

Parallelism and time in hierarchical self-assembly*

Ho-Lin Chen[†]

David Doty[‡]

Abstract

We study the role that parallelism plays in time complexity of variants of Winfree’s abstract Tile Assembly Model (aTAM), a model of molecular algorithmic self-assembly. In the “hierarchical” aTAM, two assemblies, both consisting of multiple tiles, are allowed to aggregate together, whereas in the “seeded” aTAM, tiles attach one at a time to a growing assembly. Adleman, Cheng, Goel, and Huang (*Running Time and Program Size for Self-Assembled Squares*, STOC 2001) showed how to assemble an $n \times n$ square in $O(n)$ time in the seeded aTAM using $O(\frac{\log n}{\log \log n})$ unique tile types, where both of these parameters are optimal. They asked whether the hierarchical aTAM could allow a tile system to use the ability to form large assemblies in parallel before they attach to break the $\Omega(n)$ lower bound for assembly time. We show that there is a tile system with the optimal $O(\frac{\log n}{\log \log n})$ tile types that assembles an $n \times n$ square using $O(\log^2 n)$ parallel “stages”, which is close to the optimal $\Omega(\log n)$ stages, forming the final $n \times n$ square from four $n/2 \times n/2$ squares, which are themselves recursively formed from $n/4 \times n/4$ squares, etc. However, despite this nearly maximal parallelism, the system requires superlinear time to assemble the square. We extend the definition of *partial order tile systems* studied by Adleman et al. in a natural way to hierarchical assembly and show that *no* hierarchical partial order tile system can build *any* shape with diameter D in less than time $\Omega(D)$, demonstrating that in this case the hierarchical model affords no speedup whatsoever over the seeded model. We also strengthen the $\Omega(D)$ time lower bound for deterministic seeded systems of Adleman et al. to nondeterministic seeded systems. Finally, we show that for infinitely many n , a tile system can assemble an $n \times n'$ rectangle, with $n > n'$, in time $O(n^{4/5} \log n)$, breaking the linear-time lower bound that applies to all seeded systems and partial order hierarchical systems.

1 Introduction

Tile self-assembly is an algorithmically rich model of “programmable crystal growth”. It is possible to design molecules (square-like “tiles”) with specific binding sites so that, even subject to the chaotic nature of molecules floating randomly in a well-mixed chemical soup, they are guaranteed to bind so as to deterministically form a single target shape. This is despite the number of different types of tiles possibly being much smaller than the size of the shape and therefore having only “local information” to guide their attachment. The ability to control nanoscale structures and machines to atomic-level precision will rely crucially on sophisticated self-assembling systems that automatically control their own behavior where no top-down externally controlled device could fit.

* A preliminary version of this article appeared as [14].

[†]National Taiwan University, Taipei, Taiwan, holinc@gmail.com. This author was supported by the Molecular Programming Project under NSF grant 0832824.

[‡]University of California, Davis, Davis, CA, USA, doty@ucdavis.edu. This author was supported by NSF grants 1219274, 1162589, 1619343, a Computing Innovation Fellowship under NSF grant 1019343, and the Molecular Programming Project under NSF grants 0832824 and 1317694.

A practical implementation of self-assembling molecular tiles was proved experimentally feasible in 1982 by Seeman [49] using DNA complexes formed from artificially synthesized strands. Experimental advances have delivered increasingly reliable assembly of algorithmic DNA tiles with error rates of 10% per tile in 2004 [44], 1.4% in 2007 [26], 0.13% in 2009 [8], and 0.05% in 2014 [24]. Erik Winfree [55] introduced the abstract Tile Assembly Model (aTAM) – based on a constructive version of Wang tiling [53, 54] – as a simplified mathematical model of self-assembling DNA tiles. Winfree demonstrated the computational universality of the aTAM by showing how to simulate an arbitrary cellular automaton with a tile assembly system. Building on these connections to computability, Rothmund and Winfree [45] investigated a self-assembly resource bound known as *tile complexity*, the minimum number of tile types needed to assemble a shape. They showed that for most n , the problem of assembling an $n \times n$ square has tile complexity $\Omega(\frac{\log n}{\log \log n})$, and Adleman, Cheng, Goel, and Huang [3] exhibited a construction showing that this lower bound is asymptotically tight. Under natural generalizations of the model [1, 6, 9, 13, 16, 18, 19, 33, 34, 37, 50, 51], tile complexity can be reduced for tasks such as square-building and assembly of more general shapes. See [20, 39, 58] for more background.

The authors of [3] also investigated assembly time for the assembly of $n \times n$ squares in addition to tile complexity. They define a plausible model of assembly time based (implicitly) on the standard stochastic model of well-mixed chemical kinetics [27–29] and show that under this model, an $n \times n$ square can be assembled in expected time $O(n)$, which is asymptotically optimal, in addition to having optimal tile complexity $O(\frac{\log n}{\log \log n})$. Intuitively, the optimality of the $O(n)$ assembly time for an $n \times n$ square results from the following informal description of self-assembly. The standard “seeded” aTAM stipulates that one tile type is designated as the seed from which growth nucleates, and all growth occurs by the accretion of a single tile to the assembly containing the seed. The set of locations on an assembly α where a tile could attach is called the *frontier*. An assembly with a frontier of size k could potentially have $\Theta(k)$ attachment events occur in parallel in the next “unit” of time, meaning that a speedup due to parallelism is possible in the seeded aTAM. The geometry of 2D assembly enforces that any assembly with N points has an “average frontier size” throughout assembly of size at most $O(\sqrt{N})$.¹ Therefore, the parallelism of the seeded aTAM grows at most linearly with time. To create an $n \times n$ square of size n^2 , the best parallel speedup that one could hope for would use an “average frontier size” of $O(n)$, which in $O(n)$ “parallel steps” of time assembles the entire square. This is precisely the achievement of [3].

A variant of the aTAM known as the *hierarchical* (a.k.a. *two-handed*, *recursive*, *multiple tile*, *q-tile*, *aggregation*, *polyomino*) aTAM allows non-seed tiles to aggregate together into an assembly, allows this assembly to then aggregate to other assemblies, and possibly (depending on the model) dispenses completely with the idea of a seed. Variants of the hierarchical aTAM have recently received extensive theoretical study [1, 2, 4, 6, 16–18, 22, 25, 36, 38, 40, 57]. It is intuitively conceivable that by allowing two large assemblies to form in parallel and combine in one step, it may be possible to recursively build an $n \times n$ square in $o(n)$ time, perhaps even $O(\log n)$ or $O(\text{polylog}(n))$ time. In the terminology of Reif [41], such parallelism is “distributed” rather than “local.” Determining the optimal time lower bound for uniquely self-assembling an $n \times n$ square in the hierarchical aTAM was stated as an open problem in [3].

We achieve three main results. We prove that no “partial order hierarchical system” (defined

¹For intuition, picture the fastest growing assembly: a single tile type able to bind to itself on all sides, filling the plane starting from a single copy at the origin. After t “parallel steps”, with high probability it has a circumference, and hence frontier size, of $O(t)$, while occupying area $O(t^2)$.

below) can break the $\Omega(D)$ lower bound for assembling any shape of diameter D . Next, we show that a hierarchical system violating the “partial order” property is able to assemble a rectangle of diameter D in time $o(D)$. Finally, we show a highly parallel (but surprisingly, slow) assembly of an $n \times n$ square in a hierarchical system. We now discuss these results in more detail.

Section 3 defines our model of assembly time for hierarchical tile systems. To obtain a fair comparison between our main result, Theorem 4.6, and the results for assembly time in the seeded model [3], it is necessary to introduce a definition of assembly time applicable to both seeded and hierarchical tile systems. Defining this concept is nontrivial and constitutes one of the contributions of this paper. We define such an assembly time model based on chemical kinetics. When applied to seeded systems, the model results in (nearly) the same definition used in [3], in the limit of low concentration of seed tiles.²

Section 4 shows our main result, Theorem 4.6, a linear-time lower bound on a certain class of hierarchical tile systems. In [3] the authors define a class of deterministic seeded tile systems known as *partial order systems*, which intuitively are those systems that enforce a precedence relationship (in terms of time of attachment) between any neighboring tiles in the unique terminal assembly that bind with positive strength. We extend the definition of partial order systems in a natural way to hierarchical systems, and for this special case of systems, we answer the question of [3] *negatively*, showing that $\Omega(D)$ time is required to assemble any structure with diameter D . Thus, for the purpose of speeding up self-assembly of partial order systems, the parallelism of the hierarchical assembly model is of no use whatsoever.

Section 5 shows that the partial order hypothesis is necessary to obtain a linear-time lower bound. There, we describe a hierarchical tile system that, according to our model of assembly time, can assemble a rectangle in time sublinear in its diameter. More precisely, we show that for infinitely many n , there is a hierarchical tile system that assembles an $n \times n'$ rectangle, where $n > n'$, in time $O(n^{4/5} \log n)$. The key idea is the use of both “assembly parallelism” and “binding parallelism.” By “assembly parallelism,” we mean the form of parallelism discussed above: the ability of the hierarchical model to form multiple large assemblies independently in parallel. By “binding parallelism,” we mean the (much more modest) parallelism already present in the seeded model: the ability of a single tile or assembly to have multiple potential binding sites to which to attach on the “main” growing assembly. If there are k such binding sites, the *first* such attachment will occur in expected time $\frac{1}{k}$ times that of the expected time for any *fixed* binding site to receive an attachment, a fact exploited in our tile system to achieve a speedup. We note that Theorem 3.1 implies that “binding parallelism” alone — i.e., the seeded model — cannot achieve assembly time sublinear in the diameter of the shape.

Finally, in Section 6, we show that in the hierarchical aTAM, it is possible to assemble an $n \times n$ square using nearly maximal “parallelism,” so that the full $n \times n$ square is formed from four $n/2 \times n/2$ sub-squares, which are themselves each formed from four $n/4 \times n/4$ sub-squares, etc.³ Informally,

²Low seed concentration is required to justify the assumption used in [3] of constant concentration of non-seed tiles, so we are not “cheating” by using this assumption to argue that the models nearly coincide on seeded systems. The one sense in which the models are different for seeded systems is that tile concentrations are allowed to deplete in our model. As we argue in Section 3.2, this difference does not account for our time lower bound. Furthermore, this difference makes our model strictly more realistic than the model of [3]. Tile systems in which this difference would be noticeable are those in which large assemblies not containing the seed can form, which are assumed away in the seeded model. Such systems are precisely those for which the assumptions of the seeded model are not justified. This is the sense in which our model of assembly time coincides with that of [3] when applied to the seeded model: it coincides with a slightly more realistic generalization of the model used in [3].

³If one were to assume a constant time for any two producible assemblies to bind once each is produced, this would

if tile system \mathcal{T} uniquely self-assembles a shape S , define $\text{depth}^{\text{da}}(\mathcal{T})$ to be the worst-case “number of parallel assembly steps” (depth of the tree that decomposes the final assembly recursively into the subassemblies that combined to create it) required by the tile system to reach its final assembly. (A formal definition is given in Section 6.) Clearly $\text{depth}^{\text{da}}(\mathcal{T}) \geq \log |S|$ if S is the shape assembled by \mathcal{T} . Our construction is quadratically close to this bound in the case of assembling an $n \times n$ square S_n , showing that $\text{depth}^{\text{da}}(\mathcal{T}) \leq O(\log^2 n)$. Furthermore, this is achievable using $O(\frac{\log n}{\log \log n})$ tile types, which is asymptotically optimal.⁴ That is, not only is it the case that every producible assembly can assemble into the unique terminal assembly (by the definition of unique assembly), but in fact every producible assembly is at most $O(\log^2 n)$ attachment events from becoming the terminal assembly.

Section 3 is required to understand Sections 4 and 5, but Sections 4 and 5 can be read independently of each other. Section 6 can be read independently of Sections 3, 4, and 5.

2 Informal description of the abstract tile assembly model

This section gives a brief informal sketch of the seeded and hierarchical variants of the abstract Tile Assembly Model (aTAM). See Section A for a formal definition of the aTAM.

A *tile type* is a unit square with four sides, each consisting of a *glue label* (often represented as a finite string) and a nonnegative integer *strength*. We assume a finite set T of tile types, but an infinite number of copies of each tile type, each copy referred to as a *tile*. An *assembly* (a.k.a., *supertile*) is a positioning of tiles on the integer lattice \mathbb{Z}^2 ; i.e., a partial function $\alpha : \mathbb{Z}^2 \dashrightarrow T$. Write $\alpha \sqsubseteq \beta$ to denote that α is a *subassembly* of β , which means that $\text{dom } \alpha \subseteq \text{dom } \beta$ and $\alpha(p) = \beta(p)$ for all points $p \in \text{dom } \alpha$. In this case, say that β is a *superassembly* of α . We abuse notation and take a tile type t to be equivalent to the single-tile assembly containing only t (at the origin if not otherwise specified). Two adjacent tiles in an assembly *interact* if the glue labels on their abutting sides are equal and have positive strength. Each assembly induces a *binding graph*, a grid graph whose vertices are tiles, with an edge between two tiles if they interact. The assembly is τ -*stable* if every cut of its binding graph has strength at least τ , where the weight of an edge is the strength of the glue it represents. That is, the assembly is stable if at least energy τ is required to separate the assembly into two parts. The *frontier* $\partial\alpha \subseteq \mathbb{Z}^2 \setminus \text{dom } \alpha$ of α is the set of empty locations adjacent to α at which a single tile could bind stably.

A *seeded tile assembly system* (seeded TAS) is a triple $\mathcal{T} = (T, \sigma, \tau)$, where T is a finite set of tile types, $\sigma : \mathbb{Z}^2 \dashrightarrow T$ is a finite, τ -stable *seed assembly*, and τ is the *temperature*. An assembly α is *producible* if either $\alpha = \sigma$ or if β is a producible assembly and α can be obtained from β by the stable binding of a single tile. In this case write $\beta \rightarrow_1 \alpha$ (α is producible from β by the attachment of one tile), and write $\beta \rightarrow \alpha$ if $\beta \rightarrow_1^* \alpha$ (α is producible from β by the attachment of zero or more

imply a polylogarithmic time complexity of assembling the final square. But accounting for the effect of assembly *concentrations* on binding rates in our assembly time model, the construction takes superlinear time. This is because some sub-square has concentration at most $\tilde{O}(1/n^2)$, so the time for even a single step of hierarchical assembly is at least $\tilde{\Omega}(n^2)$ by standard models of chemical kinetics. We note, however, that there are other theoretical advantages to the hierarchical model, for instance, the use of steric hindrance to enable algorithmic fault-tolerance [22]. For this reason, our highly parallel square construction may be of independent interest despite the fact that the parallelism does not confer a speedup.

⁴Without any bound on tile complexity, the problem would be trivialized by using a unique tile type for each position in the shape, each equipped with specially placed strength-1 bonds, similar to the “inter-block” bonds of Figure 7, to ensure a logarithmic-depth assembly tree.

tiles). An assembly is *terminal* if no tile can be τ -stably attached to it.

A *hierarchical tile assembly system* (hierarchical TAS) 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. An assembly α is *terminal* if for every producible assembly β , α and β cannot be τ -stably attached. The restriction on overlap is a model of a chemical phenomenon known as *steric hindrance* [52, Section 5.11] or, particularly when employed as a design tool for intentional prevention of unwanted binding in synthesized molecules, *steric protection* [30–32].

In either the seeded or hierarchical model, let $\mathcal{A}[\mathcal{T}]$ be the set of producible assemblies of \mathcal{T} , and let $\mathcal{A}_\square[\mathcal{T}] \subseteq \mathcal{A}[\mathcal{T}]$ be the set of producible, terminal assemblies of \mathcal{T} . A TAS \mathcal{T} is *directed* (a.k.a., *deterministic, confluent*) if $|\mathcal{A}_\square[\mathcal{T}]| = 1$.

3 Time complexity in the hierarchical model

In this section we define a formal notion of time complexity for hierarchical tile assembly systems. The model we use applies to both the seeded aTAM and the hierarchical aTAM.

For hierarchical systems, our assembly time model may not be completely suitable since we make some potentially unrealistic assumptions. In particular, we ignore diffusion rates of molecules based on size and assume that large assemblies diffuse as fast as individual tiles. We also assume that the binding energy τ necessary for a small tile t to attach stably to an assembly α is the same as the binding energy required for a large assembly β to attach stably to α , even though one would expect such large assemblies to have a higher reverse rate of detachment (slowing the net rate of forward growth) if bound with only strength τ . However, from the perspective of our lower bound on assembly time, Theorem 4.6, these assumptions have the effect of making hierarchical self-assembly appear *faster*. We show that even with these extra assumptions, the time complexity of hierarchical partial order systems is *still* no better than the seeded aTAM. However, caution is warranted in interpreting the upper bound result, Theorem 5.1, of a sublinear time assembly of a shape. As we discuss in Section 7, a plausible treatment of diffusion rates – together with our lower bound techniques based on low concentrations of large assemblies – may yield an absolute linear-time (in terms of diameter) lower bound on assembly time of hierarchical systems, so that Theorem 5.1 may owe its truth entirely to the heavily exploited assumption of equally fast diffusion of all assemblies. A reasonable interpretation of Theorem 5.1 is that the partial order assumption is necessary to prove Theorem 4.6 and that concentration arguments alone do not suffice to establish linear-time time lower bounds in general hierarchical systems. The techniques that weave together both “assembly parallelism” and “binding parallelism”, as discussed in Section 1 and Section 5, may prove useful in other contexts, even though their attained speedup is modest.

3.1 Definition of time complexity of seeded tile systems

We now review the definition of time complexity of seeded self-assembly proposed in [3]. A *concentrations function* on a tile set T is a subprobability measure $C : T \rightarrow [0, 1]$ (i.e., $\sum_{r \in T} C(r) \leq 1$). Each tile type r is assumed to be held at a fixed concentration $C(r)$ throughout the process of assembly.⁵ The assembly time for a seeded TAS $\mathcal{T} = (T, \sigma, \tau)$ is defined by picking a copy of

⁵For singly-seeded tile systems in which the seed tile $s \in T$ appears only once at the origin, this assumption is valid in the limit of low seed concentration $C(s)$ compared to all other concentrations $C(r)$ for $r \in T \setminus \{s\}$. This is because

the seed arbitrarily and measuring the expected time before the seed grows into some terminal assembly, when assembly proceeds according to the following stochastic model. The assembly process is described as a continuous-time Markov process in which each state represents a producible assembly, and the initial state is the seed assembly σ . For each pair of producible assemblies α, β such that $\alpha \rightarrow_1 \beta$ via the addition of tile type r , there is a transition in the Markov process from state α to state β with transition rate $C(r)$.⁶ The sink states of the Markov process are precisely the terminal assemblies. The time to reach some terminal assembly from σ is a random variable $\mathbf{T}_{\mathcal{T}, C}$, and the assembly time complexity of the seeded TAS \mathcal{T} with concentrations C is defined to be $\mathbf{T}(\mathcal{T}, C) = \mathbb{E}[\mathbf{T}_{\mathcal{T}, C}]$.

The requirement that the tile concentrations function C be a subprobability measure, rather than an arbitrary measure taking values possibly greater than 1, reflects a physical principle known as the *finite density constraint*, which stipulates that a given unit volume of solution may contain only a bounded number of molecules (if for no other reason than to avoid forming a black hole). By normalizing so that one “unit” of volume is the volume required to fit one tile, the total concentration of tiles (concentration defined as number or mass per unit volume) cannot exceed 1.⁷

We have the following time complexity lower bound for seeded systems. This theorem says that even for non-directed systems, a seeded TAS can grow its diameter only linearly with time. It strengthens and implies Lemma 4.6 of the full version of [3], which applied only to directed systems.

Let $d \in \mathbb{Z}^+$. Let $\mathcal{T} = (T, \sigma, \tau)$ be a singly-seeded TAS (meaning $|\sigma| = 1$), and let $C : T \rightarrow [0, 1]$ be a concentrations function. Since it takes only constant time for the assembly to grow to any constant radius, restricting attention to singly-seeded systems does not asymptotically affect the result for tile systems with a finite seed assembly of size larger than 1. Assume below that with probability 1, \mathcal{T} eventually places a tile at distance d (in the L_1 norm) from the seed. Define $\mathbf{D}(\mathcal{T}, C, d)$ be the random variable representing the time that any tile is first placed at distance d .

Theorem 3.1. *For each $d \in \mathbb{Z}^+$, each singly-seeded TAS \mathcal{T} , and each concentrations function $C : T \rightarrow [0, 1]$, $\mathbb{E}[\mathbf{D}(\mathcal{T}, C, d)] = \Omega(d)$.*

Proof. The intuition of the proof is as follows. We divide the plane into concentric “layers”, with layer i being the set of points at L_1 -distance i from the origin. We examine at the rate at which tiles are added to layer i , noting that such additions can only happen because of attachment to existing tiles in adjacent layers $i - 1$ and $i + 1$. Therefore, the rate of attachment in layer i is

the number of terminal assemblies (if each is of size at most K) will be limited by $C(s)$, implying the percentage change in every other tile type r 's concentration is at most $K \cdot C(s)/C(r)$; therefore “low” seed concentration means setting $C(s) \ll C(r)/K$ for all $r \in T \setminus \{s\}$. In fact, to obtain an assembly time asymptotically as fast, one need only ensure that for all r , $C(r) \geq 2\#\hat{\alpha}(r)C(s)$, where $\#\hat{\alpha}(r)$ is the number of times r appears in the terminal assembly $\hat{\alpha}$. This guarantees that the concentration of r is always at least half of its start value, which means that the assembly time, each step of which is proportional to the concentration of the tile type attaching at that step, is at most doubled compared to the case when the concentrations are held constant.

⁶That is, the expected time until the next attachment of a tile to α is an exponential random variable with rate $\sum_{r \in T} \sum_{p \in \partial^r \alpha} C(r)$, where $\partial^r \alpha$ is the r -frontier of α , the set of empty locations at which tile r could stably attach to α . Note that if r could attach at more than one location then this corresponds to separate terms in the sum; similarly, if one location p can have multiple tile types attach to it, these also correspond to separate terms in the sum.

⁷When our goal is to obtain only an asymptotic result concerning a family of tile systems assembling a family of assemblies of size/diameter D , we may relax the finite density constraint to the requirement that the concentrations sum to a constant $c \in \mathbb{R}_{\geq 0}$ independent of D , since these concentrations could be divided by c to sum to 1 while affecting the assembly time results by the same constant c , leaving the asymptotic results unaffected.

proportional to the number of tiles in layers $i - 1$ and $i + 1$. This turns out to be a process with the property that the time at which layer d gets its first tile is $\Omega(d)$, which we prove by solving some differential equations that bound the attachment process.

Since we care only about the first time at which a tile is attached at distance d (before which there are no tiles at distance d' for any $d' \geq d$), we can restrict the assembly process to the region of radius d around the seed. Therefore we model the assembly process as if it proceeds normally until the first tile attaches at distance d from the seed, at which point all growth immediately halts.

Define $\mathbb{R}_{\geq 0} = [0, \infty)$. Given $i \in \{0, \dots, d\}$ and $t \in \mathbb{R}_{\geq 0}$, let $\mathbf{X}_i(t)$ be a random variable denoting the number of tiles attached at locations with distance exactly i from the seed at time t , under the restriction stated above that all assembly halts the moment that a tile is placed at distance d . Then for all $t \in \mathbb{R}_{\geq 0}$, the event $\mathbf{X}_d(t) = 0$ (no tile is at distance d by the time t) is equivalent to the event $\mathbf{D}(\mathcal{T}, C, d) > t$ (the time of the first attachment at distance d strictly exceeds t).

In a seeded TAS, tiles can attach at a location only when there is another tile adjacent to the location. Locations at L_1 -distance i to the seed are only adjacent to locations at distance either $i + 1$ or $i - 1$ to the seed. Off the x - and y -axes, each location at distance i has two neighbors at distance $i - 1$ and two neighbors at distance $i + 1$, and for the 4 locations at distance i on either axis, every location has one neighbor at distance $i - 1$ and three neighbors at distance $i + 1$. Therefore, at time t , tiles are attachable to at most $2\mathbf{X}_{i-1}(t) + 3\mathbf{X}_{i+1}(t)$ different locations with distance i to the seed. Since the total concentration of any single tile type is at most 1, the rate at which tiles attach at any given location is at most 1. For all $i \in \{0, \dots, d\}$, define the function $f_i : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ for all $t \in \mathbb{R}_{\geq 0}$ by $f_i(t) = \mathbb{E}[\mathbf{X}_i(t)]$. Then for $i \in \{1, \dots, d - 1\}$ and $t \in \mathbb{R}_{\geq 0}$,

$$\begin{aligned} \frac{df_i(t)}{dt} &\leq 2f_{i-1}(t) + 3f_{i+1}(t), \\ \frac{df_0(t)}{dt} &= 0, \text{ and} \\ \frac{df_d(t)}{dt} &\leq 2f_{d-1}(t). \end{aligned}$$

The lack of a $3f_{d+1}(t)$ term in the latter inequality is due to our modification of the assembly process to immediately halt once the first tile attaches at distance d , which implies that $f_{d+1}(t) = 0$ for all $t \in \mathbb{R}_{\geq 0}$ since no tile is ever placed at distance $d + 1$. Since the assembly process always starts with a single seed tile, $f_0(t) = 1$ for all $t \in \mathbb{R}_{\geq 0}$, and $f_i(0) = 0$ for all $i \in \{1, \dots, d\}$. For all $t \in \mathbb{R}_{\geq 0}$ and all $i \in \{1, \dots, d\}$, $f_i(t) \leq 4i$ since there are exactly $4i$ locations at distance exactly i to the seed.

Let $t_0 \in \mathbb{R}_{\geq 0}$ be the unique time at which $f_d(t_0) = \frac{1}{2}$. This time is unique since f_d is monotonically increasing (since tiles cannot detach). Since $\mathbb{E}[\mathbf{X}_d(t_0)] = f_d(t_0) = \frac{1}{2}$, by Markov's inequality, $\Pr[\mathbf{X}_d(t_0) \geq 1] \leq \frac{1}{2}$, implying that $\Pr[\mathbf{X}_d(t_0) < 1] > \frac{1}{2}$. Since \mathbf{X}_d is integer-valued and nonnegative, this is equivalent to stating that $\Pr[\mathbf{X}_d(t_0) = 0] > \frac{1}{2}$. Recall that $\mathbf{X}_d(t_0) = 0 \iff \mathbf{D}(\mathcal{T}, C, d) > t_0$, whence $\Pr[\mathbf{D}(\mathcal{T}, C, d) > t_0] > \frac{1}{2}$. By Markov's inequality, $\mathbb{E}[\mathbf{D}(\mathcal{T}, C, d)] > \frac{t_0}{2}$. Thus it suffices to prove that $t_0 \geq \Omega(d)$. To do this, we define a simpler function that is an upper bound for f_d and solve its differential equations.

For all $i \in \{0, \dots, d\}$, define the function $g_i : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ (which will serve as an upper bound

for f_i) as follows. For all $1 \in \{1, \dots, d-1\}$ and $t \in \mathbb{R}_{\geq 0}$,

$$\begin{aligned}\frac{dg_i(t)}{dt} &= 2g_{i-1}(t) + 3g_{i+1}(t), \text{ when } g_i(t) < 4d, \\ \frac{dg_d(t)}{dt} &= 2g_{d-1}(t), \text{ when } g_d(t) < 4d, \\ \frac{dg_0(t)}{dt} &= 0,\end{aligned}$$

and for all $i \in \{1, \dots, d\}$,

$$\frac{dg_i(t)}{dt} = 0, \text{ when } g_i(t) = 4d,$$

with the boundary conditions $g_0(t) = 1$ for all $t \in \mathbb{R}_{\geq 0}$, $g_i(0) = 0$ for all $i \in \{1, \dots, d\}$. Other than the inequalities governing f_i being changed to equality with g_i , the other difference with f_i is that g_i is allowed to grow larger than $4i$ (but no larger than $4d$, which applies also to f_i whenever $i \leq d$). As a result, $g_i(t) \geq f_i(t)$ for all $i \in \{0, \dots, d\}$ and $t \in \mathbb{R}_{\geq 0}$.

Furthermore, if $g_i(t_0) > g_{i+1}(t_0)$ for all $i \in \{0, \dots, d-1\}$ at some time $t_0 \in \mathbb{R}_{\geq 0}$, then

$$\frac{dg_i(t)}{dt} \geq \frac{dg_{i+1}(t)}{dt} \text{ at time } t_0.$$

Since $g_i(0) \geq g_{i+1}(0)$ for all $i \in \{0, \dots, d\}$ by definition, the above inequality implies that $g_i(t) \geq g_{i+1}(t)$ for all $i \in \{0, \dots, d\}$ and all $t \in \mathbb{R}_{\geq 0}$. Using the fact that $g_{i-1}(t) \geq g_{i+1}(t)$, we have $2g_{i-1}(t) + 3g_{i+1}(t) \leq 5g_{i-1}(t)$. Thus, we can define a set of functions $h_i(t)$ that are upper bounds for $g_i(t)$ by the following:

$$\frac{dh_i(t)}{dt} = 5h_{i-1}(t), \text{ for all } i \in \{1, \dots, d\}, \text{ and } \frac{dh_0(t)}{dt} = 0,$$

with boundary conditions $h_0(t) = 1$ for all $t \in \mathbb{R}_{\geq 0}$, $h_i(0) = 0$ for all $i \in \{1, \dots, d\}$. Solving these differential equations, we obtain $h_d(t) = \frac{1}{d!}(5t)^d$. Letting $t' = \frac{d}{10e}$, by Stirling's inequality $d! > \sqrt{2\pi d} \left(\frac{d}{e}\right)^d e^{1/(12d+1)} > \left(\frac{d}{e}\right)^d$, we have

$$f_d(t') \leq g_d(t') \leq h_d(t') = \frac{1}{d!} \cdot (5t')^d = \frac{1}{d!} \cdot \left(\frac{d}{2e}\right)^d < \frac{1}{\left(\frac{d}{e}\right)^d} \cdot \left(\frac{d}{2e}\right)^d = \frac{1}{2^d}.$$

Since f_d is monotonically increasing, $f_d(t_0) = \frac{1}{2}$ by definition, and $\frac{1}{2^d} \leq \frac{1}{2}$ for $d \geq 1$, this implies that $t_0 \geq t' = \frac{d}{10e}$. \square

3.2 Definition of time complexity of hierarchical tile systems

3.2.1 Issues with defining hierarchical time complexity

To define time complexity for hierarchical systems, we employ more explicitly the chemical kinetics that implicitly underlie the time complexity model for seeded systems stated in Section 3.1. We treat each assembly as a single molecule. If two assemblies α and β can attach to create an assembly γ , then we model this as a chemical reaction $\alpha + \beta \rightarrow \gamma$, in which the rate constant is assumed to

be equal for all reactions (and normalized to 1). In particular, if α and β can be attached in two different ways, this is modeled as two different reactions, even if both result in the same assembly.⁸

At an intuitive level, the model we define can be explained as follows. We imagine dumping all tiles into the solution at once, and at the same time, we grab one particular tile and dip it into the solution as well, pulling it out of the solution when it has assembled into a terminal assembly. Under the seeded model, the tile we grab will be a seed, assumed to be the only copy in solution (thus requiring that it appear only once in any terminal assembly). In the seeded model, no reactions occur other than the attachment of individual tiles to the assembly we are holding. In the hierarchical model, other reactions are allowed to occur in the background (we model this using the standard mass-action model of chemical kinetics [23]), but only those reactions with the assembly we are “holding” move it closer to completion. The other background reactions merely change concentrations of other assemblies (although these indirectly affect the time it will take our chosen assembly to complete, by changing the rate of reactions with our chosen assembly).

We now discuss some intuitive justification of our model of assembly time. One reason for choosing this model is that we would like to analyze the assembly time in such a way as to facilitate direct comparison with the results of [3]. In particular, we would like the assembly time model proposed in [3] to be derived as a special case of the model we propose, when only single-tile reactions with the seed-containing assembly are allowed.⁹ With a model such as Gillespie’s algorithm [27–29] using finite molecular counts, it is possible that no copy of the terminal assembly forms, so it is not clear how to sensibly ask how long it takes to form.¹⁰ The mass-action model of kinetics [23] describes concentrations as a dynamical system that evolves continuously over time according to ordinary differential equations derived from reaction rates. This is an accepted model of kinetics when molecular counts are very large, which is already an implicit assumption in the standard aTAM. In the mass-action model, all possible terminal assemblies (assuming there are a finite number of different terminal assemblies) are guaranteed to form, which solves one issue with the purely stochastic model. But the solution goes too far: some (infinitesimal) concentration of all terminal assemblies form in the first infinitesimal amount of time, making the first appearance of a terminal assembly a useless measure of the time required to produce it. A sensible way to handle this may be to measure the time to half-completion (time required for the concentration of a terminal assembly to exceed half of its steady-state concentration). But this model is potentially subject to “cheats” such as systems that “kill” all but the fastest growing assemblies, so as to artificially inflate the average time to completion of those that successfully assemble into the terminal assembly. Furthermore, it would not necessarily be fair to directly compare such a deterministic model with the stochastic model of [3].

The model of assembly time that we define is a continuous-time, discrete-state stochastic model similar to that of [3]. However, rather than fixing transition rates at each time $t \in \mathbb{R}_{\geq 0}$ as constant,

⁸The fact that some directed systems may not require at least one of these attachments to happen in every terminal assembly tree is the reason we impose the partial order requirement when proving our time complexity lower bound.

⁹As discussed in Section 1, the model of [3] is not *exactly* a special case of our model, since we assume tile concentrations deplete. However, the assumption of constant tile concentrations is itself a simplifying assumption of [3] that is approximated by a more realistic model in which tile concentrations deplete, but seed tile types have very low concentration compared to other tile types, implying that non-seed concentrations do not deplete too much. Under this more realistic assumption, if attachments not involving the seed are disallowed, then our definition of assembly time coincides with that of [3].

¹⁰This problem is easily averted in a seeded system by setting the seed count sufficiently low to ensure that the terminal assembly is guaranteed to form at least one copy. In a hierarchical system it is not clear how to avoid this problem.

we use mass-action kinetics to describe the evolution over time of the concentration of producible assemblies, including individual tile types, which in turn determine transition rates. To measure the time to complete a terminal assembly, we use the same stochastic model as [3], which fixes attention on one particular tile¹¹ and asks what is the expected time for it to grow into a terminal assembly, where the rate of attachment events that grow it are time-dependent, governed by the continuous mass-action evolution of concentration of assemblies that could attach to it. Unlike the seeded model, we allow the tile concentrations to deplete, since it is no longer realistic (or desirable for nontrivial hierarchical constructions) to assume that individual tiles do not react until they encounter an assembly containing the seed.¹²

3.2.2 Formal definition of hierarchical time complexity

We first formally define the dynamic evolution of concentrations by mass-action kinetics. Let $\mathcal{T} = (T, \tau)$ be a hierarchical TAS, and let $C : T \rightarrow [0, 1]$ be a concentrations function. Let $\mathbb{R}_{\geq 0} = [0, \infty)$, and let $t \in \mathbb{R}_{\geq 0}$. For $\alpha \in \mathcal{A}[\mathcal{T}]$, let $[\alpha]_C(t)$ (abbreviated $[\alpha](t)$ when C is clear from context) denote the concentration of α at time t with respect to initial concentrations C , defined as follows.¹³ We often omit explicit mention of C and use the notation $[r](0)$ to mean $C(r)$, for $r \in T$, to emphasize that the concentration of r is not constant with respect to time. Given two assemblies α and β that can attach to form γ , we model this event as a chemical reaction $R : \alpha + \beta \rightarrow \gamma$. Say that a reaction $\alpha + \beta \rightarrow \gamma$ is *symmetric* if $\alpha = \beta$. Define the *propensity* (a.k.a., *reaction rate*) of R at time $t \in \mathbb{R}_{\geq 0}$ to be $\rho_R(t) = [\alpha](t) \cdot [\beta](t)$ if R is not symmetric, and $\rho_R(t) = \frac{1}{2} \cdot [\alpha](t)^2$ if R is symmetric.¹⁴

If α is consumed in reactions $\alpha + \beta_1 \rightarrow \gamma_1, \dots, \alpha + \beta_n \rightarrow \gamma_n$ and produced in asymmetric reactions $\beta'_1 + \gamma'_1 \rightarrow \alpha, \dots, \beta'_m + \gamma'_m \rightarrow \alpha$ and symmetric reactions $\beta''_1 + \beta''_1 \rightarrow \alpha, \dots, \beta''_p + \beta''_p \rightarrow \alpha$, then the concentration $[\alpha](t)$ of α at time t is described by the differential equation

$$\frac{d[\alpha](t)}{dt} = \sum_{i=1}^m [\beta'_i](t) \cdot [\gamma'_i](t) + \sum_{i=1}^p \frac{1}{2} \cdot [\beta''_i](t)^2 - \sum_{i=1}^n [\alpha](t) \cdot [\beta_i](t), \quad (3.1)$$

with boundary conditions $[\alpha](0) = C(r)$ if α is an assembly consisting of a single tile r , and $[\alpha](0) = 0$ otherwise. In other words, the propensities of the various reactions involving α determine

¹¹In the seeded model, the seed tile is the only tile to receive this attention. In our model, the tile to choose is a parameter of the definition.

¹²However, this depletion of individual tiles is *not* the source of our time lower bound. Suppose that we used a transition rate of 1 for each attachment of an individual tile (which is an upper bound on the attachment rate even for seeded systems due to the finite density constraint) and dynamic transition rates only for attachment of larger assemblies. Then the assembly of hierarchical partial order systems *still* would proceed asymptotically no faster than if single tile attachments were the only reactions allowed (as in the seeded assembly case), despite the fact that all the single-tile reactions at the intersection of the seeded and hierarchical model would be at least as fast in the modified hierarchical model as in the seeded model.

¹³More precisely, $[\alpha](t)$ denotes the concentration of the equivalence class of assemblies that are equivalent to α up to translation. We have defined assemblies to have a fixed position only for mathematical convenience in some contexts, but for defining concentration, it makes no sense to allow the concentration of an assembly to be different from one of its translations.

¹⁴That is, all reaction rate constants are equal to 1. To the extent that a rate constant models the “reactivity” of two molecules (the probability that a collision between them results in a reaction), it seems reasonable to model the rate constants as being equal. To the extent that a rate constant also models diffusion rates (and therefore rate of collisions), this assumption may not apply; we discuss the issue in Section 7. Since we are concerned mainly with asymptotic results, if rate constants are assumed equal, it is no harm to normalize them to be 1.

its rate of change, negatively if α is consumed, and positively if α is produced.

The definitions of the propensities of reactions deserve an explanation. Each propensity is proportional to the average number of collisions between copies of reactants per unit volume per unit time. For a symmetric reaction $\beta'' + \beta'' \rightarrow \alpha$, this collision rate is half that of the collision rate compared to the case where the second reactant is a distinct type of assembly, assuming it has the same concentration as the first reactant.¹⁵ Therefore the amount of α produced per unit volume per unit time is half that of a corresponding asymmetric reaction. The reason that terms of symmetric reactions $\alpha + \alpha \rightarrow \gamma$ that *consume* α are not corrected by factor $\frac{1}{2}$ is that, although the number of such reactions per unit volume per unit time is half that of a corresponding asymmetric reaction, each such reaction consumes two copies of α instead of one. This constant 2 cancels out the factor $\frac{1}{2}$ that would be added to correct for the symmetry of the reaction. Therefore, the term $[\alpha](t) \cdot [\beta_i](t)$ representing the rate of consumption of α is the proper value whether or not $\alpha = \beta_i$.

This completes the definition of the dynamic evolution of concentrations of producible assemblies; it remains to define the time complexity of assembling a terminal assembly. Although we have distinguished between seeded and hierarchical systems, for the purpose of defining a model of time complexity in hierarchical systems and comparing them to the seeded system time complexity model of [3], it is convenient to introduce a seed-like “timekeeper tile” into the hierarchical system, in order to stochastically analyze the growth of this tile when it reacts in a solution that is itself evolving according to the continuous model described above. The seed does not have the purpose of nucleating growth, but is introduced merely to focus attention on a single molecule that has not yet assembled anything, in order to ask how long it will take to assemble into a terminal assembly.¹⁶ The choice of which tile type to pick will be a parameter of the definition, so that a system may have different assembly times depending on the choice of timekeeper tile.

Fix a copy of a tile type s to designate as a “timekeeper seed”. The assembly of s into some terminal assembly $\hat{\alpha}$ is described as a time-dependent continuous-time Markov process in which each state represents a producible assembly containing s , and the initial state is the size-1 assembly with only s . For each state α representing a producible assembly with s at the origin, and for each pair of producible assemblies β, γ such that $\alpha + \beta \rightarrow \gamma$ (with the translation assumed to happen only to β so that α stays “fixed” in position), there is a transition in the Markov process from state α to state γ with transition rate $[\beta](t)$.¹⁷ Unlike the seeded model, the transition rates vary over time since the assemblies (including assemblies that are individual tiles) with which α could

¹⁵For intuition, consider finite counts: with n copies of γ and n copies of $\beta \neq \gamma$, there are n^2 distinct pairs of molecules of respective type γ and β , but with only n copies of γ , there are $\frac{n(n-1)}{2}$ distinct pairs of molecules of type γ , which approaches $\frac{1}{2}n^2$ as $n \rightarrow \infty$.

¹⁶For our lower bound result, Theorem 4.6, it will not matter which tile type is selected as the timekeeper, except in the following sense. We define partial order systems, the class of directed hierarchical TAS’s to which the bound applies, also with respect to a particular tile type in the unique terminal assembly. A TAS may be a partial order system with respect to one tile type but not another, but for all tile types s for which the TAS is a partial order system, the time lower bound of Theorem 4.6 applies when s is selected as the timekeeper. The upper bound of Theorem 5.1 holds with respect to only a single tile type.

¹⁷That is, for the purpose of determining the continuous dynamic evolution of the concentration of assemblies, including α , in solution at time t , the rate of the reaction $\alpha + \beta \rightarrow \gamma$ at time t is assumed to be proportional to $[\alpha](t)[\beta](t)$ (or half this value if the reaction is symmetric). However, for the purpose of determining the stochastic dynamic evolution of one particular copy of s , the rate of this reaction at time t is assumed to be proportional only to $[\beta](t)$. This is because we want to describe the rate at which *this particular copy* of α , the one containing the copy of s that we fixed at time 0, encounters assemblies of type β . This instantaneous rate is independent of the number of *other* copies of α at time t (although after ϵ seconds the rate will change to $[\beta](t + \epsilon)$, which of course *will* depend on $[\alpha]$ over that time interval).

interact are themselves being produced and consumed.

We define $\mathbf{T}_{\mathcal{T},C,s}$ to be the random variable representing the time taken for the copy of s to assemble into a terminal assembly via some sequence of reactions as defined above. We define the time complexity of a directed hierarchical TAS \mathcal{T} with concentrations C and timekeeper s to be $\mathsf{T}(\mathcal{T}, C, s) = \mathbb{E}[\mathbf{T}_{\mathcal{T},C,s}]$, and the time complexity with respect to s as $\mathsf{T}(\mathcal{T}, s) = \min_C \mathsf{T}(\mathcal{T}, C, s)$.¹⁸

We note in particular that our construction of Theorem 6.1 is composed of $(\frac{n}{\log n})^2$ different types of $O(\log n) \times O(\log n)$ “blocks” that can each grow via only one reaction. At least one of these blocks β must obey $[\beta](t) \leq \frac{\log^2 n}{n^2}$ for all $t \in \mathbb{R}_{\geq 0}$ (this can be seen by applying Lemma 4.5, proven in Section 4.3). This implies that the rate of the slowest such reaction has rate at most $\frac{\log^2 n}{n^2}$, hence expected time at least the inverse of that quantity. Thus our square construction assembles in at least $\Omega(\frac{n^2}{\log^2 n})$ time, slower than the optimal seeded time of $O(n)$ [3]. Proving this formally requires more details that we omit. However, it is simple to modify the system to have a tile type appearing in exactly one position in the terminal assembly, for example, by attaching such a tile type only to the block at coordinate $(0, 0)$. It is routine to check that this would make the system a partial order system with respect to that tile type as defined in Section 4.1. Then Theorem 4.6 implies a time complexity lower bound of $\Omega(n)$, much slower than the polylogarithmic time one might naïvely expect due to the polylogarithmic depth of the assembly tree.

4 Time complexity lower bound for hierarchical partial order systems

In this section we show that the subset of hierarchical TAS’s known as *partial order systems* cannot assemble any shape of diameter D in faster than time $\Omega(D)$.

4.1 Definition of hierarchical partial order systems

Seeded partial order systems were first defined by Adleman, Cheng, Goel, and Huang [3] for the purpose of analyzing the running time of their optimal square construction. Intuitively, a seeded directed TAS with unique terminal assembly $\hat{\alpha}$ is a partial order system if every pair of adjacent positions p_1 and p_2 in $\hat{\alpha}$ that interact with positive strength have the property that either p_1 always receives a tile before p_2 , or vice versa. We extend the definition of partial order systems to hierarchical systems in the following way.

Let $\mathcal{T} = (T, \tau)$ be a hierarchical directed TAS with unique terminal assembly $\hat{\alpha} \in \mathcal{A}_{\square}[\mathcal{T}]$. A *terminal assembly tree* of $\hat{\alpha}$ is a full binary tree with $|\hat{\alpha}|$ leaves, in which each leaf is labeled with an assembly consisting of a single tile, the root is labeled with $\hat{\alpha}$, and each internal node is labeled with an assembly producible in one step from the τ -stable attachment of its two child assemblies.¹⁹ Let Υ be any terminal assembly tree of \mathcal{T} . Let $p \in \text{dom } \hat{\alpha}$ and let $s = \hat{\alpha}(p)$. The *assembly sequence with respect to Υ starting at p* is the sequence of assemblies $\vec{\alpha}_{p,\Upsilon} = (\alpha_1, \dots, \alpha_k)$ that represent the path from the leaf corresponding to p to the root of Υ , so that α_1 is the single tile s at position

¹⁸It is worth noting that this expected value could be infinite. This would happen if some partial assembly α , in order to complete into a terminal assembly, requires the attachment of some assembly β whose concentration is depleting quickly.

¹⁹Note that even a directed hierarchical TAS may have more than one assembly tree of the terminal assembly $\hat{\alpha}$.

p , and $\alpha_k = \hat{\alpha}$.²⁰ An *assembly sequence starting at p* is an assembly sequence with respect to Υ starting at p , for some valid assembly tree Υ .

An *attachment quasiorder* with respect to $p \in \text{dom } \hat{\alpha}$ is a quasiorder (a reflexive, transitive relation) \preceq on $\text{dom } \hat{\alpha}$ such that the following holds:

1. For every $p_1, p_2 \in \text{dom } \hat{\alpha}$, $p_1 \preceq p_2$ if and only if for every assembly sequence $\vec{\alpha} = (\alpha_1, \dots, \alpha_k)$ starting at p , for all $1 \leq i \leq k$, $\alpha_i(p_2)$ is defined $\implies \alpha_i(p_1)$ is defined. In other words, p_1 must always have a tile by the time p_2 has a tile. (Perhaps they arrive at the same time, if they are both part of some assembly that attaches in a single step.)
2. For every pair of adjacent positions $p_1, p_2 \in \text{dom } \hat{\alpha}$, if the tiles at positions p_1 and p_2 interact with positive strength in $\hat{\alpha}$, then $p_1 \preceq p_2$ or $p_2 \preceq p_1$ (or both).

If two tiles always arrive at the same time to the assembly containing p , then they will be in the same equivalence class induced by \preceq . Given an attachment quasiorder \preceq , we define the *attachment partial order \prec induced by \preceq* to be the strict partial order on the quotient set of equivalence classes induced by \preceq . In other words, if some subassembly $\alpha \sqsubseteq \hat{\alpha}$ always attaches to the assembly containing p all at once, then all positions $p' \in \text{dom } \alpha$ will be equivalent under \preceq .²¹ It is these equivalence classes of positions that are related under \prec . Each attachment partial order \prec induces a directed acyclic graph $G = (V, E)$, where $V = \{\beta_1, \dots, \beta_k\}$, each β_i represents the subassembly corresponding to some equivalence class (under \preceq) of positions in $\text{dom } \hat{\alpha}$, and $(\beta_i, \beta_j) \in E$ if $\text{dom } \beta_i \prec \text{dom } \beta_j$.²² Note that the first assembly α_1 of the assembly sequence containing p is always size 1, since by definition p is the only position with a tile at time 0.

We say that a directed hierarchical TAS \mathcal{T} with unique terminal assembly $\hat{\alpha}$ is a *hierarchical partial order system with respect to p* if it has an attachment quasiorder with respect to p . Given a tile type $s \in T$ that appears exactly once in the terminal assembly $\hat{\alpha}$ at position p (i.e., $\hat{\alpha}(p) = s$ and $(\forall q \in \text{dom } \hat{\alpha} \setminus \{p\}) \hat{\alpha}(q) \neq s$), we say that \mathcal{T} is a *hierarchical partial order system with respect to s* if \mathcal{T} is a hierarchical partial order system with respect to p .

Remark. The condition that s appears exactly once in $\hat{\alpha}$ allows us to talk interchangeably of a partial order with respect to a *position* and a partial order with respect to a *tile type*. If instead we allowed s to appear in multiple positions in $\hat{\alpha}$, then the sequence of attachments would nondeterministically choose one of them. Since the partial order imposed on $\text{dom } \hat{\alpha}$ is different depending on the position chosen to define the partial order, we would lose the ability to focus on “the” partial order imposed by s (since there would be more than one possible). It is conceivable that a tile type s could appear in multiple positions in $\hat{\alpha}$, and \mathcal{T} could be a partial order with respect to all of these positions. It is an open question whether Theorem 4.6 would apply to such a system. Our proof technique relies fundamentally on s appearing in a single fixed position in $\hat{\alpha}$ and defining a single partial order on $\text{dom } \hat{\alpha}$ that is then used to establish the assembly time lower bound. The

²⁰That is, $\vec{\alpha}$ is like a seeded assembly sequence in that each α_i is a *subassembly* of α_{i+1} (written $\alpha_i \sqsubseteq \alpha_{i+1}$, meaning $\text{dom } \alpha_i \subseteq \text{dom } \alpha_{i+1}$ and $\alpha_i(p) = \alpha_{i+1}(p)$ for all $p \in \text{dom } \alpha_i$). The difference is that α_i and α_{i+1} may differ in size by more than one tile, since $\text{dom } \alpha_{i+1} \setminus \text{dom } \alpha_i$ will consist of all points in the domain of α_i 's sibling in Υ .

²¹More generally, if there is a subset $X \subset \text{dom } \hat{\alpha}$ such that all assemblies α attaching to the assembly containing p have the property that $X \cap \text{dom } \alpha \neq \emptyset \implies X \subseteq \text{dom } \alpha$; i.e., any position in X attaching implies all positions in X attach with it, then this implies all positions in X are equivalent under \preceq .

²²i.e., if the positions are nodes on a directed graph with an edge from p_1 to p_2 if $p_1 \preceq p_2$, then each equivalence class is a strongly connected component of the graph, and \prec describes the *condensation* directed acyclic graph obtained by contracting each strongly connected component into a single vertex.

definition of seeded partial order systems [3, 5] has a similar requirement, that the seed tile type appear exactly once in the terminal assembly.

In the case of seeded assembly, in which each attachment is of a “subassembly” containing a single tile to a subassembly containing the seed, this definition of partial order system is equivalent to the definition of partial order system given in [3]. In the seeded case, since no tiles may attach simultaneously, the attachment quasiorder we have defined induces equivalence classes that are singletons, and the resulting induced strict partial order is the same as the strict partial order used in [3].

4.2 Repetitious assemblies

This section shows that hierarchical partial order systems are well-behaved in a certain technical sense that will be useful in the proof of Theorem 4.6.

Definition 4.1. 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.

Definition 4.2. Let α be a producible assembly, let $\vec{v} \in \mathbb{Z}^2$ be a vector, and let $\alpha + \vec{v}$ denote the translation of α by \vec{v} , i.e., an assembly β such that $\text{dom } \beta = \text{dom } \alpha + \vec{v}$ and $\beta(p) = \alpha(p - \vec{v})$ for all $p \in \text{dom } \beta$. 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.

The following theorem shows an important property of hierarchical systems that will be useful.

Theorem 4.3 ([15]). *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} .*

Let α, β be producible assemblies that can attach. Imagine fixing the position of α so that β must be translated to attach to α . Let $V_{\alpha, \beta}$ be the set of all vectors \vec{v} such that $\text{dom } \alpha \cap (\text{dom } \beta + \vec{v}) = \emptyset$ and $\alpha \cup (\beta + \vec{v})$ is a stable assembly; i.e., $V_{\alpha, \beta}$ describes the set of all ways to attach β to α . We say that α and β are *disjointly attachable* if, for every $\vec{v}_1, \vec{v}_2 \in V_{\alpha, \beta}$ such that $\vec{v}_1 \neq \vec{v}_2$, it holds that $\text{dom } (\beta + \vec{v}_1) \cap \text{dom } (\beta + \vec{v}_2) = \emptyset$; in other words, no two translations of β that allow it to attach to α overlap each other.

Corollary 4.4. *Let $\mathcal{T} = (T, \tau)$ be a hierarchical partial order system with respect to tile type $s \in T$ appearing at position $p \in \text{dom } \hat{\alpha}$ in the terminal assembly $\hat{\alpha}$, and let $(\alpha_1, \dots, \alpha_k)$ be an assembly sequence starting at p . Then for every $i \in \{1, \dots, k\}$, every producible assembly β that can attach to α_i is disjointly attachable to α_i .*

Proof. Because s appears only at position $p \in \text{dom } \hat{\alpha}$ and $\hat{\alpha}$ is the unique terminal assembly, every position relative to the position of s has a fixed tile type appearing there.

Suppose for the sake of contradiction that there is some producible assembly β that is attachable to α_i , but not disjointly, so that there are two vectors $\vec{v}_1 \neq \vec{v}_2$ such that, defining $\beta_1 = \beta + \vec{v}_1$ and $\beta_2 = \beta + \vec{v}_2$, we have that $\alpha \cup \beta_1$ and $\alpha \cup \beta_2$ are both stable assemblies (and $\text{dom } \alpha \cap \text{dom } \beta_1 = \emptyset$ and $\text{dom } \alpha \cap \text{dom } \beta_2 = \emptyset$), but $\text{dom } \beta_1 \cap \text{dom } \beta_2 \neq \emptyset$. Therefore, for every position in the overlap $q \in \text{dom } \beta_1 \cap \text{dom } \beta_2$, $\beta_1(q) = \beta_2(q)$, otherwise this would imply one terminal assembly

producible from α_i that has one tile type at position q and another that has a different tile type at position q . But this is precisely what it means for β to be a repetitious assembly. Theorem 4.3 then implies that arbitrarily large assemblies are producible in \mathcal{T} , contradicting the fact that it produces a unique terminal assembly. \square

4.3 Linear time lower bound for partial order systems

Theorem 4.6 establishes that hierarchical partial order systems, like their seeded counterparts, cannot assemble a shape of diameter D in less than $\Omega(D)$ time. This is potentially counterintuitive, since an attaching assembly of size K is able to increase the size of the growing assembly by K tiles in a single attachment step. Intuitively, the time lower bound is proven by using the fact that such an assembly can have concentration at most $\frac{1}{K}$ by conservation of mass, slowing down its rate of attachment (compared to the rate of a single tile) by factor at least K , precisely enough to cancel out the potential speedup over a single tile due to its size.

This simplistic argument is not quite accurate and must be amortized — using our Conservation of Mass Lemma (Lemma 4.5) — over all assemblies that could extend the growing assembly. The growing assembly may be extended at a single attachment site by more than one assembly. However, by Lemma 4.5, these assemblies must *collectively* have limited total concentration. Intuitively, the property of having a partial order on binding subassemblies ensures that the assembly of each path in the partial order graph proceeds by a series of rate-limiting steps. We prove upper bounds on each of these rates using this concentration argument.²³ Since the rate-limiting steps must occur in order, we can then use linearity of expectation to bound the total expected time.

The following is a “conservation of mass lemma” that will be helpful in the proof of Theorem 4.6. Note that it applies to any hierarchical system.

Lemma 4.5 (Conservation of Mass Lemma). *Let $\mathcal{T} = (T, \tau)$ be a hierarchial TAS and let $C : T \rightarrow [0, 1]$ be a concentrations function. Then for all $t \in \mathbb{R}_{\geq 0}$,*

$$\sum_{\alpha \in \mathcal{A}[\mathcal{T}]} [\alpha](t) \cdot |\alpha| = \sum_{r \in T} C(r) \quad (\leq 1).$$

Proof. For all $t \in \mathbb{R}_{\geq 0}$, define $f(t) = \sum_{\alpha \in \mathcal{A}[\mathcal{T}]} [\alpha](t) \cdot |\alpha|$. According to our model, $[\alpha](0) = C(r)$ if α consists of a single tile type r and $[\alpha](0) = 0$ otherwise, so $f(0) = \sum_{r \in T} C(r)$. Therefore it is sufficient (and necessary) to show that $\frac{df}{dt} = 0$. For all $\alpha \in \mathcal{A}[\mathcal{T}]$ and $t \in \mathbb{R}_{\geq 0}$, define $f_\alpha(t) = [\alpha](t) \cdot |\alpha|$. Then by equation (3.1), and recalling from that equation the definitions of m , n , p , β'_i , β''_i , γ'_i , and β_i , annotated as $m(\alpha)$, $n(\alpha)$, etc. to show their dependence on α , we have

$$\frac{df_\alpha}{dt} = |\alpha| \cdot \left(\sum_{i=1}^{m(\alpha)} [\beta'_i(\alpha)](t) \cdot [\gamma'_i(\alpha)](t) + \sum_{i=1}^{p(\alpha)} \frac{1}{2} \cdot [\beta''_i(\alpha)](t)^2 - \sum_{i=1}^{n(\alpha)} [\alpha](t) \cdot [\beta_i(\alpha)](t) \right).$$

²³The same assembly α could attach to many locations p_1, \dots, p_n . In a TAS that is not a partial order system, it could be the case that there is not a fixed attachment location that is necessarily required to complete the assembly. In this case completion of the assembly might be possible even if only one of p_1, \dots, p_n receives the attachment of α . Since the minimum of n exponential random variables with rate $1/K$ is itself exponential with rate n/K , the very first attachment of α to any of p_1, \dots, p_n happens in expected time K/n , as opposed to expected time K for α to attach to a *particular* p_i . This prevents our technique from applying to such systems, and it is the fundamental speedup technique in our proof of Theorem 5.1.

Then

$$\begin{aligned} \frac{df}{dt} &= \frac{d}{dt} \sum_{\alpha \in \mathcal{A}[\mathcal{T}]} f_\alpha(t) = \sum_{\alpha \in \mathcal{A}[\mathcal{T}]} \frac{df_\alpha}{dt} \\ &= \sum_{\alpha \in \mathcal{A}[\mathcal{T}]} \left(\begin{aligned} &\sum_{i=1}^{m(\alpha)} |\alpha| \cdot [\beta'_i(\alpha)](t) \cdot [\gamma'_i(\alpha)](t) + \sum_{i=1}^{p(\alpha)} |\alpha| \cdot \frac{1}{2} \cdot [\beta''_i(\alpha)](t)^2 \\ &- \sum_{i=1}^{n(\alpha)} |\alpha| \cdot [\alpha](t) \cdot [\beta_i(\alpha)](t) \end{aligned} \right). \end{aligned}$$

Let \mathcal{R} denote the set of all attachment reactions of \mathcal{T} , writing $R(\alpha, \beta, \gamma)$ to denote the reaction $\alpha + \beta \rightarrow \gamma$. For each such reaction, $|\alpha| + |\beta| = |\gamma|$. In particular, if $\alpha = \beta$, then $|\gamma| = 2|\alpha|$. Each such asymmetric reaction contributes precisely three unique terms in the right hand side above: two negative (of the form $-|\alpha| \cdot [\alpha](t) \cdot [\beta](t)$ and $-|\beta| \cdot [\alpha](t) \cdot [\beta](t)$) and one positive (of the form $|\gamma| \cdot [\alpha](t) \cdot [\beta](t)$). Each such symmetric reaction contributes two unique terms: one negative (of the form $-|\alpha| \cdot [\alpha](t)^2$) and one positive (of the form $|\gamma| \cdot \frac{1}{2} \cdot [\alpha](t)^2$).

Then we may rewrite the above sum as

$$\begin{aligned} \frac{df}{dt} &= \sum_{\substack{R(\alpha, \beta, \gamma) \in \mathcal{R} \\ \alpha \neq \beta}} (|\gamma| \cdot [\alpha](t) \cdot [\beta](t) - |\alpha| \cdot [\alpha](t) \cdot [\beta](t) - |\beta| \cdot [\alpha](t) \cdot [\beta](t)) \\ &\quad + \sum_{R(\alpha, \alpha, \gamma) \in \mathcal{R}} \left(\frac{1}{2} |\gamma| \cdot [\alpha](t)^2 - |\alpha| \cdot [\alpha](t)^2 \right) \\ &= \sum_{\substack{R(\alpha, \beta, \gamma) \in \mathcal{R} \\ \alpha \neq \beta}} (|\gamma| - |\alpha| - |\beta|) \cdot [\alpha](t) \cdot [\beta](t) + \sum_{R(\alpha, \alpha, \gamma) \in \mathcal{R}} \left(\frac{1}{2} |\gamma| - |\alpha| \right) \cdot [\alpha](t)^2 \\ &= \sum_{\substack{R(\alpha, \beta, \gamma) \in \mathcal{R} \\ \alpha \neq \beta}} 0 \cdot [\alpha](t) \cdot [\beta](t) + \sum_{R(\alpha, \alpha, \gamma) \in \mathcal{R}} 0 \cdot [\alpha](t)^2 = 0. \quad \square \end{aligned}$$

The following is the main theorem of this paper, and it shows that hierarchical partial order systems require time linear in the diameter of a shape in order to produce it.

Theorem 4.6. *Let $\mathcal{T} = (T, \tau)$ be a hierarchical partial order system with respect to $s \in T$, with unique terminal assembly $\hat{\alpha}$ of L_1 diameter D . Then $\mathsf{T}(\mathcal{T}, s) = \Omega(D)$.*

Proof. Let $C : T \rightarrow [0, 1]$ be a concentrations function. Let $\hat{\alpha} \in \mathcal{A}_\square[\mathcal{T}]$ be the unique terminal assembly of \mathcal{T} , and let $p \in \text{dom } \hat{\alpha}$ be the unique position such that $\hat{\alpha}(p) = s$. Let \preceq be the attachment quasiorder testifying to the fact that \mathcal{T} is a partial order system with respect to p . Let \prec be the strict partial order induced by \preceq . Let $G = (V, E)$ be the directed acyclic graph induced by \prec . Assign weights to the edges of E by $w(\alpha_i, \alpha_j) = |\alpha_j|$.

If $P' = (\alpha'_1, \dots, \alpha'_l)$ is any path in G , define the (weighted) length of P' to be $w(P') = \sum_{i=1}^{l-1} w(\alpha'_i, \alpha'_{i+1})$. Let $q, r \in \text{dom } \hat{\alpha}$ be two points at L_1 distance D , which must exist since the diameter of $\text{dom } \hat{\alpha}$ is D . The distance d from p to one of these points — without loss of generality, call it q — in $\text{dom } \hat{\alpha}$ is at least $D/2$ by the triangle inequality. Let $\alpha_n \in V$ be such that $q \in \text{dom } \alpha_n$, and let $P = (\alpha_1, \dots, \alpha_n)$ be any path in G starting with α_1 and ending with α_n , where α_1 is the assembly consisting just of the tile at position p . The weight of each edge in P is an

upper bound on the diameter of the assembly it represents, and by the triangle inequality, the sum of these diameters for all α_i in P is itself an upper bound on the distance from p to q . Therefore $w(P) \geq D/2$.

Let \mathbf{T}_P be the random variable representing the time taken for the complete path P to assemble. Since the tiles on P represent a subassembly of $\hat{\alpha}$, $\hat{\alpha}$ cannot completely form until the path P forms. Therefore $\mathbf{T}_P \leq \mathbf{T}_{\mathcal{T},C,p}$. Since $\mathbb{T}(\mathcal{T}, C, p) = \mathbb{E}[\mathbf{T}_{\mathcal{T},C,p}]$, it suffices to show that $\mathbb{E}[\mathbf{T}_P] = \Omega(w(P))$.

Because of the precedence relationship described by \prec , no portion of the path P can form until its immediate predecessor on P is present. After some amount of time, some prefix P' of the path P has assembled (possibly with some other portions of $\hat{\alpha}$ not on the path P). Given $t \in \mathbb{R}_{\geq 0}$, let $\mathbf{L}(t)$ be the random variable indicating the weighted length of this prefix after t units of time.

We claim that for all $t \in \mathbb{R}_{\geq 0}$, $\mathbb{E}[\mathbf{L}(t)] \leq t$; this claim is proven below. Assuming the claim, by Markov's inequality, $\Pr[\mathbf{L}(t) \geq 2t] \leq \frac{1}{2}$. Letting $t = w(P)/2$, the event $\mathbf{L}(w(P)/2) \geq w(P)$ is equivalent to the event $\mathbf{T}_P \leq w(P)/2$. Thus $\Pr[\mathbf{T}_P \leq w(P)/2] \leq \frac{1}{2}$. By Markov's inequality, $\mathbb{E}[\mathbf{T}_P] \geq w(P)/4 \geq D/8 = \Omega(D)$, which proves the theorem, assuming that $\mathbb{E}[\mathbf{L}(t)] \leq t$.

The remainder of the proof shows the claim that for all $t \in \mathbb{R}_{\geq 0}$, $\mathbb{E}[\mathbf{L}(t)] \leq t$. Define the function $f : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ for all $t \in \mathbb{R}_{\geq 0}$ by $f(t) = \mathbb{E}[\mathbf{L}(t)]$, noting that $f(0) = 0$. Let $f' = \frac{df}{dt}$. Let $P' = (\alpha_1, \dots, \alpha_m)$ be the prefix of P formed after t seconds. Let $\beta_1, \beta_2, \dots, \beta_k$, with $m+k = n$, be the individual subassemblies remaining on the path, in order, so that $P = (\alpha_1, \dots, \alpha_m, \beta_1, \dots, \beta_k)$. For all $1 \leq i \leq k$, let $\gamma_i = \bigcup_{j=1}^i \beta_j$ be the union of the next i such subassemblies on the path (representing each of the amounts by which P could grow in the next attachment event). Let $s_i = |\gamma_i|$ be the size of the i^{th} subassembly, and let $c_i(t) = \sum_{\alpha \in A_i(t)} [\alpha](t)$, where $A_i(t)$ is the set of subassemblies (possibly containing tiles not on the path P) at time t that contain γ_i but do not contain γ_{i+1} . $A_i(t)$ represents the set of all assemblies that could attach to grow P by exactly the tiles in γ_i .

However, the set of *reactions* that could grow P is what matters, but Corollary 4.4 implies that no assembly extending P' could attach via two different reactions that both intersect P at the location directly succeeding P' . Thus, summing over assemblies is equivalent to summing over reactions. Our argument uses the conservation of mass property (Lemma 4.5) to show that no matter the concentration of assemblies in each $A_i(t)$, the rate of growth is at most one tile per unit of time.

For each $1 \leq i \leq k$, in the next instant dt , with probability $c_i dt$ the prefix will extend by total weighted length s_i by attachment of (a superassembly containing) γ_i . This implies that $f'(t) \leq \sum_{i=1}^k c_i(t) \cdot s_i$. Invoking Lemma 4.5, it follows that for all $t \in \mathbb{R}_{\geq 0}$,

$$f'(t) \leq \sum_{i=1}^k c_i(t) \cdot s_i \leq \sum_{\alpha \in \mathcal{A}[\mathcal{T}]} [\alpha](t) \cdot |\alpha| = \sum_{r \in \mathcal{T}} C(r) \leq 1.$$

Since $f(0) = 0$, this implies that $f(t) \leq t$ for all $t \in \mathbb{R}_{\geq 0}$, which completes the proof of the claim that $\mathbb{E}[\mathbf{L}(t)] \leq t$. \square

Although our assembly time model describes concentrations as evolving according to the standard mass-action kinetic differential equations, Lemma 4.5 is the only property of this model that is required for our proof of Theorem 4.6. Even if concentrations of attachable assemblies (and thus their associated attachment rates in the Markov process defining assembly time) could be magically adjusted throughout the assembly process so as to minimize the assembly time, so long as the concentrations obey Lemma 4.5 at all times, Theorem 4.6 still holds.

For example, staged assembly [16] is a relaxation of the mass-action model that obeys Lemma 4.5, in which certain assemblies are artificially prevented from interacting by being kept in separate bins before being mixed. Theorem 4.6 implies that staged assembly gives no time speedup on partial order systems if the completion time in each bin is taken into account in measuring the time complexity.

5 Assembly of a shape in time sublinear in its diameter

This section is devoted to proving the following theorem, which shows that relaxing the partial order assumption on hierarchical tile systems allows for assembly time sublinear in the diameter of the shape.

Theorem 5.1. *For infinitely many $n \in \mathbb{N}$, there is a (non-directed) hierarchical TAS $\mathcal{T} = (T, 2)$ that strictly self-assembles an $n \times n'$ rectangle, where $n' = o(n)$, such that $|T| = O(\log n)$ and there is a tile type $s \in T$ such that $\Upsilon(\mathcal{T}, s) = O(n^{4/5} \log n)$.*

In our proof, $n' \approx n^{3/5}$, but we care only that $n' \leq n$ so that the diameter of the shape is $\Theta(n)$. As discussed in Section 3, we interpret the upper bound of Theorem 5.1 more cautiously than the lower bound of Theorem 4.6, since some of our simplifying assumptions concerning diffusion rates and binding strength thresholds, discussed in Section 7, may cause the assembly time to appear artificially faster in our model than in reality. A reasonable conclusion to be drawn from Theorem 5.1 is that concentration arguments alone do not suffice to show a linear-time lower bound on assembly time in the hierarchical model.

Although we use mass-action kinetics to model changing concentrations, we occasionally use discrete language to describe the intuition behind reactions – e.g., “a copy of A is consumed and two copies of B are produced” – despite the fact that concentrations model continuously evolving real-valued concentrations.

The hierarchical model will permit a speed-up over the seeded model. However, when viewed as a *programming language* for tile assembly, the hierarchical model is more unwieldy to program and to analyze. Therefore we prove a number of lemmas showing that careful design of hierarchical tiles will cause them to “behave enough like” seeded tiles to remain tractable for analysis, and to ensure that the assembly proceeds sufficiently quickly. Much like parallel programming, in which critical regions are segregated into a few well-characterized parts of the program, we largely employ “seeded-like assembly” for most subcomponents of the construction, combining them using hierarchical parallelism at a small number of well-understood points.

5.1 Warm-up: A thin bar

We first “warm up” by analyzing in detail the assembly time of a simple but non-trivial system. The following lemma, Lemma 5.2 (more precisely, its corollary, Corollary 5.3 that assigns concrete concentrations to the tile types), shows that it is possible to grow a “substantial” concentration of a “hard-coded thin bar” in quadratic time under the mass-action model. This structure will be the first subassembly formed in many other subcomponents of the tile system of Theorem 5.1. Furthermore, the proof of Lemma 5.2 will illustrate several techniques for analyzing hierarchical assembly time (and “programming tricks” to ensure that this time is fast). Section 5.2 generalizes these techniques to apply to more complex tile systems used in the full construction. However,

these techniques are easier to understand by first reading this section with its simple, concrete tile system.

To achieve quadratic time we use “polyomino-safe” tiles that grow a $2 \times n$ bar in a zig-zag fashion to enforce that no substantial growth nucleates except at the “seed”, similar to the zig-zag tile set described by Schulman and Winfree [46, 47] (which prevented spurious nucleation with high probability under the more permissive *kinetic tile assembly model* [56] that allows reversible attachments and strength-1 attachments). This enforces that no large overlapping subassemblies grow that would compete to consume tiles without being able to attach to each other (which happens with the n tile types required to grow a $1 \times n$ bar).²⁴

Lemma 5.2. *Let $\mathcal{T} = (T, 2)$ be the hierarchical TAS shown in Figure 1, and let $\hat{\alpha}$ be its unique terminal assembly of a $2 \times n$ bar. Let the initial concentrations be defined by $[s_1](0) = \delta > 0$, $[r_i](0) = 2\delta$ for all $i \in \{1, \dots, n\}$, and $[s_i](0) = 3\delta$ for all $i \in \{2, \dots, n\}$. Then for all $t \geq \frac{4n}{\delta}$, $[\hat{\alpha}](t) \geq \frac{\delta}{2}$.*

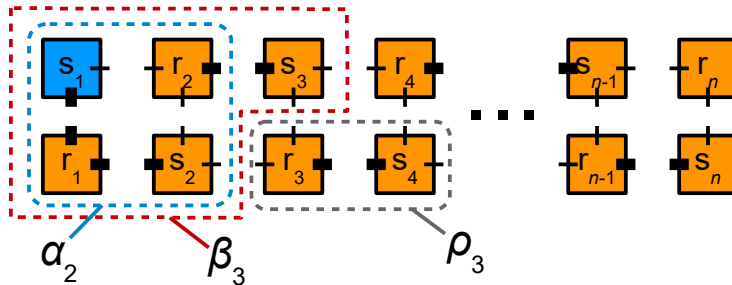


Figure 1: Tile types that grow a $2 \times n$ bar. Thick lines represent strength-2 glues, and thin lines represent strength-1 glues. Glues are not labeled, but each is hard-coded to represent its position in the final assembly. Zig-zag growth enforces that no overlapping subassemblies larger than size 2 can occur. All growth other than size-2 dimers (which will attach just as correctly as if they had stayed monomers) must nucleate from the “seed” labeled s_1 . Examples are shown of subassemblies α_i , β_i , and ρ_i , as defined in the proof of Lemma 5.2.

Intuitively, the reason for the choice of concentrations is to approximate the speed of seeded single-tile addition assembly, by enforcing that the concentrations of individual tiles (or dimers) other than s_1 remain for all time above at least a fixed constant δ , to keep high their rate of reaction with a larger assembly containing “preceding” tile types.

Proof. For $i \in \{1, \dots, n\}$, there are two types of producible assemblies containing s_1 : the assembly with exactly $2i$ tiles, which we call α_i (its single frontier location is where s_{i+1} binds), and the assembly with exactly $2i - 1$ tiles, which we call β_i (its single frontier location is where either r_i or ρ_i can bind, where ρ_i is the assembly consisting of r_i and s_{i+1}). Hence β_1 contains only s_1 , and $\alpha_n = \hat{\alpha}$.

For all $i \in \{1, \dots, n - 1\}$, let ρ_i denote the dimer (2-tile assembly) consisting of just r_i and s_{i+1} . Since β_i contains s_1 , for all $t \in \mathbb{R}_{\geq 0}$, $[\beta_i](t) \leq [s_1](0) = \delta$. Thus at most δ of the individual r_i ’s can

²⁴Assembling a $1 \times n$ bar provably requires $\Theta(n \log n)$ time to reach half of its steady-state concentration [4], even when all n tile types are allowed to have concentration as high as 1, exceeding the bound of the finite density constraint. Enforcing the finite density constraint and assigning each tile type a concentration of $1/n$ gives a time bound of $\Theta(n^2 \log n)$.

bind to β_i . The remainder must stay unbound or bind to s_{i+1} to form ρ_i . For all $i \in \{1, \dots, n-1\}$ and all $t \in \mathbb{R}_{\geq 0}$, by the fact that $[r_i](0) = 2\delta$,

$$[\rho_i](t) + [r_i](t) \geq \delta, \quad (5.1)$$

Since there is no ρ_n , we have

$$[r_n](t) \geq \delta \quad (5.2)$$

for all $t \in \mathbb{R}_{\geq 0}$ by the same reasoning. By similar reasoning, for $i \in \{2, \dots, n\}$, since $[s_i](0) = 3\delta$, no more than δ of the individual s_i 's can bind to α_{i-1} to form β_i , and no more than δ of the remaining can bind to copies of r_{i-1} that never attach to β_{i-1} . Thus for all $i \in \{2, \dots, n\}$ and all $t \in \mathbb{R}_{\geq 0}$,

$$[s_i](t) \geq \delta. \quad (5.3)$$

Let

$$\begin{aligned} S(t) &= \sum_{i=1}^n |\alpha_i| \cdot [\alpha_i](t) + \sum_{i=1}^n |\beta_i| \cdot [\beta_i](t) \\ &= \sum_{i=1}^n 2i[\alpha_i](t) + \sum_{i=1}^n (2i-1)[\beta_i](t). \end{aligned}$$

$S(t)$ can be thought of as the total ‘‘mass’’ of tiles that belong to an assembly containing s_1 at time t . Observe that

$$\sup_{t \in \mathbb{R}_{\geq 0}} S(t) = 2n\delta, \quad (5.4)$$

with the supremum $2n\delta$ attained only in the limit as $t \rightarrow \infty$, when all s_1 belong to terminal assembly α_n .

The reactions $r_i + s_{i+1} \rightarrow \rho_i$ do not change $S(t)$. All other reactions increase $S(t)$. Each reaction that increases $S(t)$ is of the form $\alpha_i + s_{i+1} \rightarrow \beta_{i+1}$, $\beta_i + r_i \rightarrow \alpha_i$ (each of which increases $S(t)$ by 1 per unit concentration of the product produced), or $\beta_i + \rho_i \rightarrow \beta_{i+1}$ (which increases $S(t)$ by 2 per unit of product). Therefore, summing the propensities of all these reactions, we obtain

$$\frac{dS(t)}{dt} = \sum_{i=1}^{n-1} [\alpha_i](t) \cdot [s_{i+1}](t) + \sum_{i=1}^n [\beta_i](t) \cdot [r_i](t) + \sum_{i=1}^{n-1} 2 \cdot [\beta_i](t) \cdot [\rho_i](t).$$

Note that for all $t \in \mathbb{R}_{\geq 0}$,

$$[s_1](0) = \sum_{i=1}^n ([\alpha_i](t) + [\beta_i](t)). \quad (5.5)$$

Each right-hand side term represents an assembly that could contain a copy of s_1 . Then

$$\begin{aligned}
\frac{dS(t)}{dt} &= \sum_{i=1}^{n-1} [\alpha_i](t) \cdot [s_{i+1}](t) + \sum_{i=1}^n [\beta_i](t) \cdot [r_i](t) + \sum_{i=1}^{n-1} 2 \cdot [\beta_i](t) \cdot [\rho_i](t) \\
&= [\beta_n](t) \cdot [r_n](t) + \sum_{i=1}^{n-1} [\alpha_i](t) \cdot [s_{i+1}](t) + \sum_{i=1}^{n-1} [\beta_i](t) \cdot ([r_i](t) + 2 \cdot [\rho_i](t)) \\
&\geq [\beta_n](t) \cdot [r_n](t) + \sum_{i=1}^{n-1} [\alpha_i](t) \cdot [s_{i+1}](t) + \sum_{i=1}^{n-1} [\beta_i](t) \cdot ([r_i](t) + [\rho_i](t)) \\
&\geq [\beta_n](t) \cdot \delta + \sum_{i=1}^{n-1} [\alpha_i](t) \cdot \delta + \sum_{i=1}^{n-1} [\beta_i](t) \cdot \delta \quad \text{by (5.1), (5.2), and (5.3)} \\
&= \delta \left([\beta_n](t) + \sum_{i=1}^{n-1} ([\alpha_i](t) + [\beta_i](t)) \right) \\
&= \delta([s_1](0) - [\alpha_n](t)) \quad \text{by (5.5)} \\
&= \delta^2 - \delta[\alpha_n](t). \tag{5.6}
\end{aligned}$$

Since α_n is not a reactant in any reaction, $[\alpha_n](t)$ is monotonically increasing. Thus it suffices to prove that $[\alpha_n](\frac{4n}{\delta}) > \frac{\delta}{2}$. Suppose for the sake of contradiction that $[\alpha_n](\frac{4n}{\delta}) \leq \frac{\delta}{2}$. Then $[\alpha_n](t) \leq \frac{\delta}{2}$ for all $t \leq \frac{4n}{\delta}$ by the monotonicity of $[\alpha_n](t)$. By this bound and (5.6), $\frac{dS(t)}{dt} \geq \frac{\delta^2}{2}$ for all $t \leq \frac{4n}{\delta}$. Since $S(0) = \delta$, this means that $S(\frac{4n}{\delta}) \geq \delta + \frac{\delta^2}{2} \cdot \frac{4n}{\delta} = (2n+1)\delta$, which contradicts (5.4). \square

The following corollary shows that if we pick the initial concentrations to be maximal subject to the finite density constraint and the constraints of Lemma 5.2, then quadratic time is sufficient to obtain terminal assembly concentration that is at least inversely linear. By Lemma 4.5, any producible assembly α obeys $[\alpha](t) \leq \frac{1}{|\alpha|}$ for all $t \in \mathbb{R}_{\geq 0}$, so this concentration bound is optimal to within a constant factor. The time bound is asymptotically suboptimal²⁵ but sufficient for our purposes, since we only use hard-coded thin bars that are logarithmically smaller than the final assembly; hence their contribution to the assembly time is negligible.

Corollary 5.3. *Let $\mathcal{T} = (T, 2)$ be the hierarchical TAS shown in Figure 1, and let $\hat{\alpha}$ be its unique terminal assembly of a $2 \times n$ bar. Let the initial concentrations be defined by $[s_1](0) = \delta = \frac{1}{cn}$ for some constant $c \in \mathbb{R}_{>0}$, $[r_i](0) = 2\delta$ for all $i \in \{1, \dots, n\}$, and $[s_i](0) = 3\delta$ for all $i \in \{2, \dots, n\}$.²⁶ Then for all $t \geq 4cn^2$, $[\hat{\alpha}](t) \geq \frac{1}{2cn}$.*

The next lemma is a discrete version of Corollary 5.3, which shows that selecting s_1 as the timekeeper results in quadratic-time assembly of the bar under our stochastic assembly time model.

Lemma 5.4. *Let $\mathcal{T} = (T, 2)$ be the hierarchical TAS shown in Figure 1, and define a concentrations function $C : T \rightarrow [0, 1]$ as in the statement of Corollary 5.3, setting $C(s_1) = \delta = \frac{1}{cn}$ for some*

²⁵We assign concentrations of $\Theta(\frac{1}{n})$ to obey the finite density constraint since there are $\Theta(n)$ distinct tile types. However, only $O(\sqrt{n})$ tile types are required to assemble a $2 \times n$ bar [6, Theorem 3.2]. The concentrations of these tile types could be set to $\Theta(\frac{1}{\sqrt{n}})$, lowering the half-completion time from $O(n^2)$ to $O(n^{1.5})$, if our goal in this section were to assemble a $2 \times n$ bar as quickly as possible (which it isn't). However, since we later use the length- n bar to encode n bits, necessitating that each tile type on the top row be unique, we could not use $O(\sqrt{n})$ tile types anyway.

²⁶We must choose $\delta \leq \frac{1}{5n-1}$ to obey the finite density constraint, hence $c \approx 5$.

constant $c \in \mathbb{R}_{\geq 0}$, $C(r_i) = 2\delta$ for $i \in \{1, \dots, n\}$, and $C(s_i) = 3\delta$ for $i \in \{2, \dots, n\}$. Then $\mathsf{T}(\mathcal{T}, C, s_1) \leq 2cn^2$.

Proof. For all $j \in \{1, \dots, 2n - 1\}$, let t_j denote the expected time until the assembly containing s_1 grows by at least one tile, conditioned on the event that the current assembly is size at least j but less than $2n$. By (5.1), (5.2), and (5.3) and the model of Markov process transition rates we employ to determine $\mathsf{T}(\mathcal{T}, C, s_1)$, it holds that $t_j \leq 1/\delta = cn$ for all $j \in \{1, \dots, 2n - 1\}$. By linearity of expectation, $\mathsf{T}(\mathcal{T}, C, s_1) \triangleq \mathbb{E}[\mathbf{T}_{\mathcal{T}, C, s_1}] \leq \sum_{j=1}^{2n-1} t_j \leq (2n - 1)cn < 2cn^2$. \square

5.2 General techniques for bounding assembly time

Techniques from the proofs of Lemmas 5.2 and 5.4 can be generalized in the following way to ease analysis of assembly time of well-behaved hierarchical systems. The results of this section will be our main technical tools used to bound the assembly time of the shape of Theorem 5.1. Intuitively, if the tile system is “polyomino-robust” (defined below), in the sense that the seeded and hierarchical models result in essentially the same producible assemblies and are well-behaved in other ways, then we can bound the hierarchical assembly time in terms of the size of the structure and the number of tile types needed to assemble it, if concentrations are set appropriately.

The TAS of Figure 1 has the following useful properties:

- 1) It is directed.
- 2) There is a constant q ($q = 2$ in Figure 1) such that producible assemblies not containing the “seed” s_1 are of size at most q . We term such assemblies *polyominos*. For mathematical convenience, we treat individual tile types that are not part of any polyomino as if they are polyominos of size 1, and we call larger polyominos *nontrivial polyominos*.
- 3) The set of producible assemblies containing s_1 is precisely the same in the seeded model as in the hierarchical model, and furthermore every terminal producible assembly contains s_1 . (The tile system is *polyomino-safe*, in the sense defined by Winfree [57].) In particular this implies that $\mathcal{A}_{\square}[\mathcal{T}] = \mathcal{A}_{\square}[\mathcal{T}_{s_1}]$, where $\mathcal{T}_{s_1} = (T, \sigma, \tau)$ is the seeded version of $\mathcal{T} = (T, \tau)$ with σ containing only s_1 .
- 4) The polyominos that attach to (an assembly containing) s_1 are a “total order (sub)system with respect to assemblies containing the seed”. More formally, define a *maximal* polyomino α to be a polyomino such that is not attachable to any assembly not containing s_1 . For each maximal polyomino α , there is a strict total order \prec on $\text{dom } \alpha$ such that if $p_1 \prec p_2$, then the tile at position p_1 always attaches to an assembly containing s_1 by at least the time that the tile at position p_2 attaches.²⁷
- 5) Every tile type belongs to at most one type of maximal polyomino (which may appear in multiple locations in the terminal assembly), and appears exactly once in the polyomino. Here we include maximal polyominos of size 1, which means any tile type in a polyomino does not appear outside of the polyomino. More formally, for each maximal polyomino α , $|T(\alpha)| = |\alpha|$, where $T(\alpha)$ is the set of tile types in α , and for each pair of maximal polyominos α and β ,

²⁷Each polyomino is a “chain” with a well-defined tile “closest” to s_1 . Therefore, while p_1 and p_2 may attach at the same time, because the order is strict (implying $p_2 \not\prec p_1$) it is always possible for the tile at p_1 to attach strictly sooner. In particular hierarchical growth is not *required* for assembly to proceed.

$\alpha \neq \beta \implies T(\alpha) \cap T(\beta) = \emptyset$. Given a tile type r , we write $\rho(r)$ to denote the unique maximal polyomino in which r is contained. This implies in particular that any tile type in a non-trivial polyomino appears in the terminal assembly equally often as any other tile type in the same polyomino.

Say that a tile system (possibly a subset of a larger tile system) that satisfies these properties is *polyomino-robust*.²⁸ Most useful seeded tile systems, when analyzed in the hierarchical model, tend to have these properties or are easily modified to have them. The two main tile subsystems that we analyze, shown in Figures 3 and 4, can be verified by inspection to obey these constraints.

The property of polyomino-robustness allows us to reason about the system, in certain senses, as if it were a seeded system. Properties (2) and (4), in particular, allow us to set concentrations in such a way that we may assume that the concentration of individual tiles or polyominos that can extend an intermediate assembly are always at least a certain value bounded away from 0 (δ in Lemma 5.2, and δ_1 in Lemma 5.5). The trick is that tiles “further from the seed” (under the ordering \prec) are always at least δ greater concentration than tiles “closer to the seed”, so that there will always be at least a δ excess of them in solution, no matter what combinations of partial polyominos form before attaching to the seed. Property (2) implies that we may use a bounded interval of concentrations (from δ to 3δ in Figure 1) to achieve this. Given a maximal polyomino α and a tile type r in α , define $\text{dist}_\alpha(r)$ to be the distance of r from the minimal position (under \prec) in the polyomino. (Since the polyomino is a linear chain, this number is well-defined). In Figure 1, for example, $\text{dist}_{\rho_i}(0,0) = 0$ and $\text{dist}_{\rho_i}(1,0) = 1$, where the polyomino ρ_i is defined as in the proof of Lemma 5.2: a size-2 polyomino containing r_i at position $(0,0)$ (within the polyomino, assuming it is translated to the lower-left corner of the first quadrant) and s_i at position $(1,0)$.

Given an assembly β and a tile type $r \in T$, define $\#\beta(r) = |\{p \in \text{dom } \beta \mid \beta(p) = r\}|$ as the number of times r appears in β . If β, ζ are assemblies such that $\beta \sqsubseteq \zeta$, define $\zeta \setminus \beta$ to be the unique assembly γ such that $\text{dom } \gamma = \text{dom } \zeta \setminus \text{dom } \beta$ and $\gamma \sqsubseteq \zeta$. If β is an assembly and T is a tile set, define $T(\beta) = \text{range } \beta = \{r \in T \mid (\exists p \in \text{dom } \beta) \beta(p) = r\}$ to be the set of tile types in β . For any polyomino-robust TAS $\mathcal{T} = (T, \tau)$ and $s_1 \in T$, let $\mathcal{T}_{s_1} = (T, \sigma, \tau)$ denote the seeded version of the hierarchical system \mathcal{T} , where σ is the single-tile initial assembly consisting of only the tile s_1 . For any producible (in the seeded model) assembly $\beta \in \mathcal{A}[\mathcal{T}_{s_1}]$, let $\mathcal{A}[\beta] = \{\alpha \in \mathcal{A}[\mathcal{T}_{s_1}] \mid \beta \sqsubseteq \alpha\}$ denote the set of producible (in the seeded model) superassemblies of β , and let $[\mathcal{A}[\beta]](t) = \sum_{\alpha \in \mathcal{A}[\beta]} [\alpha](t) \cdot \#\alpha(\beta)$, where $\#\alpha(\beta)$ denotes the number of times that β appears as a subassembly of α . That is, $[\mathcal{A}[\beta]](t)$ is the total concentration of β or of assemblies containing β , where each duplicate appearance of β in a single superassembly contributes to the concentration separately.²⁹ Note that, since assemblies can attach but not detach and $[\mathcal{A}[\beta]](t)$ takes into account not only the assembly β but any superassembly of it, $[\mathcal{A}[\beta]](t)$ is monotonically nondecreasing with t : assemblies can attach to create new copies of β , but once formed β cannot be broken apart.

The next lemma shows conditions under which a partial assembly β grows into a superassembly $\zeta \sqsupseteq \beta$. Informally, the lemma says that if we have “substantial” (at least δ_0) concentration of β

²⁸Our full tile system assembling a rectangle is not directed, hence it does not satisfy Property 1. However, we will apply the lemmas proven in this section to subsets of the full tile system that *are* directed, and in fact that satisfy all of the properties of polyomino-robustness.

²⁹For this definition to make sense, we must weight the sum by the number of times β appears in α because each time an assembly containing β binds to another assembly containing β , the number of assemblies containing β decreases by one, even though the total concentration of “completed β ’s” has stayed the same. However, whenever we actually apply this definition, it will be the case that $\#\alpha(\beta) = 1$ for any producible α such that $\beta \sqsubseteq \alpha$.

(and its superassemblies), but the total concentration of seeds is “small” (at most δ_{s_1}) compared to tile types that assemble an extension γ of β (to create a superassembly of β called ζ), and if the concentrations of those tile types are set to ensure that all individual tile types have “excess” concentration (at least δ_1), then a “substantial” concentration of ζ will assemble in time linear in $|\gamma|$.

Lemma 5.5. *Let $\mathcal{T} = (T, \tau)$ be a polyomino-robust hierarchical TAS with unique terminal assembly $\hat{\alpha}$. Let $s_1 \in T$. Let $\beta, \zeta \in \mathcal{A}[\mathcal{T}_{s_1}]$ such that $\beta \sqsubseteq \zeta$. Let $\gamma = \zeta \setminus \beta$, and suppose that $T(\gamma) \cap T(\hat{\alpha} \setminus \gamma) = \emptyset$. (i.e., tile types within γ appear only within γ), and that all polyominos contained in ζ are completely contained in β or completely contained in γ . Suppose also that for all $\alpha \in \mathcal{A}[\beta]$, $\#_\alpha(\beta) = 1$, and for all $\alpha \in \mathcal{A}[\zeta]$, $\#_\alpha(\zeta) = 1$. Suppose that there exist $t_0, c, \delta_0, \delta_{s_1} \in \mathbb{R}_{\geq 0}$ such that the following hold.*

- $[s_1](0) \leq \delta_{s_1}$.
- $[\mathcal{A}[\beta]](t_0) \geq \delta_0$.
- $\delta_{s_1} \leq \delta_0 c$.

Set initial concentrations of all $r \in T(\gamma)$ as follows. Let $\delta_1 > 0$.³⁰ Set $[r](0) = \delta_{s_1} \cdot \#_\gamma(r) + (\text{dist}_{\rho(r)}(r) + 1) \cdot \delta_1$.

Then for all $t \geq \frac{2|\gamma|c}{\delta_1} + t_0$, $[\mathcal{A}[\zeta]](t) \geq \frac{\delta_0}{2}$.

Proof. To prove the lemma, we will first argue that for every producible assembly containing β and containing a frontier location within γ , the total concentration of producible assemblies that could attach to this location is at least δ_1 . This provides a lower bound on the rate of reactions that grow the assembly into ζ (or a superassembly of ζ).

Let $\rho \sqsubseteq \gamma$ be a maximal polyomino consisting of tile types r_1, r_2, \dots, r_k at positions p_1, \dots, p_k , with $p_1 \prec p_2 \prec \dots \prec p_k$. Then we have an increasing sequence of polyominos $\rho_1 \sqsubseteq \rho_2 \sqsubseteq \dots \sqsubseteq \rho_k = \rho$, where each ρ_i is ρ_{i-1} with the tile r_i added at position p_i .

Let $j \in \{1, \dots, k\}$, and consider the polyominoes $\rho_1^j \sqsubseteq \rho_2^j \sqsubseteq \dots \sqsubseteq \rho_{k-j+1}^j \in \mathcal{A}[\mathcal{T}]$, with $|\rho_i^j| = |\rho_{i-1}^j| + 1$, that all have r_j at their minimal position under \prec . In other words, ρ_i^j contains exactly the tiles r_j, \dots, r_{j+i-1} . In particular, for all $i \in \{1, \dots, k\}$, $\rho_i = \rho_i^1$. All of these are attachable to some producible assembly containing s_1 via the tile type r_j . We allow $k = 1$ so that ρ_1^1 can also represent any individual tile type $r \in T(\gamma)$ that is not part of a nontrivial polyomino.

Call such a sequence $(\rho_1^j, \dots, \rho_{k-j+1}^j)$ a *polyomino attachment class*. Each such polyomino ρ_i^j is consumed in a reaction in one of only two ways:

1. in a reaction that binds r_j to r_{j-1} (which have concentration at most $\delta_{s_1} \cdot \#_\gamma(r_{j-1}) + \text{dist}_\rho(r_{j-1}) \cdot \delta_1 = \delta_{s_1} \cdot \#_\gamma(r_{j-1}) + (j-1) \cdot \delta_1$, which equals $\delta_{s_1} \cdot \#_\gamma(r_j) + (j-1) \cdot \delta_1$ by our assumption of equal counts of tiles that are part of the same polyomino, which is less than $[r_j](0)$ by δ_1), or
2. in a reaction producing another member of the same polyomino attachment class (thus not altering the sum of (5.7), just shifting its terms). This corresponds to the attachment of tiles “further from β under \prec ”.

³⁰Think of $\delta_1 = \frac{1}{|\mathcal{T}(\gamma)|}$ for common usage of the lemma. Intuitively, by our choice to concentrations, δ_1 excess of each tile type is ensured even after they have been maximally consumed in attachment events, meaning we can think of δ_1 as a lower bound on the concentration of tile types in the seeded model.

Therefore, there will always be at least an excess of δ_1 concentration of polyominoes in the attachment class $(\rho_1^j, \dots, \rho_{k-j+1}^j)$, i.e., we have the following for all $t \in \mathbb{R}_{\geq 0}$:

$$\sum_{i=1}^{k-j+1} [\rho_i^j](t) \geq \delta_1. \quad (5.7)$$

In particular, for any producible assembly α containing β , with frontier location $p \in \text{dom } \gamma$ at which some tile type r_j can attach in the seeded model, (5.7) ensures that the total concentration of polyominoes that can attach to position p in the hierarchical model is at least δ_1 .

Now we must argue that this lower bound on polyomino concentration implies that the claimed lower bound on concentration of assemblies containing ζ (i.e., $[\mathcal{A}[\zeta]]$). The next inequality we derive (inequality (5.8)) places a lower bound on the concentration of the *other reactant* that reacts with the polyomino (the other reactant could be any assembly in the sum on the left side of (5.8)) of the reactions that increase $[\mathcal{A}[\zeta]]$.

Recall that for all $\alpha \in \mathcal{A}[\beta]$, $\#_\alpha(\beta) = 1$ and for all $\alpha \in \mathcal{A}[\zeta]$, $\#_\alpha(\zeta) = 1$. Then for all $t \in \mathbb{R}_{\geq 0}$, $[\mathcal{A}[\beta]](t) = \sum_{\alpha \in \mathcal{A}[\beta]} [\alpha](t)$ and $[\mathcal{A}[\zeta]](t) = \sum_{\alpha \in \mathcal{A}[\zeta]} [\alpha](t)$. Since $[\mathcal{A}[\beta]](t)$ is monotonically increasing, for all $t \geq t_0$,

$$\begin{aligned} [\mathcal{A}[\beta]](t_0) &= \sum_{\alpha \in \mathcal{A}[\beta]} [\alpha](t_0) \\ &\leq \sum_{\alpha \in \mathcal{A}[\beta]} [\alpha](t) \\ &= \sum_{\alpha \in \mathcal{A}[\zeta]} [\alpha](t) + \sum_{\alpha \in \mathcal{A}[\beta] \setminus \mathcal{A}[\zeta]} [\alpha](t) \\ &= [\mathcal{A}[\zeta]](t) + \sum_{\alpha \in \mathcal{A}[\beta] \setminus \mathcal{A}[\zeta]} [\alpha](t). \end{aligned}$$

Recall that the hypothesis of the lemma supposes that $[\mathcal{A}[\beta]](t_0) \geq \delta_0$. Combined with the previous inequality, this gives us for all $t \geq t_0$,

$$\sum_{\alpha \in \mathcal{A}[\beta] \setminus \mathcal{A}[\zeta]} [\alpha](t) \geq \delta_0 - [\mathcal{A}[\zeta]](t). \quad (5.8)$$

For all $\alpha \in \mathcal{A}[\mathcal{T}_{s_1}]$, define $|\alpha|_\gamma = |\text{dom } \alpha \cap \text{dom } \gamma|$ to be the number of tiles in α that are part of γ . For all $t \in \mathbb{R}_{\geq 0}$, define

$$S_\beta^\gamma(t) = \sum_{\alpha \in \mathcal{A}[\beta]} |\alpha|_\gamma \cdot [\alpha](t).$$

Think of $S_\beta^\gamma(t)$ as the total “mass within γ ” (concentration \cdot (size within $\text{dom } \gamma$)) of assemblies that contain β , noting that $|\alpha|_\gamma \leq |\gamma|$ for all $\alpha \in \mathcal{A}[\mathcal{T}_{s_1}]$. Since β contains s_1 , this implies that $[\mathcal{A}[\beta]](t) \leq \delta_{s_1}$ for all $t \in \mathbb{R}_{\geq 0}$ by the bound on $[s_1](0)$. $S_\beta^\gamma(t)$ is maximized when all seed tiles have been incorporated into ζ or one of its superassemblies, which implies that

$$\max_{t \in \mathbb{R}_{\geq 0}} S_\beta^\gamma(t) = \max_{t \in \mathbb{R}_{\geq 0}} \sum_{\alpha \in \mathcal{A}[\beta]} |\alpha|_\gamma \cdot [\alpha](t) \leq |\gamma| \max_{t \in \mathbb{R}_{\geq 0}} \sum_{\alpha \in \mathcal{A}[\beta]} [\alpha](t) = |\gamma| \max_{t \in \mathbb{R}_{\geq 0}} [\mathcal{A}[\beta]](t) \leq \delta_{s_1} |\gamma|. \quad (5.9)$$

Since every incomplete (with respect to completing ζ) assembly $\alpha \in \mathcal{A}[\beta] \setminus \mathcal{A}[\zeta]$ that contains β can react with at least one polyomino attachment class as defined in (5.7) (in the inequality below,

let $(\rho_1^{j,\alpha}, \dots, \rho_{k-j+1}^{j,\alpha})$ be the polyomino attachment class associated to α , increasing the number of tiles occupying $\text{dom } \gamma$ by at least 1, it follows that for all $t \geq t_0$,

$$\begin{aligned} \frac{dS_\beta^\gamma(t)}{dt} &\geq \sum_{\alpha \in \mathcal{A}[\beta] \setminus \mathcal{A}[\zeta]} \left(\sum_{i=1}^{k-j+1} [\rho_i^{j,\alpha}](t) \right) \cdot [\alpha](t) \\ &\geq \delta_1 \left(\sum_{\alpha \in \mathcal{A}[\beta] \setminus \mathcal{A}[\zeta]} [\alpha](t) \right) \quad \text{by (5.7)} \\ &\geq \delta_1 \delta_0 - \delta_1 [\mathcal{A}[\zeta]](t). \quad \text{by (5.8)} \end{aligned} \tag{5.10}$$

The first inequality above is an equality if there is only a single frontier location of α within $\text{dom } \gamma$, to which any polyomino in $(\rho_1^{j,\alpha}, \dots, \rho_{k-j+1}^{j,\alpha})$ could attach, but if there are multiple frontier locations in $\text{dom } \alpha$, then $\frac{dS_\beta^\gamma(t)}{dt}$ would be even larger.

It suffices to show that for $t^* = \frac{2|\gamma|c}{\delta_1} + t_0$, $[\mathcal{A}[\zeta]](t^*) \geq \frac{\delta_0}{2}$. (By the monotonicity of $[\mathcal{A}[\zeta]](t)$, this will imply $[\mathcal{A}[\zeta]](t) \geq \frac{\delta_0}{2}$ for all $t \geq t^*$ as well.) Suppose for the sake of contradiction that $[\mathcal{A}[\zeta]](t^*) < \frac{\delta_0}{2}$. Then by (5.10), for all $t \in [t_0, t^*]$, $\frac{dS_\beta^\gamma(t)}{dt} \geq \delta_1 \delta_0 - \delta_1 [\mathcal{A}[\zeta]](t) > \frac{\delta_1 \delta_0}{2}$. This implies that $S_\beta^\gamma(t^*) > \frac{\delta_1 \delta_0}{2} (t^* - t_0) = \frac{\delta_1 \delta_0}{2} \frac{2|\gamma|c}{\delta_1} = \delta_0 c |\gamma| \geq \delta_{s_1} |\gamma|$, which contradicts (5.9). \square

Intuitively, Lemma 5.5 shows that one “stage” of assembly (the stage that goes from β to ζ) is “fast.” The following lemma extends the analysis of Lemma 5.5 to the case where we want to analyze multiple stages of assembly (p is the number of stages in the statement of Lemma 5.6), where each stage may involve the addition of tile types of asymptotically different concentrations than the other stages. Intuitively, this is required to prove fast assembly time because some stages (such as hard-coding the seed row of a counter with $\omega(1)$ tile types) proceed slowly relative to their size since they require many tile types (hence each tile type has low concentration). This is not a problem since the size of such stages is small, but it implies that we cannot apply the “average attachment time per tile” of such a slow stage uniformly across the entire assembly. Some stages (such as completing a counter with a complete seed row with $O(1)$ tile types) proceed very quickly, and such stages account for “most” of the size of the final assembly, so that the total assembly speed is fast despite a few “bottleneck stages” in which the assembly process slows down for a short time. The statement of the lemma is quite intricate, but the proof is simple, setting up each stage of growth to match the hypothesis of Lemma 5.5 and then applying that lemma iteratively.

Lemma 5.6. *Let $\mathcal{T} = (T, \tau)$ be a polyomino-robust hierarchical TAS with unique terminal assembly $\hat{\alpha}$. Let $\beta_1, \beta_2, \dots, \beta_p \in \mathcal{A}[\mathcal{T}_{s_1}]$ such that $\beta_{i-1} \sqsubseteq \beta_i$ for all $i \leq p$, where β_1 is the assembly with a single seed s_1 . Let $\gamma_i = \beta_{i+1} \setminus \beta_i$ for each $i < p$, and supposed that $T(\gamma_i) \cap T(\gamma_j) = \emptyset$ for all $i, j < p$, and that all polyominos contained in β_p are completely contained in γ_i for some i . (i.e., suppose that each β_p and β_{p+1} satisfy the hypothesis of Lemma 5.5, where β_{p+1} in this lemma is interpreted to be ζ in Lemma 5.5.) Suppose also that $\delta_{s_1} \leq \frac{1}{|\beta_p|}$. Set initial concentrations of all $r \in T(\gamma_i)$, where $i \leq p$ as $[r](0) = \delta_{s_1} \cdot \#\gamma_i(r) + (\text{dist}_{\rho(r)}(r) + 1) \cdot \frac{1}{|T(\gamma_i)|}$.*

Then for all $t \geq \sum_{i=1}^{p-1} 2^i |\gamma_i| |T(\gamma_i)|$, $[\mathcal{A}[\beta_p]](t) \geq \frac{\delta_{s_1}}{2^{p-1}}$. Furthermore, $\sum_{r \in T} [r](0) \leq 1 + p(q+1)$, where q is the maximum size of any polyomino as in the definition of polyomino-robustness.

Proof. We will show the statement $[\mathcal{A}[\beta_m]](t) \geq \frac{\delta_{s_1}}{2^{m-1}}$ for all $t \geq \sum_{i=1}^{m-1} 2^i |\gamma_i| |T(\gamma_i)|$ is true for all $m \leq p$ by induction on m .

This statement is trivially true for $m = 1$ since $[\mathcal{A}[\beta_1]](t) = [\mathcal{A}[s_1]](t) = \delta_{s_1}$ for all t .

Assume that the statement is true for $m = m_0$. In other words, at time $t = \sum_{i=1}^{m_0-1} 2^i |\gamma_i| |T(\gamma_i)|$, $[\mathcal{A}[\beta_{m_0}]](t) \geq \frac{\delta_{s_1}}{2^{m_0-1}}$. Invoking Lemma 5.5 with $\beta = \beta_{m_0}$, $\gamma = \gamma_{m_0}$, $t_0 = \sum_{i=1}^{m_0-1} 2^i |\gamma_i| |T(\gamma_i)|$, $c = 2^{m_0-1}$, and $\delta_1 = \frac{1}{|T(\gamma_i)|}$, we know that $[\mathcal{A}[\beta_{m_0+1}]](t) \geq \frac{\delta_{s_1}}{2^{m_0}}$ at $t = t_0 + 2^{m_0} |\gamma_{m_0}| |T(\gamma_{m_0})| = \sum_{i=1}^{m_0} 2^i |\gamma_i| |T(\gamma_i)|$. Therefore, the statement is true for $m = m_0 + 1$.

The total concentration of tiles whose type is in $T(\beta_p)$ is

$$\begin{aligned} \sum_{r \in T(\beta_p)} [r](0) &= \sum_{i=1}^p \sum_{r \in T(\gamma_i)} [r](0) \\ &= \sum_{i=1}^p \left(\delta_{s_1} \sum_{r \in T(\gamma_i)} \#\gamma_i(r) + \sum_{r \in T(\gamma_i)} (\text{dist}_{\rho(r)}(r) + 1) \cdot \frac{1}{|T(\gamma_i)|} \right) \\ &\leq \sum_{i=1}^p \delta_{s_1} |\gamma_i| + \sum_{i=1}^p \max_{r \in T(\gamma_i)} [\text{dist}_{\rho(r)}(r) + 1] \\ &\leq \delta_{s_1} |\beta_p| + p \max_{r \in T(\gamma_i)} [\text{dist}_{\rho(r)}(r) + 1]. \end{aligned}$$

Since $\delta_{s_1} \leq \frac{1}{|\beta_p|}$, and all polyominoes have size at most q , the total initial concentration of all tile types is at most $1 + p(q + 1)$. \square

In particular, if the number of stages p and polyomino size bound q are constant with respect to the size of the terminal assembly (call this parameter n ; as in our diameter $\Theta(n)$ rectangle of Theorem 5.1), then the total initial concentration of all tile types is constant with respect to n . In particular, we can scale these concentrations to obey the finite density constraint without affecting the asymptotic time and concentration bounds derived in Lemma 5.6.

The following is a discrete version of Lemma 5.5 that can be used to analyze polyomino-robust systems in the discrete assembly time model.

Lemma 5.7. *Let $\mathcal{T} = (T, \tau)$ be as in Lemma 5.5, and let $\alpha \in \mathcal{A}[\beta]$ be the current state of the Markov process defining the assembly time of \mathcal{T} , and suppose that the current time is t_0 and that $\beta \sqsubseteq \alpha$. Define $\mathbf{T}_{\alpha, \zeta}$ to be the random variable representing the first time at which α grows into a superassembly of ζ . Define $\mathbb{T}(\alpha, \zeta) = \mathbb{E}[\mathbf{T}_{\alpha, \zeta}]$. Then $\mathbb{T}(\alpha, \zeta) \leq \frac{|\zeta|}{\delta_1} + t_0$.*

Proof. For all $j \in \{0, \dots, |\gamma|\}$, let t_j denote the expected time until the assembly adds at least one more tile to $\text{dom } \gamma$, conditioned on the event that current size of the assembly within $\text{dom } \gamma$ is j . By (5.7) and the model of Markov process transition rates we employ to determine $\mathbb{T}(\mathcal{T}, C, s_1)$, it holds that $t_j \leq 1/\delta_1$ for all $j \in \{0, \dots, |\gamma| - 1\}$. By linearity of expectation, $\mathbb{T}(\alpha, \zeta) \triangleq \mathbb{E}[\mathbf{T}_{\alpha, \zeta}] \leq t_0 + \sum_{j=0}^{|\gamma|-1} t_j \leq t_0 + \frac{|\gamma|}{\delta_1}$. \square

5.3 Construction of a fast-assembling shape

This section describes the main components of the construction of Theorem 5.1.

Figure 2 shows an overview of the assembly described in Theorem 5.1. It consists of an initial (blue) “vertical bar”, and m copies of another type of (blue) vertical bar, each adjacent pair connected by a pair of (two different types of gray) “horizontal bars”. The leftmost vertical bar forms,

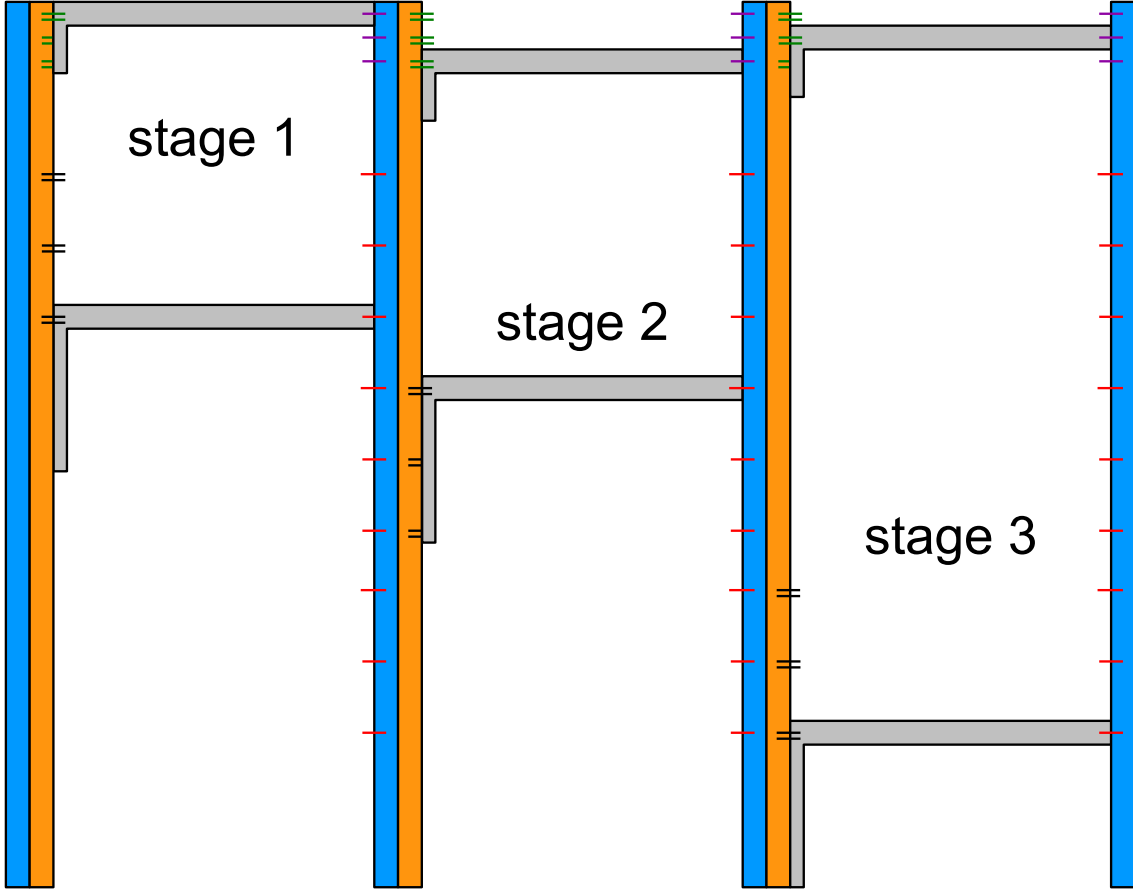


Figure 2: High-level overview of interaction of “vertical bars” and “horizontal bars” to create the rectangle of Theorem 5.1. Filler tiles fill in the empty regions as shown in Figure 6. If glues overlap two regions then represent a formed bond. If glues overlap one region but not another, they are glues from the former region but are mismatched (and thus “covered and protected”) by the latter region.

then the two horizontal bars attach, after which their right-side single strength glues cooperate to attach a new vertical bar to the right. This continues until the entire m “stages” are complete. In the meantime, filler tiles fill in the gaps to complete the rectangle. Figures 3 and 4 show some more detail of the tile types that create the horizontal and vertical bars, and Lemmas 5.8 and 5.9 respectively show that these subassemblies have “substantial” concentration “quickly enough” to be useful to prove the time bound of Theorem 5.1.

Figures 3 and 4 show details of the tile types that assemble the “horizontal bars” and “vertical bars” of Figure 2, and Lemmas 5.8 and 5.9 bound their assembly time. It may be beneficial for the reader first to skim the proof of Theorem 5.1, prior to examining Figures 3 and 4 in detail, in order to understand the intuitive purpose of the shape and outer glue placement of the horizontal and vertical bars.

For all $n \in \mathbb{Z}^+$, define $\log' n = \lfloor \log n \rfloor + 1$, the number of bits required to represent n in binary, so that $2^{\log' n}$ is the next power of 2 greater than n .




Figure 3: “Horizontal bar with a vertical arm” for the construction of a fast-assembling square. The arm (assembled by a downward-growing counter similar to the horizontal counter, and set up to grow only after the horizontal counter is complete) is intended to block other horizontal bars from binding after this one has bound to a vertical bar. The arm has height either k (for type “A” horizontal bars of Figure 5) or k^2 (for type “B” horizontal bars). Since we choose $w = k^4$, in either case the tile complexity and assembly time are dominated by the horizontal bar. The circles on the left and right indicate single-strength glues that are used when and after the bar binds hierarchically. The two west-facing single-strength glues cannot cooperate until the entire assembly is complete, so in particular the right glue must already be present. The east-facing single-strength glue is designed to cooperate with a glue from a different horizontal bar to control placement of a vertical bar, as in Figure 5. The north and south facing single-strength glues are designed to help with filler tiles or stage-counting as in Figure 6.

Lemma 5.8. *Let $\mathcal{T} = (T, 2)$ be the “horizontal bar with an arm” hierarchical TAS shown (by example) in Figure 3, and let $\hat{\alpha}$ be its unique terminal assembly of an $O(w) \times O(\log w)$ horizontal bar with an $O(\log k) \times O(k)$ “vertical arm” on its left, where w and k are as in the proof of Theorem 5.1 (so that $w \geq k$ in particular). Then there is an assignment of concentrations obeying $\sum_{r \in T} [r](0) = O(1)$ (with respect to k and w) such that for all $t \geq O(w \log w)$, $[\hat{\alpha}](t) \geq \Omega(\frac{1}{w \log w})$.*

Proof. Let $\beta_1 \sqsubseteq \beta_2 \sqsubseteq \beta_3 \sqsubseteq \beta_4 \sqsubseteq \beta_5$ respectively represent the assemblies encircled by dotted lines of Figure 3, so that β_1 is just the tile s_1 , and $\beta_5 = \hat{\alpha}$. Set the initial concentrations of tile types in T as in the statement of Lemma 5.6. Defining $\gamma_1, \dots, \gamma_5$ as in Lemma 5.6, note that $|T(\gamma_1)| = O(\log w)$, $|T(\gamma_2)| = O(1)$, $|T(\gamma_3)| = O(\log k)$, $|T(\gamma_4)| = O(1)$, $|\gamma_1| = O(\log w)$, $|\gamma_2| = O(w \log w)$, $|\gamma_3| = O(\log k)$, and $|\gamma_4| = O(k \log k)$. Therefore Lemma 5.6 tells us that for all $t \geq 2 \cdot O(\log^2 w) + 4 \cdot O(w \log w) + 8 \cdot O(\log^2 k) + 16 \cdot O(k \log k) = O(w \log w)$ (since $w \geq k$), $[\mathcal{A}[\hat{\alpha}]](t) \geq \delta_{s_1}/16$. Setting $\delta_{s_1} = \frac{1}{2w \log' w}$ satisfies $\delta_{s_1} \leq \frac{1}{|\beta_5|}$, so that the total concentration of tile types is at most $1 + p(q+1)$, where $p = 5$ and $q = 4$ in the tile system of Figure 3, where q is defined to be the maximum size of any polyomino as in the definition of polyomino-robustness, giving the required constant concentration bound. \square




Figure 4: Tiles to assemble a “vertical bar” as in Figures 2, 5, and 6 for the fast self-assembly of a rectangle. There are two types of vertical bars, the leftmost one with no glues on its west side (not shown), and the other vertical bars with glues each side (shown).

Lemma 5.9. *Let $\mathcal{T} = (T, 2)$ be the vertical bar hierarchical TAS shown (by example) in Figure 4, and let $\hat{\alpha}$ be its unique terminal assembly of an $O(mk^2) \times O(\log k)$ rectangle, where m and k are as in the proof of Theorem 5.1. $\hat{\alpha}$ places mk^2 single strength “type B” (as in Figure 5) glues on its left and right (all glues on the left identical, and all glues on the right identical to each other but different from the left glues), spaced every k vertical rows, and another k “type A” glues (as in Figure 5) spaced 2 rows apart on the top left and right. Also, there is an assignment of concentrations obeying $\sum_{r \in T} [r](0) = O(1)$ (with respect to k and w), such that for all $t \geq \Omega(mk^2 \log(mk))$, $[\hat{\alpha}](t) \geq \Omega(\frac{1}{mk^2 \log(mk)})$.*

Proof. The tiles are essentially zig-zag counters as described in [45]. Since single-strength glues must be placed in precise locations as required in Figures 2, 5, and 6, some modifications of the counter are necessary. There are three embedded counters 1, 2, and 3, counting to k , k , and m ,

respectively ($m = k$ in our construction), respectively. Counter 1 bits are shown on the left side of each tile, counter 2 bits on the right, and counter 3 bits are omitted. Counter 1 increments each row, and when it rolls over, it resets and counter 2 increments once. Similarly, when counter 2 rolls over, it and counter 1 reset and counter 3 increments. The values m and k are embedded in the first two rows and carried through each subsequent row to enable the resets. These place single strength glues on each side (the left glues to allow the vertical bar to bind to two horizontal bars as show in Figure 5, and the right glues to help the orange counter in Figure 6 to correctly place glues on the right side of the vertical bar once it has attached. These single-strength glues are the “group B ” glues of Figure 5. Finally, when counter 3 rolls over, a new counter (using new tile types) is initiated to count to $2k$, placing a glue every other row, which are the “group A ” glues of Figure 5. To ensure that the vertical arm of the bottommost horizontal bar does not protrude below the bottom of the vertical bar, it is necessary to first count k^2 rows without placing side glues, but for space reasons this is not shown. Also, the first vertical bar must count an additional $\log m$ rows, since each subsequent vertical bar will add this many rows to the bottom, as shown in Figure 6, when the stage computation counter must “crawl” below the bottom of the vertical bar in order to place the east-facing glues needed for the next-stage horizontal bars to bind to the vertical bar.

Let $\beta_1 \sqsubseteq \beta_2 \sqsubseteq \beta_3$ respectively represent the assemblies encircled by dotted lines of Figure 4, so that β_1 is just the tile s_1 , and $\beta_3 = \hat{\alpha}$. Set the initial concentrations of tile types in T as in the statement of Lemma 5.6. Defining $\gamma_1, \dots, \gamma_3$ as in Lemma 5.6, note that $|T(\gamma_1)| = O(\log k + \log m)$, $|T(\gamma_2)| = O(1)$, $|\gamma_1| = O(\log k + \log m)$, and $|\gamma_2| = O(mk^2(\log k + \log m))$. Therefore Lemma 5.6 tells us that for all $t \geq 2 \cdot O((\log k + \log m)^2) + 4 \cdot O(mk^2(\log k + \log m)) = O(mk^2 \log(mk))$, $[\mathcal{A}[\hat{\alpha}]](t) \geq \delta_{s_1}/4$. Setting $\delta_{s_1} = \frac{1}{2mk^2(\log m + 2 \log k)}$ satisfies $\delta_{s_1} \leq \frac{1}{|\beta_3|}$ and gives a total concentration of tile types at most $1 + p(q + 1)$, where $p = 3$ and $q = 3$ in the tile system of Figure 3, where q is defined to be the maximum size of any polyomino as in the definition of polyomino-robustness, giving the required constant concentration bound. \square

We now prove the main theorem of this section, Theorem 5.1.

Proof of Theorem 5.1. Based on the primitives introduced in Figures 1, 3, and 4 and their assembly time analysis, Figures 2, 5, and 6 outline the remainder of the construction.

Intuitively, the construction proceeds as follows. Fix positive integers k , m , and w to be defined later. The rectangle grows rightward in m “stages”, each stage of width w and height $h = O(mk^2)$. The speedup is obtained by using “binding parallelism”: the ability of a single (large) assembly β to bind to multiple sites on another assembly α . Think of α as the structure built so far, with a vertical bar on its right end, and think of β as one of the horizontal bars shown in Figures 5 and 2. This “binding parallelism” is in addition to the “assembly parallelism” described in Section 6: the ability for α to assemble in parallel with β so that (a large concentration of) β is ready to bind as soon as α is assembled. The number k controls the amount of “binding parallelism”: it is the number of binding sites on α to which β may bind, the *first* of which binds in expected time $\frac{1}{k}$ times that of the expected time before any *fixed* binding site binds (since the minimum of k exponential random variables of expected value t has expected value $\frac{t}{k}$). Actually two different versions of β bind to one of two different regions on α , each with k binding sites. Because assembly may proceed as soon as each of the two regions has a β bound (so that no individual binding site is required before assembly can proceed), the system is not a partial order system; in fact it is not

even directed since different filler tiles will fill in the other $k - 1$ regions where copies of β could have gone but did not.




Figure 5: “Vertical bars” for the construction of a fast-assembling square, and their interaction with horizontal bars of Figure 3. “Type B” horizontal bars have a longer vertical arm than “Type A” since the glues they must block are farther apart.

The timekeeper seed is contained in a height- h “vertical bar” with a region of k pairs of single-strength glues on its top right, and another region of k (more widely spaced) pairs of single-strength glues on its right below the first region. Within each pair of glues, the two glues are different (despite




Figure 6: Detailed view of how a single partial vertical bar attaches to complete a full vertical bar after attaching to the previous stage's two horizontal bars. This enables the correct vertical placement of the two horizontal bars in the next stage. This is needed to communicate which stage is the current one based on the position of the previous lower horizontal bar, since there is only one type of vertical bar and it doesn't "know" the current stage.

being the same color in Figures 5 and 2). However, within the first region, each pair is the same. Within the second region, each pair of glues is also identical to all other pairs in the second region, although each is distinct from the pair of glues of the first region. This vertical bar assembles as in Figure 4, but rotated 90 degrees. The placement of the glues on the right is simple to calculate from the numbers k and m , which are embedded into the tile types. The width of the vertical bar is therefore $O(\log k + \log m)$ (although most figures show the vertical bar as "thin", it is more than

one tile wide) and requires at most $O(\log k + \log m)$ tile types to encode. The leftmost vertical bar, having no west-facing glues, has distinct tile types from the other vertical bars, but all other vertical bars are identical to each other.

The horizontal bar of Figure 3 also has two types, but unlike the vertical bar, both types appear in each stage: type A to bind in the top region and type B to bind in the bottom region of the vertical bar, each with different single-strength glues on each end from the other type. The vertical “arm” on the left side below the bar is intended to prevent other horizontal bars from binding to another pair of glues in the same region of a vertical bar. An order of growth is chosen that enforces that the entire horizontal bar and the arm must be assembled before the two strength-1 glues on the left are present. This is required for speed; if a horizontal bar could bind to the right of a vertical bar before the horizontal bar is complete, then it would have to complete (taking at least time w in the worst case) before growth of the larger assembly could continue. By ensuring that only complete horizontal bars can attach to vertical bars (and ensuring that many horizontal bars will be complete by the time “most” vertical bars require them), the assembly grows rightward quickly. Since the right side of a horizontal bar has only a single strength-1 glue, two horizontal bars, one in each of regions A and B , are required to bind before the next stage’s vertical bar may attach by using cooperative binding with each of these strength-1 glues.

The following is a key idea in the construction: since there are only two kinds of horizontal bars, the glues on the right side cannot “communicate” the current stage. The natural solution to this, letting single-tile addition propagate the stage information from left to right along the horizontal bar, defeats the purpose of letting the horizontal bar attach hierarchically: such a solution would slow down the assembly process to be no faster than the seeded model. To enforce that stage $i + 1$ properly follows stage i (so that we deterministically stop after m stages), the relative *position* of the horizontal bars with respect to each other communicates the current stage. This is the reason that each stage “stagger” the vertical position of the group B region of k east-facing glues on the vertical bar, depending on the stage.

In more detail, examine the single-strength glues on the left side of the vertical bar in Figure 6. The group A glues on the top left side of the vertical bar are each 2 spaces apart (to make room for their counterpart glues that will be placed on the *right* side, which are twice as numerous since they cooperate in adjacent pairs). The group B glues are each k spaces apart. This ensures that every vertical distance between a glue from group A and a glue from group B is unique. Note that all vertical bars except the initial vertical bar have more single-strength glues on the left than on the right. These allow the vertical bars to bind at *any* stage. However, by the distance-uniqueness property just explained, there is only one vertical position in which the vertical bar can bind, so that the vertical bar will be evenly lined up with the vertical bar from all the previous stages. Furthermore, it is possible, as shown in Figure 6, to use a counter to *measure* the relative height of the bottom horizontal bar in order to determine the current stage. The counter increments once for each group of k single-strength glues that it passes (this can be implemented by marking boundaries between groups of k glues with a special glue).

As Figure 6 shows, once this counter reaches the bottom of the vertical bar and has value h to indicate that the previous stage was stage $m - h$, it uses h to determine where to place pairs of glues on the right side of the vertical bar for the next stage: they should be placed at relative height $h - 1$. This construction processing of the left side of the vertical bar to assemble the right side of the vertical bar can be done with a constant number of tile types. When the value h is calculated to be 0, no right side of the vertical bar is constructed, since we have reached the final stage. As

Figure 6 shows, filler tiles fill in the gaps above and below the horizontal bars after the vertical bar binds to the right.

We now analyze the assembly time of this construction. Although the TAS we describe is not directed, the subcomponents are deterministic, so we may apply Lemmas 5.5, 5.6 and 5.7 to them.

We choose the timekeeper seed to be the same s_1 as shown in Figure 4, for the alternate (not shown) version of the tiles that create the stage 1 vertical bar, with no glues on the left side. By our design in Figure 6, the only way a horizontal bar can attach to any other assembly is to attach to a complete vertical bar to its left that part of an assembly containing s_1 .

Let T_{h-A} , T_{h-B} , T_v , and T_{v_0} be the tile types to create horizontal type A bars, horizontal type B bars, vertical bars (partial, those colored blue in Figures 5 and 6), and leftmost vertical bar, respectively. By the construction, we have $|T_{h-A}| = O(\log w + \log k)$, $|T_{h-B}| = O(\log w + \log k)$, and $|T_v| = |T_{v_0}| = O(\log w + \log k + \log m)$.

For the general system, we add all four systems T_{h-A} , T_{h-B} , T_v , and T_{v_0} . We also add a constant number of counter tile types that wrap around the vertical bar as in Figure 6, with each tile type having concentration $\Omega(1)$ and a constant number of filler tile types also with total concentration $\Omega(1)$. In particular, assign the tile types concentrations as in the statement of Lemma 5.5, so that each tile type in the counter tiles and filler tiles (or polyomino attachment class) is guaranteed to have $\Omega(1)$ *excess* concentration by assigning extra concentration to tile types in polyominos. (Although the details are not shown, it is trivial to implement the stage counter tiles with a polyomino-robust system, and to implement the filler tiles with no polyominos at all.) This allows us to apply Lemma 5.7 to these subsystems with $\delta_1 = \Omega(1)$.

Let s_1 denote the seed of the whole tile system. Let $\delta_{s_1} = [s_1](0) = \frac{1}{2nm}$. For our choice of k and m , this will mean $\delta_{s_1} = O(\frac{1}{k^8})$. Let α_{h-A} denote the type A ‘‘horizontal bar with an arm’’ assembly. Let α_{h-B} denote the type B ‘‘horizontal bar with an arm’’ assembly. Let α_{v_0} denote the leftmost vertical bar assembly. Let α_v denote the other vertical bar assembly. By Lemma 5.8, for all $t \geq O(w \log w)$, the total concentration of type A horizontal bars that have been produced is at least $\Omega(\frac{1}{w \log w})$ (although some may be incorporated into assemblies containing the seed, and similarly for type B vertical bars). Since at most $\delta_{s_1} m$ concentration of each horizontal bar can be attached to an assembly containing s_1 , this implies that for all $t \geq O(w \log w)$, $[\alpha_{h-A}](t) \geq \Omega(\frac{1}{w \log w}) - \delta_{s_1} m$, which is $\Omega(\frac{1}{w \log w})$ by our choice of δ_{s_1} . By the same reasoning, for all $t \geq O(w \log w)$, $[\alpha_{h-B}](t) \geq \Omega(\frac{1}{w \log w})$. Similarly, by Lemma 5.9 and similar reasoning regarding the scarcity of s_1 -containing assemblies that could attach to α_v , for all $t \geq \Omega(mk^2 \log(mk))$, $[\alpha_v](t) \geq \Omega(\frac{1}{mk^2 \log(mk)})$.

Having derived these concentration bounds in the mass-action model, we now analyze the stochastic assembly time of s_1 into a terminal assembly.

By Lemma 5.7, the leftmost vertical bar’s bottom two rows assemble in expected time $O(\log k + \log m)$, using the fact that each tile type in the bottom two rows can be assigned concentration $\Omega(\frac{1}{\log k + \log m})$, to ensure that $\delta_1 = \Omega(\frac{1}{\log k + \log m})$ in Lemma 5.7. Again applying Lemma 5.7, the expected time to assemble the complete leftmost vertical bar from its seed row is at most $O(mk^2 \log(mk))$, using the fact that there are $O(1)$ tile types needed to complete the remaining rows and can be therefore be assigned concentration $\Omega(1)$ and in particular ensuring excess δ_1 at least $\Omega(1)$ in Lemma 5.7.

Once the leftmost vertical bar completes, and if the current time $t = \Omega(w \log w)$, then by our above-derived bound on the concentration of α_{h-A} , the expected time before a type A horizontal bar attaches to some binding site in the group A glues on the right side is at most $O(w \log w)/k$. This holds similarly for α_{h-B} , so the expected time before both attachments happen is also at most

$O(w \log w)/k$. Once two horizontal bars have attached, and if the current time $t = \Omega(mk^2(\log k + \log m))$, by our above-derived bound on the concentration of α_v , the expected time before the next vertical bar attaches is at most $O(mk^2(\log k + \log m))$. Once this occurs, by Lemma 5.7, the expected time before the constant-size tile set shown in Figure 6 to complete the placement of glues on the east side of the just-attached vertical bar is at most $O(mk^2(\log k + \log m))$. At this point the first stage is complete, requiring time at most $O(w \log w) + O(mk^2(\log k + \log m))$ (to wait for sufficient concentration of horizontal and vertical bars) $+ O(mk^2(\log k + \log m))$ (to grow the leftmost vertical bar) $+ O(w \log w)/k$ (to attach two horizontal bars) $+ O(mk^2(\log k + \log m))$ (to attach the second vertical bar) by linearity of expectation. Simplified, this is $O(w \log w) + O(mk^2(\log k + \log m)) + O(w \log w)/k$.

Repeating this analysis for each of the remaining stages, the total time for the complete “skeleton” of Figure 5 to complete is at most m times the previous bound, $O(w \log w) + O(m^2k^2(\log k + \log m)) + O(mw \log w)/k$, by linearity of expectation.

Finally, the filler tiles must tile the $3m$ empty regions left in the skeleton. Although this assembly process likely begins before the full skeleton is complete, we analyze it as if no filler tiles attach until the full skeleton is complete, as an upper bound for the actual assembly time. Each of these regions is a rectangle (minus the vertical arms of the horizontal bars) of diameter at most $O(w + mk^2)$, which is tiled by a constant-size rectilinear tile set (tile set in which each tile cooperates using north and east glues to grow towards the southwest), where each tile type has concentration $\Omega(1)$ at all times. By Theorem 4.4 of [3], the expected time for one of these regions to be completely tiled is at most $O(w + mk^2)$, with an exponentially decaying tail on the time distribution. Since there are $3m$ such regions assembling in parallel, and each has an exponentially decaying tail, the time for all regions to completely fill is at most $O((w + mk^2) \log m)$.

Therefore the entire expected assembly time is at most

$$O(w \log w) + O(m^2k^2(\log k + \log m)) + O(mw \log w)/k + O((w + mk^2) \log m).$$

We choose $k = m = n^{1/5}$ and $w = k^4$. Simplifying the above expression, this gives an expected assembly time of $O(n^{4/5} \log n)$. \square

By using the base-conversion technique of [3] for all counters, the number of tile types required could be reduced from $O(\log n)$ to the information-theoretically optimal $O(\frac{\log n}{\log \log n})$. However, this is a now-standard technique for obtaining optimal tile complexity of structures that “encode” a natural number n . Since the primary contribution of our construction is the bound on assembly time, we have presented a simpler (but larger than optimal) tile system for illustrative purposes. Unlike Theorem 6.1, where the problem of obtaining small assembly depth for a shape S is trivialized by allowing tile complexity $|S|$, obtaining fast assembly *time* is nontrivial whether tile complexity is large or not. In fact, small tile complexity *helps* to obtain fast assembly time, since with fewer tile types, one can distribute to each tile type a greater share of the $O(1)$ concentration allowed by the finite density constraint, which in turn reduces the expected time for each tile to attach.

Also, it is possible to shave log factors from the assembly time analysis by using the “optimal counter” tiles of [3], which grow an $n \times \log n$ counter in the seeded model in time $O(n)$, compared to the suboptimal $O(n \log n)$ time required by the zig-zag counters we use. However, our Lemma 5.5 does not take the “binding parallelism” of the seeded model into account, but instead implicitly assumes in the worst case that the frontier is always size 1. A more careful analysis could remove some of these log factors, but we have allowed the log factors in order to simplify the analysis, since our main goal is to obtain a sublinear time bound.

6 Nearly maximally parallel hierarchical assembly of a square with optimal tile complexity

In this section we show that under the hierarchical model of tile assembly, it is possible to self-assemble an $n \times n$ square, for arbitrary $n \in \mathbb{Z}^+$, using the asymptotically optimal $O(\frac{\log n}{\log \log n})$ number of tile types. Furthermore, the square assembles using nearly the maximum possible parallelism in the hierarchical model, building the final square out of four assembled sub-squares of size $n/2 \times n/2$, which are themselves each assembled from four sub-squares of size $n/4 \times n/4$, etc. The sub-optimality stems from the need for us to construct the smallest sub-squares of size $O(\log n) \times O(\log n) = O(\log^2 n)$ without parallelism.

We formalize the notion of “parallelism through hierarchical assembly” as follows.

Let $\mathcal{T} = (T, \tau)$ be a directed hierarchical TAS. Let $\alpha \in \mathcal{A}[\mathcal{T}]$ be a producible assembly. An *assembly tree* Υ of α is a full binary tree whose nodes are labeled by producible assemblies, with α labeling the root, individual tile types labeling the $|\alpha|$ leaves, and node v having children u_1 and u_2 with the requirement that u_1 and u_2 can attach to assemble v . That is, Υ represents one possible pathway through which α could be produced from individual tile types in \mathcal{T} . Let $\Upsilon(\mathcal{T})$ denote the set of all assembly trees of \mathcal{T} . Say that an assembly tree is *terminal* if its root is a terminal assembly. Let $\Upsilon_{\square}(\mathcal{T})$ denote the set of all terminal assembly trees of \mathcal{T} . Note that even a directed hierarchical TAS can have multiple terminal assembly trees that all have the same root terminal assembly. The *assembly depth* of \mathcal{T} is $\text{depth}^{\text{da}}(\mathcal{T}) = \max_{\Upsilon \in \Upsilon_{\square}(\mathcal{T})} \text{depth}(\Upsilon)$, where $\text{depth}(\Upsilon)$ denotes the standard depth of the tree Υ , the length of the longest path from any leaf to the root.

It is clear by the definition that for any shape S with N points strictly self-assembled by a tile system \mathcal{T} , $\text{depth}^{\text{da}}(\mathcal{T}) \geq \log N$. Our construction achieves $\text{depth}^{\text{da}}(\mathcal{T}) \leq O(\log^2 n)$ in the case of assembling an $n \times n$ square S_n , while simultaneously obtaining optimal tile complexity $O(\frac{\log n}{\log \log n})$.³¹ In other words, not only is it the case that every producible assembly can assemble into the terminal assembly (by the definition of directed), but in fact every producible assembly is at most $O(\log^2 n)$ attachment events from becoming the terminal assembly.

Demaine, Demaine, Fekete, Ishaque, Rafalin, Schweller, and Souvaine [16] studied a complexity measure similar to assembly depth called *stage complexity* for another variant of the aTAM known as *staged assembly*. In the staged assembly model, a hierarchical model of attachment is used, with the added ability to prepare different assemblies in separate test tubes. The separate test tubes are allowed to reach a terminal state, after which any produced nonterminal assemblies (including individual tile) are assumed to be washed away, before combining the tubes. The stage complexity of a tile system is similarly defined to be the depth of the “mixing tree” describing the order of test tube mixing steps. Our model is more restrictive by permitting only one test tube (“bin complexity 1” in the language of [16]). In a sense, Theorem 6.1 “automates” the highly selective mixing that is assumed to be externally controlled in the staged assembly model, while paying only a quadratic price in the number of parallel assembly stages required. It naturally pays a price in tile complexity as well, since unlike the staged model in which both the tile types and the mixing order can encode information, the construction of Theorem 6.1 must encode the size n of the square entirely in the tile

³¹In [6], the authors prove that whenever $n \in \mathbb{N}$ is algorithmically random, at least $\Omega(\log n / \log \log n)$ tile types are required to strictly self-assemble an $n \times n$ square in the hierarchical model. Actually, that paper states only that this holds for the *q-tile model*, in which some constant q exists that limits the size of attachable assemblies other than those containing a special seed tile, and the authors claim that the proof requires the bound q , but in fact their proof does not use the bound q and works for the general hierarchical model [48]. Thus the tile complexity obtained in Theorem 6.1 is asymptotically optimal.

types. The primary challenge in achieving a highly parallel square construction in the hierarchical model — a challenge not present in the staged assembly model — is the prevention of overlapping subassemblies.³²

Theorem 6.1. *For all $n \in \mathbb{N}$, there is a hierarchical TAS $\mathcal{T} = (T, 2)$ such that \mathcal{T} strictly self-assembles an $n \times n$ square, $|T| = O(\frac{\log n}{\log \log n})$, and $\text{depth}^{\text{da}}(\mathcal{T}) = O(\log^2 n)$.*

Proof. A high-level outline of the construction is shown in Figure 7. We assemble a number of blocks of width $O(\log n) \times O(\log n)$, each of which represents in its tile types an address indicating its position in the square, and the block binds only to (some of) its neighboring blocks. The blocks assemble using standard single-tile accretion (actually we cannot directly enforce this in the model, but each block will nonetheless assemble the same structure in either model). Since each block is $O(\log^2 n)$ total tiles, this is the source of the suboptimal $O(\log^2 n)$ assembly depth. Once the blocks are assembled, however, they assemble into the full square using $O(\log n)$ assembly depth. All blocks (x, y) with x even bind to $(x + 1, y)$ to create the two-block assembly $(x, y) : (x + 1, y)$, then all blocks $(x, y) : (x + 1, y)$ with y even bind to $(x, y + 1) : (x + 1, y + 1)$ to create the four-block assembly $(x, y) : (x + 1, y) : (x, y + 1) : (x + 1, y + 1)$, etc.

The construction will actually control the width of the square only to within an additive logarithmic factor by bring together blocks of width and length $\Theta(\log n)$; standard techniques can be used to make the square precisely $n \times n$. For instance, we could add $O(\log n)$ total filler tiles to the leftmost and bottommost blocks, while adding only $O(\log^2 n)$ to the assembly depth and $O(\log \log n)$ to the tile complexity since such filler tiles could be assembled from a counter that counts to $\log n$ using $O(\log \log n)$ tile types. For simplicity we describe the desired width n as the number of blocks instead of the desired dimensions of the square and omit the details of this last step of filling in the logarithmic gap.

Figure 7 outlines the construction of a square when the number of blocks n is a power of two. Figure 8 shows how to modify the blocks so that some of them are double in width, double in height, or both, to achieve a total square width that is an arbitrary positive integer. Each block contains the same $O(\frac{\log n}{\log \log n})$ tile types that encode n , and as the block assembles it randomly chooses x and y -coordinates, which represent an index in the square. This random choice is implemented through competition between tile types that share the same “input” glues but represent different bits of x or y . These are used to determine the block’s own size and to determine what series of bumps and dents to place on its perimeter to enforce that the only blocks that can bind are adjacent in Figure 7. The coordinates are also used to determine where to place strength-1 glues. The same strength-1 glue is used uniformly throughout Figure 7. The bumps and dents ensure that no two blocks can attach unless they are adjacent in the figure.

The growth of an individual block is shown in Figure 9. We describe the assembly as if it grows only by single-tile accretion. There are some strength-2 glues so this is not completely accurate, but the growth of the block is “polyomino-safe”, to borrow a term of Winfree [57]. By design, no assembly larger than four can form except by attachment to the growing block, and even if these assemblies attach at once to the block rather than by single-tile accretion, the correct operation of

³²Adleman [2] showed a $\Omega(n)$ lower bound (in a much different and more permissive model of assembly time than in the present paper; later improved to $\Omega(n \log n)$ by Adleman, Cheng, Goel, Huang, and Wasserman [4]) for the problem of assembling a $1 \times n$ line from n distinct tile types t_1, \dots, t_n . The intuitive reason that the time is not $O(\log n)$ is that if assemblies $\alpha_1 = t_i \dots t_j$ and $\alpha_2 = t_{i'} \dots t_{j'}$ form, with $i < i' < j < j'$, then α_1 can never attach to α_2 because they overlap. Staged assembly can be used to control the overlap directly by permitting only the growth of lines covering dyadic intervals.

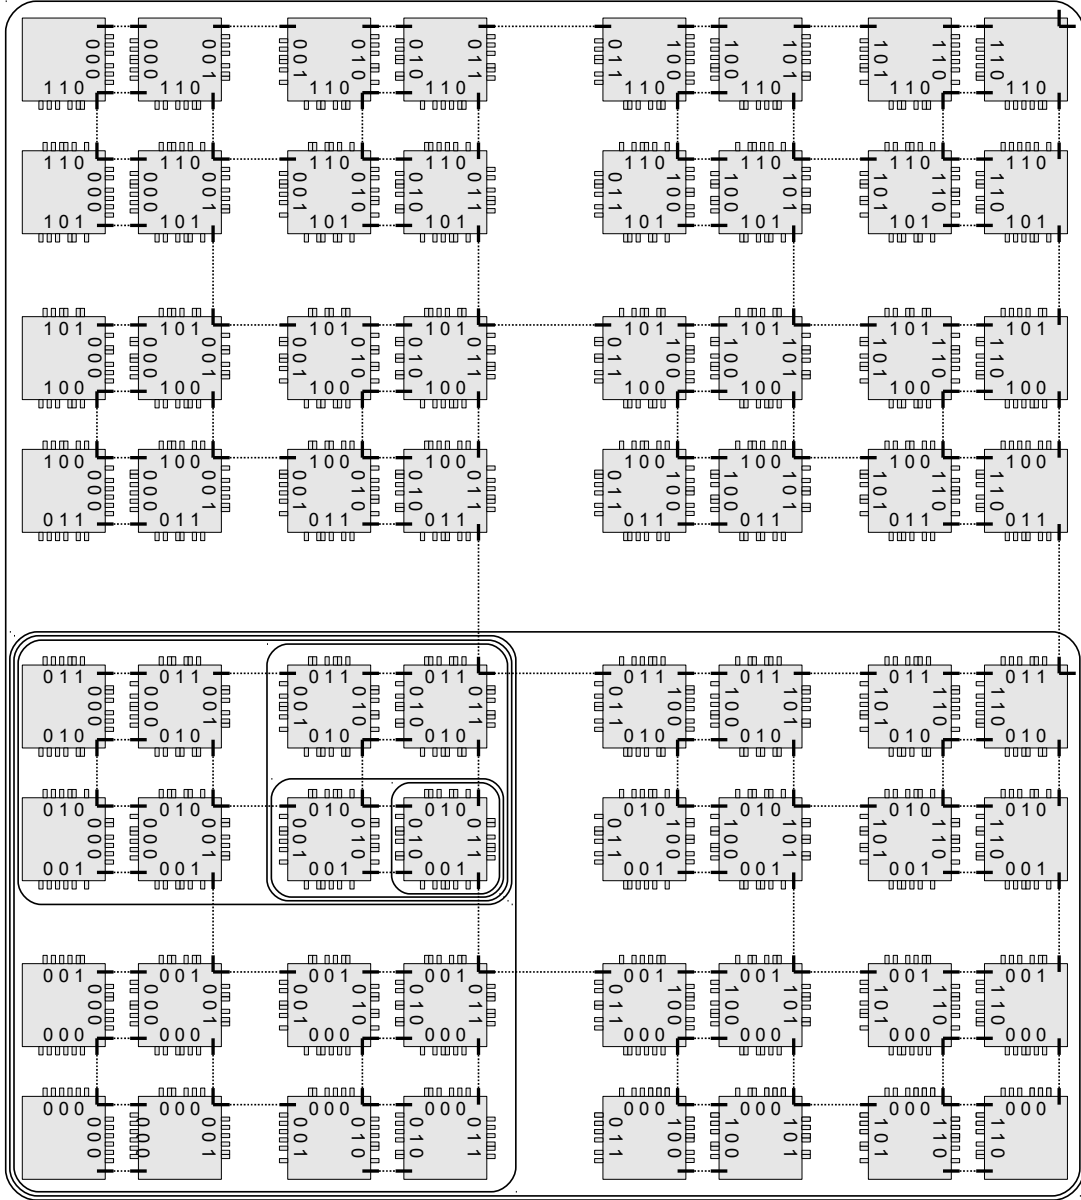


Figure 7: Overview of the hierarchical TAS that assembles an $n \times n$ square with $O(\log^2 n)$ assembly depth and $O(\frac{\log n}{\log \log n})$ tile complexity. Each square in the figure represents a block of width $O(\log n)$ with each side of each block encoding its (x, y) -address in the square. (The encoding scheme is shown in more detail in Figure 9.) Each of the thin solid lines is a strength-1 glue intended to connect the block to other blocks. Dotted lines are drawn between those glues that are intended to bind to each other. The circled subassemblies show the order of growth of one particular block (at coordinates $(3, 2)$) into the final square.

the block growth is unaffected. This is due to the fact that all strength-2 glues are “one-to-one”; no strength-2 glue is shared as an “output” (in the direction of growth in the seeded model) by two different tile types. This implies that no partial assembly occurring away from the main “seeded”




Figure 8: Design of block sizes to handle values of n that are not a power of two. There are always exactly $2^k \times 2^k$ blocks, where $2^k \leq n < 2^{k+1}$. Each block doubles its length along the x -axis (resp. y -axis) if $n - 2^k$ exceeds its x -coordinate (resp. y -coordinate).

assembly can grow “backward” and place an incorrect tile.

To form a square of size $n \times n$ “units” (where a unit is $O(\log n)$, the width and height of a small block), we choose the largest power of two $2^k \leq n$ and assemble exactly $2^k \times 2^k$ different types of blocks, doubling the width (resp. height) of the first $n - 2^k$ of them in the x -direction (resp. y -direction), as in Figure 8. The orange (medium darkness in gray-scale) tile types and the base-conversion tile types that attach to them in Figure 9 are the only non-constant set of tile types. Borrowing a technique from [3], we will represent n in base b , where $b \approx \log n / \log \log n$, using $\approx \log n / \log \log n$ unique tile types, and we use $O(\log n / \log \log n)$ tile types to convert $n - 2^k$ to binary and $O(1)$ tile types to accomplish all the other tasks needed to assemble the block.

Choose $b = 2^m$ to be a power of two such that $\log n / \log \log n \leq b < 2 \log n / \log \log n$. Each digit in base b can represent m bits of $n - 2^k$. $n - 2^k$ is encoded in exactly $m \cdot \lceil \frac{k}{m} \rceil = O(\frac{\log n}{\log \log n})$ base- b digits. The blue (dark in gray-scale) tile types in Figure 9 convert $n - 2^k$ from base b to binary and at the same time represent $n - 2^k$ with its “bit-quadrupled” version (e.g., 0110 \mapsto 0000111111110000), since each bit along the edge will eventually require width four to make room to place the bumps properly.³³ The set of base-conversion tile types from [3] of cardinality $O(\log n / \log \log n)$ can be easily modified to achieve this “bit-quadrupling” without increasing the asymptotic tile complexity. The potential doubling of height and/or width can be achieved with a constant set of tile types since the unit width is implicitly encoded in the width of the block and a constant set of “rotator” tiles similar to those in Figure 9 can be used to add extra unit-width blocks when needed. The counterclockwise order of growth ensures that if not all of the bumps and dents are formed, then at least one of the four strength-1 glues necessary for an attachment event to occur is not yet present in one of the blocks. To ensure that the TAS is directed, we do not include base-conversion tiles for

³³The bumps cannot simply be placed with strength-2 glues above a width-1 or even width-2 representation of a bit in the obvious way, otherwise there would be nothing to force that the bumps are present before the inter-block strength-1 glues. If the bumps are allowed to grow in parallel with the rest of the assembly then they may not complete fast enough. Width four is required to create a “linear assembly path” for the bumps and dents tiles to follow, ensuring that growth of the block continues only once the path is complete.




Figure 9: Assembly of $O(\log n) \times O(\log n)$ size block from $O(\frac{\log n}{\log \log n})$ tile types. Every block starts from the same tile types that encode $n - 2^k$, using exactly k bits; in this example, $n = 22$ so $k = 4$ and $n - 2^k$ is 0110 in binary using 4 bits. Thick solid lines represent strength-2 glues. For clarity, strength-1 glues are shown selectively to help verify that a certain order of growth is possible to enforce. The tiles encode $n - 2^k$ in base b chosen to be a power of two such that $\frac{\log n}{\log \log n} \leq b < 2 \frac{\log n}{\log \log n}$ (labeled “seed” for intuition, although those tiles start unattached). $n - 2^k$ is first converted to binary and each of its bits quadrupled to make room for the bumps and dents. A constant-size set of tile types does the rest. Then x and y coordinates are randomly guessed and simultaneously compared to $n - 2^k$; if either is smaller, that dimension is doubled in length (in this example the height is doubled but not the width). At the same time, the values of x and y are compared as described in the proof of Theorem 6.1 to determine where to place strength-1 glues.

any digit $d \in \{0, 1, \dots, b-1\}$ that does not appear in the base- b expansion of $n - 2^k$, otherwise those tiles will form unused terminal assemblies. Each glue in a tile type representing a bit is “marked” indicating whether it is the most significant bit, least significant bit, or interior bit, as well as being marked with its relative position among the four copies of the bit.

Once $n - 2^k$ is converted to binary, we use nondeterministic attachment of tiles to the north of this value to randomly guess $2k$ bits that represent the x - and y -coordinates of the block, meaning the binary numbers represented on the top and right, respectively, of each block in Figure 7. To be

precise, we must actually choose each of x and y to be a random bit string that is not all 1's, since each represents a connection between two blocks, of which there are $2^k - 1$ along each dimension. It is straightforward to encode into the tile types the logic that if the first k bits were 1, then the final bit must be 0. A number of additional computations are done on these values (some computations are possible to do as the values are guessed). The results of these computations will be stored in the rightmost tile type and propagated to all subsequent tile types. First, each of x and y is compared to $n - 2^k$ to determine how large to make each dimension of the block. In the example of Figure 9, $y < n - 2^k$ and $x \geq n - 2^k$, so the block is one “unit” wide and two “units” high. Also, the binary expansions of x and y are themselves compared to determine where to place strength-1 glues. After x and y are determined, to place bumps and dents on the left and bottom of the block, the values $x - 1$ and $y - 1$ must be computed, which requires assembling from least significant to most significant, so this is delayed until after the first 90-degree rotation shown in Figure 9. Once these values are computed, they are also used to determine placement of glues. The entire block is created by rotating either counter-clockwise (in the case of x even, as shown in the bottom right of Figure 9) or clockwise (in the case of x odd, not shown but the exact mirror image of the bottom right of Figure 9), placing bumps and dents and single-strength glues. The glues are placed in the order shown by the rotations, so that the last glue to be (potentially) placed is the top east-facing glue in the case of x even, or the top west-facing glue in the case of x odd.

By inspection of Figure 7, it is routine to verify that the following rules can be used to determine placement of strength-1 glues. If x is even, then place two single-strength glues on the right edge. If x is odd, then place two single-strength glues on the left edge. For a natural number n , define $t(n)$ to be the number of trailing 1's in n 's binary expansion. If $t(x) \geq t(y)$, then place exactly one strength-1 glue on the top edge. If $t(x) \geq t(y - 1)$, then place exactly one strength-1 glue on the bottom edge. If $t(x - 1) \leq t(y) + 1$ and x is even, then place exactly one strength-1 glue on the left edge. If $t(x) \leq t(y) + 1$ and x is odd, then place exactly one strength-1 glue on the right edge.

Each of these computations (for placement of glues and for determining block dimensions) can be computed by a deterministic finite automaton whose input symbols represent tuples of bits from $n - 2^k$, x , $x - 1$, $y - 1$, and y . These automata can then be combined in a product construction and embedded into the tile types that accrete in the row above $n - 2^k$ if only x and y are needed, and embedded into tile types that are placed after the first rotation if $x - 1$ or $y - 1$ is needed. Since the decision for placing glue on the top edge requires only x and y , this ensures that the decision for each glue placement can be made before the region containing the potential glue site is assembled.

As shown in Figure 9, some padding with filler tiles is necessary to make the block a perfect rectangle. Also, some padding is needed in the case of a doubling of height or width, to ensure that the resulting assembly has height or width precisely twice that of the non-doubled version. \square

7 Open Questions

There are some interesting questions that remain open. Say that a tile system \mathcal{T} *strictly self-assembles* a shape $S \subseteq \mathbb{Z}^2$ if all terminal assemblies $\hat{\alpha}$ of \mathcal{T} , appropriately translated, satisfy $\text{dom } \hat{\alpha} = S$.

1. What upper or lower bound can be placed on the quantity $\text{depth}^{\text{da}}(\mathcal{T})$ for \mathcal{T} a hierarchical TAS that strictly self-assembles an $n \times n$ square with optimal tile complexity $O(\frac{\log n}{\log \log n})$ (or even with nearly-optimal tile complexity $O(\log n)$)? It is not obvious how to show either

$\text{depth}^{\text{da}}(\mathcal{T}) = o(\log^2 n)$ for some such \mathcal{T} or $\text{depth}^{\text{da}}(\mathcal{T}) = \omega(\log n)$ for all such \mathcal{T} . Obtaining bounds for more general shapes would also be interesting.

2. What is the complexity of the following decision problems?

$$\text{HIERDIRECTEDASSEMBLY} = \left\{ \langle \alpha, \mathcal{T} \rangle \mid \begin{array}{l} \mathcal{T} \text{ is a directed hierarchical TAS with} \\ \text{unique producible terminal assembly } \alpha \end{array} \right\},$$

$$\text{HIERDIRECTEDSHAPE} = \left\{ \langle S, \mathcal{T} \rangle \mid \begin{array}{l} \mathcal{T} \text{ is a directed hierarchical TAS that} \\ \text{strictly self-assembles finite shape } S \end{array} \right\},$$

$$\text{HIERUNIQUESHAPE} = \left\{ \langle S, \mathcal{T} \rangle \mid \begin{array}{l} \mathcal{T} \text{ is a hierarchical TAS that} \\ \text{strictly self-assembles finite shape } S \end{array} \right\}.$$

In the case of the seeded aTAM, the seeded variants of these problems are known to be in P [5] for the first two, and coNP-complete [6] for the last.

For the case of 3D hierarchical tile systems, HIERDIRECTEDASSEMBLY was shown to be coNP-complete by Cannon, Demaine, Demaine, Eisenstat, Patitz, Schweller, Summers, and Winslow [12].³⁴ Furthermore, their proof shows that the 3D version of HIERDIRECTEDSHAPE is coNP-hard.³⁵

See [7] for definitions of complexity classes Σ_i^P and Π_i^P , where $\text{NP} = \Sigma_1^P$ and $\text{coNP} = \Pi_1^P$. The ‘‘obvious’’ containments are $\text{HIERDIRECTEDASSEMBLY} \in \text{coNP}$ (proven in [12, Lemma 4.3]), $\text{HIERUNIQUESHAPE} \in \Pi_2^P$ ³⁶ and $\text{HIERDIRECTEDSHAPE} \in \Pi_2^P$.³⁷ It is open whether HIERDIRECTEDASSEMBLY is coNP-hard (in 2D), and whether HIERUNIQUESHAPE or HIERDIRECTEDSHAPE are Π_2^P -hard. The proof of HIERUNIQUESHAPE in the ‘‘multiple-tile’’ model of [6] can be used to show that HIERUNIQUESHAPE (in the hierarchical aTAM as defined in this paper) is coNP-hard [6, 48].

³⁴In that paper, the problem is called the *Unique Assembly Verification* problem.

³⁵Their technique to reduce the complement of SAT to the problem is such that, if the formula is unsatisfiable, then \mathcal{T} has a unique terminal assembly α (hence strictly self-assembles the shape $\text{dom } \alpha$), and if the formula is satisfiable, then \mathcal{T} produces multiple terminal assemblies, and at least two of them are guaranteed to have different shapes. Therefore their proof also shows that the 3D version of HIERDIRECTEDSHAPE is coNP-hard. The proof that HIERDIRECTEDASSEMBLY \in coNP (which applies to any number of dimensions) does not so easily apply to HIERDIRECTEDSHAPE, so the computational complexity of the 2D shape version of the problem is still open.

³⁶The producibility of α is decidable in polynomial time [21]. For $k \in \mathbb{N}$, Let $\mathcal{A}[\mathcal{T}]_{\leq k} = \{\alpha \in \mathcal{A}[\mathcal{T}] \mid |\alpha| \leq k\}$. Then $\langle S, \mathcal{T} \rangle \in \text{HIERUNIQUESHAPE}$ if and only if for all $\alpha \in \mathcal{A}[\mathcal{T}]_{\leq 2|S|}$:

- $|\alpha| \leq |S|$ (if this is verified for all $\alpha \in \mathcal{A}[\mathcal{T}]_{\leq 2|S|}$ then no assembly larger than $2|S|$ is producible either),
- if $|\alpha| < |S|$ then there exists $\gamma \in \mathcal{A}[\mathcal{T}]_{\leq |S|}$ attachable to α (so $\alpha \notin \mathcal{A}_{\square}[\mathcal{T}]$), and
- if $|\alpha| = |S|$ (so $\alpha \in \mathcal{A}_{\square}[\mathcal{T}]$ since nothing larger than S is producible) then $S = \text{dom } \alpha$.

The second condition is a $\forall \exists$ quantifier that makes the problem in Π_2^P ; the other conditions have only one \exists or \forall . Note that the second check guarantees no assembly strictly smaller than $|S|$ is terminal, and the first check guarantees that no assembly strictly larger than $|S|$ is producible. Therefore there must be at least one $\alpha \in \mathcal{A}_{\square}[\mathcal{T}]$ with $|\alpha| = |S|$ (and the third check guarantees that it has shape S).

³⁷Using similar reasoning as above, we have $\langle S, \mathcal{T} \rangle \in \text{HIERDIRECTEDSHAPE}$ if and only if for all $\alpha \in \mathcal{A}[\mathcal{T}]_{\leq 2|S|}$:

- $|\alpha| \leq |S|$,
- if $|\alpha| < |S|$ then there exists $\gamma \in \mathcal{A}[\mathcal{T}]_{\leq |S|}$ attachable to α (so $\alpha \notin \mathcal{A}_{\square}[\mathcal{T}]$), and
- for all $\beta \in \mathcal{A}[\mathcal{T}]_{\leq |S|}$, if $|\alpha| = |\beta| = |S|$ (so $\alpha, \beta \in \mathcal{A}_{\square}[\mathcal{T}]$) then $\alpha = \beta$ and $S = \text{dom } \alpha$.

3. What is the complexity of the following decision problems?

$$\text{HIERMINTILESET} = \left\{ \langle S, c \rangle \mid \begin{array}{l} (\exists \mathcal{T} = (T, \tau)) \mathcal{T} \text{ is a hierarchical TAS with} \\ |T| \leq c \text{ and } \mathcal{T} \text{ strictly self-assembles finite shape } S \end{array} \right\},$$

$$\text{HIERDIRECTEDMINTILESET} = \left\{ \langle S, c \rangle \mid \begin{array}{l} (\exists \mathcal{T} = (T, \tau)) \mathcal{T} \text{ is a directed hierarchical TAS with} \\ |T| \leq c \text{ and } \mathcal{T} \text{ strictly self-assembles finite shape } S \end{array} \right\}.$$

In the case of the seeded aTAM, the seeded variants of these problems are known to be Σ_2^P -complete [11] and NP-complete [5], respectively.

4. What is the optimal time complexity of strictly self-assembling an $n \times n$ square with a hierarchical TAS? Any shape with diameter n ? What if we require the TAS to be directed?
5. Two asymptotically unrealistic aspects of the model are the assumption of a constant rate of diffusion of assemblies and a constant binding strength threshold required to bind two assemblies together. Large assemblies will diffuse more slowly in a well-mixed solution; some simple models predict that the diffusion rate of a molecule is inversely proportional to its diameter [10, 42]. It is conceivable that an assembly model properly accounting for diffusion rates could enforce an absolute lower bound of $\Omega(D)$ on the assembly time required to assemble any shape of diameter D .

The binding strength threshold of the seeded aTAM is a simplified model of a more complicated approximation in the *kinetic* tile assembly model (kTAM, [55]). Tiles in reality will occasionally detach, but so long as their concentration is sufficient, another tile will reattach after not too much time. While our model accounts directly for concentrations of large assemblies, it only accounts for this concentration up to the moment of first binding. A more realistic model might require a larger binding strength threshold to balance the fact that, if a large assembly detaches, it may take a long time to reattach. In particular, the seeded aTAM is justifiable as a model, despite its lack of reverse reactions or modeling of strength-1 attachments (which happen in reality but have a higher reverse rate than higher strength attachments), in part due to Winfree’s proof [55, 56] that under suitable conditions (in particular setting concentrations and binding energies such that the rate of forward attachments is just barely larger than the rate of backward detachments of strength-2-bound tiles), the kTAM “simulates the aTAM with high probability.” It is an open question whether there is any similar theorem that can be proven in the hierarchical aTAM, showing that detachment reactions may be safely ignored under certain conditions.

Incorporating these and other physical phenomena into the hierarchical assembly model would be an interesting challenge.

Acknowledgement. The authors are especially grateful to David Soloveichik for contributions to this paper, including the proof of Theorem 3.1, discussion and insights on other proofs, and generally for indispensable help with discovering and solidifying the results. The authors also thank Adam Marblestone, Robbie Schweller, Matt Patitz, and the members of Erik Winfree’s group, particularly Joe Schaeffer, Erik Winfree, Damien Woods, and Seung Woo Shin, for insightful discussion and comments, and to Bernie Yurke and Rizal Hariadi for discussing diffusion rates and pointing the

authors to [10, 42]. We are grateful to anonymous reviewers for identifying important flaws in the original proofs of some theorems, one of which spurred significant followup work [15].

References

- [1] Zachary Abel, Nadia Benbernou, Mirela Damian, Erik D. Demaine, Martin L. Demaine, Robin Flatland, Scott Kominers, and Robert Schweller. Shape replication through self-assembly and RNase enzymes. In *SODA 2010: Proceedings of the Twenty-first Annual ACM-SIAM Symposium on Discrete Algorithms*, Austin, Texas, 2010. Society for Industrial and Applied Mathematics.
- [2] Leonard M. Adleman. Toward a mathematical theory of self-assembly (extended abstract). Technical Report 00-722, University of Southern California, 2000.
- [3] Leonard M. Adleman, Qi Cheng, Ashish Goel, and Ming-Deh Huang. Running time and program size for self-assembled squares. In *STOC 2001: Proceedings of the thirty-third annual ACM Symposium on Theory of Computing*, pages 740–748, Hersonissos, Greece, 2001. ACM.
- [4] Leonard M. Adleman, Qi Cheng, Ashish Goel, Ming-Deh Huang, and Hal Wasserman. Linear self-assemblies: Equilibria, entropy and convergence rates. In *In Sixth International Conference on Difference Equations and Applications*. Taylor and Francis, 2001.
- [5] Leonard M. Adleman, Qi Cheng, Ashish Goel, Ming-Deh A. Huang, David Kempe, Pablo Moisset de Espanés, and Paul W. K. Rothmund. Combinatorial optimization problems in self-assembly. In *STOC 2002: Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing*, pages 23–32, 2002.
- [6] 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.
- [7] Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, 2009.
- [8] 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.
- [9] Florent Becker, Ivan Rapaport, and Eric Rémila. Self-assembling classes of shapes with a minimum number of tiles, and in optimal time. In *FSTTCS 2006: Foundations of Software Technology and Theoretical Computer Science*, pages 45–56, 2006.
- [10] Otto G. Berg and Peter H. von Hippel. Diffusion-controlled macromolecular interactions. *Annual Review of Biophysics and Biophysical Chemistry*, 14(1):131–160, 1985.
- [11] Nathaniel Bryans, Ehsan Chiniforooshan, David Doty, Lila Kari, and Shinnosuke Seki. The power of nondeterminism in self-assembly. *Theory of Computing*, 9(1):1–29, 2013. Preliminary version appeared in SODA 2011.

- [12] 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.
- [13] Harish Chandran, Nikhil Gopalkrishnan, and John H. Reif. The tile complexity of linear assemblies. In *ICALP 2009: 36th International Colloquium on Automata, Languages and Programming*, volume 5555, pages 235–253. Springer, 2009.
- [14] 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.
- [15] Ho-Lin Chen, David Doty, Ján Manuch, Arash Rafiey, and Ladislav Stacho. Pattern overlap implies runaway growth in hierarchical tile systems. *JoCG 2015: Journal of Computational Geometry*, 7(2):3–18, 2015. Special issue of invited papers from SoCG 2015.
- [16] Erik D. Demaine, Martin L. Demaine, Sándor P. Fekete, Mashhood Ishaque, Eynat Rafalin, Robert T. Schweller, and Diane L. Souvaine. Staged self-assembly: Nanomanufacture of arbitrary shapes with $O(1)$ glues. *Natural Computing*, 7(3):347–370, 2008. Preliminary version appeared in DNA 2007.
- [17] 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.
- [18] Erik D. Demaine, Matthew J. Patitz, Robert T. Schweller, and Scott M. Summers. Self-assembly of arbitrary shapes using RNase enzymes: Meeting the Kolmogorov bound with small scale factor. In *STACS 2011: Proceedings of the 28th International Symposium on Theoretical Aspects of Computer Science*, 2011.
- [19] David Doty. Randomized self-assembly for exact shapes. *SIAM Journal on Computing*, 39(8):3521–3552, 2010. Preliminary version appeared in FOCS 2009.
- [20] David Doty. Theory of algorithmic self-assembly. *Communications of the ACM*, 55(12):78–88, December 2012.
- [21] David Doty. Producibility in hierarchical self-assembly. *Natural Computing*, 15(1):41–49, 2016. Special issue of invited papers from UCNC 2014.
- [22] 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.
- [23] Irvin R. Epstein and John A. Pojman. *An introduction to nonlinear chemical dynamics: Oscillations, waves, patterns, and chaos*. Oxford University Press, USA, 1998.

- [24] Constantine Evans. *Crystals that count! Physical principles and experimental investigations of DNA tile self-assembly*. PhD thesis, California Institute of Technology, 2014.
- [25] Bin Fu, Matthew J. Patitz, Robert T. Schweller, and Robert Sheline. Self-assembly with geometric tiles. In *ICALP 2012: Proceedings of the 39th International Colloquium on Automata, Languages and Programming*, pages 714–725, July 2012.
- [26] Kenichi Fujibayashi, Rizal Hariadi, Sung Ha Park, Erik Winfree, and Satoshi Murata. Toward reliable algorithmic self-assembly of DNA tiles: A fixed-width cellular automaton pattern. *Nano Letters*, 8(7):1791–1797, 2007.
- [27] Michael A. Gibson and Jehoshua Bruck. Efficient exact stochastic simulation of chemical systems with many species and many channels. *Journal of Physical Chemistry A*, 104(9):1876–1889, March 2000.
- [28] Daniel T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *Journal of Physical Chemistry*, 81(25):2340–2361, 1977.
- [29] Daniel T. Gillespie. Stochastic simulation of chemical kinetics. *Annual review of physical chemistry*, 58(1):35–55, 2007.
- [30] 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.
- [31] 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.
- [32] 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.
- [33] Ming-Yang Kao and Robert T. Schweller. Reducing tile complexity for self-assembly through temperature programming. In *SODA 2006: Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 571–580, 2006.
- [34] Ming-Yang Kao and Robert T. Schweller. Randomized self-assembly for approximate shapes. In *ICALP 2008: International Colloquium on Automata, Languages, and Programming*, volume 5125 of *Lecture Notes in Computer Science*, pages 370–384. Springer, 2008.
- [35] James I. Lathrop, Jack H. Lutz, and Scott M. Summers. Strict self-assembly of discrete Sierpinski triangles. *Theoretical Computer Science*, 410:384–405, 2009. Preliminary version appeared in CiE 2007.
- [36] Chris Luhrs. Polyomino-safe DNA self-assembly via block replacement. In Ashish Goel, Friedrich C. Simmel, and Petr Sosík, editors, *DNA14*, volume 5347 of *Lecture Notes in Computer Science*, pages 112–126. Springer, 2008.
- [37] Ján Maňuch, Ladislav Stacho, and Christine Stoll. Step-assembly with a constant number of tile types. In *ISAAC 2009: Proceedings of the 20th International Symposium on Algorithms and Computation*, pages 954–963, Berlin, Heidelberg, 2009. Springer-Verlag.

- [38] Jennifer E. Padilla, Wenyan Liu, and Nadrian C. Seeman. Programmed hierarchical self-assembly of a quasiperiodic tiling pattern. In *FNANO 2011: 8th Annual Conference on Foundations of Nanoscience: Self-Assembled Architectures and Devices (Snowbird, Utah, USA, April 11-15, 2011)*, pages 49–50. Sciencetechnica, 2011.
- [39] Matthew J. Patitz. An introduction to tile-based self-assembly and a survey of recent results. *Natural Computing*, 13(2):195–224, 2014.
- [40] Matthew J. Patitz and Scott M. Summers. Identifying shapes using self-assembly. *Algorithmica*, 64(3):481–510, 2012. Preliminary version appeared in ISAAC 2010.
- [41] John H. Reif. Local parallel biomolecular computation. In *DNA 3: Proceedings of the 3rd International Meeting on DNA-Based Computers*, Series in Discrete Mathematics and Theoretical Computer Science, pages 217–254. AMS/DIMACS, 1999.
- [42] Jacob Riseman and John G. Kirkwood. The intrinsic viscosity, translational and rotatory diffusion constants of rod-like macromolecules in solution. *The Journal of Chemical Physics*, 18:512–516, 1950.
- [43] Paul W. K. Rothmund. *Theory and Experiments in Algorithmic Self-Assembly*. PhD thesis, University of Southern California, December 2001.
- [44] Paul W. K. Rothmund, Nick Papadakis, and Erik Winfree. Algorithmic self-assembly of DNA Sierpinski triangles. *PLoS Biology*, 2(12):2041–2053, 2004.
- [45] 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.
- [46] 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.
- [47] 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.
- [48] Robert Schweller, 2016. personal communication.
- [49] Nadrian C. Seeman. Nucleic-acid junctions and lattices. *Journal of Theoretical Biology*, 99:237–247, 1982.
- [50] David Soloveichik and Erik Winfree. Complexity of self-assembled shapes. *SIAM Journal on Computing*, 36(6):1544–1569, 2007. Preliminary version appeared in DNA 2004.
- [51] Scott M. Summers. Reducing tile complexity for the self-assembly of scaled shapes through temperature programming. *Algorithmica*, 63(1-2):117–136, 2012.
- [52] Leroy G. Wade. *Organic Chemistry*. Prentice Hall, 2nd edition, 1991.
- [53] Hao Wang. Proving theorems by pattern recognition – II. *The Bell System Technical Journal*, XL(1):1–41, 1961.

- [54] Hao Wang. Dominoes and the AEA case of the decision problem. In *Proceedings of the Symposium on Mathematical Theory of Automata (New York, 1962)*, pages 23–55. Polytechnic Press of Polytechnic Inst. of Brooklyn, Brooklyn, N.Y., 1963.
- [55] Erik Winfree. *Algorithmic Self-Assembly of DNA*. PhD thesis, California Institute of Technology, June 1998.
- [56] Erik Winfree. Simulations of computing by self-assembly. Technical Report CaltechC-STR:1998.22, California Institute of Technology, 1998.
- [57] 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.
- [58] Andrew Winslow. A brief tour of theoretical tile self-assembly. In *International Workshop on Cellular Automata and Discrete Complex Systems*, pages 26–31. Springer, 2016.

A Appendix: Formal Definition of Abstract Tile Assembly Model

This section gives a terse definition of the abstract Tile Assembly Model (aTAM, [55]). This is not a tutorial; for readers unfamiliar with the aTAM, [45] gives an excellent introduction to the model.

Fix an alphabet Σ . Σ^* is the set of finite strings over Σ . Given a discrete object O , $\langle O \rangle$ denotes a standard encoding of O as an element of Σ^* . \mathbb{Z} , \mathbb{Z}^+ , \mathbb{N} , \mathbb{R}^+ denote the set of integers, positive integers, nonnegative integers, and nonnegative real numbers, respectively. For a set A , $\mathcal{P}(A)$ denotes the power set of A . Given $A \subseteq \mathbb{Z}^2$, the *full grid graph* of A is the undirected graph $G_A^f = (V, E)$, where $V = A$, and for all $u, v \in V$, $\{u, v\} \in E \iff \|u - v\|_2 = 1$; i.e., if and only if u and v are adjacent on the integer Cartesian plane. A *shape* is a set $S \subseteq \mathbb{Z}^2$ such that G_S^f is connected.

A *tile type* is a tuple $t \in (\Sigma^* \times \mathbb{N})^4$; i.e., a unit square with four sides listed in some standardized order, each side having a *glue label* (a.k.a. *glue*) $\ell \in \Sigma^*$ and a nonnegative integer *strength*. For a set of tile types T , let $\Lambda(T) \subset \Sigma^*$ denote the set of all glue labels of tile types in T . Let $\{\mathbf{N}, \mathbf{S}, \mathbf{E}, \mathbf{W}\}$ denote the *directions* consisting of unit vectors $\{(0, 1), (0, -1), (1, 0), (-1, 0)\}$. Given a tile type t and a direction $d \in \{\mathbf{N}, \mathbf{S}, \mathbf{E}, \mathbf{W}\}$, $t(d) \in \Lambda(T)$ denotes the glue label on t in direction d . We assume a finite set T of tile types, but an infinite number of copies of each tile type, each copy referred to as a *tile*. An *assembly* is a nonempty connected arrangement of tiles on the integer lattice \mathbb{Z}^2 , i.e., a partial function $\alpha : \mathbb{Z}^2 \dashrightarrow T$ such that $G_{\text{dom } \alpha}^f$ is connected and $\text{dom } \alpha \neq \emptyset$. The *shape of α* is $\text{dom } \alpha$. Write $|\alpha|$ to denote $|\text{dom } \alpha|$. Given two assemblies $\alpha, \beta : \mathbb{Z}^2 \dashrightarrow T$, we say α is a *subassembly* of β , and we write $\alpha \sqsubseteq \beta$, if $\text{dom } \alpha \subseteq \text{dom } \beta$ and, for all points $p \in \text{dom } \alpha$, $\alpha(p) = \beta(p)$.

Given two assemblies α and β , we say α and β are *equivalent up to translation*, written $\alpha \simeq \beta$, if there is a vector $\vec{z} \in \mathbb{Z}^2$ such that $\text{dom } \alpha = \text{dom } \beta + \vec{z}$ (where for $A \subseteq \mathbb{Z}^2$, $A + \vec{z}$ is defined to be $\{p + \vec{z} \mid p \in A\}$) and for all $p \in \text{dom } \beta$, $\alpha(p + \vec{z}) = \beta(p)$. In this case we say that β is a *translation* of α . Given a shape $S \subseteq \mathbb{Z}^2$, we say that S is *canonical* if $S \subseteq \mathbb{N}^2$, $(x, 0) \in S$ for some $x \in \mathbb{N}$, and $(0, y) \in S$ for some $y \in \mathbb{N}$. In other words, S is located entirely in the first quadrant, but as far to the left and down as possible. We say an assembly α is *canonical* if $\text{dom } \alpha$ is canonical. For each finite assembly α , there is exactly one canonical assembly $\tilde{\alpha}$ such that $\alpha \simeq \tilde{\alpha}$. Given such a finite α , we say $\tilde{\alpha}$ is the *canonical assembly of α* . For brevity and clarity, we will tend to abuse notation and

speak of assemblies equivalent up to translation as if they are the same object, often taking α to mean $\tilde{\alpha}$, particularly when discussing concentrations. We have fixed assemblies at certain positions on \mathbb{Z}^2 only for mathematical convenience in some contexts, but of course real assemblies float freely in solution and do not have a fixed position.

Let α be an assembly and let $p \in \text{dom } \alpha$ and $d \in \{\text{N, S, E, W}\}$ such that $p + d \in \text{dom } \alpha$. Let $t = \alpha(p)$ and $t' = \alpha(p + d)$. We say that the tiles t and t' at positions p and $p + d$ *interact* if $t(d) = t'(-d)$ and $g(t(d)) > 0$, i.e., if the glue labels on their abutting sides are equal and have positive strength. Each assembly α induces a *binding graph* G_α^b , a grid graph $G = (V_\alpha, E_\alpha)$, where $V_\alpha = \text{dom } \alpha$, and $\{p_1, p_2\} \in E_\alpha \iff \alpha(p_1)$ interacts with $\alpha(p_2)$.³⁸ Given $\tau \in \mathbb{Z}^+$, α is τ -*stable* if every cut of G_α^b has weight at least τ , where the weight of an edge is the strength of the glue it represents. That is, α is τ -stable if at least energy τ is required to separate α into two parts. When τ is clear from context, we say α is *stable*.

A.1 Seeded aTAM

A *seeded tile assembly system* (seeded TAS) is a triple $\mathcal{T} = (T, \sigma, \tau)$, where T is a finite set of tile types, $\sigma : \mathbb{Z}^2 \dashrightarrow T$ is the finite, τ -stable *seed assembly*, and $\tau \in \mathbb{Z}^+$ is the *temperature*. Given two τ -stable assemblies $\alpha, \beta : \mathbb{Z}^2 \dashrightarrow T$, we write $\alpha \rightarrow_1^{\mathcal{T}} \beta$ if $\alpha \sqsubseteq \beta$ and $|\text{dom } \beta \setminus \text{dom } \alpha| = 1$. In this case we say α \mathcal{T} -*produces* β in one step.³⁹ If $\alpha \rightarrow_1^{\mathcal{T}} \beta$, $\text{dom } \beta \setminus \text{dom } \alpha = \{p\}$, and $t = \beta(p)$, we write $\beta = \alpha + (p \mapsto t)$. The \mathcal{T} -*frontier* of α is the set $\partial^{\mathcal{T}} \alpha = \bigcup_{\alpha \rightarrow_1^{\mathcal{T}} \beta} \text{dom } \beta \setminus \text{dom } \alpha$, the set of empty locations at which a tile could stably attach to α .

A sequence of $k \in \mathbb{Z}^+ \cup \{\infty\}$ assemblies $\vec{\alpha} = (\alpha_0, \alpha_1, \dots)$ is a \mathcal{T} -*assembly sequence* if, for all $1 \leq i < k$, $\alpha_{i-1} \rightarrow_1^{\mathcal{T}} \alpha_i$. We write $\alpha \rightarrow^{\mathcal{T}} \beta$, and we say α \mathcal{T} -*produces* β (in 0 or more steps) if there is a \mathcal{T} -assembly sequence $\vec{\alpha} = (\alpha_0, \alpha_1, \dots)$ of length $k = |\text{dom } \beta \setminus \text{dom } \alpha| + 1$ such that 1) $\alpha = \alpha_0$, 2) $\text{dom } \beta = \bigcup_{0 \leq i < k} \text{dom } \alpha_i$, and 3) for all $0 \leq i < k$, $\alpha_i \sqsubseteq \beta$. In this case, we say that β is the *result* of $\vec{\alpha}$, written $\beta = \text{res}(\vec{\alpha})$. If k is finite then it is routine to verify that $\text{res}(\vec{\alpha}) = \alpha_{k-1}$.⁴⁰ We say α is \mathcal{T} -*producible* if $\sigma \rightarrow^{\mathcal{T}} \alpha$, and we write $\mathcal{A}[\mathcal{T}]$ to denote the set of \mathcal{T} -producible canonical assemblies. The relation $\rightarrow^{\mathcal{T}}$ is a partial order on $\mathcal{A}[\mathcal{T}]$ [35, 43].⁴¹ A \mathcal{T} -assembly sequence $\alpha_0, \alpha_1, \dots$ is *fair* if, for all i and all $p \in \partial^{\mathcal{T}} \alpha_i$, there exists j such that $\alpha_j(p)$ is defined; i.e., no frontier location is “starved”.

An assembly α is \mathcal{T} -*terminal* if α is τ -stable and $\partial^{\mathcal{T}} \alpha = \emptyset$. It is easy to check that an assembly sequence $\vec{\alpha}$ is fair if and only $\text{res}(\vec{\alpha})$ is terminal. We write $\mathcal{A}_{\square}[\mathcal{T}] \subseteq \mathcal{A}[\mathcal{T}]$ to denote the set of \mathcal{T} -producible, \mathcal{T} -terminal canonical assemblies.

A seeded TAS \mathcal{T} is *directed* (a.k.a., *deterministic, confluent*) if the poset $(\mathcal{A}[\mathcal{T}], \rightarrow^{\mathcal{T}})$ is directed; i.e., if for each $\alpha, \beta \in \mathcal{A}[\mathcal{T}]$, there exists $\gamma \in \mathcal{A}[\mathcal{T}]$ such that $\alpha \rightarrow^{\mathcal{T}} \gamma$ and $\beta \rightarrow^{\mathcal{T}} \gamma$.⁴² We say that

³⁸For $G_{\text{dom } \alpha}^f = (V_{\text{dom } \alpha}, E_{\text{dom } \alpha})$ and $G_\alpha^b = (V_\alpha, E_\alpha)$, G_α^b is a spanning subgraph of $G_{\text{dom } \alpha}^f$: $V_\alpha = V_{\text{dom } \alpha}$ and $E_\alpha \subseteq E_{\text{dom } \alpha}$.

³⁹Intuitively $\alpha \rightarrow_1^{\mathcal{T}} \beta$ means that α can grow into β by the addition of a single tile; the fact that we require both α and β to be τ -stable implies in particular that the new tile is able to bind to α with strength at least τ . It is easy to check that had we instead required only α to be τ -stable, and required that the cut of β separating α from the new tile has strength at least τ , then this implies that β is also τ -stable.

⁴⁰If we had defined the relation $\rightarrow^{\mathcal{T}}$ based on only finite assembly sequences, then $\rightarrow^{\mathcal{T}}$ would be simply the reflexive, transitive closure $(\rightarrow_1^{\mathcal{T}})^*$ of $\rightarrow_1^{\mathcal{T}}$. But this would mean that no infinite assembly could be produced from a finite assembly, even though there is a well-defined, unique “limit assembly” of every infinite assembly sequence.

⁴¹In fact it is a partial order on the set of τ -stable assemblies, including even those that are not \mathcal{T} -producible.

⁴²The following two convenient characterizations of “directed” are routine to verify. \mathcal{T} is directed if and only if $|\mathcal{A}_{\square}[\mathcal{T}]| = 1$. \mathcal{T} is *not* directed if and only if there exist $\alpha, \beta \in \mathcal{A}[\mathcal{T}]$ and $p \in \text{dom } \alpha \cap \text{dom } \beta$ such that $\alpha(p) \neq \beta(p)$.

a TAS \mathcal{T} *strictly self-assembles* a shape $S \subseteq \mathbb{Z}^2$ if, for all $\alpha \in \mathcal{A}_\square[\mathcal{T}]$, $\text{dom } \alpha = S$; i.e., if every terminal assembly produced by \mathcal{T} has shape S . If \mathcal{T} strictly self-assembles some shape S , we say that \mathcal{T} is *strict*. Note that the implication “ \mathcal{T} is directed \implies \mathcal{T} is strict” holds, but the converse does not hold.

A.2 Hierarchical aTAM

A *hierarchical tile assembly system* (hierarchical TAS) is a pair $\mathcal{T} = (T, \tau)$, where T is a finite set of tile types, and $\tau \in \mathbb{Z}^+$ is the *temperature*. Let $\alpha, \beta : \mathbb{Z}^2 \dashrightarrow T$ be two (possibly non-canonical) assemblies. Say that α and β are *nonoverlapping* if $\text{dom } \alpha \cap \text{dom } \beta = \emptyset$. If α and β are nonoverlapping assemblies, define $\alpha \cup \beta$ to be the assembly γ defined by $\gamma(p) = \alpha(p)$ for all $p \in \text{dom } \alpha$, $\gamma(p) = \beta(p)$ for all $p \in \text{dom } \beta$, and $\gamma(p)$ is undefined for all $p \in \mathbb{Z}^2 \setminus (\text{dom } \alpha \cup \text{dom } \beta)$. An assembly γ is *singular* if $\gamma(p) = t$ for some $p \in \mathbb{Z}^2$ and some $t \in T$ and $\gamma(p')$ is undefined for all $p' \in \mathbb{Z}^2 \setminus \{p\}$. Given a hierarchical TAS $\mathcal{T} = (T, \tau)$, an assembly γ is \mathcal{T} -*producible* if either 1) γ is singular, or 2) there exist producible nonoverlapping assemblies α and β such that $\gamma = \alpha \cup \beta$ and γ is τ -stable. In the latter case, write $\alpha + \beta \rightarrow \gamma$. An assembly α is \mathcal{T} -*terminal* if for every producible assembly β such that α and β are nonoverlapping, $\alpha \cup \beta$ is not τ -stable.⁴³ Define $\mathcal{A}[\mathcal{T}]$ to be the set of all \mathcal{T} -producible canonical assemblies. Define $\mathcal{A}_\square[\mathcal{T}] \subseteq \mathcal{A}[\mathcal{T}]$ to be the set of all \mathcal{T} -producible, \mathcal{T} -terminal canonical assemblies. A hierarchical TAS \mathcal{T} is *directed* (a.k.a., *deterministic*, *confluent*) if $|\mathcal{A}_\square[\mathcal{T}]| = 1$. We say that a TAS \mathcal{T} *strictly self-assembles* a shape $S \subseteq \mathbb{Z}^2$ if, for all $\alpha \in \mathcal{A}_\square[\mathcal{T}]$, $\text{dom } \alpha = S$; i.e., if every terminal assembly produced by \mathcal{T} has shape S .

Let \mathcal{T} be a hierarchical TAS, and let $\alpha \in \mathcal{A}[\mathcal{T}]$ be a \mathcal{T} -producible assembly. An *assembly tree* Υ of $\hat{\alpha}$ is a full binary tree with $|\hat{\alpha}|$ leaves, whose nodes are labeled by \mathcal{T} -producible assemblies, with $\hat{\alpha}$ labeling the root, singular assemblies labeling the leaves, and node u labeled with γ having children u_1 labeled with α and u_2 labeled with β , with the requirement that $\alpha + \beta \rightarrow \gamma$. That is, Υ represents one possible pathway through which $\hat{\alpha}$ could be produced from individual tile types in \mathcal{T} . Let $\Upsilon(\mathcal{T})$ denote the set of all assembly trees of \mathcal{T} . Say that an assembly tree is \mathcal{T} -*terminal* if its root is a \mathcal{T} -terminal assembly. Let $\Upsilon_\square(\mathcal{T})$ denote the set of all \mathcal{T} -terminal assembly trees of \mathcal{T} . Note that even a directed hierarchical TAS can have multiple terminal assembly trees that all have the same root terminal assembly.

When \mathcal{T} is clear from context, we may omit \mathcal{T} from the notation above and instead write \rightarrow_1 , \rightarrow , $\partial\alpha$, *frontier*, *assembly sequence*, *produces*, *producible*, and *terminal*. We also assume without loss of generality that every positive-strength glue occurring in some tile type in some direction also occurs in some tile type in the opposite direction, i.e., there are no “effectively null” positive-strength glues.

⁴³The restriction on overlap is a model of a chemical phenomenon known as *steric hindrance* [52, Section 5.11] or, particularly when employed as a design tool for intentional prevention of unwanted binding in synthesized molecules, *steric protection* [30–32].