

The Complexity of Fixed-Height Patterned Tile Self-Assembly

Shinnosuke Seki*

*Department of Computer and Network Engineering, The University of Electro-Communications
1-5-1, Chofugaoka, Chofu, Tokyo, 182-8585, Japan
s.seki@uec.ac.jp*

Andrew Winslow†

*Department of Computer Science, University of Texas Rio Grande Valley
1201 W University Drive, Edinburg, TX 78541, USA
andrew.winslow@utrgv.edu*

Received (Day Month Year)
Accepted (Day Month Year)
Communicated by (xxxxxxxxxx)

We characterize the complexity of the PATS problem for patterns of fixed height and color count in variants of the model where seed glues are either chosen or fixed and identical (so-called *non-uniform* and *uniform* variants). We prove that both variants are NP-complete for patterns of height 2 or more and admit $O(n)$ -time algorithms for patterns of height 1. We also prove that if the height and number of colors in the pattern is fixed, the non-uniform variant admits a $O(n)$ -time algorithm while the uniform variant remains NP-complete. The NP-completeness results use a new reduction from a constrained version of a problem on finite state transducers.

1. Introduction

Winfree [14] introduced the *abstract tile assembly model (aTAM)* to capture nanoscale systems of DNA-based particles aggregating to form intricate crystals, leading to an entire field devoted to understanding the theoretical limits of such systems (see surveys by Doty [3] and Patitz [11]). Ma and Lombardi [10] introduced the *patterned self-assembly tile set synthesis (PATS)* problem, of designing a tile set of minimum size that assembles into a given $n \times h$ colored pattern by attaching to an L-shaped seed.

Czeizler and Popa [2] were the first to provide a proof that the PATS problem is NP-hard, thus establishing the problem as NP-complete. Subsequent work

*Work supported in part by JST Program to Disseminate Tenure Tracking System, MEXT, Japan, No. 6F36 and by JSPS Grant-in-Aids for Research Activity Start-up No. 15H06212 and for Young Scientists (A) No. 16H05854.

†A substantial portion of the work of this author was performed while at Université Libre de Bruxelles.

studied the hardness of the constrained version where the patterns have at most c colors, called the c -PATS problem. This line of work proved the 60-PATS [12], 29-PATS [7], 11-PATS [8], and finally the 2-PATS [9] problems NP-complete.

Here we study the complexity of parameterized *height- h* PATS and c -PATS problems where patterns have a specified fixed height h and increasing width n . We consider both *uniform* and *non-uniform* model variants, where the glues along the seed are fixed and identical or chosen in tandem with the tile set, respectively. We characterize the computational complexity of these problems via the following results:

- The height-2 PATS problem is NP-complete in both models (Sec. 4).
- The uniform height-2 3-PATS problem is NP-complete (Sec. 5).
- The non-uniform height- h c -PATS problem and uniform height-1 PATS problems admit $c^{O(h)}$ n -time and $O(n)$ -time algorithms, respectively (Sec. 6).

The NP-completeness results also apply to patterns of height greater than 2. Thus the complexity of the PATS problem for all combinations of height, color, and uniformity are characterized, except uniform height-2 2-PATS.

The NP-hardness reductions are based on a reduction for a new variant of the minimum-state finite state transducer problem, originally proved NP-hard by Angluin [1] and by Vazirani and Vazirani [13]. In this variant, any solution transducer is also promised to satisfy additional constraints on its transitions. The reduction is also substantially simpler than the reduction given in [1] and uses input and output strings of just two symbols, rather than the three of [13].

2. Preliminaries

Patterns, tiles, assemblies, and seeds. Define $\mathbb{N}_k = \{1, 2, \dots, k\}$. A *pattern* is a partial function $P : \mathbb{N}^2 \rightarrow C$, i.e. a function that maps a rectangular region of lattice points to a set of colors C (see Figure 1). If $\text{dom}(P) = \mathbb{N}_w \times \mathbb{N}_h$, then P is a *width- w height- h* pattern. The codomain of P , i.e. the colors seen in the pattern, is denoted $\text{color}(P)$. A pattern P is *c -color* provided $|\text{color}(P)| \leq c$.

A *tile type* t is a colored unit square with each edge labeled; these labels are called *glues*. A tile type's color is denoted $\text{color}(t)$. For a direction $d \in \{N, W, S, E\}$, $t[d]$ denotes the glue assigned to side d of t . A tile type is non-rotatable, and thus is uniquely identified by its color and four glues. Instances of tile types, called *tiles*, are placed with their centers in \mathbb{N}^2 .

An *assembly* is an arrangement of tiles from a set of tile types T ; formally a partial function $A : \mathbb{N}^2 \rightarrow T \cup \{\emptyset\}$. A *seed* is an ‘‘L-shaped’’ assembly with domain $\{(0, 0)\} \cup \{(x, 0) : x \in \mathbb{N}_w\} \cup \{(0, y) : y \in \mathbb{N}_h\}$ for some $w, h \in \mathbb{N}$. The *pattern* of an assembly A is defined as $P_A((x, y)) = \text{color}(A((x, y)))$ for $(x, y) \in \text{dom}(A) \cap \mathbb{N}^2$, i.e. the color pattern of A , *excluding* the seed.

RTASs. A *rectilinear tile assembly system (RTAS)* is a pair $\mathcal{T} = (T, \sigma)$, where T is a set of tile types and σ is a *seed*. An assembly A *yields* an assembly A' with $\text{dom}(A') = \text{dom}(A) \cup \{(x, y)\}$ provided $(x - 1, y), (x, y - 1) \in \text{dom}(A)$ and $A((x - 1, y))[\text{E}] = A'((x, y))[\text{W}]$, $A((x, y - 1))[\text{N}] = A'((x, y))[\text{S}]$. The *producible assemblies* of an RTAS are those that can be yielded, starting with the seed assembly σ . That is:

RTAS Tiling Rule: A tile of type t can be added to an assembly A at location (x, y) provided $(x - 1, y), (x, y - 1) \in \text{dom}(A)$ and the east and north glues of the tiles at $(x - 1, y)$ and $(x, y - 1)$ are the same as the west and south glues of t , respectively.

As a result, tiling proceeds from southwest to northeast, i.e., a tile is first placed at $(1, 1)$, then at either $(1, 2)$ or $(2, 1)$, etc. The *terminal assemblies* of an RTAS are the producible assemblies that do not yield other (larger) assemblies. If every terminal assembly of the system has pattern P , the system is said to *uniquely self-assemble* P . An RTAS (T, σ) is *directed*, i.e. deterministic, provided for any distinct tile types $t_1, t_2 \in T$, either $t_1[\text{W}] \neq t_2[\text{W}]$ or $t_1[\text{S}] \neq t_2[\text{S}]$.

Uniform RTASs. We also define a practical variant of an RTAS called a *uniform RTAS*. An RTAS (T, σ) is *uniform* provided there exist two glues $\ell_{\text{E}}, \ell_{\text{N}}$ such that $\sigma((x, 0))[\text{E}] = \ell_{\text{E}}$ for all $x \in \mathbb{N}_w$ and $\sigma((0, y))[\text{N}] = \ell_{\text{N}}$ for all $y \in \mathbb{N}_h$. In other words, the seed glues cannot be programmed and are generic.

The PATS problem. The *pattern self-assembly tile set synthesis problem (PATS)* [10] asks for the minimum-size RTAS that uniquely self-assembles a given rectangular color pattern, where the size of an RTAS (T, σ) is $|T|$, the number of tile types. Bounding the number of colors or height of the input pattern yields the following practically motivated special cases of PATS:




Fig. 1. A height-2 2-color pattern and minimum-size RTASs uniquely assembling the pattern in the uniform and non-uniform models.

Problem 1 (*c*-colored PATS or *c*-PATS) Given a *c*-colored pattern P and integer t , does there exist an RTAS of size $\leq t$ that uniquely self-assembles P ?

Problem 2 (Height-*h* PATS) Given a height-*h* pattern P and integer t , does there exist an RTAS of size $\leq t$ that uniquely self-assembles P ?

Restricting the system to be uniform gives rise to *uniform* variants as well, contrasting with the conventional *non-uniform* variants.

3. Minimum-State Finite State Transducer is NP-hard

We start with a reduction from a well-known NP-complete problem on integers, called 3-PARTITION, to a problem on *finite state transducers* or *FSTs*: deterministic finite automata where each transition is augmented with an output symbol and thus *transduces* an input string into an output string of the same length.

Problem 3 (3-PARTITION) *Given a set of integers $A = \{a_1, a_2, \dots, a_{3n}\}$ with $\sum_{a_i \in A} a_i/n = p$ and $p/4 < a_i < p/2$, does there exist a partition of A into n sets, each with sum p ?*

Theorem 4 ([5]) *3-PARTITION is strongly NP-hard, i.e., is NP-hard when the elements of A are given in unary.*

Formally, a FST is a 4-tuple $T = \langle \Sigma, Q, s_0, \delta \rangle$, where Σ is the *alphabet*, Q is a finite set of *states* of T , $s_0 \in Q$ is the *start state* of T , and $\delta : Q \times \Sigma \rightarrow Q \times \Sigma$ is the *transition function* of T . An input-output quadruple $\delta(s_i, b) = (s_j, b')$ is a *transition*, specifically a (b, b') -*transition* or a *b-transition*. The size of T is equal to $|Q|$.

Problem 5 (ENCODING BY FST) *Given two strings S, S' and integer K , does there exist a FST with at most K states that transduces S to S' ?*

The ENCODING BY FST problem was previously shown to be NP-hard by Angluin [1] and by Vazirani and Vazirani [13]. Here we prove a constrained variant is also NP-hard:

Problem 6 (PROMISE ENCODING BY FST) *Given two strings S, S' and an integer K with the following promises about any FST T with at most K states transducing S to S' , does such a T exist?*

- *Each state of T has exactly one incoming 0-transition.*
- *Each state of T has exactly one incoming 1-transition.*
- *When transducing S to S' :*
 - *$K - 1$ distinct $(0, 0)$ -transitions are used.*
 - *K distinct $(1, 1)$ -transitions are used.*
 - *1 distinct $(0, 1)$ -transition is used.*
 - *The order that transitions are traversed is given as part of the input.^a*

Lemma 7. *The PROMISE ENCODING BY FST problem is NP-hard.*

^aThe order is given by a sequence $t_1, t_2, \dots, t_{|S|}$ with each $t_i \in \{1, 2, \dots, 2K\}$ specifying which transition is traversed. Thus the order that transitions are revisited is specified, but the states visited are not.

Proof. The reduction is initially presented as a reduction from 3-PARTITION to ENCODING BY FST. The final paragraph extends the reduction to PROMISE ENCODING BY FST.

We borrow from [13] the approach of constructing S and S' by concatenating *segments*: pairs of input and output substrings of equal length that enforce specific structure in a solution FST. An input string X and output string Y paired as a segment is denoted $X \rightarrow Y$.

The integer output by the reduction is $K = 3pn + n + 1$, where n is the number of parts in the partition and p the size of each part. The first segment is $0^{K-1}00^{K-1} \rightarrow 0^{K-1}10^{K-1}$. This segment enforces that a solution FST must have K states; label them s_1, s_2, \dots, s_K . Then for all $i < n$, $\delta(s_i, 0) = (s_{i+1}, 0)$ and $\delta(s_K, 0) = (s_1, 1)$.




Fig. 2. A solution FST for a toy reduction from 3-PARTITION to PROMISE ENCODING BY FST with integers $a_1 = 1, a_2 = 1, a_3 = 2, a_4 = 1, a_5 = 1, a_6 = 2$ (invalid due to duplicate elements and violating $p/4 < a_i < p/2$, but used for illustrative purposes). The left-to-right states are s_1 to s_K , colored by their half-fixed interval. Transitions above the states are $(1, 1)$ -transitions. All others are $(0, 0)$ -transitions except the lowermost, a $(0, 1)$ -transition. The solution corresponds to the partition $\{\{a_2, a_4, a_6\}, \{a_1, a_3, a_5\}\}$, encoded in the assignment of the right halves of half-fixed intervals to states between the rightmost three fixed singletons (gray).

The problem of partitioning integers of A into sets of size p is implemented in the collection of 1-transitions that leave states. Each state has a 1-transition that either points to itself (a *fixed singleton*) or is one edge in a 3-cycle formed by two consecutive specified states and an unspecified third state (a *half-fixed triple*). Half-fixed triples are further organized into *half-fixed intervals*, each consisting of a set of $2a_i$ consecutive specified states and a set of a_i consecutive unspecified states for some a_i . The states are partitioned into three subsets:

- (1) States s_1 through s_{2pn} consist of the specified halves of the half-fixed intervals ($2a_i$ consecutive states for each integer $a_i \in A$).
- (2) $n + 1$ equally-spaced fixed singletons in states s_{2pn+1}, \dots, s_K (partitioning the remaining pn states into n sets of p consecutive states each).
- (3) The remaining pn states in states s_{2pn+1}, \dots, s_K partitioned into n sets of

p consecutive states.

See Figure 2 for a toy example of the reduction.

The reduction works by leaving one aspect of the FST unspecified: the assignment of the unspecified halves of half-fixed intervals (sequences of a_i consecutive states for each $a_i \in A$) to the n sequences of p consecutive states each (the third subset mentioned above).

These unspecified halves must be assigned to these sequences of p consecutive states, since all other states are either fixed halves of half-fixed intervals (first subset above) or fixed singletons (second subset above). Also, since no state can lie on multiple 3-cycles of 1-transitions, no state can be in multiple half-fixed triples. So no state can be in multiple half-fixed intervals and thus the unspecified halves of half-fixed intervals cannot overlap.

Thus the unspecified halves of half-fixed intervals may be viewed as indivisible “rods” of consecutive states (of lengths a_i for each $a_i \in A$) that must be placed into n state “boxes”, each of length p , without overlap. Since the sum of the rod lengths is exactly pn , this problem is equivalent to partitioning the rods into n subsets, each of sum p , i.e. 3-PARTITION.

In Figure 2, the input (invalid) 3-PARTITION instance $A = \{a_1 = 1, a_2 = 1, a_3 = 2, a_4 = 1, a_5 = 1, a_6 = 2\}$ has a solution $\{\{a_2, a_4, a_6\}, \{a_1, a_3, a_5\}\}$. This corresponds to a solution FST obtained by assigning the unspecified halves of half-fixed intervals corresponding to $a_2, a_4,$ and a_6 (of lengths 1, 1, and 2, respectively) into the first sequence of $p = 4$ consecutive states, and similarly the unspecified halves of the half-fixed intervals corresponding to $a_1, a_3,$ and a_5 (of lengths 1, 2, and 1, resp.) into the second sequence of $p = 4$ consecutive states. All that remains is to describe the segments that force the construction of a fixed singleton, half-fixed triple, and half-fixed interval.

Fixed singleton. The fixed singleton segment ensures that a given state s_i has $\delta(s_i, 1) = (s_i, 1)$. This is done by moving the current state to s_i , transducing a 1 to a 1, and checking whether the current state is still s_i (see Figure 3).




Fig. 3. The fixed single segment and corresponding FST structure enforced.

Half-fixed triple. The half-fixed triple segment forces two specified *fixed* states s_i, s_{i+1} and an unspecified *free* third state s_j to have $\delta(s_i, 1) = (s_{i+1}, 1)$, $\delta(s_{i+1}, 1) = (s_j, 1)$, and $\delta(s_j, 1) = (s_i, 1)$ (see Figure 4).




Fig. 4. The half-fixed triple segment and corresponding FST structure enforced.

The segment consists of two subsegments that each ensures a portion of the structure. The first, $0^{i-1}10^{K-i-1}0 \rightarrow 0^{i-1}10^{K-i-1}1$, ensures that $\delta(s_i, 1) = (s_{i+1}, 1)$. The second, $0^i110^{K-i}0 \rightarrow 0^i110^{K-i}1$, ensures that $\delta(s_{i+1}, 1) = (s_j, 1)$ and $\delta(s_j, 1) = (s_i, 1)$. The state s_j cannot be in a fixed state of another half-fixed triple segment with fixed states s'_i, s'_{i+1} and free state s'_j , as then either:

- $s_j = s'_i$ and thus $\delta(s_j, 1) = (s'_{i+1}, 1) \neq (s_i, 1)$ (and thus the segment $0^{i-1}10^{K-i-1}0 \rightarrow 0^{i-1}10^{K-i-1}1$ is not transduced).
- $s_j = s'_{i+1}$ and $(s'_j, 1) = \delta(s'_{i+1}, 1) = (s_i, 1)$, so $\delta(s'_j, 1) = (s_{i+1}, 1) \neq (s'_i, 1)$ (and thus the segment $0^{i'-1}10^{K-i'-1}0 \rightarrow 0^{i'-1}10^{K-i'-1}1$ is not transduced).




Fig. 5. The half-fixed interval segment for three consecutive free states and corresponding FST structure enforced.

Half-fixed interval. The half-fixed interval forces a collection of half-fixed triples with consecutive fixed states to also have consecutive free states. It does so by a simple traversal of the free states, checking that each has the expected pair of consecutive fixed states (see Figure 5).

Promises. Any solution transduction for the previous reduction uses an FST where each state has at most one incoming 1-transition and 0-transition, since every transition lies on a cycle (of length 1, 3, or K). Also, any solution transduction by an

FST with K states traverses $2K$ distinct transitions (with $K - 1$ $(0, 0)$ -transitions, 1 $(0, 1)$ -transition, and K $(1, 1)$ -transitions). Any other solution FST must have at least $2K$ states and traverse at least $2K + 1$ distinct transitions: $2K$ 0-transitions and at least one 1-transition. Thus the reduction proves that not only ENCODING BY FST, but also PROMISE ENCODING BY FST is NP-hard. \square

4. Height-2 PATS is NP-complete

Göös and Orponen [6] establish that all the variations of the PATS problem considered here are in NP. So we need only consider their NP-hardness.

Theorem 8. *The non-uniform height-2 PATS problem is NP-hard.*


Proof. The pattern output by the reduction consists of a bottom row encoding S and a top row encoding the sequence of transitions traversed when transducing S to S' (provided as part of the PROMISE ENCODING BY FST instance). The bottom row encoding uses two colors, pink and red, corresponding to the two symbols in S . The top row encoding uses $2K$ colors, one for each transition used in the transduction of S to S' . The number of tile types permitted is $T = 2K + 2$: one type per color.

The north glues of the bottom row either encode S (distinct north glues for the pink and red tile types) or $0^{|S|}$ (same glue). The latter is impossible, since then the leftmost $|S|$ locations of the top row are filled by many repetitions of the same K transitions. So the north glues of the bottom row encode S .

A set of $2K$ tile types that assemble the top row is equivalent to a set of $2K$ transitions transducing S to S' , with source and destination states corresponding to west and east glues. So the top row can be assembled using $2K$ tile types exactly when S can be transduced to S' using $2K$ transitions of the specified types traversed in the specified order. Note that these tile types may share the east glue in common, corresponding to that a state has more than one incoming 0-transition or more than one incoming 1-transition. Nevertheless, we can say that the pattern can be assembled using a tile set of at most $2K$ types exactly when there is a solution transducer for the corresponding instance of a harder variant of PROMISE ENCODING BY FST that is obtained by dropping the first two promises, on incoming transitions, from PROMISE ENCODING BY FST (the harder variant remains NP-hard). \square

Theorem 9. *The uniform height-2 PATS problem is NP-hard.*

Proof. The pattern output by the reduction is the following, and consists of a left *input half* and right *transduction half*:



The color patterns in the top rows of the input and transduction halves are identical to the bottom and top rows of the pattern used in the proof of Theorem 8,

respectively. That is, they consist of two colors encoding S , and $2K$ colors encoding the sequence of transitions.

The bottom row of the input half consists of $|S|$ orange tiles; these must be of $|S|$ distinct types, otherwise the entire bottom row consists of orange tiles. The number of tile types output by the reduction is $T = |S| + 2K + 4$, thus any solution set of tile types has exactly the following tile types:

- $|S|$ orange tile types.
- 2 types used in the top row of the input half.
- $2K$ types used in the top row of the transduction half.
- White and black types used for the two tiles between the halves.

Since only one tile type per color is used in the top row of the input half, the north glues of the length- $|S|$ orange tile sequence must encode S . Also, since the bottom row of the transduction half consists of a length- $|S|$ sequence of orange tiles and these tiles are all contained in the bottom row of the input half, they must be the same sequence of tiles. So the north glues of the bottom row of the transduction half encode S , and thus any set of $2K$ tile types used in the top row of the right half corresponds exactly to a set of transitions in a solution FST for the PROMISE ENCODING BY FST instance. \square

The addition of more rows with a new common color and increasing T by 1 suffices to prove both the uniform and non-uniform variants NP-hard for greater heights.

5. Uniform Height-2 3-PATS is NP-complete

The next problem is a variation of PROMISE ENCODING BY FST, which we prove to be NP-complete by modifying the reduction from 3-PARTITION to PROMISE ENCODING BY FST. This reduction serves as an introduction and warm-up for the subsequent reduction from 3-PARTITION to MODIFIED PROMISE ENCODING BY FST.

Problem 10 (MODIFIED PROMISE ENCODING BY FST) *Given two strings S , S' and an integer $K \not\equiv 0 \pmod{3}$ with the following promises about any FST T with at most K states transducing S to S' , does such a T exist?*

- Each state of T has exactly one incoming 0-transition.
- Each state of T has exactly one incoming 1-transition.
- Every $(1, 1)$ -transition lies on a 1-cycle or 3-cycle of $(1, 1)$ -transitions.
- When transducing S to S' :
 - $K - 1$ distinct $(0, 0)$ -transitions are used.
 - $K - 1$ distinct $(1, 1)$ -transitions are used.
 - 1 distinct $(0, 1)$ -transition is used.
 - 1 distinct $(1, 2)$ -transition is used.

- Every $(0, 1)$ -transition traversal is followed immediately by a $(1, 2)$ -transition traversal.
- The order that transitions are traversed is given as part of the input.
- The first and last $K + 2$ symbols of S are $10^K 1$.
- The first and last $K + 2$ symbols of S' are $20^{K-1} 12$.

Lemma 11. *The MODIFIED PROMISE ENCODING BY FST problem is NP-hard.*

Proof. The proof uses a modified version of the reduction from 3-PARTITION to PROMISE ENCODING BY FST (Lemma 7). A new state s_{new} is added to the solution transducers with the following transitions as seen in Figure 6:

- A $(0, 0)$ -transition from s_{new} to s_1 .
- $(1, 2)$ -transition from s_{new} to itself.
- A $(0, 1)$ -transition from s_K to s_{new} , replacing the $(0, 1)$ -transition from s_K to s_1 .




Fig. 6. The modified min-state FST instances used in the reduction of Theorem 12.

If the number of states in this transducer is a multiple of 3, add another fixed singleton state to the transducer. Reassign K to be the number of states in the resulting transducer and label the states s_1 to s_K in the order they are traversed by $(0, 0)$ -transitions, starting with the state with a $(1, 2)$ -transition to itself (formerly called s_{new}).


To enforce such solution transducers, replace a segment $0 \rightarrow 1$ in the strings S and S' output by the reduction with $01 \rightarrow 12$. Also add the segment $10^K 1 \rightarrow 20^{K-1} 12$ to the beginning of the strings to satisfy the last two constraints of the FST instances in the problem formulation. The resulting modified strings S , S' and

integer K yield an instance of MODIFIED PROMISE ENCODING BY FST that is “Yes” if and only if the 3-PARTITION instance was also “Yes”. \square

We expand the reduction above further towards the uniform height-2, 3-PATS problem now.

Theorem 12. *The uniform height-2, 3-PATS problem is NP-hard.*

Proof. Let (S, S', K) be a triple of strings $S \in \{0, 1\}^*$, $S' \in \{0, 1, 2\}^*$ and an integer K reduced from an instance of 3-PARTITION in the previous proof. We reduce it to an instance $(P, |S'| + 2K + 2)$ of the uniform height-2, 3-PATS problem, where P is the following width- $(1+|S'|+K^2)$, height-2 pattern over 3 colors $\{\blacksquare, \blacksquare, \blacksquare\}$:



where $\phi : \{0, 1, 2\}^* \rightarrow \{\blacksquare, \blacksquare, \blacksquare\}^*$ is a homomorphism that maps 0 to \blacksquare , 1 to \blacksquare , and 2 to \blacksquare , respectively, and for $1 \leq i \leq K-1$,

$$w_i = \begin{cases} \blacksquare^{(3i \bmod K)-1} \blacksquare \blacksquare^{K-3-(3i \bmod K)} & \text{if } 3i \leq K-3 \pmod{K} \\ \blacksquare^{K-3} & \text{otherwise.} \end{cases}$$


Notice that, for any $1 \leq i < j \leq K$, w_i and w_j differ in the location of their singleton gray tiles. Split the pattern P into the leftmost $|S'|+1$ columns and the remainder, called the *transduction* and *FST-constructor* gadgets, respectively. The FST-constructor gadget is further partitioned into K rectangular width- K , height-2 blocks.

Next, consider the constraints on RTASs with at most $|S'|+2K+2$ tile types that uniquely self-assemble P . Lemma 1 of Göös and Orponen [6] states that any smallest RTAS that uniquely self-assembles a pattern is directed. As we will prove, directed RTASs uniquely self-assembling P have size at least $|S'|+2K+2$ tile types; thus we need only consider directed systems.

Let the north and east glues of the seed be 0. The leftmost $|S'|+1$ locations in the bottom row of P are orange, with a cyan location following. So these positions must be tiled with orange tiles of pairwise-distinct type; the need for $|S'|+1$ distinct orange tile types thus arises. Similarly, the leftmost $K-1$ cyan locations in the bottom row must use $K-1$ distinct cyan tile types. These tile types share the south glue 0, and since the system is directed, their west glues are pairwise distinct. Label these $K-1$ cyan tile types left-to-right $t_{0,2}, t_{0,3}, \dots, t_{0,K}$ and the gray tile type immediately right $t_{0,1}$, as seen below.^b The cyan tile in the northwest corner of P cannot have the same type as any of these $K-1$ types, since otherwise this tile can

^bIn these later labels, the first subscript indicates the kind of transition of the FST that the tile type will be shown to simulate, e.g., $t_{0,i}$ is a $(0,0)$ -transition, $t_{0,1}$ and $(0,1)$ -transition, etc.

also appear in the southwest corner of P . Call this type t_0 . There are K tile types to be colored yet (illustrated as a dotted square).



These K tile types will turn out to be necessary, implying $(|S'|+1)+(K-1)+2+K-1+1 = |S'|+2K+2$ types total with $K-1$ colored gray and one colored orange. For this, we claim that the bottom row of all blocks but the first assemble identically by establishing that the gray tiles attaching to the southeast corner of the first two blocks are identical. Suppose not. Then the bottom row of the second block cannot reuse cyan tile types used in the bottom row of the first block. So the uncolored K tile types must consist of one gray and $K-1$ cyan types with south glue 0. Thus the complete tile set includes only two gray tile types both of which have the south glue 0.

Consider the gray tile attaching at the northeast corner of the first block. Its south glue is 0 and its west glue is equal to the east glue of the gray tile attaching to its immediate left. This contradicts the directedness of the system, since a cyan tile is provided with the same pair of west and south glues. Indeed, both gray tile types appear at the southeast corner of a block and to their east are cyan tiles attaching.

The verified claim brings following properties for all but the first block:

Property 1: For any $1 \leq i \leq K-1$, tiles attaching at the i -th top-row position of any two blocks but the first one have the same south glue; tiles attaching at the K -th top-row position (northeast corner) of any two blocks including the first one have the same south glue.


Property 2: Any such pair of tiles have pairwise-distinct east glues (and types).

Property 3: The assembly of the bottom row is provided with at least two different kinds of north glues.

Property 2 holds since an orange tile is placed in the northeast corner of only the last block. Thus without Property 3, $\Omega(K^2)$ tile types would be necessary to place the orange tile. Observe that for each $1 \leq i \leq K-3$, the i -th position of exactly one block is gray and the counterpart of all other blocks are cyan; for each $K-2 \leq i \leq K$, the i -th position of only the last block is orange and the counterpart of all others are gray. Thus, Properties 1 and 2 imply that the tile type set must contain one orange and $K-1$ gray tile types whose south glue is equal to the north glue of t_{01} and one gray and $K-2$ cyan tile types with a common south glue.

We claim these requirements enforce that the north glue of t_{01} is not 0. Suppose otherwise, then the former requirement implies $K-2$ extra gray tile types with south glue 0. So at most 3 tile types, including t_0 , have the non-0 south glue, and


Property 3 cannot be satisfied. Thus, the north glue of t_{01} is not 0; call it 1. Tiles attaching at the northeast corner of the blocks must all have distinct types due to Property 2, and now also their south glues must be 1. The K uncolored tile types thus have south glue 1, and one is colored orange and all the others are colored gray.




Note that a t_0 tile cannot attach anywhere in the blocks. Indeed, it causes glue mismatch with the seed being placed on the bottom row, and in order for it to attach on the top row, it must share its south glue with $K-3$ cyan tile types due to Properties 1 and 2. In summary, any minimum tile set uniquely assembling P consists of K cyan tile types, K gray ones, and $|S'|+2$ orange ones.

Now we prove constraints on the glues of these types. With only $K-1$ cyan tile types with south glue 0, even the first block must assemble its bottom row as other blocks do. That is, the bottom row of all blocks assemble as $t_{00,2}t_{00,3}\cdots t_{00,K}t_{01}$. Thus, the east glue of t_{01} is equal to the west glue of $t_{00,2}$, that is, s_1 . Since t_0 does not appear in any block, Property 2 implies that the north glues of $t_{00,2}, t_{00,3}, \dots, t_{00,K-2}$ are 0 and that the north glues of $t_{00,K-1}$ and $t_{00,K}$ are 1.

The top row of the last block is $w_K \blacksquare\blacksquare\blacksquare = \blacksquare^{K-4}\blacksquare\blacksquare\blacksquare$. The corresponding bottom row exposes the glues $0^{K-3}111$ northwards because it is assembled as $t_{00,2}\cdots t_{00,K}t_{01}$. Since the east glue of all orange tiles but t_2 is 0, the last four positions $\blacksquare\blacksquare\blacksquare\blacksquare$ of the top row are assembled as $t_{01}t_2t_2t_2$. This imposes that both the east and west glues of t_2 must be equal to the east glue of t_{01} , that is, s_1 . Since S' begins with 2, the east glue of t_0 is s_1 .



Since S' ends with 2, no tile (necessarily of type $t_{11,2}, t_{11,3}, \dots, t_{11,K}$ by Properties 1 and 2) appearing at the northeast corner of a block has east glue s_1 . Moreover, tiles attaching to their east are of type $t_{00,2}, \dots, t_{00,K}$ or t_{01} , thus their east glues are in $\{s_2, s_3, \dots, s_K\}$. Without loss of generality, assign them as follows:



So the east glues of all tile types in the FST-construction gadget are in $\{s_1, \dots, s_K\}$. The east glues of $t_{11,2}, \dots, t_{11,K}$ are distinct and selected from $\{s_1, \dots, s_K\}$. Since $t_{11,2}, \dots, t_{11,K}$ share the south glue 1 with t_2 , the west glue of t_2 is s_1 , and the system is directed, the west glues are distinct and from $\{s_2, s_3, \dots, s_K\}$.

The glue 0 is not in $\{s_1, \dots, s_K\}$, as otherwise a cyan or gray tile could appear in the southwest corner of P . So none of the cyan, gray, or t_2 tile types has east glue 0 and thus a t_0 tile cannot attach anywhere but the northwest corner of P . In order to assemble the rest of the top row, none of the orange tile types but t_2 can be used. These orange tile types share the south glue 0 with $t_{00,2}, t_{00,3}, \dots, t_{00,K}, t_{01}$ so that for the sake of directedness, their west glue cannot be taken from $\{s_1, s_2, \dots, s_K\}$ but the east glue of other tile types, which are illustrated above, is taken from this set. In summary, the tile types usable to assemble the top row except its leftmost position are all the cyan tile types, all the gray ones, and the orange tile type t_2 . We can see that these tile types simulate a transition of an FST with K states s_1, s_2, \dots, s_K by interpreting their west, south, and east glues as the source of transition, a letter read, and the target of transition, and their color as an output (via $\{\blacksquare, \blacksquare, \blacksquare\} \rightarrow \{0, 1, 2\}$).

Until the west glues $\{s_2, s_3, \dots, s_K\}$ are assigned to $t_{11,2}, t_{11,3}, \dots, t_{11,K}$, the simulated FST is not determined. Nevertheless, the other tiles imply the following properties of the FST:

- The FST has a K -cycle of $K - 1$ $(0, 0)$ -transitions and one $(0, 1)$ -transition starting at the initial state s_1 .
- The state s_1 has a $(1, 2)$ -loop.
- When making a transition from s_1 to another state, 1 is never output.
- The only transition to s_1 with output 1 is the $(0, 1)$ -transition from s_K .

Now we will prove that the FST must be the one in Figure 6. First, we will prove that the north glues of the orange tiles below $\phi(S')$ in P must be S . Observe that the tile types imply several constraints on the north glues exposed:

- A 1 north glue cannot be below a cyan tile (no cyan tile has south glue 1).
- A 0 north glue cannot be below an orange tile (no orange tile has south glue 0).
- Among consecutive gray tiles, only the eastmost can have a 0 south glue (any gray tile with south glue 0 is not followed by another gray tile, since

1 is never output at s_1).

- The eastmost gray tile in a sequence of consecutive gray tiles followed by an orange tile must have south glue 0 due to the $(0, 1)$ -transition (t_{01} is the unique gray tile type with east glue s_1 and has south glue 0).

The only remaining scenario to check is consecutive gray tiles followed by a cyan tile. Recall that S and S' consist of three kinds of segments: fixed singletons, half-fixed triples, and half-fixed intervals, and each such segment is preceded by the segment $01 \rightarrow 12$ at the end of the previous such segment (or added initial segment $10^K 1 \rightarrow 20^{K-1} 12$). The output (“ S' -part”) of a fixed singleton has the form $0^{i-1} 10^{K-i} 12$. The first 1 cannot be output by the $(0, 1)$ -transition (i.e., cannot be tile type t_{01} and must have south glue 1) as otherwise, the transduction does not reach the state s_K after processing the following 0’s and 12 is not subsequently output. Thus, the input (“ S -part”) of a fixed singleton with output $0^{i-1} 10^{K-i} 12$ must be $0^{i-1} 10^{K-i} 01$ and in order to assemble $\phi(S')$ in the pattern P , the sequence of north glues of the orange tiles below must be $0^{i-1} 10^{K-i} 01$.

The same argument implies that in order to assemble a portion of $\phi(S')$ encoding the output of a half-fixed triple, the sequence of north glues of orange tiles must encode the corresponding input of a half-fixed triple. Because S and S' always contain many fixed singleton segments, tile type $t_{11,K}$ has the west glue s_K ; moreover, no other tile type $t_{11,i}$ has west glue s_K , i.e., the FST has no other transition to s_K with output 1.

Finally, we analyze the half-fixed interval, whose output is of the form $(0^{i-1} 11010^{K-i-2} 1)(20^{i+1} 11010^{K-i-4} 1)$. The 1 in the second occurrence of 10 cannot be output by the $(0, 1)$ -transition, by the previous argument (the transduction does not reach state s_K). In order for the 1 in the first occurrence of 10 to be output by the $(0, 1)$ -transition, the preceding 1 must be the output of the $(1, 1)$ -loop at s_K . However, there are insufficient preceding 0s ($i \leq \lceil 2K/3 \rceil$ always) following the previous 2 (at the end of the previous segment) to have reached state s_K . Thus the 1’s in both occurrences of 10 in the first parenthesis are not output by the $(0, 1)$ -transition, and instead by $(1, 1)$ -transitions (i.e., are not placements of tile type t_{01} and so have south glues 1). The same arguments applies to occurrences in the second parenthesis.

The only remaining flexibility in the design of the tile set is assigning west glues to $t_{11,2}, \dots, t_{11,K-1}$. The fixed singletons, half-fixed triples, and half-fixed intervals, whose inputs are a sequence of (hidden) north glues of the bottom (orange) row of the transduction gadget that encode S , allow for only the intended assignment illustrated in Figure 6. Consequently, the assignment is possible if and only if the given 3-partition has a solution. \square

6. Efficiently Solvable PATS Problems

The non-uniform height-1 PATS problem is trivially solvable using one tile type for each color. This idea can be generalized for all patterns of fixed height:

The non-uniform height- h c -PATS problem can be solved in $c^{c^{O(h)}}$ n time.

Proof. Observe that the minimum-size tile set that assembles any given pattern has size at most $T = hc^h$, since there are c^h possible columns and a distinct set of h tiles for each column (whose appearance is programmed by the seed) can be used to assemble the pattern. The algorithm is a brute-force search for a smallest solution tile set (of size at most T), using dynamic programming to check each tile set in $O(hT^{h+2}n)$ time. Since a tile can be specified by a binary string of length $4 \log T \log c$, there are at most $2^{4T \log cT} = (cT)^{4T}$ such tile sets. Thus the algorithm runs in $O((cT)^{4T} hT^{h+2}n) = T^{O(T)}n = c^{c^{O(h)}}n$ time.

All that remains is to describe the dynamic programming algorithm. Lemma 1 of Göös and Orponen [6] states that any smallest tile set is directed. The subproblems solved have the form: “Does the tile set deterministically assemble the first i columns of the pattern with top-to-bottom sequence of tile types t_1, t_2, \dots, t_h in column i ?”

All subproblems for column i can be solved by checking each of at most $T \cdot T^h = T^{h+1}$ combinations of north seed glue in column i and “Yes” subproblems for column $i-1$, and recording “Yes” for the top-to-bottom tile sequences resulting from deterministic assembly of column i of the pattern and “No” for all other sequences. For each column, computing the “Yes” sequences takes $O(T^{h+1} \cdot hT)$ time and recording the solutions to all subproblems takes $O(T^h)$. So across all columns the algorithm takes $O(hT^{h+2}n)$ time. \square

As established in Section 5, a similar algorithm for the uniform model is impossible unless $P = NP$. Nevertheless, the uniform height-1 PATS problem can be solved in linear time using a pigeonhole argument and a DFS-based search for the longest repetitive suffix of a given height-1 pattern:

Theorem 14. *The uniform height-1 PATS problem can be solved in $O(n)$ time.*

Proof. Consider an input width- n height-1 pattern P for the uniform height-1 PATS problem as a word of length n over an alphabet of colors C . Let x, y be distinct suffixes of P with y a prefix of x and y as long as possible. Let $x = zy$ for some nonempty word z . Since y is a prefix of x , z is a period of x and thus $x = z^i z_p$ for some $i \geq 1$ and prefix z_p of z . Moreover, z is primitive (not a power of another word) since y is as long as possible.

Let $m = n - |y|$. We prove that m tile types are necessary and sufficient for a uniform RTAS to uniquely self-assemble P . For sufficiency, use a tile set that hardcodes the prefix of P preceding x with $m - |x|$ tile types and uses a repeating set of $|x| - |y|$ tile types to assemble the repetitions of z . For necessity, suppose P can be uniquely self-assembled using strictly less than m tile types. By the pigeonhole principle, there must exist $1 \leq i < j \leq m$ such that the tiles at $(i, 1)$ and $(j, 1)$ have the same type. Then the suffixes x' and y' of P starting at positions $(i, 1)$ and $(j, 1)$ are distinct suffixes of P with $y' > y$, a contradiction with the previous choice of x and y . So m is the minimum number of tile types to uniquely assemble P .

All that remains is to prove that y (and x) can be computed in $O(n)$ time. This can be done by DFS in a suffix tree, searching for the longest suffix ending at a non-leaf node. The suffix tree can be constructed from P in $O(n)$ time [4], and the DFS operates in the same running time. \square

7. Conclusion

Our work here extends the extensive prior work on the parameterized c -PATS problem to also incorporate pattern height and uniformity, and finds a more delicate complexity landscape: limited height and colors do not make the PATS problem tractable, except when combined in the non-uniform model, or in degenerate cases (height-1 or 1-PATS). A single combination of parameters and model remains unresolved; we conjecture the following:

Conjecture 15. *The uniform height-2 2-PATS problem is NP-hard.*

We encourage further parameterized analysis of problems in tile self-assembly in support of recent efforts in developing a more complete understanding of the structural complexity of tile self-assembly (see [15]).

Acknowledgements

We thank Yo-Sub Han for very fruitful discussions about finite automata and tile self-assembly, and anonymous reviewers for comments that improved the paper.

References

- [1] D. Angluin, On the complexity of minimum inference of regular sets, *Information and Control* **39** (1978) 337–350.
- [2] E. Czeizler and A. Popa, Synthesizing minimal tile sets for complex patterns in the framework of patterned DNA self-assembly, *Theoretical Computer Science* **499** (2013) 23–37.
- [3] D. Doty, Theory of algorithmic self-assembly, *Communications of the ACM* **55**(12) (2012) 78–88.
- [4] M. Farach, Optimal suffix tree construction with large alphabets, *Proceedings of the 38th Annual Symposium on Foundations of Computer Science (FOCS)*, (IEEE Computer Society, 1997), pp. 137–143.
- [5] M. R. Garey and D. S. Johnson, Complexity results for multiprocessor scheduling under resource constraints, *SIAM Journal on Computing* **4**(4) (1975) 397–411.
- [6] M. Göös and P. Orponen, Synthesizing minimal tile sets for patterned DNA self-assembly, *Proceedings of the 16th International Conference on DNA Computing and Molecular Programming (DNA 16)*, *LNCS 6518* (Springer, 2010), pp. 71–82.
- [7] A. C. Johnsen, M.-Y. Kao and S. Seki, Computing minimum tile sets to self-assemble color patterns, *Proceedings of the 24th International Symposium on Algorithms and Computation (ISAAC)*, *LNCS 8283* (Springer, 2013), pp. 699–710.
- [8] A. C. Johnsen, M.-Y. Kao and S. Seki, A manually-checkable proof for the NP-hardness of 11-color pattern self-assembly tile set synthesis, *Journal of Combinatorial Optimization* **33**(2) (2017) 496–529.

- [9] L. Kari, S. Kopecki and S. Seki, Binary pattern tile set synthesis is NP-hard, *Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming (ICALP)*, LNCS **9134** (Springer, 2015), pp. 1022–1034.
- [10] X. Ma and F. Lombardi, Synthesis of tile sets for dna self-assembly, *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **27**(5) (2008) 963–967.
- [11] M. J. Patitz, An introduction to tile-based self-assembly, *Proceedings of Unconventional Computation and Natural Computation (UCNC) 2012*, LNCS **7445** (Springer, 2012), pp. 34–62.
- [12] S. Seki, Combinatorial optimization in pattern assembly, *Proceedings of Unconventional Computation and Natural Computation (UCNC) 2013*, LNCS **7956** (Springer, 2013), pp. 220–231.
- [13] U. V. Vazirani and V. V. Vazirani, A natural encoding scheme proved probabilistic polynomial complete, *Theoretical Computer Science* **24**(3) (1983) 291–300.
- [14] E. Winfree, Algorithmic self-assembly of DNA, PhD thesis, Caltech (1998).
- [15] D. Woods, Intrinsic universality and the computational power of self-assembly, *Philosophical Transaction of the Royal Society A* **373**(2046) (2015).