

Loosely defined, a system simulates another if there is a mapping between the producible assemblies of both systems, such that a producible assembly A yields another producible assembly A' via a single-tile addition in one system if and only if, in the other system, the analogous assembly to A yields an analogous assembly to A' via a single-tile addition. In Sections 3.2 and 3.3 the simulating systems use a *block* of tiles to represent a single tile in the simulated system, and each single-tile addition in the simulated system is equivalent to a short sequence of single-tile additions in the simulated system, where the final addition completes the simulation of the single-tile addition in the simulated system.

3.1 Simulating hTAM systems with pfbTAM systems

A pfbTAM system $\Gamma_p = (T_p, \Sigma_p, \tau_p, \sigma_p)$ *simulates* an hTAM system $\Gamma_h = (T_h, \Sigma_h, \tau_h, \sigma_h)$ if there exists a mapping $\phi : T_p \times [0, 2\pi) \rightarrow T_h$ of orientations (specified by an angle in $[0, 2\pi)$) of tile in T_p to tiles in T_h such that there exists a bond graph G_{A_p} generated by a producible assembly A_p of Γ_p if and only if mapping the label of each node v of G_{A_p} to $\phi(p)$ yields a bond graph G_{A_h} of a producible assembly A_h of Γ_h.

Also, for each producible assembly A'_p produced by Γ_p via a single-tile addition to assembly A_p, an assembly A'_h in Γ_h equivalent to A'_p via ϕ can be produced by Γ_h via a single-tile addition to assembly A_h equivalent to A_f via ϕ and vice versa. In other words, equivalent assemblies are producible in both systems by equivalent sequences of tile additions.

3.2 Simulating aTAM systems with hTAM systems

Our definition of pfbTAM systems simulating hTAM systems uses a strict one-to-one-correspondence between tiles in the simulated and simulating systems. Here and in Section 3.3, our definition of hTAM systems simulating aTAM systems has a slightly weaker correspondence called a *c-block representation* where each tile in the simulated aTAM system corresponds to a $c \times c$ grid of tiles in the simulating hTAM system.

Let A_h be an assembly of an hTAM system $\Gamma_h = (T_h, \Sigma_h, \tau_h, \sigma_h)$ and A_a be an assembly of an aTAM system $\Gamma_a = (T_a, \Sigma_a, \tau_a, \sigma_a)$. Then A_h is a *valid c-block representation of* A_a for an odd, positive integer c and partial function $\phi : T_h \rightarrow T_a$ if two conditions hold. First, that A_h is evenly divisible in $c \times c$ blocks of tiles, as shown in Figure 2. Second, that x is in the domain of ϕ if and only if x is at the center of a $c \times c$ block.

Given a valid c-block representation A_h, define the *c-bond graph of* A_h to be a graph with a labeled node for each center tile x of a $c \times c$ block with label $\phi(x)$. The c-bond graph of A_h has an edge between two nodes corresponding to tiles x and x' if the bond graph of A_h has a length-c path between x and x' consisting of edges between tiles exclusively at angles $90°$ and $-90°$, or $120°$ and $-30°$.

Now we are ready to define simulation. We say that Γ_h *simulates* Γ_a *at scale* c, if there exists a partial function $\phi : T_h \rightarrow T_a$ such that three conditions hold. First, every tile in any producible assembly of Γ_h of size greater than $c^2 - 1$ is within distance at most c from a tile x for which $\phi(x)$ is defined. Second, there exists a producible assembly A_h of Γ_h that is a valid c-block representation for function $\phi(x)$, if and only if mapping the label of each node v in the c-block bond graph of A_h yields a bond graph of a producible assembly of Γ_a. Third, for each producible assembly A'_a of Γ_a produced by Γ_a via a single-tile addition to assembly A_a, there are equivalent c-block representation assemblies A'_h and A_h of Γ_h, such that A'_h is producible from A_h via a sequence of tile additions, for which each producible assembly created during this sequence of tile additions, the thing in Γ_h is A_h.

3.3 Simulating aTAM systems with pfbTAM systems

We define a c-scaled simulation of an aTAM system by a pfbTAM system by mapping $c \times c$ blocks within pfbTAM assemblies to aTAM tiles, where this mapping reads rotations of pfbTAM tiles in the blocks. A pfbTAM system $\Gamma_p = (T_p, \Sigma_p, \tau_p, \sigma_p)$ *simulates* an aTAM system $\Gamma_a = (T_a, \Sigma_a, \tau_a, \sigma_a)$ at scale $c \in \mathbb{N}$ if the following conditions hold, based on the more formal definition of [7].

First, there exists a mapping $\phi : ((T_p \cup \{\varnothing\}) \times [0, 2\pi))^{c^2} \rightarrow T_a \cup \{\varnothing\}$ of $c \times c$ blocks of tiles from T_p (with the output of ϕ depends on the orientations of those tiles, specified by a rotation angle in $[0, 2\pi)$) and empty locations (denoted $\varnothing$) to tiles in T_a and empty locations such that for every producible assembly A_p of Γ_p there is a producible assembly A_a in Γ_a, where $A_a = \phi^*(A_p)$ and for every producible assembly A_a of Γ_a there exists a producible assembly A_p in Γ_p, where $A_p = \phi^*(A_a)$ (here ϕ^* denotes the function ϕ applied to an entire assembly, in the most obvious block-wise way). We also require that A_p maps *cleanly* to A_a under ϕ^*, that is, for all non-empty $c \times c$ blocks b in A_p it is the case that at least one neighbor of $\phi(b)$ in $\phi^*(A_p)$ is non-empty, or else A_p has at most one non-empty $c \times c$ block. In other words, π may have tiles in $c \times c$ blocks representing empty space in α, but only if that position is adjacent to a tile in α.

Second, there exist producible assemblies A_a and A'_a of Γ_a such that $A_a \to_1 A'_a$ (growth by single tile addition), then for every producible assembly A_p of Γ_f, where $A_a = \phi^*(A_p)$ it is the case that there exists A'_p such that $A_p \to_* A'_p$ (growth by one or more tile additions) in Γ_f, where $A'_a = \phi^*(A'_a)$. Furthermore, for every pair of producible assemblies A_p, A'_p of Γ_f, if $A_p \to_* A'_p$, and $A_a = \phi^*(A_p)$ and $A'_p = \phi^*(A'_p)$, then $A_a \to_* A'_a$ for assemblies A_a, A'_a of Γ_a.

3.4 Simulating pyramid aTAM systems with single-tile translation-only pfbTAM systems

In Section 6 we simulate *pyramid aTAM systems* (defined in Section 2), a special class of aTAM systems that have significant computational power but limited enough to permit simulation by translation-only, single-tile pfbTAM systems. Recall that a pyramid aTAM system is a restricted aTAM system in which each tile must attach to the growing assembly using exactly the southwest and southeast sides. The key idea of the simulation is to place an imaginary grid of boxes over a given assembly in the simulating translation-only system to define the position each tile (and specifically the tile's north/south position) and thus the tile of the aTAM system this tile is simulating; see Figure 3 for the idea.

We now define the mapping of an assembly in a single-tile, translation-only pfbTAM system $\Gamma_p = (T_p, \Sigma_p, 3, \sigma_p)$ to an assembly in the simulated pyramid aTAM system $\Gamma_a = (T_a, \Sigma_a, 2, \sigma_a)$. Consider a 2-stable assembly A consisting of translations of tiles of type p. Now consider the westmost, southmost tile in A. Assume this tile sits at coordinate position $(0, -x_1)$. Define a partial mapping $f : Z \times Z \to Z \times Z \times T$ that maps tile coordinate locations within an assembly to both a 2D coordinate position and a tile type in T.

Given the partial mapping f, for an assembly A produced by Γ_p, we say A simulates the assembly A' in Γ_a if A' is the assembly obtained by including each tile of type t at position (w, y), such that $f(x, y) = (w, u, t)$ for some tile in A at position (x, y). If any tile in A is at a position at which f is not defined, then A does not have a defined mapping to a square aTAM tile assembly over T.

We say that Γ_p *terminally simulates*, or simply *simulates*, Γ_a if the set of terminal assemblies of Γ_p maps exactly to the set of terminal assemblies of Γ_a. [7]

4 Low-Strength hTAM Systems Simulate aTAM Systems

In this section we prove that any temperature-τ aTAM system can be simulated by a temperature-τ hTAM system that uses only glues of strength *less than*

[7] This is a weaker definition of simulation than what is defined for the other pairs of models. While our construction actually satisfies an equivalent stronger definition, we omit the more involved simulation definition for simplicity.

τ (called *low-strength glues*). We call these hTAM systems *low-strength hTAM systems*. This is used later in Section 5 to simulate temperature-τ aTAM systems with single-tile pfbTAM systems (Lemma 2). See Figure 2 for an example.

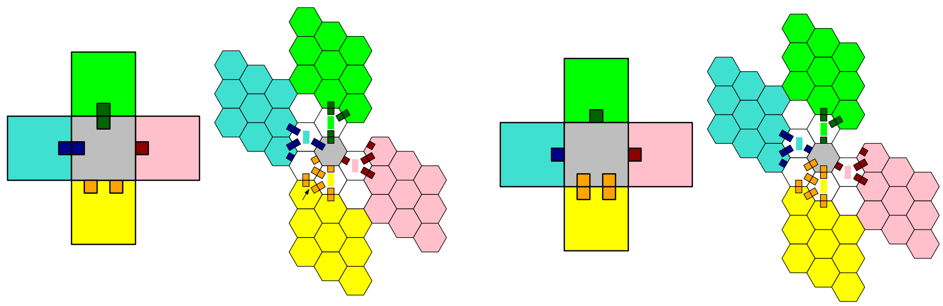

Fig. 2. Simulating a non-deterministic attachment in the aTAM via a low-strength hTAM system. The center location in the 3×3 block is a location of contention; multiple blocks compete and/or cooperate to claim it.

Lemma 1. *For any aTAM system $\Gamma_a = (T_a, \Sigma_a, \tau, \sigma_a)$ with $|\sigma_a| = 1$ and $\tau \geq 2$, there exists an hTAM system $\Gamma_h = (T_h, \Sigma_h, \tau, \sigma_h)$ that simulates Γ_a at scale 3 and has the property that all glues in Σ_h are of strength less than τ. Also, $|T_h| = O(|T_a|^2)$ and $|\sigma_h| = 3$.*

Furthermore, it is straightforward to see from the above construction that an aTAM system with seed σ_a and $|\sigma_a| \geq 1$, i.e. a seed assembly consisting of multiple tiles, can be simulated by an hTAM system, where the $9|\sigma_a|$ hexagonal tiles simulating the aTAM seed assembly are appropriately placed to represent that seed assembly. This gives the following corollary:

Corollary 1. *For any aTAM system $\Gamma_a = (T_a, \Sigma_a, \tau, \sigma_a)$ with $|\sigma_a| \geq 1$ and $\tau \geq 2$, there exists a low-strength hTAM system $\Gamma_h = (T_h, \Sigma_h, \tau, \sigma_h)$ that simulates Γ_a at scale 3. Also, $|T_h| = O(|T_a|^2)$ and $|\sigma_h| = 9|\sigma_a|$.*

5 Single-Tile pfbTAM Systems Simulate Low-Strength hTAM Systems

In this section we show that pfbTAM systems with a single tile type can simulate low-strength hTAM systems. Combining this result with Lemma 1 proves that single-tile pfbTAM systems can simulate all aTAM systems. Two-step simulation enables independent resolution of two main difficulties: using rotation as an encoding mechanism and eliminating uncontrolled growth of a single tile type with strength-τ glues.

Lemma 2. *For any low-strength hTAM system $\mathit{\Gamma}_h = (T_h, \Sigma_h, \tau, \sigma_h)$ with $|\sigma_h| = 3$, there is a pfbTAM system $\mathit{\Gamma}_p = (T_p, \Sigma_p, \tau, \sigma_p)$ with $|T_p| = 1$ and $|\sigma_p| = 3$ that simulates $\mathit{\Gamma}_h$ at scale 1.*

Theorem 1. *(Universal Single-Tile Simulation) There exists a polygonal tile t_U such that for any aTAM system $\mathit{\Gamma}_a = (T_a, \Sigma_a, \tau_a, \sigma_a)$ with $|\sigma_a| = 1$ and $\tau_a \geq 2$, there exists a pfbTAM system $\mathit{\Gamma}_p = (\{t_U\}, \Sigma_p, 2, \sigma_p)$ simulating $\mathit{\Gamma}_a$.*

Here U denotes the intrinsically universal tile set [7]. As U is a fixed tile set, t_U is a tile with a constant number of sides. By adapting ideas from [4], we can eliminate the need for a multi-tile seed assembly, making the system *self-seeding*.

Theorem 2. *(Self-Seeding Single-Tile Simulation) For any aTAM system $\mathit{\Gamma}_a = (T_a, \Sigma_a, \tau, \sigma_a)$ with $|\sigma_a| = 1$ and $\tau \geq 2$, there exists a pfbTAM system $\mathit{\Gamma}_p = (T_p, \Sigma_p, \tau, \sigma_p)$ with $|T_p| = 1$ and $|\sigma_p| = 1$ that simulates $\mathit{\Gamma}_a$.*

6 Single-Tile Translation-Only pfbTAM Systems Simulate Cellular Automata

First we prove that single-tile pfbTAM systems where rotation is forbidden, called *translation-only* systems, are capable of arbitrary computation given an appropriately large seed. See Figure 3 for an example.

Theorem 3. *For any double-checkerboarded pyramid aTAM system $\mathit{\Gamma}_a = (T_a, \Sigma_a, 2, \sigma_a)$, there exists a translation-only pfbTAM $\mathit{\Gamma}_p = (\{t_p\}, \Sigma_p, 3, \sigma_2)$ that simulates $\mathit{\Gamma}_a$. Furthermore, t_p has $O(|T_a|^5)$ sides.*

Next, we prove that any single-tile, translation-only pfbTAM system with a seed assembly consisting of fewer than four tiles either produces only the seed assembly, or produces an infinite assembly.

Lemma 3. *Let S be a two-dimensional, bounded, connected, regular closed set S and $\boldsymbol{v}$ be a two-dimensional vector. Define $S + \boldsymbol{v} = \{p + \boldsymbol{v} : p \in S\}$. If $S + \boldsymbol{v} \cap S = \emptyset$, then $S + c \cdot \boldsymbol{v} \cap S = \emptyset$ for any non-zero integer c.*

Theorem 4. *For any translation-only pfbTAM system $\mathit{\Gamma} = (T, \Sigma, \tau, \sigma)$ with $|T| = 1$ and $|\sigma| = 1$, the set of producible assemblies of $\mathit{\Gamma}$ is either $\{\sigma\}$ or contains assemblies of unbounded size.*

Corollary 2. *There are aTAM systems that cannot be simulated by any single-tile, translation-only pfbTAM system.*

Theorem 4 utilizes Lemma 3 and a simple observation about self-seeding systems: to form a two-tile assembly requires a strength-τ attachment between two individual tiles.

Theorem 5. *For any translation-only pfbTAM system $\mathit{\Gamma} = (T, \Sigma, \tau, \sigma)$ with $|T| = 1$ and $|\sigma| = 3$, the set of producible assemblies of $\mathit{\Gamma}$ is either $\{\sigma\}$ or contains assemblies of unbounded size.*

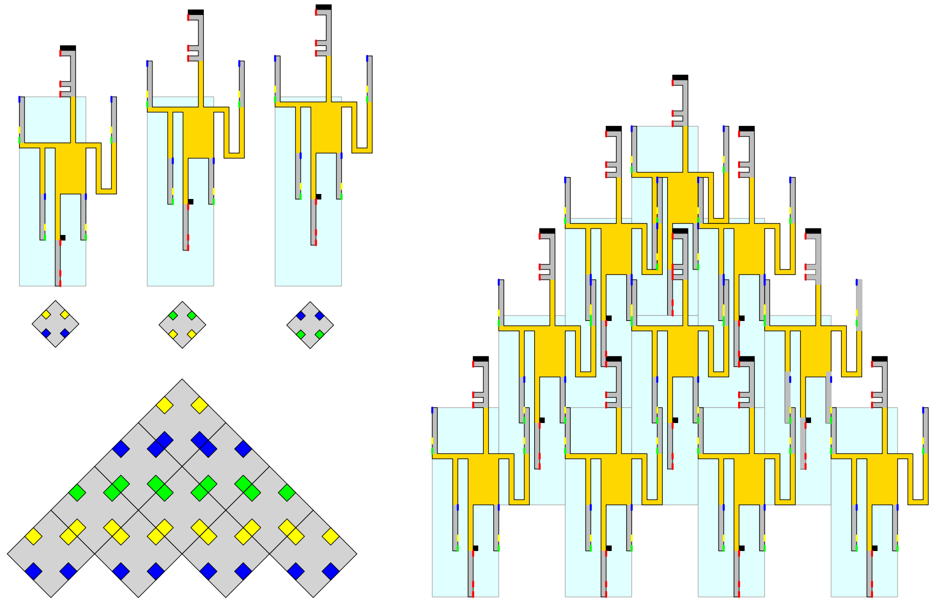

Fig. 3. An assembly of sliders (right) is mapped to a corresponding square tile assembly (left) by placing an imaginary grid (light blue).

A detailed proof is in the full paper. Matters get more involved with arbitrary seeds, and we conjecture polynomially large seeds are needed in general.

Conjecture 1. Let $\Gamma = (T, \Sigma, \tau, \sigma)$ be a translation-only pfbTAM system with $|T| = 1$ and $|\sigma| = n$. If $|\sigma| = n$, then the set of producible assemblies of Γ either contains exclusively assemblies of size $O(n^2)$ or contains assemblies of unbounded size.

Acknowledgements

This work was initiated at the 27th Bellairs Winter Workshop on Computational Geometry held on February 11-17, 2012 in Holetown, Barbados. We thank the other participants of that workshop for a fruitful and collaborative environment. In particular, we thank Brad Ballinger and Anna Lubiw for important discussions regarding Lemma 3. In addition, we thank Jarkko Kari for interesting and fruitful discussions on aperiodic tilings of the plane with few tile types.

References

1. L. Adleman, Q. Cheng, A. Goel, and M.-D. Huang. Running time and program size for self-assembled squares. In *Proceedings of 33rd Annual ACM Symposium on Theory of Computing*, pages 740–748, 2001.

2. R. D. Barish, P. W. Rothemund, and E. Winfree. Two computational primitives for algorithmic self-assembly: Copying and counting. *Nano Letters*, 5(12):2586–2592, 2005.
3. R. D. Barish, R. Schulman, P. W. Rothemund, and E. Winfree. An information-bearing seed for nucleating algorithmic self-assembly. *Proceedings of the National Academy of Sciences*, 106(15):6054–6059, March 2009.
4. S. Cannon, E. D. Demaine, M. L. Demaine, S. Eisenstat, M. J. Patitz, R. T. Schweller, S. M. Summers, and A. Winslow. Two hands are better than one (up to constant factors): Self-assembly in the 2HAM vs. aTAM. In *STACS 2013*, volume 20 of *LIPIcs*, pages 172–184. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2013.
5. H.-L. Chen, R. Schulman, A. Goel, and E. Winfree. Reducing facet nucleation during algorithmic self-assembly. *Nano Letters*, 7(9):2913–2919, September 2007.
6. M. Cook, Y. Fu, and R. Schweller. Temperature 1 self-assembly: Deterministic assembly in 3D and probabilistic assembly in 2D. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 570–589, 2011.
7. D. Doty, J. H. Lutz, M. J. Patitz, R. T. Schweller, S. M. Summers, and D. Woods. The tile assembly model is intrinsically universal. In *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 302–310, 2012.
8. B. Fu, M. J. Patitz, R. Schweller, and R. Sheline. Self-assembly with geometric tiles. In *ICALP 2012: Proceedings of the 39th International Colloquium on Automata, Languages and Programming*, pages 714–725, Warwick, UK, 2012.
9. L. Kari, S. Seki, and Z. Xu. Triangular and hexagonal tile self-assembly systems. In *Proceedings of the 2012 International Conference on Theoretical Computer Science*, pages 357–375, 2012.
10. P.-E. Meunier, M. J. Patitz, S. M. Summers, G. Theyssier, A. Winslow, and D. Woods. Intrinsic universality in tile self-assembly requires cooperation. In *SODA 2014: Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 752–771, Portland, Oregon, 2014. SIAM.
11. P. W. K. Rothemund. Design of DNA origami. In *ICCAD '05: Proceedings of the 2005 IEEE/ACM International Conference on Computer-Aided Design*, pages 471–478, Washington, DC, USA, 2005. IEEE Computer Society.
12. P. W. K. Rothemund and E. Winfree. The program-size complexity of self-assembled squares (extended abstract). In *Proceedings of 32nd Annual ACM Symposium on Theory of Computing*, pages 459–468, 2000.
13. R. Schulman and E. Winfree. Synthesis of crystals with a programmable kinetic barrier to nucleation. *Proceedings of the National Academy of Sciences*, 104(39):15236–15241, 2007.
14. D. Soloveichik and E. Winfree. Complexity of self-assembled shapes. *SIAM Journal on Computing*, 36(6):1544–1569, 2007.
15. E. Winfree. *Algorithmic Self-Assembly of DNA*. PhD thesis, California Institute of Technology, June 1998.
16. E. Winfree, F. Liu, L. A. Wenzler, and N. C. Seeman. Design and self-assembly of two-dimensional DNA crystals. *Nature*, 394(6693):539–544, 1998.
17. S. Woo and P. W. Rothemund. Stacking bonds: Programming molecular recognition based on the geometry of DNA nanostructures. *Nature Chemistry*, 3:620–627, 2011.
18. D. Woods. Intrinsic universality and the computational power of self-assembly. In *MCU: Proceedings of Machines, Computations and Universality*, volume 128, pages 16–22. Electronic Proceedings in Theoretical Computer Science, 2013.