Electronic Colloquium on Computational Complexity, Report No. 195 (2010)

Parallelism, Program Size, Time, and Temperature in Self-Assembly*

Ho-Lin Chen† David Doty† Shinnosuke Seki‡ David Soloveichik§

Abstract

We settle a number of questions in 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, showed that both of these parameters are optimal, and 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 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 $n \times n$ 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 leave open the question of whether some hierarchical tile system can break the $\Omega(n)$ assembly time lower bound for assembling an $n \times n$ 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.

We then investigate the relationship between the temperature of a tile system and its size. We show that a tile system can in general require temperature that is exponentially greater than its number of tile types. On the other hand, for the special case of *2-cooperative* systems, in which all binding events involve at most 2 sides of tiles, it suffices to use temperature linear in the number of tile types. We show that there is a polynomial-time algorithm that, given any tile system $\mathcal{T}$ specified by its desired binding behavior, finds a temperature and binding energies (at most exponential in the number of tile types of $\mathcal{T}$) that realize this behavior or reports that no such energies exist. This result is applied to show that there is a polynomial-time algorithm that, given an $n \times n$ square S_n, determines the smallest (non-hierarchical "seeded") system (at any temperature) that is deterministic and self-assembles S_n. This answers an open question of Adleman, Cheng, Goel, Huang, Kempe, Moisset de Espanés, and Rothemund (*Combinatorial Optimization Problems in Self-Assembly*, STOC 2002).

*The first author was supported by the Molecular Programming Project under NSF grant 0832824, the second and fourth authors were supported by NSF Computing Innovation Fellowships, and the third author was supported by NSERC Discovery Grant R2824A01 and the Canada Research Chair in Biocomputing to Lila Kari.

†California Institute of Technology, Pasadena, CA, USA, holinc@gmail.com, ddoty@caltech.edu

‡University of Western Ontario, Dept. of Computer Science, London, ON, Canada, N6A 5B7, sseki@csd.uwo.ca.

§University of Washington, Dept. of Computer Science, Seattle, WA, USA, dsolov@u.washington.edu.

ISSN 1433-8092

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 [34] 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 [32], 1.4% in 2007 [17], 0.13% in 2009 [7], and 0.05% in 2010 [16]. Erik Winfree [41] introduced the abstract Tile Assembly Model (aTAM) – based on a constructive version of Wang tiling [39, 40] – 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, Rothemund and Winfree [31] 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 [1,6,8,10–13,25,26,29,36,37], 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 on the standard stochastic model of chemical kinetics [19–21], and show that under this model, an $n \times n$ square can be assembled in expected time $O(n)$, which is asymptotically optimal, in addition to having optimal tile complexity $O(\frac{\log n}{\log \log n})$. Intuitively, the optimality of the $O(n)$ assembly time for an $n \times n$ square results from the following informal description of self-assembly. The standard "seeded" aTAM stipulates that one tile type (or pre-fabricated assembly of tiles) 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 all 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 assemble the entire square, which is precisely what is achieved in [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,

[1] For 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 $O(t^2)$.

allows this assembly to then aggregate to other assemblies, and possibly (depending on the model) dispenses completely with the idea of a seed assembly. Variants of the hierarchical aTAM have recently received extensive theoretical study [1, 2, 4, 6, 11, 12, 14, 28, 30, 42]. It is 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. Breaking the $\Omega(n)$ time lower bound for uniquely self-assembling an $n \times n$ square using in the hierarchical aTAM was stated as an open question in [3]. We show that in the hierarchical aTAM, it is indeed possible to assemble an $n \times n$ square S_n using nearly maximal "parallelism," so that the full $n \times n$ square is formed from four $n/2 \times n/2$ sub-squares, which are themselves each formed from four $n/4 \times n/4$ sub-squares, etc. If one were to assume a constant concentration of all producible assemblies, this would imply a polylogarithmic time complexity of assembling the final square. However, taking concentrations into account, our construction takes superlinear time to assemble the square, since some sub-square has concentration at most $\widetilde{O}(1/n^2)$, so the time for even a single step of hierarchical assembly is at least $\widetilde{\Omega}(n^2)$ by standard models of chemical kinetics.

We leave open the question of [3], although we do achieve a partial result toward showing a lower bound. 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 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 [14]. 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, define $\text{depth}^{\text{da}}(S)$ 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 any tile system to uniquely assemble shape S. (A formal definition is given in Section 3.) For any shape S with N points, clearly $\text{depth}^{\text{da}}(S) \geq \log N$. Our construction nearly achieves this bound in the case of assembling an $n \times n$ square S_n, showing that $\text{depth}^{\text{da}}(S_n) \leq O(\log^2 n)$. Furthermore, this is achievable using $O(\frac{\log n}{\log \log n})$ tile types, which is asymptotically optimal. 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 [11] 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

[2]Low 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 coincide on seeded systems.

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 [11]). 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 $t_1, \ldots, t_n$. The main intuitive reason that the time is not $O(\log n)$ is that if assemblies $\alpha_1 = t_i \ldots t_j$ and $\alpha_2 = t_{i'} \ldots t_{j'}$ 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.

We also investigate the relationship between the temperature τ of a tile system and its number of tile types. This is partially motivated by the following. The algorithm of Adleman, Cheng, Goel, Huang, Kempe, Moisset de Espanés, and Rothemund [5] for finding a minimum size temperature-2 tile system to assemble an $n \times n$ square works by brute-force search over the set of all temperature-2 tile systems with at most $O(\frac{\log n}{\log \log n})$ tile types, using the fact proven by Adleman, Cheng, Goel, and Huang [3] that such an upper bound on tile complexity suffices to assemble any $n \times n$ square. A simple counting argument shows that for any constant τ, the number of tile systems with glue strengths and temperature at most τ and $O(\frac{\log n}{\log \log n})$ tile types is polynomial in n. One conceivable approach to extending the algorithm to arbitrary temperature is to prove that for any tile system with K tile types, the strengths and temperature can be re-assigned so that they are upper-bounded by a constant or slow-growing function of K, without affecting the behavior of the tile system. However, we show that this approach cannot work, by demonstrating that for each K, there is a tile system with K tile types whose behavior cannot be preserved using any temperature less than $2^{K/4}$. The proof crucially uses *3-cooperative* binding, meaning attachment events that require three different glues of a tile to match the assembly. On the other hand, we show that any 2-cooperative tile system with K tile types requires temperature at most $2K + 2$.

Of course, integer temperatures and glue strengths are an artifact of the model, since actual temperatures and binding energies do not necessarily come in discrete quantities. Nonetheless, our investigation does reflect fundamental questions about how finely divided molecular binding energies must be in a real molecular self-assembly system. The requirement of integer strengths is simply one way of "quantizing" the minimum distinction we are willing to make between energies and then re-scaling so that this quantity is normalized to 1.[3] Our 3-cooperative lower bound therefore shows that in general, certain self-assembling systems that have very large gaps between some of their binding energies nonetheless require other binding energies to be extremely close (exponentially small in terms of the larger gaps) and yet still unequal. This can be interpreted as an infeasibility result if one defines "exponentially fine control" of binding energies as "infeasible"

[3]Indeed, our proof does not require that strengths be integer, merely that the distance between the smallest energy strong enough to bind and the largest energy too weak to bind be at least 1.

to execute in any real laboratory, since no implementation of the specified tile behavior can use courser energies.

As a converse to the temperature lower bound stated above, we show that there is an algorithm that, given any tile system $\mathcal{T}$ with K tile types specified by its desired binding behavior, finds a temperature and glue strengths at most $2^{O(K^2)}$ that realize this behavior or reports that no such strengths exist. This result is used to show that there is a polynomial-time algorithm that, given an $n \times n$ square S_n, determines the smallest tile assembly system (at any temperature) that is deterministic and strictly self-assembles S_n, answering an open question of Adleman, Cheng, Goel, Huang, Kempe, Moisset de Espanés, and Rothemund [5], who demonstrated such an algorithm for the special case of temperature 2 systems.

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 [31] for a formal definition of the standard aTAM and [6,14] for a formal definition of the hierarchical aTAM.

A *tile type* is a unit square with four sides, each having a *glue label* (often represented as a finite string). 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 $\mathbb{Z}^2 \dashrightarrow T$. For a set of tile types T, let $\Lambda(T)$ denote the set of all glue labels of tile types in T. A *strength function* is a function $g : \Lambda(T) \to \mathbb{N}$ indicating, for each glue label ℓ, the strength $g(\ell)$ with which it binds. Two adjacent tiles in an assembly *interact* if the glue labels on their abutting sides are equal and have positive strength according to g. Each assembly induces a *binding graph*, a grid graph whose vertices are tiles, with an edge between two tiles if they interact. The assembly is *τ-stable* if every cut of its binding graph has strength at least τ, where the weight of an edge is the strength of the glue it represents. That is, the assembly is stable if at least energy τ is required to separate the assembly into two parts.

A *seeded tile assembly system* (seeded TAS) is a quadruple $\mathcal{T} = (T, \sigma, g, \tau)$, where T is a finite set of tile types, $\sigma : \mathbb{Z}^2 \dashrightarrow T$ is a finite, τ-stable *seed assembly*, $g : \Lambda(T) \to \mathbb{N}$ is the *strength function*, and τ is the *temperature*. Given a seeded TAS $\mathcal{T} = (T, \sigma, g, \tau)$, an assembly α is *producible* if either $\alpha = \sigma$ or if β is a producible assembly and α can be obtained from β by placing a single tile type t on empty space (a position $p \in \mathbb{Z}^2$ such that $\beta(p)$ is undefined), such that the resulting assembly α is τ-stable. In this case write $\beta \to_1 \alpha$, and write $\beta \to \alpha$ if $\beta \to_1^* \alpha$. An assembly is *terminal* if no tile can be τ-stably attached to it.

A *hierarchical tile assembly system* (hierarchical TAS) is a triple $\mathcal{T} = (T, g, \tau)$. Given a hierarchical TAS $\mathcal{T} = (T, g, \tau)$, 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* [38, Section 5.11] or, particularly when employed as a design tool for intentional prevention of unwanted binding in synthesized molecules, *steric protection* [22–24].

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 Hierarchical Self-Assembly Parallelism and Time

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 then show that, somewhat counterintuitively, this construction is *slower* than the construction of [3], which assembles an $n \times n$ square in time $O(n)$ in the seeded aTAM. The authors of [3] define a class of seeded TAS's called *partial order systems*, whose definition we extend in a natural way to hierarchical systems. We show, using basic principles of chemical kinetics, that *no* partial order system can assemble any shape of size N in faster than $\Omega(\sqrt{N})$ time. Thus hierarchical partial order systems achieve no speedup at all over seeded systems.

3.1 Nearly Maximally Parallel Hierarchical Assembly of an $n \times n$ Square with Optimal Tile Complexity

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

Let $\mathcal{T} = (T, g, \tau)$ be a directed hierarchical TAS. Let $\alpha \in \mathcal{A}[\mathcal{T}]$ be a producible assembly. An *assembly tree* A of α is a binary tree whose nodes are labeled by producible assemblies, with α labeling the root, individual tile types labeling the $|\text{dom } \alpha|$ leaves, and node u having children u_1 and u_2 with the requirement that u_1 and u_2 can attach to assemble v. That is, A 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 $\text{depth}^{\text{da}}(\mathcal{T}) = \max_{A \in \Upsilon_{\square}(\mathcal{T})} \text{depth}(A)$, where $\text{depth}(A)$ denotes the standard depth of the tree A, the length of the longest path from any leaf to the root. Let $S \subset \mathbb{Z}^2$ be a finite shape. The *directed assembly depth* of S is

$$\text{depth}^{\text{da}}(S) = \min \left\{ \text{depth}^{\text{da}}(\mathcal{T}) \;\middle|\; \begin{array}{l} \mathcal{T} = (T, g, \tau) \text{ is a directed hierarchical} \\ \text{TAS and } \mathcal{T} \text{ strictly self-assembles } S \end{array} \right\}.$$

It is clear by the definition that for any shape S with N points, $\text{depth}^{\text{da}}(S) \geq \log N$. Our construction nearly achieves this bound in the case of assembling an $n \times n$ square S_n, showing that $\text{depth}^{\text{da}}(S_n) \leq O(\log^2 n)$. 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.

Furthermore, we achieve nearly optimal tile complexity in additional to optimal assembly depth. 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 [33].

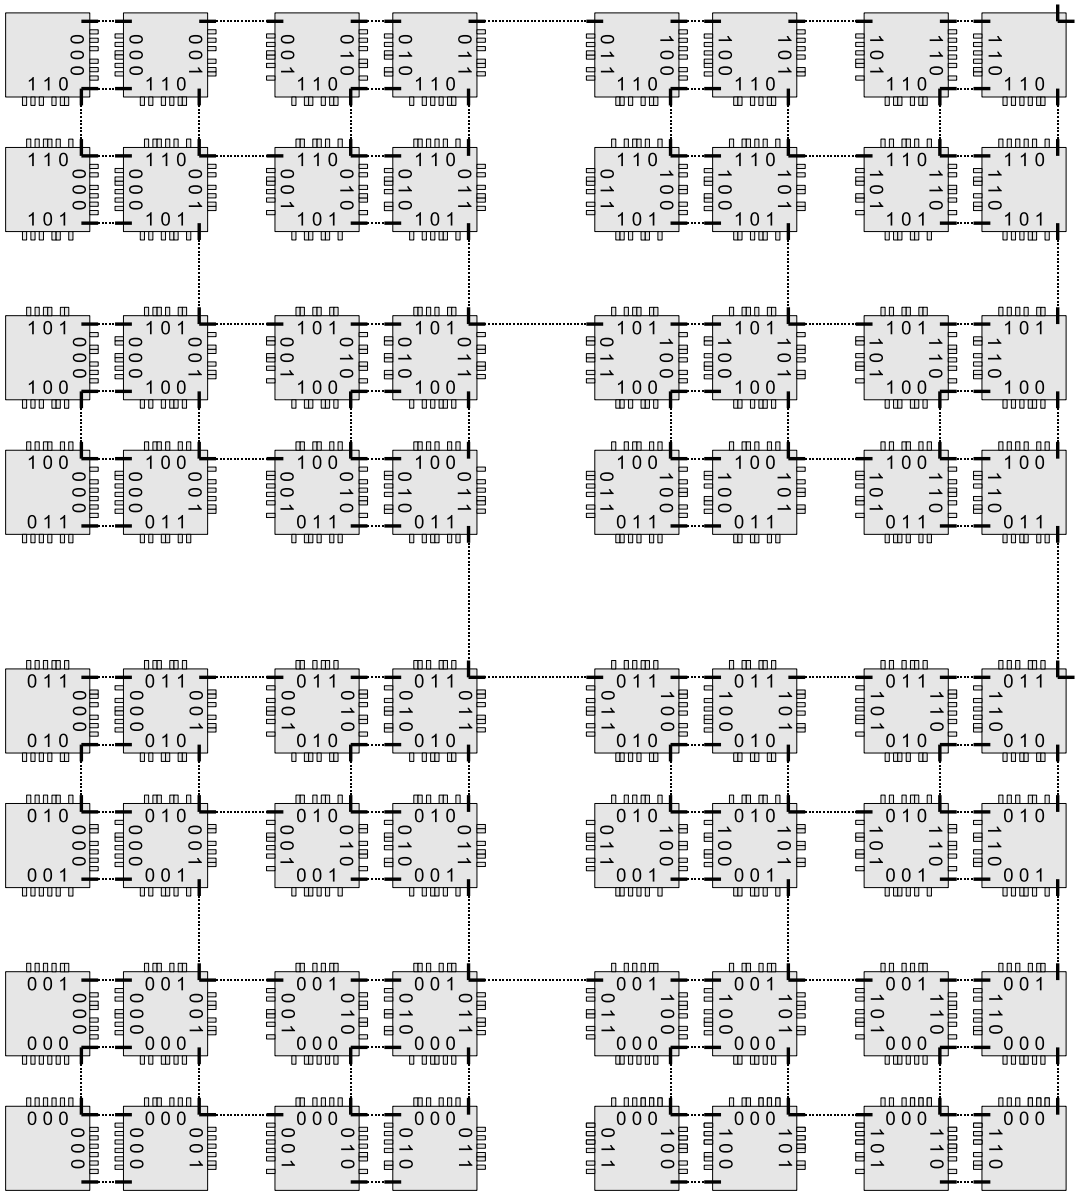

Figure 1: Overview of the hierarchical TAS that assembles an $n \times n$ square with $O(\log^2 n)$ assembly depth and $O(\frac{\log n}{\log \log n})$ tile complexity. Each square in the figure represents a block of width $O(\log n)$ with the all sides of each block encoding its (x, y)-address in the 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 supposed to bind to each other.

Theorem 3.1. *There are constants $c_1, c_2 \in \mathbb{N}$ such that, for all $n \in \mathbb{N}$, there is a directed hierarchical TAS $\mathcal{T} = (T, g, 2)$ such that $\mathcal{T}$ strictly self-assembles the $n \times n$ square, $|T| \leq c_1 \log n / \log \log n$, and $\text{depth}^{\text{da}}(\mathcal{T}) \leq c_2 \log^2 n$.*

Proof. A high-level outline of the construction is shown in Figure 1. We assemble a number of

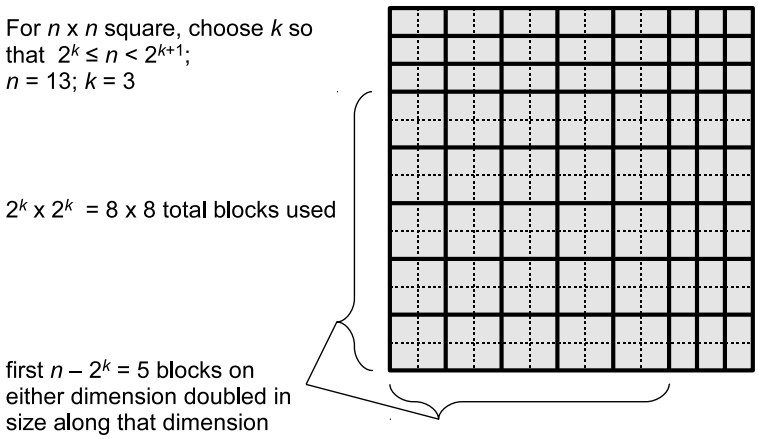

Figure 2: Design of block sizes to handle values of n that are not a power of two. There are always exactly $2^k \times 2^k$ blocks, where $2^k \leq n < 2^{k+1}$. Each block doubles the length of its along the x-axis (resp. y-axis) if $n - 2^k$ exceeds its x-coordinate (resp. y-coordinate). The solid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

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

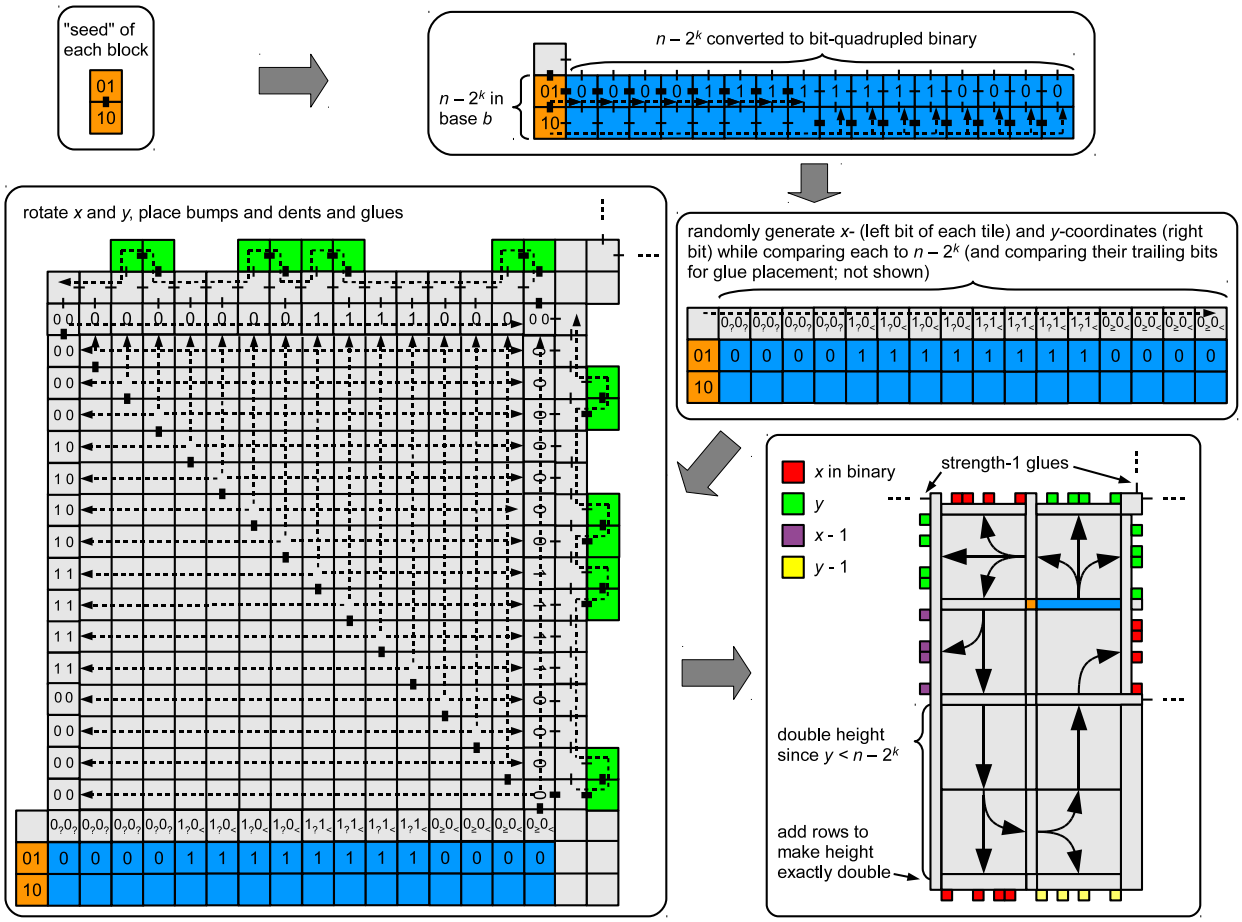

Figure 3: Assembly of $O(\log n) \times O(\log n)$ size block from $O(\frac{\log n}{\log \log n})$ tile types. Every block starts from the same tile types that encode $n - 2^k$, using exactly k bits; in this example, $n = 22$ so $k = 4$ and $n - 2^k$ is 0110 in binary using 4 bits. Thick solid lines 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). A constant-size set of tile types does the rest. $n - 2^k$ is first converted to binary and each of its bits quadrupled to make room for the bumps and dents. 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.

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 [42]. 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 "position-unique"; no strength-2 glue is shared between two points in Figure 3, and by inspection 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 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 a constant number of tile types to convert $n - 2^k$ to binary and 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 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 $d \in \{0, 1, \ldots, b-1\}$ 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 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

[4] The 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) 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. □

3.2 Time Complexity in the Hierarchical Model

In this section we define a formal notion of time complexity for tile assembly systems, in order to show that the subset of systems known as partial order systems cannot assemble any shape of diameter N in faster than time $\Omega(N)$. The model we use applies to both the seeded aTAM and the hierarchical aTAM. For hierarchical systems, our model may not be completely suitable since we make some unrealistic assumptions.[5] However, our entire purpose for defining time complexity in the hierarchical aTAM is to show a lower bound, demonstrating that hierarchical partial order systems can proceed no faster than seeded assembly despite the apparent extra parallelism. Our assumptions will have the effect of making hierarchial self-assembly appear *faster* than if we modeled the underlying chemical kinetics more carefully. We show that even with these extra assumptions, the time complexity of hierarchical partial order systems is *still* no better than the seeded aTAM.

3.2.1 Definition of Time Complexity of Directed Seeded Systems

We recall the definition of time complexity of seeded self-assembly proposed in [3]. A *concentrations function* on a tile set T is a probability measure $C : T \to [0,1]$. Each tile type t is assumed to be held at a fixed concentration $C(t)$ throughout the process of assembly.[6] The assembly time for

[5] For example, 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.

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

producing a terminal assembly $\widehat{\alpha}$ in a seeded TAS $\mathcal{T} = (T, \sigma, g, \tau)$ is defined by picking a copy of the seed arbitrarily and taking the expected time before the seed grows into $\widehat{\alpha}$, when assembly proceeds according to the following stochastic model. The assembly of $\widehat{\alpha}$ 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 $\alpha \to_1 \beta$ via the addition of tile type t, there is a transition in the Markov process from state α to state β with transition rate $C(t)$.[7] If $\mathcal{T}$ is directed, then $\widehat{\alpha}$ is the unique sink state of the Markov process. The time to reach $\widehat{\alpha}$ from σ is a random variable $\mathbf{T}_{\mathcal{T},C}$, and the assembly time complexity of $\mathcal{T}$ with concentrations C is defined to be $\mathsf{T}_C(\mathcal{T}) = \mathrm{E}[\mathbf{T}_{\mathcal{T},C}]$.

The requirement that the tile concentrations function C be a probability 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, we obtain that the total concentration of tiles (concentration defined as number or mass per unit volume) cannot exceed 1. By letting their total concentration be exactly 1, we are essentially requiring that the assembly process is carried out in the smallest volume (which leads to the fastest reaction time) required to satisfy the finite density constraint.

First, we have the following time complexity lower bound for seeded systems, which strengthens and implies Lemma 4.6 of the full version of [3], which applies only to directed systems. This theorem says that even for non-directed systems, a seeded TAS can still grow its diameter only linearly. Let $d \in \mathbb{Z}^+$. Let $\mathcal{T} = (T, \sigma, g, \tau)$ be a singly-seeded TAS (meaning $|\mathrm{dom}\ \sigma| = 1$). 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 $\mathbf{D}(\mathcal{T}, d)$ be the random variable representing the time that any tile is first placed at distance d from the seed, conditioned on the event that a tile is eventually placed at least distance d (in the L_1 norm) from the seed, with $\mathbf{D}(\mathcal{T}, d) = \infty$ if all producible, terminal assemblies of $\mathcal{T}$ are completely contained in radius $d-1$ around the seed.

Theorem 3.2. *For each $d \in \mathbb{Z}^+$ and each singly-seeded TAS $\mathcal{T}$, $\mathrm{E}[\mathbf{D}(\mathcal{T}, d)] \geq \Omega(d)$.*

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 $i \in \{0, \ldots, d\}$ 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.

the number of terminal assemblies (if each is of size at most K) will be limited by $C(t_s)$, implying the percentage change in every other tile type t's concentration is at most $K \cdot C(t_s)/C(t)$; therefore "low" seed concentration means setting $C(t_s) \ll C(t)/K$ for all $t \in T \setminus \{t_s\}$. In fact, to obtain an assembly time asymptotically as fast, one need only ensure that for all t, $C(t) \geq 2\#(t)C(t_s)$, where $\#(t)$ is the number of times t appears in the terminal assembly α. This guarantees that the concentration of t 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.

[7] That is, the expected time until the next attachment of a tile to α is an exponential random variable with rate $r = \sum_{p \in \partial\alpha} C(\widehat{\alpha}(p))$, 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 $C(\widehat{\alpha}(p'))/r$.

Then for all $t \in \mathbb{R}^+$, the event $\mathbf{X}_i(t) = 0$ (no tile is at distance d by the time t) is equivalent to the event $\mathbf{D}(\mathcal{T}, d) > t$ (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 single tiles is 1, the rate at which tiles attach at any given location is at most 1. For all $i \in \{0, \ldots, d\}$, define the function $f_i : \mathbb{R}^+ \to \mathbb{R}^+$ for all $t \in \mathbb{R}^+$ by $f_i(t) = \mathrm{E}[\mathbf{X}_i(t)]$. Then for $i \in \{0, \ldots, d-1\}$,

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

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 $i \in \{0, \ldots, d\}$. For all $t \in \mathbb{R}^+$ and all $i \in \{1, \ldots, d\}$, $f_i(t) \leq 4i$ 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 $f_d(t_0) = \mathrm{E}[\mathbf{X}_d(t_0)]$, 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 $\mathbf{X}_d(t_0) = 0 \iff \mathbf{D}(\mathcal{T}, d) > t_0$, whence $\Pr[\mathbf{D}(\mathcal{T}, d) > t_0] > \frac{1}{2}$. By Markov's inequality, $\mathrm{E}[\mathbf{D}(\mathcal{T}, d)] > \frac{t_0}{2}$. 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 $i \in \{0, \ldots, d\}$, define the function $g_i : \mathbb{R}^+ \to \mathbb{R}^+$ (which will serve as an upper bound for f_i) by

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

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

and

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

with the boundary conditions $g_0(t) = 4d$ for all $t \in \mathbb{R}^+$, $g_i(0) = 0$ for all $i \in \{1, \ldots, d\}$. It is routine to check that $g_i(t) \geq f_i(t)$ for all $i \in \{0, \ldots, d\}$ and $t \in \mathbb{R}^+$. Furthermore, if $g_i(t_0) > g_{i+1}(t_0)$ for all $i \in \{0, \ldots, d\}$ 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 $i \in \{0, \ldots, d\}$ by definition, we know that $g_i(t) \geq g_{i+1}(t)$ for all $i \in \{0, \ldots, d\}$ 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.

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

with boundary conditions $h_0(t) = 4d$ for all $t \in \mathbb{R}^+$, $h_i(0) = 0$ for all $i \in \{0, \ldots, d\}$. 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

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

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 $t_0 > t' = \frac{d}{10e}$. □

3.2.2 Definition of Time Complexity of Directed Hierarchical Systems

To define time complexity for hierarchical systems, we employ more explicitly the chemical kinetics implicitly underlying the above time complexity model for seeded systems. We treat each assembly as a single molecule. If two assemblies α and β can attach to create a assembly γ, then we model this as a chemical reaction $\alpha + \beta \to \gamma$, in which the rate constant is assumed to be equal for all reactions (and normalized to 1). Note in particular that 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]

We would like to analyze the assembly time in such as way as to facilitate direct comparison with the results of [3]. With a stochastic model such as Gillespie's algorithm [19–21] using finite tile 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.[9] The mass-action model of kinetics [15] describes concentrations as a dynamical system that evolves continuously over time according to ordinary differential equations derived from reaction rates. In the mass-action model, the terminal assembly is guaranteed to form, which resolves one issue with the purely stochastic model; however, some (low) concentration of the terminal assembly forms in the first infinitesimal amount of time. It is not clear how to define the time to assemble the terminal assembly $\widehat{\alpha}$. Some choices such as measuring the time to half-completion (time required for the concentration of $\widehat{\alpha}$ to exceed half of its steady-state concentration) are 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.

The model of assembly time that we define is a continuous-time stochastic model similar to that of [3], but rather than fixing transition rates at each time $t \in \mathbb{R}^+$ as constant, uses 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 the 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. Note in particular that 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 the single assembly containing the seed. Like the model of [3], we define assembly

[8] The 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.

[9] This problem is easily averted in a seeded system by setting the seed concentration sufficiently low (but still $\Omega(|T|)$). In a hierarchical system it is not clear how to avoid this problem.

time only for directed systems, in order to meaningfully ask the question "How long until terminal assembly $\widehat{\alpha}$ forms?" without having to deal with the possibility that $\widehat{\alpha}$ does not form.

Let $\mathcal{T} = (T, g, \tau)$ be a (possibly non-directed) hierarchical TAS, and let $C : T \to [0, 1]$ be a concentrations function. Let $\mathbb{R}^+ = [0, \infty)$, and let $t \in \mathbb{R}^+$. Let $\alpha \in \mathcal{A}[\mathcal{T}]$ be a producible assembly, and let $[\alpha](t)$ denote the concentration of α at time t, defined as follows. If α is consumed in reactions $\alpha + \beta_1 \to \gamma_1, \ldots, \alpha + \beta_n \to \gamma_n$ and produced in reactions $\beta'_1 + \gamma'_1 \to \alpha, \ldots, \beta'_m + \gamma'_m \to \alpha$, then the concentration $[\alpha](t)$ of α at time t is described by the differential equation

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

and similarly for all other producible assemblies, with $[\alpha](0) = C(r)$ if α is an assembly consisting of a single tile r, and $[\alpha](0) = 0$ otherwise. This completes the definition of the dynamics of concentrations of producible assemblies; it remains to define the time complexity of assembling the terminal assembly of directed systems.

Suppose $\mathcal{T}$ is directed. Then let $\widehat{\alpha} \in \mathcal{A}_\square[\mathcal{T}]$ be its unique producible, 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" 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 the final terminal assembly. Let $\widehat{\alpha}$ be the unique terminal assembly of the hierarchical TAS. We choose a copy of a tile type s at position $p \in \operatorname{dom} \widehat{\alpha}$ to designate as a "timekeeper seed". The assembly of s into $\widehat{\alpha}$ is described as a time-dependent continuous-time Markov process in which each state represents a producible assembly defined at p, and the initial state is the assembly with only s at position p. For each state α representing a producible assembly with s at the origin, and for each pair of producible assemblies β, γ such that $\alpha + \beta \to \gamma$ (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 $[\beta](t)$.[10] Note that 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. In particular, tile concentrations are not assumed to be held constant, since we can no longer assume that tiles have no interaction except with the single assembly containing the seed.

We define $\mathbf{T}_{\mathcal{T},C,p}$ to be the random variable representing the time taken for the copy of s at position p to assemble into $\widehat{\alpha}$ via some sequence of reactions as defined above. Note in particular that unlike the seeded model, there are other reactions occurring outside the sequence, between two assemblies α and β that do not contain the timekeeper tile. However, these (continuously

[10] That 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 $\alpha + \beta \to \gamma$ at time t is assumed to be proportional to $[\alpha](t)[\beta](t)$. 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 at position p 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).

modeled) reactions will determine the concentration of tiles and assemblies that in turn determine the propensities of (stochastically modeled) reactions that *do* involve the timekeeper. We define the time complexity of a directed hierarchical TAS $\mathcal{T}$ to be $\mathsf{T}_{C,p}(\mathcal{T}) = \mathrm{E}[\mathbf{T}_{\mathcal{T},C,p}]$.

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].

3.2.3 Definition of Hierarchical Partial Order Systems

We will show a lower bound on time complexity for a special class of hierarchical directed TAS's known as *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 $\widehat{\alpha}$ is a partial order system if every pair of adjacent positions p_1 and p_2 in $\widehat{\alpha}$ 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 $\mathcal{T} = (T, g, \tau)$ be a hierarchical directed TAS with unique terminal assembly $\widehat{\alpha} \in \mathcal{A}_\square[\mathcal{T}]$. Let Υ be any terminal assembly tree of $\mathcal{T}$. Let $p \in \mathrm{dom}\ \widehat{\alpha}$ and let $s = \widehat{\alpha}(p)$. The *assembly sequence with respect to* Υ *starting at* p is the sequence of assemblies $\vec{\alpha}_{p,\Upsilon} = (\alpha_1, \ldots, \alpha_k)$ 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 $\alpha_k = \widehat{\alpha}$.[11] 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 preorder* with respect to $p \in \mathrm{dom}\ \widehat{\alpha}$ is a preorder $\preceq$ on $\mathrm{dom}\ \widehat{\alpha}$ such that the following holds:

1. For every $p_1, p_2 \in \mathrm{dom}\ \widehat{\alpha}$, if and only if $p_1 \preceq p_2$, then for every assembly sequence $\vec{\alpha} = (\alpha_1, \ldots, \alpha_k)$ 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 $p_1, p_2 \in \mathrm{dom}\ \widehat{\alpha}$, if the tiles at positions p_1 and p_2 interact with positive strength in $\widehat{\alpha}$, then $p_1 \preceq p_2$ or $p_2 \preceq p_1$ (or both).

We say that a directed hierarchical TAS $\mathcal{T}$ with unique terminal assembly $\widehat{\alpha}$ is a *hierarchical partial order system with respect to* p if it has an attachment preorder with respect to p. 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 preorder $\preceq$, we define the *attachment (strict) partial order* $\prec$ *induced by* $\preceq$ to be the partial order on the quotient set of equivalence classes induced by $\preceq$. In other words, if some subassembly $\alpha \sqsubseteq \widehat{\alpha}$ always attaches to the assembly containing p all at once, then all positions $p' \in \mathrm{dom}\ \alpha$ will be equivalent under $\preceq$.[12] It is these equivalence classes

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

[12] More generally, if there is a subset $X \subset \mathrm{dom}\ \widehat{\alpha}$ such that all assemblies α attaching to the assembly containing p have the property that $\mathrm{dom}\ \alpha \cap X \implies X \subseteq \mathrm{dom}\ \alpha$.

of positions that are related under $\prec$.

Each attachment partial order $\prec$ induces a directed acyclic graph $G = (V, E)$, where $V = \{\alpha_1, \ldots, \alpha_k\}$, where each α_i represents the subassembly corresponding to some set of positions in dom α that are equivalent under $\preceq$, and $(\alpha_i, \alpha_j) \in E$ if dom $\alpha_i \prec$ dom α_j. 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]. It is routine to check that the TAS described in Section 3.1 is a partial order system, using one of the orange tiles in the upper-left of Figure 3 as a "seed."

3.2.4 Time Complexity Lower Bound for Hierarchical Partial Order Systems

Theorem 3.4 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 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 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, which allows to us prove the theorem by proving upper bounds on each of these rates using this concentration argument.[13]

The following is a "mass conservation lemma" that will be helpful in the proof of Theorem 3.4. Note that it applies to any hierarchical system.

Lemma 3.3. *Let $\mathcal{T} = (T, g, \tau)$ be a hierarchial TAS and let $C : T \to [0, 1]$ be a concentrations function. Then for all $t \in \mathbb{R}^+$,*

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

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 $f(0) = \sum_{r \in T} C(r) = 1$. 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 (3.1), and recalling the definitions of m, n, β_i', γ_i', and β_i, annotated as $m(\alpha), n(\alpha)$, etc. to show their dependence on α, we have

$$\frac{df_\alpha}{dt} = |\alpha| \left(\sum_{i=1}^{m(\alpha)} [\beta_i'(\alpha)](t)[\gamma_i'(\alpha)](t) - \sum_{i=1}^{n(\alpha)} [\alpha](t)[\beta_i(\alpha)](t) \right).$$

[13] The same assembly α could attach to many locations $p_1, \ldots, p_n$. 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 complete assembly might be possible even if only one of $p_1, \ldots, p_n$ 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 $p_1, \ldots, p_n$ 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. It is open whether TAS's that are not partial order systems are subject to the same time complexity lower bound.

Then

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

Let $\mathcal{R}$ denote the set of all attachment reactions of $\mathcal{T}$, writing $R(\alpha,\beta,\gamma)$ to denote the reaction $\alpha+\beta\to\gamma$. Note that for each such reaction, $|\alpha|+|\beta| = |\gamma|$. Note also that each such reaction contributes precisely three unique terms in the last expression above, two negative (of the form $-|\alpha|[\alpha](t)[\beta](t)$ and $-|\beta|[\alpha](t)[\beta](t)$) and one positive (of the form $|\gamma|[\alpha](t)[\beta](t)$).

Then

$$\begin{aligned}\frac{df}{dt} &= \sum_{R(\alpha,\beta,\gamma)\in\mathcal{R}} (|\gamma|[\alpha](t)[\beta](t) - |\alpha|[\alpha](t)[\beta](t) - |\beta|[\alpha](t)[\beta](t)) \\ &= \sum_{R(\alpha,\beta,\gamma)\in\mathcal{R}} (|\gamma| - |\alpha| - |\beta|)\cdot[\alpha](t)[\beta](t) = 0.\end{aligned}$$

□

Theorem 3.4. *Let $\mathcal{T} = (T, g, \tau)$ be a hierarchial partial order system starting at $p \in \mathbb{Z}^2$, with unique terminal assembly of diameter N. Then for all concentration functions $C : T \to [0,1]$, $\mathsf{T}_{C,p}(\mathcal{T}) \geq \Omega(N)$.*

Proof. Let $\widehat{\alpha} \in \mathcal{A}_\square[\mathcal{T}]$ be the unique terminal assembly of $\mathcal{T}$. Let $\preceq$ be the attachment preorder testifying to the fact that $\mathcal{T}$ is a partial order system. Let $\prec$ be the strict partial order induced by $\preceq$. Let $G = (V, E)$ be the directed acyclic graph induced by $\prec$, with $V = \{\alpha_1, \ldots, \alpha_k\}$. Assign weights to the edges of E by $w(\alpha_i, \alpha_j) = |\text{dom } \alpha_j|$. If P is a longest weighted path in G, it must be the case that $w(P) \geq N/2$.

Let $P = (\alpha_1', \ldots, \alpha_l')$ be any path in G, with weight $w(P) = \sum_{i=2}^{l} |\alpha_i'|$, with α_1' being the single tile at position 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. It suffices to show that $w(P) \leq O(\mathrm{E}[\mathbf{T}_{\mathcal{T},C,p}])$. Let $\mathbf{T}_P$ be the random variable representing the time taken for the complete path P to form. It is clear that $\mathbf{T}_P \leq \mathbf{T}_{\mathcal{T},C,p}$. We will show that $\mathrm{E}[\mathbf{T}_P] \geq \Omega(w(P))$.

To show that $\mathrm{E}[\mathbf{T}_P] \geq \Omega(w(P))$, we will show that for all $t \in \mathbb{R}^+$, $\mathrm{E}[\mathbf{L}(t)] \leq t$. This will suffice because of the following argument. After some amount of time, some prefix P' of the path P has assembled (possibly with some other portions of $\widehat{\alpha}$ 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 seconds. 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, $\mathrm{E}[\mathbf{T}_P] \geq w(P)/4 = \Omega(w(P))$, which establishes the theorem.

It remains to show that for all $t \in \mathbb{R}^+$, $\mathrm{E}[\mathbf{L}(t)] \leq t$. Define the function $f : \mathbb{R}^+ \to \mathbb{R}^+$ for all $t \in \mathbb{R}^+$ by $f(t) = \mathrm{E}[\mathbf{L}(t)]$, and let $f' = \frac{df}{dt}$. It suffices to show that for all $t \in \mathbb{R}^+$, $f'(t) \leq 1$.

Let $P' = (\alpha_1', \ldots, \alpha_{m'}')$ be the prefix of P formed after t seconds. Let $\beta_1, \beta_2, \ldots, \beta_m$, with $m' + m = l$, be the individual subassemblies remaining on the path, in order, so that $P = (P'(1), \ldots, P'(m'), \beta_1, \ldots, \beta_m)$. For all $1 \leq i \leq m$, let $\gamma_i = \bigcup_{j=1}^{i} \beta_j$ be the next i such subassemblies on the path. 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_0 that contain γ_i but do not contain γ_{i+1}. 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 for all $1 \leq i \leq m$. Therefore, invoking Lemma 3.3, the instantaneous rate of change through time at time t in the expected weighted length of the assembled prefix of the path is

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

4 Bounds on Temperature Relative to Number of Tile Types in Seeded Systems

This section shows two bounds relating the number of tile types in a seeded TAS to its temperature. We first show in Theorem 4.1 a polynomial-time algorithm that, given a desired behavior of a tile system, can find strengths to implement that behavior that are at most exponential in the number of tile types, or report that no such strengths exist. The first bound, Theorem 4.4, shows that there are TAS's that require temperature *exponential* in the number of tile types, if any combination of sides may be used for binding. This result can be interpreted to mean that the algorithm of [3] to find the minimum temperature-2 TAS for assembling an $n \times n$ square, which searches over all possible assignments of strengths to the glues, cannot be extended in a straightforward manner to handle larger temperatures, which is why it is necessary for the algorithm of Theorem 5.1 to "shortcut" through the behaviors of tile types rather than enumerating strengths. The second bound, Theorem 4.5, on the other hand, shows that if we restrict attention to those (quite prevalent) classes of tile systems that use only one or two sides of tiles to bind, then *linear* temperature always suffices.

4.1 Behaviors and local equivalence of seeded tile assembly systems

Let T be a set of tile types, and let $t \in T$. We formalize describe the notion of the "behavior" of t as follows. Given a strength function $g : \Lambda(T) \to \mathbb{N}$ and a temperature $\tau \in \mathbb{Z}^+$, define the *cooperation set of t with respect to g and τ* to be the collection $\mathcal{D}_{g,\tau}(t) = \{D_1, D_2, \ldots, D_r\}$ consisting of exactly the subsets $D_j \subseteq \{\mathsf{N}, \mathsf{S}, \mathsf{E}, \mathsf{W}\}$ such that $\sum_{d \in D_j} g(t, d) \geq \tau$, i.e., those collections of sides of t whose glues have sufficient strength to bind cooperatively. Let $\sigma : \mathbb{Z}^2 \dashrightarrow T$ be a seed assembly, let $\tau_1, \tau_2 \in \mathbb{Z}^+$ be temperatures, and let $g_1, g_2 : \Lambda(T) \to \mathbb{N}$ be strength functions. We say that the seeded TAS's $\mathcal{T}_1 = (T, \sigma, g_1, \tau_1)$ and $\mathcal{T}_2 = (T, \sigma, g_2, \tau_2)$ are *locally equivalent* if, for each tile type $t \in T$, $\mathcal{D}_{g_1,\tau_1}(t) = \mathcal{D}_{g_2,\tau_2}(t)$.[14] The behavior of an individual tile type during assembly is completely determined by its cooperation set, in the sense that if $\mathcal{T}_1$ and $\mathcal{T}_2$ are locally equivalent, then $\mathcal{A}[\mathcal{T}_1] = \mathcal{A}[\mathcal{T}_2]$ and $\mathcal{A}_\square[\mathcal{T}_1] = \mathcal{A}_\square[\mathcal{T}_2]$.[15]

Even without any strength function or temperature, by specifying a cooperation set for each tile type, one can describe a "behavior" of a TAS in the sense that its dynamic evolution can be

[14] Note that the definition of equivalence is independent of the seed assembly; we include it only to be able to talk about the equivalence of TAS's rather than the more cumbersome "equivalence of triples of the form (T, g, τ)."

[15] The converse does not hold, however. For instance, some tile types may have a subset of sides whose glues never appear at a frontier location during assembly, so it would be irrelevant to the definition of $\mathcal{A}[\mathcal{T}]$ and $\mathcal{A}_\square[\mathcal{T}]$ whether or not that combination of glues have enough strength to bind.

simulated knowing only the cooperation set of each tile type. We call a system thus specified a *strength-free* TAS. More formally, a *seeded strength-free TAS* is a triple $(T, \sigma, \mathcal{D})$, where T is a finite set of tile types, $\sigma \in T$ is the seed, and $\mathcal{D} : T \to \mathcal{P}(\{\mathsf{N}, \mathsf{S}, \mathsf{E}, \mathsf{W}\})$ is a function from a tile type $t \in T$ to a cooperation set $\mathcal{D}(t)$. For a standard seeded TAS $\mathcal{T} = (T, \sigma, g, \tau)$ and a seeded strength-free TAS $\mathcal{T}_{\text{sf}} = (T, \sigma', \mathcal{D})$, we say that they are *locally equivalent* if $\sigma = \sigma'$ and $\mathcal{D}(t) = \mathcal{D}_{g,\tau}(t)$ for each tile type $t \in T$.

Note that every seeded TAS has a unique cooperation set for each tile type, and hence, has a locally equivalent strength-free TAS. However, not every seeded strength-free TAS is realizable by a TAS. This is because cooperation sets for tile types in T could be contradictory; for instance, two tile types t_1 and t_2 could share a north glue with $\{\mathsf{N}\}$ in the cooperation set of t_1 but not that of t_2.

Theorem 4.1. *For a given seeded strength-free TAS with k tile types, it is decidable in polynomial time whether there exists a locally equivalent seeded TAS. Moreover, such a TAS whose temperature is at most $2^{O(k^2)}$ can be output in polynomial time.*

Proof. Let $\mathcal{T}_{\text{sf}}$ be a strength-free TAS with a tile set T of k tile types with $u \leq k$ different glues.[16] We would like to decide whether $\mathcal{T}_{\text{sf}}$ is in fact realizable by a TAS. To have the tightest upper bound on the temperature, ideally we would like to solve the problem of finding the minimum temperature TAS that is locally equivalent to $\mathcal{T}_{\text{sf}}$. This optimization problem can be cast as an integer linear program on a temperature variable τ and a set of glue-strength variables $s_1, s_2, \ldots, s_u$ as in the following example:

$$\begin{array}{lllll} \text{Minimize} & \tau & & & \\ \text{subject to} & \tau, s_1, s_2, \ldots, s_u & \in & \mathbb{N} \\ & s_1 + s_3 - \tau & \geq & 0 \\ & s_1 + s_4 + s_6 - \tau & \geq & 0 \\ & & \cdots & \\ & s_1 + s_2 + s_4 - \tau & \leq & -1 \\ & s_2 + 2s_3 - \tau & \leq & -1 \\ & & \cdots & \end{array}$$

The "≥ 0" inequalities correspond to the union of all cooperation sets of all tile types, and the "≤ -1" inequalities correspond to the union of all complements of the cooperation sets; i.e., each set $\mathcal{P}(\{\mathsf{N}, \mathsf{S}, \mathsf{E}, \mathsf{W}\}) \setminus \mathcal{D}(t)$, where $\mathcal{D}(t)$ is the cooperation set of t. Since we require each strength to be an integer, "≤ -1" is equivalent to "< 0". Since each tile type has $|\mathcal{D}(t)|$ "≥ 0" inequalities (one for each subset of sides in its cooperation set) and $16 - |\mathcal{D}(t)|$ "≤ -1" inequalities, there are $16k$ inequalities in the integer linear program.

Our goal will not be to find the smallest temperature TAS that satisfies the constraints above (which remains an open problem), but simply to find any feasible integer solution with temperature and strengths at most $2^{O(k^2)}$. Call the above system of constraints (including the integer constraint) S_1. Consider the real-valued system of linear inequalities S_2 defined as the above inequalities with the integer constraint $\tau, s_1, s_2, \ldots, s_u \in \mathbb{N}$ relaxed to simply $\tau, s_1, s_2, \ldots, s_u \geq 0$. Then we have

[16] Each tile type has 4 sides so it might seem that there could be $4k$ total glues if there are k tile types. However, in a nontrivial system (one that has no "effectively null" glues that appear on only one tile type and therefore do not bind to any other tile types), for each side of a tile type, the choice of glue for that side is limited to those glues on the opposite side of the $k - 1$ other tile types, or alternately we could choose the null strength-0 glue.

the implication "S_1 has a solution" $\implies$ "S_2 has a solution". Conversely, any rational-valued solution to S_2 can be converted to an integer-valued solution to S_1 by multiplying each value by the least common multiple L of the denominators of the rational numbers.[17] Furthermore, since the input coefficients are integers, if the feasible polytope of S_2 is non-empty, then all of its vertices are rational. Therefore S_2 has a solution if and only if it has a rational solution. Therefore we have the full bidirectional implication "S_1 has a solution" $\iff$ "S_2 has a solution". We can pick any n linearly independent inequalities of S_2, interpret them as equalities, and use Gaussian elimination (with exact rational arithmetic) to obtain some vertex of the feasible polytope described by the inequalities, and convert these to integer solutions to S_1 through multiplication as described above. If we cannot find n linearly independent inequalities (testable by computing the rank of the matrix defining the inequalities) then there is no TAS implementing the behavior of $\mathcal{T}_{\text{sf}}$. It remains to show that in case there is a solution, the integers we obtain by this method obey the stated upper bound $2^{O(k^2)}$.

Each coefficient has absolute value at most 2 (since we may assume N/S glues are disjoint from E/W glues), and each equation has at most 5 nonzero terms since each tile type has only 4 sides (together with the -1 coefficient for τ). Applying Lemma 4.2 (stated and proven after the current proof) with $n = u + 1$, $c_1 = 2$, and $c_2 = 5$, each vertex is a rational vector $\vec{x} = (\frac{p_1}{q_1}, \ldots, \frac{p_{u+1}}{q_{u+1}})$ such that, for each i, $|p_i|, |q_i| \leq 2^{u+1}6^{(u+1)/2} = 2^{u+1+u\log 6+0.5} < 2^{4u+2}$. Since we enforce nonnegativity, $p_i = |p_i|$ and $q_i = |q_i|$. We multiply this vector by $L = \text{LCM}(q_1, \ldots, q_{u+1}) \leq \prod_{j=1}^{u+1} q_j$ to obtain integer values $x'_i = x_i \cdot L \leq \frac{p_i}{q_i} \cdot \prod_{j=1}^{u+1} q_j < (2^{4u+2})^{u+1} \leq 2^{(4k+2)(k+1)} = 2^{O(k^2)}$. □

Lemma 4.2. *Let $c_1, c_2 \in \mathbb{Z}^+$ be constants, and let $\vec{b}$ be an $n \times 1$ integer column vector and $A = (a_{ij})$ be a nonsingular $n \times n$ integer matrix such that for each i, j, $|a_{ij}| \leq c_1$ and $|b_j| \leq c_1$, and each row of A contains at most c_2 nonzero entries. Then the solution to the linear system $A\vec{x} \leq \vec{b}$ is a rational vector $\vec{x} \in \mathbb{Q}^n$ such that, if each component $x_i = \frac{p_i}{q_i}$ is written in lowest terms with $p_i, q_i \in \mathbb{Z}$, then $|p_i| \leq c_1^n(c_2+1)^{n/2}$ and $|q_i| \leq c_1^n c_2^{n/2}$.*

Proof. Recall Hadamard's inequality $|\det A| \leq \prod_{i=1}^n \|v_i\|_2$, where v_i is the i^{th} row of A.[18] Since v_i has at most c_2 nonzero entries that are each at most absolute value c_1, Hadamard's inequality tells us that $|\det A| \leq \prod_{i=1}^n \sqrt{c_2 \cdot c_1^2} = c_1^n c_2^{n/2}$. Similarly, letting A_i be A with column vector $\vec{b}$ replacing A's i^{th} column, A_i has at most $c_2 + 1$ nonzero entries per row, so a similar argument gives $|\det A_i| \leq c_1^n(c_2+1)^{n/2}$. The i^{th} solution is $x_i = \frac{\det A_i}{\det A}$ by Cramer's rule. Since A and A_i are integer-valued, so are $\det A$ and $\det A_i$, whence the upper bounds on $|\det A|$ and $|\det A_i|$ also apply to $|q_i|$ and $|p_i|$, respectively, since they are x_i's lowest terms representation. □

We conclude this section by counting the number of different strength-free TAS's with k tile types.

Proposition 4.3. *For each $k \in \mathbb{Z}^+$, there are at most $168^k k^{4k+1}$ strength-free TAS's with k tile types.*

[17] This actually enforces the stronger condition that each "≤ -1" inequality is actually "$\leq -L$". This is possible because we have no *upper* bound on the variables, which would prevent multiplication from preserving the inequalities. Such an upper bound is key to reductions that show the NP-hardness of detecting whether a given system of linear inequalities has an integer solution.

[18] Hadamard's inequality is typically stated for v_i a column of A, but the determinant of a matrix and its transpose are equal so the bound holds when taking the product over rows as well.

Proof. For $t \in T$, the cooperation set of t is a collection of subsets of $\{\mathsf{N}, \mathsf{S}, \mathsf{E}, \mathsf{W}\}$ that is closed under the superset operation. The closure property under the superset operation means that if $D \subseteq D' \subseteq \{\mathsf{N}, \mathsf{S}, \mathsf{E}, \mathsf{W}\}$ and D is in a cooperation set of t, then D' is also in the cooperation set. This is due to the fact that having strictly more sides available to bind cannot inhibit binding that would otherwise occur, as long as strengths are assumed to be nonnegative. The number of different possible behaviors of a tile type is at most the number of cooperation sets. Each cooperation set $\mathcal{D}$ is defined by a unique *antichain*, which is a subcollection $\mathcal{D}' = \{D'_1, \ldots, D'_m\} \subseteq \mathcal{D}$ such that, for all $1 \leq i, j \leq m$, $D'_i \not\subseteq D'_j$, whose closure under the superset operation is equal to $\mathcal{D}$. The antichain consists of the minimal elements of $\mathcal{D}$ under the relation $\subseteq$. The number of antichains of subsets of $\{\mathsf{N}, \mathsf{S}, \mathsf{E}, \mathsf{W}\}$ is given by the fourth Dedekind number $M(4) = 168$ [35]. Thus, each tile type has at most 168 different behaviors (cooperation sets) depending on the strengths assigned to its glues and on the temperature.

For each side of a tile type, there are at most k glue labels to choose, so there are at most k^4 ways to assign these labels to each side of a tile type. Therefore, encoding each tile type as a list of 4 glue labels and cooperation set, and encoding a TAS as a list of tile types, there are at most $(168k^4)^k = 168^k k^{4k}$ different seeded strength-free TAS's with k tile types, and hence at most $168^k k^{4k+1}$ different seeded strength-free TAS's with k tile types, which have the choice of their seed among k tile types. □

4.2 Tile Assembly Systems Requiring Temperature Exponential in Number of Tile Types

In this section, we prove that a temperature that is exponential in the number of tile types given by Theorem 4.1 is optimal, although there is a gap between the exponents ($2^{|T|/4}$ for Theorem 4.4 below versus $2^{O(|T|^2)}$ for Theorem 4.1).

Theorem 4.4. *For every $n \in \mathbb{Z}^+$, there is a TAS $\mathcal{T} = (T, \sigma, g, \tau)$ such that $|T| = 4n$ and for every TAS $\mathcal{T}' = (T, \sigma, g', \tau')$ that is locally equivalent to $\mathcal{T}$, $\tau' \geq 2^n$.*

Proof. The tile set T is shown in Figure 4.[19] In each stage, each of the top two light tile types represents a triple (a pair in stage 1) of glues whose sum is at least τ. Each of the bottom two dark tiles represents a triple of glues whose sum is less than τ. For the dark tile types to be nontrivial, we could imagine that the (unlabeled) north glue is strong enough to cooperate with some of the other glues. The actual strengths of the glues are left as variables, but the caption of Figure 4 gives one example of strengths that would satisfy the inequalities that the tile types represent.

We prove by induction on n that $A''_n \geq A_n + 2^n$. For the base case, in the first stage, the top light tile type and top dark tile type enforce that $A'_1 + B'_1 \geq \tau > A_1 + B'_1$, so $A_1 < A'_1$. Similarly, the bottom light tile type and the bottom dark tile type enforce that $A'_1 < A''_1$. Therefore $A''_1 \geq A_1 + 2$.

For the inductive case, assume that $A''_{n-1} \geq A_{n-1} + 2^{n-1}$. The top dark tile type enforces that $\tau > A''_{n-1} + A_n + B'_n$, and by the induction hypothesis, $A''_{n-1} + A_n + B'_n \geq A_{n-1} + 2^{n-1} + A_n + B'_n$. The top light tile type enforces that $A'_n + A_{n-1} + B'_n \geq \tau$, which combined with the previous two inequalities shows that $A'_n \geq 2^{n-1} + A_n$. A similar analysis with the bottom light tile type and

[19] We do not specify the seed assembly since we are concerned only with the local behavior of the tiles. We would need to add a small number of tile types to the TAS to ensure that each tile shown is actually attachable at some point during assembly, but this would not affect the asymptotic size of the tile set as $n \to \infty$, so the exponential lower bound on the temperature would still hold.

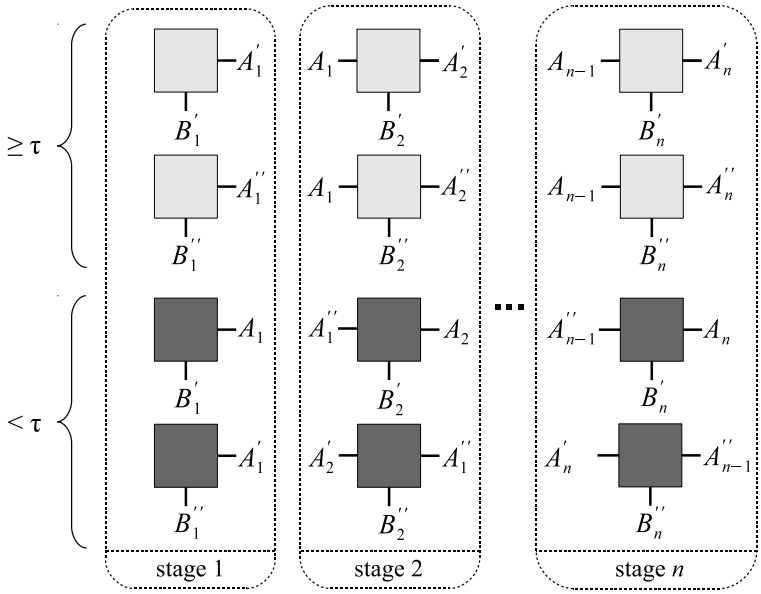

Figure 4: A set of tile types requiring temperature that is exponentially larger than the number of tile types. There are stages $1, 2, \ldots, n$, with stage i containing 4 tiles, and stage i ensuring that the gap between the largest and smallest strength in the stage is at least 2^i. In each stage, each of the top two light tiles represents a triple (a pair in stage 1) of glues whose sum is at least τ. Each of the bottom two dark tiles represents a triple of glues whose sum is less than τ. The inequalities are satisfiable, for instance, by setting $A_n = 3^{n-1}, A'_n = 2A_n, A''_n = 3A_n, B_n = \tau - A_n, B'_n = \tau - A'_n, B''_n = \tau - A''_n$.

bottom dark tile type shows that $A''_n \geq 2^{n-1} + A'_n$, whence $A''_n \geq 2^{n-1} + 2^{n-1} + A_n$, establishing the inductive case.

Since it can be assumed without loss of generality that strengths are at most τ, this shows that the tile set consisting of n stages, having $|T| = 4n$ tile types, requires $\tau \geq 2^n$ to be realized. □

4.3 Temperature Linear in the Number of Tile Types Suffices for 2-Cooperative Equivalence

Theorem 4.4 shows that temperature exponentially larger than the number of tile types is sometimes necessary for a TAS's behavior to be realized by integer strengths. However, the definition of locally equivalent assumes that all possible combinations of sides of a tile type may be present in an assembly. Many TAS's are more constrained than this. There is a wide class of TAS's that we term *2-cooperative*, meaning that all binding events during all assembly sequences use only 1 or 2 sides that bind with positive strength. Nearly all theoretical TAS's found in the literature are 2-cooperative. In this section we show that the 3-cooperativity of Figure 4 is necessary, by showing that 2-cooperative systems can always be realized by strengths linear in the number of tile types.

Theorem 4.5. *Let $\mathcal{T} = (T, \sigma, g, \tau)$ be a TAS, and let $\mathcal{D}^{(2)}_{g,\tau}(t) \subseteq \mathcal{D}_{g,t}(t)$ be the cooperation set of t with respect to g and τ restricted to containing only subsets of $\{\mathsf{N}, \mathsf{S}, \mathsf{E}, \mathsf{W}\}$ of cardinality 1 or 2. Then there is a TAS $\mathcal{T}' = (T, \sigma, g', \tau')$ with $\tau' \leq 2|T| + 2$ such that, for each $t \in T$, $\mathcal{D}^{(2)}_{g,\tau}(t) = \mathcal{D}^{(2)}_{g',\tau'}(t)$.*

That is, $\mathcal{T}'$ is equivalent in behavior to $\mathcal{T}$, so long as all attachments involve only 1 or 2 sides.

Proof. Let $K = |T|$. Let $\Lambda(T)$ denote the set of all glue labels on tile types in T. Let $G = \{\ g(\sigma)\ \mid \sigma \in \Lambda(T)\ \} \setminus \{0, \tau\}$. That is, G is the set of all positive but insufficient glue strengths used in this system. $|G| \leq 2K$ since there are at most $|T|$ north-south glues and at most $|T|$ east-west glues.

We split G into two subsets $L = \{\ g \in G \mid 0 < g < \tau/2\ \}$ and $H = \{\ g \in G \mid \tau/2 \leq g < \tau\ \}$. Let $L = \{\ell_1, \ell_2, \ldots, \ell_n\}$ such that $\ell_1 < \ell_2 < \ldots < \ell_n < \tau/2$, and let $H = \{h_1, h_2, \ldots, h_m\}$ with $\tau/2 \leq h_1 < h_2 \ldots h_m < \tau$. For descriptive purposes, define $\ell_0 = 0$ and $\ell_{n+1} = \tau/2$ (although these numbers may or may not be glue strengths).

We aim at designing an algorithm to find a glue function g' satisfying:

1. for any label $\sigma \in \{\sigma_1, \ldots, \sigma_{2K}\}, g(\sigma) \geq \tau \iff g'(\sigma) \geq 2n + 2$;
2. for any pair of labels $\{\sigma, \sigma'\} \subset \{\sigma_1, \ldots, \sigma_{2K}\}, g(\sigma) + g(\sigma') \geq \tau \iff g'(\sigma) + g'(\sigma') \geq 2n + 2$.

Then for $\tau' = 2n + 2$, the TAS $\mathcal{T}' = (T, \sigma, g', \tau')$ satisfies the 2-cooperative equivalence with $\mathcal{T}$ that we seek.

First, we define an equivalence relation $\equiv$ on H defined as: for $h, h' \in H$, $h \equiv h'$ if $(\forall\ 1 \leq i \leq n)(h + \ell_i \geq \tau \iff h' + \ell_i \geq \tau)$. This partitions H into subsets $H_{n+1}, \ldots, H_2, H_1$ such that $h \in H_j$ if and only if $\tau - \ell_j \leq h < \tau = \ell_{j-1}$. In other words, $h \in H_j$ if $h + \ell_j \geq \tau$ implies $i \geq j$.

A glue function $g' : \Lambda(T) \to \mathbb{N}$ is defined as follows: for a label $\sigma \in \Lambda(T)$,

$$g'(\sigma) = \begin{cases} 0, & \text{if } g(\sigma) = 0; \\ i, & \text{if } g(\sigma) = \ell_i\ (1 \leq i \leq n); \\ 2n + 2 - j, & \text{if } g(\sigma) \in H_j\ (1 \leq j \leq n + 1); \\ 2n + 2, & \text{if } \tau \leq g(\sigma). \end{cases}$$

It is trivial that this satisfies condition (1) above. Let $\sigma, \sigma' \in \Lambda(T)$ with $0 < g(\sigma) \leq g(\sigma') < \tau$. There are two cases.

1. Suppose that $\tau \leq g(\sigma) + g(\sigma')$. Note that $\tau/2 \leq g(\sigma')$. If $\tau/2 \leq g(\sigma)$, then by definition $n + 1 \leq g'(\sigma)$ and $n + 1 \leq g'(\sigma')$, whence their sum it at least $2n + 2$. If $g(\sigma) < \tau/2$, then let $g(\sigma) = \ell_i$ for some $1 \leq i \leq n$. Since $\tau \leq g(\sigma) + g(\sigma')$, this means that $\tau = \ell_i \leq g(\sigma')$, which in turn implies $g(\sigma') \in H_i \cup H_{i-1} \cup \ldots \cup H_1$. By definition, $g'(\sigma) = i$ and $g'(\sigma') \geq 2n + 2 - i$. Consequently, $2n + 2 \leq g'(\sigma) + g'(\sigma')$.
2. Suppose that $g(\sigma) + g(\sigma') < \tau$. In this case, $g(\sigma) < \tau/2$. If $g(\sigma') < \tau/2$, then by the definition of g', both $g(\sigma)$ and $g(\sigma')$ are at most n so that their sum cannot reach $2n + 2$. Otherwise, let $g(\sigma) = \ell_i$. An argument similar to the one above gives $g(\sigma') \in H_{n+1} \cup \ldots \cup H_{i+1}$, whence $g'(\sigma') < 2n + 2 - i$. Thus, $g'(\sigma) + g'(\sigma') < 2n + 2$.

This verifies that g' satisfies condition (2). □

5 Polynomial-Time Algorithm to Find the Minimum Seeded Tile Assembly System for a Square at any Temperature

In this section we show that there is a polynomial-time algorithm that, given an $n \times n$ square S_n, computes the smallest directed seeded TAS that strictly self-assembles S_n. Adleman, Cheng, Goel,

Huang, Kempe, Moisset de Espanés, and Rothemund [5] showed that this problem is polynomial-time solvable when the temperature is restricted to be 2, and asked whether there is an algorithm that works when the temperature is unrestricted, which we answer affirmatively.

Throughout this section, for any seeded TAS $\mathcal{T} = (T, \sigma, g, \tau)$, assume that $|\text{dom } \sigma| = 1$, so that all TAS's under consideration are assumed to be singly-seeded. Let $S \subseteq \mathbb{Z}^2$ be a shape. The *directed tile complexity* of S is

$$\mathrm{C}^{\mathrm{dtc}}(S) = \min\{\ |T| \mid \mathcal{T} = (T, \sigma, g, \tau) \text{ is a directed seeded TAS that strictly self-assembles } S\ \}.$$

Let $\mathcal{FS} \subset \mathcal{P}(\mathbb{Z}^2)$ denote the set of all finite shapes, and let $\mathcal{SQ} \subset \mathcal{FS}$ denote the set of all squares. To avoid ambiguity between shapes that are equivalent up to translation, assume that $\mathcal{FS}$ and $\mathcal{SQ}$ consist only of shapes confined to the first quadrant with at least one x-coordinate equal to 0 and at least one y-coordinate equal to 0 (i.e., translated as far left and down as possible while remaining in the first quadrant). Define the *minimum tile set problem restricted to squares* as

$$\textsc{MinDirectedTileSetSquare} = \left\{\ \langle S, c\rangle\ \middle|\ S \in \mathcal{SQ}, c \in \mathbb{Z}^+, \text{ and } \mathrm{C}^{\mathrm{dtc}}(S) \leq c\ \right\}.$$

Theorem 5.1. $\textsc{MinDirectedTileSetSquare} \in \mathsf{P}$. *Furthermore, given a square S, it is possible in polynomial time to output a minimal directed seeded TAS $\mathcal{T} = (T, \sigma, g, \tau)$ that strictly self-assembles S.*

Proof. In [5], the authors study the variant of the minimum directed tile set problem restricted to squares where the temperature is fixed at $\tau = 2$. They show this variant of the problem is in P by the following argument. For all $n \in \mathbb{N}$, let $S_n = [n]^2$ denote the $n \times n$ square. Adleman, Cheng, Goel, and Huang [3] showed that for all $n \in \mathbb{N}$, $\mathrm{C}^{\mathrm{dtc}}(S_n) \leq O(\frac{\log n}{\log\log n})$. The proof of [5] first shows by a simple counting argument that there are at most a polynomial in n number of temperature-2 TAS's with at most $O(\frac{\log n}{\log\log n})$ tile types, using the fact that all strengths may without loss of generality be assumed to be 0, 1, or 2. They then make use of a polynomial-time algorithm Unique-Shape devised in the same paper [5] that, given any shape S and any TAS $\mathcal{T}$, determines whether $\mathcal{T}$ is directed and strictly self-assembles S.[20] Finding the minimum directed TAS for S_n then amounts to iterating over every "small" ($O(\frac{\log n}{\log\log n})$ tile types) TAS $\mathcal{T}$ and using the algorithm Unique-Shape to check which systems assemble S_n, reporting the size of the smallest.

Let $c \in \mathbb{Z}^+$ be the constant, shown to exist in [3], such that, for all $n \in \mathbb{N}$, $\mathrm{C}^{\mathrm{dtc}}(S_n) \leq c \log n / \log\log n$. Let $k = c \log n / \log\log n$. By enumerating every strength-free TAS with k tile types, we are guaranteed to enumerate one of those that testifies to the value of $\mathrm{C}^{\mathrm{dtc}}(S_n)$.[21] For each of these strength-free TAS's, we test whether the strength-free system strictly self-assembles S_n and is directed. The algorithm Unique-Shape of [5] can be executed just as easily with a strength-free TAS as with a standard TAS, so Unique-Shape may be used for this test.

Proposition 4.3 tells us that the number of strength-free TAS's with k tile types is at most

[20] The conjunction "and" is crucial. To determine whether a given $\mathcal{T}$ is directed is undecidable in general, and to determine whether a given (possibly non-directed) $\mathcal{T}$ strictly self-assembles a given shape S is coNP-complete, even in the case when S is restricted to be a square [6, 18, 27].

[21] We need not search smaller TAS's since for every $k < k'$, every directed TAS with k tile types has a locally equivalent TAS with k' tile types, by padding with junk tile types that have all strength-0 glues.

$168^k k^{4k+1}$. We have that $168^k \leq 168^{c \log n / \log \log n} \leq 168^{c \log n} \leq (2^{\log n})^{c \log 168} \leq n^{8c}$, and

$$\begin{aligned} k^{4k+1} &\leq k(c \log n / \log \log n)^{4c \log n / \log \log n} = k(2^{\log(c \log n / \log \log n)})^{4c \log n / \log \log n} \\ &= k(2^{\log c + \log \log n - \log \log \log n})^{4c \log n / \log \log n} = k 2^{4c \log n (\log c + \log \log n - \log \log \log n)/ \log \log n} \\ &= k(2^{\log n})^{4c(\log c + \log \log n - \log \log \log n)/ \log \log n} = k n^{4c \log c / \log \log n + 4c - 4c \log \log \log n / \log \log n} \\ &\leq n^{4c \log c + 4c + 1}, \end{aligned}$$

whence the number of strength-free TAS's to search is at most $n^{4c \log c + 12c + 1} = \text{poly}(n)$. Once we obtain a smallest such strength-free TAS $\mathcal{T}_{\text{sf}}$, we run the algorithm described in Theorem 4.1 to obtain a minimal standard directed seeded TAS $\mathcal{T}$ that strictly self-assembles S_n. □

We note that while we have stated the theorem for the family of square shapes, our method, as well as that of [5], works for any family of shapes $S_1, S_2, \ldots$ where $|S_n| = \text{poly}(n)$ and the directed tile complexity of S_n is at most $O(\frac{\log n}{\log \log n})$. This includes, for instance, the family $\{T_1, T_2, \ldots\}$, where T_n is a width-n right triangle, and for each $q \in \mathbb{Q}^+$ the family $\{R_{q,1}, R_{q,2}, \ldots\}$, where $R_{q,n}$ is the $n \times \lfloor qn \rfloor$ rectangle.

6 Open Questions

There are some interesting questions that remain open:

1. What upper or lower bound can be placed on the quantity $\text{depth}^{\text{da}}(S_n)$ for S_n 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 $\text{depth}^{\text{da}}(S_n) < o(\log^2 n)$ or $\text{depth}^{\text{da}}(S_n) > \omega(\log n)$. Obtaining bounds for more general shapes would also be interesting.

2. What is the status of the following problems?

$$\textsc{HierarchicalUniqueShape} = \left\{ \langle S, \mathcal{T} \rangle \;\middle|\; \begin{array}{l} \mathcal{T} \text{ is a hierarchical TAS and} \\ \text{strictly self-assembles finite shape } S \end{array} \right\},$$

$$\textsc{HierarchicalDirectedUniqueShape} = \left\{ \langle S, \mathcal{T} \rangle \;\middle|\; \begin{array}{l} \mathcal{T} \text{ is a directed hierarchical TAS and} \\ \text{strictly self-assembles finite shape } S \end{array} \right\}.$$

   In the case of the seeded aTAM, the seeded variants of these problems are known to be coNP-complete [6, 18, 27] and in P [5], respectively.

3. What is the status of the following problems?

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

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

   In the case of the seeded aTAM, the seeded variants of these problems are known to be Σ_2^{P}-complete [9] and NP-complete [5], respectively.

4. What is the time complexity of strictly self-assembling an $n \times n$ square with a directed hierarchical TAS? Any shape with diameter n? What if we do not require the TAS to be directed? (Note that this requires defining a reasonable notion of assembly time for general hierarchical systems.)

5. Is there a polynomial-time algorithm that, given a strength-free TAS $\mathcal{T}_{\mathsf{sf}}$, outputs a TAS of minimal temperature that is locally equivalent to $\mathcal{T}$?

Acknowledgement. The authors thank Ehsan Chiniforooshan especially, for many fruitful and illuminating discussions that led to the results on temperature, and also the members of Erik Winfree's group, particularly Damien Woods and Erik Winfree, for insightful discussion and comments. The third author is grateful to Aaron Meyerowitz (via the website http://mathoverflow.net) for pointing out the Dedekind numbers as a way to count the number of collections of subsets of a given set that are closed under the superset operation.

References

[1] 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.

[2] Leonard M. Adleman. Toward a mathematical theory of self-assembly (extended abstract). Technical Report 00-722, University of Southern California, 2000.

[3] Leonard M. Adleman, Qi Cheng, Ashish Goel, and Ming-Deh Huang. Running time and program size for self-assembled squares. In *STOC 2001: Proceedings of the thirty-third annual ACM Symposium on Theory of Computing*, pages 740–748, Hersonissos, Greece, 2001. ACM.

[4] Leonard M. Adleman, Qi Cheng, Ashish Goel, Ming-Deh Huang, and Hal Wasserman. Linear self-assemblies: Equilibria, entropy and convergence rates. In *In Sixth International Conference on Difference Equations and Applications*. Taylor and Francis, 2001.

[5] Leonard M. Adleman, Qi Cheng, Ashish Goel, Ming-Deh A. Huang, David Kempe, Pablo Moisset de Espanés, and Paul W. K. Rothemund. Combinatorial optimization problems in self-assembly. In *STOC 2002: Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing*, pages 23–32, 2002.

[6] Gagan Aggarwal, Qi Cheng, Michael H. Goldwasser, Ming-Yang Kao, Pablo Moisset de Espanés, and Robert T. Schweller. Complexities for generalized models of self-assembly. *SIAM Journal on Computing*, 34:1493–1515, 2005. Preliminary version appeared in SODA 2004.

[7] Robert D. Barish, Rebecca Schulman, Paul W. Rothemund, 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.

[8] 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.

[9] 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*, 2011. to appear.

[10] 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.

[11] 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.

[12] Erik D. Demaine, Matthew J. Patitz, Robert T. Schweller, and Scott M. Summers. Self-assembly of arbitrary shapes with RNA and DNA tiles. Technical Report 1004.4383, Computing Research Repository, 2010.

[13] David Doty. Randomized self-assembly for exact shapes. *SIAM Journal on Computing*, 39(8):3521–3552, 2010. Preliminary version appeared in FOCS 2009.

[14] 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.

[15] Irvin R. Epstein and John A. Pojman. *An introduction to nonlinear chemical dynamics: oscillations, waves, patterns, and chaos*. Oxford University Press, USA, 1998.

[16] Constantine Evans. Personal communication.

[17] Kenichi Fujibayashi, Rizal Hariadi, Sung Ha Park, Erik Winfree, and Satoshi Murata. Toward reliable algorithmic self-assembly of DNA riles: A fixed-width cellular automaton pattern. *Nano Letters*, 8(7):1791–1797, 2007.

[18] Michael R. Garey and David S. Johnson. *Computers and Intractability*. W. H. Freeman, New York, 1979.

[19] 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.

[20] Daniel T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *Journal of Physical Chemistry*, 81(25):2340–2361, 1977.

[21] Daniel T. Gillespie. Stochastic simulation of chemical kinetics. *Annual review of physical chemistry*, 58(1):35–55, 2007.

[22] 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.

[23] 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.

[24] 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.

[25] 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.

[26] Ming-Yang Kao and Robert T. Schweller. Randomized self-assembly for approximate shapes. In Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldrsson, Anna Ingólfsdóttir, and Igor Walukiewicz, editors, *ICALP 2008: International Colloqium on Automata, Languages, and Programming*, volume 5125 of *Lecture Notes in Computer Science*, pages 370–384. Springer, 2008.

[27] Harry R. Lewis and Christos H. Papadimitriou. *Elements of the Theory of Computation*. Prentice-Hall, Inc., Upper Saddle River, New Jersey, 1998.

[28] 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.

[29] 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.

[30] Matthew J. Patitz and Scott M. Summers. Identifying shapes using self-assembly (extended abstract). In *ISAAC 2010: Proceedings of the 21st International Symposium on Algorithms and Computation*, 2010. to appear.

[31] Paul W. K. Rothemund and Erik Winfree. The program-size complexity of self-assembled squares (extended abstract). In *STOC 2000: Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing*, pages 459–468, 2000.

[32] Paul W.K. Rothemund, Nick Papadakis, and Erik Winfree. Algorithmic self-assembly of DNA Sierpinski triangles. *PLoS Biology*, 2(12):2041–2053, 2004.

[33] Robert Schweller. Personal communication.

[34] Nadrian C. Seeman. Nucleic-acid junctions and lattices. *Journal of Theoretical Biology*, 99:237–247, 1982.

[35] Neil J. Sloane. The On-Line Encyclopedia of Integer Sequences. http://oeis.org/wiki/A000372. Dedekind numbers: number of monotone Boolean functions of n variables or number of antichains of subsets of an n-set.

[36] 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.

[37] Scott M. Summers. Reducing tile complexity for the self-assembly of scaled shapes through temperature programming. Technical Report 0907.1307, Computing Research Repository, 2009.

[38] Leroy G. Wade. *Organic Chemistry*. Prentice Hall, 2nd edition, 1991.

[39] Hao Wang. Proving theorems by pattern recognition – II. *The Bell System Technical Journal*, XL(1):1–41, 1961.

[40] 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.

[41] Erik Winfree. *Algorithmic Self-Assembly of DNA*. PhD thesis, California Institute of Technology, June 1998.

[42] 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.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de