

Pattern overlap implies runaway growth in hierarchical tile systems

Ho-Lin Chen* David Doty† Ján Maňuch‡ Arash Rafiey§ Ladislav Stacho¶

Abstract

We show that in the hierarchical tile assembly model, if there is a producible assembly that overlaps a nontrivial translation of itself consistently (i.e., the pattern of tile types in the overlap region is identical in both translations), then arbitrarily large assemblies are producible. The significance of this result is that tile systems intended to controllably produce finite structures must avoid pattern repetition in their producible assemblies that would lead to such overlap.

This answers an open question of Chen and Doty (*SODA 2012*), who showed that so-called “partial-order” systems producing a unique finite assembly **and** avoiding such overlaps must require time linear in the assembly diameter. An application of our main result is that any system producing a unique finite assembly is automatically guaranteed to avoid such overlaps, simplifying the hypothesis of Chen and Doty’s main theorem.

1 Introduction

Winfrey’s abstract Tile Assembly Model (aTAM) [23] is a model of crystal growth through cooperative binding of square-like monomers called *tiles*, implemented experimentally (for the current time) by DNA [2,25]. It models the potentially algorithmic capabilities of tiles that are designed to bind if and only if the total strength of attachment (summed over all binding sites, called *glues* on the tile) is at least a threshold τ , sometimes called the *temperature*. When glue strengths are integers and $\tau = 2$, two glues of strength 1 each must cooperate to bind the tile to a growing assembly. Two assumptions are key: 1) growth starts from a single *seed* tile type, and 2) only individual tiles bind to an assembly. We refer to this model as the *seeded aTAM*.

While violations of these assumptions are often viewed as errors in implementation of the seeded aTAM [20,21], relaxing them results in a different model with its own programmable abilities. In the *hierarchical* (a.k.a. *multiple tile* [1], *polyomino* [15,24], *two-handed* [3,6,9]) *aTAM*, there is no seed tile, and an assembly is considered producible so long as two producible assemblies are able to attach to each other with strength at least τ , with all individual tiles being considered as “base case” producible assemblies. In either model, an assembly is considered *terminal* if nothing can attach to it; viewing self-assembly as a computation, terminal assemblies are often interpreted to be the output. See [7,17] for an introduction to recent theoretical work using these models.

*National Taiwan University, Taipei, Taiwan, holinc@gmail.com. Supported by NSC grant number 101-2221-E-002-122-MY3.

†University of California, Davis, Davis, CA, USA, doty@ucdavis.edu. The author was supported by the Molecular Programming Project under NSF grants 0832824 and 1317694 and by NSF grants CCF-1219274 and CCF-1162589, and by a Computing Innovation Fellowship under NSF grant 1019343.

‡University of British Columbia, Vancouver BC, Canada and Simon Fraser University, Burnaby BC, Canada, jmanuch@sfu.ca. Supported by NSERC Discovery grant.

§Simon Fraser University, Burnaby, BC, Canada, arashr@sfu.ca

¶Simon Fraser University, Burnaby, BC, Canada, lstacho@sfu.ca. Supported by NSERC Discovery grant.

1.1 Informal summary of the hierarchical tile assembly model

A *tile type* is a unit square with four sides, each side having a *glue* distinguished by a label and a nonnegative integer strength. We assume a finite set T of tile types, but an infinite number of copies of each tile type. An *assembly* is a connected positioning of tiles on the integer lattice \mathbb{Z}^2 .

Two adjacent tiles in an assembly *interact* if the glue labels on their abutting sides are equal and have positive strength. The assembly is τ -*stable* if its binding graph’s minimum cut has weight at least τ , where the weight of an edge is the strength of the glue it represents.

A *hierarchical tile assembly system* is a pair $\mathcal{T} = (T, \tau)$, where T is a finite set of tile types and $\tau \in \mathbb{N}$ is the temperature. An assembly is *producible* if either it is a single tile from T , or it is the τ -stable result of translating two producible assemblies without overlap. The restriction on overlap is a model of a chemical phenomenon known as *steric hindrance* [22, Section 5.11] or, particularly when employed as a design tool for intentional prevention of unwanted binding in synthesized molecules, *steric protection* [12–14]. An assembly α is *terminal* if for every producible assembly β , α and β cannot be τ -stably attached.

1.2 Contribution of this paper

As with other models of computation, in general it is considerably more difficult to prove negative results (*limitations* on what a tile system can do) than to prove positive results. A common line of inquiry aimed at negative results in tile self-assembly concerns the notion of “pumping”: showing that a single repetition of a certain group of tiles implies that the same group can be repeated indefinitely to form an infinite periodic structure.

The *temperature-1* problem in the seeded model of tile assembly concerns the abilities of tile systems in which every positive-strength glue is sufficiently strong to bind two tiles. It may seem “obvious” that if two tile types repeat in an assembly, then a segment of tiles connecting them could be repeated indefinitely (“pumped”) to produce an infinite periodic path (since, at temperature 1, each tile along the segment has sufficient strength for the next tile in the segment to attach). However, this argument fails if the attempt to pump the segment “crashes” into an existing part of the assembly. It is conjectured [10] that only finite unions of periodic patterns (so-called *semilinear* sets) can be assembled at temperature 1 in the seeded model, but despite considerable investigation [11,16,18,19], the question remains open. If true, temperature-1 *hierarchical* tile systems would suffer a similar limitation, due to a formal connection between producible assemblies in the seeded and hierarchical models [3, Lemma 4.1]. It *has* been established, using pumping arguments, that temperature-1 seeded tile systems are unable to simulate the dynamics of certain temperature-2 systems [11].

Moving to temperature 2, both models gain power to assemble much more complex structures; both are able to simulate Turing machines, for instance. In a certain sense, the hierarchical model is at least as powerful as the seeded model, since every seeded tile system can be simulated by a hierarchical tile system with a small “resolution loss”: each tile in the seeded system is represented by a 5×5 block of tiles in the hierarchical system [3, Theorem 4.2].

From this perspective, the main theorem of this paper, a negative result on hierarchical tile assembly that does not apply to seeded tile assembly, is somewhat surprising. We show that hierarchical systems, of *any* temperature, are forced to admit a sort of infinite “pumping” behavior if a special kind of “pattern repetition” occurs. More formally, suppose that a hierarchical tile system \mathcal{T} is able to produce an assembly α_0 such that, for some nonzero vector \vec{v} , the assembly $\alpha_1 = \alpha_0 + \vec{v}$ (meaning α_0 translated by \vec{v}) intersects α_0 , but the two translations agree on every tile type in the intersection (they are *consistent*). It is known that this implies that the union $\alpha_0 \cup \alpha_1$

is producible as well [8, Theorem 5.1]. Our main theorem, Theorem 3.11, shows that this condition implies that \mathcal{T} can produce arbitrarily large assemblies, answering the main open question of [8].




Figure 1: Example of the main theorem of this paper. (a) A producible assembly α_0 . Gray tiles are all distinct types from each other, but red, green, and blue each represent one of three different tile types, so the two blue tiles are the same type. (b) By Theorem 3.9, $\alpha_0 \cup \alpha_1$ is producible, where $\alpha_1 = \alpha_0 + (2, -2)$, because they overlap in only one position, and they both have the blue tile type there. (c) α_0 and α_2 both have a tile at the same position, but the types are different (green in the case of α_0 and red in the case of α_2). (d) However, a subassembly β_i of each new α_i can grow, enough to allow the translated equivalent subassembly β_{i+1} of α_{i+1} to grow from β_i , so an infinite structure is producible.

The assembly is not necessarily infinitely many translations of all of α_0 , since although α_0 and α_1 are consistent, which implies that α_1 must be consistent with $\alpha_2 = \alpha_0 + 2\vec{v}$, it may be that α_0 is not consistent with α_2 . However, our proof shows that a subassembly β_2 of α_2 can be assembled that is sufficient to grow another translated copy of β_2 , so that the infinite producible assembly consists of infinitely many translations of β_2 . See Figure 1 for an example illustration.

An immediate application of this theorem is to strengthen a theorem of Chen and Doty [4]. They asked whether every hierarchical tile system obeying a technical condition known as the *partial order* property and producing a unique finite terminal assembly, also obeys the condition that no producible assembly is consistent with a translation of itself. The significance of the latter condition is that the main theorem of [4] shows that systems satisfying the condition obey a time lower bound for assembly: they assemble their final structure in time $\Omega(d)$, where d is the diameter of the final assembly. Our main theorem implies that every system *not* satisfying the condition must produce arbitrarily large assemblies and therefore cannot produce a unique finite terminal assembly. Hence *all* hierarchical partial order systems are constrained by this time lower bound, the same lower bound that applies to all seeded tile systems. Thus hierarchical partial order systems, despite the ability to assemble many sub-assemblies of the final assembly in parallel, provably cannot exploit this parallelism to obtain a speedup in assembly time compared to the seeded model.

It is worthwhile to note that our main theorem does not apply to the seeded model. For instance, it is routine to design a seeded tile system that assembles a unique terminal assembly shaped like a square, which uses the same tile type in the upper right and lower left corners of the square. Translating this assembly to overlap those two positions means that this tile system satisfies the hypothesis of our main theorem. Why does this not contradict the fact that this system, like all seeded systems, can be simulated by a hierarchical tile system at scale factor 5 [3, Theorem 4.2], which would apparently satisfy the same consistent overlap condition? The answer is that the hierarchical simulating system of [3] uses different 5×5 blocks to represent the same tile type from the seeded system, depending on the *sides* of the tile that are used to bind in the seeded system.

Since the upper-right corner tile and lower-left corner tile in the seeded system must clearly bind using different sides, they are represented by different blocks in the simulating hierarchical system. Hence in the hierarchical system, the terminal assembly does *not* consistently overlap with itself.

Our argument proceeds by reducing the problem (via a simple argument) to a simple-to-state theorem in pure geometry. That theorem's proof contains almost all of the technical machinery required to prove our main theorem. Let S_0 be a discrete *shape*: a finite, connected subset of \mathbb{Z}^2 , and let $\vec{v} \in \mathbb{Z}^2$ be a nonzero vector. Let $S_1 = S_0 + \vec{v}$ ($= \{ p + \vec{v} \mid p \in S_0 \}$), and let $S_2 = S_1 + \vec{v}$. The theorem states that $S_2 \setminus S_1$ (possibly a disconnected set) contains a connected component that does not intersect S_0 . This is clear when \vec{v} is large enough that $S_0 \cap S_2 = \emptyset$, but for the general case, we encourage the reader to attempt to prove it before concluding that it is obvious. In Figure 1, $S_2 \setminus S_1$ (referring respectively to the shapes of assemblies α_2 and α_1) contains two connected components, one on top and the other on bottom. The top component intersects S_0 , but not the bottom.

This problem is in turn reduced to a more technical statement about simple curves (continuous, one-to-one functions $\varphi : [0, 1] \rightarrow \mathbb{R}^2$) whose intersection implies the shapes theorem. Although we need the curve theorem to hold only for polygonal curves on the integer grid \mathbb{Z}^2 , the result holds for general curves and may be useful in other contexts.

2 The hierarchical tile assembly model

Let \mathbb{R} , \mathbb{Z} , \mathbb{N} and \mathbb{Z}^+ denote the set of all real numbers, integers, non-negative integers and positive integers, respectively. Given a set $S \subseteq \mathbb{R}^2$ and a vector $\vec{v} \in \mathbb{R}^2$, let $S + \vec{v} = \{ p + \vec{v} : p \in S \}$. A subset $A \subseteq \mathbb{Z}^2$ is *connected* if the graph $G_A = (V, E)$ is connected, where $V = A$ and $\{u, v\} \in E$ if and only if $\|u - v\|_1 = 1$. A *shape* is a finite, connected subset of \mathbb{Z}^2 .

We will consider the square lattice, i.e., the graph L_\square with the vertex set \mathbb{Z}^2 and the edge set $\{\{u, v\} : \|u - v\|_1 = 1\}$. The directions $\mathcal{D} = \{N, E, S, W\}$ are used to indicate the natural directions in the lattice. Formally, they are functions from $\mathbb{Z} \times \mathbb{Z}$ to $\mathbb{Z} \times \mathbb{Z}$: $N(x, y) = (x, y + 1)$, $E(x, y) = (x + 1, y)$, $S(x, y) = (x, y - 1)$, and $W(x, y) = (x - 1, y)$. Note that $-E = W$ and $-N = S$.

Informally, a tile is a square with the north, east, south, and west edges labeled from some finite alphabet Σ of *glues*. Formally, a tile t is a 4-tuple $(g_N, g_E, g_S, g_W) \in \Sigma^4$, indicating the glues on the north, east, south, and west side, respectively. Each pair of glues g and g' is associated with a nonnegative integer $str(g, g')$ called the *interaction strength*.

An *assembly* on a set of tiles T is a partial map $\alpha : \mathbb{Z}^2 \dashrightarrow T$ such that the subgraph of L_\square induced by the domain of α , denoted by $L_\square[\text{dom } \alpha]$, is connected, i.e., a way to place tiles on the grid so they are stable and connected. The *weighted subgraph induced by α* , denoted by $L_\square[\alpha]$, is $L_\square[\text{dom } \alpha]$ in which every edge pq has weight equal to the interaction strength of the glues on the abutting sides of tiles at positions p and q , respectively, i.e., $str(\alpha(p)_d, \alpha(q)_{-d})$ where $d = q - p$. Given a positive integer $\tau \in \mathbb{Z}^+$, called a *temperature*, a set of edges of $L_\square[\alpha]$ of an assembly α is τ -*stable* if the sum of the weights of edges in this set is at least τ , and assembly α is τ -*stable* if every edge cut of $L_\square[\alpha]$ is τ -stable.

A *hierarchical tile assembly system* is a triple $\mathcal{T} = (T, \tau, str)$, where T is a finite set of tile types, $\tau \in \mathbb{Z}^+$ and $str : \Sigma \times \Sigma \rightarrow \mathbb{N}$ is the interaction strength function. Let $\alpha, \beta : \mathbb{Z}^2 \rightarrow T$ be two assemblies. We say that α and β are *nonoverlapping* if $\text{dom } \alpha \cap \text{dom } \beta = \emptyset$. Two assemblies α and β are *consistent* if $\alpha(p) = \beta(p)$ for all $p \in \text{dom } \alpha \cap \text{dom } \beta$. If α and β are consistent assemblies, define the assembly $\alpha \cup \beta$ in a natural way, i.e., $\text{dom } (\alpha \cup \beta) = \text{dom } \alpha \cup \text{dom } \beta$ and $(\alpha \cup \beta)(p) = \alpha(p)$ for $p \in \text{dom } \alpha$ and $(\alpha \cup \beta)(p) = \beta(p)$ for $p \in \text{dom } \beta$. If α and β are nonoverlapping, the *cut of the union* $\alpha \cup \beta$ is the set of edges of L_\square with one end-point in $\text{dom } \alpha$ and the other end-point

in $\text{dom } \beta$. An assembly γ is *singular* if $|\text{dom } \gamma| = 1$. We say that an assembly γ is \mathcal{T} -*producible* if either γ is singular or there exist \mathcal{T} -producible nonoverlapping assemblies α and β such that $\gamma = \alpha \cup \beta$ and the cut of $\alpha \cup \beta$ is τ -stable. In the latter case, we write $\alpha + \beta \xrightarrow{\mathcal{T}} \gamma$. Note that every \mathcal{T} -producible assembly is τ -stable. A \mathcal{T} -producible assembly α is \mathcal{T} -*terminal* if there are no \mathcal{T} -producible assemblies β and γ such that $\alpha + \beta \xrightarrow{\mathcal{T}} \gamma$. We say two assemblies α and β are *equivalent up to translation*, written $\alpha \simeq \beta$, if there is a vector $\vec{x} \in \mathbb{Z}^2$ such that $\text{dom } \alpha = \text{dom } \beta + \vec{x}$ and for all $p \in \text{dom } \beta$, $\alpha(p + \vec{x}) = \beta(p)$. We say that \mathcal{T} *uniquely produces* α if α is \mathcal{T} -terminal and for every \mathcal{T} -terminal assembly β , $\alpha \simeq \beta$.

A *restriction* of an assembly α to a set $D \subseteq \text{dom } \alpha$, denoted by $\alpha|_D$, is $\text{dom } \alpha|_D = D$ and for every $p \in D$, $\alpha|_D(p) = \alpha(p)$. If C is a subgraph of L_\square such that $V(C) \subseteq \text{dom } \alpha$, we define $\alpha|_C = \alpha|_{V(C)}$.

When \mathcal{T} is clear from context, we may omit \mathcal{T} from the notation above and instead write \rightarrow_1 , \rightarrow , *produces*, *producible*, and *terminal*.

Figure 2 shows an example of hierarchical attachment.




Figure 2: Typical example of hierarchical assembly, at temperature $\tau = 2$. The segments between tiles represent the bonds, the number of segments encodes the strength of the bond (here, 1 or 2). In the seeded, single tile model with seed $\sigma = t_0$, the assembly at step (b) would be terminal.

3 Main result

Section 3.1 proves a theorem about curves in \mathbb{R}^2 (Theorem 3.7) that contains most of the technical detail required for our main theorem. Theorem 3.7 states that if a finite set of simple curves $\varphi_1, \dots, \varphi_k$ do not intersect each other, and for some nonzero $\vec{v} \in \mathbb{R}^2$, for each i , there is $n_i \in \mathbb{Z}^+$ so that $\varphi_i(1) = \varphi_{i+1}(0) + n_i \vec{v}$ (where $\varphi_{k+1} = \varphi_1$, i.e., each curve ends a positive integer multiple of \vec{v} from the start of the next), then some curve φ_i intersects $\varphi_i + \vec{v}$. Section 3.2 uses Theorem 3.7 to prove a geometrical theorem about shapes in \mathbb{Z}^2 (Theorem 3.8), namely that for any shape S_0 , with $S_1 = S_0 + \vec{v}$ and $S_2 = S_0 + 2\vec{v}$, it holds that $S_2 \setminus S_1$ has a connected component that does not intersect S_0 . Section 3.3 uses Theorem 3.8 to prove our main theorem (Theorem 3.11), which is that if a tile system can produce an assembly that overlaps itself consistently, then arbitrarily large assemblies are producible.

The high-level intuition of the proofs of these results is as follows (described in reverse order). Theorem 3.11 intuitively holds by the following argument. If a producible assembly α_0 is consistent with its translation $\alpha_1 = \alpha_0 + \vec{v}$ by some nonzero vector $\vec{v} \in \mathbb{Z}^2$, then Theorem 3.8 implies that some portion C of $\alpha_2 = \alpha_0 + 2\vec{v}$ does not intersect α_0 , and C is furthermore assemblable from α_1 (by Theorem 3.9). Therefore, it is assemblable from $\alpha_0 \cup \alpha_1$ (since α_2 is consistent with α_1 , and this part C of α_2 does not intersect α_0 , ruling out inconsistency due to C clashing with α_0). Thus $\alpha_1 \cup C$ is producible and overlaps consistently with its translation by \vec{v} . Since C is nonempty, $\alpha_1 \cup C$ is strictly larger than α_0 . Iterating this argument shows that arbitrarily large producible assemblies exist.

Why does Theorem 3.8 hold? If it did not, then every connected component C_i of $S_2 \setminus S_1$ would intersect S_0 at a point p_i . Since $p_i \in S_0$, $p_i + 2\vec{v} \in S_2$. Since $p_i \in S_2$, there is a path q_i from p_i to

$p_i + 2\vec{v}$ lying entirely inside of S_2 . But Corollary 3.5 implies that q_i must intersect $q_i - \vec{v}$, which, being a path inside of S_1 , implies that $p_i + 2\vec{v}$ is in a different connected component C_{i+1} of $S_2 \setminus S_1$. But since C_{i+1} also intersects S_0 , there is a point p_{i+1} in this intersection, and there is a curve φ_i from $p_i + 2\vec{v}$ to p_{i+1} . Since every connected component of $S_2 \setminus S_1$ intersects S_0 , we can repeat this argument until we return to the original connected component C_i . But then the various curves φ_i defined within each component will satisfy the conditions of Theorem 3.7, a contradiction.

3.1 A theorem about curves

Definition 3.1. Given a nonzero vector $\vec{v} \in \mathbb{R}^2$ and a point $p \in \mathbb{R}^2$ the \vec{v} -axis through p , denoted as $L_{\vec{v},p}$, is the line parallel to \vec{v} through p .

Definition 3.2. Let $\varphi : [0, 1] \rightarrow \mathbb{R}^2$ be continuous one-to-one mapping. Then $\varphi([0, 1])$ is called a *simple* (non-self-intersecting) curve from $\varphi(0)$ to $\varphi(1)$. If $\varphi : [0, 1] \rightarrow \mathbb{R}^2$ is continuous with $\varphi(0) = \varphi(1)$ and one-to-one on $[0, 1)$, then $\varphi([0, 1])$ is called a *simple closed* curve.

Obviously, any curve $\varphi([0, 1])$ from $\varphi(0)$ to $\varphi(1)$ (being a subset of the plane) can be considered also as a curve from $\varphi(1)$ to $\varphi(0)$. Therefore, for the sake of brevity, we sometimes denote this curve simply by φ and say that φ connects points $\varphi(0)$, $\varphi(1)$. If $0 \leq t_1 \leq t_2 \leq 1$, then $\varphi([t_1, t_2])$ is a simple curve as well. If φ_1 and φ_2 are simple non-closed curves such that $\varphi_1 \cap \varphi_2 = \{\varphi_1(1)\} = \{\varphi_2(0)\}$ then their concatenation $\varphi_1 \circ \varphi_2$, defined by $(\varphi_1 \circ \varphi_2)(t) = \varphi_1(2t)$ if $t \leq \frac{1}{2}$ and $(\varphi_1 \circ \varphi_2)(t) = \varphi_2(2(t - \frac{1}{2}))$ otherwise, is also a simple curve (closed if $\varphi_2(1) = \varphi_1(0)$).

Definition 3.3. Given a subset of a plane $A \subseteq \mathbb{R}^2$ and a vector $\vec{v} \in \mathbb{R}^2$, the *shift (or translation)* of A by \vec{v} , denoted by $A + \vec{v}$, is the set $A + \vec{v} = \{p + \vec{v} : p \in A\}$.

The following lemma, due to Demaine, Demaine, Fekete, Patitz, Schweller, Winslow, and Woods [5, Lemma 6.3], states that if a curve does not intersect a translation of itself, then it also does not intersect any integer multiples of the same translation. We state the lemma in terms of curves instead of shapes as in ref.[5], and for the sake of self-containment, we provide a proof stated in these terms.

Lemma 3.4 ([5]). *Consider points $p_1, p_2 \in \mathbb{R}^2$, nonzero vector $\vec{v} \in \mathbb{R}^2$ and a simple curve φ connecting p_1 and p_2 (φ may be closed if $p_1 = p_2$) such that $\varphi \cap (\varphi + \vec{v}) = \emptyset$. Let $\varphi^{\rightarrow k} = \varphi + k\vec{v}$, $k \in \mathbb{Z}$. Then all $\varphi^{\rightarrow k}$'s are mutually disjoint.*

Proof. To every point of φ we can assign “relative distance” d from the line $L_{\vec{v},p_1}$ —positive for points left to the line and negative for points right to the line (with respect to \vec{v}). Since the function $d \circ \varphi : [0, 1] \rightarrow \mathbb{R}$ is continuous, by the extreme value theorem it attains both its minimum d_{\min} and maximum d_{\max} .

If $d_{\min} = d_{\max}$ then φ is just a line segment on the line $L_{\vec{v},p_1}$ with a length less than $|\vec{v}|$ and the statement of the lemma holds true.

If $d_{\min} < d_{\max}$, let $T_{\min} = \{t \in [0, 1] : d \circ \varphi(t) = d_{\min}\}$ and $T_{\max} = \{t \in [0, 1] : d \circ \varphi(t) = d_{\max}\}$. Since both T_{\min} and T_{\max} are closed and non-empty, we can take $t_{\min} \in T_{\min}$ and $t_{\max} \in T_{\max}$ such that $d_{\min} < d \circ \varphi(t) < d_{\max}$ for every $t \in (\min\{t_{\min}, t_{\max}\}, \max\{t_{\min}, t_{\max}\})$. Denote $p_{\min} = \varphi(t_{\min})$ and $p_{\max} = \varphi(t_{\max})$. All curves $\varphi^{\rightarrow k}$, $k \in \mathbb{Z}$, lie within the stripe between lines $L_{\vec{v},p_{\min}}$ and $L_{\vec{v},p_{\max}}$. Denote $\psi = \varphi([\min\{t_{\min}, t_{\max}\}, \max\{t_{\min}, t_{\max}\}])$ (a simple curve connecting p_{\min} and p_{\max}) and let $\psi^{\rightarrow k} = \psi + k\vec{v}$, $k \in \mathbb{Z}$, be the corresponding shifts of ψ .

Since $\psi^{\rightarrow k}$ meets neither $L_{\vec{v},p_{\min}}$ nor $L_{\vec{v},p_{\max}}$ at any point except its end-points, it splits the stripe into two disjoint regions—left and right (with respect to vector \vec{v})—let us denote the left region by L_k and the right one by R_k .

Since $\varphi \cap (\varphi + \vec{v}) = \emptyset$, we have for every $k \in \mathbb{Z}$, $\psi^{\rightarrow k} \cap \varphi^{\rightarrow k+1} \subseteq \varphi^{\rightarrow k} \cap \varphi^{\rightarrow k+1} = \emptyset$. Since the point $p_{\min} + (k+1)\vec{v} \in \varphi^{\rightarrow k+1}$ lies in R_k and $\varphi^{\rightarrow k+1} \cap \psi^{\rightarrow k} = \emptyset$, the whole curve φ_{k+1} lies in R_k . Hence $\psi^{\rightarrow k+1} \subseteq R_k$ and similarly $\psi^{\rightarrow k-1} \subseteq \varphi^{\rightarrow k-1} \subseteq L_k$. This yields $R_{k+1} \subseteq R_k$ and $L_{k-1} \subseteq L_k$ and consequently $R_\ell \subseteq R_k$ and $L_k \subseteq L_\ell$ for any $k \leq \ell$, $k, \ell \in \mathbb{Z}$.

Consider now any $k < \ell$, $k, \ell \in \mathbb{Z}$. If $\ell = k+1$ then $\varphi^{\rightarrow k} \cap \varphi^{\rightarrow \ell} = \emptyset$ by the assumption of the lemma. If $\ell > k+1$ then $\varphi^{\rightarrow k} \subseteq L_{k+1}$ and $\varphi^{\rightarrow \ell} \subseteq R_{\ell-1} \subseteq R_{k+1}$, i.e., $\varphi^{\rightarrow k}$ and $\varphi^{\rightarrow \ell}$ are disjoint. \square

The following corollary of Lemma 3.4 shows that if a curve is translated by a vector \vec{v} , and the vector between its start and end points is an integer multiple of \vec{v} , then the curve must intersect its translation by \vec{v} .

Corollary 3.5. *Consider an integer $n \geq 1$, a point $p \in \mathbb{R}^2$ and a nonzero vector $\vec{v} \in \mathbb{R}^2$. Let φ be a simple curve connecting p and $p + n\vec{v}$. Then φ intersects its translation by \vec{v} .*

Proof. Assume for the sake of contradiction that φ and $\varphi + \vec{v}$ do not intersect. By Lemma 3.4 all curves $\varphi + n\vec{v}$, $n \in \mathbb{N}$, are mutually disjoint but $(p + n\vec{v}) \in \varphi \cap (\varphi + n\vec{v})$ —a contradiction. \square

The assumption that the vector from the start point to the end point of the curve φ is an integer multiple of the vector \vec{v} is essential in Corollary 3.5. The following example provides a general construction of a curve $\varphi \subseteq \mathbb{R}^2$ connecting points p and $p + x\vec{v}$ such that $\varphi \cap (\varphi + \vec{v}) = \emptyset$, where $\vec{0} \neq \vec{v} \in \mathbb{R}^2$ and $x \in \mathbb{R} \setminus \mathbb{Z}$, $|x| > 1$. Note that for $|x| < 1$ the line segment from p to $p + x\vec{v}$ does not intersect its shift by \vec{v} .

Example 3.6. For simplicity assume that $p = (0, 0)$ and $\vec{v} = (1, 0)$. Let $n = \lfloor x \rfloor$, $y = x - n$ and choose any $\varepsilon > 0$. Let μ denote the line segment (simple curve) from $(0, 0)$ to $(y, n\varepsilon)$ and ν denote the line segment from $(y, n\varepsilon)$ to $(1, -\varepsilon)$. Denote $\mu_k = \mu + k(1, -\varepsilon)$ and $\nu_k = \nu + k(1, -\varepsilon)$ for $k \in \mathbb{Z}$. Then let $\varphi = \mu_0 \circ \nu_0 \circ \dots \circ \mu_{n-1} \circ \nu_{n-1} \circ \mu_n$ be the desired curve. Figure 3 shows an example of this construction for $x = 3.6$. Note that φ starts and ends on the x -axis and that $\varphi + \vec{v}$ does not intersect φ since for each stripe between $x = i$ and $x = i+1$, $i = 1, \dots, n$, the part of $\varphi + \vec{v}$ in this stripe lies above the part of φ in the same stripe (shifted up by ε).




Figure 3: An example of a curve φ from $(0, 0)$ to $(3.6, 0)$ (solid) that does not intersect its shift $\varphi + (1, 0)$ (dashed)

The following theorem is quite technical to state. Informally, it concerns a finite set of non-intersecting curves $\varphi_1, \dots, \varphi_k$ and a vector \vec{v} of the following form. Each curve connects two points in the plane, subject to the condition that the end point of φ_i is the start point of φ_{i+1} translated by a positive integer multiple of \vec{v} , with $\varphi_{k+1} = \varphi_1$. See Figure 4(a) for an example. An alternative way to think of these curves is as a single “mostly continuous” simple closed curve, with k discontinuities allowed, where each discontinuity is of the form “jump backwards by some positive integer multiple of \vec{v} .” The theorem states that this curve must intersect its translation by \vec{v} .

Theorem 3.7. *Let $k \in \mathbb{Z}^+$, let $p_1, \dots, p_k \in \mathbb{R}^2$ be points, let $n_1, \dots, n_k \in \mathbb{Z}^+$, and let $\vec{v} \in \mathbb{R}^2$ be a nonzero vector. Then there do not exist curves $\varphi_1, \dots, \varphi_k : [0, 1] \rightarrow \mathbb{R}^2$ satisfying the following conditions:*

1. φ_i is a simple curve from p_i to $(p_{i+1} + n_{i+1}\vec{v})$, for every $1 \leq i \leq k$, where $p_{k+1} = p_1$ and $n_{k+1} = n_1$,
2. $\varphi_i \cap (\varphi_i + \vec{v}) = \emptyset$, for every $1 \leq i \leq k$,
3. $\varphi_i \cap \varphi_j = \emptyset$, for every $1 \leq i < j \leq k$.

Proof. By induction on k . The base case $k = 1$ immediately follows by Corollary 3.5.

Intuitively, the inductive case will show that if we suppose, for the sake of contradiction, that k curves exist satisfying the conditions, then we can find a common point of intersection between two of their integer translations by \vec{v} , and we can connect two subcurves of these translations to create a set of $k' < k$ curves also satisfying the hypothesis of the theorem, without introducing an intersection. Figure 4 shows an example of three curves being reduced to two. The new curves will simply be $k' - 1$ translations of some of the original k curves (which already satisfy the conditions by hypothesis), together with one new curve ψ , so our main task will be to show that ψ , in the presence of the other pre-existing curves, satisfies the three conditions.

More formally, let $k > 1$ and suppose the theorem holds for all integers $0 < k' < k$. Assume for the sake of contradiction that there are curves $\varphi_1, \dots, \varphi_k$ satisfying conditions 1, 2, and 3, and define $\varphi_m^{\rightarrow \ell} = \varphi_m + \ell\vec{v}$ for all $m \in \{1, \dots, k\}$ and $\ell \in \mathbb{N}$. We find the first intersection of φ_1 with any of curves $\varphi_m^{\rightarrow \ell}$ for all $m \in \{2, \dots, k\}$ and $\ell \in \mathbb{N}$. Let

$$\begin{aligned} t_1 &= \min\{t \in [0, 1] : (\exists m \in \{2, \dots, k\})(\exists \ell \in \mathbb{N}) \varphi_1(t) \in \varphi_m^{\rightarrow \ell}\}, \\ M &= \text{any } m \in \{2, \dots, k\} \text{ such that } (\exists \ell \in \mathbb{N}) \varphi_1(t_1) \in \varphi_m^{\rightarrow \ell}, \\ L &= \text{the unique } \ell \in \mathbb{N} \text{ such that } \varphi_1(t_1) \in \varphi_M^{\rightarrow \ell}, \\ t_2 &= \text{the unique } t \in [0, 1] \text{ such that } \varphi_1(t_1) = \varphi_M^{\rightarrow L}(t). \end{aligned}$$

Since φ_1 intersects $\varphi_2^{\rightarrow n_2}$ at $p_2 + n_2\vec{v}$ by condition 1, t_1 , M , and L are well-defined. The uniqueness of L follows by Lemma 3.4. The uniqueness of t_2 follows from the fact that $\varphi_M^{\rightarrow L}$ is simple.

Now define the curve ψ as a concatenation

$$\psi = \varphi_1([0, t_1]) \circ \varphi_M^{\rightarrow L}([t_2, 1])$$

and consider its shift

$$\psi + \vec{v} = \varphi_1^{\rightarrow 1}([0, t_1]) \circ \varphi_M^{\rightarrow L+1}([t_2, 1]).$$

In what follows we will show that points $p_1, p_{M+1} + L\vec{v}, \dots, p_k + L\vec{v}$, integers $n_1 + L, n_{M+1}, \dots, n_k$ and curves $\psi, \varphi_{M+1}^{\rightarrow L}, \dots, \varphi_k^{\rightarrow L}$ form another instance satisfying conditions 1, 2, and 3.

Observe that ψ is a curve connecting the point p_1 to the point $p_{M+1} + (n_{M+1} + L)\vec{v}$. It consists of subcurves of two simple curves whose concatenation at the intersection point $\varphi_1(t_1) = \varphi_M^{\rightarrow L}(t_2)$, by the definition of t_1 , is the first point of intersection between φ_1 and $\varphi_M^{\rightarrow L}$. The curve $\varphi_M^{\rightarrow L}$ after that point (i.e., $\varphi_M^{\rightarrow L}([t_2, 1])$) therefore cannot intersect $\varphi_1([0, t_1])$, so ψ is simple. It follows that ψ satisfies condition 1 of the new instance.

We establish that ψ does not intersect its shift by vector \vec{v} by analyzing each of the two parts of ψ , $\varphi_1([0, t_1])$ and $\varphi_M^{\rightarrow L}([t_2, 1])$, and their translations by \vec{v} , separately:

- $\varphi_1([0, t_1]) \cap \varphi_1^{\rightarrow 1}([0, t_1]) = \emptyset$, since $\varphi_1 \cap \varphi_1^{\rightarrow 1} = \emptyset$ by condition 2.
- $\varphi_M^{\rightarrow L}([t_2, 1]) \cap \varphi_M^{\rightarrow L+1}([t_2, 1]) = \emptyset$, since it follows by condition 2 that $\varphi_M^{\rightarrow L} \cap \varphi_M^{\rightarrow L+1} = \emptyset$.

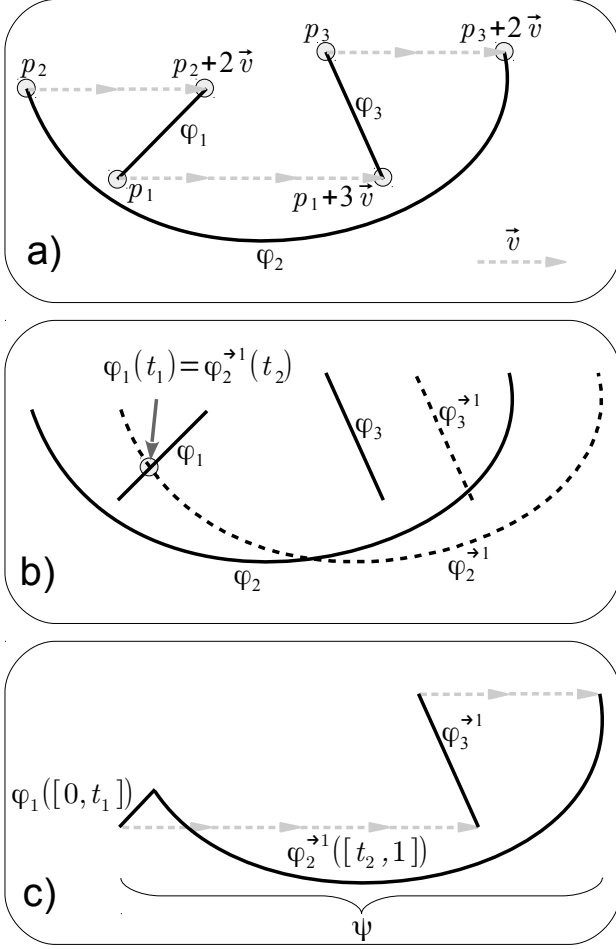


Figure 4: An example of the proof of Theorem 3.7 for $k = 3$ curves.

(a) Three curves, φ_1 , φ_2 , and φ_3 , with start and end points obeying condition 1 and also condition 3 (the curves violate condition 2, however, as Theorem 3.7 dictates they must if obeying the other two conditions). In this case, $n_1 = 3$, $n_2 = 2$, and $n_3 = 2$.

(b) Translations of curves φ_2 and φ_3 by \vec{v} , showing that φ_1 first intersects $\varphi_2^{\rightarrow 1}$, among all positive integer translations of φ_2 and φ_3 . So in this example, $M = 2$ and $L = 1$.

(c) ψ defined as the concatenation of $\varphi_1([0, t_1])$ with $\varphi_2^{\rightarrow 1}([t_2, 1])$. ψ and $\varphi_3^{\rightarrow 1}$ are the two curves produced by the proof for the inductive argument.

- $\varphi_1([0, t_1]) \cap \varphi_M^{\rightarrow L+1}([t_2, 1]) = \emptyset$, since by the definition of t_1 (in particular, the fact that it is the minimum element of the set defining it), $\varphi_1([0, t_1])$ does not intersect any $\varphi_m^{\rightarrow \ell}$, for any $m \geq 2, \ell \in \mathbb{N}$.
- $\varphi_M^{\rightarrow L}([t_2, 1]) \cap \varphi_1^{\rightarrow 1}([0, t_1]) = \emptyset$, since otherwise $\varphi_1([0, t_1])$ would intersect $\varphi_M^{\rightarrow L-1}$, violating the definition of t_1 similarly to the previous point.

This implies that ψ satisfies condition 2.




Figure 5: An example of four curves $\varphi_1, \dots, \varphi_4$ that satisfy the conditions of Theorem 3.7, except that $n_1 = \frac{3}{5}$ is not an integer.

We have $\varphi_i^{\rightarrow L} \cap \psi = \emptyset$ for every $i > M$, since $\varphi_i^{\rightarrow L}$ cannot intersect $\varphi_1([0, t_1])$ (by definition of t_1) and $\varphi_i^{\rightarrow L} \cap \varphi_M^{\rightarrow L} = \emptyset$ by condition 3. This implies that ψ satisfies condition 3 of the new instance.

Thus, the new instance with points $p_1, p_{M+1} + L\vec{v}, \dots, p_k + L\vec{v}$, integers $n_1 + L, n_{M+1}, \dots, n_k$ and curves $\psi, \varphi_{M+1}^{\rightarrow L}, \dots, \varphi_k^{\rightarrow L}$ satisfy conditions 1, 2, and 3. In addition, it has a smaller number of curves ($k+1-M = k' < k$), and hence, using the induction hypothesis we have a contradiction. \square

The example in Figure 5 shows that the theorem does not hold if we allow just one of the numbers n_1, \dots, n_k to be a non-integer.

3.2 A theorem about shapes




Figure 6: An example of a shape S_0 and its two translations. Starting at $p_1 \in (S_2 \cap S_0) \setminus S_1$, we repeat the following procedure: from point p_i in connected component C_i of $S_2 \setminus S_1$, jump to point $p_i + 2\vec{v}$, which is guaranteed to be in a different connected component C_{i+1} of $S_2 \setminus S_1$ from p_i (see proof of Theorem 3.8 to see why this is implied by Corollary 3.5). If C_{i+1} intersects S_0 at point p_{i+1} , then there is a curve φ_i in $S_2 \setminus S_1$ from $p_i + 2\vec{v}$ to p_{i+1} , and jumping to point $p_{i+1} + 2\vec{v}$ takes us to yet another connected component $C_{i+2} \neq C_{i+1}$. Repeating this must eventually result in a connected component (in this example, C_4) that does not intersect S_0 , or else the curves φ_i would contradict Theorem 3.7.

Theorem 3.7 gives rise to the following geometrical theorem about discrete shapes, which is the main technical tool to prove our main self-assembly result, Theorem 3.11. We define a *shape* to be a finite, connected subset of \mathbb{Z}^2 .

Theorem 3.8. *Let $S_0 \subset \mathbb{Z}^2$ be a shape, and let $\vec{v} \in \mathbb{Z}^2$ be a nonzero vector. Let $S_1 = S_0 + \vec{v}$ and $S_2 = S_1 + \vec{v}$. Then there is a connected component of $S_2 \setminus S_1$ that does not intersect S_0 .*

Proof. We first sketch an informal intuition of the proof, shown by example in Figure 6. The argument is constructive: it shows a way to iterate through some connected components of $S_2 \setminus S_1$ to actually find one that does not intersect S_0 .

Start with component C_1 , and suppose it intersects S_0 at point $p_1 \in C_1 \cap S_0$. Then $p_1 + 2\vec{v} \in S_2$ since $p_1 \in S_0$.¹ Let ψ be a path (simple curve) from p_1 to $p_1 + 2\vec{v}$ lying entirely within S_2 . Corollary 3.5 implies that ψ intersects $\psi - \vec{v}$, which is a curve lying entirely within S_1 . In other words, every path from p_1 to $p_1 + 2\vec{v}$ lying inside S_2 hits S_1 , i.e., $p_1 + 2\vec{v}$ and p_1 are in different connected components of $S_2 \setminus S_1$. We call $C_2 \neq C_1$ the connected component of $p_1 + 2\vec{v}$. Suppose C_2 also intersects S_0 ; then there is some curve φ_1 lying entirely within $S_2 \setminus S_1$ and going from $p_1 + 2\vec{v}$ to this new point $p_2 \in C_2 \cap S_0$. Repeating the previous argument, $p_2 + 2\vec{v}$ must be in a different connected component $C_3 \neq C_2$, and if C_3 also intersects S_0 , then there is another curve $\varphi_2 \subset C_3$ from $p_2 + 2\vec{v}$ to $p_3 \in C_3 \cap S_0$. In this example, we iterate this one more time and find that connected component $C_4 \subset S_2 \setminus S_1$ does not intersect S_0 .

For the sake of contradiction, suppose that we fail to find such a connected component, i.e., every one of the connected components C_1, \dots, C_k of $S_2 \setminus S_1$ intersects S_0 . Then eventually the above described procedure cycles back to a previously visited connected component, and the curves φ_j contained in $S_2 \setminus S_1$ satisfy condition 1 of Theorem 3.7. Since each $\varphi_i \in S_2 \setminus S_1$, we have $\varphi_i + \vec{v} \in S_3 \setminus S_2$, hence $\varphi_i \cap (\varphi_i + \vec{v}) = \emptyset$ for all $1 \leq i \leq k$, so they satisfy condition 2. Since each curve lies in a different connected component of $S_2 \setminus S_1$, they do not intersect each other, satisfying condition 3, a contradiction.

More formally, consider connected components of $S_2 \setminus S_1$, say C_1, \dots, C_k , for some $k \geq 1$. We say that C_i is *non-conflicting* if $C_i \cap S_0 = \emptyset$. We will show that there is a non-conflicting C_i . Assume for the sake of contradiction that for every $i = 1, \dots, k$, $C_i \cap S_0 \neq \emptyset$ and let $p_i \in C_i \cap S_0$. Note that $p_i + \vec{v} \in S_1$. Let n_i be the smallest positive integer such that $p_i + n_i\vec{v} \notin S_1$ (since S_1 is finite, such an n_i must exist). Since $p_i + (n_i - 1)\vec{v} \in S_1$, we have $p_i + n_i\vec{v} \in S_2 \setminus S_1$. Hence, $p_i + n_i\vec{v}$ belongs to some connected component of $S_2 \setminus S_1$. Both p_i and $p_i + n_i\vec{v}$ are in S_2 , but by Corollary 3.5, any path within S_2 connecting them must intersect its translation by $-\vec{v}$, which is a path in S_1 , so $p_i + n_i\vec{v}$ must be in a different connected component than C_i . We call this connected component C_{i+1} .²

Consider a simple curve (a self-avoiding path in the lattice) φ_i from p_i to $p_{i+1} + n_{i+1}\vec{v}$ in $C_i \subseteq S_2 \setminus S_1$. Since each φ_i lies in a different connected component, φ_i does not intersect φ_j if $i \neq j$. Furthermore, since $\varphi_i + \vec{v} \subset S_3 \setminus S_2$, φ_i does not intersect $\varphi_i \subset S_2$. But these curves $\varphi_1, \dots, \varphi_k$ contradict Theorem 3.7. \square

3.3 Implication for self-assembly

In this section we use Theorem 3.8 to prove our main theorem, Theorem 3.11. We require the following theorem from [8]. We say that two overlapping assemblies α and β are *consistent* if $\alpha(p) = \beta(p)$ for every $p \in \text{dom } \alpha \cap \text{dom } \beta$. If α and β are consistent, define their *union* $\alpha \cup \beta$ to be the assembly with $\text{dom } (\alpha \cup \beta) = \text{dom } \alpha \cup \text{dom } \beta$ defined by $(\alpha \cup \beta)(p) = \alpha(p)$ if $p \in \text{dom } \alpha$ and $(\alpha \cup \beta)(p) = \beta(p)$ if $p \in \text{dom } \beta$. Let $\alpha \cup \beta$ be undefined if α and β are not consistent.

Theorem 3.9 ([8]). *If α and β are \mathcal{T} -producible assemblies that are consistent and overlapping, then $\alpha \cup \beta$ is \mathcal{T} -producible. Furthermore, it is possible to assemble first α and then assemble the missing portions of β , i.e., $\beta|_{C_1}, \dots, \beta|_{C_k}$, where C_1, \dots, C_k are connected components of $\text{dom } \beta \setminus \text{dom } \alpha$.*

¹In this example $p_1 + 2\vec{v} \notin S_1$; in the full argument we consider $p_1 + n\vec{v}$ for $n \in \mathbb{Z}^+$ large enough to ensure this.

²Assuming we do this for every point p_i , at some point we must cycle back to a connected component already visited. It may not be that this cycle contains all connected components of $S_2 \setminus S_1$, but in this case we consider C_1, \dots, C_k to be not every connected component of $S_2 \setminus S_1$, but merely those encountered in the cycle, so that for the sake of notational convenience we can assume that C_1, \dots, C_k are all encountered, and indexed by the order in which they are encountered.

Definition 3.10. We say that assembly α is *repetitious* if there exists a nonzero vector $\vec{v} \in \mathbb{Z}^2$ such that $\text{dom } \alpha \cap \text{dom } (\alpha + \vec{v}) \neq \emptyset$ and α and $\alpha + \vec{v}$ are consistent.

Note that Theorem 3.9 implies that if a producible assembly α is repetitious with translation vector \vec{v} , then $\alpha \cup (\alpha + \vec{v})$ is also producible. The following is the main theorem of this paper.

Theorem 3.11. *Let \mathcal{T} be a hierarchical tile assembly system. If \mathcal{T} has a producible repetitious assembly, then arbitrarily large assemblies are producible in \mathcal{T} .*

Proof. It suffices to show that the existence of a producible repetitious assembly α implies the existence of a strictly larger producible repetitious assembly $\alpha' \sqsupset \alpha$. Let α be a producible repetitious assembly, with $\vec{v} \in \mathbb{Z}^2$ a nonzero vector such that α and $\alpha + \vec{v}$ overlap and are consistent. For all $i \in \{0, 1, 2\}$, let $\alpha_i = \alpha + i\vec{v}$ and $S_i = \text{dom } \alpha_i$.

By Theorem 3.8, at least one connected component $C_2 \subseteq S_2 \setminus S_1$ does not intersect S_0 . Define $C_1 = C_2 - \vec{v}$. Note that $C_1 \subseteq S_1 \setminus S_0$, which implies, since $C_2 \subseteq S_2 \setminus S_1$, that $C_2 \cap C_1 = \emptyset$. Let $\bar{\alpha} = \alpha_1|_{C_1}$. Define $\alpha' = \alpha \cup \bar{\alpha}$. By Theorem 3.9, α' is producible. Consider $\text{dom } \alpha' \cap \text{dom } (\alpha' + \vec{v})$; it suffices to show that α' and $\alpha' + \vec{v}$ are consistent on every tile type in this intersection. We have

$$\begin{aligned} \text{dom } \alpha' \cap \text{dom } (\alpha' + \vec{v}) &= (S_0 \cup C_1) \cap (S_1 \cup C_2) \\ &= (S_0 \cap S_1) \cup (S_0 \cap C_2) \cup (C_1 \cap S_1) \cup (C_1 \cap C_2) \\ &= (S_0 \cap S_1) \cup \emptyset \cup (C_1 \cap S_1) \cup \emptyset \\ &= (S_0 \cap S_1) \cup C_1, \end{aligned}$$

We handle the cases for $S_0 \cap S_1$ and C_1 separately:

$S_0 \cap S_1$: Since $C_1 \cap S_0 \cap S_1 = \emptyset$, the addition of $\bar{\alpha}$ to α_0 cannot introduce new tiles anywhere in $S_0 \cap S_1$, so only tiles from α_0 could appear here. By the hypothesis that α_0 is consistent with α_1 , α' and $\alpha' + \vec{v}$ are consistent on $S_0 \cap S_1$.

C_1 : Observe that $\alpha'|_{C_1 - \vec{v}} \sqsubset \alpha_0$ (this is the subassembly of α' that will overlap C_1 after being translated by \vec{v}) and $(\alpha' + \vec{v})|_{C_1} \sqsubset \alpha_1$, so the fact that α_0 is consistent with α_1 implies that α' and $\alpha' + \vec{v}$ are consistent on C_1 as well.

Hence α' is repetitious. Since $C_1 \subseteq S_1 \setminus S_0$ and is nonempty, $|\text{dom } \alpha'| > |\text{dom } \alpha|$. \square

Acknowledgements. The authors are extremely grateful to Jozef Haleš for the proof of Theorem 3.7. Although Jozef requested not to be a coauthor, that theorem is the keystone of the paper. The second author is also grateful to David Kirkpatrick, Pierre-Étienne Meunier, Damien Woods, Shinnosuke Seki, and Andrew Winslow for several insightful discussions. The third author would like to thank Sheung-Hung Poon for useful discussions.

References

- [1] Gagan Aggarwal, Qi Cheng, Michael H. Goldwasser, Ming-Yang Kao, Pablo Moisset de Espanés, and Robert T. Schweller. Complexities for generalized models of self-assembly. *SIAM Journal on Computing*, 34:1493–1515, 2005. Preliminary version appeared in SODA 2004.
- [2] Robert D. Barish, Rebecca Schulman, Paul W. K. Rothmund, and Erik Winfree. An information-bearing seed for nucleating algorithmic self-assembly. *Proceedings of the National Academy of Sciences*, 106(15):6054–6059, March 2009.

- [3] Sarah Cannon, Erik D. Demaine, Martin L. Demaine, Sarah Eisenstat, Matthew J. Patitz, Robert T. Schweller, Scott M. Summers, and Andrew Winslow. Two hands are better than one (up to constant factors). In *STACS 2013: Proceedings of the Thirtieth International Symposium on Theoretical Aspects of Computer Science*, pages 172–184, 2013.
- [4] Ho-Lin Chen and David Doty. Parallelism and time in hierarchical self-assembly. In *SODA 2012: Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1163–1182, 2012.
- [5] Erik D. Demaine, Martin L. Demaine, Sándor P. Fekete, Matthew J. Patitz, Robert T. Schweller, Andrew Winslow, and Damien Woods. One tile to rule them all: Simulating any Turing machine, tile assembly system, or tiling system with a single puzzle piece. In *ICALP 2014: Proceedings of the 41st International Colloquium on Automata, Languages, and Programming*, 2014.
- [6] Erik D. Demaine, Matthew J. Patitz, Trent Rogers, Robert T. Schweller, Scott M. Summers, and Damien Woods. The two-handed tile assembly model is not intrinsically universal. In *ICALP 2013: Proceedings of the 40th International Colloquium on Automata, Languages and Programming*, July 2013.
- [7] David Doty. Theory of algorithmic self-assembly. *Communications of the ACM*, 55(12):78–88, December 2012.
- [8] David Doty. Producibility in hierarchical self-assembly. In *UCNC 2014: Proceedings of 13th Unconventional Computation and Natural Computation*, 2014.
- [9] David Doty, Matthew J. Patitz, Dustin Reishus, Robert T. Schweller, and Scott M. Summers. Strong fault-tolerance for self-assembly with fuzzy temperature. In *FOCS 2010: Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science*, pages 417–426. IEEE, 2010.
- [10] David Doty, Matthew J. Patitz, and Scott M. Summers. Limitations of self-assembly at temperature 1. *Theoretical Computer Science*, 412(1–2):145–158, January 2011. Preliminary version appeared in DNA 2009.
- [11] Pierre Étienne Meunier, Matthew J. Patitz, Scott M. Summers, Guillaume Theyssier, Andrew Winslow, and Damien Woods. Intrinsic universality in tile self-assembly requires cooperation. In *SODA 2014: Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 752–771, 2014.
- [12] Kei Goto, Yoko Hinob, Takayuki Kawashima, Masahiro Kaminagab, Emiko Yanob, Gaku Yamamotob, Nozomi Takagic, and Shigeru Nagasec. Synthesis and crystal structure of a stable S-nitrosothiol bearing a novel steric protection group and of the corresponding S-nitrothiol. *Tetrahedron Letters*, 41(44):8479–8483, 2000.
- [13] Wilfried Heller and Thomas L. Pugh. “Steric protection” of hydrophobic colloidal particles by adsorption of flexible macromolecules. *Journal of Chemical Physics*, 22(10):1778, 1954.
- [14] Wilfried Heller and Thomas L. Pugh. “Steric” stabilization of colloidal solutions by adsorption of flexible macromolecules. *Journal of Polymer Science*, 47(149):203–217, 1960.
- [15] Chris Luhrs. Polyomino-safe DNA self-assembly via block replacement. *Natural Computing*, 9(1):97–109, March 2010. Preliminary version appeared in DNA 2008.

- [16] Ján Maňuch, Ladislav Stacho, and Christine Stoll. Two lower bounds for self-assemblies at temperature 1. *Journal of Computational Biology*, 17(6):841–852, 2010.
- [17] Matthew J. Patitz. An introduction to tile-based self-assembly. In *UCNC 2012: Proceedings of the 11th international conference on Unconventional Computation and Natural Computation*, pages 34–62, Berlin, Heidelberg, 2012. Springer-Verlag.
- [18] John H. Reif and Tianqi Song. Complexity and computability of temperature-1 tilings. Technical report, 2013.
- [19] Paul W. K. Rothmund and Erik Winfree. The program-size complexity of self-assembled squares (extended abstract). In *STOC 2000: Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing*, pages 459–468, 2000.
- [20] Rebecca Schulman and Erik Winfree. Synthesis of crystals with a programmable kinetic barrier to nucleation. *Proceedings of the National Academy of Sciences*, 104(39):15236–15241, 2007.
- [21] Rebecca Schulman and Erik Winfree. Programmable control of nucleation for algorithmic self-assembly. *SIAM Journal on Computing*, 39(4):1581–1616, 2009. Preliminary version appeared in DNA 2004.
- [22] Leroy G. Wade. *Organic Chemistry*. Prentice Hall, 2nd edition, 1991.
- [23] Erik Winfree. Simulations of computing by self-assembly. Technical Report CaltechC-STR:1998.22, California Institute of Technology, 1998.
- [24] Erik Winfree. Self-healing tile sets. In Junghuei Chen, Natasa Jonoska, and Grzegorz Rozenberg, editors, *Nanotechnology: Science and Computation*, Natural Computing Series, pages 55–78. Springer, 2006.
- [25] Erik Winfree, Furong Liu, Lisa A. Wenzler, and Nadrian C. Seeman. Design and self-assembly of two-dimensional DNA crystals. *Nature*, 394(6693):539–44, 1998.