

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 N in less than time $\Omega(N)$, demonstrating that in this case the hierarchical model affords no speedup whatsoever over the seeded model. We also strengthen the $\Omega(N)$ 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 practical implementation of self-assembling molecular tiles was proved experimentally feasible in 1982 by Seeman [39] 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 [37], 1.4% in 2007 [19], 0.13% in 2009 [7], and 0.05% in 2010 [17]. Erik Winfree [45] introduced the abstract Tile Assembly Model (aTAM) – based on a constructive version of Wang tiling [43, 44] – 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 [36] 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, 8, 11–14, 27, 28, 31, 40, 41], tile complexity can be reduced for tasks such as square-building and assembly of more general shapes.

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 [21–23] 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”

*The first author was supported by the Molecular Programming Project under NSF grant 0832824, the second author was supported by a Computing Innovation Fellowship under NSF grant 1019343.

[†]National Taiwan University, Taipei, Taiwan, holinc@gmail.com

[‡]California Institute of Technology, Pasadena, CA, USA, ddoty@caltech.edu

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, 12, 13, 15, 18, 30, 32, 33, 47]. 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 [34], 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. Stated informally, we show a highly parallel (but surprisingly, slow) assembly of an $n \times n$ square in a hierarchical system. We prove that no “partial order hierarchical system” (defined below) can break the $\Omega(n)$ lower bound for assembling any shape of diameter n . Finally, we show that a hierarchical system violating the “partial order” property is able to assemble a rectangle of diameter n in time $o(n)$. We now discuss these results in more detail.

We show that in the hierarchical aTAM, it is indeed possible to assemble an $n \times n$ square using nearly maximal “parallelism,” so that the full $n \times n$ square

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 $O(t^2)$.

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. If one were to assume a constant time for any two producible assemblies to bind once each is produced, this would imply a polylogarithmic time complexity of assembling the final square. But accounting for the effect of assembly *concentrations* on binding rates, our construction takes superlinear time, since 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.

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 $\Omega(N)$ time is required to assemble any structure with diameter N . This implies in particular that the $\Omega(n)$ lower bound in the seeded model for assembling an $n \times n$ square applies to partial order systems in the hierarchical model. To obtain this result, it is necessary to introduce a definition of assembly time applicable to both seeded and hierarchical tile systems. We define such a 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

Thus, for the purpose of speeding up the process of self-assembly, the parallelism of the hierarchical assembly model is of no use whatsoever in partial order systems. 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 [15]. For this reason, our highly parallel square construction may be of independent interest despite the fact that the parallelism does not confer a speedup. Informally, if tile system \mathcal{T} uniquely

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 4.2, 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.

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 3.) 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.³ 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.

Demaine, Demaine, Fekete, Ishaque, Rafalin, Schweller, and Souvaine [12] studied a similar complexity measure 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 individual tile types 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 [12]). In a sense, we are able to “automate” 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 (and naturally paying a price in tile complexity as well, since our system, unlike the staged model, must encode the size n of the square entirely in the tile types). The primary challenge in achieving a highly parallel square construction in the hierarchical model is the prevention of overlapping subassemblies. Adleman [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 $\Omega(n \log n)$ within the same model 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 main 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.

Finally, to show that our partial order hypothesis is necessary to obtain a linear-time lower bound, we construct 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.

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 [36] for a formal definition of the seeded aTAM and [6, 15] for a formal definition of the hierarchical 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

³Without 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 bonds as in Figure 1 to ensure a logarithmic-depth assembly tree.

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 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* [42, Section 5.11] or, particularly when employed as a design tool for intentional prevention of unwanted binding in synthesized molecules, *steric protection* [24–26].

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 $|\mathcal{A}_\square[\mathcal{T}]| = 1$. Given a connected shape $X \subseteq \mathbb{Z}^2$, a TAS \mathcal{T} *strictly self-assembles* S if every producible, terminal assembly places tiles exactly on those positions in X .

3 Nearly Maximally Parallel Hierarchical Assembly of an $n \times n$ 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 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 nearly achieves this bound in the case of assembling an $n \times n$ square S_n , while simultaneously obtaining optimal tile complexity $O(\frac{\log n}{\log \log n})$. 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.

In [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 [38].

THEOREM 3.1. *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 1. 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

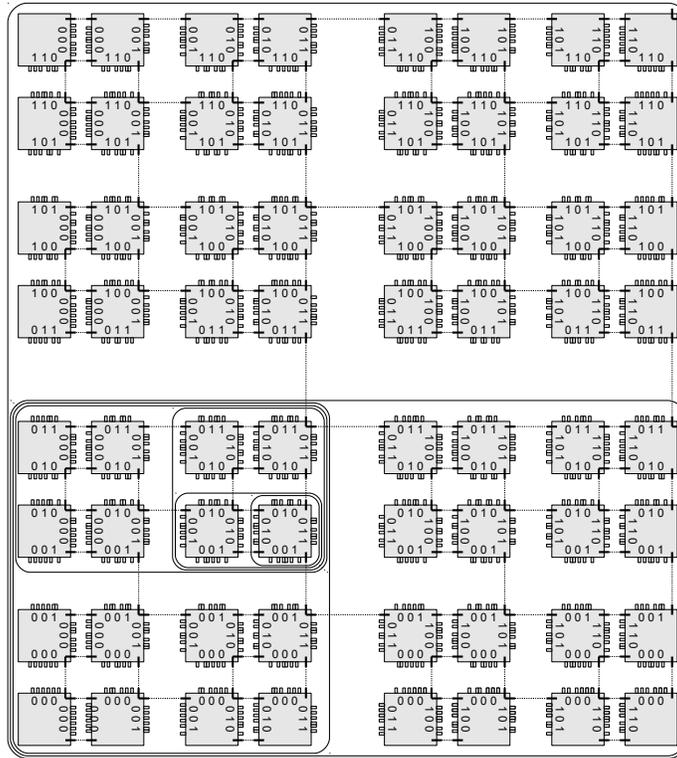


Figure 1: Overview of the hierarchical TAS that assembles an $n \times n$ square with $O(\log^2 n)$ assembly depth and $O(\frac{\log n}{\log \log n})$ tile complexity. Each square in the figure represents a block of width $O(\log n)$ with each side of each block encoding its (x, y) -address in the square. (The encoding scheme is shown in more detail in Figure 3.) Each of the thin solid lines is a strength-1 glue; in fact they are all identical. Dotted lines connect 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.

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 1 outlines the construction of a square when the number of blocks n is a power of two. Figure 2 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 1. The coordinates are also used to determine where to place strength-1 glues. The same

For $n \times n$ square, choose k so that $2^k \leq n < 2^{k+1}$;
 $n = 13$; $k = 3$

$2^k \times 2^k = 8 \times 8$ total blocks used

first $n - 2^k = 5$ blocks on either dimension doubled in size along that dimension

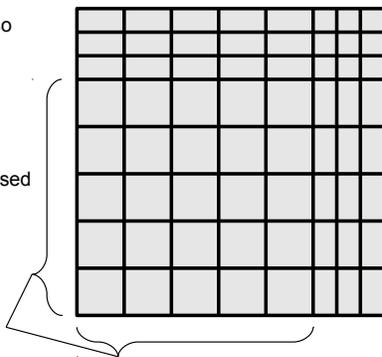


Figure 2: Design of block sizes to handle values of n that are not a power of two. There are always exactly $2^k \times 2^k$ blocks, where $2^k \leq n < 2^{k+1}$. Each block doubles the length of its along the x -axis (resp. y -axis) if $n - 2^k$ exceeds its x -coordinate (resp. y -coordinate).

strength-1 glue is used uniformly throughout Figure 1. 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 3. 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 [47]. 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 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” 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 2. The orange (medium darkness in gray-scale) tile types and the base-conversion tile types that attach to them in Figure 3 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 3 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.⁴ 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 3 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 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 1. 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

⁴The 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.

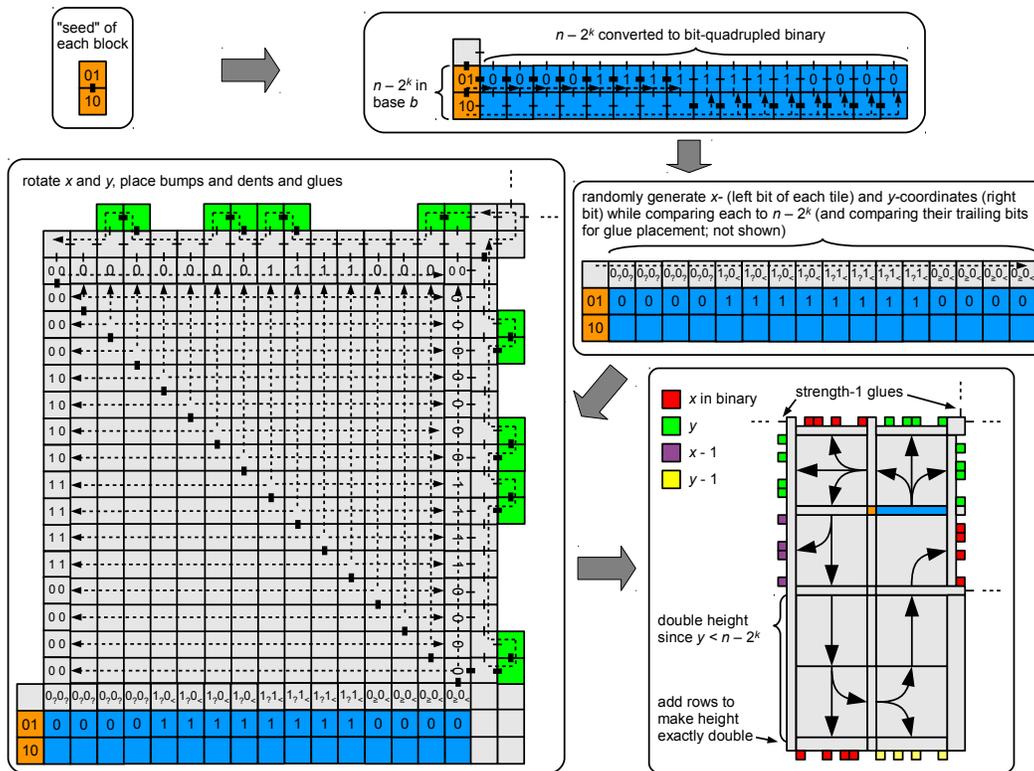


Figure 3: Assembly of $O(\log n) \times O(\log n)$ size block from $O(\frac{\log n}{\log \log n})$ tile types. Every block starts from the same tile types that encode $n - 2^k$, using exactly k bits; in this example, $n = 22$ so $k = 4$ and $n - 2^k$ is 0110 in binary using 4 bits. Thick solid lines 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 3.1 to determine where to place strength-1 glues.

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 3, $y < n - 2^k$ 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 3. 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 3) or clockwise (in the case of x odd, not shown but the exact mirror image of the bottom right of Figure 3), 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 1, 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 3, 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

4 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 5.1, 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 6.1, 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 6.1 may owe its truth entirely to the heavily exploited assumption of equally fast diffusion of all assemblies. A reasonable interpretation of Theorem 6.1 is that the partial order assumption is necessary to prove Theorem 5.1 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 6, may prove useful in other contexts, even though their attained speedup is modest.

4.1 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 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 $\alpha \rightarrow_1 \beta$

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 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.

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. Let $\mathbf{D}(\mathcal{T}, C, d)$ be the random variable representing the time that any tile is first placed at distance d (in the L_1 norm) from the seed, conditioned on the event that a tile is eventually placed at least distance d from the seed, with $\mathbf{D}(\mathcal{T}, C, d) = \infty$ if all producible, terminal assemblies of \mathcal{T} are completely contained in radius $d - 1$ around the seed.

THEOREM 4.1. *For each $d \in \mathbb{Z}^+$, each singly-seeded TAS \mathcal{T} , and each concentrations function $C : T \rightarrow$*

⁶That is, the expected time until the next attachment of a tile to α is an exponential random variable with rate $\sum_{p \in \partial\alpha} C(\hat{\alpha}(p))$, where $\partial\alpha$ is frontier of α , the set of empty locations at which a tile could stably attach to α . The probability that a particular $p' \in \partial\alpha$ leading to assembly β is the next location of attachment is $C(\hat{\alpha}(p')) / \sum_{p \in \partial\alpha} C(\hat{\alpha}(p))$.

⁷When our goal is to obtain only an asymptotic result concerning a family of tile systems assembling a family of assemblies of size/diameter N , we may relax the finite density constraint to the requirement that the concentrations sum to a constant $c \in \mathbb{R}^+$ independent of N , 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.

$[0, 1]$, $\mathbb{E}[\mathbf{D}(\mathcal{T}, C, d)] = \Omega(d)$.

Proof. 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 $\mathbb{R}^+ = [0, \infty)$. Given $i \in \{0, \dots, d\}$ and $t \in \mathbb{R}^+$, 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 $t \in \mathbb{R}^+$, 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 $3\mathbf{X}_{i-1}(t) + 2\mathbf{X}_{i+1}(t)$ 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 $f_i : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ for all $t \in \mathbb{R}^+$ by $f_i(t) = \mathbb{E}[\mathbf{X}_i(t)]$. Then for $i \in \{1, \dots, d - 1\}$,

$$\frac{df_i(t)}{dt} \leq 3f_{i-1}(t) + 2f_{i+1}(t), \text{ and } \frac{df_d(t)}{dt} \leq 3f_{d-1}(t).$$

The lack of a $2f_{d+1}(t)$ 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 . Since the assembly process always starts with a single seed tile, $f_0(t) = 1$ for all $t \in \mathbb{R}^+$, and $f_i(0) = 0$ for all $i \in \{0, \dots, d\}$. For all $t \in \mathbb{R}^+$ 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 $t_0 \in \mathbb{R}^+$ be the unique time at which $f_d(t_0) = \frac{1}{2}$. This time is unique since f_d is monotonically increasing. Since $f_d(t_0) = \mathbb{E}[\mathbf{X}_d(t_0)]$, 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, $E[\mathbf{D}(\mathcal{T}, C, d)] > \frac{t_0}{2}$. 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 $g_i : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ (which will serve as an upper bound for f_i) as follows. For all $1 \in \{1, \dots, d-1\}$,

$$\frac{dg_i(t)}{dt} = 3g_{i-1}(t) + 2g_{i+1}(t), \text{ when } g_i(t) < 4d,$$

$$\frac{dg_d(t)}{dt} = 3g_{d-1}(t), \text{ when } g_d(t) < 4d,$$

and for all $i \in \{0, \dots, d\}$,

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

with the boundary conditions $g_0(t) = 4d$ for all $t \in \mathbb{R}^+$, $g_i(0) = 0$ for all $i \in \{1, \dots, d\}$. It is routine to check that $g_i(t) \geq f_i(t)$ for all $i \in \{0, \dots, d\}$ and $t \in \mathbb{R}^+$. Furthermore, if $g_i(t_0) > g_{i+1}(t_0)$ for all $i \in \{0, \dots, d\}$ and a specific time point $t_0 \in \mathbb{R}^+$, 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, we know that $g_i(t) \geq g_{i+1}(t)$ for all $i \in \{0, \dots, d\}$ and all $t \in \mathbb{R}^+$. Using this monotonicity, 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 $h_0(t) = 4d$ for all $t \in \mathbb{R}^+$, $h_i(0) = 0$ for all $i \in \{0, \dots, d\}$. Solving these differential equations, we obtain $h_d(t) = \frac{4d}{d!} (5t)^d$. 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

$$\begin{aligned} f_d(t') &\leq g_d(t') \leq h_d(t') = \frac{4d}{d!} \cdot (5t')^d = \frac{4d}{d!} \cdot \left(\frac{d}{2e}\right)^d \\ &< \frac{4d}{\left(\frac{d}{e}\right)^d} \cdot \left(\frac{d}{2e}\right)^d = \frac{4d}{2^d}. \end{aligned}$$

Since f_d is monotonically increasing, $f_d(t_0) = \frac{1}{2}$ by definition, and $\frac{4d}{2^d} < \frac{1}{2}$ for sufficiently large d , this implies that $t_0 > t' = \frac{d}{10e}$. \square

4.2 Definition of Time Complexity of Hierarchical Tile Systems

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 4.1. 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 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 [16]), but only those reactions with the assembly we are holding move us "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

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 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, our definition of assembly time coincides with that of [3] on seeded systems.

model such as Gillespie’s algorithm [21–23] 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 [16] 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 resolves 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. 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 stochastic model similar to that of [3]. However, rather than fixing transition rates at each time $t \in \mathbb{R}^+$ as constant, 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 tile 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.¹¹

¹⁰This problem is easily averted in a seeded system by setting the seed concentration sufficiently low (but still $\Omega(|T|)$) 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.

¹¹However, 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

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 $\mathbb{R}^+ = [0, \infty)$, and let $t \in \mathbb{R}^+$. 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.¹² 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 $t \in \mathbb{R}^+$ 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.¹³

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

$$(4.1) \quad \begin{aligned} \frac{d[\alpha](t)}{dt} &= \sum_{i=1}^m [\beta'_i](t) \cdot [\gamma'_i](t) + \sum_{i=1}^p \frac{1}{2} \cdot [\beta''_i](t)^2 \\ &- \sum_{i=1}^n [\alpha](t) \cdot [\beta_i](t), \end{aligned}$$

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

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.

¹²More 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.

¹³That 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.

various reactions involving α determine 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.¹⁴ 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.¹⁵ The choice of which tile type to pick will

¹⁴For 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 $n \rightarrow \infty$.

¹⁵For our lower bound result, Theorem 5.1, 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 5.1 applies when s is selected as the timekeeper. The upper bound of

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 $[\beta](t)$.¹⁶ Unlike the seeded model, the transition rates vary over time since the assemblies (including assemblies that are individual tiles) with which s could 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 $\mathbf{T}(\mathcal{T}, C, s) = \mathbf{E}[\mathbf{T}_{\mathcal{T},C,s}]$. It 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.

We note in particular that our construction of Theorem 3.1 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 $t \in \mathbb{R}^+$. This implies that the rate of the slowest such reaction is at most $\frac{\log^2 n}{n^2}$. 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

Theorem 6.1 holds with respect to only a single tile type, which will appear at a unique position in the final assembly.

¹⁶That 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).

this formally requires more details that we omit, but the fact that the system is a partial order system as defined in Section 5.1 means that Theorem 5.1 implies a time complexity lower bound of $\Omega(n)$.

5 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 N in faster than time $\Omega(N)$.

5.1 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 $\hat{\alpha} \in \mathcal{A}_{\square}[\mathcal{T}]$. 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 p , and $\alpha_k = \hat{\alpha}$.¹⁷ 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 always arrive at the same time in every assembly sequence starting at p .)
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).

¹⁷That 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 .

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 .¹⁸ It is these equivalence classes of positions that are related under \prec . Note that the assembly α_1 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 . 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 $\text{dom } \beta_i \prec \text{dom } \beta_j$. Given a tile type $s \in T$, we say that \mathcal{T} is a *hierarchical partial order system with respect to s* if, for all $p \in \text{dom } \hat{\alpha}$ such that $\hat{\alpha}(p) = s$, \mathcal{T} is a hierarchical partial order system with respect to p , and furthermore if there is no producible assembly β with the property that β can attach to a producible assembly α containing p in two different overlapping positions. The latter is a technical condition required for the proof of Theorem 5.1 that we believe holds for any partial order system, but we have been unable to prove this.

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]. It is routine to check that the TAS described in Section 3 is a partial order system, using one of the orange tiles in the upper-left of Figure 3 as a "seed."

5.2 Linear Time Lower Bound for Partial Order Systems Theorem 5.1 establishes that hierarchical partial order systems, like their seeded counterparts, cannot assemble a shape of diameter N in less than $\Omega(N)$ time. Intuitively, this is proven by using the fact

¹⁸More 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 $\text{dom } \alpha \cap X \implies X \subseteq \text{dom } \alpha$.

that an attaching assembly of size K , which is able to increase the size of the growing assembly by K tiles in a single step, 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 5.1) – over all assemblies that could extend the growing assembly since the growing assembly may be extended at a single attachment site by more than one assembly, but the more assemblies that could extend it, the less the total concentration of the assemblies that could attach to create the extension. 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.¹⁹ 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 5.1. Note that it applies to any hierarchical system.

LEMMA 5.1. (Conservation of Mass Lemma) *Let $\mathcal{T} = (T, \tau)$ be a hierarchical TAS and let $C : T \rightarrow [0, 1]$ be a concentrations function. Then for all $t \in \mathbb{R}^+$,*

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

Proof. For all $t \in \mathbb{R}^+$, 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 $t \in \mathbb{R}^+$, define $f_\alpha(t) = [\alpha](t) \cdot |\alpha|$. Then by equation (4.1), 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

¹⁹The 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 6.1.

their dependence on α , we have

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

Then

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

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

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

□

THEOREM 5.1. *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 N . Then for all concentration functions $C : T \rightarrow [0, 1]$, $\mathbb{T}(\mathcal{T}, C, s) = \Omega(N)$.*

Proof. 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 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 N , which must exist since the diameter of $\text{dom } \hat{\alpha}$ is N . Without loss of generality, the distance d from p to q in $\text{dom } \hat{\alpha}$ is at least $N/2$. 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 the sum of these diameters for all α_i in P is itself an upper bound on the distance from p to q . Therefore $w(P) \geq N/2$.

Let \mathbf{T}_P be the random variable representing the time taken for the complete path P to form. 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 $t \in \mathbb{R}^+$, 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 $t \in \mathbb{R}^+$, $\mathbb{E}[\mathbf{L}(t)] \leq t$. Define the function $f : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ for all $t \in \mathbb{R}^+$ 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 . The set of reactions that could grow P is what matters, but since we have assumed in the definition of partial order systems that no assembly extending P' could attach via two different reactions that both intersect P at the location directly succeeding P' , summing over assemblies is equivalent to summing over reactions. Our argument uses the conservation of mass property (Lemma 5.1) 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 5.1, it follows that for all $t \in \mathbb{R}^+$,

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

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

Since $\mathbb{E}[\mathbf{L}(t)] \leq t$, 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, $\mathbb{E}[\mathbf{T}_P] \geq w(P)/4 = \Omega(w(P)) = \Omega(N)$, which proves the theorem. \square

Although our assembly time model describes concentrations as evolving according to the standard mass-action kinetic differential equations, Lemma 5.1 is the only property of this model that is required for our proof of Theorem 5.1. 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 5.1 at all times, Theorem 5.1 still holds. For example, staged assembly [12] is a relaxation of the mass-action model that obeys Lemma 5.1, in which certain assemblies are artificially prevented from interacting by being kept in separate bins before being mixed. Theorem 5.1 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.

6 Assembly of a Shape in Time Sublinear in its Diameter

This section states and outlines the ideas of the proof of 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 6.1. *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 $s \in T$ and concentrations function $C : T \rightarrow [0, 1]$ such that $\Upsilon(\mathcal{T}, C, s) = O(n^{4/5} \log n)$.*

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 4, we interpret the upper bound of Theorem 6.1 more cautiously than the lower bound of Theorem 5.1, 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 6.1 is that concentration arguments alone do not suffice to show a linear-time lower bound on assembly time in the hierarchical model.

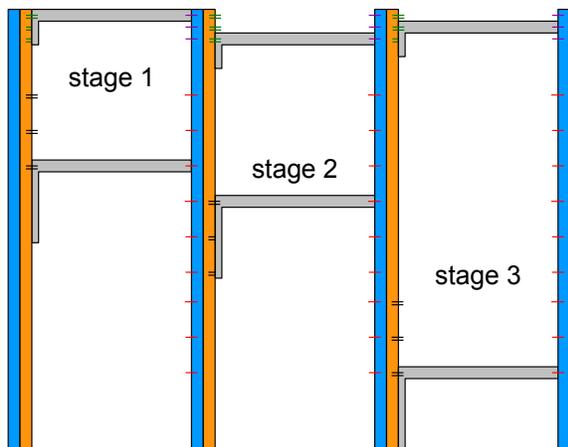


Figure 4: High-level overview of interaction of “vertical bars” and “horizontal bars” to create the rectangle of Theorem 6.1. Filler tiles fill in the empty regions. If glues overlap two regions then represent a formed bond. If glues overlap one region but not another, they are glues from the former region but are mismatched (and thus “covered”) by the latter region. This is intended to ensure that only one horizontal bar can attach in each of the top and bottom regions of a vertical bar.

A high-level overview of the construction is shown in Figure 4. The interaction of the components of this figure are shown in more detail in Figure 5. Let $k = m = n^{1/5}$, and let $w = n^{4/5}$. The rectangle grows rightward in m “stages”, each stage of width w

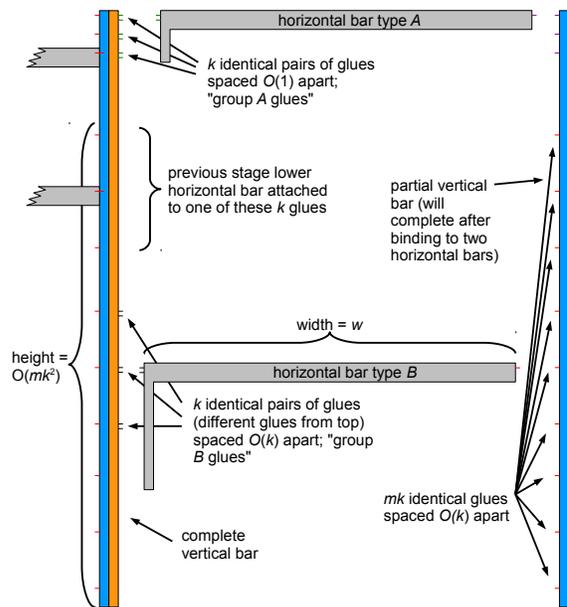


Figure 5: Vertical bars and horizontal bars for the construction of a fast-assembling square. “Type B” horizontal bars have a longer vertical arm than “Type A” since the glues they must block are farther apart.

and height $h = O(mk^2)$.²⁰ Each stage consists of the attachment of two “horizontal bars” to the right, which in turn cooperate to place a single “vertical bar”, which (after some additional growth not shown) will contain glues for binding of horizontal bars in the next stage. 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 4. This “binding parallelism” is in addition to “assembly parallelism”: 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”: k 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 $\frac{t}{k}$).

More precisely, two different versions of β bind to one of two different regions on α , each region having k

²⁰In this section, we use the term “stage” merely to mean different regions of the final assembly. This is a different usage of “stage” than used in [12].

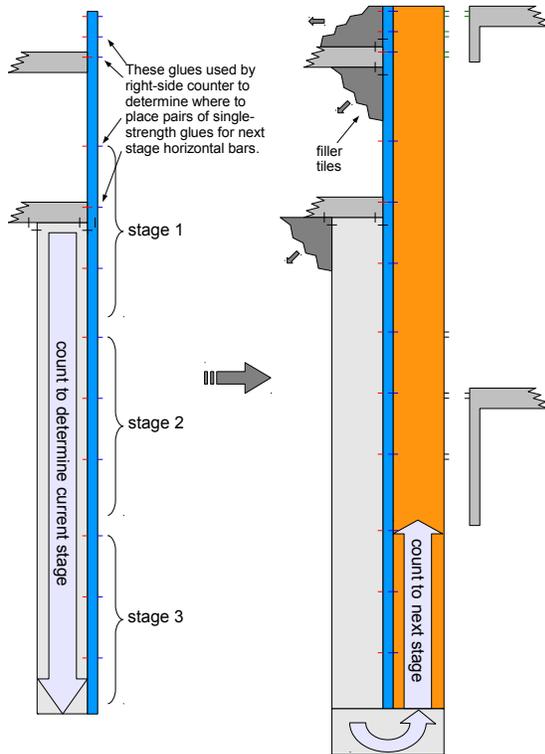


Figure 6: Detailed view of how a single partial vertical bar attaches to complete a full vertical bar after attaching to the previous stage's two horizontal bars. This enables the correct vertical placement of the two horizontal bars in the next stage. This is needed to communicate which stage is the 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.

binding sites. The choice of spacing between binding sites is to ensure that each pair of vertical positions where the two horizontal bars could go are separated by a unique distance. This means that the vertical bar, despite having all possible binding sites available on its left side, will always bind in the same vertical position relative to the vertical bar to its left, ensuring that the *shape* assembled is always a rectangle. These distances are also used to communicate and increment the current stage, encoded in the vertical position of the bottom horizontal bar, since the glues on its right are not specific to the stage. This process is shown in Figure 6. 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.

7 Open Questions

There are some interesting questions that remain open:

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?

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

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

In the case of the seeded aTAM, the seeded variants of these problems are known to be coNP-complete [6, 20, 29] and in P [5], respectively.

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

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

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

In the case of the seeded aTAM, the seeded variants of these problems are known to be Σ_2^{P} -complete [10] 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?

- 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 [9, 35]. 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, [45]). 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 [45, 46] 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 4.1, 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 [9, 35]. We are grateful to anonymous reviewers for identifying important flaws in the original proofs of some theorems.

References

- Zachary Abel, Nadia Benbernou, Mirela Damian, Erik D. Demaine, Martin 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.
- Leonard M. Adleman. Toward a mathematical theory of self-assembly (extended abstract). Technical Report 00-722, University of Southern California, 2000.
- 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.
- 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.
- 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.
- 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.
- Robert D. Barish, Rebecca Schulman, Paul W. 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.
- 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.
- Otto G. Berg and Peter H. von Hippel. Diffusion-controlled macromolecular interactions. *Annual Review of Biophysics and Biophysical Chemistry*, 14(1):131–160, 1985.
- Nathaniel Bryans, Ehsan Chiniforooshan, David Doty, Lila Kari, and Shinnosuke Seki. The power of nondeterminism in self-assembly. In *SODA 2011: Proceedings of*

- the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms, pages 590–602. SIAM, 2011.
- [11] 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.
- [12] 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 13.
- [13] 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.
- [14] David Doty. Randomized self-assembly for exact shapes. *SIAM Journal on Computing*, 39(8):3521–3552, 2010. Preliminary version appeared in FOCS 2009.
- [15] 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.
- [16] Irvin R. Epstein and John A. Pojman. *An introduction to nonlinear chemical dynamics: oscillations, waves, patterns, and chaos*. Oxford University Press, USA, 1998.
- [17] Constantine Evans. Personal communication.
- [18] Bin Fu, Matthew J. Patitz, Robert T. Schweller, and Robert Sheline. Self-assembly with geometric tiles. Technical Report 1104.2809, Computing Research Repository, 2011.
- [19] 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.
- [20] Michael R. Garey and David S. Johnson. *Computers and Intractability*. W. H. Freeman, New York, 1979.
- [21] 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.
- [22] Daniel T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *Journal of Physical Chemistry*, 81(25):2340–2361, 1977.
- [23] Daniel T. Gillespie. Stochastic simulation of chemical kinetics. *Annual review of physical chemistry*, 58(1):35–55, 2007.
- [24] 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.
- [25] 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.
- [26] 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.
- [27] 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.
- [28] Ming-Yang Kao and Robert T. Schweller. Randomized self-assembly for approximate shapes. In Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Haldrsson, Anna Ingólfssdóttir, and Igor Walukiewicz, editors, *ICALP 2008: International Colloquium on Automata, Languages, and Programming*, volume 5125 of *Lecture Notes in Computer Science*, pages 370–384. Springer, 2008.
- [29] Harry R. Lewis and Christos H. Papadimitriou. *Elements of the Theory of Computation*. Prentice-Hall, Inc., Upper Saddle River, New Jersey, 1998.
- [30] 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.
- [31] 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.
- [32] 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. Scientetchnica, 2011.
- [33] Matthew J. Patitz and Scott M. Summers. Identifying shapes using self-assembly. *Algorithmica*. to appear. Preliminary version appeared in ISAAC 2010.
- [34] 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.
- [35] 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.
- [36] Paul W. K. Rothmund and Erik Winfree. The program-size complexity of self-assembled squares (ex-

- tended abstract). In *STOC 2000: Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing*, pages 459–468, 2000.
- [37] Paul W.K. Rothmund, Nick Papadakis, and Erik Winfree. Algorithmic self-assembly of DNA Sierpinski triangles. *PLoS Biology*, 2(12):2041–2053, 2004.
 - [38] Robert Schweller. Personal communication.
 - [39] Nadrian C. Seeman. Nucleic-acid junctions and lattices. *Journal of Theoretical Biology*, 99:237–247, 1982.
 - [40] David Soloveichik and Erik Winfree. Complexity of self-assembled shapes. *SIAM Journal on Computing*, 36(6):1544–1569, 2007. Preliminary version appeared in DNA 10.
 - [41] Scott M. Summers. Reducing tile complexity for the self-assembly of scaled shapes through temperature programming. *Algorithmica*. to appear.
 - [42] Leroy G. Wade. *Organic Chemistry*. Prentice Hall, 2nd edition, 1991.
 - [43] Hao Wang. Proving theorems by pattern recognition – II. *The Bell System Technical Journal*, XL(1):1–41, 1961.
 - [44] 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.
 - [45] Erik Winfree. *Algorithmic Self-Assembly of DNA*. PhD thesis, California Institute of Technology, June 1998.
 - [46] Erik Winfree. Simulations of computing by self-assembly. Technical Report CaltechCSTR:1998.22, California Institute of Technology, 1998.
 - [47] 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.