

University of Texas Rio Grande Valley

ScholarWorks @ UTRGV

Computer Science Faculty Publications and
Presentations

College of Engineering and Computer Science

9-2023

Complexity of verification in self-assembly with prebuilt assemblies

David Caballero

The University of Texas Rio Grande Valley

Timothy Gomez

The University of Texas Rio Grande Valley

Robert Schweller

The University of Texas Rio Grande Valley, robert.schweller@utrgv.edu

Tim Wylie

The University of Texas Rio Grande Valley, timothy.wylie@utrgv.edu

Follow this and additional works at: https://scholarworks.utrgv.edu/cs_fac



Part of the [Computer Sciences Commons](#)

Recommended Citation

Caballero, David, Timothy Gomez, Robert Schweller, and Tim Wylie. "Complexity of verification in self-assembly with prebuilt assemblies." *Journal of Computer and System Sciences* 136 (2023): 1-16.
<https://doi.org/10.1016/j.jcss.2023.03.002>

This Article is brought to you for free and open access by the College of Engineering and Computer Science at ScholarWorks @ UTRGV. It has been accepted for inclusion in Computer Science Faculty Publications and Presentations by an authorized administrator of ScholarWorks @ UTRGV. For more information, please contact justin.white@utrgv.edu, william.flores01@utrgv.edu.

Complexity of Verification in Self-Assembly with Prebuilt Assemblies

David Caballero^a, Timothy Gomez^{a,b}, Robert Schweller^a, Tim Wylie^a

^a*Department of Computer Science, University of Texas Rio Grande Valley, 1201 W
University Dr., Edinburg, 78539, TX, USA*

^b*Department of Electrical Engineering and Computer Science, Massachusetts Institute of
Technology, 77 Massachusetts Ave., Cambridge, 02139, MA, USA*

Abstract

We analyze the complexity of two fundamental verification problems within a generalization of the two-handed tile self-assembly model (2HAM) where initial system assemblies are not restricted to be singleton tiles, but may be larger prebuilt assemblies. Within this model we consider the *producibility* problem, which asks if a given tile system builds, or produces, a given assembly, and the *unique assembly verification* (UAV) problem, which asks if a given system *uniquely* produces a given assembly. We show that producibility is NP-complete and UAV is coNP^{NP} -complete even when the initial assembly size and temperature threshold are both bounded by a constant. This is in stark contrast to results in the standard model with singleton input tiles where producibility is in P and UAV is coNP-complete with constant temperature. We further provide preliminary polynomial time results for producibility and UAV in the case of 1-dimensional *linear* assemblies with pre-built assemblies, as well as extend our results to the abstract Tile Assembly Model (aTAM) with constant-size attachable assemblies.

Keywords: self-assembly, 2HAM, two-handed assembly, hierarchical assembly, producibility, assembly verification, prebuilt assemblies

*This research was supported in part by National Science Foundation Grant CCF-1817602.

1. Introduction

Self-assembly is the process by which a system of simple particles autonomously come together to form complex structures. *Algorithmic* self-assembly studies scenarios in which dynamics of system molecules encode computation, allowing for algorithmic control of the self-assembly of matter. A premiere model for the study of algorithmic self-assembly is the *tile self-assembly model* [1, 2], in which system monomers are modeled as four-sided Wang tiles that randomly collide and combine based on matching tile edges and a given bonding threshold called the *temperature*. Tile self-assembly has received substantial theoretical consideration (see [3, 4, 5] for surveys and recent results) as well as various experimental DNA implementations [6, 7, 8].

In this paper we focus on a specific generalization of the standard 2-handed tile self-assembly model (2HAM) in which we permit initial assemblies to consist of *prebuilt* assemblies of more than one tile. The motivation for studying such a generalization is strong. First, some of the most successful implementations of algorithmic DNA self-assembly utilize a combination of singleton DNA tiles mixed with larger prebuilt assemblies. For example, the experimental implementations of DNA tile counters [7] and the 21 DNA tile circuits implemented in [8] both utilize a combination of single tiles seeded with a larger prebuilt DNA origami structure encoding the program input. In [7], they additionally include a mix of singleton and prebuilt domino assemblies among the system's tile set. Second, the inclusion of prebuilt shapes of different geometries and sizes allows for the potential application of *steric hindrance*, in which geometric blocking of potential attachments is used to control algorithmic growth, as seen in the theoretical works of [9, 10, 11], and the experimental works of [12, 13]. These examples suggest that consideration of shapes more general than uniform single squares has the promise to allow for improved computational power and efficiency of self-assembled systems. We use geometric blocking techniques in our prebuilt construction, and have more discussion on it there.

Given the importance of understanding self-assembly with prebuilt initial assemblies, we consider the complexity of two fundamental computational questions related to verifying the correctness of such systems. First is the *Producibility* problem, which asks if a given tile system can build/produce a given assembly. The second is the *Unique Assembly Verification* (UAV) problem, which asks if a given tile system *uniquely* produces a given assembly, i.e., produces only one terminal assembly. For the producibility problem,

the 2HAM with just single-tile initial assemblies has a polynomial time solution [14], whereas we show NP-completeness when prebuilt assemblies are permitted. In the case of the UAV problem with singleton tile initial assemblies and a constant temperature, the problem was recently shown to be coNP-complete [15]. Previously, membership was known to be in coNP [1], but was only known to be coNP-complete for larger temperature thresholds [16].

1.1. Previous work

The model used here differs from the polyTAM model of [10] in that the set of starting elements in our system are defined as assemblies made up of multiple tiles. In the polyTAM, the set of elements are single tiles that are allowed to be larger than a unit square. We are attempting to model the situations where smaller components may have preassembled into larger structures prior to being introduced to the system, which is beneficial since production of some individual elements may be more difficult and costly, or due to standard practices where the entire assembly process may be implemented through multiple experiments. This is similar to the staged model of self-assembly [9], however, in that model all the bins are usually assumed to have the same temperature. Thus, any assemblies built in early stages must be producible. Here, we only require that the input assemblies are stable.

Verification problems have been well-studied in many models of Tile Self-Assembly. In the Abstract Tile Assembly model (aTAM), both the producibility and UAV problem are solvable in polynomial time [17]. When allowing negative (or repulsive) glues the UAV problem becomes undecidable due to detachment [18], but when restricted to growth-only systems (no detachment can occur) the problem is coNP-complete [19]. Producibility verification in the 2-handed assembly model at any temperature, along with UAV for temperature 1, are both solvable in polynomial time [14]. Membership of producibility for general 2HAM systems was shown to be in the class coNP in [1]. They also showed that UAV is coNP-completeness when one step into the 3^{rd} dimension is allowed. By allowing the temperature to be a part of the input, UAV was shown to be coNP-complete [16] even in 2 dimensions. The most recent results resolves this line of inquiry and shows that UAV in the 2HAM, even with a constant temperature of two, is coNP-complete [15].

More powerful generalizations have shown an increase in complexity of the UAV problem, such as the aforementioned staged model, which for arbitrary

stages is PSPACE-complete [20, 21] and is coNP-complete even with a single stage [15]. Multiple-handed assembly (the k -handed model), allows for k different assemblies to be combined in a single step as long as the final assembly was stable. UAV in this generalization was shown to be coNP-complete if the number of hands, k , is encoded in unary as input, and coNEXP-complete if encoded in binary [22].

Another generalization is Tile Automata [23], which merges ideas from Cellular Automata and the 2HAM. The UAV problem was shown to be coNP^{NP}-complete even with the restrictions of freezing (A tile only changes states a finite number of times) and growth only (no detachment) [24]. Note the class coNP^{NP} is the class of problems solvable in coNP with access to an NP oracle. This is also referred to as Π_2^P , or the 2nd layer of the polynomial hierarchy, and more details can be found in [25].

1.2. Contributions

Table 1 and 2 highlight our results in relation to previous research with producibility (Table 1) and Unique Assembly Verification (Table 2). With prebuilt assemblies, we show that UAV with a constant temperature is actually coNP^{NP}-complete. In both scenarios, our hardness results hold even for prebuilt assemblies of a bounded $\mathcal{O}(1)$ size and $\mathcal{O}(1)$ -bounded temperature thresholds. We accompany these results with a preliminary exploration of the producibility and UAV problems when restricted to 1-dimensional *linear* assemblies with pre-built assemblies, and provide polynomial time solutions. Finally, we show that producibility is NP-complete in the aTAM with pre-built assemblies and that UAV is coNP-complete.

A version of this paper was originally published in [26]. However, this version has several important improvements and additions. We have included proofs that were previously omitted. Specifically, the proofs related to 1D UAV and producibility. Besides additional text and providing missing details in the rest of the paper, we also prove some interesting results related to the aTAM with prebuilt assemblies. We provide some smaller prebuilt assemblies, and motivate investigation into the minimal prebuilt assembly size that maintains the same complexity for the problem. This is further motivated by the recent result showing that with singleton tiles, UAV is only coNP-complete instead of coNP^{NP}-complete with prebuilt assemblies.

Model	Input Assembly Size	Temp.	Result	Reference
aTAM	Single Tiles	τ	P	[17]
2HAM	Single Tiles	τ	P	[14]
2HAM	Constant	2	NP-complete	Thm. 3.1

Table 1: Complexity of verifying producibility of an assembly in various models. The Assembly Size column indicates the size of the assemblies in the initial assembly set. Previous work has studied the case when only single tiles are allowed. Our results allow for up to constant sized assemblies.

Model	Input Assembly	Temp.	Result	Reference
aTAM	Single Tiles	τ	P	[17]
Neg. aTAM G.O.*	Single Tiles	2	coNP-complete	[19]
aTAM	Constant	2	coNP-complete	Cor. 6.2
2HAM	Single Tiles	1	P	[14]
2HAM	Single Tiles	2	coNP-complete	[15]
2HAM	Single Tiles	τ	coNP-complete	[16]
2HAM 3D	Single Tiles	2	coNP-complete	[1]
2HAM	Constant	2	coNP ^{NP} -complete	Thm. 4.3

Table 2: Complexity results of Unique Assembly Verification in the aTAM and the 2HAM. The results for 2HAM with variable temperature and in 3D follow from the result in [15], but are included since they were proven first. UAV is undecidable in the negative aTAM, but coNP-complete with negative glues if the system never allows detachment (*growth only).

1.3. Outline

The paper is organized as follows. In Section 2, we overview the model and problem definitions. The main results then follow with Section 3 covering the complexity of producibility and Section 4 discussing the hardness of unique assembly verification. Following, we discuss the problems for one-dimensional 2HAM systems in Section 5. Section 6 discusses an extension of the results to look at the aTAM with prebuilt assemblies. We wrap the paper up by discussing some interesting directions for future work in Section 7 and a conclusion in Section 8.

2. Preliminaries

In this section we overview the basic definitions related to the two-handed self-assembly model and the verification problems under consideration. When used in reference to an assembly, all cardinal directions are assumed to be in standard orientation with north being the top.

2.1. Self-Assembly Model

Tiles. A *tile* is a non-rotatable unit square with each edge labeled with a *glue* from a set Σ . Each pair of glues $g_1, g_2 \in \Sigma$ has a non-negative integer *strength* $\text{str}(g_1, g_2)$.

Configurations, bond graphs, and stability. A *configuration* is a partial function $A : \mathbb{Z}^2 \rightarrow T$ for some set of tiles T , i.e., an arrangement of tiles on a square grid. For a given configuration A , define the *bond graph* G_A to be the weighted grid graph in which each element of the domain of A , $\text{dom}(A)$, is a vertex, and the weight of the edge between a pair of tiles is equal to the strength of the coincident glue pair. A configuration is said to be τ -*stable* for positive integer τ if every edge cut of G_A has strength at least τ , and is τ -*unstable* otherwise.

Assemblies. For a configuration A and vector $\vec{u} = \langle u_x, u_y \rangle$ with $u_x, u_y \in \mathbb{Z}$, $A + \vec{u}$ denotes the configuration $A \circ f$, where $f(x, y) = (x + u_x, y + u_y)$. For two configurations A and B , B is a *translation* of A , written $B \simeq A$, provided that $B = A + \vec{u}$ for some vector \vec{u} . For a configuration A , the *assembly* of A is the set $\tilde{A} = \{B : B \simeq A\}$. An assembly \tilde{A} is a *subassembly* of an assembly \tilde{B} , denoted $\tilde{A} \sqsubseteq \tilde{B}$, provided that there exists an $A \in \tilde{A}$ and $B \in \tilde{B}$ such that $A \subseteq B$. An assembly is τ -*stable* provided the configurations it contains are τ -stable. Assemblies \tilde{A} and \tilde{B} are τ -*combinable* into an assembly \tilde{C} provided there exist $A \in \tilde{A}$, $B \in \tilde{B}$, and $C \in \tilde{C}$ such that $A \cup B = C$, $A \cap B = \emptyset$, and \tilde{C} is τ -stable.

Two-handed assembly. A Two-handed assembly system is an ordered tuple (S, τ) where S is a set of *initial* assemblies and τ is a positive integer parameter called the *temperature*. Each assembly in S must be τ -stable. For a system (S, τ) , the set of *producible* assemblies $P'_{(S, \tau)}$ is defined recursively as follows:

1. $S \subseteq P'_{(S, \tau)}$.
2. If $A, B \in P'_{(S, \tau)}$ are τ -combinable into C , then $C \in P'_{(S, \tau)}$.

A producible assembly is *terminal* if it is not τ -combinable with any other producible assembly, and the set of all terminal assemblies of a system (S, τ)

is denoted $P_{(S,\tau)}$. Intuitively, $P'_{(S,\tau)}$ represents the set of all possible assemblies that can self-assemble from the initial set S , whereas $P_{(S,\tau)}$ represents only the set of assemblies that cannot grow any further. An assembly A is *uniquely produced* if $P_{(S,\tau)} = \{A\}$ and for each $B \in P'_{(S,\tau)}$ $B \sqsubseteq A$.

2.2. Problems

As discussed in the introduction, we focus on two fundamental problems for self-assembly: producibility and unique assembly verification. The formal definitions are below.

Problem 2.1 (Producibility). Given a 2HAM system $\Gamma = (S, \tau)$ and an assembly A , is A a producible assembly of Γ ?

Problem 2.2 (Unique Assembly Verification (UAV)). Given a 2HAM system $\Gamma = (S, \tau)$ and an assembly A , is A uniquely produced by Γ ?

3. Producibility Hardness

In this section, we show that the producibility problem is NP-complete if the initial set of assemblies may include assemblies larger than singleton tiles. The hardness is derived by reducing from 3SAT and holds even if assembly size and system temperature is bounded by a constant. Our construction extends to the seeded *abstract tile assembly model* where there is a seed tile, and elements of the tile set (in this case assembly set) attach one at a time to the growing seed.

3.1. Overview

We reduce from 3SAT¹, which asks whether a given 3CNF formula ϕ is satisfiable. Let $|V|$ and $|C|$ be the number of variables and clauses in ϕ , respectively. The reduction creates an instance of producibility (Γ, A) with $\tau = 2$, such that Γ produces the target assembly A iff ϕ is satisfiable. The target assembly is a rectangle shown and described in Figure 4a. The assemblies in the reduction can be divided into two groups: edge assemblies (Figure 1a) and macroblocks (Figure 1b). There are two variable assemblies for each bit that each correspond to an assignment of 0 or 1 based on the

¹Note this reduction technique works for general SAT but we reduce from 3SAT because it is well known.

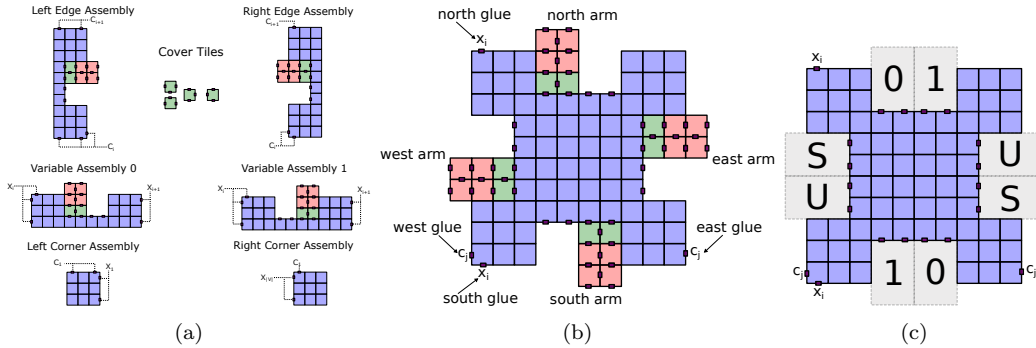


Figure 1: (a) Edge Assemblies used to construct the frame of the assembly. For each clause we include a left and right edge assembly. For each variable we include two variable assemblies representing 0 and 1. We include a single left corner assembly along with a right corner assembly. The filler tiles are used to fill in holes between attached macroblocks. The glue labels x_i and C_j are unique to each possible choice of variable and clause for the given SAT formula. (b) A single macroblock $m_{i,j}(0,U,U)$ with outer glues labeled. The north and south glues connect to other macroblocks that represent the same variable. The east and west glues connects to other macroblocks that represent the same clause. (c) Arm position labels on a macroblock. Opposite sides have complementary values to allow for attachment.

position of the 3×2 *arm* section of the assembly (the green and red tiles constitute the 3×2 arm assemblies in the figures). As shown in Figure 1c, each arm position represents a certain value. For vertical arms, the position represents either 0 or 1. For horizontal arms, their position is either U or S , meaning unsatisfied or satisfied.

These assemblies will combine to form an assembly for each possible assignment to ϕ (Figure 2a). We include a unique left edge assembly for each clause whose arm is always in the top position representing that the clause is currently not yet satisfied. A left edge assembly can attach to an assembly that encodes an assignment to ϕ . This starts to form an L shaped frame assembly as shown in Figure 2b. Macroblocks may attach to this frame if it has complementary arms to the currently growing frame assembly. As shown in Figure 2c, a macroblock can attach using two strength-1 glues on its east and south sides. If the arms have complementary positions the attachment will be able to take place. However, if the arms overlap, the macroblock is geometrically blocked from attaching, and thus not allowed.

The final challenge for designing this reduction is there must exist a single assembly that is produced regardless of the satisfying assignment. This




Figure 2: (a) Variable gadgets can non-deterministically build a frame assembly for each possible assignment to ϕ . (b) Example of the assembly made for assignment 1101 with a single left edge assembly attached. (c) If the macroblock has complementary arm positions it will be able to attach to the frame assembly. If the macroblocks do not have matching arm positions the attachment will be geometrically blocked and cannot occur.

means we must hide the values passed between macroblocks. We do this by including a certain subset of the tiles which will fill any spaces left between macroblocks.

3.2. Macroblocks

A single macroblock can be seen in Figure 1b and has two parts: the body that contains glues to allow attachment (blue), and four arms which encode ϕ (green/red). Each arm on the macroblock encodes a single bit of information by being in one of two positions. We call these positions “0” and “1” for the north/south arms and “U” and “S” (unsatisfied/satisfied) for the east/west arms respectively (Figure 1c).

We denote a macroblock representing a variable/clause pair (v_i, c_j) by its glues and arm positions as $m_{i,j}(b, w, e)$ where $b \in \{0, 1\}$ is the position of the north/south arm, $w \in \{U, S\}$ is the position of the west arm, and $e \in \{U, S\}$ is the position of the east arm. Each macroblock has a single strength-1 glue on each side (macroblocks representing the last clause do not contain glues or arms on their north side). The glues indicate which variable and clause pair this macroblock represents. The north and south glues relate to the variable, and the east and west glues relate to the clause. The south and west glues allow for cooperative attachment to an assembly that already contains macroblocks $m_{i-1,j}$ and $m_{i,j-1}$. The north and east glues allow for attachment of the next macroblocks.




Figure 3: Possible Macroblocks that make up $\mathcal{M}_{i,j}$. Arm positions represent the value assigned to x_i and whether or not c_j has been satisfied. There will always be 4 macroblocks in each set. The left pair of macroblocks are always included and will attach if a clause is already satisfied. The remaining macroblocks attach if the clause is not yet satisfied, and their arm positions depend on ϕ . If the positive literal x_i is in c_j , then $m_{i,j}(1, U, S) \in \mathcal{M}_{i,j}$, otherwise $m_{i,j}(1, U, U) \in \mathcal{M}_{i,j}$. If the negative literal \bar{x}_i is in c_j , then $m_{i,j}(0, U, S) \in \mathcal{M}_{i,j}$, otherwise $m_{i,j}(0, U, U) \in \mathcal{M}_{i,j}$. The final clause macroblocks do not contain the north arms and just have the area filled in.

Each variable/clause pair (v_i, c_j) has a set $\mathcal{M}_{i,j}$ of four macroblocks associated with it (shown in Figure 3). The exact macroblocks that are included depends on whether x_i or \bar{x}_i is present in the j^{th} clause. The macroblocks $m_{i,j}(0, S, S)$ and $m_{i,j}(1, S, S)$ are always included since the assignment of a variable can not change a clause from being satisfied to unsatisfied. If the the positive literal v_i appears in c_j , then we include the macroblocks $m_{i,j}(1, U, S)$, or $m_{i,j}(1, U, U)$ if it does not. If the negative literal \bar{v}_i appears in c_j , include the macroblock $m_{i,j}(0, U, S)$, or $m_{i,j}(0, U, U)$ if it does not.

3.3. Computing Clauses

The left edge assembly starts with an arm in the U position. Macroblocks maintain this position (Figure 4b) until a macroblock attaches that satisfies the clause, and changes the arm to an S position (Figure 4c). Once a row of the assembly is complete (Figure 4d), if the horizontal arm is in the satisfied position, the right edge assembly can attach cooperatively and complete the row (Figure 4e). For a right edge assembly to attach, the clause must be satisfied and the assembly below it (either another right edge assembly or the right corner assembly) must attach as well. Thus, the right edge assembly only attaches if the clause it represents is satisfied.

Each row has multiple size-2 holes between macroblocks. Filler tiles cannot attach to macroblocks on their own due to the temperature of the system.

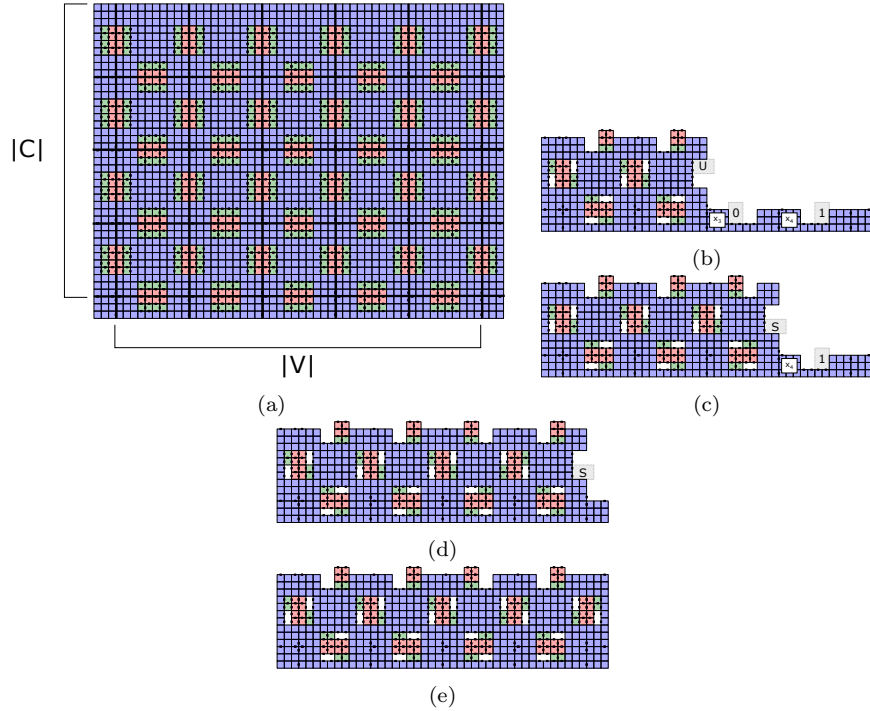


Figure 4: (a) Target assembly for producibility in the 2HAM with prebuilt assemblies. Target assembly is a $10|C| + 3$ by $10|V| + 6$ rectangle. Smaller rectangles between tiles denote strength-1 glues. Glues between blue tiles are not shown. Each blue tile shares a strength-2 glue with neighboring blue tiles. The exceptions are tiles separated by the thicker borders that do not share a glue unless shown. (b) A frame assembly with macroblocks attached. Here, $c_1 = \bar{x}_1 \vee \bar{x}_3 \vee x_4$. (c) Here $x_3 = 0$ satisfies the clause so this macroblock that attaches has its arm in the S position. (d) The macroblocks that attach after always have their arms in the S position. (e) The right edge assembly can attach to the last macroblock since its arm position is in the S position completing the row.

However, once two macroblocks are attached, the tiles can cooperatively bind across the hole with strength-1 glues from each side (Figure 5b). As shown in Figure 5c, this hides the previously used arm positions. The exposed northern arms also continue to encode the input string so the next clause can be computed after the attachment of the next left edge assembly. Note again that the final row of a macroblock does not contain north arms.

3.4. Complexity

Using the construction above, we can show that even with constant-sized prebuilt assemblies, producibility is NP-complete in the 2HAM. This means

even determining if a tile system with prebuilt assemblies can ever build a specific assembly is hard.

Theorem 3.1. The producibility problem in the 2HAM with prebuilt assemblies of size $\mathcal{O}(1)$ is NP-complete with $\tau = 2$ and an initial assembly set containing $\mathcal{O}(|V||C|)$ distinct assemblies.

Proof. For membership in the class NP consider an assembly tree v for a producible assembly A . To verify A is producible we may use v as a certificate. If v is a valid tree (each combination is legal) and each leaf is in the input set I , then A is producible.

To show NP-hardness we reduce from 3SAT. Given a formula ϕ in 3CNF form with $|V|$ variables and $|C|$ clauses. Our target assembly A is the rectangle described above made from macroblocks and edge assemblies with each of the spaces between arms completely filled with tiles. The input assembly set includes I , $|C|$ left edge assemblies with their arms in the unsatisfied position, and $|C|$ right edge assemblies in their satisfied position. For each variable, two input assemblies representing 0 and 1 assignments are included. The clauses are encoded by the selection of macroblock arm combinations. A set of 4 macroblocks are included in I for each variable and clause combination, so there are a total of $\mathcal{O}(|V||C|)$ macroblocks. This results in a total initial assembly set size of $\mathcal{O}(|V||C|)$, and each assembly is constant sized.

The starting assemblies, I , will grow into A if and only if ϕ is satisfiable. If ϕ is satisfiable by some assignment x , then the ‘L’ shaped frame assembly representing x will grow by attaching macroblocks. Since x satisfies ϕ each clause will eventually have its arm position changed to satisfied allowing for all the right edge assemblies to attach. The single tiles will then fill in the spaces to complete A . If A is producible then there exists some ‘L’ shaped frame assembled that grew into A . The only way this frame could have grown into A is if the position of the arms on the input assemblies represented a string x that satisfied each clause meaning ϕ is satisfiable. \square

4. Unique Assembly Verification is coNP^{NP} -complete

In this section we show coNP^{NP} -completeness of the Unique Assembly verification problem in the 2HAM with constant-sized prebuilt assemblies. We start by proving UAV is in the class of problems solvable by a nondeterministic algorithm with access to an oracle for a problem in the class NP. We then show hardness by reducing from $\forall\text{ESAT}$. This is an extension of

the reduction shown in the previous section, and further, we utilize similar techniques used in previous reductions in the 2HAM and Tile Automata [24, 20].

Lemma 4.1. The Unique Assembly Verification problem in the 2HAM with prebuilt assemblies is in coNP^{NP} .

Proof. Given an instance (Γ, A) , refer to a “rogue assembly” R as a producible assembly in the system Γ that is either (1) not a subassembly of the target assembly A , $R \not\sqsubseteq A$, or (2) a strict subassembly of A and terminal, i.e., $R \sqsubset A$ and $R \in P_{(S,\tau)}$. The following nondeterministic algorithm solves UAV.

1. Nondeterministically build an assembly B of size $|B| \leq 2|A|$.
2. If B is a rogue assembly, reject.

It suffices to check all assemblies B up to size $2|A|$ since any assembly of size $> 2|A|$ must have been built from at least one other assembly B' , s.t. $|A| < |B'| \leq 2|A|$. B' is a rogue assembly itself and will be accounted for in a different branch of the computation. B must still be checked to see if it is a rogue assembly. The first condition can be verified in polynomial time by checking if B is a subassembly of A by simply checking if each tile in B exists in A at the same location. The second condition can be checked using an NP oracle that answers the following: “Does there exist an assembly C , $|C| \leq |A|$, such that C can attach to B ?”. This problem can be solved by an NP machine that nondeterministically builds an assembly C up to size A and attempts to attach it to B . If any C can attach to B , B is not terminal. In other words, for “yes” instances such a C serves as a certificate for NP membership. If any branch finds a rogue assembly, the co-nondeterministic machine will reject. \square

4.1. Reduction Overview

Definition 4.2 ($\forall\text{ESAT}$). Given an n -bit Boolean formula $\phi(x_1, x_2 \dots x_n)$ with the inputs divided into two sets X and Y , for every assignment to X , does there exist an assignment to Y such that $\phi(X, Y) = 1$?

We show this problem is coNP^{NP} -hard by reducing from $\forall\text{ESAT}$. An overview of the important assemblies and processes are shown in Figure 5a. The same construction used in the previous reduction is used to create exponentially many ‘SAT assemblies’, each of which evaluates the Boolean formula

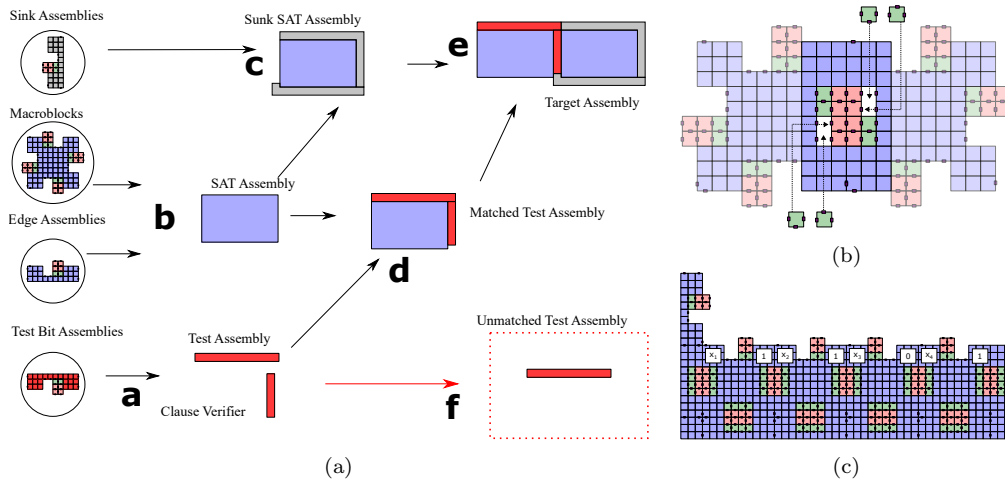


Figure 5: (a.a) Test bit assemblies come together to build a test assembly for all possible assignments of the variables in X . Clause Verifier assemblies may attach to SAT assemblies that satisfied all clauses. (a.b) Macroblocks and edge assemblies from the previous reduction are used to create SAT assemblies that evaluate the formula for every assignment of all the variables. (a.c) The sink assemblies begin attaching to SAT assemblies, ensuring they all grow into the target assembly. (a.d) A test assembly will attach to a SAT assembly that satisfies the formula and has matching assignments to the variables in X . (a.e) Matched test assemblies and sunk SAT assemblies attach to each other forming the target assembly. (a.f) Any test assemblies that do not find a SAT assembly to attach to (and are thus unmatched) are terminal. If any unmatched test assemblies exist, the instance of UAV is false. (b) Once two macroblocks attach, the green filler tiles are able to cooperatively attach using one glue on the macroblock, and the other glue from the red tiles of the arms from the other macroblock. The filling process hides the information that was passed. (c) Another left edge assembly may attach above the first. Since the north arms of macroblocks encode the variable assignment, the second clause may be computed in the same way as the first.

on one of its input assignments. A SAT assembly is shown in Figure 6b. We do not include right edge assemblies, so SAT assemblies that have finished computing have exposed arms on their right side denoting whether or not each clause was satisfied. The assembly has exposed arms to the north, above variables in X , that represent their assigned values. Variables in Y still have no arms on their north side. We construct a test assembly (Figure 7a) for each possible assignment to the variables in X (Figure 8a) along with a *clause verifier* assembly. The clause verifier assemblies are made of right edge assemblies with their arms in the S position. Each test assembly can attach to any SAT assembly with matching assignments to X (complementary arm

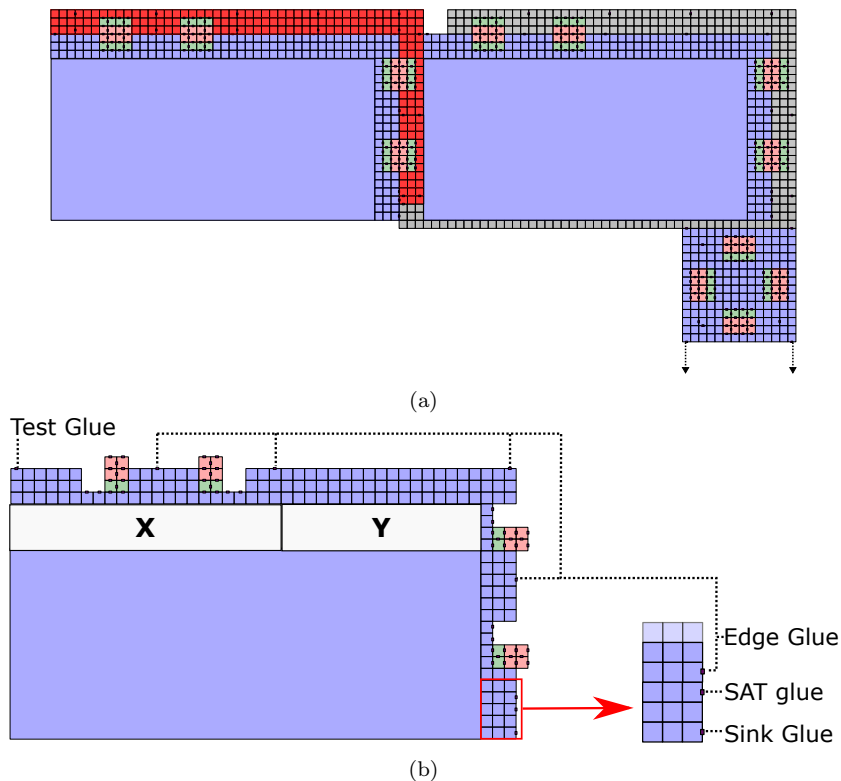


Figure 6: (a) Target Assembly for UAV. Assembly can be divided into three parts, Test assembly with satisfied SAT assembly, Sunk SAT assembly, and a column of clean up frames (b) SAT assembly built using the previous reduction with additional northern arms for the variables in X . Note we do not add right edge assemblies, and an arm is exposed that shows whether a clause has been satisfied.

positions) if the clause verifier has already attached. Other assemblies are included that ‘sink’ all assemblies to the target except for test assemblies that did not attach to a SAT assembly. Test assemblies may build into the target assembly if and only if they find a matching SAT assembly. The key question about the system is “For all test assemblies, does there exist a compatible SAT assembly where all clauses evaluated to true?”

This system uniquely constructs the target assembly if and only if the instance of $\forall\exists\text{SAT}$ is true. If the instance is false, there exists an assignment to the variables in X where no satisfying assignment of Y exists. In this case, the test assembly representing that assignment will be terminal, which means the system does not uniquely produce the target assembly.

4.2. SAT Assembly

We use the assemblies from the previous reduction to compute all satisfying assignments to ϕ . We use the same macroblocks, input assemblies, and left edge assemblies, however, we do not include any right edge assemblies or the right corner assembly. A frame assembly is constructed for each assignment to ϕ . By attaching macroblocks with matching geometry, the assembly process computes whether or not the assignment satisfies the clauses. In this construction, we mark the assignment to variables in X by including northern arms to the top most macroblocks for those variables instead of omitting them as in the previous construction. This results in a final assembly that has its assignment to X and the computed value of each clause being encoded in the exposed arm positions.

The assembly can be seen in Figure 6b with glues labeled. The exposed glues all have strength-2. In the bottom right corner of the assembly, the last variable assembly has two glues. The bottom glue is called the sink glue and is one of the glues the sink assembly uses to attach to a SAT assembly. The glue above it, called the SAT glue, is used by both the left sink base assembly and the test assembly. The next glue appears on each macroblock with an exposed arm, and on the northern side of the rightmost macroblock. This glue allows for sink assemblies to attach and cover exposed arms to reach the target assembly. The test glue is the last glue on this assembly and appears in the top left corner. The test assembly uses this glue with the SAT glue to attach to a SAT assembly with matching geometry.

4.3. Test Assembly

A test assembly is constructed for each possible assignment to X starting from constant sized test bit assemblies shown in Figure 7a. For each variable in X , we include two test bit assemblies. For each variable in Y , we include a blank test bit assembly, which is a 3×10 rectangle. This row is capped by left and right corner test assemblies. The left and right corner assemblies have a strength-1 glue. The clause verifier assembly is built using right edge assemblies. These assemblies all have their arms in the satisfied position. The north most assembly has a strength-1 glue on its left side to attach to SAT assemblies and another on its north side to allow test assemblies to attach. They connect downward to a 1×3 test cap assembly. The test cap has the strength-1 SAT glue on its left side to allow a test assembly to cooperatively bind to a SAT assembly with all arms in the satisfied position (Figure 7b). On its south side it has the test-sink glue which will be used to attach to a

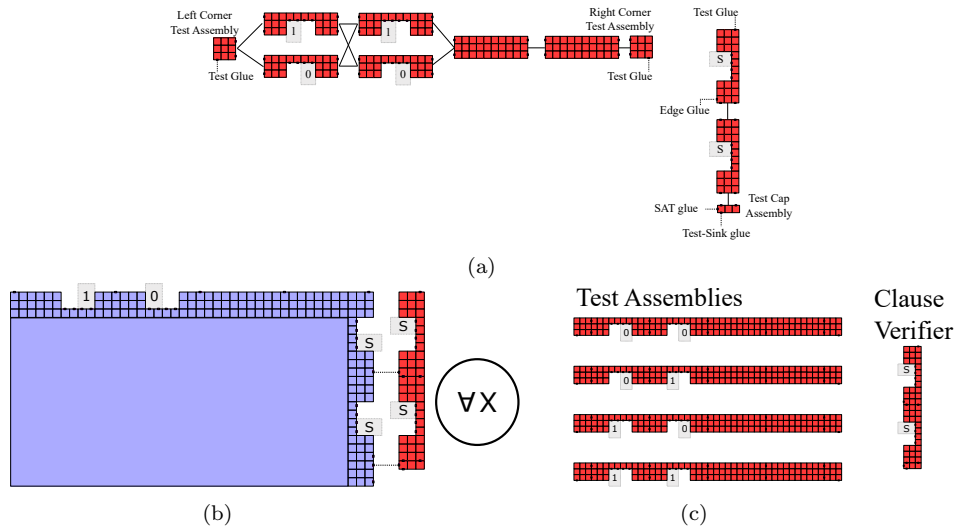


Figure 7: (a) Test bit assemblies nondeterministically construct a test assembly for each assignment to the variables in X . Right edge assemblies with arms in the satisfied positions are used to build the clause verifier. Only the north most edge assembly has an edge glue to allow a clause verifier to attach to a SAT assembly. A test assembly may attach to a SAT assembly with a clause checker attached using the test glues on their south side. The Test-Sink glue is used to cooperatively attach to a sink assembly once the test assembly is matched with a SAT assembly. (b) A clause verifier may attach to SAT assemblies that have their arms in the satisfied position. (c) Test Assemblies that are created for an X that contains 2 variables along with the single clause verifier.

sink assembly cooperatively once a SAT assembly is found. This assembly process creates a test assembly for each possible assignment to the variables in X . The example of test assemblies for an instance where $|X| = 2$ and $|Y| = 2$ can be seen in Figure 7c.

The test assembly has two exposed glues on opposite ends of the assembly. Thus, the assembly cannot attach to a SAT assembly until it is completely constructed and a clause verifier assembly has attached. A clause verifier assembly may only attach to SAT assembly with matching arm positions, i.e., SAT assemblies that satisfied all clauses. A test assembly may only attach to SAT assemblies with a clause verifier that has the same assignment to X . Figure 8a shows an example test assembly and the four possible SAT assemblies it could attach to. The test assembly is terminal if there does not exist a SAT assembly with the same assignment to X that satisfies all clauses. A terminal test assembly can be seen in Figure 8b. We call these

terminal test assemblies *unmatched* test assemblies. Since we construct a SAT assembly for each possible assignment to the formula ϕ , if the test assembly representing a partial assignment X is terminal, then there does not exist a remaining assignment to the variables in Y that satisfies the formula. Thus, the instance of UAV would be false.

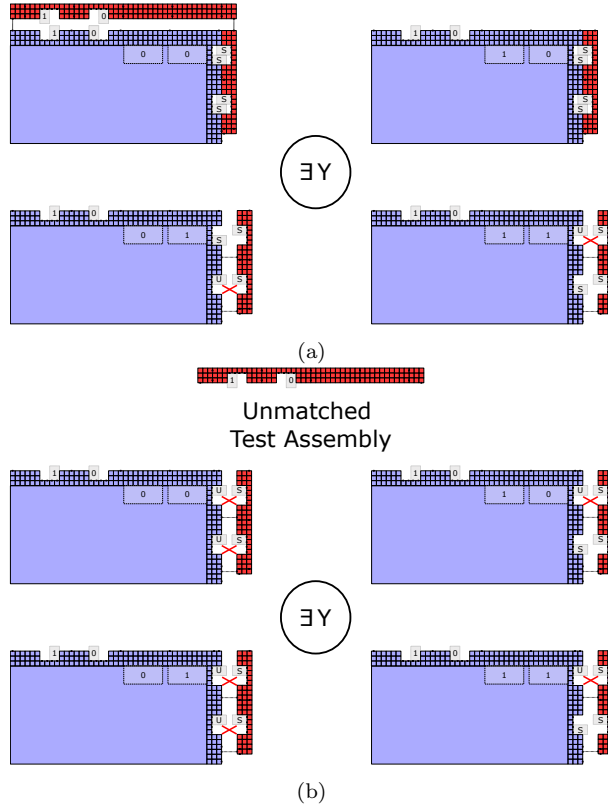


Figure 8: (a) If there exists a satisfying assignment to Y when $X = 10$, the test assembly with those arms can attach to a SAT assembly. (b) If there does not exist a satisfying assignment, all the SAT assemblies will have at least one arm not in the S position that will geometrically block the test assembly's arm. This test assembly will be terminal so the answer to UAV will be no.

4.4. Sink

Since our goal is to design a system that uniquely constructs an assembly when the instance of $\forall\exists\text{SAT}$ is true, we will *sink* assemblies representing a non-satisfying assignment to the target assembly. This ensures that each

assembly (besides unmatched test assemblies) must eventually grow into the target assembly thus sinking all other producible assemblies to our target. We do so via sink assemblies and macroblock frames, as shown in Figure 9a. The first sink assembly, the right sink base assembly, is similar to the right corner assembly of the previous section but is of height 4. Sink base tiles are single tiles that may attach to the right base assembly on its left side bottom row, which eventually connects to the left base assembly.

The right sink base assembly attaches to any SAT assembly that has not attached to a test assembly using the sink and test-sink glues. We call the attachments that occur after this the *Sink Process*. During the Sink Process, sink edge assemblies attach cooperatively with the right sink base assembly and with the SAT assembly. This allows for the sink edge assemblies to attach one-by-one and ‘cover’ each exposed arm on the SAT assembly. The last sink edge (northern most) is slightly longer to allow for the horizontal sink edges to attach across the top of the assembly. The left sink base assembly cooperatively attaches to test assemblies that have already attached to SAT assemblies using the SAT glue on its right side and the test-sink glue on its north side.

The last assembly type that has not been accounted for in the final assembly are any unused macroblocks. In the previous reduction, any unused macroblocks are terminal since they are never used in the computation. In this reduction, we cannot have any other terminal assembly, so these must be included in our target assembly. We do this by adding frames to store the macroblocks. For each macroblock we include in our input set, we also include a clean up frame (Figure 9b). Any macroblock is now not terminal as it can attach to the clean up frame. These frames are enumerated and attach in order to the south side of the sink assembly in a single column (Figure 9c).

Theorem 4.3. The Unique Assembly Verification problem in the 2HAM with prebuilt assemblies of size $\mathcal{O}(1)$ is coNP^{NP} -complete with an initial assembly set size of $\mathcal{O}(|V||C|)$ and $\tau = 2$.

Proof. We show membership in Lemma 4.1. Given an instance of $\forall\text{ESAT}$ with a formula $\phi(x_1, x_2, \dots, x_n)$ we create a 2HAM system S that uniquely assembles a target assembly S if and only if the instance of $\forall\text{ESAT}$ is true. If the instance of $\forall\text{ESAT}$ is false then S will also produce a test assembly that is terminal.

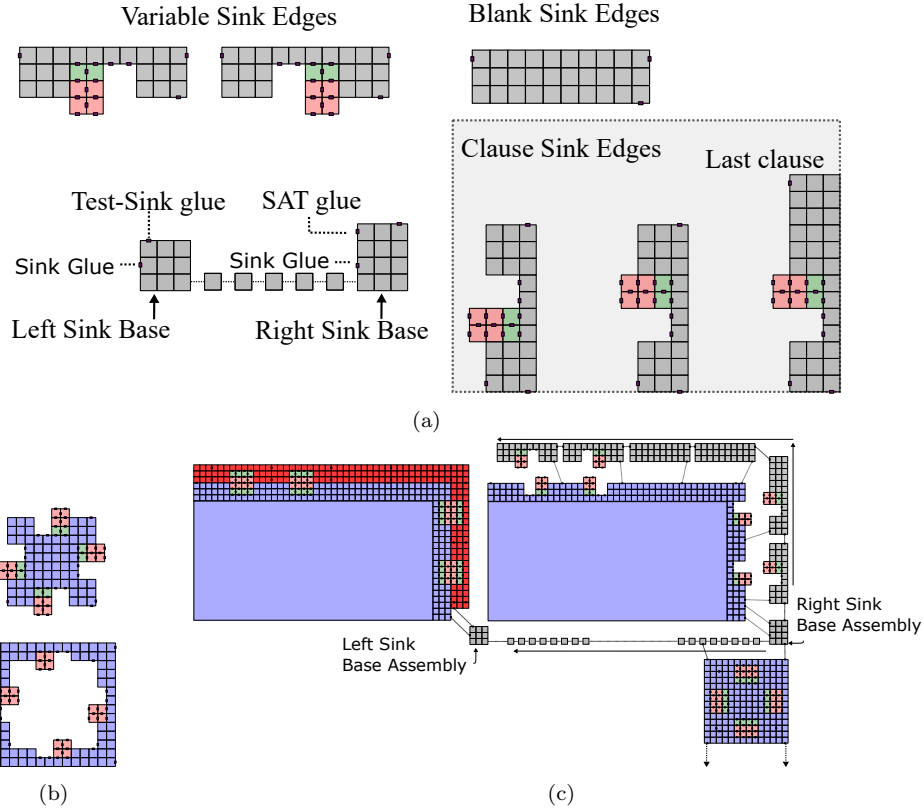


Figure 9: (a) The set of sink assemblies. The sink glue and test-sink glue are utilized for cooperative attachment. (b) For each macroblock we also include a frame so none of the macroblocks are terminal. These frames attach to each other in order at the bottom of the sink assembly. (c) The process of sinking all assemblies towards the target assembly.

The construction of SAT assemblies begins with combinations of variable and edge assemblies to produce an L shaped frame assembly for each possible assignment to ϕ . The output of ϕ on each assignment is then computed by the attachment of macroblocks that encode the clauses of ϕ , and producing a SAT assembly. The SAT assemblies expose arms representing the assignment of the variables in X , as well as arms that represent whether each clause has been satisfied. From the included prebuilt assemblies, a test assembly is produced for each possible assignment to X . Due to their arm positions, they may only attach to SAT assemblies that have a matching assignment to X and that represent an assignment that satisfies every clause.

Each assembly besides unmatched test assemblies will sink to the target

assembly. Every prebuilt assembly that was designed for building a SAT assembly will be used in the construction of at least one SAT assembly, besides the macroblocks. In some cases it is possible for a macroblock to not be able to attach to any SAT assembly. To account for this we include a macroblock frame for each macroblock to attach to, ensuring that no macroblock is terminal. The right sink base assembly may attach to any SAT assembly regardless of arm position so none of the SAT assemblies are terminal. Each sink assembly is used in the process of reaching the target assembly so none are terminal.

A test assembly is only terminal if there does not exist a SAT assembly with matching X arm positions that has each clause satisfied. This means for that assignment to X there does not exist an assignment to Y that satisfied ϕ , and the instance of $\forall\exists\text{SAT}$ is false.

The new assemblies in this reduction only increase the number of assemblies by a constant factor. The test and sink assemblies only add $\mathcal{O}(|V|+|C|)$ to the input assembly size, and the added macroblock frames are equal to the number of macroblocks, which is constant. \square

5. 1D Verification

In this section, we show that the producibility and the UAV problems in one-dimensional 2HAM (all assemblies are of height-1, and tiles only have glues on their left and right sides) with prebuilt assemblies is solvable in polynomial time.

5.1. Producibility Verification

We first present the algorithm for verifying whether a given assembly is producible in a given system. At a high level, the algorithm constructs a graph of initial assemblies that may be combined to form the target assembly A . An example graph is shown in Figure 10. We add a starting node s that connects to possible leftmost assemblies and a target node t that connects to possible rightmost assemblies. There exists an edge between two assemblies if they may attach as part of an assembly sequence building A . Thus, a path from s to t implies an assembly sequence using those nodes to build the target assembly.

Theorem 5.1. The producibility problem in the one-dimensional 2HAM with prebuilt assemblies is solvable in polynomial time.




Figure 10: Example instance of Producibility with the graph G created. The path between s and t implies a build sequence to produce the target assembly.

Proof. Given a 2HAM system $\Gamma = (S, \tau)$, note that τ is part of the input system, but does not affect the system since there is no cooperative binding. Thus, a one-dimensional system works for any positive temperature. Now given Γ , create a graph G where each node represents an initial assembly $i \in S$. For each pair of initial assemblies, create an edge going from assembly a to b if the leftmost tile of b shares a glue with the rightmost tile of a , meaning that b may attach to a .

For two assemblies A and B , we say B is consistent with A by an offset of i , if $B \subset A$ and for each $0 < j \leq |B|$, the j^{th} tile of B is the same as the $(i + j)^{\text{th}}$ tile of A . Informally, B is consistent with A if B has the same tiles as A starting at the i^{th} tile. We say the consistency ends at position $i + |B|$.

The producibility problem can be solved in polynomial time by performing the following algorithm. Create a graph G initialized with two nodes called s and t . For $0 \leq n < |A|$, add a node for each initial assembly $B \in S$ where B is consistent with A offset by n . Add an edge from assemblies a to b if b attaches on the right side of a to form assembly c and c is consistent with A by the same offset of a . That is, there exists an edge between two assemblies if they may combine to form an assembly that is still consistent with A . Finally, for each assembly that ends at the last tile of A , add an edge from that node to the t node.

If there exists a path between s and t , the assemblies represented by the path can be used as an assembly sequence. The first node will be the left most tiles of A . Each consecutive attachment will remain consistent with A until t is reached and all the tiles of A have been placed. Creating G takes $\mathcal{O}(|A||S|)$ steps so this algorithm runs in $\mathcal{O}(n^2)$ time. \square

5.2. Unique Assembly Verification

We now present an algorithm for the unique assembly verification problem in the same model. Given an instance of UAV (Γ, A) , we first present

three conditions that, if checked, provide the answer. If and only if all these conditions are true, the answer to UAV is ‘yes’: 1) A is producible. 2) There does not exist a subassembly $B \sqsubseteq A$ that is terminal. 3) There does not exist a producible assembly $R \not\sqsubseteq A$. Here, R is a rogue assembly since it will never grow into A . As a given system can have an exponential number of assemblies of size $\leq |A|$, we present a lemma showing a limit on the search for a witness that condition 3 is false.

Lemma 5.2. Let (Γ, A) be an instance of UAV in the one-dimensional 2HAM with prebuilt assemblies such that every initial assembly is a subassembly of A , and A is producible. If there exists a rogue assembly $R \not\sqsubseteq A$, then there exists a rogue assembly R' that can be produced by combining two assemblies R'_1, R'_2 that are subassemblies of A .

Proof. Consider a rogue assembly $R \not\sqsubseteq A$. Let I be the set of initial assemblies of Γ . Let a decomposition tree T_R of R be a binary tree that describes a valid assembly process of R starting from initial assemblies in I . A node of T_R is an assembly, e.g., the root node is R . The child nodes n_1, n_2 of a node n represent that assemblies n_1 and n_2 can come together in one production step to form the assembly n . Every leaf node of T is an initial assembly from I . With the assumed condition that every initial assembly is a subassembly of the target A , every leaf of T is a subassembly of A .

Mark every node of the tree that is not a subassembly of A with label ‘ r ’ (rogue assembly), and every other node with label ‘ s ’ (subassembly). The root node R has label ‘ r ’ while every leaf has label ‘ s ’.

We are now searching for a witness node R' with label ‘ r ’, which must have two children with label ‘ s ’ that prove the lemma true. Set the current candidate node B to R , and let B_1 and B_2 be the children of B . If B_1 and B_2 have label ‘ s ’, then the witness node has been found. Otherwise, w.l.o.g., let B_1 have label ‘ r ’. Set the candidate node B to the node R_1 and repeat the process. Since every leaf has label ‘ s ’, eventually the witness node R' is found. \square

Utilizing this lemma, we search only the subassemblies of our target assembly A to verify the third condition. 1D assemblies only have a polynomial number of subassemblies, so this allows us to check both the second and third condition. Combining these two steps with the polynomial time producibility algorithm from above, we solve UAV in polynomial time.

Result: Given an instance of UAV (Γ, A) in the 1D 2HAM with prebuilt assemblies, does Γ uniquely assemble A ?

if A is not producible in Γ **then** reject;

if there is an initial assembly $a \not\sqsubseteq A$ **then** reject;

for every subassembly a of A **do**

 | **if** a is producible and a is terminal **then** reject;

for every pair of subassemblies a, b of A **do**

 | **if** a is producible and b is producible **then**

 | $c_1 \leftarrow$ the assembly created by attaching a to the left of b ;

 | $c_2 \leftarrow$ the assembly created by attaching a to the right of b ;

 | **if** (c_1 is stable and $c_1 \not\sqsubseteq A$) or (c_2 is stable and $c_2 \not\sqsubseteq A$) **then**

 | reject;

 | accept;

Algorithm 1: A $\mathcal{O}(n^6)$ time algorithm for UAV in the one-dimensional 2HAM with prebuilt assemblies.

Theorem 5.3. The Unique Assembly Verification problem in the one-dimensional 2HAM with prebuilt assemblies is solvable in polynomial time.

Proof. Algorithm 1 solves the UAV problem in $\mathcal{O}(n^6)$ time. The algorithm first checks if the target assembly A is producible in $\mathcal{O}(n^2)$ time (from Theorem 5.1). The algorithm then checks for every producible subassembly of the target, and whether that subassembly is terminal. This requires $\mathcal{O}(n^2) \times \mathcal{O}(n^2) \times \mathcal{O}(n) = \mathcal{O}(n^5)$ time, as there are at most $\mathcal{O}(|A|^2)$ subassemblies of A . For each subassembly, terminality is checked by comparing the endpoints of the subassembly to endpoints of all other initial assemblies. These endpoints can be stored for this process to require $\mathcal{O}(n)$ time.

The last step of the algorithm searches for a witness that condition 3 is false, and therefore a witness to a ‘no’ answer. This witness would be a rogue assembly R that is producible, but not a subassembly of the target A . Since R can never grow into A , the answer to UAV must be no. By Lemma 5.2, if a rogue assembly exists, then there exists some rogue assembly R' that can be built by combining two subassemblies of A . The algorithm checks every pair of subassemblies of A (of which there are $|A|^4$), and checks if they can be combined to form a rogue assembly. \square

6. aTAM with Prebuilt Assemblies

Here, we show that the hardness reduction holds in the case where growth happens in a seeded fashion, both in the 2HAM and in the aTAM. In the aTAM there is a set of tiles that may attach one at a time to the seed assembly to create the set of producible assemblies. Here, we start with a seed assembly and a set of initial assemblies that may attach to it. We may only attach one assembly at a time to the seed or to any other producible assembly.

Abstract Tile Assembly Model with Prebuilt Assemblies. An aTAM system is an ordered triple (S, σ, τ) where S is a set of *initial* assemblies, σ is the seed assembly, and τ is a positive integer parameter called the *temperature*. Each assembly in S must be τ -stable. For a system (S, σ, τ) , the set of *producible* assemblies $P'_{(S, \sigma, \tau)}$ is defined recursively as follows:

1. σ .
2. If $A \in P'_{(S, \sigma, \tau)}$ and $B \in S$ are τ -combinable into C , then $C \in P'_{(S, \sigma, \tau)}$.

The definitions and notations carry over from the standard definition for terminal assemblies and unique production.

Corollary 6.1. The producibility problem in the aTAM, with prebuilt assemblies of size $\mathcal{O}(1)$, is NP-complete with $\tau = 2$ and an initial assembly set size of $\mathcal{O}(|V||C|)$.

Proof. In the aTAM, growth begins from a designated seed assembly. All producible assemblies are formed by starting with this seed assembly and attaching single assemblies from the input set one-by-one.

The same set of edge assemblies and macroblocks from Section 3 can be used to show hardness of the producibility problem in this version of the model. The left corner assembly can be used as the seed. The edge assemblies are able to attach to this assembly to build the frame. Macroblocks can still attach one-by-one, which computes the clauses. Each assembly used in the reduction attaches to the growing frame assembly. No attachments that are needed occur separate from this assembly, so the target assembly is still producible if and only if there exists a satisfying assignment to ϕ . \square

Corollary 6.2. The UAV problem in the aTAM, with prebuilt assemblies of size $\mathcal{O}(1)$, is coNP-complete with $\tau = 2$ and an initial assembly set size of $\mathcal{O}(|V||C|)$.




Figure 11: Smaller assemblies that can be used for the reduction with modifications to the other assemblies.

Proof. We can show UAV with constant-sized assemblies by reducing from \forall SAT, which takes as input a formula ϕ and outputs yes if all assignments evaluate to true and no if any assignment is false. The target assembly will be the rectangle assembly that is the target in Corollary 6.1. The bottom row of attachments represents an assignment to the formula, and any satisfying assignment grows to the target assembly.

If all assignments evaluate to true, then all build sequences will result in the target assembly and thus it is uniquely produced. If the target assembly is uniquely produced then all build sequences must represent a satisfying assignment. \square

7. Future Work

The macroblocks used above have constant size, but in this section we provide smaller alternate assemblies that have size 38. This is a substantial improvement over the size 92 blocks of the original design, and points to interesting future work in determining at what size of prebuilt assembly the complexity of verification fundamentally changes.

7.1. Minimizing Assembly Size

In Figure 11, we show equivalent assemblies that can be used in the reduction described above. The arms are now 2×2 subassemblies and only have tiles required to connect them.

Theorem 7.1. The Producibility problem in the 2HAM and aTAM with prebuilt assemblies of max size 38 is NP-complete with an initial assembly set size of $\mathcal{O}(|V||C|)$ and $\tau = 2$.

Proof. The new assemblies in Figure 11 have the same glues as the assemblies from the previous proof. The arms also still enforce that only macroblocks with matching assignments may attach. \square

Theorem 7.2. The UAV problem in the 2HAM with prebuilt assemblies of max size 38 is coNP^{NP} -complete with an initial assembly set size of $\mathcal{O}(|V||C|)$ and $\tau = 2$.

Proof. Even the test and sink assemblies can be minimized to be made of fewer than 38 tiles. Since these assemblies have an identical structure to the assemblies used in the reduction for production, we can use the same methods to shrink the assemblies. \square

The producibility problem in the 2HAM with only single tiles is solvable in polynomial time. This leaves a gap for max assembly sizes $2 \leq n \leq 35$. This points to several open questions related to producibility and UAV with prebuilt assemblies. Is the producibility problem solvable in polynomial time with only dominoes, or what is the smallest size of prebuilt assemblies necessary for it to be NP-complete?

With the recent work in [15], UAV in the 2HAM with constant temperature and singleton tiles is coNP -complete, and we know that with prebuilt assemblies of size 38, UAV is coNP^{NP} -complete. What is the smallest size of prebuilt assemblies to maintain coNP^{NP} -completeness? It also is of interest to find the minimum number of glues needed for hardness of these problems. Can we show hardness for these problems with a small number of glues? What is the trade off between number of glues and size of the assembly?

8. Conclusion

In this paper we explored the fundamental self-assembly problems of Producibility and Unique Assembly Verification (UAV) in the 2HAM with prebuilt assemblies. These problems have been well studied in the classic scenario in which starting assemblies are singleton tiles, with the producibility problem being in P , and the UAV problem being coNP -complete. By extending these models to include prebuilt assemblies, even of constant size, we have shown that each problem rises one level in the polynomial hierarchy, with producibility becoming NP-complete, and UAV becoming coNP^{NP} -complete. We further showed that in the case of 1-dimensional linear assemblies, both of these problems have polynomial time solutions. We also extended these results to the aTAM with constant-size attachable assemblies,

showing NP-completeness for producibility and coNP-completeness for UAV. Understanding the complexity of these problems is motivated as these are two of the key problems related to verifying the behavior of self-assembling systems, and the inclusion of prebuilt assemblies is easily motivated by practical laboratory procedures. That these hardness results hold even for constant sized assemblies further leads to an intriguing question: For what size assemblies do these verification questions switch complexity classes? Our initial construction utilizes size 92 prebuilt assemblies, but we show this can be decreased down to 38. Can it be pushed further? What about the case of size-2 assemblies?

References

- [1] S. Cannon, E. D. Demaine, M. L. Demaine, S. Eisenstat, M. J. Patitz, R. T. Schweller, S. M. Summers, A. Winslow, Two Hands Are Better Than One (up to constant factors): Self-Assembly In The 2HAM vs. aTAM, in: 30th International Symposium on Theoretical Aspects of Computer Science (STACS 2013), Vol. 20 of Leibniz International Proceedings in Informatics (LIPIcs), Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2013, pp. 172–184.
- [2] E. Winfree, Algorithmic self-assembly of DNA, Ph.D. thesis, California Institute of Technology (June 1998).
- [3] P.-E. Meunier, D. Regnault, D. Woods, The program-size complexity of self-assembled paths, in: STOC: Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Association for Computing Machinery, 2020, pp. 727–737. doi:10.1145/3357713.3384263.
URL <https://doi.org/10.1145/3357713.3384263>
- [4] M. J. Patitz, An introduction to tile-based self-assembly and a survey of recent results, *Natural Computing* 13 (2) (2014) 195–224. doi:10.1007/s11047-013-9379-4.
- [5] D. Woods, Intrinsic universality and the computational power of self-assembly, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 373 (2046) (2013) 16–22.

- [6] D. Doty, Theory of algorithmic self-assembly, *Communications of the ACM* 55 (12) (2012) 78–88.
- [7] C. Evans, Crystals that count! physical principles and experimental investigations of DNA tile self-assembly, Ph.D. thesis, California Inst. of Tech. (2014).
- [8] D. Woods, D. Doty, C. Myhrvold, J. Hui, F. Zhou, P. Yin, E. Winfree, Diverse and robust molecular algorithms using reprogrammable DNA self-assembly, *Nature* 567 (2019) 366–372. doi:10.1038/s41586-019-1014-9.
- [9] E. D. Demaine, M. L. Demaine, S. P. Fekete, M. Ishaque, E. Rafalin, R. T. Schweller, D. L. Souvaine, Staged self-assembly: nanomanufacture of arbitrary shapes with $o(1)$ glues, *Natural Computing* 7 (3) (2008) 347–370.
- [10] S. P. Fekete, J. Hendricks, M. J. Patitz, T. A. Rogers, R. T. Schweller, Universal computation with arbitrary polyomino tiles in non-cooperative self-assembly, in: *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, SIAM, 2014, pp. 148–167.
- [11] B. Fu, M. J. Patitz, R. T. Schweller, R. Sheline, Self-assembly with geometric tiles, in: *International Colloquium on Automata, Languages, and Programming*, Springer, 2012, pp. 714–725.
- [12] M. Endo, T. Sugita, Y. Katsuda, K. Hidaka, H. Sugiyama, Programmed-assembly system using DNA jigsaw pieces, *Chemistry: A European Journal* (2010) 5362–5368.
- [13] S. Woo, P. W. Rothemund, Stacking bonds: Programming molecular recognition based on the geometry of DNA nanostructures, *Nature Chemistry* 3 (2011) 620–627.
- [14] D. Doty, Producibility in hierarchical self-assembly, in: O. H. Ibarra, L. Kari, S. Kopecki (Eds.), *Unconventional Computation and Natural Computation*, Springer International Publishing, Cham, 2014, pp. 142–154.
- [15] D. Caballero, T. Gomez, R. Schweller, T. Wylie, Unique assembly verification in two-handed self-assembly, in: M. Bojańczyk, E. Merelli, D. P.

- Woodruff (Eds.), 49th International Colloquium on Automata, Languages, and Programming (ICALP 2022), Vol. 229 of ICALP'22, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2022, pp. 34:1–34:21. doi:10.4230/LIPIcs.ICALP.2022.34.
- [16] R. Schweller, A. Winslow, T. Wylie, Complexities for high-temperature two-handed tile self-assembly, in: R. Brijder, L. Qian (Eds.), DNA Computing and Molecular Programming, Springer International Publishing, Cham, 2017, pp. 98–109.
- [17] L. M. Adleman, Q. Cheng, A. Goel, M.-D. A. Huang, D. Kempe, P. M. de Espanés, P. W. K. Rothmund, Combinatorial optimization problems in self-assembly, in: Proceedings of the 34th Annual ACM Symposium on Theory of Computing, 2002, pp. 23–32.
- [18] D. Doty, L. Kari, B. Masson, Negative interactions in irreversible self-assembly, *Algorithmica* 66 (1) (2013) 153–172. doi:10.1007/s00453-012-9631-9.
- [19] A. A. Cantu, A. Luchsinger, R. Schweller, T. Wylie, Covert computation in self-assembled circuits, *Algorithmica* 83 (2) (2021) 531–552.
- [20] R. Schweller, A. Winslow, T. Wylie, Verification in staged tile self-assembly, *Natural Computing* 18 (1) (2019) 107–117.
- [21] D. Caballero, T. Gomez, R. Schweller, T. Wylie, Covert Computation in Staged Self-Assembly: Verification Is PSPACE-Complete, in: P. Mutzel, R. Pagh, G. Herman (Eds.), 29th Annual European Symposium on Algorithms (ESA 2021), Vol. 204 of Leibniz International Proceedings in Informatics (LIPIcs), Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2021, pp. 23:1–23:18. doi:10.4230/LIPIcs.ESA.2021.23.
URL <https://drops.dagstuhl.de/opus/volltexte/2021/14604>
- [22] D. Caballero, T. Gomez, R. Schweller, T. Wylie, The complexity of multiple handed self-assembly, in: I. Kostitsyna, P. Orponen (Eds.), Unconventional Computation and Natural Computation, Springer International Publishing, Cham, 2021, pp. 1–18.

- [23] C. Chalk, A. Luchsinger, E. Martinez, R. Schweller, A. Winslow, T. Wylie, Freezing simulates non-freezing tile automata, in: *DNA Computing and Molecular Programming*, 2018, pp. 155–172.
- [24] D. Caballero, T. Gomez, R. Schweller, T. Wylie, Verification and Computation in Restricted Tile Automata, in: C. Geary, M. J. Patitz (Eds.), *26th International Conference on DNA Computing and Molecular Programming (DNA 26)*, Vol. 174 of *Leibniz International Proceedings in Informatics (LIPIcs)*, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2020, pp. 10:1–10:18. doi:10.4230/LIPIcs.DNA.2020.10.
- [25] S. Arora, B. Barak, *Computational complexity: a modern approach*, Cambridge University Press, 2009.
- [26] D. Caballero, T. Gomez, R. Schweller, T. Wylie, Complexity of verification in self-assembly with prebuilt assemblies, in: *Proceedings of the Symposium on Algorithmic Foundations of Dynamic Networks*, Vol. 221 of *SAND’22*, 2022, pp. 8:1–8:15.

Graphical Abstract

Complexity of Verification in Self-Assembly with Prebuilt Assemblies*

David Caballero, Timothy Gomez, Robert Schweller, Tim Wylie

Highlights

Complexity of Verification in Self-Assembly with Prebuilt Assemblies

David Caballero, Timothy Gomez, Robert Schweller, Tim Wylie

- 2HAM producibility with prebuilt assemblies coNP-complete
- 2HAM unique assembly verification with prebuilt assemblies is coNP^{NP}-complete.
- 2HAM 1D producibility and UAV are polynomial.
- aTAM producibility with prebuilt assemblies is NP-complete.