Tile Self-Assembly Simulations
Senior Honors Thesis

Sarah Cannon *
Tufts University

April 9th, 2012

Abstract

Tile self-assembly models describe both mathematically and computationally the ways in which small square tiles can attach to each other to form larger assemblies. Two such models are the abstract tile assembly model, in which all tiles attach to one main assembly containing a seed tile, and the two-handed assembly model, in which there is no seed tile and tiles can attach at any time. This thesis presents results showing that the seed-based assembly process of any tile set that assembles under the conditions of the abstract tile assembly model can be simulated by the assembly process of another tile set in the two-handed model. The simulation is only a constant scale factor larger than the original system and actually requires a lower temperature, a characteristic that describes how easily tiles attach to each other. This result is surprising and interesting, and provides insight into the relative fundamental computational power of both models. This work has direct applications in DNA computing, where biologists are using these models to self-assemble pieces of DNA (represented abstractly by tiles) into structures at the nanoscale level.

*Department of Mathematics, Tufts University, Medford, MA 02155, USA, scanno01@cs.tufts.edu. Research supported in part by National Science Foundation grants CCF-0830734 and CBET-0941538.

1 Acknowledgements

These results also appear in [7]. I would like to thank all of my coauthors, with whom this was joint research: Erik D. Demaine, Martin L. Demaine, Sarah Eisenstat, Matthew J. Patitz, Robert Schweller, Scott W. Summers, and Andrew Winslow. I would especially like to thank Andrew Winslow for all of his contributions to writing and editing Section 4, including creating many of the figures. Finally, I would like to thank Ben Hescott, Diane Souvaine, and Todd Quinto, my thesis advisors, for all of their help in writing and editing this document, as well as their support and encouragement throughout all of my research endeavors.

2 Introduction

In biology and engineering on the nanoscale level, scientists are using small pieces of DNA to assemble nanostructures via the self-assembly process. That is, they create several DNA pieces with "sticky ends" that can attach to each other; then, some set of these DNA pieces are mixed together in a solution and the pieces self-assemble to form a larger assembly by attaching at matching sticky ends. Such an assembly can then be used to organize and assemble additional molecules, such as metals or proteins. In addition, the same DNA self-assembly process has been used to create DNA computers, which can perform a variety of calculations on a much smaller scale and more efficiently than traditional computers. These practical explorations suggest the necessity of developing a theoretical framework for DNA self-assembly. To accomplish this goal, the DNA self-assembly process has been abstracted using square tiles to represent DNA pieces and glues to represent "sticky ends." Several models of the tile self-assembly process have developed, including seeded models and two-handed models. The relationship between these different models is not yet completely understood, but this paper seeks to clarify it. I present a simulation result that shows any assembly process that occurs in a seeded model can be mimicked by an assembly process in the two-handed model. This provides fundamental insights into the two models, indicating that the two-handed model is at least as powerful as the seeded model and suggesting that biological exploration of DNA self-assembly should perhaps focus more on the exploration of two-handed systems.

2.1 Motivations

An ability to create precise structures at the nanoscale level is a current goal of many engineers, including those within electronics and biology. Some high density nanoelectronic devices require precise spacing of components, down to the nanometer scale; biologists may wish to control the spacing between and the density of certain kinds of proteins. Because there are a variety of nanoparticles that can be chemically attached to DNA, assembled DNA nanostructures can provide a precise scaffold for achieving nanometer precision in such endeavors [16].

Nanoscale structures can be created with extreme precision via DNA self-assembly. In this process, small exact pieces of DNA of some standard shape are created. These pieces all have "sticky ends," exposed single helices of DNA with certain base sequences. When allowed to mix together, single helices with complementary base sequences can bond together to form a double helix. Through this process, large DNA assemblies can be created. When each DNA piece is designed so that a certain nanoparticle can chemically attach to it, this large DNA assembly serves as a scaffolding to assemble nanoparticles into particular shapes. Among other work, Yan et al. used DNA aptamer binding to link proteins to a periodic self-assembled DNA array [18], and Kiehl et al. used a DNA lattice to organize particles of gold into parallel lines spaced only about 63 nm apart [15].

In addition to its practical uses in creating nanoassemblies, DNA self-assembly has also been used as an alternate method of computation. For example, Braich et al. designed and implemented a DNA self-assembly system which solves a twenty-variable instance of the three-satisfiability problem [6]. The computation was performed by exhaustively trying all possible variable truth assignments, over a million possibilities, a task that would be infeasible on a traditional computer. However, because of the high degree of parallelism achievable in DNA self-assembly, this computation was possible.

A systematic study of DNA self-assembly is clearly merited in order to develop a theoretical framework to guide practical exploration. In order to understand the abilities and limitations of such processes, multiple mathematical models have been developed and explored.

2.2 Tile self-assembly

The first comprehensive algorithmic study of the DNA self-assembly process was Erik Winfree's PhD thesis in 1998 [24]. Winfree abstracted several processes and techniques of DNA self-assembly, and used the properties of these abstracted computational and mathematical models to make conclusions about and suggestions for the DNA self-assembly process. While Winfree considered several different DNA pieces, a main result is his Tile Assembly Model, formalized in [21] as the abstract tile assembly model (aTAM). Winfree proposed the use of square tiles to represent a certain piece of DNA, the double-crossover molecule, which is especially useful in the self-assembly process because of its rigidity and cooperative binding capabilities. Each square tile represents a double-crossover molecule with four sticky ends, where each sticky end is represented by a glue of a certain type on one side of the tile. A tile sticks to a tile assembly precisely when the glues on their abutting edges match and are of sufficient strength for the attachment to occur. Of note, more recently researchers have created self-assembling plus shapes made from two protein nanorods attached at their midpoints, and the self-assembly process of these particles can also be modeled by tiles with glues [11].

Winfree's tile assembly model is similar to the mathematical problem of Wang tiling [23]. A Wang tile is a unit square with a color on each edge, and tiles must be placed so that abutting edges have the same colors. Mathematical explorations of Wang tiles generally consider whether, given a finite set of Wang tiles, it is possible to tile the plane with copies of the tiles in the set. While tile self-assembly is similar in that tiles bond when the glues on their abutting edges match, researchers consider the tile assembly process and the way in which large tile assemblies are formed instead of just the pattern of tiles produced. In Wang tiling, the only concern is the position of each tile in a final assembly, and the process by which this tile assembly was constructed is generally unimportant.

2.3 aTAM and 2HAM

In order to simplify his Tile Assembly Model, Winfree made several assumptions, including the restriction that all attachments consist of a single tile joining a larger assembly. This means it is not possible for two tile assemblies both containing two or more tiles to attach to each other, even if the glues between them are of sufficient strength. That is, any aTAM system has a single *seed tile*, and all attachments consist of single tiles joining the tile assembly containing the seed tile. This convention persisted through many later investigations within the field of tile self-assembly; for example, consider [3], [4], [5], [10], [14], [21], and [22], among others.

Winfree gives little justification for this assumption; while it makes his tile assembly model mathematically simpler and easier to analyze, even he notes that it may not be entirely accurate. There are still doubts within the tile self-assembly community as to whether such an assumption is valid. Consequently, another tile assembly model has developed. In the two-handed tile assembly model (2HAM) [11], also called the multiple tile model [9] or hierarchical self-assembly [8], there is no seed tile. Instead, two tile assemblies with matching glues between them of sufficient strength can attach at any time, regardless of the size of each assembly.

Discussion of a tile self-assembly model in which there is not a seed first appeared in [2], while the 2HAM as it is used in current research was first defined in [9], where the authors studied the minimum number of distinct tiles required to self-assemble squares and rectangles. Since then, many authors have used the capabilities of the 2HAM to develop models and constructions that would otherwise likely not have been possible, including staged self-assembly, in which certain tile assemblies are created in separate bins and then mixed together [11]; self-assembly systems in which certain tiles can be destroyed at some step by an RNAse enzyme [1, 12]; and self-assembly systems in which can identify a certain shape [20]. While tile systems in the 2HAM have often led to efficient production of shapes and to solutions to otherwise difficult problems, it is also much harder to design well-behaved 2HAM systems, since attachment can occur between tiles or tile assemblies at any time.

While the performance of both models on particular problems has been analyzed and there are certain constructions developed in the 2HAM that do not have a clear analogue in the aTAM, it has been remarkably difficult to prove a distinct separation between the models; [7] was the first comprehensive exploration into such a question. A main result of that paper, and the result presented in this thesis, is a simulation result proving that anything that can be done in the aTAM can also be done in the 2HAM.

2.4 Simulations

From the beginning of the field of tile self-assembly, one area of inquiry has been into simulation results. That is, researchers have considered how to use tile sets to simulate things such as Turing machines, blocked cellular automata [24], and other tile sets [11, 13, 19, 25]. To simulate a Turing machine or cellular automata means to perform the same computation, while for tile sets, simulate means to follow the same assembly process. Such results provide insights into the computational power of tile self-assembly models. For example, Winfree originally showed that the aTAM is Turing-universal [24], meaning it can perform any computation that a Turing machine can. He does so by showing that aTAM can simulate a blocked cellular automata, which was already well-known to be Turing universal by a proof analogous to that in [17]. This same reduction can be easily modified to show that 2HAM is also Turing-universal.

Simulation results between tile sets have provided insights into the relationships between different tile self-assembly models. Often, if a tile system $\mathcal{T}$ is simulated by another tile system $\mathcal{T}'$, then each tile in $\mathcal{T}$ is actually represented by a $k \times k$ block of tiles in $\mathcal{T}'$. Demaine et al. showed that a certain subclass of tile assembly systems could be simulated by a tile set with a constant number of glues and tile types, with $k = \log |T|$ and requiring $\log(\log |T|)$ stages, where in each stage tile assemblies created in two or more separate bins are combined together and allowed to assemble [11]. This demonstrates that the staged tile assembly model is capable of performing some

aTAM systems	Simulating 2HAM systems
$\tau = 1$	$\tau = 2$, $k = 5$
$\tau = 2$	$\tau = 2$, $k = 5$
$\tau = 2$	$\tau = 3$, $k = 5$
$\tau \geq 4$	$\tau = 4$, $k = 5$

Table 1: Summary of simulation results, where τ = temperature and k = scale factor.

of the same assembly processes as a traditional one-bin model, providing insight into how these two tile assembly models are related. More recently, Doty et al. showed that there exists a single aTAM system at temperature 2 that can simulate any aTAM system at any temperature [13], where the temperature of a tile assembly system is the total glue strength needed for attachment to occur. Such a result means that the aTAM is intrinsically universal, and shows that raising the temperature of any aTAM system above 2 does not allow the production of any additional shapes, up to scale factor k. However, in this simulation $k = O(g^4 \log(g))$, where g is the number of unique glues in the original tile assembly system, an extremely large and impractical scale factor. Regardless, simulation results such as these can provide insights into relationships and similarities between models, helping to focus the direction of study on those models that are the most powerful.

2.5 Results

In this paper, I present a construction via which any aTAM system can be simulated by a 2HAM system. The *temperature* of a tile assembly system is the strength needed for attachment between tiles to occur, and lower temperature systems are much simpler, easier both to implement and to study, and are generally preferred. See Table 1 for a summary of simulation results for varying temperatures. In general, these results show that anything that can be done in the aTAM can also be done in the 2HAM with only a small constant scale factor of five. Such conclusions suggest that the aTAM has no more computational power than the 2HAM (up to small constants). While scientists are still in the process of developing the tools needed to create large error-free self-assembled nanostructures successfully, this result suggests that perhaps a focus should be placed on developing the tools to implement two-handed systems rather than seeded systems.

Perhaps the most surprising and interesting result is listed in the bottom row of Table 1, indicating that an aTAM tile assembly system at an arbitrarily high temperature can be simulated by a 2HAM system with just temperature 4 and a small constant scale factor of 5. This result suggests that above temperature 4, no additional computational power is gained by any aTAM system, to within a much smaller scale factor that the similar result by Doty et al [13]. This is believed to be false for the 2HAM.

3 Background

I will now define the abstract tile assembly model and the two-handed tile assembly model formally, beginning with the fundamental components of a tile assembly system: tiles, glues, and temperature.

3.1 Tiles, Glues, and Temperature

A *tile* t is an axis parallel unit square with corners at grid points of the integer grid $\mathbb{Z} \times \mathbb{Z}$ and a glue on each of its four sides. A *glue* is a pair $g = (a, k)$, where a is the *glue type*, an element of some label set Σ, and k is the *glue strength*, $k \in \mathbb{Z}_{\geq 0}$. Two glues of the same type a will always have the same strength k, and two glues are equal if they are of the same type.

Consider a tile in $\mathbb{Z} \times \mathbb{Z}$ with corners at (x, y), $(x+1, y)$, $(x+1, y+1)$, and $(x, y+1)$, generally given in that order. Define the *south face* of the tile to be the side of the tile between (x, y) and $(x + 1, y)$; call the glue on this face the *south glue* g_S. Define the *east face* to be the side of the tile between $(x + 1, y)$ and $(x + 1, y + 1)$; call the glue on this face the *east glue* g_E. Similarly, the side between $(x + 1, y + 1)$ and $(x, y + 1)$ is the *north face* and its glue is the *north glue*, while the side between $(x, y + 1)$ and (x, y) is the *west face* and its glue is the *west glue* of the tile. A *tile type* is specified by four glues (g_S, g_E, g_N, g_W), listed in south-east-north-west order; there may be (and often are) an infinite number of tiles of a certain tile type in a given tile assembly system. A tile's location is specified by the location of its bottom left corner in $\mathbb{Z} \times \mathbb{Z}$. Both tile assembly models discussed assume that these tiles can move horizontally and vertically to other locations on the integer grid, but do not rotate. In fact, tiles move randomly throughout $\mathbb{Z} \times \mathbb{Z}$ according to Brownian motion.

The *temperature* $\tau \in \mathbb{N}$ of a tile system is the minimum total glue strength needed for attachment between tiles to occur. Informally, two tiles *match* if their locations are one unit apart on the integer grid and the glues on the interface between the two tiles are the same; these tiles *attach* or *bond* if this glue is of strength greater that or equal to the temperature τ. That is, tile $t = (g_S, g_E, g_N, g_W)$ at position (x, y) matches tile $t' = (g'_S, g'_E, g'_N, g'_W)$ at position $(x + 1, y)$ precisely if $g_E = g'_W = (a, k)$, for some glue a of strength k; these tiles attach if $k \geq \tau$. If t' is at position $(x, y + 1)$, the two tiles attach precisely when $g_N = g'_S = (a, k)$ and $k \geq \tau$.

When tiles attach, they form connected *tile assemblies*. Tile assemblies are rigid; if one tile in the assembly translates one unit in a given direction, then the entire assembly translates in that direction. The *size* of a tile assembly is the number of tiles in the assembly. Trivially, single tiles are tile assemblies of size one. Tile assemblies can also attach to each other. Two tile assemblies A and A' *overlap* if some $t \in A$ and some $t' \in A'$ are both at location $(x, y) \in \mathbb{Z} \times \mathbb{Z}$; assemblies A and A' *attach* if they do not overlap and the sum of the strengths of all matches between tiles in A and a tiles in A' is greater than or equal to τ. For an example of tile assembly attachment, see Figure 1.

3.2 Abstract Tile Assembly Model

A *tile assembly system* $\mathcal{T} = (T, s, \tau)$ in the abstract tile assembly model (aTAM) consists of a finite set of tile types T, a single special *seed tile* s, and a temperature τ. Attachment in this model

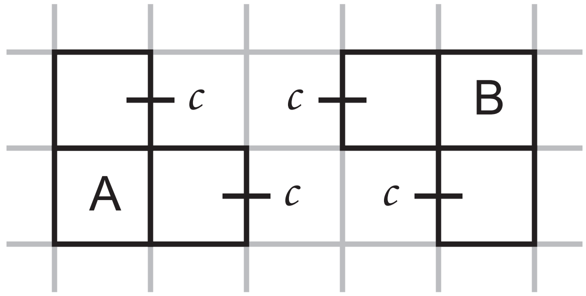

Figure 1: Let the temperature $\tau = 2$, and let all glues c be of strength one. If tile assembly B translates two units to the left, then tile assemblies A and B will attach.

is defined as above, with the additional restriction that any tile assembly of size greater than one must contain the seed tile; such an assembly is called the *seed assembly*, and it is unique in any aTAM system. Because of this restriction, all attachments consist of a single tile joining the seed assembly.

The assembly process begins by fixing the single seed tile at location $(0, 0)$, while infinitely many copies of all other tile types are randomly distributed over $\mathbb{Z} \times \mathbb{Z}$. Over time, the seed tile remains fixed at $(0, 0)$, while other tiles move arbitrarily horizontally and vertically (irrotational Brownian motion restricted to the integer grid). When a tile is in a location adjacent to the seed assembly such that attachment is possible, that attachment occurs. Note that if there is a location on the boundary of the seed assembly to which two different tiles could possibly attach, there is no way of determining which one would do so, because tile motion is random and unpredictable. However, we assume one of the two possible tiles will eventually attach.

3.3 Two-handed Tile Assembly Model

A *tile assembly system* $\mathcal{T} = (T, \tau)$ in the two-handed tile assembly model (2HAM) consists of a finite set of tile types T and a temperature τ. Attachment is defined as above, without the additional restriction of a seed tile. The assembly process simply begins by placing infinitely many copies of all tiles types at arbitrary locations in $\mathbb{Z} \times \mathbb{Z}$ and allowing them to randomly translate horizontally or vertically while remaining on the integer grid (Brownian motion). When two tile assemblies are positioned such that attachment is possible, they attach. Again, note that if there are multiple attachments possible for a given tile assembly, there is in general no way of determining which one(s) will occur.

3.4 Termination, Unique Assembly, Planarity

The assembly process of a tile assembly system under one of the two models defined above *terminates* when there are no possible attachments that can be made for any positions of the tiles and tiles assemblies. A tile system is *terminating* if its assembly process always terminates. Any tile assemblies to which no other tiles or tile assemblies can attach are *terminal* or *final* assemblies.

A tile assembly system is *unique* if it always produces final assemblies of the same shapes, though there may be different tiles at certain locations in the assembly. A tile assembly system is *strongly unique* if it always produces the same final assembly or assemblies, with the same tiles at the same locations in each assembly.

Two tile assemblies overlap if they each have a tile at the same location (x, y) in the integer grid. In either tile assembly system described above, tile assemblies are permitted to overlap, but no attachments are possible between overlapping tiles or tile assemblies. This ensures that any tile assemblies remain planar, without overlaps. Permitting overlaps allows tile systems to fill in holes in tile assemblies, for example. This is perhaps a more realistic model of DNA tile attachment, since in practice these planar tiles exist in three-dimensional space. One can consider restricting tile motion to disallow overlaps, however; in a *planar* tile assembly system, two tiles never overlap.

3.5 Simulations

Let $\mathcal{T} = (T, s, \tau)$ be a tile assembly system in the aTAM model, and let $\mathcal{T}' = (T', \tau')$ be a tile assembly system in the 2HAM model. Informally, $\mathcal{T}'$ simulates $\mathcal{T}$ if each tile $t \in T$ is represented in $\mathcal{T}'$ by a set of $k \times k$ tiles that, as a whole, act like tile t and build the same assemblies as in $\mathcal{T}$.

For $k \in \mathbb{Z}^+$, a *k-block supertile* over tile set T' is a tile assembly contained within a $k \times k$ square (a *k-block*). Let $B_k^{T'}$ denote the set of all k-block supertiles over T'. A partial function $R : B_k^{T'} \to T$ is a *k-block supertile replacement function* from tile set T' to tile set T if for any $\alpha, \beta \in B_k^{T'}$ such that $\alpha \subseteq \beta$ and $\alpha \in dom(R)$, then $\beta \in dom(R)$ and $R(\alpha) = R(\beta)$. That is, R maps certain supertiles in T' to tiles in T; if R sends part or all of a k-block supertile to a given tile $t \in T$, then R cannot map any subassemblies of this supertile to a different tile in T.

Let the set of all assemblies produced by a given tile set T be denoted by $\mathcal{A}[T]$. For $x, y \in \mathbb{Z}$, $A(x, y)$ is the tile in assembly A at location (x, y) in $\mathbb{Z} \times \mathbb{Z}$, if such a tile exists; otherwise, $A(x, y) = \varnothing$ if there is no tile in assembly A at location (x, y). For assembly A' in tile assembly system $\mathcal{T}'$, define the supertile $A'_k(x, y)$ to be the assembly contained in the k-block that has lower left corner (kx, ky).

Given a valid k-block supertile replacement function R from $\mathcal{T}'$ to $\mathcal{T}$, one can also define a *k-block assembly replacement function* $R^* : \mathcal{A}[T'] \to \mathcal{A}[T]$. R^* is a k-block assembly replacement function if for all assemblies $A' \in \mathcal{A}[T']$, $R^*(A') = A$ if and only if it is possible to position A and A' in $\mathbb{Z} \times \mathbb{Z}$ such that $A(x, y) = R(A'_k(x, y))$ for all tile positions (x, y) in A.

For the purposes of this paper, a k-block supertile in $\mathcal{T}'$ is *empty*, $A'_k(x, y) = \varnothing$, if the only tiles present within the k-block are adjacent to the boundary of the k-block; see Figure 2. Note this is not a strict definition of an empty supertile, which lends a *fuzziness* of one tile to our simulation definition; see section 4.1 for a discussion of fuzziness in simulations and to see why such a definition of empty is necessary for the purposes of this result. We say R^* maps A' to A *cleanly* if for all $x, y \in \mathbb{Z}$, $A'_k(x, y) = \varnothing$ precisely when $R(A'_k(x, y)) = A(x, y) = \varnothing$.

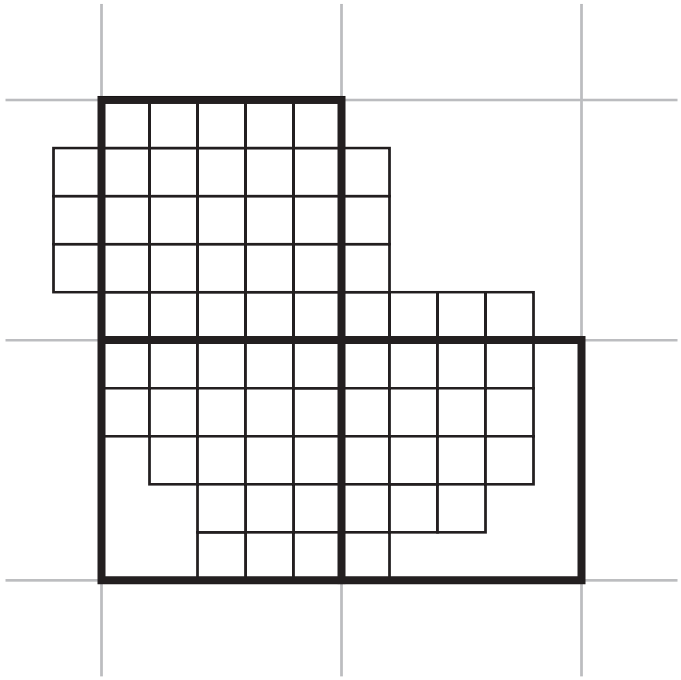

Figure 2: A tile assembly in $\mathcal{T}'$. The three bold boxes represent non-empty 5-blocks, while all other adjacent 5-blocks are empty. The gray lines are the grid $5\mathbb{Z} \times 5\mathbb{Z}$.

Tile assembly system $\mathcal{T}'$ in the 2HAM simulates tile assembly system $\mathcal{T}$ in the aTAM at scale $k \in \mathbb{Z}^+$ if there exists a k-block replacement function $R : B_k^{T'} \to T$ such that the following hold:

1. Equivalent Production
   (a) $\{R^*(A') | A' \in \mathcal{A}[T']\} = \mathcal{A}[T]$
   (b) For all $A' \in \mathcal{A}[T']$, A' maps cleanly to $R^*(A')$.
2. Equivalent Dynamics
   (a) If assembly A produces assembly X via a sequence of tile attachments in $\mathcal{T}$ for some $A, X \in \mathcal{A}[T]$, then for all $A' \in \mathcal{A}[T']$ such that $R^*(A') = A$, A' can produce some assembly $X' \in \mathcal{A}[T']$ with $R^*(X') = X$.
   (b) If assembly A' produces assembly X' via a sequence of attachments in $\mathcal{T}'$ for some $A', X' \in \mathcal{A}[T']$, then $R^*(X') \in \mathcal{A}[T]$ can be produced by some sequence of attachments in $\mathcal{T}$ from $R^*(A') \in \mathcal{A}[T]$.

The equivalent production condition ensures that the same assemblies form in $\mathcal{T}$ and $\mathcal{T}'$ up to a scale factor of k, while equivalent dynamics ensures that the same assembly process is followed in both models.

Lemma 3.1. *If tile assembly system $\mathcal{T}'$ simulates tile assembly system $\mathcal{T}$, then $\mathcal{T}$ is terminating if and only if $\mathcal{T}'$ is terminating.*

Proof. Whether a system is terminating depends only on whether all produced assemblies are final assemblies. By equivalent production, assembly A' forms in $\mathcal{T}'$ precisely when assembly $R^*(A')$ forms in $\mathcal{T}$. So, assembly A' is final in $\mathcal{T}'$ precisely when assembly A is final in $\mathcal{T}$. Since this is true for all assemblies A' and A, $\mathcal{T}'$ is terminating precisely when $\mathcal{T}$ is terminating. $\square$

No conclusions can be made about the uniqueness of tile system $\mathcal{T}'$ based on the uniqueness of tile system $\mathcal{T}$ without additional information about the specifics of the simulation.

4 Simulating aTAM at $\tau \geq 4$ with 2HAM at $\tau = 4$

It is possible to simulate aTAM at temperature $\tau \geq 4$ using 2HAM at temperature 4 with a constant scale factor of 5. Let $\mathcal{T} = (T, s, \tau)$ be an aTAM system at arbitrary temperature, and let $\mathcal{T}' = (T', 4)$ be the 2HAM system that simulates it. I will now describe how to construct $\mathcal{T}'$ from $\mathcal{T}$ such that the two systems have equivalent production and dynamics.

Overview Given any aTAM system $\mathcal{T}$, each tile t in the aTAM system is represented by 25 tiles forming a 5×5 supertile in $\mathcal{T}'$; that is, we have a 5-block supertile replacement function from $\mathcal{T}'$ to $\mathcal{T}$. A supertile in $\mathcal{T}'$ consists of a 3×3 center *brick* assembly, surrounded on all sides by a *mortar* one tile thick. These tiles are designed such that bricks and certain mortar pieces can assemble independently (utilizing the seedless nature of $\mathcal{T}'$), but bricks cannot attach to mortar pieces or other bricks unless certain other conditions are met.

We mimic the seeded nature of $\mathcal{T}'$ by strengthening the bonds between brick and mortar for a single supertile type, called the seed brick, which corresponds to the seed tile in $\mathcal{T}$. This enables the mortar to complete around the seed brick. Once a complete (or nearly complete) mortar has attached around a brick, adjacent mortar pieces corresponding to other bricks can begin to attach. This then allows new bricks to attach, followed by the completion of the mortar around these new brick; the assembly process continues in this manner. Because this process can only begin with the attachment of the mortar to a seed brick, we ensure that bricks can only attach to partially built assemblies containing a seed brick, mimicking the seeded nature of an aTAM system.

Additionally, we divide instances of glues into *inward* and *outward* glue sets, such that an outward glue g can only attach to an inward glue of the same type. Throughout the assembly process, the invariant that all exposed glues in any assembly containing a seed brick are outward glues is maintained; this prevents partially built seeded assemblies from attaching to each other. An example of the construction in which 3×3 bricks, 3×1 mortar rectangles, and individual mortar tiles attach to form 5×5 supertiles can be seen in Figure 3.

Inward and Outward Glues In order to prevent unwanted attachment, every instance of a glue g in $\mathcal{T}'$ is assigned one of two labels, "inward" or "outward." Inward and outward glues appear as arrows pointing inward or outward from a tile in the figures throughout this section. We enforce that all glues in $\mathcal{T}'$ only attach in complementary inward-outward pairs; for example, an outward west glue will attach to an inward east glue of the same type but not to an outward east glue. This can be easily implemented for each glue g in $\mathcal{T}'$ using four glues corresponding to the four directions each glue arrow may "point." For instance, an outward glue a on the west side of a tile is a west-pointing a_W glue, while an inward glue a on the north side of a tile is a south-pointing a_S glue. The following lemma shows that this correctly implements inward-outward glue pairs.

Lemma 4.1. *Replacing each instance of a glue a in some tile assembly system $\mathcal{S}$ with one of the four direction-based glues a_S, a_N, a_E, a_W corresponding to the inward/outward pointing direction assigned to the instance results in exactly inward-outward glue pair attachment.*

Proof. Let $\mathcal{S}$ be a tile system in the 2HAM with inward/outward glues, and let $\mathcal{S}'$ be the corresponding tile system, also in the 2HAM, where inward/outward glues are implemented by four

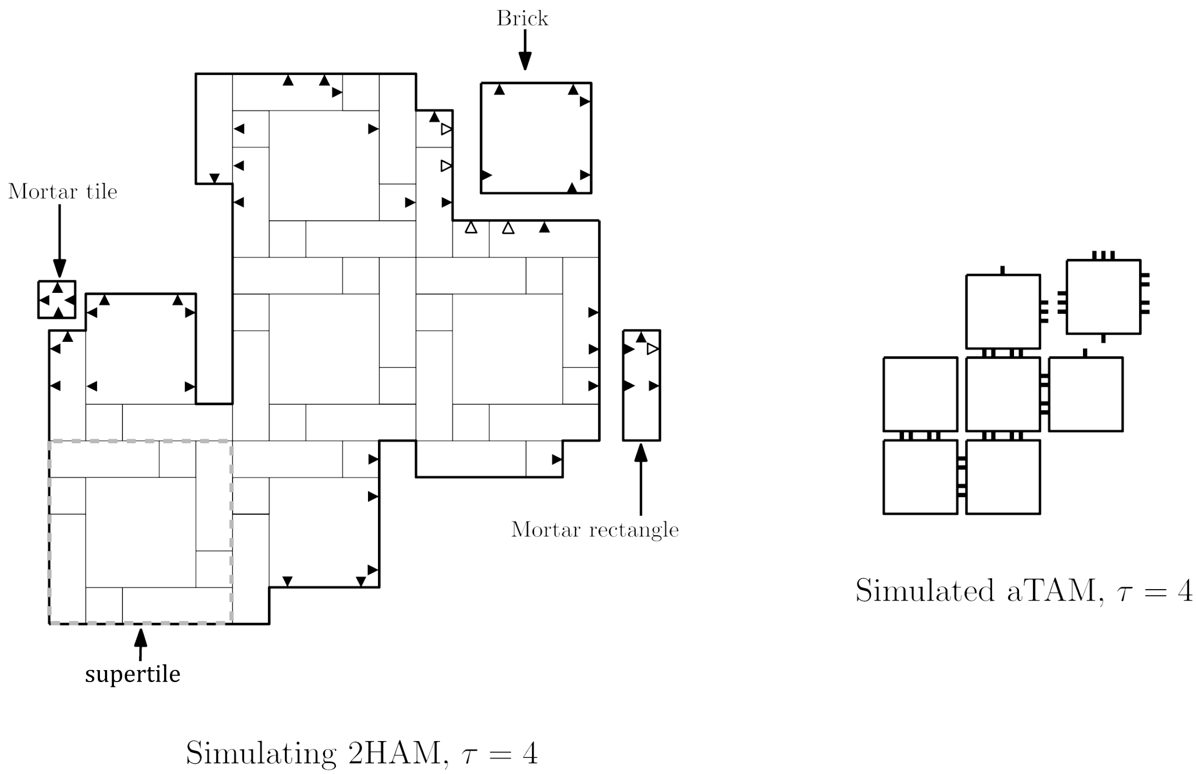

Figure 3: The simulation of an assembly in an aTAM system simulated using a 2HAM system. The filled and unfilled arrows represent glues of strength 2 and 1 respectively in the 2HAM system, while the dashes each represent a bond of strength 1 in the aTAM system (i.e. 4 dashes on the north side of a tile is a glue of strength 4).

direction-based glues. Any inward-outward glue pair in $\mathcal{S}$ (i.e. an outward east glue and an inward west glue) point in the same direction (i.e. east) and so are the same glue (i.e. a_E), meaning the direction-based glue system $\mathcal{S}'$ bonds whenever the inward/outward tile system $\mathcal{S}$ was able to bond. Any pair of glues not on opposite sides of tiles cannot bond by geometry because tiles do not rotate; for example, an east glue of one tile can only bond to the west glue of another tile. A pair of glues on opposite sides of tiles whose glues are either both inward or both outward in $\mathcal{S}$ point in opposite directions and thus are different glues in $\mathcal{S}$ and also cannot bond. This implies the direction-based system $\mathcal{S}'$ bonds precisely when $\mathcal{S}$ does. □

Since inward-outward glue pairs can be easily implemented using four direction-based glues, we assume for the remainder of this section that all glues in 2HAM system $\mathcal{T}'$ are assigned a direction, inward or outward. Intuitively, in 2HAM system $\mathcal{T}'$ simulating aTAM system $\mathcal{T}$, each piece (brick or mortar) attaches to a partially completed assembly at its own inward glues, leaving only exposed outward glues to which more pieces can attach.

Bricks For each tile $t \in T$, we can simulate t in $\mathcal{T}'$ by a set of 3×3 *brick* assemblies, one for each minimal set of glues of t of total strength greater than or equal to the temperature τ of $\mathcal{T}$. Note all glues between tiles within a brick are unique across $\mathcal{T}'$. Figure 4 depicts the gluing pattern

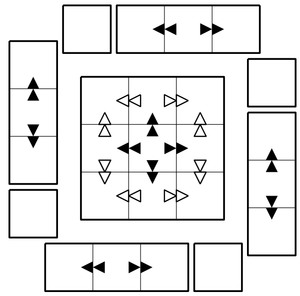

Figure 4: The internal gluing pattern for bricks and mortars. Dark arrows represent glues of strength 4, and light arrows represent glues of strength 2.

for the interior of any brick, which clearly implies the following two lemmas. Define a brick as *partially assembled* if it consists of at least two attached tiles.

Lemma 4.2. *If a brick B in $\mathcal{T}'$ is partially assembled, then the center tile of B is present in this partial brick assembly.*

Proof. See Figure 4. Between any two non-center tiles within B, there is at most one glue of strength at most two, which is insufficient for attachment. Because of this, B cannot contain only non-center tiles, because two non-center tiles cannot attach to each other. So, B must contain its center tile. □

Lemma 4.3. *For brick B that is partially assembled in $\mathcal{T}'$, all exposed glues internal to B are outward glues.*

Proof. See Figure 4. The glues on B's center tile are all outward glues. Any tile in the middle of one side of B must use its strength four inward glue to attach to the center tile, leaving two exposed strength two glues, both of which are outward glues. Any corner tile must use both of its inward strength two glues to attach to a partially completed brick, leaving no exposed inward glues internal to B. Since the assembly process of any partially assembled brick starts with the center tile, then any partially assembled version of B does not contain any exposed inward glues internal to B. □

We will see that any brick attaches to the rest of the assembly at two different tiles. Lemma 4.2 then will be used to ensure the uniqueness of any supertile; once the center tile of the brick is present, it completely determines the identity of the supertile. Lemma 4.3 will later be used to prove the invariant that all exposed glues on any partially completed assembly are outward glues.

Given any tile $t \in T$, consider every subset S of glues on t with total strength greater than τ such that the removal of any glue from S yields a total glue strength less than τ. For each such *minimal glue set* S, a brick B_S in $\mathcal{T}'$ is created such that all glues on the sides of B_S corresponding to glues in S are inward glues while all glues on other sides are outward glues. Note that if glue set S is contained in glue set S', then S and S' cannot both be minimal glue sets. Because of this, there will never be more than six distinct minimal glue sets for a given tile. If there are more than six minimal glue sets, then one must be contained within another, a contradiction. A tile where

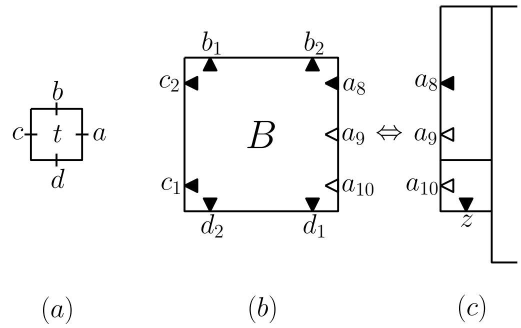

Figure 5: Filled arrows represent glues of strength 2, unfilled arrows represent glues of strength 1. (a) an aTAM tile $t \in T$ with minimal glue set $S = \{a\}$; (b) the brick B in $\mathcal{T}'$ generated by S; (c) a location where B could attach to a partially built assembly.

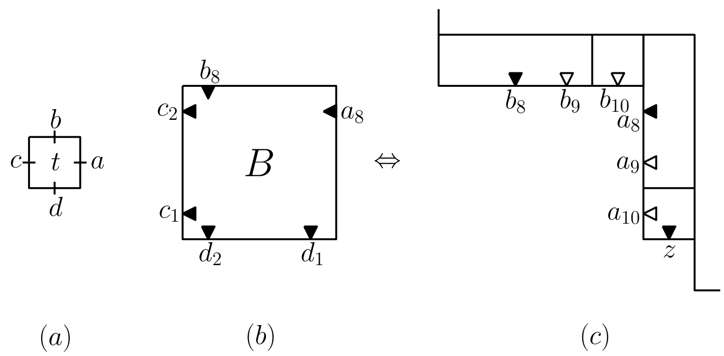

Figure 6: (a) an aTAM tile $t \in T$ with minimal glue set $S = \{a, b\}$; (b) the brick B in $\mathcal{T}'$ generated by S; (c) a location where B could attach to a partially built assembly; note not all possible attachment points are used.

all four sides have glues of strength $\lceil \tau/2 \rceil$ yields exactly six minimal glue sets, all of which have exactly two elements.

The specific types and strengths of external glues on bricks in $\mathcal{T}'$ are constructed as follows. Let $t \in T$. For each glue p on one side of t in a minimal glue set S corresponding to B_S, there are inward glues p_8, p_9, and p_{10} in clockwise order on the corresponding side of B_S. All other glues q on t yield outward glues q_1 and q_2 on the two corner tiles of the corresponding side of B_S, both with strength 2, with q_2 clockwise from q_1. For a minimal glue set $\{a\}$ of size 1, the glue a_8 has strength 2, while a_9 and a_{10} have strength 1 (see Figure 5). For a minimal glue set $\{a, b\}$ of size 2, glues a_8 and b_8 have strength 2, while glues a_9, a_{10}, b_9, and b_{10} have strength 0, i.e. do not exist (see Figure 6). For a minimal glue set $\{a, b, c\}$ of size 3, glues a_9, b_9, c_9, and c_{10} have strength 1, while a_8, a_{10}, b_8, b_{10}, and c_8 have strength 0 (see Figure 7). For a minimal glue set $\{a, b, c, d\}$ of size 4, glues a_9, b_9, c_9, and d_9 have strength 1, while a_8, a_{10}, b_8, b_{10}, c_8, c_{10}, d_8, and d_{10} have strength 0 (see Figure 8). See Figure 9 for an example of a tile $t \in T$ and the two bricks it generates in $\mathcal{T}$ based upon its two minimal glue sets. In these and all subsequent figures in this section, filled arrows represent glues of strength 2, unfilled arrows represent glues of strength 1, while glues of strength 0 are not shown.

Note that the set of inward glues on any brick is minimal; that is, if a brick attaches to an assembly at its own inward glues, then all exposed glues are outward glues.

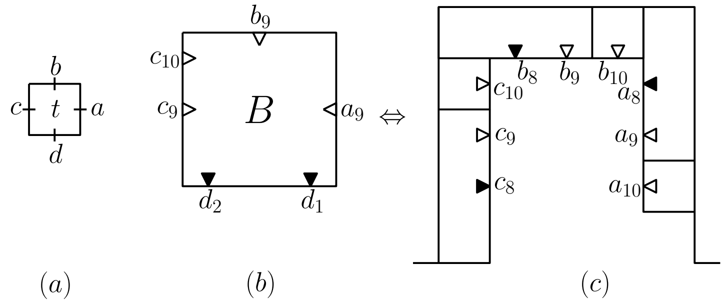

Figure 7: (a) an aTAM tile $t \in T$ with minimal glue set $S = \{a, b, c\}$; (b) the brick B in $\mathcal{T}'$ generated by S; (c) a location where B could attach to a partially built assembly; note not all possible attachment points are used.

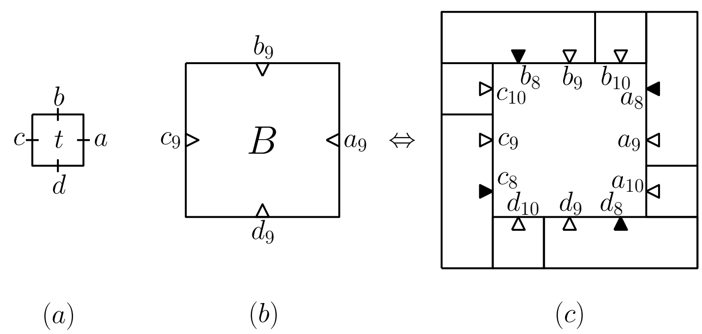

Figure 8: (a) an aTAM tile $t \in T$ with minimal glue set $S = \{a, b, c, d\}$; (b) the brick B in $\mathcal{T}'$ generated by S; (c) a location where B could attach to a partially built assembly; note not all possible attachment points are used.

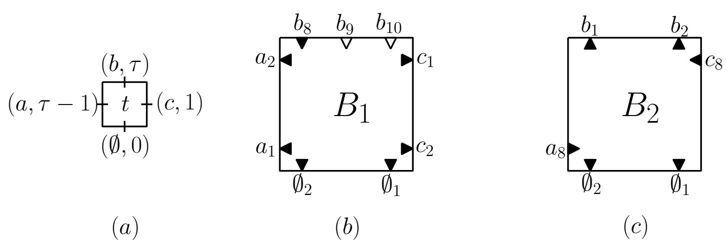

Figure 9: Ordered pairs denote a glue and its strength. (a) A tile $t \in T$, with minimal glue sets $S_1 = \{(b, \tau)\}$ and $S_2 = \{(a, \tau - 1), (c, 1)\}$ (b) Brick B_1 in $\mathcal{T}'$ representing tile t, generated by the minimal glue set S_1 (c) Brick B_2 in $\mathcal{T}'$ representing tile t, generated by the minimal glue set S_2

Mortar Pieces In any 5×5 supertile assembly in $\mathcal{T}'$ representing an aTAM tile $t \in T$, the brick is surrounded by a *mortar* one tile thick. This mortar consists of both single-tile *mortar tile* assemblies and 3×1 and 1×3 *mortar rectangle* assemblies with internal glues of strength 4. However, *mortar pieces* - both tiles and rectangles - cannot attach to each other or to bricks unless other tiles are present. See Figure 4 for the general structure of the mortar assemblies around any brick. Any mortar rectangle must attach to an assembly at exactly two glues, and the following lemma will later be used to prove that even if a partially completed rectangle attaches to an assembly, all exposed glues are outward glues. A partially assembled mortar rectangle is defined as a mortar rectangle that consists of at least two attached tiles

Lemma 4.4. *If a mortar rectangle is partially completed, then all exposed glues internal to the rectangle are outward glues.*

Proof. See Figure 4. A completed mortar rectangle has no exposed internal glues. A partially but not fully completed mortar rectangle consists of two tiles, one of which must be the center tile of the rectangle. The only exposed glue on such an assembly is a strength four glue on the rectangle's center tile, which is an outward glue. □

The construction of exterior glues for mortar pieces adjacent to a brick with a null glue is shown in part (a) of Figure 10. Note null glues will never be part of any minimal glue set S, so will always be represented by outward glues on a brick B_S. Outward glue z and its complementary inward glue are generic glues that appear on many mortar pieces.

The glue structure of adjacent mortar pieces for a glue g of strength $k \geq 1$ on the right face of an aTAM tile t is shown in part (b) of Figure 10 if g is an outward glue on a generated brick B_1, and in part (c) of Figure 10 if g is an inward glue on a generated brick B_2. In part (c), the image assumes the minimal glue set of B_2 is {g}, though this will not always be the case; if there are more elements in the minimal glue set generating brick B_2, then the glue on the mortar pieces remains the same but there will be mismatches between the outward glues on the mortar pieces and the inward glues on the adjacent brick. For glues on other faces of tile t, this construction is simply rotated.

Lemma 4.5. *A brick or mortar rectangle must be partially assembled before it can attach to any other tile structures.*

Proof. By inspection, any tiles within bricks or mortar rectangles have at most one inward glue external to the brick or mortar rectangle, and this glue is of strength strictly less than four. This means no single tile of a brick or mortar rectangle can attach to any external tiles without first attaching to another tile via a glue internal to the brick or mortar rectangle. Such an attachment yields an assembly of at least two tiles, which by definition means the brick or mortar rectangle is partially assembled. □

Lemma 4.6. *No pair of brick, mortar rectangle, or mortar tile assemblies either partially or fully assembled can attach to each other unless one of the assemblies is a proper subassembly of some larger assembly.*

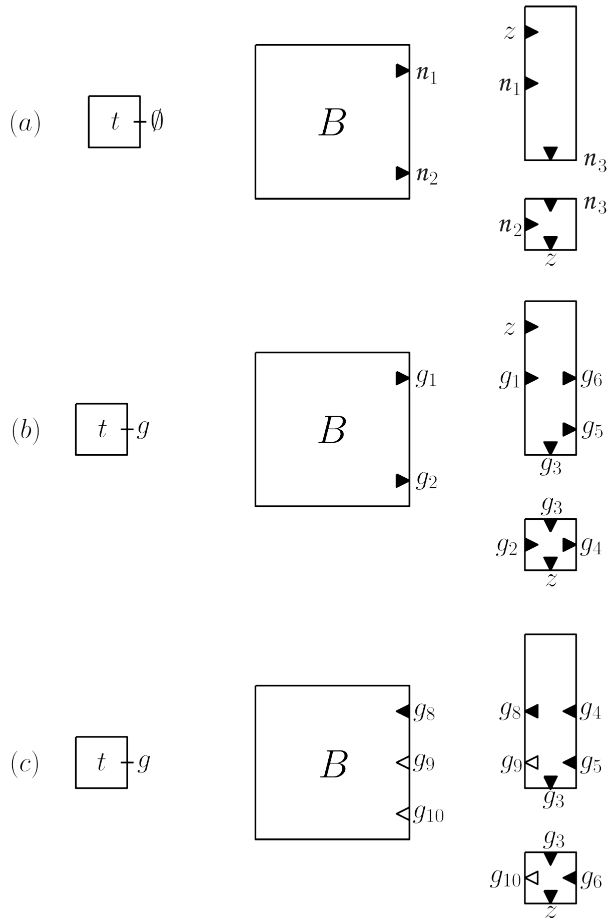

Figure 10: (a) A tile $t \in T$ with a null (strength 0) glue on its right face and corresponding brick and mortar pieces in the 2HAM simulation. (b) A tile $t \in T$ with a glue not in the minimal glue set S on its right face and corresponding brick and mortar pieces in the 2HAM simulation. (c) A tile $t \in T$ with a glue in the minimal glue set S on its right face and corresponding brick and mortar pieces in the 2HAM simulation. For this example, the minimal glue set is the singleton set containing g.

Proof. Consider attachment involving only glues on the interior of bricks or mortar rectangles; such interior glues are exposed only in partially assembled bricks and mortar rectangles. By Lemmas 4.3 and 4.4, any exposed interior glues on partially assembled bricks and mortar rectangles are exclusively outward glues. Since two outward glues cannot attach, this means glues internal to partially completed bricks or mortar rectangles do not bond with any glues found on the interior of other partially completed bricks or mortar rectangles.

Next, note that because the glues used on the interior of bricks and mortar rectangles are distinct from glues used on the exterior of bricks, mortar rectangles, and mortar tiles, exposed interior glues on partially completed bricks and mortar rectangles cannot attach to such exterior glues. Since all glues are either interior or exterior to bricks, mortar rectangles, or mortar squares and exposed interior glues cannot attach to either, then exposed interior glues on partially completed bricks or mortar rectangles cannot attach to any other fully or partially completed bricks, mortar rectangles, or mortar tiles.

Consider the attachment of a pair of brick, mortar rectangle, or mortar tile assemblies, where neither of the pair is a proper subassembly of some larger assembly. By the previous paragraph, any attachments must occur between glues on the exterior of bricks, mortar rectangles, and mortar tiles. All such glues have strength less than three, which implies that two matching glues are necessary for attachment. By considering the six possible pairs of bricks, mortar rectangles, and mortar squares, we will see that no pair has sufficient matching glues for attachment to occur.

- *Two Mortar Rectangles:* Between any two mortar rectangles, there is at most one common glue g_5 of strength two.
- *Two Mortar Squares:* Any translation of two mortar squares has the potential for at most one matching glue of strength at most two.
- *Mortar Rectangle and Mortar Square:* Any translations of a mortar rectangle and mortar square has the potential for at most one matching glue of strength at most two.
- *Mortar Rectangle and Brick:* Between a brick and a mortar rectangle, there is at most one matching glue of strength 2 (n_2, g_2, or g_8) and one matching glue of strength 1 (g_9), giving a total maximum strength of three.
- *Mortar Square and Brick:* Between a brick a a mortar square, there is at most one common glue (n_2, g_2, or g_{10}) of strength at most two.
- *Bricks:* Consider attachment between two bricks. An inward (outward) glue on a brick only has a complementary outward (inward) glue on mortar rectangles and mortar tiles. Since the inward glues g_1, g_2 and the outward glues g_8, g_9, g_{10} (the complements to the glues that appear on the exterior of any bricks) only appear on mortar rectangles and mortar tiles, this means no translation of a pair of bricks can have any positive strength bonds.

In any case, there is never a glue set of strength at least four between any pair of mortar squares, mortar rectangles, or bricks. Consequently, no pair of these assemblies can attach unless one of the assemblies is a proper subassembly of some larger assembly, that is, unless there are more tiles present in a superassembly of either piece to facilitate attachment. □

Lemmas 4.5 and 4.6 imply tiles or partially completed supertile pieces (bricks or mortar) in $\mathcal{T}'$ cannot attach to each other unless one piece is part of a larger assembly. Since we will see next that larger assemblies only begin to form around a special seed brick, this ensures that any larger assembly contains a seed brick, mimicking the seeded nature of $\mathcal{T}$.

The Assembly Process of $\mathcal{T}'$ The seed tile s of $\mathcal{T}$ is represented by a brick B_s in $\mathcal{T}'$ with all outward glues, where outward glues and adjacent mortar pieces are created as above; there will be infinite copies of this seed brick, since in any 2HAM system there are infinite copies of every tile type. However, this brick (and all of its infinite copies) is modified slightly so that one glue g_1 or g_2 is of strength 4 instead of strength 2, and the corresponding mortar piece is modified as well. This forms the only exception to lemma 4.6, meaning a single mortar tile or rectangle can attach to B_s

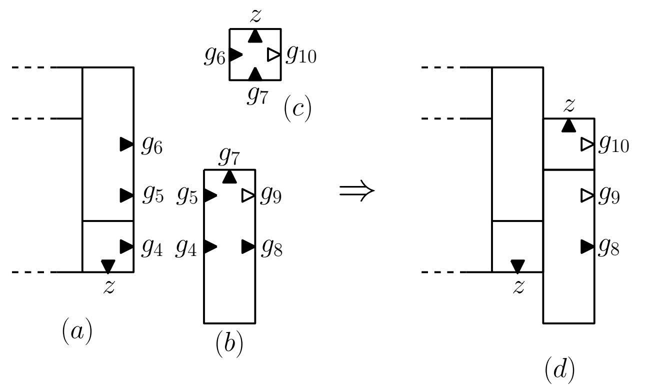

Figure 11: (a) A brick B in $\mathcal{T}'$ corresponding to tile $t \in T$ with mortar completed on its right side. (b) At this point, a mortar rectangle can attach to the assembly. (c) Next, a mortar square can attach. (d) At this point, a new center brick could attach with a strength 4 attachment via glues g_8, g_9, and g_{10}.

without either the brick or the mortar piece being a proper subassembly of some larger assembly. This starts the process of assembling supertiles.

Once one side of the mortar surrounding a brick is completed, the mortar pieces for the adjacent brick can attach; see Figure 11. After this process, there are exposed outward glues available for a new center brick to possibly attach, simulating an exposed glue in $\mathcal{T}$. A brick will attach to the assembly precisely when all inward glues (i.e. a complete minimal glue set) on a brick match the exposed glues on adjacent mortar pieces. Once one brick has attached to the assembly, all remaining mortar pieces adjacent to this brick can attach in clockwise order, completing the supertile; see Figure 12. Once one outward side of the supertile is completed, new adjacent mortar pieces can then begin to attach, and this process repeats.

Note that if a brick attaches to a partially completed assembly, then it must have attached at two or more tiles, meaning the center of the brick is also present by Lemma 4.2. This uniquely determines which supertile is present at this location in the assembly. Moreover, if a mortar rectangle attaches to the partially completed assembly, then it must attach at exactly two of its tiles, including the middle of its three tiles, and the rectangle is uniquely determined.

Maintaining Invariants in $\mathcal{T}'$

Lemma 4.7. *All exposed glues on any assembly containing a seed brick B_s are outward glues.*

Proof. This proof will proceed by induction. The exposed glues on all seed bricks B_s are all outward glues by construction. Suppose that all exposed glues are oriented outward in some partially completed assembly containing a seed brick. Lemma 4.5 implies that only a brick or mortar rectangle that is at least partially completed or a mortar tile can attach to the assembly, and it must attach at its own inward glues. Inspection shows that the set of exposed inward glues on any fully assembled mortar piece or brick is minimal, meaning the removal of any inward glue from this set gives a total inward glue strength of less that τ; lemmas 4.3 and 4.4 imply that this holds even if a brick or mortar rectangle is only partially completed. Thus, if a mortar piece or brick attaches to the assembly, it must attach at all of its inward glues, leaving only outward glues exposed on its exterior.

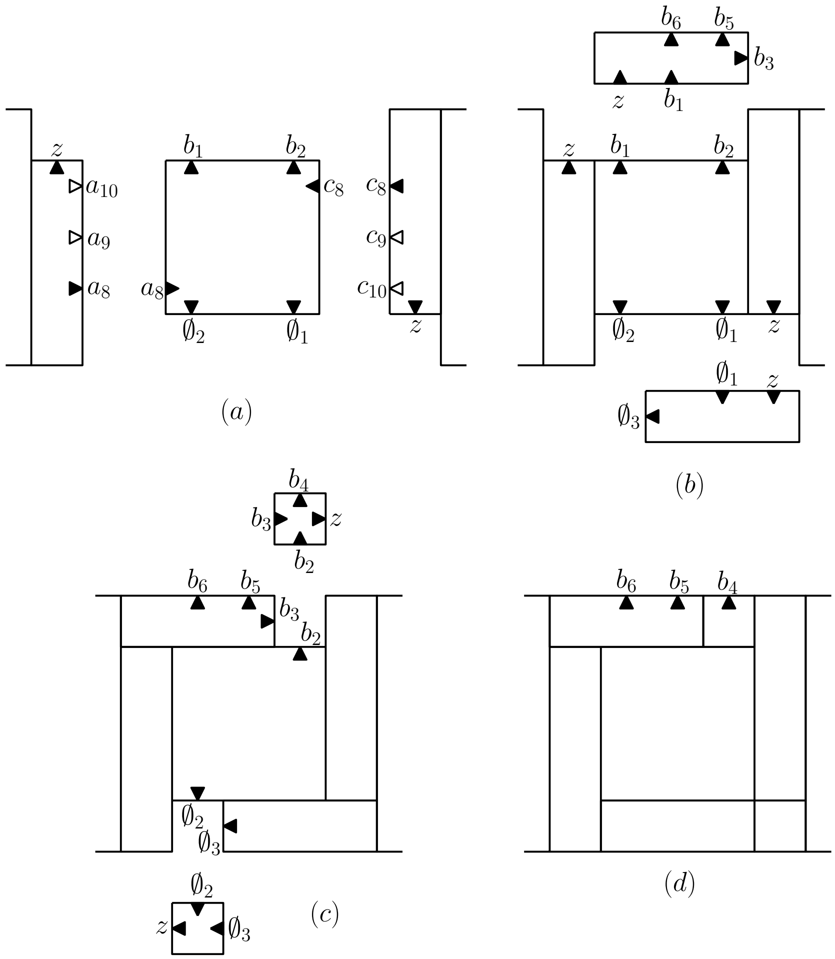

Figure 12: (a) Partially completed mortar attaching to a center brick; (b) additional adjacent mortar rectangles attach. (c) Next, mortar squares attach; note there may be outward glues that are blocked by other pieces in the assembly. (d) The supertile is completed.

□

Lemma 4.8. *If a mortar piece is attached to a brick or to another mortar piece, then both pieces are part of a partially built assembly containing a seed brick B_s.*

Proof. Lemma 4.6 states that bricks and mortar pieces cannot attach to each other unless other tiles are present. The only exception is a seed brick, which can attach to one adjacent mortar piece without any other tiles present. So, every partially completed assembly involving more than one mortar piece or one brick must include a seed brick B_s. □

Theorem 4.9. *Any aTAM system $\mathcal{T}$ at temperature $\tau \geq 4$ can be simulated by a 2HAM system $\mathcal{T}'$ at temperature $\tau = 4$.*

Proof. Transform aTAM system $\mathcal{T}$ into 2HAM system $\mathcal{T}'$ using the construction described in this section. By defining a 5-block supertile replacement function between the simulating tile system $\mathcal{T}'$

and the original tile system $\mathcal{T}$, it is possible to see that the two systems have equivalent production and dynamics, the necessary characteristics of a simulation.

Without loss of generality, any assembly A' in $\mathcal{T}'$ can be translated translated so that the supertile boundaries align with the grid $k\mathbb{Z} \times k\mathbb{Z}$, where $k\mathbb{Z} = \{... - 2k, -k, 0, k, 2k, ...\}$. Recall that $A'_k(x, y)$ is the $k \times k$ supertile in $\mathcal{T}'$ with its bottom left corner at (kx, ky). Define a 5-block supertile replacement function R that maps each supertile in $\mathcal{T}'$ that contains a partially completed brick and that is part of an assembly containing a seed brick to the tile $t \in T$ that generated that brick. The domain of partial function R is precisely the $A'_k(x, y)$ that contain a partially completed brick that is connected to an assembly containing a seed brick. Extend R so it maps all other supertiles $A'_k(x, y)$ to the empty tile. Clearly R is a valid 5-block supertile replacement function. Let R^* be the k-block assembly replacement function corresponding to this supertile replacement function, that is, the function that maps supertile assemblies in $\mathcal{T}'$ to assemblies in $\mathcal{T}$ according to which tiles R maps individual supertiles to.

These functions R and R^* will now be used to prove that tile assembly systems $\mathcal{T}$ and $\mathcal{T}'$ have equivalent production and dynamics via induction on the number of tiles in an assembly in $\mathcal{T}$. First, note any assembly of size one in $\mathcal{A}[T]$ is simply the seed tile s of $\mathcal{T}$. By Lemmas 4.8 and 4.7, any assembly in $\mathcal{T}'$ containing a combination of more than one brick or mortar piece contains exactly one seed brick, so any assembly in $\mathcal{A}[T']$ containing exactly one nonempty supertile consists of only the seed brick B_s and possibly some adjacent mortar tiles. For assemblies in $\mathcal{A}[T]$ of size one, equivalent production holds because any produced assemblies in $\mathcal{T}$ consist of only one nonempty supertile containing the seed brick, which under R^* is mapped to the seed tile s in $\mathcal{T}$, the only possible produced assembly in $\mathcal{T}$. Equivalent dynamics also holds. No attachments have occurred in $\mathcal{T}$ because the assembly of size one only contains only one tile. Consider any assembly X' consisting of exactly one nonempty supertile formed via some sequence of attachments of pieces to a partially completed seed brick B_s, the smallest assembly in $\mathcal{T}'$ that contains a nonempty supertile. Since both X' and B_s consist of only one nonempty supertile, then $R^*(X') = R^*(B_s) = s$, the seed tile in $\mathcal{T}$. So, $R^*(X')$ is trivially produced from $R^*(B_s)$ by an (empty) sequence of attachments. This means that $\mathcal{T}$ and $\mathcal{T}'$ have equivalent production and dynamics for tile assemblies in $\mathcal{T}$ of size one.

Now, suppose equivalent production and dynamics hold for all assemblies in $\mathcal{A}[T]$ of size up to n tiles. By equivalent production, this is equivalent to the same properties holding for all assemblies in $\mathcal{T}'$ having up to n non-empty supertiles.

Equivalent Production: Let $X' \in \mathcal{A}[\mathcal{T}']$ have $n + 1$ nonempty supertiles; consider all empty supertiles $X'_k(x, y)$, which by the definition of empty given in section 3.5 may contain mortar tiles and rectangles attached to adjacent supertiles but do not contain any tiles in their center 3×3 region. Empty supertiles map to the empty tile in $\mathcal{T}$ under R by the definition of a k-block replacement function. However, if the supertile is non-empty, i.e. it contains non-empty tiles in the center 3×3 subassembly, then the center tile must be present by Lemma 4.2, and the supertile contains a partially completed brick which maps to a tile $t \in T$ under R. So supertile $X'_k(x, y)$ is empty precisely when $R(X'_k(x, y))$ is the empty tile, and consequently R^* maps cleanly for all assemblies and simulation property 1(b) holds.

Moreover, any partially completed brick added to an assembly A' in $\mathcal{A}[\mathcal{T}']$ that has n nonempty supertiles to form assembly X' with $n + 1$ nonempty supertiles must be generated by a tile $t \in T$

that can attach to $R^*(A') = A$ to form $R^*(X') = X$. The minimal glue set at which the brick attaches to A' corresponds to a glue set on t via which t can attach to A, and this brick attaches precisely when t can attach to A at this minimal glue set to form X. This means $R^*(X') \in \mathcal{A}[T]$. Additionally, for X in $\mathcal{A}[T]$ produced by adding t to assembly A, because A' exists in $\mathcal{A}[T']$ by the inductive hypothesis, then X' can be produced by adding some mortar pieces and a center brick to A', meaning $X = R^*(X')$ for some $X' \in \mathcal{A}[T']$. Since both directions of containment hold, $\{R^*(A')|A' \in \mathcal{A}[T']\} = \mathcal{A}[T]$ for tile assemblies of size up to $n+1$, and because R^* maps cleanly R has equivalent production.

Equivalent Dynamics: Define X to be an $(n+1)$-tile assembly in $\mathcal{A}[T]$, produced from A by the addition of one tile t. Let A' be an assembly in $\mathcal{A}[T']$ with $R^*(A') = A$, guaranteed to exist by the equivalent production of $\mathcal{T}'$. Let A'' be A' with all possible mortar tiles and rectangles added to the 5-block supertile corresponding to the location where t will attach. Then X' with $(n+1)$ non-empty supertiles can be generated by adding to A'' a single partially completed brick containing the center tile of this 5-block, and $R^*(X') = X$. So, simulation property 2(a) holds, because X' could be produced from A'.

Moreover, for any assemblies $A', X' \in \mathcal{A}[T']$, consider adding any tiles or tile assemblies to A' to produce X'; by lemmas 4.7 and 4.8, all such attachments consist of adding only a single mortar tile, partially completed mortar rectangle, or partially completed brick, or the attachment of single tiles to complete a partially completed rectangle or brick. For any such addition, at most one 5-block supertile has the identity of its center brick modified, so $R^*(X')$ is produceable from $R^*(A')$ by the addition of at most one tile, the tile generating the partially completed brick that possibly attached to form X'. So simulation propert 2(b) holds, and R has equivalent dynamics.

Because R has equivalent production and equivalent dynamics, this construction is a simulation and any aTAM tile assembly system $\mathcal{T}$ at arbitrary temperature can be simulated by a 2HAM tile assembly system $\mathcal{T}'$ at temperature 4. $\square$

4.1 Limitations

Fuzziness In the definition of a simulation and within the previous proof, it was assumed that a supertile was empty even if it actually contains a single layer of mortar tiles. This gives our simulation a "fuzziness" of one, meaning any final shape in $\mathcal{T}'$ may not be an exact copy of a final shape in $\mathcal{T}$, but rather may have in places an extra layer of tiles around its boundary. This "fuzziness" in in fact necessary for this simulation. In a final assembly in $\mathcal{T}$, there might be a glue exposed on the exterior of the assembly even though there is no tile that can possibly attach to this glue. Correspondingly, for glue g on the exterior of a tile t in a final assembly in $\mathcal{T}$, the supertile corresponding to t in $\mathcal{T}'$ has exposed glues g_4, g_5, and g_6; to these glues, another mortar rectangle and mortar tile can attach; see figure 13. However, no brick can then attach to glues g_8, g_9, and g_{10}, because no tile t could have attached to g in $\mathcal{T}$, meaning this assembly is terminal. Because any final assemblies in $\mathcal{T}'$ necessarily have these extra mortar pieces around their boundaries whenever the corresponding final assembly in $\mathcal{T}$ has an unused exposed glue, any definition of a simulation must capture the presence of these extra tiles. If $\mathcal{T}$ is restricted such that the exposed glues on all final assemblies are null glues, then a simulation can be defined without fuzziness.

Additionally, the definition of *simulate* used here allows small assemblies of size less than 5^2

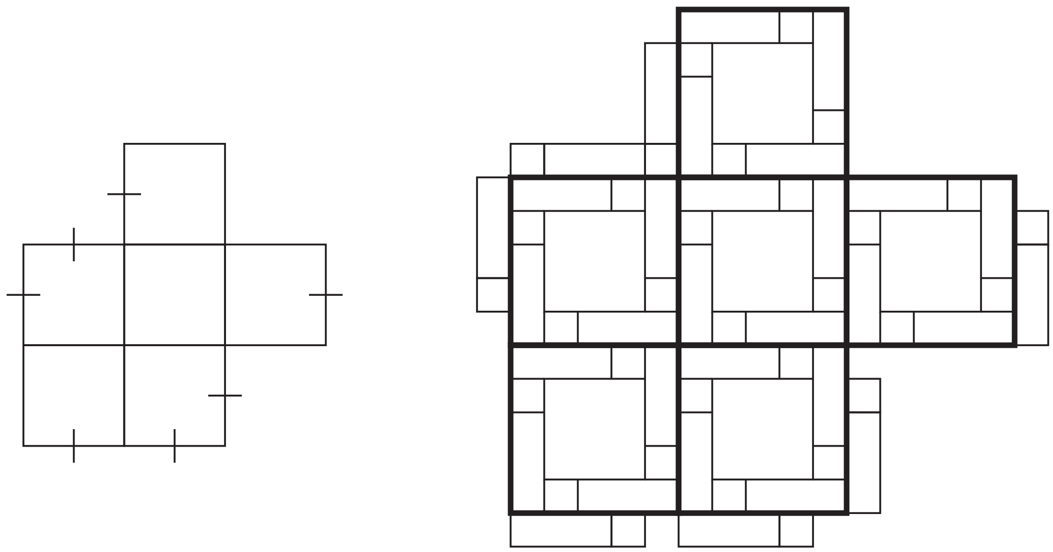

(a) (b)

Figure 13: **(a)** a terminal tile assembly in $\mathcal{T}$, with exposed external glues marked; **(b)** the terminal shape of the simulation of this tile assembly in $\mathcal{T}'$. Because no tiles can attach to the exposed glues in part (a), no bricks can attach to the exterior mortar pieces in part (b)

that are not part of any larger assembly to form; because these assemblies are not large enough to necessarily be in the domain of R, they are not mapped to any tile in $\mathcal{T}$, and in fact may never be used in any assembly in $\mathcal{T}'$. This arises because not all minimal glue sets of a tile t may be used for attachment of t to the seed assembly during the assembly process of $\mathcal{T}$. Consequently, there may be some bricks in $\mathcal{T}'$ corresponding to these minimal glues sets that are never used in an assembly. Similarly, there may also be some mortar pieces that never attach to a larger assembly as well.

Because of these nuances, the simulation result described in this section is not a true simulation in the strictest sense of the term. However, when building large scale assemblies in $\mathcal{T}$, the presence of a few additional extra boundary tiles or some additional small constant-sized tile assemblies is insignificant when compared to the size and structure of the assembly and the degree to which it successfully mimics the behavior of $\mathcal{T}$.

Connectivity The simulation described in this section also does not necessarily preserve the full connectivity of tile assembly system $\mathcal{T}$. In order to preserve the invariant that all exposed glues on seed assemblies in $\mathcal{T}'$ are outward glues, all bricks attach to the assemblies only at minimal glue sets. However, a tile in $\mathcal{T}$ need only attach to the seed assembly with a set of glues of strength greater that τ, and this set of glues may not be minimal. By restricting attachment in $\mathcal{T}'$ to minimal glue sets, the additional connectivity gained when t attaches at a glue set that is not minimal is not preserved in $\mathcal{T}'$. For an example of a tile assembly for which this simulation does not preserve connectivity, see figure 14.

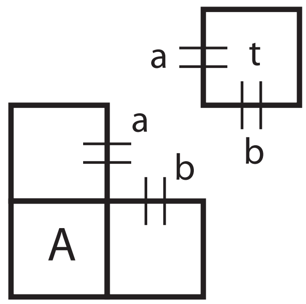

Figure 14: A tile t that could potentially attach to assembly A to form a fully connected square in $\mathcal{T}$, with $\tau = 2$. In $\mathcal{T}'$, the simulation of this tile assembly would not be fully connected, as either the brick B_t corresponding to minimal glue set {a} or minimal glue set {b} would attach, in both cases producing a glue mismatch on the other supertile face.

Uniqueness If tile assembly system $\mathcal{T}$ is unique, then tile assembly system $\mathcal{T}'$ is unique, with fuzziness. That is, if $\mathcal{T}$ only ever assembles one unique shape, then $\mathcal{T}'$ assembles infinitely many copies of this shape(scaled up by 5), all of which differ from each other by at most one layer of tiles around their boundary. This occurs because while $\mathcal{T}$ only ever produces assemblies of one shape, each time this assembly is produced it may have different tiles in different locations; in particular, there may be glues exposed at different places around the exterior of this assembly. This means the attachment of the extra mortar pieces to exposed external glues in $\mathcal{T}'$ which necessitates the fuzziness of this simulation may occur at different places around the boundary of the final assembly in $\mathcal{T}$, producing slightly different shapes.

If $\mathcal{T}$ is a strongly unique tile assembly system, then $\mathcal{T}'$ is a unique tile assembly system. Because the same tiles are always in the same location in the unique assembly A assembled in $\mathcal{T}$, then the extra mortar tiles that attach to exposed external glues on assembly $F(A)$ as described above always attach in the same locations, and thus any assembly produced by $\mathcal{T}'$ will always have the same shape. System $\mathcal{T}'$ is not strongly unique because of the same issue discussed in the previous subsection. If a tile t attaches to the seed assembly in $\mathcal{T}$ at more than a minimal glue set, then there may be multiple bricks corresponding to t in $\mathcal{T}'$. For example, if tile t attaches to the seed assembly at glues a, b, and c, all of strength $\tau - 1$, then t has at least three minimal glue sets, $S_1 = \{a, b\}$, $S_2 = \{a, c\}$, and $S_3 = \{b, c\}$. Any one of the corresponding bricks B_1, B_2, or B_3 could attach to the adjacent mortar pieces in $\mathcal{T}'$, meaning assemblies in $\mathcal{T}'$ are not strongly unique.

Recall that lemma 3.1 says $\mathcal{T}'$ is a terminal tile assembly system precisely when $\mathcal{T}$ is.

4.2 Scale Factors

Each tile in $\mathcal{T}$ is representile by a 5-tile by 5-tile block in $\mathcal{T}'$. In the construction described above, if aTAM system $\mathcal{T}$ has t tile types and g glues, then a careful count shows 2HAM system $\mathcal{T}'$ has at most $54t + 32g$ tiles and $72t + 26g$ glues, where an inward glue g and an outward glue g are considered the same glue and only counted once. This is only an upper bound, however; this count assumes all tiles in $\mathcal{T}$ have the maximum six minimum glue sets and thus generate six bricks in $\mathcal{T}'$, but in practice this will often not be the case. A careful design of blocks and tiles can yield much lower constants; one such modification yields the same number of tiles but only $18t + 18g$ glues.

However, in any such modifications the glues within bricks and mortar rectangles are no longer unique, and the proofs of lemmas 4.5 and 4.6 require much more detail.

Although these scale factors may seem high, they are still constant. Because many simulations require a scale factor with some dependence on the size of the assemblies built, as in [11] and [13], the fact that this simulation has constant scale factor is much more interesting than how large that constant is.

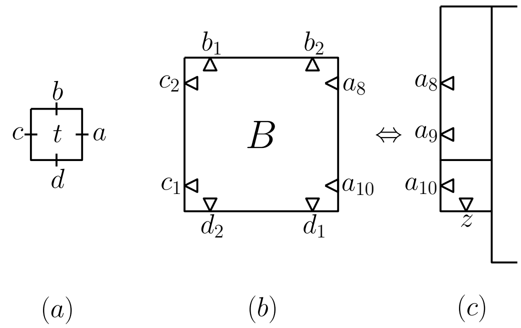

Figure 15: Unfilled arrows represent glues of strength 1. (a) an aTAM tile t with minimal glue set $\{a\}$; (b) the brick B in the 2HAM system generated by this minimal set; (c) a location where B could attach to a partially built assembly.

5 Extensions and Generalizations

5.1 Simulating aTAM at $\tau \in \{1, 2\}$ with 2HAM $\tau = 2$

The construction described in the previous section can be modified to also enable simulating aTAM systems at $\tau = \{1, 2\}$ with the 2HAM at $\tau = 2$ with scale factor 5.

Theorem 5.1. *Any aTAM system at $\tau \in \{1, 2\}$ can be simulated by a 2HAM system at $\tau = 2$.*

Since minimal glue sets have at most 2 glues, $\tau = 2$ is sufficient for determining when a minimal glue set is sufficient to bond two assemblies.

Modifying the construction involves changing all strength 2 and 4 glues to strength 1 and 2 respectively, and modifying how bricks for minimal glue sets are generated. Because minimal glue sets at $\tau = 2$ contain at most 2 glues there are 2 cases, rather than 4, for generating a brick based on a minimal glue set. See Figures 15 and 16 for constructing the bricks in these cases, as well as for the structure of adjacent mortar pieces corresponding to inward glues. Adjacent mortar pieces corresponding to outward glues are the same as in the previous section, with glue strengths adjusted accordingly.

We model an aTAM system at $\tau = 1$ with equivalent aTAM system at $\tau = 2$ where each glue is strength two instead of strength one; we then apply the same construction to simulate any aTAM system at $\tau = 1$ with a 2HAM system at $\tau = 2$.

5.2 Simulating aTAM at $\tau = 3$ with 2HAM $\tau = 3$

The construction used to simulate the $\tau \geq 4$ aTAM model with the $\tau = 4$ 2HAM model can also be modified to simulate the $\tau = 3$ aTAM model with the $\tau = 3$ 2HAM model.

The modification only changes the bricks generated for each tile. Since the aTAM system being simulated is $\tau = 3$, minimal glue sets have size at most 3. The three cases for generating bricks for minimal glue sets of sizes 1,2, and 3 are seen in Figures 17, 18, and 19.

Theorem 5.2. *Any aTAM system at $\tau = 3$ can be simulated by a 2HAM system at $\tau = 3$.*

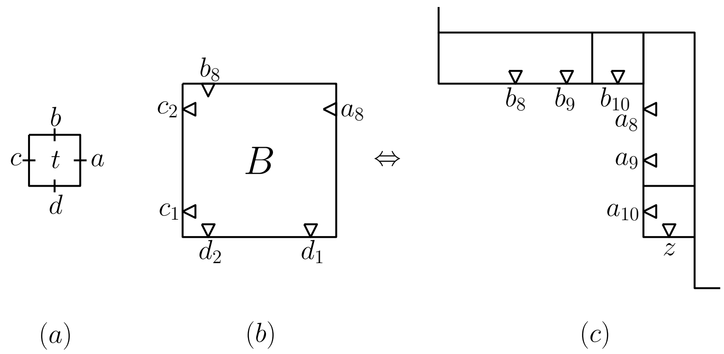

Figure 16: (a) an aTAM tile t with minimal glue set $\{a, b\}$; (b) the brick B in the 2HAM system generated by this minimal set; (c) a location where B could attach to a partially built assembly.

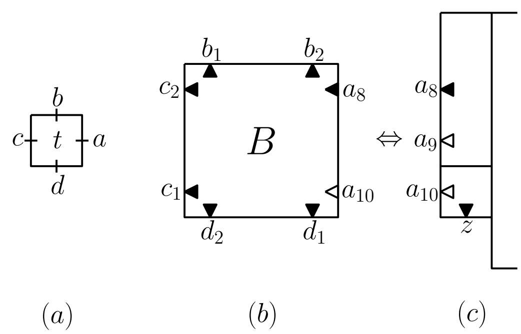

Figure 17: (a) an aTAM tile t with minimal glue set $\{a\}$; (b) the brick B in the 2HAM system generated by this minimal set; (c) a location where B could attach to a partially built assembly.

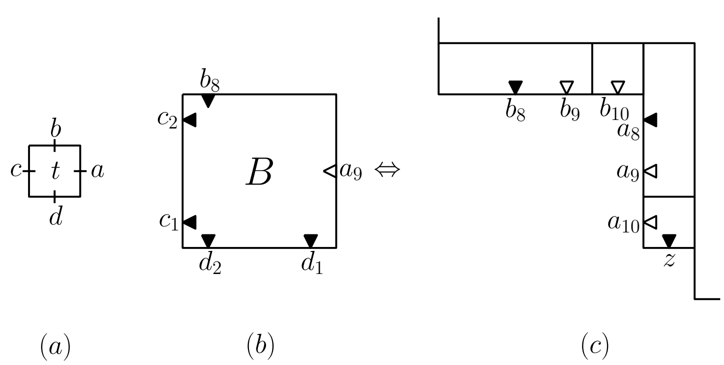

Figure 18: (a) an aTAM tile t with minimal glue set $\{a, b\}$; (b) the brick B in the 2HAM system generated by this minimal set; (c) a location where B could attach to a partially built assembly.

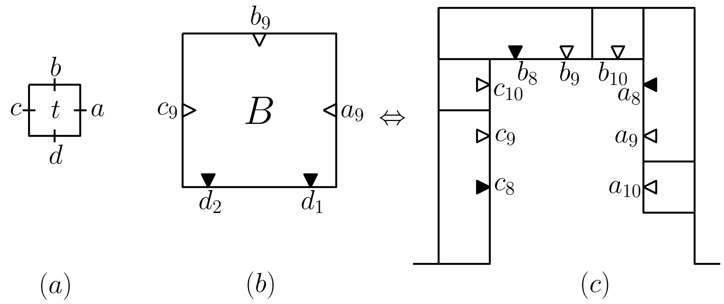

Figure 19: (a) an aTAM tile t with minimal glue set $\{a, b, c\}$; (b) the brick B in the 2HAM system generated by this minimal set; (c) a location where B could attach to a partially built assembly.

5.3 Simulating Planar aTAM at τ with 2HAM at $\tau = 3$

Recall that a planar tile assembly system is a system in which two tiles can never overlap. In this case, all attachments consist of tile assemblies attaching along their exteriors. Because of this, in any planar aTAM system, any tile attaches to the exterior of the seed assembly using at most three of its glues, meaning all minimal glue sets are of size at most three. Consequently, the construction given in the previous subsection also simulates any planar aTAM tile system $\mathcal{T}$ at temperature τ with a 2HAM system $\mathcal{T}'$ at temperature 3. If we add the additional restriction that $\mathcal{T}'$ must be a planar tile assembly system as well, then while this simulation is still valid, there may be some supertiles in $\mathcal{T}'$ that are never fully completed. Because tiles can never overlap, it might not be possible for the tiles needed to complete certain supertiles to move to locations adjacent to those partially completed supertiles. If system $\mathcal{T}'$ is not required to be planar, however, then all supertiles can be fully completed, but there may be some assemblies A that are terminal in $\mathcal{T}$, because planarity requirements restrict the motions of tiles, while $F(A)$ is not terminal in $\mathcal{T}'$.

6 Open Problems

Our construction has a fuzziness of one, does not preserve full connectivity of an assembled shape, and does not preserve uniqueness or strict uniqueness. A simulation without even one of these limitations would be an even stronger result, though it is not immediately evident how to achieve this.

Another open problem is extending these simulation results to three dimensions. In three-dimensional tile self-assembly, tiles are actually cubes with six glues, one on each face.

Conjecture 6.1. *Any three-dimensional tile system $\mathcal{T}$ in the abstract tile assembly model at temperature τ can be simulated by a three-dimensional tile system $\mathcal{T}'$ at temperature 6 in the two-handed tile assembly model.*

Another natural question to ask is whether it is possible to simulate any 2HAM system with an aTAM system. I conjecture the answer is no, though I expect such a conclusion will be extremely hard to prove. Additionally, this result does not simulate aTAM systems at temperature one with a 2HAM system at temperature one; I again conjecture such a simulation is not possible because of the difficulty of creating well-behaved temperature one 2HAM systems.

7 Conclusion

This thesis describes a simulation that can be used to simulate any aTAM system at any temperature with a 2HAM system at an equal or lower temperature. Such a result provides fundamental insights into the relative computational power of the aTAM and the 2HAM, the two main models of tile self-assembly in use today. The existence of such a simulation can also be applied to other results, as well. Doty et al. showed that there exists a single aTAM system at temperature 2 that simulates any other aTAM tile assembly system [13]; our result then implies there also exists a single 2HAM tile assembly system at temperature 2 that can simulate any aTAM system. Additionally, distinct glues are relatively expensive to create in practical experiments; a new result by Allen, Cannon, Damian, Flatland, and Silveira (not yet published, tentatively titled ”Self-Assembly using Few Glues per Tile”) shows that any 2HAM system can be simulated by another 2HAM system with at most two unique non-null glues per tile. This simulation result then implies that any aTAM system can be simulated by a 2HAM system with at most two unique non-null glues per tile as well. Because of the existence of this simulation, any future simulation results can be extended to apply to both aTAM and 2HAM systems as well.

With the plethora of tile assembly models explored recently by researchers, many differing from each other only slightly, a formal hierarchy of these models would be extremely useful. DNA self-assembly systems offer a new way to compute the answers to complex problems, and just as there is a hierarchy of computational classes such as P, NP, and P-SPACE , hopefully there will soon be a hierarchy of tile-self assembly models. Just as we know $P \subseteq NP$, this result suggests the set of assembly processes producible by aTAM systems are contained within the set of assembly processes producible by 2HAM systems, albeit with fuzziness and a scale factor of five. Perhaps, among the many tile self-assembly models, there is one definitive model which can simulate all other tile self-assembly models; perhaps certain tile self-assembly models are equivalent. I hope these results inspire other researchers to continue to investigate this problem of the relationships between different self-assembly models, which is one that will help the tile self-assembly community to better understand the self-assembly process as a whole.

References

[1] Zachary Abel, Nadia Benbernou, Mirela Damian, Erik Demaine, Martin Demaine, Robin Flatland, Scott Kominers, and Robert Schweller, *Shape replication through self-assembly and RNase enzymes*, SODA 2010: Proceedings of the Twenty-first Annual ACM-SIAM Symposium on Discrete Algorithms (Austin, Texas), Society for Industrial and Applied Mathematics, 2010.

[2] Leonard Adleman, *Toward a mathematical theory of self-assembly (extended abstract)*, Tech. Report 00-722, University of Southern California, 2000.

[3] Leonard Adleman, Qi Cheng, Ashish Goel, and Ming-Deh Huang, *Running time and program size for self-assembled squares*, Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (Hersonissos, Greece), 2001, pp. 740–748.

[4] Leonard M. Adleman, Qi Cheng, Ashish Goel, Ming-Deh A. Huang, David Kempe, Pablo Moisset de Espanés, and Paul W. K. Rothemund, *Combinatorial optimization problems in self-assembly*, Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing, 2002, pp. 23–32.

[5] Florent Becker, Ivan Rapaport, and Eric Rémila, *Self-assembling classes of shapes with a minimum number of tiles, and in optimal time*, Foundations of Software Technology and Theoretical Computer Science (FSTTCS), 2006, pp. 45–56.

[6] Ravinderjit S. Braich, Nickolas Chelyapov, Cliff Johnson, Paul W. K. Rothemund, and Leonard Adleman, *Solution of a 20-variable 3-sat problem on a dna computer*, Science **296** (2002), no. 5567, 499–502.

[7] Sarah Cannon, Erik D. Demaine, Martin L. Demaine, Sarah Eisenstat, Matthew J. Patitz, Robert Schweller, Scott M. Summers, and Andrew Winslow, *Two hands are better than one (up to constant factors)*, Proceedings of the 30th International Symposium on Theoretical Aspects of Computer Science (Kiel, Germany), 2013.

[8] Ho-Lin Chen and David Doty, *Parallelism and time in hierarchical self-assembly*, SODA 2012: Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, 2012.

[9] Qi Cheng, Gagan Aggarwal, Michael H. Goldwasser, Ming-Yang Kao, Robert T. Schweller, and Pablo Moisset de Espanés, *Complexities for generalized models of self-assembly*, SIAM Journal on Computing **34** (2005), 1493–1515.

[10] Matthew Cook, Yunhui Fu, and Robert T. Schweller, *Temperature 1 self-assembly: Deterministic assembly in 3D and probabilistic assembly in 2D*, SODA 2011: Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, 2011.

[11] Erik D. Demaine, Martin L. Demaine, Sándor P. Fekete, Mashhood Ishaque, Eynat Rafalin, Robert T. Schweller, and Diane L. Souvaine, *Staged self-assembly: nanomanufacture of arbitrary shapes with $O(1)$ glues*, Natural Computing **7** (2008), no. 3, 347–370.

[12] Erik D. Demaine, Matthew J. Patitz, Robert T. Schweller, and Scott M. Summers, *Self-assembly of arbitrary shapes using rnase enzymes:. meeting the kolmogorov bound with small scale factor (extended abstract)*, STACS 2011: Proceedings of the 28th International Symposium on Theoretical Aspects of Computer Science, 2010.

[13] David Doty, Jack H. Lutz, Matthew J. Patitz, Scott M. Summers, and Damien Woods, *Intrinsic universality in self-assembly*, Proceedings of the 27th International Symposium on Theoretical Aspects of Computer Science, 2009, pp. 275–286.

[14] Ming-Yang Kao and Robert T. Schweller, *Reducing tile complexity for self-assembly through temperature programming*, Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2006), Miami, Florida, Jan. 2006, pp. 571-580, 2007.

[15] John D. Le, Yariv Pinto, Nadrian C. Seeman, Karin Musier-Forsyth, T. Andrew Taton, and Richard A. Kiehl, *Dna-templated self-assembly of metallic nanocomponent arrays on a surface*, Nano Letters **4** (2004), no. 12, 2343–2347.

[16] Chenxiang Lin, Yan Liu, Sherri Rinker, and Hao Yan, *Dna tile based self-assembly: Building complex nanoarchitectures*, ChemPhysChem **7** (2006), 1641–1647.

[17] Kristian Lindgren and Mats G. Nordahl, *Universal computation in simple on-dimensional cellular automata*, Complex Systems **4** (1990), 299–318.

[18] Yan Liu, Chenxiang Lin Hanying Li, and Hao Yan, *Aptamer-directed self-assembly of protein arrays on a dna nanostructure*, Angewandte Chemie **117** (2005), 4407–4412.

[19] Chris Luhrs, *Polyomino-safe DNA self-assembly via block replacement*, DNA14 (Ashish Goel, Friedrich C. Simmel, and Petr Sosík, eds.), Lecture Notes in Computer Science, vol. 5347, Springer, 2008, pp. 112–126.

[20] Matthew J. Patitz and Scott M. Summers, *Identifying shapes using self-assembly (extended abstract)*, Algorithmica (2010), to appear. Special issue of invited papers from the Twenty First International Symposium on Algorithms and Computation.

[21] Paul W. K. Rothemund 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 (Portland, Oregon, United States), ACM, 2000, pp. 459–468.

[22] David Soloveichik and Erik Winfree, *Complexity of self-assembled shapes*, SIAM Journal on Computing **36** (2007), no. 6, 1544–1569.

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

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

[25] ———, *Self-healing tile sets*, Nanotechnology: Science and Computation (Junghuei Chen, Natasa Jonoska, and Grzegorz Rozenberg, eds.), Natural Computing Series, Springer, 2006, pp. 55–78.