

Title:	Hierarchical self-assembly
Name:	David Doty ¹
Affil./Addr.	California Institute of Technology, Pasadena, CA, USA. Supported by NSF grants CCF-1219274, CCF-1162589, and 1317694.
Keywords:	self-assembly; hierarchical assembly, intrinsic universality, running time, verification
SumOriWork:	2005; Aggarwal, Cheng, Goldwasser, Kao, Espanes, Schweller 2012; Chen, Doty 2013; Cannon, Demaine, Demaine, Eisenstat, Patitz, Schweller, Summers, Winslow

Hierarchical self-assembly

DAVID DOTY¹

California Institute of Technology, Pasadena, CA, USA. Supported by NSF grants CCF-1219274, CCF-1162589, and 1317694.

Years and Authors of Summarized Original Work

2005; Aggarwal, Cheng, Goldwasser, Kao, Espanes, Schweller

2012; Chen, Doty

2013; Cannon, Demaine, Demaine, Eisenstat, Patitz, Schweller, Summers, Winslow

Keywords

self-assembly; hierarchical assembly, intrinsic universality, running time, verification

Problem Definition

The general idea of *hierarchical* self-assembly (a.k.a., *multiple tile* [2], *polyomino* [8,10], *two-handed* [3,5,6]) is to model self-assembly of tiles in which attachment of two multi-tile assemblies is allowed, as opposed to all attachments being that of a single tile onto a larger assembly. Several problems concern comparing hierarchical self-assembly to its single-tile-attachment variant (called the “seeded” model of self-assembly), so we define both models here. The model of hierarchical self-assembly was first defined (in a slightly different form that restricted the size of assemblies that could attach) by Aggarwal, Cheng, Goldwasser, Kao, Moisset de Espanes, and Schweller [2]. Several generalizations of the model exist that incorporated staged mixing of test tubes, “dissolvable” tiles, active signaling across tiles, etc., but here we restrict attention to the model closest to the seeded model of Winfree [9], different from that model only in the absence of a seed and the ability of two large assemblies to attach.

Definitions

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

We now define both the seeded and hierarchical variants of the tile assembly model. A *seeded tile system* is a triple $\mathcal{T} = (T, \sigma, \tau)$, where T is a finite set of tile types, $\sigma : \mathbb{Z}^2 \dashrightarrow T$ is a finite, τ -stable *seed assembly*, and τ is the *temperature*. If \mathcal{T} has a single seed tile $s \in T$ (i.e., $\sigma(0, 0) = s$ for some $s \in T$ and is undefined elsewhere), then we write $\mathcal{T} = (T, s, \tau)$. Let $|\mathcal{T}|$ denote $|T|$. An assembly α is *producible* if either $\alpha = \sigma$ or if β is a producible assembly and α can be obtained from β by the stable binding of a single tile. In this case write $\beta \rightarrow_1 \alpha$ (α is producible from β by the attachment of one tile), and write $\beta \rightarrow \alpha$ if $\beta \rightarrow_1^* \alpha$ (α is producible from β by the attachment of zero or more tiles). An assembly is *terminal* if no tile can be τ -stably attached to it.

A *hierarchical tile system* is a pair $\mathcal{T} = (T, \tau)$, where T is a finite set of tile types and $\tau \in \mathbb{N}$ is the temperature. An assembly is *producible* if either it is a single tile from T , or it is the τ -stable result of translating two producible assemblies without overlap. Therefore, if an assembly α is producible, then it is produced via an *assembly tree*, a full binary tree whose root is labeled with α , whose $|\alpha|$ leaves are labeled with tile types, and each internal node is a producible assembly formed by the stable attachment of its two child assemblies. An assembly α is *terminal* if for every producible assembly β , α and β cannot be τ -stably attached. If α can grow into β by the attachment of zero or more assemblies, then we write $\alpha \rightarrow \beta$.

In either 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 $|\mathcal{A}_\square[\mathcal{T}]| = 1$. If \mathcal{T} is directed with unique producible terminal assembly α , we say that \mathcal{T} *uniquely produces* α . It is easy to check that in the seeded aTAM, \mathcal{T} uniquely produces α if and only if every producible assembly $\beta \sqsubseteq \alpha$. In the hierarchical model, a similar condition holds, although it is more complex since hierarchical assemblies, unlike seeded assemblies, do not have a “canonical translation” defined by the seed position. \mathcal{T} uniquely produces α if and only if for every producible assembly β , there is a translation β' of β such that $\beta' \sqsubseteq \alpha$. In particular, if there is a producible assembly $\beta \neq \alpha$ such that $\text{dom } \alpha = \text{dom } \beta$, then α is not uniquely produced. Since $\text{dom } \beta = \text{dom } \alpha$, every nonzero translation of β has some tiled position outside of $\text{dom } \alpha$, whence no such translation can be a subassembly of α , implying α is not uniquely produced.

Power of hierarchical assembly compared to seeded

One sense in which we can conclude that one model of computation M is at least as powerful as another model of computation M' is to show that any machine defined by

M' can be “simulated efficiently” by a machine defined by M . In self-assembly, there is a natural definition of what it means for one tile system \mathcal{S} to “simulate” another \mathcal{T} . We now discuss intuitively how to define such a notion. There are several intricacies to the full formal definition that are discussed in further detail in [3, 5].

First, we require that there is a constant $k \in \mathbb{Z}^+$ (the “resolution loss”) such that each tile type t in \mathcal{T} is “represented” by one or more $k \times k$ blocks β of tiles in \mathcal{S} . In this case, we write $r(\beta) = t$, where $\beta : \{1, \dots, k\}^2 \dashrightarrow S$ and S is the tile set of \mathcal{S} . Then β represents a $k \times k$ block of such tiles, possibly with empty positions at points \mathbf{x} where $\beta(\mathbf{x})$ is undefined. We call such a $k \times k$ block in \mathcal{S} a “macrotile.” We can extend r to a function R that, given an assembly $\alpha_{\mathcal{S}}$ partitioned into $k \times k$ macrotiles, outputs an assembly $\alpha_{\mathcal{T}}$ of \mathcal{T} such that, for each macrotile β of $\alpha_{\mathcal{S}}$, $r(\beta) = t$, where t is the tile type at the corresponding position in $\alpha_{\mathcal{T}}$.

Given such a representation function R indicating how to interpret assemblies of \mathcal{S} as representing assemblies of \mathcal{T} , we now define what it means to say that \mathcal{S} *simulates* \mathcal{T} . For each producible assembly $\alpha_{\mathcal{T}}$ of \mathcal{T} , there is a producible assembly $\alpha_{\mathcal{S}}$ of \mathcal{S} such that $R(\alpha_{\mathcal{S}}) = \alpha_{\mathcal{T}}$, and furthermore, for every producible assembly $\alpha_{\mathcal{S}}$, if $R(\alpha_{\mathcal{S}}) = \alpha_{\mathcal{T}}$, then $\alpha_{\mathcal{T}}$ is producible in \mathcal{T} . Finally, we require that R respects the “single attachment” dynamics of \mathcal{T} : there is a single tile that can be attached to $\alpha_{\mathcal{T}}$ to result in $\alpha'_{\mathcal{T}}$ if and only if there is some sequence of attachments to $\alpha_{\mathcal{S}}$ that results in assembly $\alpha'_{\mathcal{S}}$ such that $R(\alpha'_{\mathcal{S}}) = \alpha'_{\mathcal{T}}$.

With such an idea in mind, we can ask, “Is the hierarchical model at least as powerful as the seeded model?”

Problem 1. For every seeded tile system \mathcal{T} , design a hierarchical tile system \mathcal{S} that simulates \mathcal{T} .

Another interpretation of a solution to Problem 1 is that, to the extent that the hierarchical model is more realistic than the seeded model by incorporating the reality that tiles may aggregate even in the absence of a seed, such a solution shows how to enforce seeded growth even in such an unfriendly environment that permits non-seeded growth.

Assembly time

We now define time complexity for hierarchical systems (this definition first appeared in [4], where it is explained in more detail). We treat each assembly as a single molecule. If two assemblies α and β can attach to create an assembly γ , then we model this as a chemical reaction $\alpha + \beta \rightarrow \gamma$, in which the rate constant is assumed to be equal for all reactions (and normalized to 1). In particular, if α and β can be attached in two different ways, this is modeled as two different reactions, even if both result in the same assembly.

At an intuitive level, the model we define can be explained as follows. We imagine dumping all tiles into solution at once, and at the same time, we grab one particular tile and dip it into the solution as well, pulling it out of the solution when it has assembled into a terminal assembly. Under the seeded model, the tile we grab will be a seed, assumed to be the only copy in solution (thus requiring that it appear only once in any terminal assembly). In the seeded model, no reactions occur other than the attachment of individual tiles to the assembly we are holding. In the hierarchical model, other reactions are allowed to occur in the background (we model this using the standard mass-action model of chemical kinetics [7]), but only those reactions with the assembly we are holding move it “closer” to completion. The other background reactions merely change concentrations of other assemblies (although these indirectly

affect the time it will take our chosen assembly to complete, by changing the rate of reactions with our chosen assembly).

More formally, let $\mathcal{T} = (T, \tau)$ be a hierarchical TAS, and let $\rho : T \rightarrow [0, 1]$ be a concentrations function, giving the *initial* concentration of each tile type (we require that $\sum_{t \in T} \rho(t) = 1$, a condition known as the “finite density constraint”). Let $\mathbb{R}^+ = [0, \infty)$, and let $t \in \mathbb{R}^+$. For $\alpha \in \mathcal{A}[\mathcal{T}]$, let $[\alpha]_\rho(t)$ (abbreviated $[\alpha](t)$ when ρ is clear from context) denote the concentration of α at time t with respect to initial concentrations ρ , defined as follows. Given two assemblies α and β that can attach to form γ , we model this event as a chemical reaction $R : \alpha + \beta \rightarrow \gamma$. Say that a reaction $\alpha + \beta \rightarrow \gamma$ is *symmetric* if $\alpha = \beta$. Define the *propensity* (a.k.a., *reaction rate*) of R at time $t \in \mathbb{R}^+$ to be $\rho_R(t) = [\alpha](t) \cdot [\beta](t)$ if R is not symmetric, and $\rho_R(t) = \frac{1}{2} \cdot [\alpha](t)^2$ if R is symmetric.

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

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

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

This completes the definition of the dynamic evolution of concentrations of producible assemblies; it remains to define the time complexity of assembling a terminal assembly. Although we have distinguished between seeded and hierarchical systems, for the purpose of defining a model of time complexity in hierarchical systems and comparing them to the seeded system time complexity model of [1], it is convenient to introduce a seed-like “timekeeper tile” into the hierarchical system, in order to stochastically analyze the growth of this tile when it reacts in a solution that is itself evolving according to the continuous model described above. The seed does not have the purpose of nucleating growth, but is introduced merely to focus attention on a single molecule that has not yet assembled anything, in order to ask how long it will take to assemble into a terminal assembly. The choice of which tile type to pick will be a parameter of the definition, so that a system may have different assembly times depending on the choice of timekeeper tile.

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

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

For a shape $S \subset \mathbb{Z}^2$ (finite and connected), define the *diameter* of S to be $\text{diam}(S) = \max_{\mathbf{u}, \mathbf{v} \in S} \|\mathbf{u} - \mathbf{v}\|_1$, where $\|\mathbf{w}\|_1$ is the L_1 norm of \mathbf{w} .

Problem 2. Design a hierarchical tile system $\mathcal{T} = (T, \tau)$ such that every producible terminal assembly $\hat{\alpha}$ has the same shape S , and for some $s \in T$ and concentrations function $\rho : T \rightarrow [0, 1]$, $\Upsilon(\mathcal{T}, \rho, s) = o(\text{diam}(S))$.

It is provably impossible to achieve this with the seeded model [1, 4], since all assemblies in that model require expected time at least proportional to their diameter.

Key Results

Power of hierarchical assembly compared to seeded

Cannon, Demaine, Demaine, Eisenstat, Patitz, Schweller, Summers, and Winslow [3] showed a solution to Problem 1. (They also showed several other ways in which the hierarchical model is more powerful than the seeded model, but we restrict attention to simulation here.) For the most part, temperature 2 seeded systems are as powerful as those at higher temperatures, but the simulation results of [3] hold for higher temperatures as well. In particular, they showed that every seeded temperature ≥ 4 tile system \mathcal{T} can be simulated by a hierarchical temperature 4 tile system (as well as showing it is possible for temperature τ hierarchical tile systems to simulate temperature τ seeded tile systems for $\tau \in \{2, 3\}$, using similar logic to the higher-temperature construction). The definition of simulation has a parameter k indicating the resolution loss of the simulation. In fact, the simulation described in [3] requires only resolution loss $k = 5$.

Figure 1 shows an example of \mathcal{S} simulating \mathcal{T} . The construction enforces the “simulation of dynamics” constraint that if and only if a single tile can attach in \mathcal{T} , then a 5×5 macrotile representing it in \mathcal{S} can assemble. It is critical that each tile type in \mathcal{T} is represented by *more than one* type of macrotile in \mathcal{S} : each different type of macrotile represents a different subset of sides that can cooperate to allow the tile to bind. To achieve this,

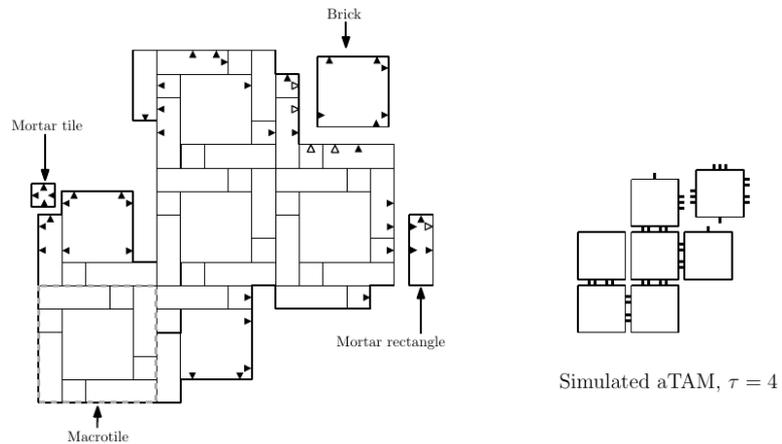


Fig. 1. Simulation of a seeded tile system \mathcal{T} of temperature ≥ 4 by a hierarchical tile system \mathcal{S} of temperature 4. Figure taken from [3]. Filled arrows represent glues of strength 2, and unfilled arrows represent glues of strength 1. In the seeded tile system, the number of dashes on the side of a tile represent its strength.

each macrotile consists of a central “brick” (itself a 3×3 block composed of 9 unique tile types with held together with strength-4 glues) surrounded by “mortar” (forming a ring around the central brick). Figure 1 shows “mortar rectangles” but, similarly to the brick, these are just 3×1 assemblies of 3 individual tile types with strength-4 glues. The logic of the system is such that, if a brick B designed for a subset of cooperating sides $C \subseteq \{N, S, E, W\}$, then only if the mortar for all sides in C is present can B attach. Its attachment is required to fill in the remaining mortar representing the other sides in $\{N, S, E, W\} \setminus C$ that may not be present. Finally, those tiles enable the assembly of mortar in *adjacent* 5×5 blocks, to be ready for possible cooperation to bind bricks in those blocks.

Assembly time

Chen and Doty [4] showed a solution to Problem 2, by proving that for infinitely many $n \in \mathbb{N}$, there is a (non-directed) hierarchical TAS $\mathcal{T} = (T, 2)$ that strictly self-assembles an $n \times n'$ rectangle S , where $n' = o(n)$ (hence $\text{diam}(S) = \Theta(n)$), such that $|T| = O(\log n)$ and there is a tile type $s \in T$ and concentrations function $\rho : T \rightarrow [0, 1]$ such that $\Upsilon(\mathcal{T}, \rho, s) = O(n^{4/5} \log n)$.

The construction consists of $m = n^{1/5}$ stages shown in Figure 2, where each stage consists of the attachment of two “horizontal bars” to a single “vertical bar” as shown in Figure 3. The vertical bar of the next stage then attaches to the right of the two horizontal bars, which cooperate to allow the binding because they each have a single strength 1 glue. All vertical bars are identical when they attach, but attachment triggers the growth of some tiles (shown in orange in Figures 2 and 3) that make the attachment sites on the right side different from their locations in the previous stage, which is how the stages “count down” from m to 1.

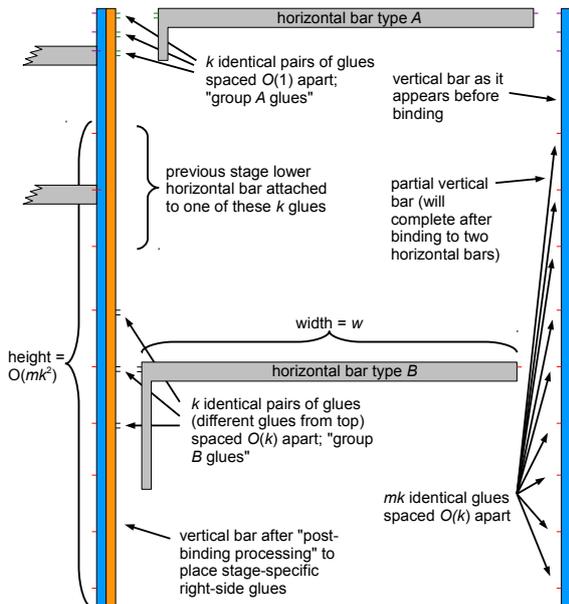


Fig. 3. “Vertical bars” for the construction of a fast-assembling square, and their interaction with horizontal bars, as shown for a single stage of Figure 2. “Type B” horizontal bars have a longer vertical arm than “Type A” since the glues they must block are farther apart.

rectangle, there are many such terminal assemblies. Chen and Doty [4] also showed that for a class of directed systems called “partial order tile systems,” no solution to Problem 2 exists: provably any such tile system assembling a shape of diameter d requires expected time $\Omega(d)$.

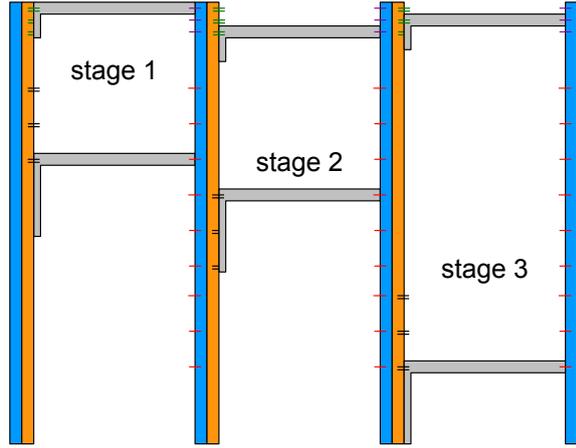


Fig. 2. High-level overview of interaction of “vertical bars” and “horizontal bars” to create the rectangle in the solution to Problem 2 that assembles in time sublinear in its diameter. Filler tiles fill in the empty regions. If glues overlap two regions then represent a formed bond. If glues overlap one region but not another, they are glues from the former region but are mismatched (and thus “covered and protected”) by the latter region.

The bars themselves are assembled in a “standard” way that requires time linear in the diameter of the bar, which is $w = n^{4/5}$ for a horizontal bar and $mk^2 = n^{3/5}$ (where k is a parameter that we set to be $n^{1/5}$) for a vertical bar. The speedup comes from the fact that each horizontal bar can attach to one of k different binding sites on a vertical bar, so the expected time for this to happen is factor k lower than if there were only a single binding site. The vertical “arm” on the left of each horizontal bar has the purpose of preventing any other horizontal bars from binding near it. Each stage also requires filler tiles to fill in the gap regions, but the time required for this is negligible compared to the time for all vertical and horizontal bars to attach.

Note that this construction is not directed: although every producible terminal assembly has the shape of an $n \times n'$

Open problems

It is known [2] that the tile complexity of assembling an $n \times k$ rectangle in the seeded aTAM, if $k < \frac{\log n}{\log \log n - \log \log \log n}$, is asymptotically lower-bounded by $\Omega\left(\frac{n^{1/k}}{k}\right)$ and upper-bounded by $O(n^{1/k})$. For the hierarchical model, the upper bound holds as well [2], but the strongest known lower bound is the information-theoretic $\Omega\left(\frac{\log n}{\log \log n}\right)$.

Question 1. What is the tile complexity of assembling an $n \times k$ rectangle in the hierarchical model, when $k < \frac{\log n}{\log \log n - \log \log \log n}$?

Recommended Reading

1. 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.
2. 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.
3. Sarah Cannon, Erik D. Demaine, Martin L. Demaine, Sarah Eisenstat, Matthew J. Patitz, Robert T. Schweller, Scott M. Summers, and Andrew Winslow. Two hands are better than one (up to constant factors). In *STACS 2013: Proceedings of the Thirtieth International Symposium on Theoretical Aspects of Computer Science*, pages 172–184, 2013.
4. Ho-Lin Chen and David Doty. Parallelism and time in hierarchical self-assembly. In *SODA 2012: Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1163–1182, 2012.
5. Erik D. Demaine, Matthew J. Patitz, Trent Rogers, Robert T. Schweller, Scott M. Summers, and Damien Woods. The two-handed tile assembly model is not intrinsically universal. In *ICALP 2013: Proceedings of the 40th International Colloquium on Automata, Languages and Programming*, July 2013.
6. 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.
7. Irvin R. Epstein and John A. Pojman. *An introduction to nonlinear chemical dynamics: oscillations, waves, patterns, and chaos*. Oxford University Press, USA, 1998.
8. Chris Luhrs. Polyomino-safe DNA self-assembly via block replacement. *Natural Computing*, 9(1):97–109, March 2010. Preliminary version appeared in DNA 2008.
9. Erik Winfree. *Algorithmic Self-Assembly of DNA*. PhD thesis, California Institute of Technology, June 1998.
10. 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.