

The two-handed tile assembly model is not intrinsically universal

Erik D. Demaine* Matthew J. Patitz† Trent A. Rogers‡
 Robert T. Schweller§ Scott M. Summers¶ Damien Woods||

Abstract

The well-studied Two-Handed Tile Assembly Model (2HAM) is a model of tile assembly in which pairs of large assemblies can bind, or self-assemble, together. In order to bind, two assemblies must have matching glues that can simultaneously touch each other, and stick together with strength that is at least the temperature τ , where τ is some fixed positive integer. We ask whether the 2HAM is intrinsically universal, in other words we ask: is there a single universal 2HAM tile set U which can be used to simulate any instance of the model? Our main result is a negative answer to this question. We show that for all $\tau' < \tau$, each temperature- τ' 2HAM tile system does not simulate at least one temperature- τ 2HAM tile system. This impossibility result proves that the 2HAM is not intrinsically universal, in stark contrast to the simpler (single-tile addition only) abstract Tile Assembly Model which is intrinsically universal (*The tile assembly model is intrinsically universal*, FOCS 2012). However, on the positive side, we prove that, for every fixed temperature $\tau \geq 2$, temperature- τ 2HAM tile systems are indeed intrinsically universal: in other words, for each τ there is a single universal 2HAM tile set U that, when appropriately initialized, is capable of simulating the behavior of any temperature- τ 2HAM tile system. As a corollary of these results we find an infinite set of infinite hierarchies of 2HAM systems with strictly

*Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, 32 Vassar St., Cambridge, MA 02139, USA, edemaine@mit.edu

†Department of Computer Science and Computer Engineering, University of Arkansas, patitz@uark.edu This author's research was supported in part by National Science Foundation Grants CCF-1117672 and CCF-1422152.

‡Department of Computer Science and Computer Engineering, University of Arkansas, tar003@uark.edu. This author's research was supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-1450079, and National Science Foundation grants CCF-1117672 and CCF-1422152.

§Department of Computer Science, University of Texas–Pan American, Edinburg, TX, 78539, USA. rtsweller@utpa.edu. This author's research was supported in part by National Science Foundation Grant CCF-1117672.

¶Department of Computer Science and Software Engineering, University of Wisconsin–Platteville, Platteville, WI 53818, USA. summer@uwplatt.edu.

||Computer Science, California Institute of Technology, woods@caltech.edu. This author's research was supported by National Science Foundation grants 0832824 (The Molecular Programming Project), CCF-1219274, and CCF-1162589.

increasing simulation power within each hierarchy. Finally, we show that for each τ , there is a temperature- τ 2HAM system that simultaneously simulates all temperature- τ 2HAM systems.

1 Introduction

Self-assembly is the process through which unorganized, simple, components automatically coalesce according to simple local rules to form some kind of target structure. It sounds simple, but the end result can be extraordinary. For example, researchers have been able to self-assemble a wide variety of structures experimentally at the nanoscale, such as regular arrays [33], fractal structures [15, 28], smiling faces [27], DNA tweezers [34], logic circuits [25], neural networks [26], and molecular robots [21]. These examples are fundamental because they demonstrate that self-assembly can, in principle, be used to manufacture specialized geometrical, mechanical and computational objects at the nanoscale. Potential future applications of nanoscale self-assembly include the production of smaller, more efficient microprocessors and medical technologies that are capable of diagnosing and even treating disease at the cellular level.

Controlling nanoscale self-assembly for the purposes of manufacturing atomically precise components will require a bottom-up, hands-off strategy. In other words, the self-assembling units themselves will have to be “programmed” to direct themselves to do the right thing—efficiently and correctly. Thus, it is necessary to study the extent to which the process of self-assembly can be controlled in an algorithmic sense.

In 1998, Erik Winfree [32] introduced the abstract Tile Assembly Model (aTAM), an over-simplified discrete mathematical model of nanoscale DNA self-assembly pioneered by Seeman [30]. The aTAM essentially augments classical Wang tiling [31] with a mechanism for sequential “growth” of a tiling (in Wang tiling, only the existence of a valid, mismatch-free tiling is considered and not the order of tile placement). In the aTAM, the fundamental components are un-rotatable, but translatable square “tile types” whose sides are labeled with (alpha-numeric) glue “colors” and (integer) “strengths”. Two tiles that are placed next to each other *interact* if the glue colors on their abutting sides match, and they *bind* if the strengths on their abutting sides match and sum to at least a certain (integer) “temperature”. Self-assembly starts from a “seed” tile type and proceeds nondeterministically and asynchronously as tiles bind to the seed-containing-assembly. Despite its deliberate over-simplification, the aTAM is a computationally expressive model. For example, Winfree [32] proved that it is Turing universal, which implies that self-assembly can be directed by a computer program.

In this paper, we work in a generalization of the aTAM, called the *two-handed* [3] (a.k.a., hierarchical [5], q-tile [6], polyomino [20]) abstract Tile Assembly Model (2HAM). A central feature of the 2HAM is that, unlike the aTAM, it allows two “supertile” assemblies, each consisting of one or more tiles, to fuse together. For two such assemblies to bind, they should not “sterically hinder”

each other, and they should have a sufficient number of matching glues distributed along the interface where they meet. Hence the model includes notions of local interactions (individual glues) and non-local interactions (large assemblies coming together). In the 2HAM, an assembly of tiles is producible if it is either a single tile, or if it results from the stable combination of two other producible assemblies.

We study the *intrinsic universality* in the 2HAM. Intrinsic universality uses a special notion of simulation, where the simulator preserves the dynamics of the simulated system. For tile assembly systems this means that, modulo spatial rescaling, a simulator self-assembles the same assemblies as any simulated system, and even does this in the same way (via the same assembly sequences). In the field of cellular automata, the topic of intrinsic universality has given rise to a rich theory [2, 4, 8, 9, 14, 23, 24] and indeed has also been studied in Wang tiling [16–18] and tile self-assembly [11–13, 22]. The aTAM has been shown to be intrinsically universal [12], meaning that there is a single set of tiles U that works at temperature 2, and when appropriately initialized, is capable of simulating the behavior of an arbitrary aTAM tile assembly system. Modulo rescaling, this single tile set U represents the full power and expressivity of the entire aTAM model, at any temperature. On the other hand, it has been shown that at temperature-1, there is no tile set that can simulate the aTAM [22]. Interestingly, the latter negative result holds for 3D temperature-1 systems, which are known to be Turing universal [7]. Here, we ask whether there is such a universal tile set for the 2HAM.

The theoretical power of non-local interaction in the 2HAM has been the subject of recent research. For example, Doty and Chen [5] proved that, surprisingly, $N \times N$ squares do not self-assemble any faster in so-called *partial order* 2HAM systems than they do in the aTAM, despite being able to exploit massive parallelism. More recently, Cannon, et al. [3], while comparing the abilities of the 2HAM and the aTAM, proved three main results, which seem to suggest that the 2HAM is at least as powerful as the aTAM: (1) non-local binding in the 2HAM can dramatically reduce the tile complexity (i.e., minimum number of unique tile types required to self-assemble a shape) for certain classes of shapes; (2) the 2HAM can simulate the aTAM in the following sense: for any aTAM tile system \mathcal{T} , there is a corresponding 2HAM tile system \mathcal{S} , which simulates the exact behavior—modulo connectivity—of \mathcal{T} , at scale factor 5; (3) the problem of verifying whether a 2HAM system uniquely produces a given assembly is coNP-complete (for the aTAM this problem is decidable in polynomial time [1]).

Main results. In this paper, we ask if the 2HAM is *intrinsically universal*: does there exist a “universal” 2HAM tile set U that, when appropriately initialized, is capable of simulating the behavior of an arbitrary 2HAM tile system? A positive answer would imply that such a tile set U has the ability to model the capabilities of all 2HAM systems.¹ Our first main result, Theorem 3.1, says

¹Note that the above simulation result of Cannon et al. does not imply that the 2HAM is intrinsically universal because (a) it is for 2HAM simulating aTAM, and (b) it is an example of a “for all, there exists...” statement, whereas intrinsic universality is a “there exists, for all...” statement.

that the 2HAM is *not* intrinsically universal, which means that the 2HAM is incapable of simulating itself. This statement stands in stark contrast to the case of the aTAM, which was recently shown to be intrinsically universal by Doty, Lutz, Patitz, Schweller, Summers and Woods [12]. Specifically, we show that for any sufficiently large temperature τ , there is a temperature τ 2HAM system that cannot be simulated by any temperature $\tau' < \tau$ 2HAM system. It is worthy of note that, in order to prove this result, we use a simple, yet novel combinatorial argument, which as far as we are aware of, is the first lower bound proof in the 2HAM that does not use an information-theoretic argument. In our proof of Theorem 3.1 we show that the 2HAM cannot simulate massively cooperative binding, where the number of cooperative bindings is larger than the temperature of the simulator.

Our second main result, Theorem 4.1, is positive: we show, via constructions, that the 2HAM *is* intrinsically universal for fixed temperature, that is, the temperature τ 2HAM can simulate the temperature τ 2HAM. So although our impossibility result tells us that the 2HAM can not simulate “too much” cooperative binding, our positive result tells us it can indeed simulate *some* cooperative binding: an amount exactly equal to the temperature of the simulator.

As a corollary of these results, we get a separation between classes of 2HAM tile systems based on their temperatures. That is, we exhibit an infinite hierarchy of 2HAM systems, of strictly-increasing temperature, that cannot be simulated by lesser temperature systems but can downward simulate lower temperature systems. Moreover, we exhibit an infinite number of such hierarchies in Theorem 4.2. Thus, as was suggested as future work in [12], and as has been shown in the theory of cellular automata [9], we use the notion of intrinsic universality to classify, and separate, these tile assembly systems via their simulation ability.

As noted above, we show that temperature τ 2HAM systems are intrinsically universal. We actually show this for two different, seemingly natural, notions of simulation (called *simulation* and *strong simulation*), showing tradeoffs between, and even within, these notions of simulation. For both notions of simulation, we show tradeoffs between scale factor, number of tile types, and complexity of the initial configuration. Finally, we show how to construct, for each τ , a temperature- τ 2HAM system that simultaneously simulates all temperature- τ 2HAM systems. We finish with a conjecture:

Conjecture 1. *There exists $c \in \mathbb{N}$, such that for each $\tau \geq c$, temperature τ 2HAM systems do not strongly simulate temperature $\tau - 1$ 2HAM systems.*

2 Definitions

2.1 Informal definition of the 2HAM

The 2HAM [6, 10] is a generalization of the aTAM in that it allows for two assemblies, both possibly consisting of more than one tile, to attach to each other. Since we must allow that the assemblies might require translation before

they can bind, we define a *supertile* to be the set of all translations of a τ -stable assembly, and speak of the attachment of supertiles to each other, modeling that the assemblies attach, if possible, after appropriate translation. We now give a brief, informal, sketch of the 2HAM.

A *tile type* is a unit square with four sides, each having a *glue* consisting of a *label* (a finite string) and *strength* (a non-negative integer). We assume a finite set T of tile types, but an infinite number of copies of each tile type, each copy referred to as a *tile*. A *supertile* is (the set of all translations of) a positioning of tiles on the integer lattice \mathbb{Z}^2 . Two adjacent tiles in a supertile *interact* if the glues on their abutting sides are equal and have positive strength. Each supertile induces a *binding graph*, a grid graph whose vertices are tiles, with an edge between two tiles if they interact. The supertile is τ -*stable* if every cut of its binding graph has strength at least τ , where the weight of an edge is the strength of the glue it represents. That is, the supertile is stable if at least energy τ is required to separate the supertile into two parts. A 2HAM *tile assembly system* (TAS) is a pair $\mathcal{T} = (T, \tau)$, where T is a finite tile set and τ is the *temperature*, usually 1 or 2. Given a TAS $\mathcal{T} = (T, \tau)$, a supertile is *producible*, written as $\alpha \in \mathcal{A}[\mathcal{T}]$ if either it is a single tile from T , or it is the τ -stable result of translating two producible assemblies without overlap. A supertile α is *terminal*, written as $\alpha \in \mathcal{A}_{\square}[\mathcal{T}]$ if for every producible supertile β , α and β cannot be τ -stably attached. A TAS is *directed* if it has only one terminal, producible supertile.²

2.2 Formal definition of the 2HAM

We now formally define the 2HAM.

Two assemblies α and β are *disjoint* if $\text{dom } \alpha \cap \text{dom } \beta = \emptyset$. For two assemblies α and β , define the *union* $\alpha \cup \beta$ to be the assembly defined for all $\vec{x} \in \mathbb{Z}^2$ by $(\alpha \cup \beta)(\vec{x}) = \alpha(\vec{x})$ if $\alpha(\vec{x})$ is defined, and $(\alpha \cup \beta)(\vec{x}) = \beta(\vec{x})$ otherwise. Say that this union is *disjoint* if α and β are disjoint.

The *binding graph* of an assembly α is the grid graph $G_{\alpha} = (V, E)$, where $V = \text{dom } \alpha$, and $\{\vec{m}, \vec{n}\} \in E$ if and only if (1) $\vec{m} - \vec{n} \in U_2$, (2) $\text{label}_{\alpha(\vec{m})}(\vec{n} - \vec{m}) = \text{label}_{\alpha(\vec{n})}(\vec{m} - \vec{n})$, and (3) $\text{str}_{\alpha(\vec{m})}(\vec{n} - \vec{m}) > 0$. Given $\tau \in \mathbb{N}$, an assembly is τ -*stable* (or simply *stable* if τ is understood from context), if it cannot be broken up into smaller assemblies without breaking bonds of total strength at least τ ; i.e., if every cut of G_{α} has weight at least τ , where the weight of an edge is the strength of the glue it represents. In contrast to the model of Wang tiling, the nonnegativity of the strength function implies that glue mismatches between adjacent tiles do not prevent a tile from binding to an assembly, so long as sufficient binding strength is received from the (other) sides of the tile at which the glues match.

For assemblies $\alpha, \beta : \mathbb{Z}^2 \dashrightarrow T$ and $\vec{u} \in \mathbb{Z}^2$, we write $\alpha + \vec{u}$ to denote the assembly defined for all $\vec{x} \in \mathbb{Z}^2$ by $(\alpha + \vec{u})(\vec{x}) = \alpha(\vec{x} - \vec{u})$, and write $\alpha \simeq \beta$ if there exists \vec{u} such that $\alpha + \vec{u} = \beta$; i.e., if α is a translation of β . Given two

²We do not use this definition in this paper but have included it for the sake of completeness.

assemblies $\alpha, \beta : \mathbb{Z}^2 \dashrightarrow T$, we say α is a *subassembly* of β , and we write $\alpha \sqsubseteq \beta$, if $S_\alpha \subseteq S_\beta$ and, for all points $p \in S_\alpha$, $\alpha(p) = \beta(p)$. Define the *supertile* of α to be the set $\tilde{\alpha} = \{ \beta \mid \alpha \simeq \beta \}$. A supertile $\tilde{\alpha}$ is τ -*stable* (or simply *stable*) if all of the assemblies it contains are τ -stable; equivalently, $\tilde{\alpha}$ is stable if it contains a stable assembly, since translation preserves the property of stability. Note also that the notation $|\tilde{\alpha}| \equiv |\alpha|$ is the size of the supertile (i.e., number of tiles in the supertile) is well-defined, since translation preserves cardinality (and note in particular that even though we define $\tilde{\alpha}$ as a set, $|\tilde{\alpha}|$ does not denote the cardinality of this set, which is always \aleph_0).

For two supertiles $\tilde{\alpha}$ and $\tilde{\beta}$, and temperature $\tau \in \mathbb{N}$, define the *combination* set $C_{\tilde{\alpha}, \tilde{\beta}}^\tau$ to be the set of all supertiles $\tilde{\gamma}$ such that there exist $\alpha \in \tilde{\alpha}$ and $\beta \in \tilde{\beta}$ such that (1) α and β are disjoint (steric protection), (2) $\gamma \equiv \alpha \cup \beta$ is τ -stable, and (3) $\gamma \in \tilde{\gamma}$. That is, $C_{\tilde{\alpha}, \tilde{\beta}}^\tau$ is the set of all τ -stable supertiles that can be obtained by “attaching” $\tilde{\alpha}$ to $\tilde{\beta}$ stably, with $|C_{\tilde{\alpha}, \tilde{\beta}}^\tau| > 1$ if there is more than one position at which β could attach stably to α .

It is common with seeded assembly to stipulate an infinite number of copies of each tile, but our definition allows for a finite number of tiles as well. Our definition also allows for the growth of infinite assemblies and finite assemblies to be captured by a single definition, similar to the definitions of [19] for seeded assembly.

Given a set of tiles T , define a *state* S of T to be a multiset of supertiles, or equivalently, S is a function mapping supertiles of T to $\mathbb{N} \cup \{\infty\}$, indicating the multiplicity of each supertile in the state. We therefore write $\tilde{\alpha} \in S$ if and only if $S(\tilde{\alpha}) > 0$.

A (*two-handed*) *tile assembly system* (TAS) is an ordered triple $\mathcal{T} = (T, S, \tau)$, where T is a finite set of tile types, S is the *initial state*, and $\tau \in \mathbb{N}$ is the temperature. If not stated otherwise, assume that the initial state S is defined $S(\tilde{\alpha}) = \infty$ for all supertiles $\tilde{\alpha}$ such that $|\tilde{\alpha}| = 1$, and $S(\tilde{\beta}) = 0$ for all other supertiles $\tilde{\beta}$. That is, S is the state consisting of a countably infinite number of copies of each individual tile type from T , and no other supertiles. In such a case we write $\mathcal{T} = (T, \tau)$ to indicate that \mathcal{T} uses the default initial state. For notational convenience we sometimes describe S as a set of supertiles, in which case we actually mean that S is a multiset of supertiles with infinite count of each supertile. We also assume that, in general, unless stated otherwise, the count for any supertile in the initial state is infinite.

Given a TAS $\mathcal{T} = (T, S, \tau)$, define an *assembly sequence* of \mathcal{T} to be a sequence of states $\vec{S} = (S_i \mid 0 \leq i < k)$ (where $k = \infty$ if \vec{S} is an infinite assembly sequence), and S_{i+1} is constrained based on S_i in the following way: There exist supertiles $\tilde{\alpha}, \tilde{\beta}, \tilde{\gamma}$ such that (1) $\tilde{\gamma} \in C_{\tilde{\alpha}, \tilde{\beta}}^\tau$, (2) $S_{i+1}(\tilde{\gamma}) = S_i(\tilde{\gamma}) + 1$,³ (3) if $\tilde{\alpha} \neq \tilde{\beta}$, then $S_{i+1}(\tilde{\alpha}) = S_i(\tilde{\alpha}) - 1$, $S_{i+1}(\tilde{\beta}) = S_i(\tilde{\beta}) - 1$, otherwise if $\tilde{\alpha} = \tilde{\beta}$, then $S_{i+1}(\tilde{\alpha}) = S_i(\tilde{\alpha}) - 2$, and (4) $S_{i+1}(\tilde{\omega}) = S_i(\tilde{\omega})$ for all $\tilde{\omega} \notin \{\tilde{\alpha}, \tilde{\beta}, \tilde{\gamma}\}$. That is, S_{i+1} is obtained from S_i by picking two supertiles from S_i that can attach to each other, and attaching them, thereby decreasing the count of the two reactant

³with the convention that $\infty = \infty + 1 = \infty - 1$

supertiles and increasing the count of the product supertile. If $S_0 = S$, we say that \vec{S} is *nascent*.

Given an assembly sequence $\vec{S} = (S_i \mid 0 \leq i < k)$ of $\mathcal{T} = (T, S, \tau)$ and a supertile $\tilde{\gamma} \in S_i$ for some i , define the *predecessors* of $\tilde{\gamma}$ in \vec{S} to be the multiset $\text{pred}_{\vec{S}}(\tilde{\gamma}) = \{\tilde{\alpha}, \tilde{\beta}\}$ if $\tilde{\alpha}, \tilde{\beta} \in S_{i-1}$ and $\tilde{\alpha}$ and $\tilde{\beta}$ attached to create $\tilde{\gamma}$ at step i of the assembly sequence, and define $\text{pred}_{\vec{S}}(\tilde{\gamma}) = \{\tilde{\gamma}\}$ otherwise. Define the *successor* of $\tilde{\gamma}$ in \vec{S} to be $\succ_{\vec{S}}(\tilde{\gamma}) = \tilde{\alpha}$ if $\tilde{\gamma}$ is one of the predecessors of $\tilde{\alpha}$ in \vec{S} , and define $\succ_{\vec{S}}(\tilde{\gamma}) = \tilde{\gamma}$ otherwise. A sequence of supertiles $\vec{\alpha} = (\tilde{\alpha}_i \mid 0 \leq i < k)$ is a *supertile assembly sequence* of \mathcal{T} if there is an assembly sequence $\vec{S} = (S_i \mid 0 \leq i < k)$ of \mathcal{T} such that, for all $1 \leq i < k$, $\succ_{\vec{S}}(\tilde{\alpha}_{i-1}) = \tilde{\alpha}_i$, and $\vec{\alpha}$ is *nascent* if \vec{S} is nascent.

The *result* of a supertile assembly sequence $\vec{\alpha}$ is the unique supertile $\text{res}(\vec{\alpha})$ such that there exist an assembly $\alpha \in \text{res}(\vec{\alpha})$ and, for each $0 \leq i < k$, assemblies $\alpha_i \in \tilde{\alpha}_i$ such that $\text{dom } \alpha = \bigcup_{0 \leq i < k} \text{dom } \alpha_i$ and, for each $0 \leq i < k$, $\alpha_i \sqsubseteq \alpha$. For all supertiles $\tilde{\alpha}, \tilde{\beta}$, we write $\tilde{\alpha} \rightarrow_{\mathcal{T}} \tilde{\beta}$ (or $\tilde{\alpha} \rightarrow \tilde{\beta}$ when \mathcal{T} is clear from context) to denote that there is a supertile assembly sequence $\vec{\alpha} = (\tilde{\alpha}_i \mid 0 \leq i < k)$ such that $\tilde{\alpha}_0 = \tilde{\alpha}$ and $\text{res}(\vec{\alpha}) = \tilde{\beta}$. It can be shown using the techniques of [29] for seeded systems that for all two-handed tile assembly systems \mathcal{T} supplying an infinite number of each tile type, $\rightarrow_{\mathcal{T}}$ is a transitive, reflexive relation on supertiles of \mathcal{T} . We write $\tilde{\alpha} \rightarrow_{\mathcal{T}}^1 \tilde{\beta}$ ($\tilde{\alpha} \rightarrow^1 \tilde{\beta}$) to denote an assembly sequence of length 1 from $\tilde{\alpha}$ to $\tilde{\beta}$ and $\tilde{\alpha} \rightarrow_{\mathcal{T}}^{\leq 1} \tilde{\beta}$ ($\tilde{\alpha} \rightarrow^{\leq 1} \tilde{\beta}$) to denote an assembly sequence of length 1 from $\tilde{\alpha}$ to $\tilde{\beta}$ if $\tilde{\alpha} \neq \tilde{\beta}$ and an assembly sequence of length 0 otherwise.

A supertile $\tilde{\alpha}$ is *producible*, and we write $\tilde{\alpha} \in \mathcal{A}[\mathcal{T}]$, if it is the result of a nascent supertile assembly sequence. A supertile $\tilde{\alpha}$ is *terminal* if, for all producible supertiles $\tilde{\beta}$, $C_{\tilde{\alpha}, \tilde{\beta}}^{\tau} = \emptyset$.⁴ Define $\mathcal{A}_{\square}[\mathcal{T}] \subseteq \mathcal{A}[\mathcal{T}]$ to be the set of terminal and producible supertiles of \mathcal{T} . \mathcal{T} is *directed* (a.k.a., *deterministic*, *confluent*) if $|\mathcal{A}_{\square}[\mathcal{T}]| = 1$.

2.3 Definitions for simulation

In this subsection, we formally define what it means for one 2HAM TAS to “simulate” another 2HAM TAS. For a tileset T , let A^T and \tilde{A}^T denote the set of all assemblies over T and all supertiles over T respectively. Let $A_{<\infty}^T$ and $\tilde{A}_{<\infty}^T$ denote the set of all finite assemblies over T and all finite supertiles over T respectively.

In what follows, let U be a tile set. An *m-block assembly*, or *macrotile*, over tile set U is a partial function $\gamma : \mathbb{Z}_m \times \mathbb{Z}_m \dashrightarrow U$, where $\mathbb{Z}_m = \{0, 1, \dots, m-1\}$. Let B_m^U be the set of all m -block assemblies over U . The m -block with no

⁴Note that a supertile $\tilde{\alpha}$ could be non-terminal in the sense that there is a producible supertile $\tilde{\beta}$ such that $C_{\tilde{\alpha}, \tilde{\beta}}^{\tau} \neq \emptyset$, yet it may not be possible to produce $\tilde{\alpha}$ and $\tilde{\beta}$ simultaneously if some tile types are given finite initial counts, implying that $\tilde{\alpha}$ cannot be “grown” despite being non-terminal. If the count of each tile type in the initial state is ∞ , then all producible supertiles are producible from any state, and the concept of terminal becomes synonymous with “not able to grow”, since it would always be possible to use the abundant supply of tiles to assemble $\tilde{\beta}$ alongside $\tilde{\alpha}$ and then attach them.

domain is said to be *empty*. For an arbitrary assembly $\alpha \in A^U$ define $\alpha_{x,y}^m$ to be the m -block defined by $\alpha_{x,y}^m(i, j) = \alpha(mx + i, my + j)$ for $0 \leq i, j < m$.

For a partial function $R : B_m^U \dashrightarrow T$, define the *assembly representation function* $R^* : A^U \dashrightarrow A^T$ such that $R^*(\alpha) = \beta$ if and only if $\beta(x, y) = R(\alpha_{x,y}^m)$ for all $x, y \in \mathbb{Z}^2$. Further, α is said to map *cleanly* to β under R^* if either (1) for all non empty blocks $\alpha_{x,y}^m$, $(x+u, y+v) \in \text{dom } \beta$ for some $u, v \in \{-1, 0, 1\}$ such that $u^2 + v^2 < 2$, or (2) α has at most one non-empty m -block $\alpha_{x,y}^m$. In other words, we allow for the existence of simulator “fuzz” directly north, south, east or west of a simulator macrotile, but we exclude the possibility of diagonal fuzz.

For a given *assembly representation function* R^* , define the *supertile representation function* $\tilde{R} : \tilde{A}^U \dashrightarrow \mathcal{P}(A^T)$ such that $\tilde{R}(\tilde{\alpha}) = \{R^*(\alpha) \mid \alpha \in \tilde{\alpha}\}$. $\tilde{\alpha}$ is said to map *cleanly* to $\tilde{R}(\tilde{\alpha})$ if $\tilde{R}(\tilde{\alpha}) \in \tilde{A}^T$ and α maps cleanly to $R^*(\alpha)$ for all $\alpha \in \tilde{\alpha}$.

In the following definitions, let $\mathcal{T} = (T, S, \tau)$ be a 2HAM TAS and, for some initial configuration $S_{\mathcal{T}}$, that depends on \mathcal{T} , let $\mathcal{U} = (U, S_{\mathcal{T}}, \tau')$ be a 2HAM TAS, and let R be an m -block representation function $R : B_m^U \dashrightarrow T$.

Definition 2.1. We say that \mathcal{U} and \mathcal{T} have equivalent productions (at scale factor m), and we write $\mathcal{U} \Leftrightarrow_R \mathcal{T}$ if the following conditions hold:

1. $\{\tilde{R}(\tilde{\alpha}) \mid \tilde{\alpha} \in \mathcal{A}[\mathcal{U}]\} = \mathcal{A}[\mathcal{T}]$.
2. For all $\tilde{\alpha} \in \mathcal{A}[\mathcal{U}]$, $\tilde{\alpha}$ maps cleanly to $\tilde{R}(\tilde{\alpha})$

Definition 2.2. We say that \mathcal{T} follows \mathcal{U} (at scale factor m), and we write $\mathcal{T} \dashv_R \mathcal{U}$ if, for any $\tilde{\alpha}, \tilde{\beta} \in \mathcal{A}[\mathcal{U}]$ such that $\tilde{\alpha} \rightarrow_{\mathcal{U}}^1 \tilde{\beta}$, $\tilde{R}(\tilde{\alpha}) \rightarrow_{\mathcal{T}}^{\leq 1} \tilde{R}(\tilde{\beta})$.

Definition 2.3. We say that \mathcal{U} weakly models \mathcal{T} (at scale factor m), and we write $\mathcal{U} \models_{\tilde{R}}^- \mathcal{T}$ if, for any $\tilde{\alpha}, \tilde{\beta} \in \mathcal{A}[\mathcal{T}]$ such that $\tilde{\alpha} \rightarrow_{\mathcal{T}}^1 \tilde{\beta}$, for all $\tilde{\alpha}' \in \mathcal{A}[\mathcal{U}]$ such that $\tilde{R}(\tilde{\alpha}') = \tilde{\alpha}$, there exists an $\tilde{\alpha}'' \in \mathcal{A}[\mathcal{U}]$ such that $\tilde{R}(\tilde{\alpha}'') = \tilde{\beta}$, $\tilde{\alpha}' \rightarrow_{\mathcal{U}} \tilde{\alpha}''$, and $\tilde{\alpha}'' \rightarrow_{\mathcal{U}}^1 \tilde{\beta}'$ for some $\tilde{\beta}' \in \mathcal{A}[\mathcal{U}]$ with $\tilde{R}(\tilde{\beta}') = \tilde{\beta}$.

Definition 2.4. We say that \mathcal{U} strongly models \mathcal{T} (at scale factor m), and we write $\mathcal{U} \models_{\tilde{R}}^+ \mathcal{T}$ if for any $\tilde{\alpha}, \tilde{\beta} \in \mathcal{A}[\mathcal{T}]$ such that $\tilde{\gamma} \in C_{\tilde{\alpha}, \tilde{\beta}}^{\tau}$, then for all $\tilde{\alpha}', \tilde{\beta}' \in \mathcal{A}[\mathcal{U}]$ such that $\tilde{R}(\tilde{\alpha}') = \tilde{\alpha}$ and $\tilde{R}(\tilde{\beta}') = \tilde{\beta}$, it must be that there exist $\tilde{\alpha}'', \tilde{\beta}'', \tilde{\gamma}' \in \mathcal{A}[\mathcal{U}]$, such that $\tilde{\alpha}' \rightarrow_{\mathcal{U}} \tilde{\alpha}''$, $\tilde{\beta}' \rightarrow_{\mathcal{U}} \tilde{\beta}''$, $\tilde{R}(\tilde{\alpha}'') = \tilde{\alpha}$, $\tilde{R}(\tilde{\beta}'') = \tilde{\beta}$, $\tilde{R}(\tilde{\gamma}') = \tilde{\gamma}$, and $\tilde{\gamma}' \in C_{\tilde{\alpha}'', \tilde{\beta}''}^{\tau'}$.

Definition 2.5. Let $\mathcal{U} \Leftrightarrow_R \mathcal{T}$ and $\mathcal{T} \dashv_R \mathcal{U}$.

1. \mathcal{U} simulates \mathcal{T} (at scale factor m) if $\mathcal{U} \models_{\tilde{R}}^- \mathcal{T}$.
2. \mathcal{U} strongly simulates \mathcal{T} (at scale factor m) if $\mathcal{U} \models_{\tilde{R}}^+ \mathcal{T}$.

For simulation, we require that when a simulated supertile $\tilde{\alpha}$ may grow, via one combination attachment, into a second supertile $\tilde{\beta}$, then any simulator

supertile that maps to $\tilde{\alpha}$ must also grow into a simulator supertile that maps to $\tilde{\beta}$. The converse should also be true.

For strong simulation, in addition to requiring that all supertiles mapping to $\tilde{\alpha}$ must be capable of growing into a supertile mapping to $\tilde{\beta}$ when $\tilde{\alpha}$ can grow into $\tilde{\beta}$ in the simulated system, we further require that this growth can take place by the attachment of *any* supertile mapping to $\tilde{\gamma}$, where $\tilde{\gamma}$ is the supertile that attaches to $\tilde{\alpha}$ to get $\tilde{\beta}$.

2.4 Intrinsic universality

Let REPR denote the set of all m -block (or macrotile) representation functions. Let \mathfrak{C} be a class of tile assembly systems, and let U be a tile set. We say U is *intrinsically universal* for \mathfrak{C} if there are computable functions $\mathcal{R} : \mathfrak{C} \rightarrow \text{REPR}$ and $\mathcal{S} : \mathfrak{C} \rightarrow (A_{<\infty}^U \rightarrow \mathbb{N} \cup \{\infty\})$, and a $\tau' \in \mathbb{Z}^+$ such that, for each $\mathcal{T} = (T, S, \tau) \in \mathfrak{C}$, there is a constant $m \in \mathbb{N}$ such that, letting $R = \mathcal{R}(\mathcal{T})$, $S_{\mathcal{T}} = \mathcal{S}(\mathcal{T})$, and $\mathcal{U}_{\mathcal{T}} = (U, S_{\mathcal{T}}, \tau')$, $\mathcal{U}_{\mathcal{T}}$ simulates \mathcal{T} at scale m and using macrotile representation function R . That is, $\mathcal{R}(\mathcal{T})$ gives a representation function R that interprets macrotiles (or m -blocks) of $\mathcal{U}_{\mathcal{T}}$ as assemblies of \mathcal{T} , and $\mathcal{S}(\mathcal{T})$ gives the initial state used to create the necessary macrotiles from U to represent \mathcal{T} subject to the constraint that no macrotile in $S_{\mathcal{T}}$ can be larger than a single $m \times m$ square.

3 The 2HAM is not intrinsically universal

In this section, we prove the main result of this paper: there is no universal 2HAM tile set that, when appropriately initialized, is capable of simulating an arbitrary 2HAM system. That is, we prove that the 2HAM, unlike the aTAM, is not intrinsically universal.

Theorem 3.1. *The 2HAM is not intrinsically universal.*

We first prove Theorem 3.2, which says that, for any claimed 2HAM simulator \mathcal{U} , that runs at temperature τ' , there exists a 2HAM system, with temperature $\tau > \tau'$, that cannot be simulated by \mathcal{U} . We use this as the main tool to prove Theorem 3.3, a restatement of Theorem 3.1; our main result.

Theorem 3.2. *Let $\tau \in \mathbb{N}, \tau \geq 2$. For every tile set U , there exists a 2HAM TAS $\mathcal{T} = (T, S, \tau)$ such that for any initial configuration $S_{\mathcal{T}}$ over U and $\tau' \leq \tau - 1$, the 2HAM TAS $\mathcal{U} = (U, S_{\mathcal{T}}, \tau')$ does not simulate \mathcal{T} .*

The basic idea of the proof of Theorem 3.2 is to use Definitions 2.3 and 2.1 in order to exhibit two producible supertiles in \mathcal{T} , that do not combine in \mathcal{T} because of a lack of total binding strength, and show that the supertiles that simulate them in \mathcal{U} do combine in the (lower temperature) simulator \mathcal{U} . Then we argue that Definition 2.2 says that, because the simulating supertiles can combine in the simulator \mathcal{U} , then so too can the supertiles being simulated in the simulated system \mathcal{T} , which contradicts the fact that the two originally chosen supertiles from \mathcal{T} do not combine in \mathcal{T} .

Proof. Our proof is by contradiction. Therefore, suppose, for the sake of obtaining a contradiction, that there exists a universal tile set U such that, for any 2HAM TAS $\mathcal{T} = (T, S, \tau)$, there exists an initial configuration $S_{\mathcal{T}}$ and $\tau' \leq \tau - 1$, such that $\mathcal{U} = (U, S_{\mathcal{T}}, \tau')$ simulates \mathcal{T} . Define $\mathcal{T} = (T, \tau)$ where T is the tile set defined in Figure 1, the default initial state is used, and $\tau > 1$. Let $\mathcal{U} = (U, S_{\mathcal{T}}, \tau')$ be the temperature $\tau' \leq \tau - 1$ 2HAM system, which uses tile set U and initial configuration $S_{\mathcal{T}}$ (depending on \mathcal{T}) to simulate \mathcal{T} at scale factor m . Let \tilde{R} denote the assembly replacement function that testifies to the fact that \mathcal{U} simulates \mathcal{T} .




Figure 1: The tile set for the proof of Theorem 3.2. Black rectangles represent strength- τ glues (labeled 1-8), and black squares represent the strength-1 glue (labeled 0).

We say that a supertile $\tilde{l} \in \mathcal{A}[\mathcal{T}]$ is a *left half-ladder* of height $h \in \mathbb{N}$ if it contains h tiles of the type A2 and $h - 1$ tiles of type A3, arranged in a vertical column, plus τ tiles of each of the types A1 and A0. (An example of a left half-ladder is shown on the left in Figure 2. The dotted lines show positions at which tiles of type A1 and A0 could potentially attach, but since a half-ladder has exactly τ of each, only τ such locations have tiles.) Essentially, a left half-ladder consists of a single-tile-wide vertical column of height $2h - 1$ with an A2 tile at the bottom and top, and those in between alternating between A3 and A2 tiles. To the east of exactly τ of the A2 tiles, an A1 tile is attached and to the east of each A1 tile, an A0 tile type is attached. These A1-A0 pairs, collectively, form the τ *rungs* of the left half-ladder. We can define *right* half-ladders similarly. A *right half-ladder* of height h is defined exactly the same way but using the tile types B3, B2, B1, and B0 and with rungs growing to the left of the vertical column. The east glue of A0 is a strength-1 glue matching the west glue of B0.

Let $LEFT \subseteq \mathcal{A}[\mathcal{T}]$ and $RIGHT \subseteq \mathcal{A}[\mathcal{T}]$ be the set of all left and right half-ladders of height h , respectively. Note that there are $\binom{h}{\tau}$ half-ladders of height h in $LEFT$ ($RIGHT$). Define, for each $\tilde{l} \in LEFT$, the *mirror image* of \tilde{l} as the supertile $\tilde{\bar{l}} \in RIGHT$ such that $\tilde{\bar{l}}$ has rungs at the same positions as \tilde{l} .

For some $\tilde{l} \in LEFT$, we say that $\tilde{\hat{l}} \in \mathcal{A}[\mathcal{U}]$ is a *simulator* left half-ladder of height h if $\tilde{R}(\tilde{\hat{l}}) = \tilde{l}$. Note that $\tilde{\hat{l}}$ need not be unique. (One could even imagine $\tilde{\hat{l}}$ and $\tilde{\hat{l}'}$ satisfying $\tilde{R}(\tilde{\hat{l}}) = \tilde{l}$ and $\tilde{R}(\tilde{\hat{l}'}) = \tilde{l}$ but $\tilde{\hat{l}}$ and $\tilde{\hat{l}'}$ only differ by a single tile!) The notation $C_{\tilde{\alpha}, \tilde{\beta}}^{\tau}$ is defined as the set of all supertiles that result in the τ -stable combination of the supertiles $\tilde{\alpha}$ and $\tilde{\beta}$.

For some $\tilde{\hat{r}} \in \mathcal{A}[\mathcal{U}]$, we say that $\tilde{\hat{r}}$ is a *mate* of $\tilde{\hat{l}}$ if $\tilde{R}(\tilde{\hat{r}}) = \tilde{\hat{r}} \in RIGHT$, where

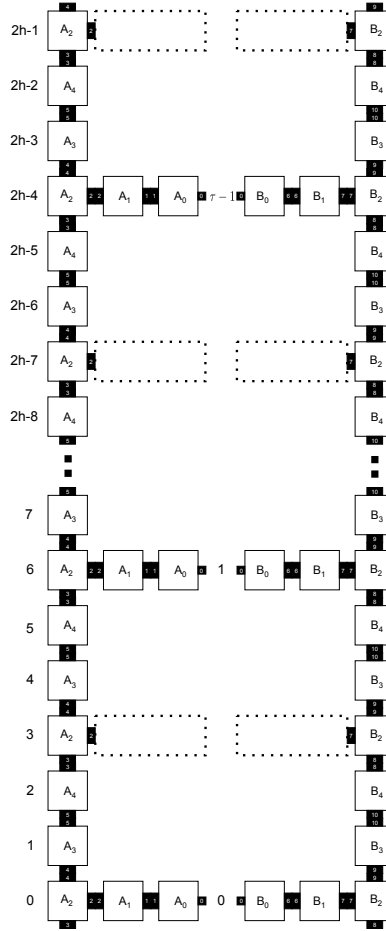


Figure 2: Example half-ladders with τ rungs.

$\tilde{r} = \tilde{l}$, $C_{\tilde{l}, \tilde{r}}^\tau \neq \emptyset$ (they combine in \mathcal{T}), and $C_{\tilde{l}, \tilde{r}}^{\tau-1} \neq \emptyset$ (they combine in \mathcal{U}). For a simulator left half-ladder \tilde{l} , we say that \tilde{l} is *combinable* if \tilde{l} has a mate. Part 1 of Definition 2.5 guarantees the existence of at least one combinable simulator left half-ladder for each left half-ladder. It is easy to see from Part 1 of Definition 2.5 that an arbitrary simulator left half-ladder need not be combinable, since by Definition 2.3, it may be a half-ladder $\tilde{l} \in \mathcal{A}[\mathcal{U}]$, which must first “grow into” a combinable left half-ladder \tilde{l}' (analogous to $\tilde{\alpha}' \rightarrow_{\mathcal{U}} \tilde{\alpha}''$ in Definition 2.3).

Denote as $LEFT'$ some set that contains exactly one combinable simulator left half-ladder for each $\tilde{l} \in LEFT$. Note that, by Definitions 2.1 and 2.3, there must be at least one combinable simulator left half-ladder \tilde{l} for each \tilde{l} , but that there also may be more than one, so the set $LEFT'$, while certainly

not empty, need not be unique. By the definition of $LEFT'$, it is easy to see that $|LEFT'| = \binom{h}{\tau}$. We know that each combinable simulator left half-ladder \tilde{l} has exactly τ rungs, and furthermore, since glue strengths in the 2HAM cannot be fractional, it is the case that τ' of these rungs bind to (the corresponding rungs of) a mate with a combined total strength of at least τ' . (Note that some, but not all, of these τ' rungs may be redundant in the sense that they do not interact with positive strength.)

There are $\binom{h}{\tau'}$ ways to position/choose τ' rungs on a (simulator) half-ladder of height h . (Note that a rung on a simulator half-ladder need not be a $m \times m$ block of tiles but merely a collection of rung-like blocks that map to rungs in the input system \mathcal{T} via \tilde{R} .) Now consider the size $\binom{h}{\tau}$ set of all possible rung positions, each denoted by a subset $X \subset \{0, 1, \dots, h-1\}$, and the size $\binom{h}{\tau}$ set $LEFT'$. For each simulated half-ladder $\tilde{l} \in LEFT'$, there must exist a set of τ' rungs X such that \tilde{l} binds to a mate via the rungs specified by X , with total strength at least τ' . As there are $\binom{h}{\tau}$ elements of $LEFT'$ and only $\binom{h}{\tau'}$ choices for X , the Generalized Pigeonhole Principle implies that there must be some set $LEFT'' \subset LEFT'$ with $|LEFT''| \geq \binom{h}{\tau} / \binom{h}{\tau'}$ such that every simulator left half-ladder in $LEFT''$ binds to a mate via the τ' rungs specified by a single choice of X , with total strength at least τ' . In the case that $h \geq 2\tau$, we have that $|LEFT''| \geq \binom{h}{\tau} / \binom{h}{\tau'} \geq \binom{h}{\tau} / \binom{h}{\tau-1} = \frac{h-\tau+1}{\tau}$.

Let $k = |U|^{4m^2}$, which is the number of ways to tile a neighborhood of four $m \times m$ squares from a set of $|U|$ distinct tile types. If $h = \tau(k^{\tau-1} + \tau)$, then $|LEFT''| \geq k^{\tau-1} + 1$. There are $k^{\tau'} \leq k^{\tau-1}$ ways to tile τ' neighborhoods that map to tiles of type A0 (plus any additional simulator fuzz that connects to simulated A0 tiles), under \tilde{R} , at the ends of the τ' rungs of a simulator left half-ladder. This tells us that there are at least two (combinable) simulator left half-ladders $\tilde{l}_1, \tilde{l}_2 \in LEFT''$ such that \tilde{l}_1 binds to a mate via the rungs specified by X , with total strength at least τ' , \tilde{l}_2 binds to a mate via the rungs specified by X , with total strength at least τ' and the rungs (along with any surrounding fuzz) specified by X of \tilde{l}_1 are tiled exactly the same as the rungs specified by X of \tilde{l}_2 are tiled. Thus, we can conclude that \tilde{r} , a mate of \tilde{l}_1 , is a mate of \tilde{l}_2 . We can conclude this because, while \tilde{l}_1 and \tilde{l}_2 agree exactly along τ' of their rungs, they also each have one rung in a unique position and since consecutive rungs in \mathcal{T} have at least two empty spaces between them, the offset simulator rungs (and even their fuzz) cannot prevent \tilde{l}_2 from matching up with the mate of \tilde{l}_1 .

However, $\tilde{R}(\tilde{r}) = \tilde{r} \in R$, $\tilde{R}(\tilde{l}_2) = \tilde{l}_2 \in L$ but $C_{\tilde{r}, \tilde{l}_2}^\tau = \emptyset$ because \tilde{r} and \tilde{l}_2 differ from each other in one rung location and therefore interact in \mathcal{T} with total strength at most $\tau - 1$. This is a contradiction to Definition 2.2, which implies $C_{\tilde{r}, \tilde{l}_2}^\tau \neq \emptyset$. \square

We now have the main tool needed to prove our main Theorem 3.1, which we restate as follows.

Theorem 3.3. *There is no universal tile set U for the 2HAM, i.e., there is no U such that, for all 2HAM tile assembly systems $\mathcal{T} = (T, S, \tau)$, there exists an initial configuration $S_{\mathcal{T}}$ over U and temperature τ' such that $\mathcal{U} = (U, S_{\mathcal{T}}, \tau')$ simulates \mathcal{T} .*

Proof. Our proof is by contradiction, so assume that U is a universal tile set. Denote as g the strength of the strongest glue on any tile type in U . Let $\mathcal{T}' = (T', 4g + 1)$ be a modified version of the TAS $\mathcal{T} = (T, \tau)$ from the proof of Theorem 3.2 with each τ -strength glue in T converted to a strength $4g + 1$ glue in T' (all other glues and labels are unmodified). By the proof of Theorem 3.2, we know that for any initial configuration $S_{\mathcal{T}}$ over U , $\mathcal{U} = (U, S_{\mathcal{T}}, \tau')$ does not simulate \mathcal{T} for any $\tau' < 4g + 1$. If $\tau' \geq 4g + 1$, then the size of the largest supertile in $\mathcal{A}[\mathcal{U}]$ is 1 (since g is the maximum glue strength in U , the supertiles in the initial state (input) $S_{\mathcal{T}}$ are not τ' -stable and indeed no tile can bind to any assembly with strength $\geq 4g + 1$), whence U is not a universal tile set. \square

4 The temperature- τ 2HAM is intrinsically universal

In this section we state our second main result, which states that for fixed temperature $\tau \geq 2$ the class of 2HAM systems at temperature τ is intrinsically universal. In other words, for such τ there is a tile set that, when appropriately initialized, simulates any temperature τ 2HAM system. Denote as $2\text{HAM}(k)$ the set of all 2HAM systems at temperature k .

Theorem 4.1. *For all $\tau \geq 2$, $2\text{HAM}(\tau)$ is intrinsically universal.*

We prove this theorem for two different, but seemingly natural notions of simulation. The first, simply called *simulation*, is where we require that when a simulated supertile $\tilde{\alpha}$ may grow, via one attachment, into a second supertile $\tilde{\beta}$, then any simulator supertile that maps to $\tilde{\alpha}$ must also grow into a simulator supertile that maps to $\tilde{\beta}$. The converse should also be true. Results for *simulation* are given in Section 6.

The second notion, called *strong simulation*, is a stricter definition where in addition to requiring that all supertiles mapping to $\tilde{\alpha}$ must be capable of growing into a supertile mapping to $\tilde{\beta}$ when $\tilde{\alpha}$ can grow into $\tilde{\beta}$ in the simulated system, we further require that this growth can take place by the attachment of *any* supertile mapping to $\tilde{\gamma}$, where $\tilde{\gamma}$ is the supertile that attaches to $\tilde{\alpha}$ to get $\tilde{\beta}$. Results for *strong simulation* are given in Section 5.

For each of the two notions of simulation we provide three results, and in all cases we provide lower scale factor for simulation relative to strong simulation. Specifically, strong simulation achieves a modest linear scale factor simulation, but a compact single input assembly is sufficient to encode the entire simulated system. In contrast, for simulation (i.e. not strong), we are able to achieve a logarithmic scale factor in the size of the simulated system. However, such small scale requires that the simulated system be encoded in a larger (linear) number

of input assemblies, in order to describe the simulated system without loss of information.

For strong simulation, Theorems 5.1, 5.2, and 5.3 provide three different proofs of our main positive result (Theorem 4.1) with each of the three providing different trade-offs between number tile types, scale factor, and complexity of initial configuration for the simulator. For simulation, Theorems 6.1, 6.2 and 6.3, provide similar trade-offs.

When we combine our negative and positive results, we get a separation between classes of 2HAM tile systems based on their temperatures.

Theorem 4.2. *There exists an infinite number of infinite hierarchies of 2HAM systems with strictly-increasing power (and temperature) that can simulate downward within their own hierarchy.*

Proof. Our first main result (Theorem 3.2) tells us that the temperature- τ 2HAM cannot be simulated by any temperature $\tau' < \tau$ 2HAM. Hence we have, for all $i > 0, c \geq 4$, $2\text{HAM}(c^i) \succ 2\text{HAM}(c^{i-1})$, where \succ is the relation “cannot be simulated by”. Moreover, Theorem 4.1 tells us that temperature τ 2HAM is intrinsically universal for fixed temperature τ . Suppose that $\tau' < \tau$ such that $\tau/\tau' \in \mathbb{N}$. Then the temperature- τ 2HAM can simulate temperature τ' (by simulating strength $g \leq \tau'$ attachments in the temperature τ' system with strength $g\tau/\tau'$ attachments in the temperature τ system). Thus, for all $0 < i' \leq i$, $2\text{HAM}(c^i)$ can simulate, via Theorem 4.1, $2\text{HAM}(c^{i'})$. The theorem follows by noting that our choice of c was arbitrary. \square

We have shown that for each $\tau \geq 2$ there exists a single set of tile types U_τ , and a set of input supertiles over U_τ , such that the 2HAM system strongly simulates any 2HAM TAS \mathcal{T} . A related question is: does there exist a tile set that can simulate, or strongly simulate, all temperature τ 2HAM TASs simultaneously? Surprisingly, the answer is yes! The proof of the following theorem is given in Section 7.

Theorem 4.3. *For each $\tau > 1$, there exists a 2HAM system $\mathcal{S} = (U_\tau, \tau)$ which simultaneously strongly simulates all 2HAM systems $\mathcal{T} = (T, \tau)$.*

5 The temperature- τ 2HAM is intrinsically universal: strong simulation

We present a total of six simulation results, three in this section and three in Section 6. In this section the three simulation results exhibit for any integer $\tau \geq 2$, a single set of tiles U_τ that at temperature τ *strongly simulates* any temperature τ 2HAM system, given a proper configuration of initial assemblies over U_τ . Each of the three results in this section depict different tradeoffs between the number of encoded input supertiles and the scale of the simulation. Theorem 5.3 achieves this while having the simulated tile set encoded as a single input assembly.

For the aTAM it is known [12] that there is a single tile set U that simulates any aTAM tile assembly system \mathcal{T} , when initialized with a *single* seed assembly $\sigma_{\mathcal{T}}$ that encodes \mathcal{T} . Assembly proceeds by additions of single tiles to this seed. In this paper, where we study the 2HAM, it makes sense to allow the simulator to be programmed with multiple copies of the seed (input), rather than a single copy. In particular, this is the case in Theorem 5.3 for strong simulation (and thus, also the weaker notion of *simulation*) where the simulator’s input consists of infinitely many copies of both a single seed supertile, as well as the simulator’s tiles. However, the definition of input configuration allows fancier input configurations: it permits us to have numerous distinct seed assemblies. By exploiting this we achieve better scaling in Theorems 5.1, 5.2, 6.1, 6.2, and 6.3 than in the single-seed case of Theorem 5.3 (the six results also give trade-offs in numbers of tile types in the simulator). However, since these improvements in scaling (and possibly number of tile types) come at the expense of having many seed assemblies in the simulator, there is an intuitive sense in which “less” self-assembly, or at least a different form of self-assembly, is happening as fewer, and larger, assemblies themselves can act as large polyomino jigsaw pieces that come together to simulate tiles. It is worth pointing out that our main result, an impossibility result, holds *despite* the fact that the simulator may try to use a large number of complicated-looking input assemblies.

Let $\|M\|$ denote the number of distinct elements in the multiset M , i.e. $\|M\| = |M'|$ is the cardinality of the set M' defined by ignoring multiplicities in the multiset M . Let T_{sup} denote the set of supertiles induced by a tile set T . By this, we simply mean T_{sup} is the set of supertiles formed by taking all tiles in T and translating them to all locations in \mathbb{Z}^2 .

5.1 Strong simulation with small scale and few tile types

Theorem 5.1. *For every $\tau \geq 2$ there exists a single set of tile types U_{τ} , with $|U_{\tau}| = O(1)$ (i.e. independent of τ), such that for all 2HAM systems $\mathcal{T} = (T, \sigma, \tau)$, there is a set $I_{\mathcal{T}}$ of $\|\sigma\|$ input supertiles such that the 2HAM system $\mathcal{U}_{\mathcal{T}} = (U_{\tau}, U_{\tau, \text{sup}} \cup I_{\mathcal{T}}, \tau)$ strongly simulates \mathcal{T} at scale $O(\sqrt{|G|(\tau + \log |G|)})$, where G is the set of glues in T .*

Construction. For a given simulated TAS $\mathcal{T} = (T, \sigma, \tau)$, the TAS $\mathcal{U}_{\mathcal{T}}$ represents the initial state of \mathcal{T} as follows.

Each *singleton* tile type $t \in T$ that is used in the initial state σ is represented as a macrotile of the form shown in Figure 3a (larger input supertiles from σ are described below). Tile type t is mapped to a unique macrotile \hat{t} as follows. First, each glue $g \in G$, from the tile set T , is uniquely encoded as a pair $\hat{g} = (i, j) \in \ell \times \ell$ where $\ell = \{x \in \mathbb{N} \mid x < \lceil \sqrt{|G|} \rceil\}$ (using the inverse of some simple pairing function). Assume that tile type t has the four glues $(g_{\text{north}}, g_{\text{west}}, g_{\text{south}}, g_{\text{east}})$. Next, these glues are encoded as the four coordinates $(\hat{g}_{\text{north}}, \hat{g}_{\text{west}}, \hat{g}_{\text{south}}, \hat{g}_{\text{east}})$ using the above encoding. The macrotile \hat{t} is composed of five parts: a square *body* (of size $k \times k$, where $k = O(\sqrt{|G|(\tau + \log |G|)})$), and four arms so that one arm

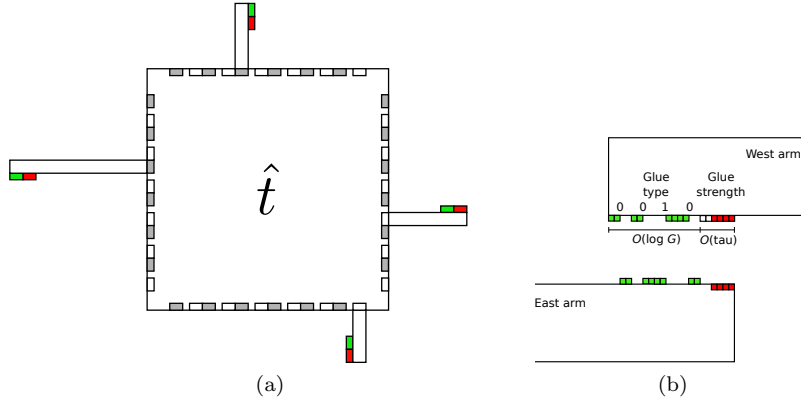


Figure 3: Strong simulation. (a) Input assembly design. For each tile type t in the simulated TAS there is a unique macrotile \hat{t} . Glue pads and binding pads are shown in green and red, respectively. North arms project from grey regions so that they are staggered relative to south arms, west and east arms act similarly. For each unique glue g in the simulated TAS, there is a unique pair of coordinates for the glue/binding pads, one denoting the position along the side of the macrotile (i.e. the particular grey/white location), and one denoting the distance that the arm projects outward from the macrotile. (b) Detail showing two arms with matching glues (encoded as binary bumps and dents using green tiles) of strength 4 (encoded using 4 red tiles), from a temperature $\tau = 6$ system. Each red tile provides unit binding strength, and so macrotiles bind with strength equal to the number of matching red tiles. Only red tiles provide non-zero binding strength.

is placed in each of four square regions, each of size $k \times k$, and adjacent to the body. The location of the red and green glue-binding pad on an arm uniquely encodes the relevant glue of tile type t : for example g_{west} is encoded by the glue-binding pad location $(ki/\lceil\sqrt{|G|}\rceil, kj/\lceil\sqrt{|G|}\rceil)$, where $\hat{g}_{\text{west}} = (i, j)$. (The arm length is $ki/\lceil\sqrt{|G|}\rceil$ and its position along the supertile side is $kj/\lceil\sqrt{|G|}\rceil$.) Arm lengths on east sides are “complimentary” to those on west sides, in the following sense. Let $g = g_{\text{west}} = g_{\text{east}}$ be some glue that appears on the east side of one tile and west side of another. As described above, the arm with the pad \hat{g}_{west} has length such that the glue pad appears at location $ki/\lceil\sqrt{|G|}\rceil$. However, the arm with the pad \hat{g}_{east} is defined to have length such that the glue pad appears at location $k - (ki/\lceil\sqrt{|G|}\rceil)$. The same complementarity trick is used for north and south arms. Finally, as can be seen in Figure 3a, arms are *staggered* relative to each other: north and west arms sit on grey patches, south and east arms sit on white patches.

Figure 3b shows the individual glue-binding pads at the end of two arms. Each glue $g \in G$ is uniquely represented as a bit sequence, which in turn is represented using bumps and dents in a $1 \times 4 \lceil\log |G|\rceil$ region shown in green

on the west arm in Figure 3b. The same bump-dent pattern would be used on a north arm. For east and south arms, the complementary bump-dent pattern is used. To illustrate this, Figure 3b shows a west and east arm that share the same glue type. It can be seen that the arms are able to be translated so that the green regions *fit* together. A glue-binding pad also encodes the binding strength $\text{str}(g) \leq \tau$ of its represented glue g , in a straightforward way as a sequence of $\tau - \text{str}(g)$ white tiles followed by $\text{str}(g)$ red tiles. Each red tile, on the west arm, exposes a single strength 1 glue to its south. The red tiles are the *only* tiles on the entire macrotile \hat{t} that expose positive strength glues. A matching east arm, as shown in Figure 3b, exposes the same number $\text{str}(g)$ of matching strength 1 glues from its $\text{str}(g)$ red tiles.⁵

Due to their complementary green regions, and their matching sequence of exactly $\text{str}(g)$ red tiles, the two arms shown in Figure 3b can be translated so that they bind together with strength $\text{str}(g)$.

This completes the description of the encoding of the *singleton* tile types $t \in T$ that are used in the initial state σ . The remaining supertiles of σ (i.e. of size > 1) are encoded as in the following paragraph.

Consider 2 macrotiles \hat{t}_1, \hat{t}_2 , that represent tiles t_1, t_2 . From the above description, it can be seen that \hat{t}_1 and \hat{t}_2 can be positioned so that their bodies' centers lie on the same horizontal line, at exactly $2k$ distance apart, such that \hat{t}_1, \hat{t}_2 do not intersect. Furthermore, if t_1 and t_2 have a matching glue on their east and west sides, respectively, then the glue-binding pads of the east arm of \hat{t}_1 and the west arm of \hat{t}_2 will be positioned so that their matching bump-dent patterns interlock, and their arms bind with whatever strength t_1 and t_2 bind. Finally, if t_1 and t_2 do not have matching glues on their east and west sides, respectively, then it is the case the glue-binding pads of the two arms do not touch, and indeed their "mismatching arms" do not intersect. This holds for the only other potential binding position (i.e. north-south) of two arbitrary tiles t_1 and t_2 . Figure 5b shows six macrotiles translated into position to simulate an assembly of 6 tiles in some simulated TAS. The figure shows 5 matching arms (simulating matching glues) and two mismatched arms (simulating mismatched glues). Supertiles of size 1 in σ are encoded in this manner.

We have now completely specified the initial state from which the self-assembly process proceeds in $\mathcal{U}_{\mathcal{T}}$.

Scale. Each glue $g \in G$ of the simulated TAS \mathcal{T} is encoded using a glue-binding pad consisting of $O(\log G)$ (green) tiles and τ (red) tiles. There are $|G|$ such glue-binding pads, which are rasterized into each of four $k \times k$ regions, where $k = O(\sqrt{G(\tau + \log G)})$. To see that the tile set size is a constant (i.e. independent of the simulated TAS \mathcal{T}), note that all tiles on the outside of the macrotile expose strength 0 glues, except for the red tiles, which each expose the same strength 1 glue. Hence, the interior of each macrotile can be filled

⁵Note that in the glue-binding pad region there are no "single tile" bumps: this ensures that the simulator tile set U does not contain strength τ glues, which in turn simplifies our construction.

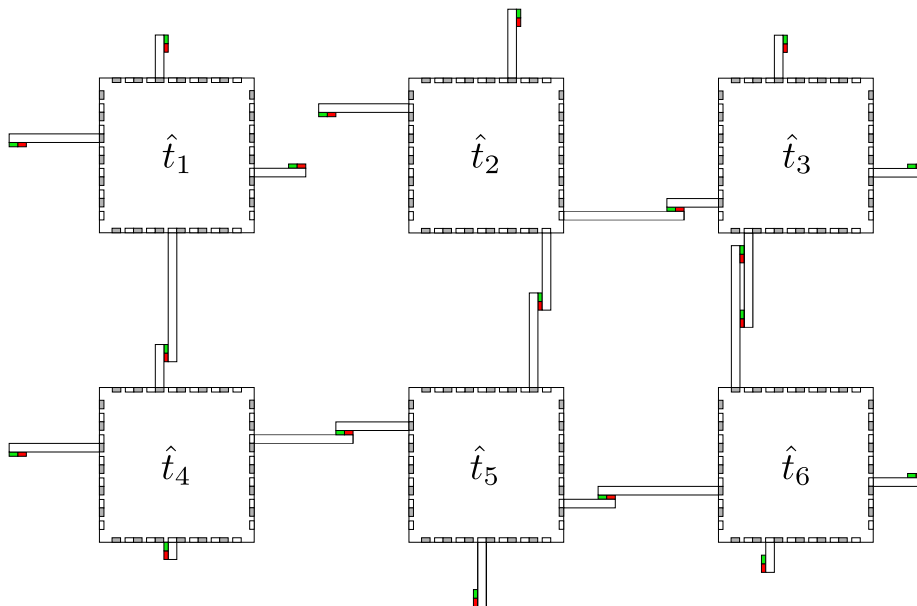


Figure 4: Strong simulation. Example assembly with 6 macrotiles and 2 simulated mismatches. The square “body” of each tile is of size $k \times k$, and the assembly sits on a $2k \times 2k$ lattice.

in using a single filler tile, with a constant size set of tile types used to fill the exterior.

Correctness of Simulation. Via the following two cases, macrotiles stay “on-grid.” (1) Due to their arm lengths, if two *matching* green glue pads bind (causing the binding of two macrotiles) then the combined (horizontal or vertical) arm length is exactly k . Thus *matching* macrotiles only bind in a way that their centers are exactly distance $2k$ apart. (2) Due to the green glue pad design, if two glue-binding pads mismatch (i.e. represent two mismatching glues in T) then they can not bind, since the two green pads sterically hinder each other. Taken together, this means that whenever two macrotiles (that encode two tile types $t_1, t_2 \in T$) bind, they are always positioned on a $2k \times 2k$ square grid.

This immediately implies that whenever larger assemblies, with multiple macrotiles, bind, they have all of their tiles positioned on a $2k \times 2k$ square grid.

It remains to show that, due to our macrotile design, $\mathcal{U}_{\mathcal{T}}$ strongly simulates \mathcal{T} (i.e. point 2 of Definition 2.5). Firstly, due the glue-binding pad design and the fact that \mathcal{T} and $\mathcal{U}_{\mathcal{T}}$ work at the same temperature τ , a pair of supertiles in \mathcal{T} bind if and only if their corresponding encoded pair of supertiles bind in $\mathcal{U}_{\mathcal{T}}$. This, taken together with the fact that the initial state of $\mathcal{U}_{\mathcal{T}}$ is an encoding

of the initial state of \mathcal{T} , implies that the two systems have the same dynamics (in the strong sense), thus satisfying Definitions 2.2 and 2.4. Definition 2.1(1) (equivalent production) is satisfied since equivalent dynamics implies equivalent production, and Definition 2.1(2) is satisfied as our choice of macrotile design directly implies that each supertile produced in $\mathcal{U}_{\mathcal{T}}$ maps cleanly to a supertile in \mathcal{T} . Taken together, these facts are sufficient to satisfy Definition 2.5(2).

This completes the proof of Theorem 5.1.

5.2 Strong simulation with smaller scale but more tile types

Theorem 5.2. *For each integer $\tau \geq 2$ there exists a single set of tile types U_{τ} , $|U_{\tau}| = O(\tau)$, such that for any 2HAM system $\mathcal{T} = (T, \sigma, \tau)$, there exists a set $I_{\mathcal{T}}$ of $\|\sigma\|$ input supertiles such that the 2HAM system $\mathcal{U}_{\mathcal{T}} = (U_{\tau}, U_{\tau, \text{sup}} \cup I_{\mathcal{T}}, \tau)$ strongly simulates \mathcal{T} at scale $O(\sqrt{|G| \log |G|})$, where G is the set of glues in T .*

Theorem 5.2 simply is a trade-off in tile types for scale factor with Theorem 5.1. The construction for Theorem 5.1 is modified so that where before each arm had $\text{str}(g)$ red tiles each exposing a strength 1 glue, now each arm has a *single* red tile that exposes one strength $\text{str}(g)$ glue. This adds an additional $\tau - 1$ tile types to the previous construction: one red tile type for each strength s where $\{1 \leq s \leq \tau\}$. However, by shrinking the red binding pads to size 1, it results in a reduction in scale factor to $O(\sqrt{|G| \log |G|})$, giving the statement of Theorem 5.2.

5.3 Strong simulation with larger scale but with fewer tile types and with only a single macrotile

The following Theorem provides an interesting tradeoff with the prior two results. It uses a larger scale factor than Theorems 5.1 and 5.2, but gives a universal 2HAM TAS for each τ where the *simulated tile set* is encoded in a single simulator supertile. This supertile has the ability to “grow into” any encoded tile in the simulated system \mathcal{T} .

Theorem 5.3. *For every $\tau \geq 2$ there exists a single set of tile types U_{τ} , with $|U_{\tau}| = O(1)$ (i.e. independent of τ), such that for all 2HAM systems $\mathcal{T} = (T, \sigma, \tau)$, there is a single supertile s_T (that represents $T_{\text{sup}} \cap \sigma$) and a set I_{σ} containing $\|\sigma - T_{\text{sup}}\|$ supertiles (that represent $\sigma - T_{\text{sup}}$) such that the 2HAM system $\mathcal{U}_{\mathcal{T}} = (U_{\tau}, U_{\tau, \text{sup}} \cup \{s_T\} \cup I_{\sigma}, \tau)$ strongly simulates \mathcal{T} at scale $O(\sqrt{|T|}(\log \tau + \log |G|) + \sqrt{|G|}(\tau + \log |G|))$, where G is the set of glues in T .*

Construction. For notation, let T_{in} be the set of singleton tiles used to make the multiset of singleton supertiles $T' = T_{\text{sup}} \cap \sigma$. The 2HAM simulator $\mathcal{U}_{\mathcal{T}}$ starts simulation from (A) a single input supertile s_T that encodes *all* singleton tiles from T_{in} , (B) $\|\sigma - T_{\text{sup}}\|$ supertiles that encode non-singleton input supertiles (i.e. they encode $\sigma - T_{\text{sup}}$), as well as (C) its own tile set U_{τ} .

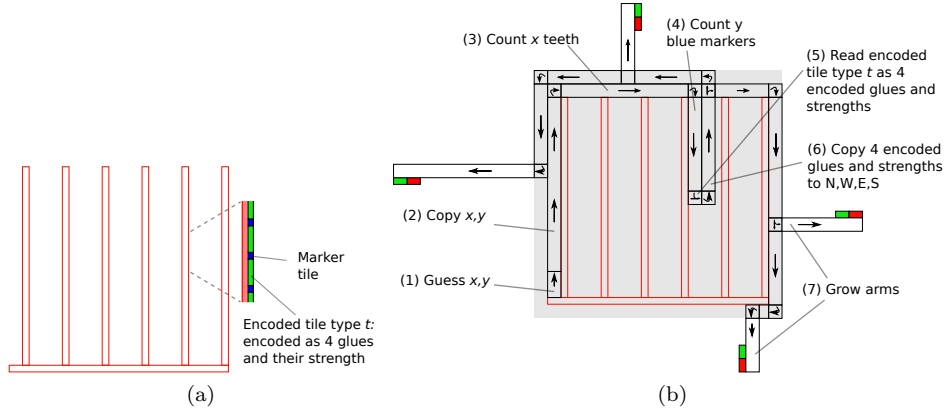


Figure 5: Overview of simulation for Theorem 5.3. (a) Input supertile. The red comb structure encodes the entire input simulated tile set T_{in} . Each of the $\lceil \sqrt{|T_{\text{in}}|} \rceil$ teeth of the comb encodes $\lceil \sqrt{|T_{\text{in}}|} \rceil$ tile types from T_{in} . Each simulated tile $t \in T_{\text{in}}$ is encoded using $O(\log |G| + \log \tau)$ tiles (shown in green in zoom-in) that encode the 4 glues of t , and their strengths, in binary. Encoded tile types are separated by a single marker tile shown in blue. (b) Growth of an input supertile into a macrotile that represents some tile type $t \in T$ from the simulated TAS \mathcal{T} . Numbers and arrows are used to indicate order and direction of growth. Growth of a crawler (black outline) begins in the lower left. (1) The crawler guesses two integers x, y , where $0 \leq x, y \leq \lfloor \sqrt{|T_{\text{in}}|} \rfloor$, and (2) copies these values to the top left corner. (3) The crawler grows right, counting x teeth, (4) grows south counting y blue tiles (see (a)), and (5) copies the encoded tile type t (as 4 encoded glues with strengths). (6) The glue information is copied to the four sides of the macrotile, and the glues are used as input to counters that (7) grow arms to the relevant lengths, including the red and green binding pads.

We begin by describing (A), the single supertile s_T . Figure 5a shows the supertile assembly s_T , which we call a “comb.” The comb encodes the entire simulated set of tile types, T_{in} , that are actually used by \mathcal{T} : each of the $\lceil \sqrt{|T_{\text{in}}|} \rceil$ teeth of the comb encodes $\leq \lceil \sqrt{|T_{\text{in}}|} \rceil$ tile types from T_{in} . Each simulated tile $t \in T_{\text{in}}$ is encoded using $O(\log |G| + \log \tau)$ tiles (shown in green in the zoom-in in Figure 5a) that encode the 4 glues of t , and their strengths, in binary. Encoded tile types are separated by a single marker tile shown in blue in Figure 5a.

The initial configuration of the simulator contains an infinite number of copies of the comb. Via the self-assembly process, each copy of the comb chooses to encode a tile type $t \in T_{\text{in}}$, as follows. Using tile types from U_τ , growth initiates from the lower left corner of the comb as shown in part (1) of Figure 5b. An $O(\log |G| + \log \tau)$ width crawler is initiated.⁶ This crawler “guesses” which tile

⁶The crawlers, counters, computational primitives (guessing strings, computing simple nu-

type the comb should represent via a nondeterministic procedure (in a way that guarantees that any $t \in T_{\text{in}}$ can be guessed), which works as follows. Via nondeterministic placement of $O(\log |G|)$ binary (0 or 1) tiles, the crawler guesses two bit strings, representing positive integers x, y . This pair of integers will act as indices to a location on the two-dimensional comb. The crawler climbs to the north of the leftmost tooth, as shown in Figure 5b(2), and then crawls to the east, counting exactly x teeth (3). The crawler heads south (4), counting y blue marker tiles (shown in Figure 5a: zoom-in). At the y^{th} marker, the crawler has found the encoding of its chosen tile type $t \in T_{\text{in}}$. Here (5), through cooperative binding, the crawler “reads” the encoding of t from the tooth. (The tile type t is encoded as $O(\log |G| + \log \tau)$ tiles, that represent the 4 glues, and their strengths in binary.) The crawler rotates and copies this information to the north (6). Upon reaching the top of the comb, the crawler splits into multiple crawlers, sending each of the 4 glue-strength pairs to their 4 respective sides (7). In this process each encoded glue g acts as an index of the position, and length, of the relevant macrotile arm (in other words, an encoded glue g acts as input to a counter that counts to an arm position, and then turns, and counts out an arm length—it can be easily seen that the $O(\log |G|)$ bits used to encode the glue in the crawler are sufficient to store the counter input). An arm generates a green and red glue-binding pad (of the same form as used the proof of in Theorem 5.1; see Figure 3b), positioned as shown in Figure 5b. Note that the green pad with the bumps and dents that encode the glue type in geometry must complete first (and therefore must be two tiles wide so that a path can grow out into each bump and then back down before continuing to the next) before the red pad forms, which is easily done by designing them to assemble as a path completely through green and then to red. This prevents the situation where the red pad with the generic glues could form first, allowing another macrotile to potentially bind without having the identity of the glue verified by the geometry (meaning that it could allow macrotiles representing mismatched glues of the simulated system to bind). Arm positions and lengths, and the position and structure of glue-binding pads, all follow the form used to prove Theorem 5.1, although the scaling is different here: specifically, as argued below, the body (outlined in grey in Figure 5b) is of size $k \times k$ where $k \in O(\sqrt{|T|}(\log \tau + \log |G|) + \sqrt{|G|(\tau + \log |G|)})$.

This completes the description of (A); the encoding of the *singleton* tile

merical functions on bit strings, and even simulating Turing machines), and geometric primitives (copying bit sequences around in two-dimensional space) used in this and later constructions are relatively straightforward implementations similar to those used in the aTAM in [12], among others. These primitives are designed to assemble on the edges of existing supertiles (or assemblies in the aTAM), and can be made (and usually already are) “2HAM-safe” (essentially, “polyomino safe” as in [20]), meaning that in the 2HAM they function identically and correctly without danger of unwanted supertiles forming which are unattached to the desired supertiles. The general technique is to limit the number of τ -strength glues on any particular tile type which assembles the primitive to 1, so that the largest unattached supertile which can form from them is a size 2 duple. All other attachments, and even the incorporation of the duples, requires cooperation provided by the surface of the supertile onto which the primitive is intended to form. Since the constructions for these primitives are standard and straightforward, we omit the details here.

types $t \in T$ that are used in the initial state σ . The remaining supertiles (B) of σ (i.e. of size > 1) are “hard-coded” as supertiles \hat{s} (using tiles from U_τ) that encode supertiles $s \in \sigma$ (analogously to Theorem 5.1). Essentially these pre-built (hardcoded) supertile assemblies are composed of fully grown comb macrotiles, complete with arms and glue-binding pads.

Scale. We first analyze the size of the *body* of a filled-out comb macrotile (shaded in grey in Figure 5b). $T_{\text{in}} \subseteq T$, so $|T_{\text{in}}| \leq |T|$. Each of the $\leq \lceil \sqrt{|T|} \rceil$ teeth of the comb encodes $\leq \lceil \sqrt{|T|} \rceil$ tile types from T_{in} . Each simulated tile $t \in T_{\text{in}}$ is encoded using $O(\log |G| + \log \tau)$ tiles (shown in green in the zoom-in in Figure 5a) that encode the 4 glues of t , and their strengths, in binary. Crawlers are of width $O(\log |G| + \log \tau)$. Together these terms sum to $\ell_1 = O(\sqrt{|T|}(\log \tau + \log |G|))$, giving the $\ell_1 \times \ell_1$ scaling for the body (shaded in grey in Figure 5b).

The square arm regions for a filled-out comb macrotile use the same values as appeared in Theorem 5.1: the glue-binding pad is of size $O(\tau + \log |G|)$, there are $|G|$ such pads, which are rasterized into an $\ell_2 \times \ell_2$ region where $\ell_2 = \sqrt{|G|(\tau + \log |G|)}$.

We take $\max(\ell_1, \ell_2) = k$ to get our final scaling of $k \times k$ where

$$k = O(\sqrt{|T|}(\log \tau + \log |G|) + \sqrt{|G|(\tau + \log |G|)}).$$

Correctness of Simulation. The simulation correctness uses the same argument as in the proof of Theorem 5.1, along with the following observations. In Theorem 5.1 the macrotiles were “hardcoded” in advance and so no actual growth takes place within the macro tile itself. Here, for Theorem 5.3, the macrotiles grow from a comb. Macrotiles bind via the red binding pads (Figure 5b). First, observe that after a comb has selected the values x, y , its (future) identity \hat{t} is completely determined. Second, observe that macrotiles bind to each other via their red binding pads only. Third, from the above description of the construction, the red binding pads form only *after* the green glue pads. These three facts together, imply that macrotiles bind in the simulator if and only if they bind in the simulated system. From here, the correctness argument proceeds as in the proof of Theorem 5.1.

6 The temperature- τ 2HAM is intrinsically universal: simulation

In this section we give three results that exhibit for any integer $\tau \geq 2$, a single set of tiles U_τ that can simulate any temperature τ 2HAM system at temperature τ given a proper configuration of initial assemblies over U_τ . Of particular focus in this section is the achievement of logarithmic scale factors in the size of the simulated tile system. Beyond this, each of the three results in this section depicts different tradeoffs between the size of U_τ in terms of τ versus the scale of the simulation in terms of τ .

As in Section 5, we adopt the following notation. Let $\|M\|$ denote the number of distinct elements in the multiset M , i.e. $\|M\| = |M'|$ is the cardinality of the set M' induced by the multiset M . Let T_{sup} denote the set of supertiles induced by a tile set T . By this, we simply mean T_{sup} is the set of supertiles formed by taking all tiles in T and translating them to all locations in \mathbb{Z}^2 .

Theorem 6.1. *For each integer $\tau \geq 2$ there exists a single set of tile types U_τ , $|U_\tau| = O(\tau)$, such that for any 2HAM system $\mathcal{T} = (T, \sigma, \tau)$, there exists a set $I_\mathcal{T}$ of $O(|T| + \|\sigma\|)$ input supertiles such that the 2HAM system $\mathcal{U}_\mathcal{T} = (U_\tau, U_{\tau, \text{sup}} \cup I_\mathcal{T}, \tau)$ simulates \mathcal{T} at scale $O(\sqrt{\log |T|})$.*

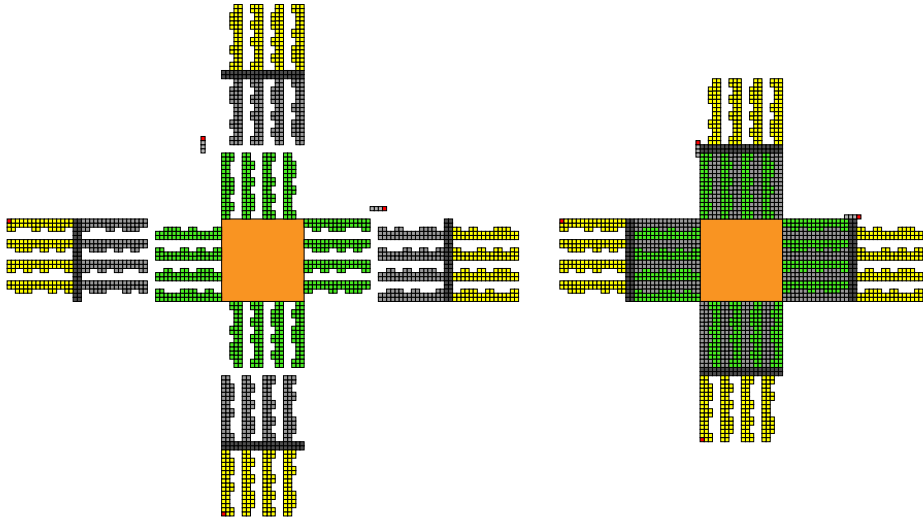


Figure 6: These input assemblies for each tile type t comprise the construction for Theorem 6.1 and are each termed the *megatile* representation of their respective tile t .

Construction. For a given tile system $\mathcal{T} = (T, \sigma, \tau)$, we first construct a collection of supertiles for each tile in T corresponding to the supertiles depicted in Figure 6. We then construct an assembly for each element of σ by placing combined instances of the assemblies for tile types according to the tiles making up the assemblies from σ . First, for each tile in T , we construct a corresponding *macrotiler* consisting of a scale $4\lceil\sqrt{\log |T|}\rceil$ orange square assembly with protruding green *teeth* growths whose geometry encodes a unique identifying binary number for the tile represented by the macrotiler. For each set of green teeth on each of the four sides of the macrotiler, there is also a matching *glue gadget* assembly shown with grey and yellow tiles. The grey portion of this gadget encodes the complement of the macrotiler's green geometry, allowing for a snug fit of the appropriate glue gadget to each face of the macrotiler. Additionally,

some τ strength glue is exposed on some portion of the green teeth that matches a glue on the grey portion of each glue gadget (note that all glue gadgets share this same glue and thus rely on the unique complementary geometry to enforce that the unique correct glue gadget is the only gadget that may attach to the macrotile face). The yellow portion of the glue gadget encodes in geometry a unique binary string representing the glue that occurs on the respective face of the tile to be simulated. Further, a given glue type represented by a gadget on an east face exposes a binary bump pattern that is the complement to the glue gadget for the same glue type occurring on any west face. The same complementary setup is used for north and south glue gadgets. A close-up of the green teeth portion of the assembly, and the attachable glue gadget, is shown in Figure 7.

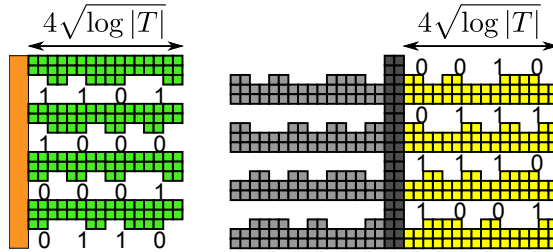


Figure 7: Each tile type can select the proper glue gadget by exposing a pattern of green binary teeth unique to the tile type. Each glue gadget exposes a binary sequence of yellow teeth encoding the glue type represented by the glue gadget.

In addition to the macrotiles for each tile in T , we also have a supertile consisting of a collection of macrotiles for each supertile in $S \in \sigma$. The construction is the natural one in which a copy of the representing macrotile is placed, at scale, for each of the corresponding tiles in S . For any adjacent tile faces in S that match glue type, the glue gadgets for the respective sides of the respective macrotiles are attaching, along with the glue assemblies, implying that the two macrotiles are connected to each other with the same net strength that the tiles in S are connected with. For mismatched adjacent glues, both glue gadgets are unattached (or, alternately, just one of the two glue gadgets is unattached), which ensures that the macro blocks can sit together without occupying the same tile positions. For all glue faces in S that are exposed, i.e., are not adjacent to another tile face, the corresponding macrotile edge does not have its glue gadget attached. This initial lack of glue gadgets is fundamentally important to ensure a proper simulation in the case of glue mismatches.

Given this set of supertiles $I_{\mathcal{T}}$ derived from T and σ , we obtain the system $\mathcal{U}_{\mathcal{T}} = (U_{\tau}, U_{\tau, \text{sup}} \cup I_{\mathcal{T}}, \tau)$ where U_{τ} is a generic set of tiles used to form the body of the construction's macrotiles, as well as a set of τ tiles used to supply a set of tiles that form the exposed red glues of strength from 1 to τ , thereby yielding a size $|U_{\tau}| = O(\tau)$ tile set for the construction. Note that U_{τ} only depends on

τ and can be used regardless of which system is being simulated.

Assembly for system $\mathcal{U}_{\mathcal{T}} = (U_{\tau}, U_{\tau, \text{sup}} \cup I_{\mathcal{T}}, \tau)$ occurs via attachment between a base macrotile's exposed macro face and any of the macrotile's corresponding glue gadgets. Upon attachment of either a north or east glue gadget, a cooperative binding site is exposed allowing for the attachment of the 4-tile assembly shown as 3 grey tiles and a red tile in Figure 6. In the case that the strength of the glue represented by the corresponding glue gadget is less than τ , this 4-tile assembly attaches a glue (exposed on the surface of the red tile) of strength equal to the strength of that glue. In the case that the glue to be represented is a τ -strength glue, the red glue has strength only $\tau-1$ and the glue gadget is assumed to expose an additional strength-1 glue somewhere along its surface. The exposed red glues on north and east macrotile glue gadgets match (any of) the red glues of equal strength on the west and south glue gadgets. Again, since the same glue is used for all distinct glues of the same strength, the geometric compatibility of the binary teeth of the gadget are relied upon to ensure only complementary glues will realize this attraction.

Scale and Tile Set Size. The scale of this construction is $O(\sqrt{\log |T|})$ with the mapping from macrotile assemblies to simulated assemblies being the natural mapping implied by the mapping from each macrotile to the tile type the macrotile is derived from. The bulk of the assemblies can be constructed from an $O(1)$ set of tile types. An additional special set of τ tile types is required to place exposed glues of strengths in the range of 1 to τ . Instances of the special set of tiles are represented in the figure by red tiles.

Correctness of Simulation. The mapping of supertile blocks to tiles in T is the natural mapping from the large orange blocks to the unique tile in T that the macrotile was designed for. Whenever two assemblies of macrotiles attach, they must do so based on a sufficient number of matching glue gadgets with a sufficient strength of at least τ . By design, then, the assemblies mapped to by these two macrotiles have the same exposed glue, and therefore must be able to attach. Thus, the system \mathcal{T} from which the macrotiles and macro assemblies are derived follows the derived system $\mathcal{U}_{\mathcal{T}}$.

We now argue that the derived system $\mathcal{U}_{\mathcal{T}}$ weakly models the original system \mathcal{T} . Suppose some assembly $\alpha \in \mathcal{A}[\mathcal{T}]$ can attach to some $\beta \in \mathcal{A}[\mathcal{T}]$ to produce some $c \in \mathcal{A}[\mathcal{T}]$. Now consider any α' macrotile assembly that maps to α . We must show that α' can grow into an α'' that also maps to α , and that there must exist a β' that can attach to α'' to form an assembly that maps to c .

The α'' that suffices is the super tile consisting of attaching to α all glue gadgets on any exposed macro surfaces of α , along with the 4-tile glue assemblies that expose the red tiles with the exposed glues. The β' that suffices is any supertile that maps to β with the added restriction that the only glue gadgets that have attached to β' 's exposed macro edges correspond to the set of glue faces that correspond to matched pairs of glues in the bonding of α and β . This ensures that α' and β' will not have mutually exclusive *mismatched* glue gadgets

that might prevent attachment via geometric hindrance. Further, the inclusion of all matching glues ensures that enough affinity for attachment will be present based on the assumption that α and β are combinable.

Theorem 6.2. *For each integer $\tau \geq 2$ there exists a single set of tile types U_τ , $|U_\tau| = O(1)$ (i.e. independent of τ), such that for all 2HAM systems $\mathcal{T} = (T, \sigma, \tau)$, there exists a set $I_\mathcal{T}$ of $O(|T| + \|\sigma\|)$ input supertiles such that the 2HAM system $\mathcal{U}_\mathcal{T} = (U_\tau, U_{\tau, \text{sup}} \cup I_\mathcal{T}, \tau)$ simulates \mathcal{T} at scale $O(\sqrt{\log |T| + \tau})$.*

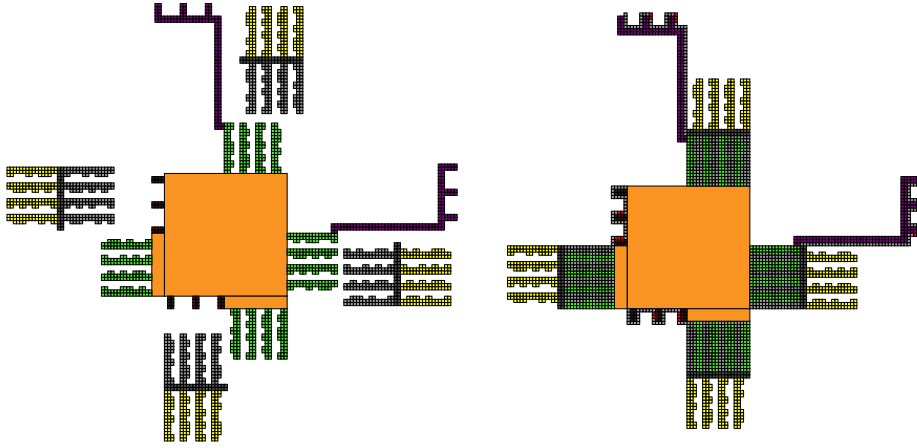


Figure 8: To simulate any 2HAM system from a set of tiles U_τ that does not scale in size with τ , the construction for Theorem 6.2 extends the construction of Theorem 6.1 to encode glue strength linearly with single strength glues. This construction increases the scale by a $\sqrt{\tau}$ factor. Alternately, Theorem 6.3 provides a much more modest scale factor increase of $O(\sqrt{\log \tau})$ by encoding strengths through a binary representation and paying the price of a size $|U_\tau| = O(\log \tau)$ set of tiles, which represents a nice compromise between the alternate extremes.

Construction. The construction for this simulation is a modification of the construction for Theorem 6.1 and the blocks utilized for each tile in T are shown in Figure 8. The key extension in this construction is that the exposed glues that provide the affinity for attachment among macrotiles with attached glue assemblies can no longer make use of a size τ set of tiles, but instead must encode τ through a smaller tile set. In this construction we utilize a linear encoding of τ that uses an optimal $O(1)$ set of tiles by stacking single strength glues to create a net force of the desired glue strength. This linear encoding is displayed in the $\sqrt{\tau} \times \sqrt{\tau}$ rectangular regions of dark tiles on the west and south block faces, and at the end of the extended purple arms on north and east faces. As with the previous construction, attachment of the glue assemblies exposes a cooperative bonding site, in this case permitting a chain of grey tiles to *coat* the

$O(\sqrt{\tau}) \times O(\sqrt{\tau})$ rectangular region with tiles that expose a number of red tiles equal to the strength of the glue represented on the macrotile face. Each red tile exposes a single strength glue, yielding the desired net force of attraction.

Scale and Tile Set Size. The scale of this construction is $O(\sqrt{\log |T| + \tau})$. Each assembly can be constructed from an $O(1)$ set of tile types as the additional $O(\tau)$ glue types are no longer needed based on the linear encoding of each glue strength.

Correction of Simulation. The argument for correct simulation is essentially the same as for Theorem 6.1. A sample figure depicting how the subtle issue of mismatched glues are handled in this simulation is shown in Figure 9.

Theorem 6.3. *For each integer $\tau \geq 2$ there exists a single set of tile types U_τ , $|U_\tau| = O(\log \tau)$, such that for any 2HAM system $\mathcal{T} = (T, \sigma, \tau)$, there exists a set I_τ of $O(|T| + \|\sigma\|)$ input supertiles such that the 2HAM system $\mathcal{U}_\tau = (U_\tau, U_{\tau, \text{sup}} \cup I_\tau, \tau)$ simulates \mathcal{T} at scale $O(\sqrt{\log |T| + \log \tau})$.*

This theorem represents an interesting compromise between the $O(\tau)$ size tileset required for the simulation for Theorem 6.1, and the substantial $\sqrt{\tau}$ scale factor increase from Theorem 6.2. This middle ground approach utilizes a construction similar to Theorem 6.2, modified to make use of $O(\log \tau)$ tile types, each of which exposes a glue strength equal to one of the numbers $1, 2, 4, \dots, 2^{\log \tau}$. The linear encoding of glue strengths via the $\sqrt{\tau} \times \sqrt{\tau}$ rectangular regions from Theorem 6.2 may now be implemented with smaller $\sqrt{\log \tau} \times \sqrt{\log \tau}$ rectangular regions by using a binary encoding of each glue strength.

7 Simultaneous simulation of all 2HAM systems at temperature τ

In previous sections, we have shown that while the entire 2HAM is not intrinsically universal, the 2HAM at each temperature τ is intrinsically universal. In this section we show that there exists a tile set in the 2HAM at each temperature $\tau > 1$ which simultaneously and in parallel simulates *every* 2HAM tile assembly system (with a default initial state) at temperature τ . These simultaneous simulations are guaranteed to occur in parallel without any interaction between the macrotiles simulating different systems, and to faithfully strongly simulate each system. Note that due to technical reasons discussed in Section 7.3, this construction finitely self-assembles (see [3]) the supertiles of the simultaneous simulations.

Before we formally can state our main result for this section, we must define the notion of *simultaneous strong* simulation.

Assume we have a standard enumeration of every 2HAM TAS at temperature τ , i.e., $\mathcal{T}_0 = (T_0, \tau), \mathcal{T}_1 = (T_1, \tau), \dots$. Let $\mathcal{U} = (U_\tau, \tau)$ be a 2HAM TAS and for

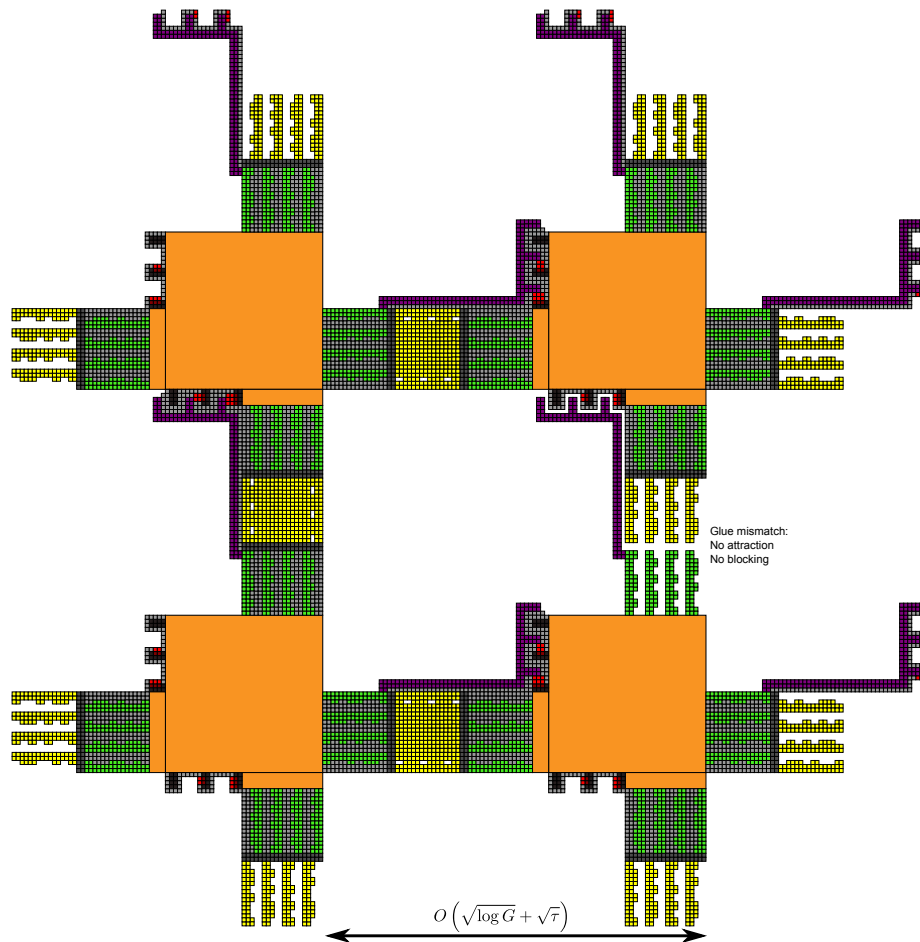


Figure 9: Megablocks that have attached glue assemblies and subsequently placed enough single-strength glue tiles may begin attaching to each other based on the combined strength of matching glues. The geometric teeth of the glue assemblies ensure that only matching glue types will be compatible, thus preventing mismatches from providing positive strength attachments. Further, let A' be any simulated assembly that combines with simulated assembly B' . Then for any assembly A that maps to A' there must exist a simulator assembly B that attaches to A , even in the presence of simulated glue mismatches (such a B must exist—it simply has not yet attached its glue assembly in the positions for which simulated glue mismatches occur).

each $i \in \mathbb{N}$, let R_i be an m -block representation function $R_i : B_m^U \dashrightarrow T_i$. Define $\mathcal{T} = \langle T_i \rangle_{i=0}^\infty$ and $R = \langle R_i \rangle_{i=0}^\infty$.

Definition 7.1. We say that \mathcal{U} and \mathcal{T} have simultaneous equivalent productions, and we write $\mathcal{U} \Leftrightarrow_R^\infty \mathcal{T}$ if the following conditions hold:

1. For all $\tilde{\alpha} \in \mathcal{A}[\mathcal{U}]$, there exists at most one value $i \in \mathbb{N}$ such that $\tilde{R}_i(\tilde{\alpha}) \in \mathcal{A}[T_i]$.
2. For all $i \in \mathbb{N}$, for every $\tilde{\beta} \in \mathcal{A}[T_i]$, there exists $\tilde{\alpha} \in \mathcal{A}[\mathcal{U}]$ such that $\tilde{R}_i(\tilde{\alpha}) = \tilde{\beta}$.
3. For all $\tilde{\alpha} \in \mathcal{A}[\mathcal{U}]$ there exists at most one value $i \in \mathbb{N}$ such that $\tilde{\alpha}$ maps cleanly to $\tilde{R}_i(\tilde{\alpha})$.

Definition 7.2. We say that \mathcal{T} follows \mathcal{U} , and we write $\mathcal{T} \dashv_R^\infty \mathcal{U}$ if, for any $\tilde{\alpha}, \tilde{\beta} \in \mathcal{A}[\mathcal{U}]$ such that $\tilde{\alpha} \rightarrow_{\mathcal{U}}^1 \tilde{\beta}$, there exists at most one value $i \in \mathbb{N}$ such that $\tilde{R}_i(\tilde{\alpha}) \rightarrow_{\tilde{T}_i}^{\leq 1} \tilde{R}_i(\tilde{\beta})$.

Definition 7.3. We say that \mathcal{U} strongly models \mathcal{T} , and we write $\mathcal{U} \models_R^\infty \mathcal{T}$ if for any $i \in \mathbb{N}$, for any $\tilde{\alpha}, \tilde{\beta} \in \mathcal{A}[T_i]$ such that $\tilde{\gamma} \in C_{\tilde{\alpha}, \tilde{\beta}}^\tau$, then for all $\tilde{\alpha}', \tilde{\beta}' \in \mathcal{A}[\mathcal{U}]$ such that $\tilde{R}_i(\tilde{\alpha}') = \tilde{\alpha}$ and $\tilde{R}_i(\tilde{\beta}') = \tilde{\beta}$, it must be that there exist $\tilde{\alpha}'', \tilde{\beta}'', \tilde{\gamma}' \in \mathcal{A}[\mathcal{U}]$, such that $\tilde{\alpha}' \rightarrow_{\mathcal{U}} \tilde{\alpha}'', \tilde{\beta}' \rightarrow_{\mathcal{U}} \tilde{\beta}'', \tilde{R}_i(\tilde{\alpha}'') = \tilde{\alpha}, \tilde{R}_i(\tilde{\beta}'') = \tilde{\beta}, \tilde{R}_i(\tilde{\gamma}') = \tilde{\gamma}$, and $\tilde{\gamma}' \in C_{\tilde{\alpha}'', \tilde{\beta}''}^\tau$.

Definition 7.4. Let $\mathcal{U} \Leftrightarrow_R^\infty \mathcal{T}$ and $\mathcal{T} \dashv_R^\infty \mathcal{U}$. We say that \mathcal{U} strongly simultaneously simulates \mathcal{T} if $\mathcal{U} \models_R^\infty \mathcal{T}$.

Theorem 7.5. For each $\tau > 1$, there exists a 2HAM system $\mathcal{S} = (U_\tau, \tau)$ which strongly simultaneously simulates all 2HAM systems $\mathcal{T} = (T, \tau)$.

In [3], it was shown that for every temperature $\tau > 1$, for every aTAM system $\mathcal{T} = (T, \sigma, \tau)$ with $|\sigma| = 1$, there exists a 2HAM system $\mathcal{S} = (S, \tau)$ which simulates it. From this and Theorem 7.5, the following corollary arises:

Corollary 7.6. For every temperature $\tau > 1$, there exists a 2HAM system $\mathcal{S} = (U_\tau, \tau)$ which simultaneously simulates every aTAM system $\mathcal{T} = (T, \sigma, \tau)$ where $|\sigma| = 1$.

7.1 Construction overview

This construction works with the singleton tiles of the tile set U_τ being the only contents of the initial state, and with them, in parallel via the process of self-assembly, forming macrotiles such that each macrotile simulates one specific tile type from one specific 2HAM system at temperature τ . (Note that such a system, since it has a default initial state, i.e. only singleton tiles, can be fully specified by its tile set since the temperature is also given.) This is done

in such a way that every tile type of every temperature τ 2HAM tile set is represented by a unique macrotile. These macrotiles are guaranteed to interact only with macrotiles representing tile types from the same tile set, and to do so in such a way that the group of macrotiles for that system strongly simulate that system. The scale factor of the simulation of each system, and thus the size of the macrotiles used to simulate it, is potentially unique, and depends upon the running time of a Turing machine and the size of the simulated tile set.

Each macrotile construction begins by randomly selecting the tile set to which the macrotile will belong, and then randomly selecting which tile in the chosen tile set that macrotile will simulate. We guarantee that, in parallel, macrotiles strongly simulate every 2HAM tile assembly system at temperature τ . Thus it is required that each tile type of each tile set has a corresponding macrotile. Also, each macrotile must encode information about not only the tile type it is simulating, but also must encode information about the tile set (namely, its number in a fixed enumeration) to which the tile type belongs. Without the information about which tile set the macrotile is simulating, it would be possible for two macrotiles that simulate tiles from different tile sets to bind with strength greater than 0, possibly violating the definition of strong simulation.




Figure 10: A high level overview of a macrotile.

7.2 Construction details

The assembly of a macrotile begins with the four tile types shown in Figure 14. These tiles nondeterministically assemble to yield a random number $r \in \mathbb{Z}^+$. Next, a standard Turing machine simulation reads r and outputs the tile set

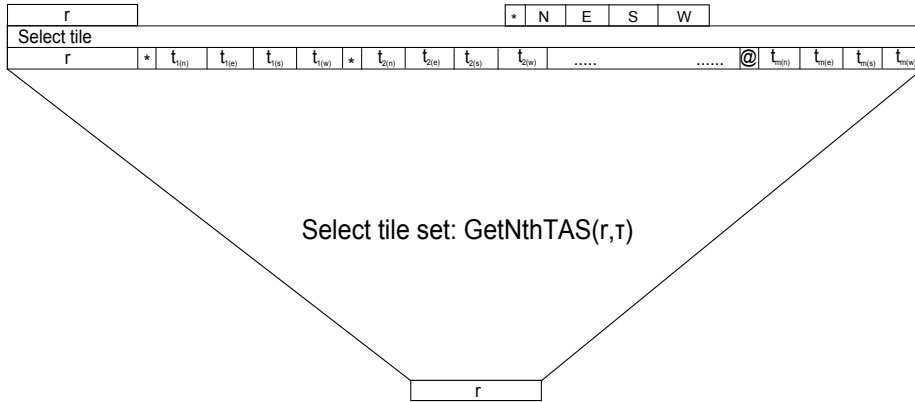


Figure 11: A high level overview of the main components involved in building a macrotile. Note that each square represents a segment of tiles rather than an individual tile.

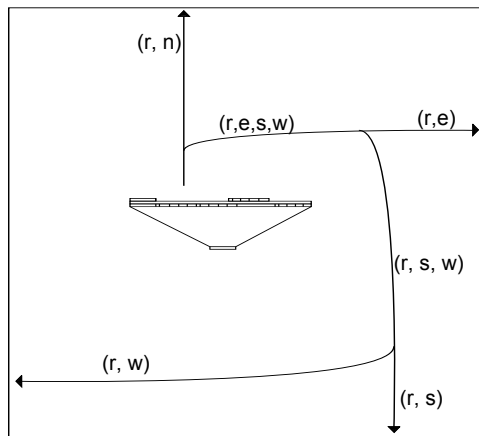


Figure 12: The information flow during the assembly of a macrotile. In this figure, r represents the random number, and n, e, s, w stand for the binary representation of the north, east, south and west glues respectively.

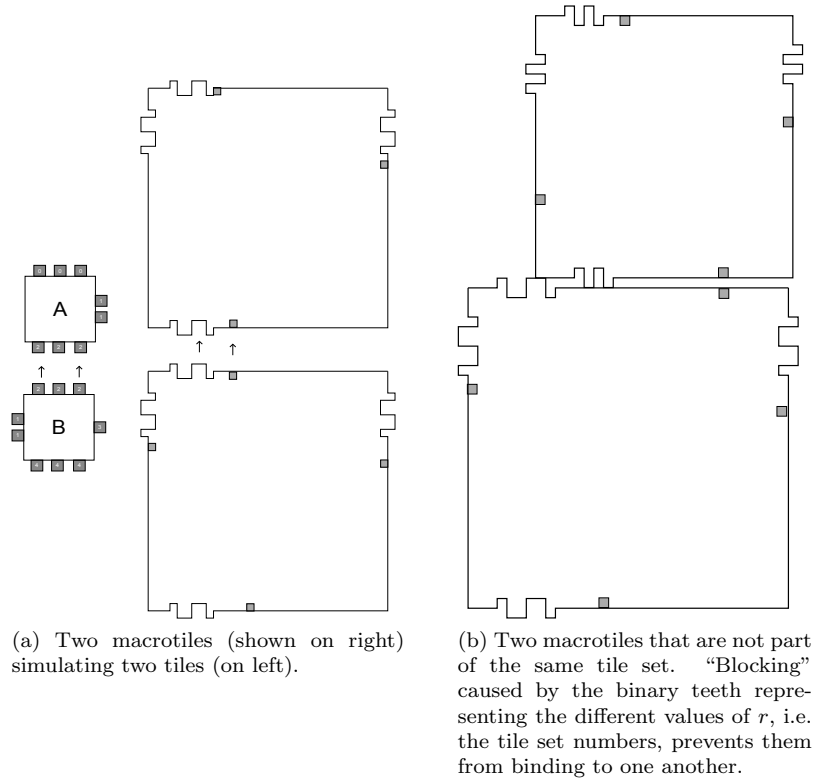


Figure 13

with that index number by running the enumeration program described below. The tile set is output as a row which encodes each tile type side by side, with spacing tiles in between, and with each tile type represented as the set of its four glues (which in turn are each represented as a pair of binary values for the label and strength of that glue). Next, exactly one tile type is nondeterministically selected and its information is copied upward. (See Figure 11.) Then, the assembly carries all of the information necessary for each macrotile side to the appropriate locations as seen in Figure 12. Next, each macrotile edge first forms a region of bumps and dents which correspond to a binary encoding of the number r , which represents the number of the tile set being simulated. Once those are complete, the remaining portion of the side, which encodes the glue information, forms. (See Figure 10.)

7.2.1 Creation of random numbers

The assembly generates random numbers using the tile set shown in Figure 14. Since all of the interior glues are of the same type and τ strength, these tiles assemble nondeterministically. Because of this nondeterministic assembly, for any positive integer there exists an assembly of the four tiles such that the assembly is a binary representation of that number. That is, these four tile types will generate the set of all positive integers.

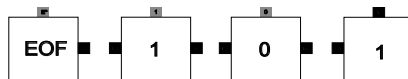


Figure 14: The four tiles that bond nondeterministically to generate the random numbers required for our construction. The black squares represent τ strength glues and the grey squares represent $\lceil \frac{\tau}{2} \rceil$ strength glues.

7.2.2 Selecting a tile set

Now that we have a random number, the assembly can decide to which tile set the macrotile belongs. In order to accomplish this, we use the program $\text{GETNTHTAS}(tas_num, \tau)$. The idea of $\text{GETNTHTAS}(tas_num, \tau)$ is to enumerate all of the tile sets at a temperature τ . (See Section 7.3 for more details.) The algorithm will take a natural number and τ as input and output the tile set associated with that number. It does so by, for every possible number of glues (in succession) and combinations of assignments of strengths to those glues from 1 to τ , creating all possible permutations of unique tile types for those glues. For each such set, it then enumerates the power set, each element of that being a tile set. The program continues enumerating tile sets until the number assigned to a tile set equals tas_num . Once this happens, it outputs the tile set associated with this number. The tile set is outputted as the glues that compose each tile in the tile set with special symbols in between each tile definition.

```

GETNTHTAS(tas_num,  $\tau$ )
1  tile_set_number  $\leftarrow$  0
2   $|G| \leftarrow 1$ 
3      do
4          glue_config  $\leftarrow$  the value of  $|G|$  1's in base  $(\tau + 1)$ 
5          while glue_config  $\leq (\tau + 1)^{|G|} - 1$ 
              (the value of  $|G|$   $\tau$ 's in base  $(\tau + 1)$ )
6              do
7                   $a_i \leftarrow$  the  $i^{\text{th}}$  digit from the least
                    significant digit of glue_config
8                   $T = \{\}$ 
9                   $n \leftarrow 1$ 
10                 while  $n \leq (|G| + 1)^4 - 1$ 
11                     do
12                          $s \leftarrow n$ , represented in base  $(|G| + 1)$  and
                            padded to length 4
13                          $s_i \leftarrow$  the  $i^{\text{th}}$  digit from the least significant
                            digit of  $s$ 
14                          $t \leftarrow ((s_3, a_{s_3}), (s_2, a_{s_2}), (s_1, a_{s_1}), (s_0, a_{s_0}))$ 
15                          $T \leftarrow T \cup \{t\}$ 
16
17                 for each  $p \in \mathcal{P}(T)$ 
18                     do
19                         if tile_set_number = tas_num
20                             then
21                                 return  $p$ 
22                             else
23                                 tile_set_number ++
24
25                                 glue_config ++
26
27                                  $|G|$  ++
28

```

7.2.3 Selecting a tile

Now that GETNTHTAS(*tas_num*, τ) has selected the tile set in which our macrotile lives, the assembly chooses the specific tile that this macrotile will become from that tile set. It achieves this by growing a row of tiles across the top of the row which defines the tile set (which is shown in Figure 11). Until it has selected a particular tile type as that which will be simulated, at each position denoting the beginning of the definition of a new tile type, two tiles are able to bind - one which selects that tile and one which passes on it. If a tile type is selected, the row which then grows above its definition copies the definition of the tile type upward. If it is not selected, that information is not propagated upward. Once a tile type has been selected, all others are ignored

(i.e. the tile type which previously could have bound to select the tile can no longer do so). Finally, if the row grows to the position of the final tile in the tile set (which is specially marked to denote that it is the last tile) and has yet to select a tile type, it is forced to choose this final type. Thus, every tile type has some probability of being selected, and it is impossible for no tile type to be selected.

7.2.4 Assembling the macrotile

As in Section 5.3, we are able to use computational and geometric primitives to disperse the information to all sides of the macrotile as shown in Figure 12. The glue portion is simply a series of $|G|\tau$ (where $|G|$ is the number of glues in the tile set being simulated) tile locations such that the first τ positions represent glue 1, the second τ positions glue 2, and so on. If a given side of a supertile is to simulate the i th glue, whose strength is j , exactly j of the τ tiles which represent that glue expose strength 1 glues to the exterior of the macrotile. (On the north and south sides of a macrotile it is the westernmost j tiles of the group, and for the east and west sides it is the southern most j tiles of the group.) All other tiles on that side of the macrotile expose 0 strength glues. By forming the binary bumps and dents before the section encoding the glue, it is ensured that only two macrotiles belonging to the same tile set may come together, since only they will have the same pattern and thus geometrically be able to fit together. See Figure 13a for an example of how the “binary bump” segment allows tiles from the same tile set to come together, and see Figure 13b for an example of the “binary bump” segment blocking two tiles from coming together that belong to different tile sets. On the other hand, the glue segment ensures that two macrotiles bind only if the tiles they are representing are able to bind. As in Section 5.3, we must ensure that the “binary bump” segment of the side of a macrotile is assemble before the glue segment. Without this measure, it would be possible for a partially assembled macrotile that has its glue segment assembled but not its teeth to bind to a macrotile simulating a different tile set.

7.3 Proof of correctness

To see that Theorem 7.5 follows from this construction we show that 1) for every tile type in every possible tile set at temperature τ , a unique macrotile self-assembles which maps to that tile type in that tile set, and 2) each system is strongly simulated by the macrotiles representing it.

We first define what it means for two tile sets to be *functionally equivalent* to each other.

Definition 7.7. *We say that tile set T_1 is functionally equivalent to tile set T_2 if there exists a one-to-one mapping function $f : T_1 \rightarrow T_2$ that maps each tile type in T_1 to a unique tile type in T_2 such that for every pair of tiles $t_a, t_b \in T_1$, if and only if side $d_1 \in \{N, E, S, W\}$ of t_1 binds to side $d_2 \in \{N, E, S, W\}$ of t_2 with strength s , side d_1 of $f(t_1)$ binds to side d_2 of $f(t_2)$ with strength s .*

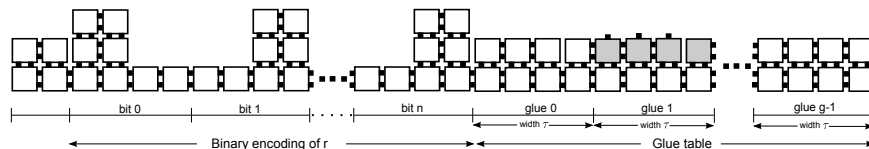


Figure 15: Example macrotile side showing the binary teeth which encode r as well as the glue table, which is simply a $\tau \cdot |G|$ (where $|G|$ is the number of glues in system r) row of tiles such that the i th consecutive group of τ tiles represent the i th glue. Since each macrotile side can simulate exactly one glue, only the tiles of a single glue segment have non-zero external glues. Those glues are all strength-1 and the number of them is equal to the strength value of the simulated glue. In this example, glue 1 is the glue being simulated where the strength of glue 1 is 3 in $\tau = 4$ system.

Therefore, if two tile sets are functionally equivalent, each tile type of each set has exactly one identical counterpart in the other tile set which is able to bind to the same exact sides of the same set of equivalent tiles. Furthermore, we say that two tile sets are *functionally distinct* if they are not functionally equivalent.

We begin with the claim that a set covering all functionally distinct tile sets in the 2HAM at temperature τ is countable. To do this we first note that regardless of the specific labels used for the glues of an arbitrary tile set T , and in fact there is an uncountably infinite set of possible labels, there is a functionally equivalent tile set which replaces each glue label with the integer g for $0 \leq g < |G|$ (represented as a string) where G is the set of glues in T . We will refer to such tile sets (i.e. those using the integers as their glue labels) as the *canonical* tile sets. Now, observe that given a maximum strength (i.e. τ) and a fixed number of glues $|G|$, the tile sets that can be produced is finite. The number of possible glues for any canonical tile set is given by a positive integer, and the set of positive integers is countable. It follows that the set of all canonical tile sets in the 2HAM at temperature τ thus consists of a countable union of finite sets. Consequently, all the set of all canonical tile sets in the 2HAM at temperature τ is countable. Since every tile set is functionally equivalent to a canonical tile set, the countable set of canonical tile sets represents a set of all functionally distinct 2HAM tile sets at temperature τ .

Now, we note that by the definition of (strong) simulation, the (strong) simulation of a system at temperature τ which includes a single element of a set of functionally equivalent tile sets E is a (strong) simulation of all systems at temperature τ which include any tile set in E . Therefore, in order to simulate every 2HAM system at temperature τ , it is sufficient to simulate the set of all systems composed of canonical tile sets. We now show that $\text{GETNTHAS}(tas_num, \tau)$ enumerates every canonical tile set at temperature τ . For every positive integer $|G|$ in succession, it uses $|G|$ as the number of glues and, using the canonical form of glue labels, creates every possible combination of mappings of glue labels

to strengths from 1 through τ . For each such mapping, signifying a set of glues and associated strengths, it creates a set of every possible tile type that can be created using those glues (and sides with no glue). Given this set, the complete set of all tiles that could be created with these glues, it treats each element of the power set as its own tile set, assigning the next tile set number to each, until it reaches the number matching its input value and then outputs that tile set. Note that, as it iterates and creates power sets, it will create power sets which contain elements (corresponding to tile sets) that are identical to those of other power sets. This does not pose a problem, as it simply means that such tile sets will be counted multiple times and therefore simultaneously simulated under different numbers. The important fact is that no canonical tile set is excluded from the enumeration. In such a way, it is guaranteed that every canonical tile set is included in the enumeration.

Since every tile in a chosen tile set r is selected to be built into a macrotile with some probability $p > 0$ (based on the nondeterministic selection described in Section 7.2.3) and there will be an infinite number of copies of the assembly signifying that tile set number r has been selected in solution, it follows that there exists a unique macrotile that self-assembles which maps to that tile type for that system. Note that an assembly sequence exists in which the line assembly built from the tile types used to nondeterministically select r could grow to infinite length (in the limit) and never terminate, preventing the supertile from ever growing into a macrotile. For this reason, we say that this construction finitely self-assembles the macrotiles, as defined in [3].

To show that the behavior of the macrotiles is consistent with the definition of strong simulation, we first define an index function $I : A^{U_\tau} \dashrightarrow \mathbb{N}$ which takes as input an assembly over U_τ and returns the index of the tile set for which it is a macrotile simulation of some tile type, or is undefined if the assembly has not yet selected a tile set to simulate. I can do this by inspecting the supertile to find the number r which was selected (simply looking for tiles from the set in Figure 14) and running the same Turing machine program used by the construction to find the tile set and scaling factor (which is based on the running time of the Turing machine and the size of the tile set). Using I , we can determine which supertile representation function R_i^* to use to map a supertile over U_τ to a supertile over tile set T_i , for $0 \leq i < \infty$. Thus we can separate all supertiles $\alpha \in \mathcal{A}[\mathcal{S}]$ into sets S_i for $0 \leq i < \infty$ such that the supertiles in S_i represent macrotiles in the simulation of tile set T_i . Now we examine the physical features of macrotiles to prove the claim that those in S_i strongly simulate the system $\mathcal{T}_i = (T_i, \tau)$. Since each macrotile is constructed such that each of its edges are encoded with a “binary bump” segment corresponding to i (the tile set it is simulating), it is only possible for macrotiles’ glue segments to come into contact if they are both in S_i (i.e. simulating tiles from the same tile set). Thus for all macrotiles $\alpha \in S_i$ and $\beta \in S_j$ where $i \neq j$, they are simulating different tile sets and their glue segments may never come into contact so they cannot bind with strength $s > 0$. Note that the guaranteed ordering of growth of the features of each side is important for this, because the binary teeth must be in place before the external glues to ensure that correct exclusion will occur for macrotiles

from different sets. Since the glues are placed uniquely on the glue segment with strengths equivalent to the tiles they are simulating, macrotiles bind with strength $s > 0$ if and only if the tiles they are simulating bind with strength $s > 0$. Thus, by using the corresponding supertile representation function R_i^* for each S_i , it is shown that tile set T_i is strongly simulated at temperature τ .

Further, while the macrotiles for a given tile set faithfully strongly simulate that tile set, they do not interfere with any of the (infinite) simulations simultaneously occurring in parallel. Finally, we note that since the exterior of each macrotile contains no glues of strength > 1 along the flat glue portions of each side, it is impossible for the individual tiles of the simulator to attach to a macrotile. In fact, only once a supertile has grown into a macrotile and represents at least the binary teeth and glue portion of some side, can it interact with any other macrotile.

References

- [1] L. M. Adleman, Q. Cheng, A. Goel, M.-D. A. Huang, D. Kempe, P. M. de Espanés, and P. W. K. Rothmund. Combinatorial optimization problems in self-assembly. In *Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing*, pages 23–32, 2002.
- [2] P. Arrighi, N. Schabanel, and G. Theyssier. Intrinsic simulations between stochastic cellular automata. Technical report, arXiv, 2012. arXiv:1208.2763.
- [3] S. Cannon, E. D. Demaine, M. L. Demaine, S. Eisenstat, M. J. Patitz, R. Schweller, S. M. Summers, and A. Winslow. Two hands are better than one (up to constant factors). In *Proceedings of the Thirtieth International Symposium on Theoretical Aspects of Computer Science*, 2013. to appear.
- [4] E. G. Ch., P.-E. Meunier, I. Rapaport, and G. Theyssier. Communication complexity and intrinsic universality in cellular automata. *Theor. Comput. Sci.*, 412(1-2):2–21, 2011.
- [5] H.-L. Chen and D. Doty. Parallelism and time in hierarchical self-assembly. In *SODA 2012: Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1163–1182. SIAM, 2012.
- [6] Q. Cheng, G. Aggarwal, M. H. Goldwasser, M.-Y. Kao, R. T. Schweller, and P. M. de Espanés. Complexities for generalized models of self-assembly. *SIAM Journal on Computing*, 34:1493–1515, 2005.
- [7] M. Cook, Y. Fu, and R. T. Schweller. Temperature 1 self-assembly: Deterministic assembly in 3D and probabilistic assembly in 2D. In *SODA 2011: Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 570–589. SIAM, 2011.

- [8] M. Delorme, J. Mazoyer, N. Ollinger, and G. Theyssier. Bulking I: an abstract theory of bulking. *Theoretical Computer Science*, 412(30):3866–3880, 2011.
- [9] M. Delorme, J. Mazoyer, N. Ollinger, and G. Theyssier. Bulking II: Classifications of cellular automata. *Theor. Comput. Sci.*, 412(30):3881–3905, 2011.
- [10] E. D. Demaine, M. L. Demaine, S. P. Fekete, M. Ishaque, E. Rafalin, R. T. Schweller, and D. L. Souvaine. Staged self-assembly: nanomanufacture of arbitrary shapes with $O(1)$ glues. *Natural Computing*, 7(3):347–370, 2008.
- [11] E. D. Demaine, M. L. Demaine, S. P. Fekete, M. J. Patitz, R. T. Schweller, A. Winslow, and D. Woods. One tile to rule them all: simulating any Turing machine, tile assembly system, or tiling system with a single puzzle piece. Technical report, arXiv, 2012. arXiv:1212.4756 [cs.DS].
- [12] 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*, pages 439–446, Oct. 2012.
- [13] D. Doty, J. H. Lutz, M. J. Patitz, S. M. Summers, and D. Woods. Intrinsic universality in self-assembly. In *Proceedings of the 27th International Symposium on Theoretical Aspects of Computer Science*, pages 275–286, 2009.
- [14] B. Durand and Z. Róka. The game of life: universality revisited. In M. Delorme and J. Mazoyer, editors, *Cellular Automata*. Kluwer, 1999.
- [15] K. Fujibayashi, R. Hariadi, S. H. Park, E. Winfree, and S. Murata. Toward reliable algorithmic self-assembly of DNA tiles: A fixed-width cellular automaton pattern. *Nano Letters*, 8(7):1791–1797, 2007.
- [16] G. Lafitte and M. Weiss. Universal tilings. In W. Thomas and P. Weil, editors, *STACS 2007, 24th Annual Symposium on Theoretical Aspects of Computer Science, Aachen, Germany, February 22-24, 2007, Proceedings*, volume 4393 of *Lecture Notes in Computer Science*, pages 367–380. Springer, 2007.
- [17] G. Lafitte and M. Weiss. Simulations between tilings. In *Conference on Computability in Europe (CiE 2008), local proceedings*, pages 264–273, 2008.
- [18] G. Lafitte and M. Weiss. An almost totally universal tile set. In J. Chen and S. B. Cooper, editors, *Theory and Applications of Models of Computation, 6th Annual Conference, TAMC 2009, Changsha, China, May 18-22, 2009. Proceedings*, volume 5532 of *Lecture Notes in Computer Science*, pages 271–280. Springer, 2009.

- [19] J. I. Lathrop, J. H. Lutz, and S. M. Summers. Strict self-assembly of discrete Sierpinski triangles. *Theoretical Computer Science*, 410:384–405, 2009.
- [20] C. Luhrs. Polyomino-safe DNA self-assembly via block replacement. In A. Goel, F. C. Simmel, and P. Sosik, editors, *DNA14*, volume 5347 of *Lecture Notes in Computer Science*, pages 112–126. Springer, 2008.
- [21] K. Lund, A. T. Manzo, N. Dabby, N. Micholotti, A. Johnson-Buck, J. Nangreave, S. Taylor, R. Pei, M. N. Stojanovic, N. G. Walter, E. Winfree, and H. Yan. Molecular robots guided by prescriptive landscapes. *Nature*, 465:206–210, 2010.
- [22] P.-E. Meunier-asd, M. J. Patitz, S. M. Summers, G. Theyssier, A. Winslow, and D. Woods. Intrinsic universality in tile self-assembly requires cooperation. Technical report, arXiv, 2013. arXiv:1304.1679 [cs.CC].
- [23] N. Ollinger. Intrinsically universal cellular automata. In *The Complexity of Simple Programs, in Electronic Proceedings in Theoretical Computer Science*, volume 1, pages 199–204, 2008.
- [24] N. Ollinger and G. Richard. Four states are enough! *Theoretical Computer Science*, 412(1):22–32, 2011.
- [25] L. Qian and E. Winfree. Scaling up digital circuit computation with DNA strand displacement cascades. *Science*, 332(6034):1196, 2011.
- [26] L. Qian, E. Winfree, and J. Bruck. Neural network computation with DNA strand displacement cascades. *Nature*, 475(7356):368–372, 2011.
- [27] P. Rothmund. Folding DNA to create nanoscale shapes and patterns. *Nature*, 440(7082):297–302, 2006.
- [28] P. W. Rothmund, N. Papadakis, and E. Winfree. Algorithmic self-assembly of DNA Sierpinski triangles. *PLoS Biology*, 2(12):2041–2053, 2004.
- [29] P. W. K. Rothmund. *Theory and Experiments in Algorithmic Self-Assembly*. PhD thesis, University of Southern California, December 2001.
- [30] N. C. Seeman. Nucleic-acid junctions and lattices. *Journal of Theoretical Biology*, 99:237–247, 1982.
- [31] H. Wang. Proving theorems by pattern recognition – II. *The Bell System Technical Journal*, XL(1):1–41, 1961.
- [32] E. Winfree. *Algorithmic Self-Assembly of DNA*. PhD thesis, California Institute of Technology, June 1998.

- [33] E. Winfree, F. Liu, L. A. Wenzler, and N. C. Seeman. Design and self-assembly of two-dimensional DNA crystals. *Nature*, 394(6693):539–44, 1998.
- [34] B. Yurke, A. Turberfield, A. Mills Jr, F. Simmel, and J. Neumann. A DNA-fuelled molecular machine made of DNA. *Nature*, 406(6796):605–608, 2000.