Chem Soc Rev

REVIEW ARTICLE

View Article Online
View Journal

Cite this: DOI: 10.1039/c6cs00745g

Physical principles for DNA tile self-assembly

Constantine G. Evans [a] and Erik Winfree *[b]

DNA tiles provide a promising technique for assembling structures with nanoscale resolution through self-assembly by basic interactions rather than top-down assembly of individual structures. Tile systems can be programmed to grow based on logical rules, allowing for a small number of tile types to assemble large, complex assemblies that can retain nanoscale resolution. Such algorithmic systems can even assemble different structures using the same tiles, based on inputs that seed the growth. While programming and theoretical analysis of tile self-assembly often makes use of abstract logical models of growth, experimentally implemented systems are governed by nanoscale physical processes that can lead to very different behavior, more accurately modeled by taking into account the thermodynamics and kinetics of tile attachment and detachment in solution. This review discusses the relationships between more abstract and more physically realistic tile assembly models. A central concern is how consideration of model differences enables the design of tile systems that robustly exhibit the desired abstract behavior in realistic physical models and in experimental implementations. Conversely, we identify situations where self-assembly in abstract models can not be well-approximated by physically realistic models, putting constraints on physical relevance of the abstract models. To facilitate the discussion, we introduce a unified model of tile self-assembly that clarifies the relationships between several well-studied models in the literature. Throughout, we highlight open questions regarding the physical principles for DNA tile self-assembly.

Received 15th October 2016

DOI: 10.1039/c6cs00745g

rsc.li/chem-soc-rev

1 Introduction

Molecules that self-assemble into macromolecular structures are found throughout nature. These structures can range from simple homogeneous crystals to complex biological structures such as cytoskeletal microtubules of tubulin, virus capsids, and actin filaments.[1–3] The simplicity and power of such self-assembly is striking: while many synthetic structures, like houses or integrated circuits, are fabricated individually by outside processes, self-assembling structures are able to form from their basic components as a consequence of physical principles, often allowing a multitude of structures to form simultaneously.

The degree of complexity possible in self-assembled structures is a theoretically compelling question. From a mathematical perspective, the study of tiling theory made seminal contributions to the understanding of periodic crystalline structures, such as the classification of all two-dimensional and three-dimensional symmetry groups.[4] Tiling theory also uncovered more complex, and often beautiful, possibilities, such as the existence of aperiodic tilings related to quasicrystals.[5,6] Perhaps more surprising, a direct correspondence between tilings and Turing machines[7,8] showed that the geometry of fitting tiles together can induce patterns of remarkable "algorithmic" complexity. There are even finite sets of tiles that can tile the infinite plane (in the sense that mathematical proof shows that a perfect tiling exists) but for which the resulting pattern is non-algorithmic (in the sense that no Turing machine or other computer can calculate how to correctly place the tiles).[9,10] This "unreasonable effectiveness" of tiling theory suggests that considerations of perfect tilings (equivalently, thermodynamic ground states) do not adequately describe what structures can and will actually form in a physical system. Consequently, complexity questions about self-assembly involve not just thermodynamic principles, but must also carefully treat the kinetics. A theory of tile self-assembly, which treats not the existence of tilings but rather their growth from a seed according to generalized crystal growth principles, has been developed and studied extensively;[11,12] while the "unreasonable effectiveness" of existential tiling theory is eliminated, the connection to Turing machine computation and algorithmic patterns remains intact, and the theory predicts that remarkably complex structures can be self-assembled even from relatively small sets of tiles.

In theories of generalized crystal growth,[16,17] we can identify (at least) three different classes of self-assembly.

Periodic self-assembly includes classical crystal growth in which the same element or set of elements are arranged

[a] Evans Foundation for Molecular Medicine and California Institute of Technology, Physics, Pasadena, CA, USA

[b] California Institute of Technology, Computer Science and California Institute of Technology, Bioengineering, Pasadena, CA, USA. E-mail: winfree@caltech.edu

periodically to create an object of unbounded size, as well as analogous examples from biology such as microtubules that self-assemble from α- and β-tubulin.[3,18] Thus, each monomer type appears within the final structure an unbounded number of times.

Uniquely-addressed self-assembly is in some sense the opposite extreme, wherein each monomer type appears exactly once and has binding interactions that each dictate a specific neighbor. An example from biology would be the eukaryotic ribosome, which self-assembles from nearly a hundred distinct RNA and protein components.

Somewhere in between would be algorithmic self-assembly, which actually encompasses the prior two classes as trivial cases, but which also includes situations where the some or all monomer types appear multiple times in the final structure. Depending on the tile set, algorithmic self-assembly may produce finite structures (which, unlike uniquely-addressed self-assembly, may be much larger than the total number of monomer species) or may produce structures of unbounded size (which, unlike periodic self-assembly, may create patterns of arrangement with no periodicity). The self-assembly of complex virus capsids (such as T4 with its icosahedral head, whiskers, tail, baseplate, and tail fibers[19]) could be considered biological examples of algorithmic self-assembly, in the sense that each distinct protein monomer species appears multiple times in the final structure and some kind of logic during self-assembly is required to determine whether a given monomer will bind to one potential neighbor species or another. The quintessential form of logic used in tile-based algorithmic self-assembly is cooperative attachment, where a certain tile type binds to a growing assembly only in locations where it can make a particular pair of binding interactions. This is sufficient for growth processes that embody Turing machine computations and other complex algorithmic growth behaviors.

Algorithmic self-assembly can be implemented physically using DNA tiles, which have been used to create numerous structures of increasing complexity,[20] as illustrated in Fig. 1. DNA tiles, which can have a variety of structures, attach to each other by short, single-stranded regions of typically 5 to 10 nucleotides, commonly referred to as "sticky ends." The complementary binding of DNA allows binding rules to be programmed by setting the sticky end sequences of each tile type, resulting in programmable self-assembly that can demonstrate many of the theoretically interesting behaviors. The systems can be made to have considerable complexity: even short sticky ends of 5 nucleotides (like those found on double-crossover (DX) tiles[21,22]) can allow 30 to 60 unique sticky end sequences, and longer sticky ends (like those found on single-stranded tiles (SSTs)[23]) can allow for considerably more, according to plausible sequence design criteria.[24] DNA tiles have been used in a wide range of large periodic structures[20] and structures with uniquely-addressed tiles that have ranged up to a thousand unique tiles in three dimensions.[14,25]

Experimental implementation of algorithmic self-assembly, however, has posed several challenges. Of primary importance, two types of assembly defects can occur that result in undesired structures: growth errors where tiles attach in locations where they have imperfect bond matches, and spurious nucleation where tiles assemble without a seed.[26–30] Considerable progress has been made in overcoming these challenges, and increasingly complex structures have been successfully built.[30,31] This progress has benefited from refinements in experimental technique, but derives more significantly from a better understanding of the physics behind DNA tile assembly: simple models of growth have motivated fundamental changes to tile system design and experimental protocols.

In this review, we examine models of DNA tile growth, both more abstract and more physically-motivated, and show how

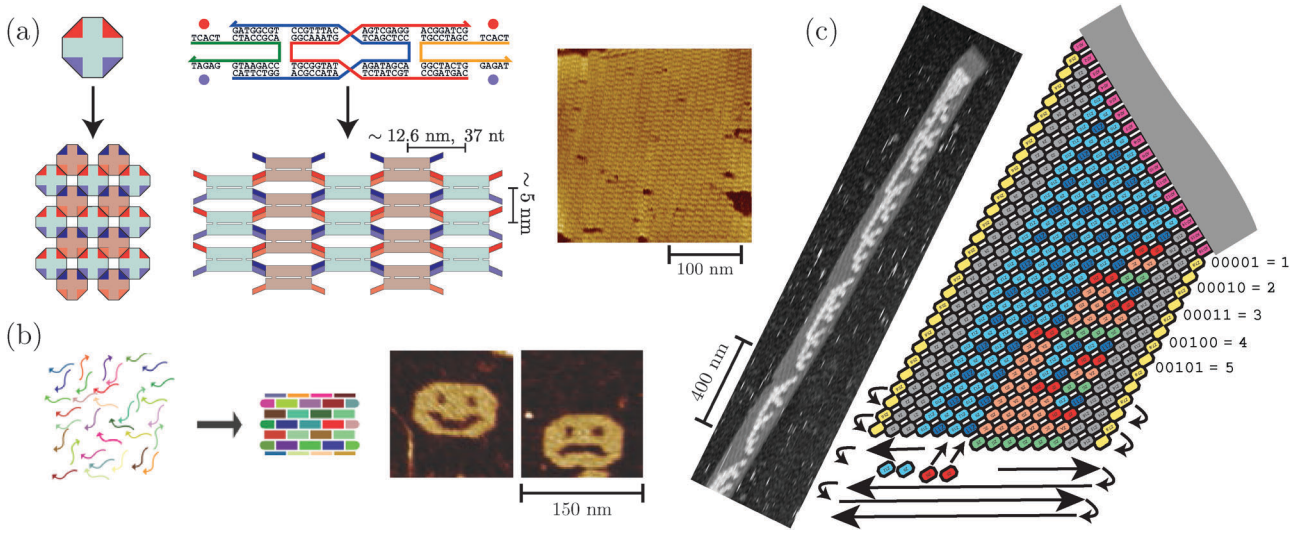

Fig. 1 Examples of types of DNA tile self-assembly. (a) Shows the abstract representation, tile structure, lattice representation and atomic force microscope (AFM) image of a periodic tile system of DAO-E tiles;[13] different shades represent end/complement sticky end pairs. (b) Shows a lattice representation and two AFM image of a uniquely-addressed tile system, where subsets of tiles that assemble into a square have been added in order to assemble different shapes.[14] (c) Shows a representation and AFM image of a large, algorithmic tile system implementing a binary counter ribbon.[15] Figure components adapted from ref. 13–15, with permission of authors.

tile system behavior changes when moving from model to model. We begin in Section 2 with the example of systems designed to grow in a specific abstract model, the abstract Tile Assembly Model (aTAM), which has been well-studied in theoretical works and experimental implementations. In subsequent sections, we then consider the phenomena that arise when transitioning through a series of more physically accurate models, and how the effects of these phenomena can be mitigated through design and growth constraints in order to implement experimental systems that behave in approximately equivalent ways to their abstract counterparts.

In Sections 2 and 3, we show how the transition from the aTAM to a kinetically-derived model, the kinetic Tile Assembly Model (kTAM), allows errors in growth to arise, and how these can be modelled through kinetic and thermodynamic methods and ameliorated through tile system design methods referred to as "proofreading." In Section 4, we consider the nucleation phenomena that must be addressed when transitioning from models that start from pre-determined seeds to models that can start from any tile–tile attachment event, and explain methods that have been developed to control tile system nucleation. In Section 5, we explain the challenges to preserving tile system behavior when transitioning to a model that considers the depletion of free tile monomers, and explain experimental methods that have been used to avoid such depletion and complexities.

We then examine more formally in Section 6 the relationships between models by presenting a unified framework for models of tile assembly. Choices of reactions and parameters in this framework can result in models equivalent to or closely related to several well-studied abstract models, like the aTAM and the Two-Handed Assembly Model (2HAM), as well as more physically accurate variations, like the kinetic Tile Assembly Model (kTAM). Using this framework, we briefly discuss some of the phenomena that arise when considering tile systems in variants of the aTAM, 2HAM, and other models, and the constraints these phenomena place on the tile systems' physical relevance.

2 Basic models

Two complementary models of tile assembly are widely studied, providing a starting point for discussion.[26] The abstract Tile Assembly Model (aTAM) derives from the basic concept of tile attachment *via* matching bonds; it provides insight into the potential computational power of algorithmic assembly, and allows for intuitive design and understanding of systems. The kinetic Tile Assembly Model (kTAM) is instead based on the kinetics and thermodynamics of tiles in solution binding by reversible bonds; this model serves as the basis for most of what we discuss in this review.

In the aTAM, tile systems consist of non-rotatable square tiles with up to four sticky ends of typically integer strengths: usually 1 ("single-strength" or "weak") or 2 ("double-strength" or "strong"). Growth begins from a designated seed tile and subsequent tiles then attach asynchronously, one at a time. Among all possible tiles, attachment is possible only if the tile's sticky ends match those on adjacent tiles in the assembly with strengths that sum to at least the system "threshold" τ. τ is often described as "temperature" in theoretical papers, but is related physically to the relative favorability of attachment and detachment, which could depend on a number of parameters, including temperature, salt concentrations, and tile concentrations. Once attached, tiles never detach.

Fig. 2 illustrates this process for a tile system at $\tau = 2$ that implements exclusive-or (XOR) logic. As the system grows from an initial seed tile at the bottom left, tiles with double-strength

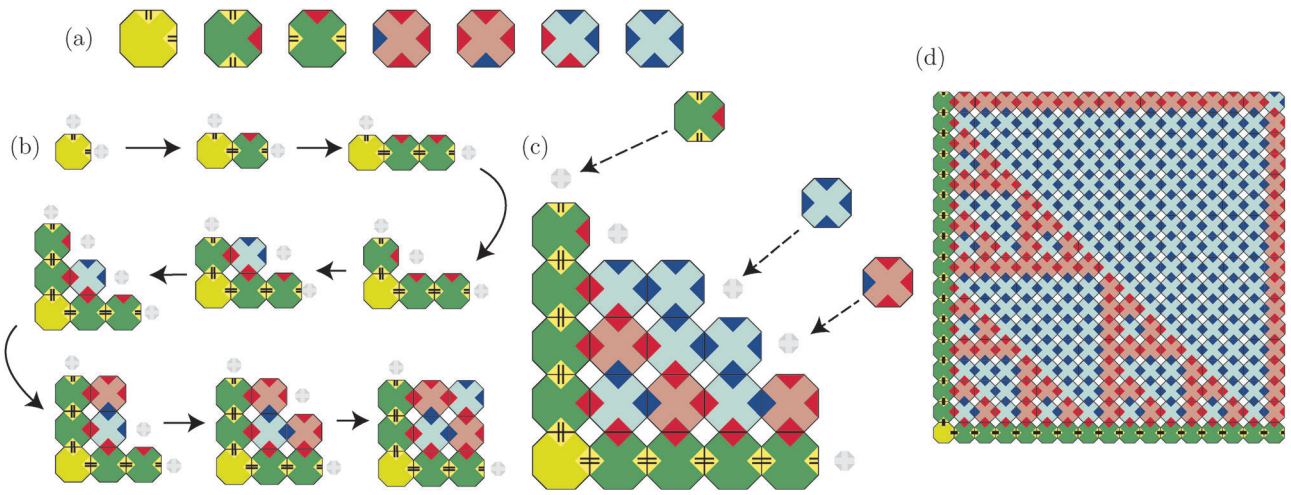

Fig. 2 A $\tau = 2$ XOR/Sierpinski pattern tile system in the abstract Tile Assembly Model (aTAM). (a) Shows the 7 tiles in the set: a yellow seed tile with double-strength ends (denoted by two lines), two boundary tiles, and four rule tiles. These are called the exclusive-or (XOR) rule tiles because the output bonds are red if and only if the input bonds are different types. Tiles are colored according to their output bonds. (b) Shows a possible series of attachments, with gray tiles indicating potential attachment sites. (c) Shows the possible tiles that can attach to a small assembly, while (d) shows a larger 17 × 17 assembly that would continue to grow.

bonds form a V-shaped edge at the left and bottom, while tiles implementing the XOR logic in the center attach by two single-strength bonds: one from the tile left of the site, and one from the tile below. The result is a Sierpinski triangle pattern that grows out infinitely to the top right—a non-periodic fractal pattern that is a consequence of the XOR function, and that could not be produced by simple periodic tiling.

This cooperative attachment of tiles by two weak bonds to two different tiles, where each bond is individually too weak to allow attachment, makes growth at $\tau = 2$ particularly interesting. $\tau = 2$ tile systems in the aTAM are capable of Turing-universal computation.[26,32] Squares[33,34] and arbitrary shapes[35] can be efficiently assembled by relatively few tile types, and numerous other properties have been shown.[11,12] These results rely on cooperative attachment: with $\tau = 1$, at least for certain classes of tile systems, significant computation cannot be performed or is much more difficult.[36]

The aTAM is useful for considering the theoretical potential of tile assembly, but ignores important aspects of physical implementations. The kTAM derives from a different approach, starting from basic physical considerations. DNA tiles, in general, are monomers with ends that form reversible bonds to matching sticky ends on other tiles. As tile systems involve large numbers of tiles in solution, tile assembly can therefore be considered from the perspective of chemical kinetics and thermodynamics for tile attachment and detachment.

Here we focus on the growth of a single crystal in solution, although later we will discuss a mass-action version of the kTAM and other variants. To simplify the analysis for now, this most basic version of the kTAM has a few fundamental assumptions:

(1) Free tiles are assumed to be in solution, at equal and constant concentrations never depleted by assembly, and to remain well-mixed, so diffusion can be ignored. Thus growth of individual crystals will be independent of the growth of other crystals in the solution.
(2) We only (for now) consider growth initiating from specific seed tiles, as in the aTAM, and only consider single free tiles (monomers) attaching to and detaching from a well-formed crystal.
(3) Tiles attach based only on diffusion. Since we are considering the growth of an individual crystal in solution, attachment of a tile to an individual lattice site will depend only on the concentration of the tile. In a marked departure from the aTAM, we assume that bonds make no difference here: the attachment rate will be the same whether the tile attaches by one bond or three, or has a bond that does not match. Tiles that could not attach in the aTAM have the same chance of attaching as any other tile.
(4) Once attached, tile detachment is assumed to depend only on the total change in free energy for all the correct bonds holding it to the lattice. We will assume that the standard free energies of correct bonds can simply be added (*i.e.* cooperative attachment does not increase or decrease the favorability of the individual attachments) and that mismatched bonds do not hinder binding.[13]

We will consider bonds in comparison to a reference standard free energy of some typical single bond, ΔG°_{se}, with se referring to "sticky end", as the term is commonly used in experimental literature to refer to a bond. A tile will thus be attached by some free energy $b\Delta G^\circ_{se}$, where b is the sum of the "strengths" of all correct bonds compared to ΔG°_{se}. In many theoretical analyses, all bonds are assumed to be identical, with strength 1 (*i.e.*, ΔG°_{se}), or bonds are assumed to have some integer multiple strength (*e.g.*, ΔG°_{se} or $2\Delta G^\circ_{se}$), but the kTAM does not require this, and there have been some investigations of systems with differing or sequence-dependent bond strengths.[24]

These assumptions result in the following rates for attachment and detachment of a single tile at a single lattice site:

$$r_f = k_f[c] \quad r_{r,b} = k_{r,b} = k_f u_0 e^{b\Delta G^\circ_{se}/RT+\alpha} \tag{1}$$

where k_f is a forward rate determined experimentally, $[c]$ is the tile concentration, b is the total strength of the correct bonds between the tile and adjoining tiles, u_0 is a standard concentration (1 M), and α is a constant unitless free energy change from other factors, such as the loss of rotational entropy during binding. These rates, illustrated in Fig. 3, describe the events occurring on a single crystal grown from a seed.

In the kTAM literature, these rate equations are typically used in an equivalent but more symmetric form. By defining $\hat{k}_f \equiv u_0 k_f e^{\alpha}$, $G_{se} \equiv -\Delta G^\circ_{se}/RT$, and G_{mc} such that $[c] = u_0 e^{-G_{mc}+\alpha}$, we can construct the symmetric rate equations

$$r_f = \hat{k}_f e^{-G_{mc}} \quad r_{r,b} = \hat{k}_f e^{-bG_{se}} \tag{2}$$

that are the core of the kTAM.[26] In this form, $\hat{k}_f$ is a reference rate for tile dissociation in units of s^{-1}. G_{se} and G_{mc} are two unitless free energy analogues. G_{se} is the sign-reversed, dimensionless reference free energy of a typical single sticky end (se) bond, where higher numbers correspond to stronger binding and thus longer times before detachment. G_{mc} is a logarithmic term for the free monomer concentration (mc), where larger numbers correspond to exponentially lower concentrations and longer times between attachments. Arranged in this manner, rate equilibrium can be easily understood as a comparison between G_{mc} and G_{se}: if, for example, G_{mc} is larger than G_{se} but less than double G_{se}, then the attachment rate for a strength 2, or $b = 2$ interaction, will be higher than its detachment rate, but the attachment rate for a $b = 1$ interaction will not.

For a given tile set, seed tile, and values for G_{mc} and G_{se}, the standard kTAM is a continuous-time Markov chain (CTMC) that satisfies detailed balance. The initial state is an assembly consisting only of the seed tile, and every subsequent step involves the reversible addition or removal of a non-seed tile. (As we are identifying the crystal by the seed, removal of the seed tile itself is not considered.) Thus, considering a tile addition to an assembly A that results in an assembly A':

$$\frac{r_f}{r_{r,b}} = e^{bG_{se}-G_{mc}} = e^{-\Delta G} \tag{3}$$

where $\Delta G = G(A') - G(A)$ with respect to the energy of an assembly, $G(A) = NG_{mc} - BG_{se}$ where N is the total number of

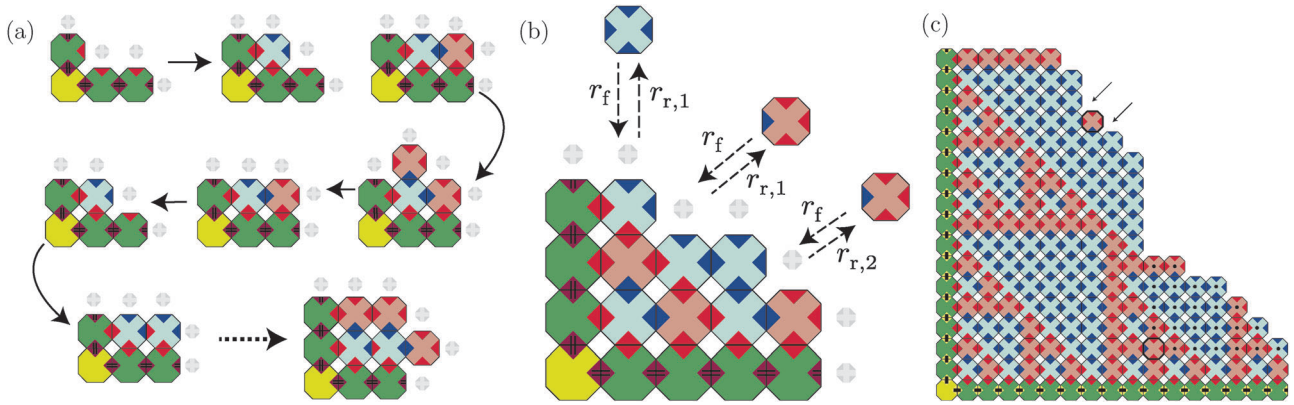

Fig. 3 The Sierpinski system from Fig. 2 in the kinetic Tile Assembly Model (kTAM). (a) Shows one possible path of attachments and detachments, disregarding the time between each change. (b) Shows attachment and detachment rates for various tiles and lattice sites. (c) Shows the a larger assembly, as in Fig. 2. Errors are shown in bold: at the top, a growth error is show, with errors indicating potential subsequent attachments that could "trap" the error in place in the kinetic trapping model, while at the bottom, a single error causes a different pattern to arise, despite later tiles (dotted) attaching by correct bonds.

tiles in the assembly and B is the summed strength of all matching bonds in the assembly. Consequently, the equilibrium probability of an assembly A, and the associated partition function, conform to the Boltzmann distribution:

$$P(A) = \frac{1}{Z} e^{-G(A)} \quad Z = \sum_A e^{-G(A)}. \tag{4}$$

While this formula is convenient and can be insightful, it must be used with caution because in many cases of interest the CTMC state space is infinite, the partition function does not converge, and there is no reachable equilibrium.

The kTAM employs assumptions that are not valid in general: for example, tiles may have flexible structures that reconfigure upon attachment, binding strength may not simply be a sum of individual bond strengths, or experiments might incorporate growth on a surface. However, results from DX tile ribbon and nanotube growth in solution largely fit predictions and simulations in the kTAM,[37,38] and many experimental results qualitatively fit the model.[28,29,39] Departures from kTAM assumptions often have effects that can be incorporated into renormalized parameter values, and in experimental measurements, these effects are often built in to measured values, particularly G_{se} and α. For example, while directly imaging growth on a surface with an atomic force microscope (AFM) involves surface binding terms and an overall force from the AFM tip, as approximately constant terms, these will both be incorporated into a significantly different α in measurements, as seen in experiments.[13] More generally, so long as binding is at least dependent primarily upon the total bond strength b and some constant free energy of a bound tile, then G_{se} and α can be chosen to make any two values of b at which the kTAM detachment rate will be a good approximation, regardless of nonlinear or constant effects. As discussed later, in many cases, tile system growth in the kTAM is most interesting when tiles attached by $b < 1$ fall off very quickly, and tiles attached by $b > 2$ fall off very slowly, with $b = 1$ and $b = 2$ involving the most relevant rates. In this regime, even if there are significant departures from the kTAM's assumptions, so long as parameters cause the $b = 1$ and $b = 2$ cases to be approximately correct, the kTAM will still give reasonable results. It is thus hoped that such an approximation will be suitable for SSTs, even though, as single strands with no rigid structure, they must significantly change configuration during attachment.

Growth on a surface, not the focus of the kTAM, could also involve multiple diffusion modes, with some tiles in solution attaching directly to crystals, and others binding to the surface itself and sliding in two dimensions. These two diffusion modes could, for example, result in attachment rates dependent upon lattice site location, with sites on the edges of crystals having more frequent attachments than those in the interior of crystals, and would require a different model. In direct tile-level atomic force microscope observation of DX tile crystal growth near equilibrium on a surface, this was not found to be a significant factor. Single-tile attachments and detachments rates fit the kTAM within error, but had α and G_{se} values incorporating surface attachment and AFM effects that were significantly different from typical values in solution.[13] With other growth conditions or tile types, surface diffusion could be significant: for example, surface diffusion effects have been modelled and seen experimentally when using larger tile structures made from T-shaped junctions or DNA origami.[40–42] However, for the purposes of this review, we consider only growth in solution.

Additionally, lattice defects and assembly–assembly interactions have been observed, motivating designs and conditions to avoid them,[13,30] and nucleation and free tile depletion are notable deviations discussed later in this review. Slightly different tile structures, such as those with hairpin labels, have also been found to potentially have significantly different energetics, making their use in systems difficult.[30,37]

The kTAM can be applied to 1, 2, and 3D lattices. Most theoretical analyses and simulations to date have not considered 3D lattices, though there has been some theoretical work on 3D error-correcting schemes.[43] Uniquely-addressed 3D structures

of SSTs have been grown experimentally,[25] and while algorithmic systems have primarily used 2D lattices, in some cases these have incorporated tube topologies.[23,28,29,44]

To more accurately model particular situations, several extensions to the kTAM have been studied. These have included models that account for DNA hybridization energetics and unequal concentrations,[24,45] and mass-action and fixed-volume stochastic models[27,46] that allow analysis of nucleation and crystal scission. However, other phenomena, such as lattice defects and interactions between assemblies, have not yet been well studied.

Simulations of the kTAM are most often discrete, Gillespie algorithm simulations of the growth of single crystals. The two most commonly-used simulators are Xgrow and ISU TAS, which support both aTAM and kTAM simulations, and have some support for kTAM extensions.[45,47,48]

3 Errors and proofreading

3.1 Connections between the aTAM and kTAM

The same tile systems can be examined in both the aTAM and kTAM, and, as discussed later, in numerous other models. That they will exhibit similar behaviors in the aTAM and kTAM is not, however, immediately apparent. As theoretical investigation of the computational power behind tile self-assembly and the intuitive logical design of tile systems is much simpler in the aTAM, while the kTAM more closely corresponds to physical implementations, it is important to understand the situations and designs where tile systems will behave similarly in both models.

Tiles in the kTAM attach with equal rates regardless of the number of correct bonds, and always have some chance of detaching. Thus, depending upon relative attachment and detachment rates, crystals can grow or shrink, unlike in the aTAM. For a given lattice site, a tile attaching by two bonds will have equal attachment and detachment rates when $G_{mc} = 2G_{se}$. Thus, if all tiles attach by two bonds, as is usual for a $\tau = 2$ aTAM system, this G_{mc} corresponds to a system at its melting temperature. To examine growth in the kTAM, we define ε such that $G_{mc} = 2G_{se} - \varepsilon$. This corresponds to the supersaturation of tiles in solution, and the extent to which tile attachment is faster than tile detachment. When ε is negative, crystals will melt if the tiles are attached on average by strength 2 or less.

The kTAM's connection to the aTAM can be seen intuitively by considering a large G_{mc} and small positive ε, corresponding to low tile concentrations at slight supersaturation, with a correspondingly large G_{se}. In this situation, tiles attach at an extremely slow rate $r_f = \hat{k}_f e^{-G_{mc}}$. A tile attaching by b bonds will fall off at a relative rate $r_{r,b}/r_f = e^{(2-b)G_{se}-\varepsilon}$. Tiles attaching by zero or one bonds will fall off much faster than the time scale of tile attachment, and can be ignored, while tiles attached by three or more bonds will be firmly attached at that time scale. Tiles attached by two bonds will detach at a relative rate $r_{r,2}/r_f = e^{-\varepsilon} < 1$, resulting in a random walk that will be slightly biased forward. As long as we consider tile systems that are deterministic in the aTAM, so that there is always one correct tile that can attach in a location, the random walk will be equivalent to growth in the aTAM.[26]

This equivalence is only in the limit of slow growth, where $r_f \to 0$. For faster growth, with smaller values of G_{mc} or larger values of ε, other behaviors can arise. A tile attaching by one or no correct bonds will still tend to fall off faster than it attaches, but will be attached for a non-negligible time. If another tile attaches adjacent to it before it detaches, as in Fig. 3(c), then the erroneous tile could end up attached by two or more bonds, and growth could continue with the error still in place. While less likely than correct tile attachment, this process causes growth errors that could not have taken place in the aTAM. As correct algorithmic self-assembly depends on prior growth, a single error can cause growth to drastically change: in the Sierpinski system, for example, a single error in a sea of zeros can initiate an entire Sierpinski pattern, as seen in Fig. 3(c).

Growth errors present a significant obstacle for algorithmic growth at practical speeds, and were found in simulations and experimental results.[28] Fortunately, growth errors in the kTAM can be simply modeled, and through clever tile system design, error rates can be drastically reduced.

As a first assessment, we note that at equilibrium, the ratio of the probability of a correct assembly A to the probability of an assembly A' of the same size but with a single mismatched bond, can be easily computed from the Boltzmann distribution (eqn (4)). Specifically, $P(A')/P(A) = e^{-G_{se}}$ provides a good approximation to the per-tile error rate. However, there is no *a priori* guarantee that growth will approach equilibrium on reasonable time scales; for example, setting $\varepsilon = 0$ dramatically slows down growth. Worse yet, for tile sets that permit unbounded growth such as the Sierpinski system, any $\varepsilon > 0$ that biases growth forward also ensures that the partition function is infinite and equilibrium does not exist. Thus, a more nuanced approach is needed in order to understand error rates at finite growth speeds.

We can construct the **kinetic trapping model** by formalizing our intuitive reasoning about how growth errors occur.[26] Consider a single, empty lattice site where a unique "correct" tile can attach by two bonds and m "almost-correct" tiles can attach by only one bond (tiles attaching by no correct bonds will detach very quickly and can be ignored to first order). The attachment rate for each tile type will be the same, $r_f = \hat{k}_f e^{-G_{mc}}$, but tiles will detach at different rates. If we consider the possible states that the lattice site can be in—empty (E), or filled with a correct tile (C) or an almost-correct (A) tile—we obtain the transition rates shown in Fig. 4.

In a filled state, there is a possibility that a further tile can attach in an adjacent site by two bonds, and growth will continue from this point, "trapping" the initial attachment in place regardless of correctness. This may not be possible—there may be no tile type that can attach by the combination of bonds present[24]—but in the worst case, such an attachment will be possible in two neighboring sites regardless of whether the correct or almost-correct tile attached originally. (This worst case occurs everywhere in direct implementations of cellular automata such as the XOR tile set, because there exists a rule tile that can bind to any given input pair.)

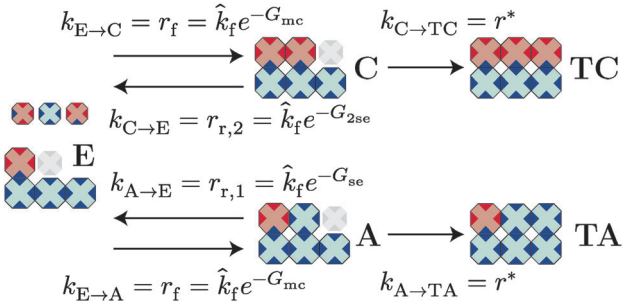

Fig. 4 Transition rates in the kinetic trapping model, as described in the text.

The rate at which this trapping will take place, r^*, is important to the derivation of error rates but is not easily determinable directly. For example, for small ε the growth front behaves as a random walk and will wash back and forth over any given site many times before never visiting it again. It is reasonable, however, to consider it as proportional to the overall growth rate of a crystal r in "layers per second", with an adjustable parameter $r^* = \beta r$, as the time for a tile to become strongly held by further growth is related to how quickly further layers grow. The overall growth rate of the crystal will be dependent upon the tile system, and particularly the shape of the growth front. For a system with only single-strength bonds, ignoring detachments of tiles held by more than two bonds and attachments by less than two bonds, the growth rate in tiles per second will be $r_{gt} = r_f \cdot$(# of $b = 2$ growth sites) $- r_{r,2} \cdot$(# of $b = 2$ detachment sites). Considering a length L growth front, if sites exist at every empty adjacent lattice point, and every tile is attached by $b = 2$, then the two numbers will be the same, and equal to the length of the front, making $r = r_{gt}/L = r_f - r_{r,2}$. If, alternatively, the growth front consists of facets of length l distributed as $P(l) = 2^{-l}$, each with one potential attachment site and one potential detachment site, then the rate will instead become $r = \frac{1}{2}(r_f - r_{r,2})$. For a system with a growth front of exactly one attachment and one detachment site (as in the zig-zag ribbon discussed later), then $r = \frac{1}{L}(r_f - r_{r,2})$, where L is the length of the layer. Thus $r = \gamma(r_f - r_{r,2})$, for some γ: for many systems $\gamma = \frac{1}{2}$ is a reasonable choice. r^* then becomes $r^* = \beta\gamma(r_f - r_{r,2})$.

Taken together, the kinetic trapping model results in transition rates, with the addition of "trapped" states for a correct tile (TC) or almost-correct tile (TA), such that

$$\frac{d}{dt}P(t) = \begin{array}{c} \\ E \\ C \\ A \\ TC \\ TA \end{array} \begin{array}{c} \begin{array}{ccccc} E & C & A & TC & TA \end{array} \\ \begin{pmatrix} -(1+m)r_f & r_{r,2} & r_{r,1} & 0 & 0 \\ r_f & -r_{r,2}-r^* & 0 & 0 & 0 \\ mr_f & 0 & -r_{r,1}-r^* & 0 & 0 \\ 0 & r^* & 0 & 0 & 0 \\ 0 & 0 & r^* & 0 & 0 \end{pmatrix} \end{array} P(t) \equiv M \cdot P(t) \tag{5}$$

where $P(t)$ is a probability vector of the site being in a state at time t. The per-site error rate for growth, according to this model, is the probability of reaching TA rather than TC when starting at E. One way of solving for this hitting probability is to consider a flow problem where probability is injected into E and we examine the accumulation of probability in TA and TC, in other words, considering $\frac{d}{dt}P_{TC}$ or $\frac{d}{dt}P_{TA}$ when $\frac{d}{dt}P = \frac{d}{dt}[0\,0\,0\,P_{TC}\,P_{TA}]^T = MP + [1\,0\,0\,0\,0]^T$. Then

$$P_{error} = \frac{d}{dt}P_{TA} = \frac{1}{1 + \frac{r_{r,1} + r^*}{m(r_{r,2} + r^*)}} \approx \frac{m(r_{r,2} + r^*)}{r_{r,1} + r^*} \tag{6}$$

with the approximation on the right assuming that the error rate is small and the system is near equilibrium.

This unwieldy result can be approximated further by assuming that $r_{r,1}$ is much larger than r^*, substituting in G_{mc} and G_{se}, and assuming growth where $G_{mc} = 2G_{se} - \varepsilon$, resulting in $P_{error} = (\beta\gamma e^{\varepsilon} - \beta\gamma + 1)e^{-G_{se}}$. It is apparent that β and γ will have no effect if $\varepsilon = 0$: β and γ change how much of an effect increasing ε has on error rates. When $\beta\gamma = 1$,

$$P_{error} \approx m e^{-G_{se}+\varepsilon}. \tag{7}$$

The kinetic trapping model allows us to obtain a global view of the growth process. The error rate will be dependent upon both G_{se} and, through ε, G_{mc}, and both of these parameters will affect the growth rate r^* of the system. If $\varepsilon < 0$, then $r^* < 0$, and no growth will occur. If $G_{mc} < G_{se}$, such that even a single bond is favorable, $\varepsilon \gg 0$ and the error rate will be extremely high, leading to unordered growth. Algorithmic growth will take place only in the region where $G_{se} < G_{mc} < 2G_{se}$, with the error rate changing throughout the region. The phase diagram for these parameters is shown in Fig. 5.

While the kinetic trapping model is derived from kinetic considerations, the $\varepsilon = 0$ case fits a thermodynamic equilibrium perspective well. Per the partition function in eqn (4), the equilibrium probability of an assembly missing one correct G_{se} bond as compared to an assembly with that bond will be $P_{incorrect} = e^{-G_{se}}P_{correct}$. If there are m such possible assemblies, then the combined probability of an incorrect assembly *versus* a correct one will be $P_{incorrect} = me^{-G_{se}}P_{correct}$, which matches the $\varepsilon = 0$ case of the kinetic trapping model. Intuitively, for a system at the $\varepsilon = 0$ thermodynamic equilibrium, the prevalence of errors is directly related to the energetic penalty those errors incur, though for an infinite system, this thermodynamic equilibrium will be unreachable. When systems are instead grown away from equilibrium, and thus faster, with $\varepsilon > 0$, the probability of erroneous assemblies correspondingly increases from the equilibrium probability.

The error rates of several tile systems have been examined in discrete Gillespie simulations, and have matched well with the kinetic trapping model.[26]

The limit of infinitely slow growth in this model is the limit, for fixed ε, as $G_{se} \to \infty$: the limit of low concentrations and high bond strengths. In this limit, the error rate goes to zero: by

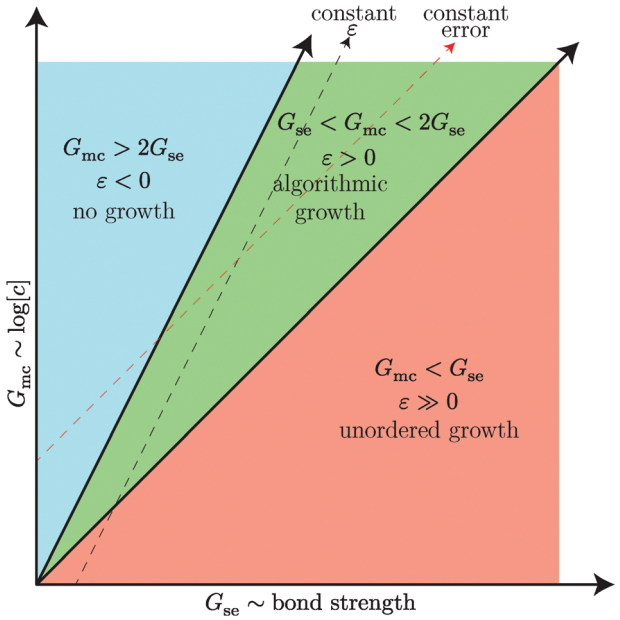

Fig. 5 Phase diagram for G_{mc} ("monomer concentration") and G_{se} ("sticky end") in the kTAM with a $\tau = 2$ system. For $G_{mc} > 2G_{se}$, detachment rates are higher than attachment rates of tiles attached total strength 2, and there is no growth. For $G_{mc} < G_{se}$, tiles can attach favorably by one single-strength bond, and thus growth is unordered. Algorithmic growth by two weak bonds or one strong bond lies in the region between these two, where attachment by total strength 2 is favorable but attachments by a single weak bond are unfavorable. Error rates and ε are explained in the text.

slowing down growth, growth errors can be reduced arbitrarily. However, the error rate does not decrease quickly with slower growth rate. The crystal growth rate is $r = \gamma e^{-2G_{se}}(e^{\varepsilon} - 1)$, so for small ε, $r \sim P_{error}^2$. Thus, close to equilibrium, growth rate and error rate are fundamentally linked: a 10-fold improvement in error would require a 100-fold slowdown of growth. To build an $N \times N$ crystal with high probability would require N^2 tiles with an error rate of $\sim N^{-2}$, necessitating a growth rate $\sim N^{-4}$. Assuming all N layers grew in parallel, the time to assemble the crystal would scale as N^5, implying a substantial slowdown for large crystals.[49,50]

The above analysis assumes, as per the basic kTAM, that mismatched bonds don't have an energy penalty or contribution, and can simply be ignored. In practice, the finite sequence space of DNA means that some partial matching is likely to be possible between mismatched sticky ends, and so mismatched bonds may add to the total bond strength of an incorrect tile.[24] If this contribution is interpreted as a fraction s of the correct G_{se}, and the same kinetic trapping analysis is used, the resulting error rate is

$$P_{error} \approx m e^{-(1-s)G_{se}+\varepsilon}. \quad (8)$$

While this contribution is dependent on sequence design and is different for each pair of sticky end sequences in a system, a value $s \approx 0.5$ is reasonable as a rough approximation.[24]

The experimental system of Fujibayashi *et al.* offers an opportunity to compare error rates.[30] The average error rate in early growth for their XOR system was 1.4% ± 1.1% per tile. In a XOR system with two possible almost-correct attachments at every site, $m = 2$. The XOR tiles were at a concentration of 50 nM, and the standard initiation energy for hybridization[51] can provide a reasonable assumption of $\alpha = \ln 20$, resulting in $G_{mc} = \alpha - \ln(u_0^{-1}[c]) \simeq 19.8$. If we assume growth near equilibrium ($\varepsilon \approx 0$), then, with no mismatch contribution, the theoretical error rate is 0.01%, much lower than the experimental result. However, with $s = 0.5$, the theoretical error rate becomes 1.4%, showing the importance of contributions from partially matching sequences, though the surprising closeness of this simple theoretical result to the experiment is clearly coincidental. From the reported protocol, we estimate that the crystals grew over a period of approximately 1 hour during annealing.

These experiments have shown that algorithmic assembly can be implemented, and a 1.4% error rate is enough to demonstrate algorithmic behavior in crystal growth. However, it is not enough to execute more sophisticated algorithmic processes. In order to build error-free, complex structures using algorithmic self-assembly at high yields, extremely long assembly times would be required: a crystal with a 0.01% error rate, for example, would require growth over roughly a year. To allow for viable, complex algorithmic self-assembly, some other way to reduce error rates is necessary.

3.2 Proofreading

The design of fault-tolerant systems is a widely studied problem. Tile systems could be designed where errors could be accommodated as faults in assembly that would not affect the logic of growth; similar designs have been studied for cellular automata.[52] Alternatively, the structure of individual tiles could be changed to depart from the "passive" behavior of tiles in the kTAM, incorporating "active" features through strand displacement, locking and unlocking of ends, or passing signals upon attachment, perhaps akin to a hybridization chain reaction system.[53,54] The behaviors and potential of several abstract designs have been studied,[55–58] but research into structures that can implement them remains in its early stages, with one-dimensional signaling tiles using strand displacement only recently being experimentally implemented.[59] These approaches lack the molecular simplicity of passive self-assembly and present more substantial challenges in the laboratory.

By carefully considering the kinetic trapping model, however, another solution can be found that is specific to the kTAM and surprising in its simplicity. Errors arise in the kinetic trapping model when erroneously attached tiles are trapped in place by continued growth. This may not always be possible: it requires the presence of a tile type that can attach by two bonds in one of the sites adjacent to the erroneous tile. If no tile can, continued growth would require a further erroneous attachment to take place, while the initial erroneous tile would have more time to detach.

The uniform proofreading construction of Winfree and Bekbolatov[60] therefore replaces each tile in a tile system with a $K \times K$ block of unique tiles that attach to the lattice individually. The ends on each external edge of the block

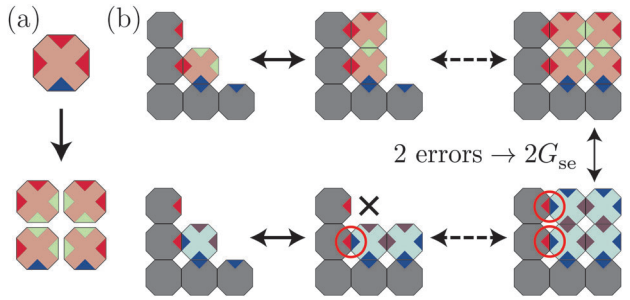

Fig. 6 The uniform proofreading construction of Winfree and Bekbolatov.[60] (a) Shows the transformation of a single tile into a 2 × 2 block of tiles, with unique internal bonds that for simplicity are not distinguished here; bonds along each edge are also not distinguished. (b) Shows the attachment of a correct *vs.* incorrect tile in a block; for the incorrect block to form, two erroneous attachments must take place.

correspond to the ends on the original tile and their locations, while the internal sticky ends are all unique to that block, as illustrated in Fig. 6.

While a single tile in a block may attach by one correct and one incorrect bond, the unique internal bonds, combined with the additional bonds on the side of the block, necessitate further incorrect attachments to fill the block. Although tiles attach one by one, the $K \times K$ area will typically fill up with a mutually consistent block: either a correct block with 0 mismatches, or an "almost-correct" block with K mismatches along one side. Near thermodynamic equilibrium—if it is reached—where the Boltzmann distribution gives a reasonable approximation of error rates, the K tiles on one side of the K^2 block will all be incorrect with a probability $\sim e^{-KG_{se}}$, resulting in an error rate, for m almost-correct blocks, of

$$P_{error} \approx me^{-KG_{se}}. \quad (9)$$

At this equilibrium level, proofreading thus significantly reduces growth errors as compared to the equivalent non-proofreading tile set, by a factor of $e^{-(K-1)G_{se}}$. This exponential decrease in error rate is gained at the cost of a K-fold increase in the linear size of the assembled object, and thus, assuming the same tile concentrations, a K-fold increase in the time required for assembly.

Unlike the kinetic trapping model for systems without proofreading, however, this model of error rates in proofreading assumes that the tile system approaches thermodynamic equilibrium within the time scale considered. The conditions where this is the case are dependent upon the tile system, and are not necessarily well understood. In simulations, Winfree and Bekbolatov found that 2 × 2 proofreading at reasonable growth rates (*e.g.*, 10^{-2} layers of tiles per second) largely followed this equilibrium model, but at higher growth rates (*e.g.*, 1 layer of tiles per second), where kinetic traps became more important, error rates became significantly higher.[60]

Their simulations of 3 × 3 and 4 × 4 proofreading, however, did not show the same improvement as would be suggested by this thermodynamic model, and did not scale with increasing G_{se} as strongly as predicted. Thus in these cases, the simulations did not approach thermodynamic equilibrium, instead tending to fall into kinetic traps for the duration of the time scales considered. It may be that, in these cases of larger proofreading, the time scales of relevant kinetic traps increase much faster, perhaps exponentially, with increasing proofreading scale. Thus, for the infinite tile systems often considered, evidence suggests that the equilibrium probability distribution will never be approached, even as $\varepsilon \to 0$. This is a surprising and poorly understood result.

One possible explanation for this behavior in 3 × 3 and 4 × 4 proofreading is the existence of a second type of error that can arise, creating a second type of kinetic trap, called facet nucleation errors.[49] These occur when a tile attaches where no tile can correctly attach, and a new layer of tiles extends from it. In a lattice site where there is only one adjacent tile with a strength-one bond, for example, along the facet seen in Fig. 7, no tile can attach in the aTAM. However, in the kTAM, a tile attaching here by one bond could allow further attachments along this new, extended facet that would be a violation of the normal growth order of the system.

As there are no correct tiles that can attach in sites allowing facet nucleation, the error rate will depend on how common those

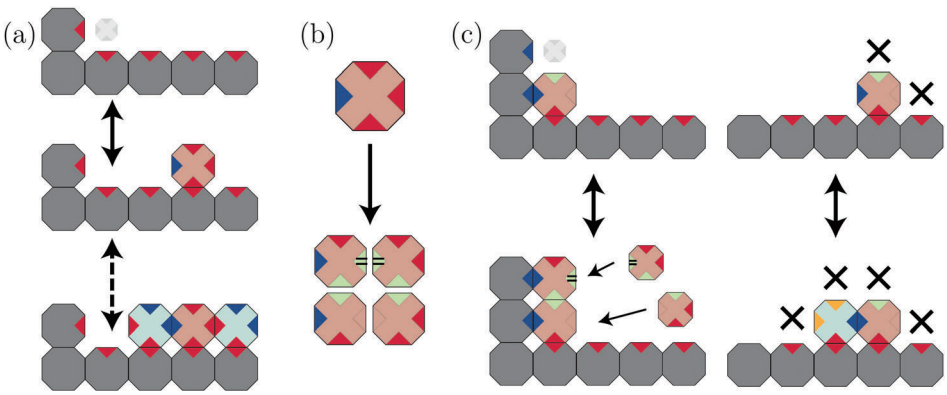

Fig. 7 Facet nucleation and the snaked proofreading construction of Chen and Goel.[49] (a) Shows a facet, along with one possible facet nucleation that can then allow for further growth. (b) Shows the construction of a snaked proofreading block (null-strength bonds are denoted by the absence of a colored triangle), while (c) shows the attachment of a correctly-placed tile (left) *vs.* a tile that attaches along a facet.

sites are, and how long they remain available. Whenever a site is available, the facet nucleation rate over time will be the rate for a tile to attach times the probability that an additional attachment occurs to lock the first tile in place before it detaches. Under $\tau = 2$ conditions, where $G_{mc} = 2G_{se} - \varepsilon$ for small ε, facet nucleation will occur only by a single weak bond, so the resulting rate is $r_{facet} \approx k_f[c]P_{locked\ in} \approx k_f[c]\frac{k_f[c]}{r_{r,1} + k_f[c]} \approx \hat{k}_f e^{-3G_{se}+2\varepsilon}$. In the worst case, where every potential growth site could allow facet nucleation, the probability of a facet nucleation error at a site would become $P_{error} \approx r_{facet}/r^* \sim e^{-G_{se}}$. This worst-case error probability is the same as the growth error probability without proofreading, limiting the effectiveness of strategies that only prevent growth errors. In practice, facet errors may not be worst-case, but are still of considerable relevance. For example, simulations of cellular automata tile sets show that facet-limited errors only come into play for tilesets with 3×3 or larger proofreading,[60] suggesting that $P_{error} < e^{-2G_{se}}$. In other words, facet errors may indeed arise at a rate r_{facet}, but all but less than $e^{-G_{se}}$ of them are cleaned up by the reversible growth near equilibrium.

To address both error types, Chen and Goel designed a "snaked" proofreading construction with differing internal bond strengths inside proofreading blocks, forcing the block to be assembled in a certain order;[49] in contrast, the "uniform" proofreading of Winfree and Bekbolatov uses all strength-one bonds. Double-strength and null-strength bonds are arranged such that an edge of a block cannot be completed favorably after a single title attaches by only one weak bond unless a further facet nucleation error takes place, reducing facet nucleation error rates by a factor of $e^{-G_{se}}$ with 2×2 blocks in one "hard" orientation, while reducing it further, in both orientations, for larger blocks.[39,49] In contrast to uniform proofreading, for which no rigorous treatment of error rates has been developed, Chen and Goel provided a proof of snaked proofreading's error rate reduction at arbitrary scale. In later research, Soloveichik, Cook, and Winfree provided an alternative proof method and applied it to a novel construction that also exhibits self-healing behavior.[61]

While the frequency of potential facet nucleation sites will be tile system dependent, for simulations of a Sierpinski system, Chen and Goel found 4×4 snaked proofreading was able to perform significantly better than 4×4 uniform proofreading to the extent that no errors occurred with snaked proofreading, whereas uniform proofreading resulted in a 25% probability of an imperfect 20×20 block assembly.[49]

Another way of considering the benefit of proofreading is considering how quickly an error-free assembly can be grown. Without proofreading, it was seen that for slightly supersaturated systems, the growth rate necessary to obtain a certain error rate as $r \sim P_{error}^2$, such that a 10-fold reduction in error rate would require a 100-fold reduction in growth rate. By making $P_{error} \sim e^{-2G_{se}}$, while only causing a two-fold decrease in growth rate of a scaled-up assembly, 2×2 proofreading will instead result in $r \sim P_{error}$. As a consequence, the time needed to assemble an $N \times N$ square now scales as only N^3, in contrast to the N^5 scaling required without proofreading. This is because a low enough error rate can now be achieved at higher concentrations. Further, Chen and Goel were able to prove that with $K \times K$ snaked proofreading, for even K, an $N \times N$ square can be assembled in time $N^{1+8/K}$. This result is not likely to be tight; it seems likely that near-equilibrium error rates are achieved by snaked proofreading, in which case $P_{error} \sim e^{-KG_{se}}$ and $r \sim e^{-2G_{se}}$ and thus $r \sim P_{error}^{2/K}$. Consequently an $N \times N$ square is assembled in time $N^{1+4/K}$, which matches the results for non-proofreading and 2×2 proofreading wherein $K = 1$ and $K = 2$ respectively. The remarkable implication is that this design, in principle, allows error rates to be decreased arbitrarily without significantly slowing down growth.

The basic mechanism of snaked proofreading has been tested experimentally, showing a reduction in facet nucleation rates along a long seed structure,[39] but has not yet been used in other tile systems. Some experiments have implemented partial 2×2 uniform proofreading, resulting in 10 to 50-fold reductions in error rates for bit-copying and binary counting systems,[29,31] and more recent works have reduced error rates even further.[15,62] In particular, Schulman *et al.*[62] demonstrated a bit-copying system with 2×1 proofreading that had an error rate of $0.017\% \pm 0.013\%$ with tiles at 25 nM. This proofreading variant is expected to have error rates similar to 2×2 proofreading, but with additional facet errors due to the zig-zag growth path in the DNA ribbons. With $K = 2$, $m = 1$ since there is only a single competing tile type, and assuming a partial matching fraction $s = 0.5$ in a generalization of eqn (9), $P_{error} \approx me^{-K(1-s)G_{se}} \approx 0.0035\%$, which is lower than but within the range of the experimental findings.

Schulman *et al.* later built similar bit-copying systems with 1, 2, 3, and 4×1 proofreading, and grew them in consistent experimental conditions:[46] a slow anneal with 50 nM tile concentration. While each increase in proofreading level decreased error rates, the error rates decreased by only around a factor of 5 from 1×1 to 2×1, and 2.5 from 2×1 to 3×1 and 3×1 to 4×1. Even with a partial matching fraction $s = 0.5$, these decreases would be expected to be around a factor of 140 for each change. However, these systems also had a zig-zag growth path, and were annealed, which introduces additional complexity and errors. In particular, the authors attributed this discrepancy to poor control over nucleation and to interference from unintended hierarchical self-assembly, both of which are discussed later in this review.

Neither uniform nor snaked proofreading have been fully implemented in a tile system, nor have the larger proofreading blocks that may be possible with SSTs been tried. Other proofreading methods are also an area with the potential for further research. Compact proofreading methods exist that avoid the scaling up of patterns at the cost of more tile types, but these have not been experimentally tested,[50,63] while proofreading in three dimensions may have the potential to be simpler and more effective.[43]

4 Nucleation

We have so far only considered individual crystals, nucleated from a chosen tile designated as a "seed tile." Growth in solution, however, takes place *via* the interaction of many tiles

and assemblies at once. Any free tile may attach to an assembly, or may attach to another free tile and form a new assembly. Assemblies may grow, or may melt away into only free tiles. Unlike the aTAM, where growth proceeds only by τ bonds and only from a seed tile, interactions between any two tiles in solution that are initially unfavorable may form an undesired seed crystal that can allow further growth. ''Spurious'' nucleation of this form poses a problem for controlling how crystals are nucleated. For algorithmic assembly, where growth may depend heavily on the seed, this spurious nucleation is particularly problematic. Rothemund *et al.*,[28] for example, saw unseeded assemblies with little resemblance to Sierpinski patterns, and Fujibayashi *et al.*[30] found numerous thin ribbon structures.

In some ways, this spurious nucleation is similar to homogeneous nucleation in single-monomer crystal nucleation,[64,65] but where homogeneous nucleation through monomer–monomer interactions is distinguished from heterogeneous nucleation at a surface interface, spurious nucleation refers only to monomer–monomer nucleation that does not start from the desired seed monomer: in one sense, all nucleation in DNA tile assembly is homogeneous. An important distinction when considering nucleation of DNA tile systems is the algorithmic nature of many of the systems: while most crystals, even when multi-component, are modelled in classical nucleation theory as having chemical potentials for each monomer being added to the nucleus from solution that don't depend on the specific nucleus,[66] with DNA tile systems, the monomers that can be added to a nucleus can completely change depending upon the exact assembly and lattice sites involved, making bulk models of limited relevance, and requiring an analysis of individual attachments and detachments.

To examine nucleation of DNA tile assemblies, we can use the rates of the kTAM as the basis for a mass-action growth model considered earlier by Schulman and Winfree.[27] Doing so will provide a basis for analyzing spurious nucleation rates in different tile systems, and motivate system designs that allow for strong nucleation control.

For simplicity, we preserve the kTAM assumption that growth takes place *via* single-tile attachments and detachments, ignoring attachment between and breakage of multi-tile assemblies. While the latter type of interactions have been seen experimentally,[30] they are unlikely in conditions where free tiles are at significantly higher concentrations than assemblies, and assemblies do not have large numbers of ends available for binding away from a growth front. Single-tile attachments and detachments can be seen as a series of chemical reactions, where an assembly A (which might be a single tile) and tile t reversibly react to form an assembly At: $A + t \rightleftharpoons At$. In a mass-action model, this results in equilibrium concentrations $[At] = [A][t](k_f/k_{r,b})$, where k_f and $k_{r,b}$ are the forward and reverse rates for the reaction, specified in eqn (1), and b is the total bond strength between the tile and the assembly.

$$[At] = [A][t]u_0^{-1}e^{-b\Delta G^\circ_{se}/RT-\alpha} \tag{10}$$

If we assume all tiles are at equal concentrations, and consider the concentrations of smaller assemblies recursively, making the same substitutions made for eqn (2), it is easily seen that the equilibrium concentration of an assembly A is

$$\begin{aligned}[A] &= [t]^N u_0^{1-N} e^{-B\Delta G^\circ_{se}/RT-(N-1)\alpha} = u_0 e^{BG_{se}-NG_{mc}+\alpha} \\ &\equiv u_0 e^{-G(A)+\alpha}\end{aligned} \tag{11}$$

where N is the total number of tiles in A, B is the total strength of all the bonds between tiles in A, and $G(A) \equiv NG_{mc} - BG_{se}$ is the total (unitless) free energy of A (with negative values being more favorable).

For simplicity, we have assumed in eqn (11) that the tile concentrations are not just equal initially, but that they remain equal and constant regardless of use in assemblies, as though free tile monomers are added and removed by some outside regulatory process, keeping the supersaturation ε constant. For such an ''open'' or ''powered'' system, the free energy, and thus concentration, of increasingly larger assemblies is unbounded. In a usual, ''closed'' system, where tile concentrations decrease as they are used in assemblies, assembly concentrations would be limited by initial tile concentrations, and the resulting equilibrium concentrations could be quite different. In practice, however, when starting near equilibrium with sufficiently low seed concentrations and high initial tile concentrations, the tile concentrations can be made to remain approximately constant over experimental time scales, as discussed later in this review.

Unseeded nucleation of crystals can be modeled using the concept of critical nuclei,[67] crystals where melting and growth are equally favorable: smaller assemblies would tend to melt, while further favorable attachments would result in a crystal that would tend to grow. These critical nuclei thus act as spurious seeds for further growth. Since it can be shown in the mass-action kTAM that assembly concentrations are bounded by equilibrium concentrations,[27] spurious nucleation rates are limited by the equilibrium concentrations of critical nuclei.

A simple example is a homogeneous tile system with a single tile type. The most favorable assemblies, with the highest number of bonds per tile, are squares, which will have differing number of tiles and bonds depending upon size, as shown in Fig. 8. For $\varepsilon \leq 0$, growing larger squares will always be unfavorable; $G(A)$ will always increase. For small $\varepsilon > 0$, small squares will be unfavorable, but at a certain point, growing larger will become favorable. Squares of this size are critical nuclei for the growth of larger squares, and their concentrations limit spurious nucleation. As ε becomes smaller, the critical nuclei of the system become larger, and have higher $G(A)$ and lower concentration, resulting in less spurious nucleation at the cost of slower growth. As long as ε is small enough, little nucleation should take place without some desired seed structure to initiate it.

Quantitatively, with the values for N and B in Fig. 8, the free energy $G(A) \equiv NG_{mc} - BG_{se}$ will be, in terms of G_{se} and $G_{mc} = 2G_{se} - \varepsilon$,

$$\frac{G(A)}{G_{se}} = \frac{N(2G_{se}-\varepsilon)}{G_{se}} - B = 2k - \frac{k^2\varepsilon}{G_{se}}. \tag{12}$$

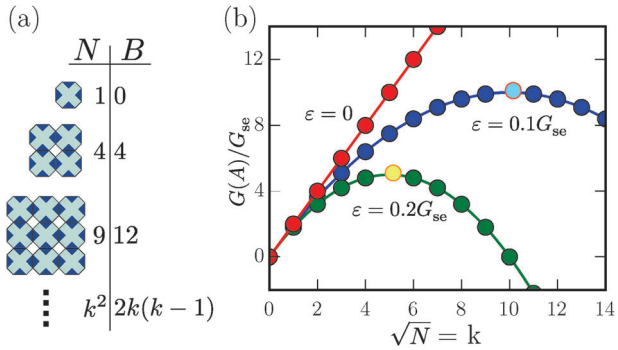

Fig. 8 Critical nuclei for uniform squares: the number of bonds and number of tiles increase at different rates as square size increases, as shown in (a). For smaller ε, squares of increasing size must form before the number of bonds *vs.* number of tiles added makes further growth favorable, as shown in (b). For $\varepsilon = 0$, growth is never favorable. The highest $G(A)$, lowest concentration squares here are critical nuclei that bound nucleation of larger crystals.

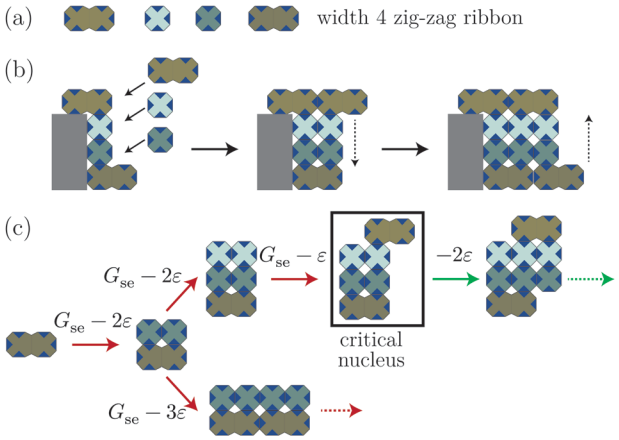

Fig. 9 The zig-zag tile system.[27] (a) Shows a width 4 tile set, which can be expanded to arbitrary width by adding tiles in the center; all top and bottom bonds are unique. Seeded growth (b) constructs a ribbon with rows that grow in alternate directions and double tiles that nucleate each subsequent row. Unseeded growth (c) has several pathways (examples are shown along with the change in $G(A)$ for the steps) but for small enough ε requires that a full row form *via* unfavorable (red) steps before growth can continue with favorable (green) steps.

where $k = \sqrt{N}$, and N is the number of tiles in the system. This free energy, seen in Fig. 8, will increase to a maximum dependent upon ε and G_{se}, creating a barrier to nucleation. As ε decreases or G_{se} increases, the height of the barrier, and the number of tile attachments required to reach it, will increase, with the maximum being reached at $N = G_{se}^2/\varepsilon^2$. At that point, the free energy will be $G(A) = G_{se}/\varepsilon$. From eqn (11), the concentration of the critical nucleus will be $[A]_{crit} = u_0 e^{-G_{se}\varepsilon^{-1}+\alpha}$. The rate of formation of larger assemblies will be limited by this critical nucleus concentration.

Thus, by decreasing supersaturation (reducing ε) or increasing bond strengths (increasing G_{se}), nucleation rates can be significantly decreased for a homogeneous tile system. For the Sierpinski system in Fig. 2, however, even just the edge tiles with double bonds form critical nuclei for long 1D polymers, for any positive ε, since adding one tile by a double bond will always be favorable. With such small, favorable critical nuclei at free tile concentrations, the system would exhibit severe spurious nucleation if implemented without modification, and such behavior has been observed experimentally.[68]

One system that allows the formation of ribbons with low spurious nucleation is the zig-zag ribbon system,[27] illustrated in Fig. 9. Ribbons grow in a zig-zag fashion, with alternating rows of tiles unique to each column growing in different directions and tiles with double bonds (or permanent "double tiles") reversing the direction of growth and nucleating the next row. Without a seed to initiate growth, continued favorable attachments by two bonds can only take place once an entire row has been formed by mostly unfavorable attachments; for small ε, this means that the formation of a critical nucleus will require $k - 1$ unfavorable attachments for a ribbon of width k. The critical nucleus shown in Fig. 9 will have $2k - 2$ tiles and a total bond strength of $3k - 5$, resulting in $G(\text{crit. nuc.}) = (k + 1)G_{se} - (2k - 2)\varepsilon$. Thus, by designing zig-zag systems of increasing width, with larger barriers to critical nucleus formation, spurious nucleation can be reduced arbitrarily.[27]

In principle, growth can be seeded relatively simply by using a structure of uniquely-addressed tiles to reliably form a large seed structure, but this can introduce a range of other experimental problems such as sticky end sequence space depletion and the difficulty of double bond construction. As an alternative, substitute structures built using other methods, such as DNA origami or assembly PCR, are often used along with adapters that allow DNA tiles to attach; this method has been used for one-dimensional seeds,[28,29] as well as rectangular and cylindrical seeds for ribbons and nanotubes.[31,37]

Schulman and Winfree examined spurious nucleation of DX tile zig-zag ribbons at several different widths.[38] Nucleation rates in the absence of a seed were measured at 50 nM tile concentration and 24 hour growth periods at fixed temperatures between 25 and 41 °C. Width 3, 4, and 6 ribbons resulted in spurious nucleation rates of $>7 \times 10^{-6}$, 1×10^{-6}, and 3×10^{-7} nM s^{-1}. Since each additional tile width will change $G(\text{crit. nuc.})$ by $G_{se} - 2\varepsilon$, for slight supersaturation conditions $\varepsilon \ll G_{se}$ and a measured $G_{se} = 7.86$, a significantly larger decrease per additional tile width, around $e^{7.86} \simeq 2600$-fold, would have been expected, though the effect would be lessened by the extent of supersaturation. Nonetheless, nucleation by origami seeds at width 4 was already found to be significantly faster than unseeded growth. Seeded nanotubes with a barrier to unseeded nucleation have recently been examined as well, with low spurious nucleation, though energetic penalties of binding to origami were noticed.[37] These forms of nucleation control have been incorporated into some experiments[31,62] that have qualitatively shown far fewer spuriously nucleated assemblies than those seen in experiments with seeded growth but a smaller barrier to unseeded nucleation.[30]

5 Tile concentrations

In the aTAM, tile concentrations are not usually considered; apart from assembly speed, they would affect final structures

only in nondeterministic systems. In the kTAM, tile concentrations are usually assumed to be equal and constant, with G_{mc} as a global, constant parameter. To have growth near equilibrium, G_{mc} needs to remain slightly smaller than $2G_{se}$, where error rates and spurious nucleation will be minimized.

Experimentally, tile concentrations are neither equal nor constant. Pipetting variation and strand synthesis can cause unintended variations in initial tile concentrations. More importantly, as crystals grow, the concentrations of free tiles in solution will be depleted, increasing G_{mc} and moving the system closer to the chemical equilibrium $G_{mc} = 2G_{se}$ where no forward growth occurs. In a system near equilibrium to begin with, small increases in G_{mc} will be enough to make growth unfavorable, and growth will never proceed very far: most tiles will remain unbound.

A simple solution to this problem is to anneal the system, gradually decreasing temperature over some period of time. While free tiles are bound, and G_{mc} decreases, the decrease in temperature will gradually raise G_{se}, and ensure that the system eventually moves back to a slightly supersaturated state favorable for growth.[22] As long as this temperature decrease is slow enough, free tiles are depleted evenly, and G_{se}s for each sticky end changes evenly with temperature, the system will remain near equilibrium, allowing for continued accurate growth.

This technique has been used in numerous experiments, and is the standard method for both non-algorithmic and algorithmic DNA tile systems.[14,28,31,45] Annealing to a sufficiently low temperature (most often room temperature) usually uses all free tiles present, potentially allowing for high yields compared to initial tile concentrations. As the system is slowly cooled, precise knowledge of equilibrium growth conditions, especially the temperature and concentration combinations required for near-equilibrium growth, is not required.

There are problems, however, with annealing. This is particularly the case for algorithmic systems; while a periodic structure, for example, will usually consist of set ratios of different tile types, algorithmic structures may use tile types at markedly differing rates, possibly at different times, and use may vary depending upon the computation or initial seed. Even if initially equal, different tile types will have their concentrations decrease at different rates. Thus attachment of some tiles may become unfavorable while attachment of others remains favorable, a situation that can result in erroneous and unpredictable growth. In another sense, in an algorithmic system that is annealed, or has tile concentrations significantly deplete, the supersaturation of each tile type will change unpredictably over time based on the algorithmic behavior of the individual assemblies as they grow.

As annealing can only act as a global control of growth rates, affecting all tile types equally, it is ineffective in handling tiles with diverging concentrations. One theoretical alternative is to use tile concentrations that are set depending upon the tile system and initial seed to result in optimal growth as tiles are depleted, but doing so in a general, experimentally-robust way remains elusive. Even without depletion, optimal choice of concentrations remains an open problem: there has been some theoretical work[24,69] and concentration choices based on simulations,[15] but concentrations used in experiments, which have often not been equal,[31] have not had any strong theoretical basis. Alternatively, if algorithmic tile systems could be designed such that all tiles in the system are consumed equally as structures grow, tile concentrations could be made to decrease at equal rates. Whether this is possible is an open and intriguing question.

Another solution, however, is simply to avoid depleting tiles significantly altogether, thus keeping the system in a nearly-constant-concentration regime approximating the standard kTAM,[62] and keeping supersaturation approximately constant, which has significantly simplifying effects even when considering simple single-component crystals.[70] With strong control of nucleation, non-seed tiles can be added at much higher concentrations than required for the desired structures, while seeds can be added at very low concentrations. Thus, the concentration of free tiles will not be significantly depleted, and G_{mc} will remain effectively homogeneous and constant. This is particularly achievable if systems have terminal assemblies, and won't continue consuming tiles indefinitely. Using this technique, structures can also be grown at constant temperature: as G_{mc} will not change significantly, there is no need to change G_{se}. Essentially, by ensuring that only very few structures are nucleated compared to the number of free tiles, growth can be made to behave in a manner largely similar to the kTAM. Experiments using this method have achieved considerably better results than earlier work using ordinary annealing. For example, error rates for bit-copying ribbons were reduced from 0.26%[31] to 0.034%[62] per 2×2 block, and similar improvements have been shown with binary counters.[15]

Structures grown at constant temperature do have the disadvantage that they must remain near their growth temperature to remain near equilibrium, especially since free tiles will remain in solution. However, many techniques and applications, especially for crystal imaging, require crystals that are stable without precise temperature control at room temperatures, significantly below typical growth temperatures. As cooling crystals would result in significant undesired growth from remaining free tiles, constant-temperature growth can be combined with "guard strands," strands that are complementary to specific tiles or portions of tiles and sticky ends.[62] When added to a system (after growth has completed) in significant excess of tile concentrations, guard strands prevent further growth by binding much more favorably than normal tile attachment, and thus deactivating free tiles. Temperatures can then be arbitrarily lowered without changing the structures that have already grown; deactivated free tiles can either be removed through purification, or be ignored, if imaging in conditions where only larger structures can bind to a surface.

The excess free tiles, however, mean that while the yield of correct structures compared to all grown structures may be high, the yield of structures compared to the concentrations of tiles used is very low. That tile concentrations are not significantly depleted, an advantage for error rates, also means that most tiles are wasted by remaining unused in solution.

Techniques or system designs that would allow most tiles in solution to be incorporated into high-accuracy structures remain an open challenge.

6 A unified model of tile assembly

6.1 Model formulation

In previous sections, we considered several related models to examine how systems designed in the $\tau = 2$ aTAM behave in other models or physical implementation, and what designs and growth methods need to be incorporated into such systems in order to allow experimental behavior that approximates ideal, abstract growth. The $\tau = 2$ aTAM and kTAM, however, are not the only models of tile assembly, and the same considerations can be made in a more general fashion by delineating more formally the relationships between more physical and more approximate or abstract models. Here, we construct a framework for understanding a number of tile assembly models as choices within a larger, unified model.

In general, all passive tile assembly can be viewed as a series of chemical reactions involving tiles and assemblies. The choice of which reactions to consider, and how the reactions should be modelled, generates many widely-studied tile assembly models that in some sense are all approximations of reality.

From a physical perspective, tile assembly taking place in solution involves tiles and assemblies that bind and unbind *via* weak sticky-end bonds. To construct a tractable model, we will assume assembly takes place slowly enough that the tiles remain well-mixed, so diffusion does not result in localized concentration gradients, and that lattices are rigid enough to prevent lattice defects from occurring. Violations of these assumptions can lead to other distinct phenomena, such as diffusion-limited aggregation[71] or lattice dislocations.[72] Both assumptions can be approximated experimentally by sufficiently slow growth and sufficiently rigid tile structures, though some tile systems, such as those with long, thin regions of assemblies, may violate them. Similarly, we will assume that only non-overlapping, lattice-preserving assembly–assembly interactions will take place, such that attachments will not cause lattice defects. This assumption is not true in general for all tile systems, but can be approximated *via* tile system design and is usually a desired property of abstract designs. With these simplifications, the following types of reactions will be possible, where capital letters denote "assemblies" of two or more bound tiles, and t_i are different tile types:

(1) Monomer–assembly interactions, where assembly B consists of a tile t_i attached to an assembly A by some number of bonds:
   (a) Attachment: $t_i + A \rightarrow B$.
   (b) Detachment: $B \rightarrow t_i + A$.

(2) Assembly–assembly (hierarchical) interactions, where assembly C consists of an assembly A bound to an assembly B in some non-overlapping fashion by some number of bonds:
   (a) Attachment: $A + B \rightarrow C$.
   (b) Detachment: $C \rightarrow A + B$.

(3) Monomer–monomer interactions, where assembly A consists only of tiles t_i and t_j bound to each other:
   (a) Attachment: $t_i + t_j \rightarrow A$.
   (b) Detachment: $A \rightarrow t_i + t_j$.

As in the kTAM, we will assume that all bimolecular attachment reactions, unless specified otherwise, have the same bimolecular rate constant, and the mass-action reaction rate depends on both concentrations. Detachment will be assumed to take place at a rate determined by the free energy change of the bonds involved: for a bimolecular forward rate constant of k_f, the detachment rate for a bond free energy change of ΔG° is $r_r = k_f e^{\Delta G^\circ/RT+\alpha}$, as in the kTAM.

It should be recognized that this assumption is an over-simplification: for example, diffusion may be a limiting factor for attachment of larger assemblies, or the kinetics of assembly–assembly interactions may depend on the specific assemblies involved rather than being the same as monomer–assembly attachment.[73,74] However, it is a reasonable approximation for monomer–assembly reactions. For assembly–assembly reactions, using the same bimolecular rate constant for all cases provides a useful upper bound on the rate of assembly, since in reality larger assemblies would react more slowly. This upper bound therefore makes it possible to formulate worst-case upper bounds on error rates: error pathways involving spurious assembly–assembly reactions will appear more likely in the model than they would be in reality. Analogously, one can obtain best-case results for hierarchical assembly speed, which is particularly relevant in considering the limitations of computation speed with hierarchical assembly.[75]

This full, "unified" tile assembly model is difficult to use or design for, owing to the diverse range of reactions. It is instead useful by virtue of encompassing many models of tile assembly, both abstract and kinetic, and illustrating relationships between them. These models can be seen as involving choices of subsets of reactions from the unified model and the choice of how to model the reactions themselves. The following choices of account for several different models of assembly:

(1) **Reversible or irreversible**: models with reversible reactions, often referred to as kinetic models, have both attachment and detachment reactions. Models with irreversible reactions contain only attachment reactions that form bonds of total strength higher than a specified threshold τ, and detachment reactions that break bonds of total bond strength less than τ. The effects of concentration and other parameters on the relative rates of attachment and detachment reactions, and thus the favorability of growth, are not considered: concentration only affects the rate of reactions that proceed in a fixed direction. Unless the threshold is changed, the set of reactions in the system does not change over time. In most cases, the detachment reactions in an irreversible model never matter, as they only involve cases where attachments will never occur in the first place, but they may be relevant in certain cases, such as systems that allow specific tiles to have destabilizing interactions with other specific tiles.[76,77]

(2) **Hierarchical or monomer-addition**: hierarchical models allow assembly–assembly interactions (2a and 2b), while monomer-addition models assume that only monomer–assembly and monomer–monomer reactions take place.
(3) **Unpowered (closed) or powered (open)**: unpowered models behave as typical systems of chemical reactions, with concentrations of monomers and assemblies changing as reactions take place. Powered models instead assume that monomer concentrations remain constant throughout assembly, whether because they are not significantly depleted or because the concentrations are controlled by some outside process.
(4) **Unseeded or seeded**: unseeded models allow any two monomers to attach to each other and form an assembly (3a). Seeded models do not allow these reactions, and thus only allow growth from preexisting assemblies. However, seeded models introduce the notion of ''seed tiles'': these are special tile types that are treated as assemblies rather than as tiles for the purposes of reaction types (1a) through (3b), and whose concentration is not preserved when using a powered model. Note that seed tiles might not be used, for example if a large assembly is part of the initial conditions (which could be considered a ''seed assembly'' but which is given no special treatment).
(5) **Discrete or continuous concentrations**: the reactions of any model can be considered using discrete stochastic chemical kinetics (as simulated by Gillespie's algorithm[78]) which is appropriate for finite volumes with finite counts, or with continuous deterministic chemical kinetics.[79]

Choices 1–4 represent choices between more complex models more closely matching physical understanding of assembly (reversible, hierarchical, unpowered, unseeded), and simplifications that assume some aspect of potential physical behavior can be ignored. Within this framework, we can define the aTAM as irreversible, monomer-addition, powered, seeded, and discrete: essentially, every ''simpler'' choice is made. This ''unified aTAM,'' in order to fit within the unified model, contains concepts such as tile concentrations, attachment rates, and time, that were not originally included in the definition of the aTAM, but the simplifications involved make these concepts irrelevant, at least for questions of reachability and growth, except when non-deterministic growth is involved. The unified kTAM differs from the unified aTAM only in being reversible, allowing tiles to detach as well as attach, while the unified 2HAM differs from the unified aTAM in being hierarchical and unseeded.

Full tile system specifications in the unified model consist of a set of tiles, a set of initial tile concentrations (or G_{mc}s), (optionally) a set of initial assemblies and their concentrations, and the set of model choices and their associated parameters. Specifically, (1) all models require the rate constant $\hat{k}_{\mathrm{f}}$ for the dimensionless rate equations; (2) irreversible models require τ; (3) reversible models require G_{se}s for each pair (possibly mismatched[24]) sticky ends; (4) seeded models require specifying which tile types are to be considered seed tiles; and (5) discrete concentration models require the reaction volume V. While the unified model concerns itself with the nondimensionalized parameters, one can easily convert to or from more realistic values, where k_{f}, α, ΔG°, and T are considered. Other experimentally relevant situations, such as temperature-annealed systems, would strictly speaking require a generalization of the unified model as presented here.

The unified tile assembly model provides a natural framework for considering how tile system behavior will change when the choice of model is changed. Any given tile set can be examined with respect to any set of other choices in the unified model framework, resulting in distinct tile systems. In many cases, tile system behavior will change completely when model choices are changed. For example, a tile system designed to assemble hierarchically will have limited growth in a monomer-addition model. Similarly, a tile system designed to grow from a single tile in an unseeded model will not grow at all in a seeded model unless that tile is added as a seed tile.

We are more interested in understanding when tile systems will continue to behave in approximately similar ways despite changes to model choices. There are a number of abstract models fitting within the unified model that have been studied for their computational properties. However, if no non-trivial tile system in one of those abstract models continues to exhibit similar behavior in a model with more physically-accurate choices, then the physical relevance of that more abstract model may be limited. Likewise, if certain properties of systems, like proofreading or nucleation control, can allow some class of system in an abstract model to assemble in an approximately similar way in a more physically relevant model, that can inform tile system design for experimental implementation. It can also suggest limitations on the situations where the more abstract model should be considered a reasonable model for tile assembly. On the other hand, designing computationally complex algorithms directly in more physically relevant systems can be challenging, and general questions of computation can become intractable, so abstract models do not become irrelevant simply because more physically accurate models can be formulated. The unified model allows an understanding of how abstract and physical models can be linked to bridge computational and physical questions.

In the case of the aTAM at $\tau = 2$, the previous sections can be interpreted as making individual changes in model choices, and providing constrained subsets of tile systems that approximately preserve growth behavior:

- When moving from an irreversible to a reversible model, the use of tile systems that implement proofreading, or that grow at extremely slow rates, allows for behavior to remain approximately the same by minimizing the growth and facet nucleation errors that arise in the reversible model.
- When moving from a reversible, seeded model to a reversible, unseeded model, tile systems that implement some form of nucleation control constrain the rates of spurious nucleation from tiles other than seed tiles acting as seeds for growth.
- When moving from a reversible, seeded, powered model to a reversible, seeded, unpowered model, behavior

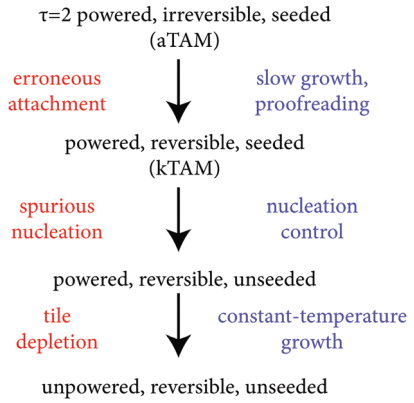

Fig. 10 Model changes from the $\tau = 2$ aTAM to more physically-accurate models. As systems are moved to each new model, problems arise that cause behavior to differ (red), which result in techniques and design constraints (blue) that allow certain classes of systems to maintain similar behavior in the more physically-accurate model.

remains the same for all tile systems in the "initial moments" of growth, before monomer concentrations have been depleted enough to significantly change attachment and detachment rates. By growing with a significant excess of monomers, and a low concentration of seeds, these "initial moments" can be extended indefinitely. However, with unbounded assembly growth, tile concentrations will eventually be depleted.

These model changes are summarized in Fig. 10. With the combination of these techniques and constraints, tile systems can be designed in the aTAM at $\tau = 2$ that should behave similarly in experiments, and have been seen to do so.[15,31,62]

The same transitions between model choices, however, may not be possible with other models, and the constraints on tile systems, designed for seeded tile systems that operate in the aTAM at $\tau = 2$, may not be applicable to other tile systems. To determine the physical viability of a design, the same analysis needs to be performed, finding ways to approximate the desired model while moving toward the choices of the more physical, general tile assembly model. We will explore two examples here: the aTAM at $\tau = 1$, which presents serious obstacles to physical implementation, and the 2HAM, which presents several challenges that may be surmountable with appropriate designs. Both the $\tau = 1$ aTAM and the 2HAM have been extensively researched within the field of theoretical computer science as models that have significantly differing computational properties.[12] Placing constraints on what systems within these models will behave similarly in more physically relevant models allows for a bridge between theoretical computer science research on the computational power of self-assembly and theoretical and experimental physical research on the physics the same self-assembly processes.

6.2 Considerations for the aTAM at $\tau = 1$

As an irreversible, monomer-interaction, powered, seeded model, the $\tau = 1$ aTAM requires consideration of the same model changes that the $\tau = 2$ aTAM requires: a design must behave similarly when the model is made reversible, hierarchical, unpowered, and unseeded. The model has reactions analogous to the $\tau = 2$ aTAM for assemblies A and B, tiles t_i, and seed tiles s_i:

(1) Monomer–assembly attachment: $t_i + A \rightarrow B$ if $\Delta b \geq 1$.
(2) Monomer–assembly detachment: $B \rightarrow t_i + A$ if $\Delta b > -1$.
(3) Seed–monomer attachment: $s_i + t_j \rightarrow A$ if $\Delta b \geq 1$.
(4) Seed–monomer detachment: $A \rightarrow t_i + t_j$ if $\Delta b > -1$.
(5) Constant monomer concentrations: $[t_i] = [t_i]_0$, $[s_i] = [s_i]_0$.

Moving to a reversible model actually imposes significantly fewer constraints for the $\tau = 1$ case than the $\tau = 2$ case. At $\tau = 1$, any single-strength bond is sufficient for correct attachment. Growth errors and facet nucleation errors of the form considered in the $\tau = 2$ case are impossible: any tile attaching by one bond in a reversible model would also attach in the irreversible model. Only when lattice deformation or sticky end sequences and spurious binding are considered can errors arise.[24]

In moving to an unseeded model, however, problems arise, even when keeping the model irreversible. As tiles can attach to each other by one bond, all correct attachments of two free tiles are favorable. There is no method, with passive tiles, to approximate seeded growth in an unseeded model: every tile that is actually used in the system will immediately serve as a spurious seed for growth, and growth will begin everywhere. The physical relevance of the seeded $\tau = 1$ aTAM is thus limited. The same problem also limits the relevance of the powered $\tau = 1$ aTAM, as an inability to control nucleation means that the excess-tile approach used to approximate powered growth in an unpowered system in the $\tau = 2$ aTAM cannot be adapted for the $\tau = 1$ case.

Even when considering systems designed in the unseeded $\tau = 1$ aTAM, problems arise when considering them in a reversible, hierarchical model. Most $\tau = 2$ tile system systems utilizing cooperative binding, where most attachments involve a tile attaching to two tiles within the lattice by two strength-1 bonds, result in assemblies where most tiles are bound to most of their adjacent tiles, as shown in Fig. 11(a). Thus, between any

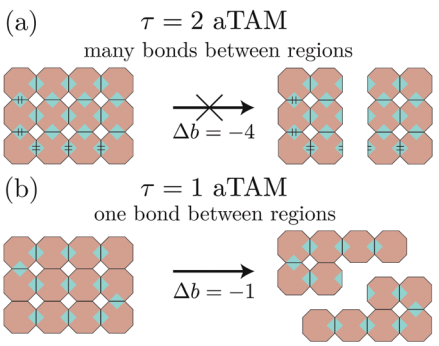

Fig. 11 Structure stability in the $\tau = 1$ and $\tau = 2$ aTAM with hierarchical detachments. In the $\tau = 2$ model, systems with cooperative assembly by strength-1 bonds naturally result in assemblies with many bonds between different regions as growth proceeds. For such an assembly to break into two smaller assemblies, many bonds would need to be broken: one example is show in (a). In the $\tau = 1$ aTAM, when each tile attaches by a single strength-1 bond, many assemblies can have only single bonds joining two regions, such that only a single bond needs to be broken to form two subassemblies: one of numerous possible examples is shown in (b).

two regions of an assembly, there are usually numerous bonds, and unintended hierarchical assembly–assembly detachments are unlikely. In $\tau = 1$ systems, this is not necessarily the case, as each tile attachment need only involve a tile attaching to a single other tile within the lattice by a strength-1 bond. If growth proceeds by a series of such attachments, regions of earlier and later growth will be attached by only one bond, as shown in Fig. 11(b). Near equilibrium in a reversible, hierarchical model, detachment into two subassemblies would therefore be only slightly slower than the monomer attachment rate: with assemblies at lower concentrations than monomers, this detachment would actually be favorable, and assemblies would constantly fall apart into smaller subassemblies while growing. Thus, for $\tau = 1$ aTAM assemblies to grow, growth would either need to take place far from equilibrium, or systems would need to be designed with "extra" bonds to ensure that assembly melting would be unfavorable.

Although the $\tau = 1$ aTAM is thus not physical for algorithmic self-assembly with simple, passive tiles, it could still be a meaningful model for physical implementations of more complex "triggered" or "signaling" tile self-assembly, where seeding and irreversibility could be enforced by the tile interactions.[54–59] It could similarly be useful for certain simple systems, like simple periodic structures, where control of nucleation location may not be important and crystals can be grown far from equilibrium. In those cases, however, the unseeded $\tau = 1$ aTAM would still be more physically relevant.

6.3 Considerations for the 2HAM

The 2HAM and its variations present an opportunity for this form of analysis on a more complex system. While the 2HAM has been defined in several different ways, it is essentially a model that treats assemblies in the same way monomers are treated in the aTAM, allowing hierarchical interaction. Monomer–monomer and monomer–assembly interactions are unchanged, but are unseeded, while assemblies, often referred to as "supertiles", are able to interact, and form larger assemblies if they can bind by a total bond strength of at least τ.[12] This makes the 2HAM within the unified model irreversible, hierarchical, powered, and unseeded. In this model, as with other variations of the 2HAM, concentrations are largely irrelevant to questions of reachability and reaction favorability: the combination of the hierarchical and irreversible choices means that every reaction between monomers or assemblies will be favorable and irreversible if the total binding strength is higher than the threshold τ, regardless of tile or assembly concentration. One physical interpretation of this would be that with equal diffusion rates, all tiles and all reachable assemblies are at constant and equal concentrations.[12] Otherwise, the motivation for the irreversible simplification—that reaction direction is predetermined and constant in a system—would not apply.

Such an interpretation, while usable for theoretical questions of abstract assembly reachability,[12] is not a reasonable model of physical growth. To have meaningful growth *via* self-assembly, at least some assembly concentrations need to change, and at least some assemblies should initially have zero concentration:

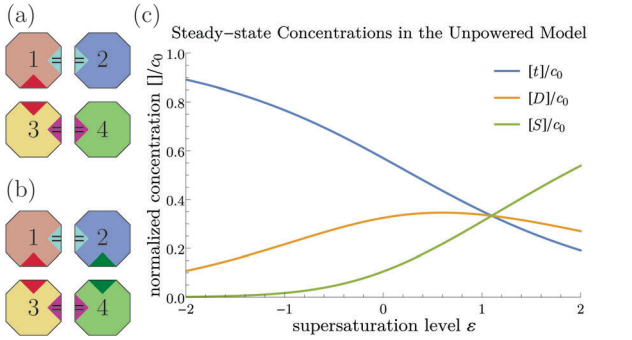

Fig. 12 Simple tile systems in the 2HAM. (a) Shows a tile system that will assemble into two dimers at $\tau = 2$, but will not assemble into a complete square, as only a bond strength of 1 can be made between the two dimers. (b) Shows the same system with one additional bond, allowing it to assemble hierarchically into a square. (c) Shows the steady-state concentrations for the system in (b) in an unpowered, reversible model, where [t] is the concentration of each monomer, [D] is the concentration of each favorable dimer, and [S] is the concentration of the square; unfavorable assemblies remain at least a factor of 10^5 below c_0 for the region in the graph. For this graph, $G_{se} = 8$.

having all assemblies present and held at equal concentrations would make the self-assembly process pointless. With our exact definition of an irreversible system, where inclusion of reactions is determined by the threshold τ without consideration of concentration, starting with zero assembly concentration, will result in similar assemblies, but with changing assembly concentrations. While moving away from the motivation for irreversibility, this choice will allow us to move to a reversible model to take into account concentration and reaction favorability. Yet moving to a reversible model, either powered or unpowered, with assembly concentrations starting at zero, presents difficulties that illustrate limitations on the physical relevance of the 2HAM.

In the unified 2HAM at $\tau = 2$, the tile system in Fig. 12(a) will assemble into two dimers, which, as they have only a total bond strength of 1 that could join them, will never assemble into a square. Moving to a reversible, hierarchical, powered, continuous model, however, the square assembly will have an initially unfavorable reaction constructing it:

(1) $t_1 + t_2 \underset{\hat{k}_f e^{-2G_{se}}}{\overset{\hat{k}_f}{\rightleftharpoons}} A_{12}$

(2) $t_3 + t_4 \underset{\hat{k}_f e^{-2G_{se}}}{\overset{\hat{k}_f}{\rightleftharpoons}} A_{34}$

(3) $t_1 + t_3 \underset{\hat{k}_f e^{-G_{se}}}{\overset{\hat{k}_f}{\rightleftharpoons}} A_{13}$

(4) $A_{12} + A_{34} \underset{\hat{k}_f e^{-G_{se}}}{\overset{\hat{k}_f}{\rightleftharpoons}} A_{1234}$

(5) Other reactions that are either unfavorable, or involve unfavorable, incomplete squares.

(6) Constant monomer concentrations $[t_i] = c_0 = u_0 e^{-G_{mc}+\alpha}$.

While not every reaction is enumerated here, with constant monomer concentrations, we can obtain steady-state concentrations of assemblies in the same manner as for the kTAM on the basis of the free energy changes involved: an assembly A will have a steady-state concentration $[A] = u_0 e^{BG_{se}-NG_{mc}+\alpha} \equiv u_0 e^{-G(A)+\alpha}$,

where B is the total bond strength and N is the total number of tiles in the assembly. If concentrations of monomers are set such that $G_{mc} = 2G_{se} - \varepsilon$, as would be done for the $\tau = 2$ kTAM, then the dimers A_{12} and A_{34} will each have steady-state concentrations $u_0 e^{2G_{se}-2G_{mc}+\alpha} = u_0 e^{-G_{mc}+\alpha} e^{\varepsilon} = c_0 e^{\varepsilon}$, higher than the monomer concentration. While monomer–monomer reactions that are initially unfavorable will remain unfavorable, the dimer–dimer interactions that are initially unfavorable will eventually become favorable as dimer concentrations increase. The steady state concentration of the square, with $B = 5$ and $N = 4$, is $[A_{1234}]_{ss} = u_0 e^{5G_{se}-4G_{mc}+\alpha} = c_0 e^{3\varepsilon-G_{se}}$. For a large enough ε, the square will assemble favorably, and eventually end up higher than the monomer concentration as well. In this basic system, $\varepsilon < \frac{G_{se}}{3}$ could be seen as an upper bound for ε; for systems with larger assemblies, however, this limit will become smaller. In general, for an assembly constructed with no hierarchical steps where each tile attaches by bond strength 2, $[A(n)]_{ss} = c_0 e^{(n-1)\varepsilon}$, where n is the number of tiles in the assembly. Two of these assemblies attaching by bond strength 1 would result in concentration $c_0 e^{(2n-1)\varepsilon-G_{se}}$. Thus, in the reversible, hierarchical, powered, continuous model with parameters set for $\tau = 2$, for any choice of ε, there may exist two sufficiently large assemblies with concentrations high enough that $\tau = 1$ growth between them becomes favorable.

This problem of high steady-state concentrations of undesired assemblies arises from the limits of the powered model as an approximation of unpowered growth. The same points could also be applied to assemblies with errors in the kTAM, but in the kTAM powered models can remain physically relevant if assembly concentrations are limited by some means, such as nucleation control. There is no need for assembly concentrations to reach or exceed tile concentrations during growth, and growth can be studied as remaining in the "initial moments" of a powered model, where assembly concentrations remain low. For hierarchical systems to assemble hierarchically, however, assemblies need to reach concentrations high enough to make assembly–assembly attachment favorable. If monomer concentrations are initially set to allow for slightly favorable monomer–monomer attachment, then assemblies must reach concentrations similar to those of tiles: no significant hierarchical assembly will take place in the initial moments. A significant concentration of monomers must be used to attain these assembly concentrations, thus limiting the relevance of a powered hierarchical model as an approximation of an unpowered model.

Two other choices exist. First, the powered model could be seen as a model of an actual powered system, where monomers are added to the system as they are used by some currently undeveloped experimental technique. As an actual powered system, however, the above problems of undesired assemblies would actually be present, and thus the 2HAM would be inappropriate for studying it.

Alternatively, unpowered hierarchical models can be considered directly. Adding an extra bond to the system in Fig. 12(a) results, in Fig. 12(b), in a working system for 2×2 squares in the 2HAM. In the reversible, hierarchical, unpowered, continuous model, the chemical reactions and rates for assemblies and monomers will remain the same as in the powered model, but rather than monomer concentrations being held constant, concentrations will be constrained by the series of equations, for each tile t_i, of

$$[t_i] + \sum_{\substack{A_j \\ t_i \in A_j}} n_{t_i,A_k} [A_j] = c_0, \tag{13}$$

where A_j is any assembly that contains tile type t_i, n_{t_i,A_k} is the number of times tile type t_i appears in assembly A_j (in this example, 1), and c_0 is the initial concentration of each monomer. There are nine possible assemblies (4 dimers, 4 trimers, and a square), of which 3 (2 dimers and a square) are energetically favorable. Using symmetry to reduce monomer concentrations to one degree of freedom, the equilibrium concentrations now become the root of the fourth-order polynomial

$$[t] + (e^{-2G_{se}} + e^{-G_{se}})[t]^2 + 3e^{-3G_{se}}[t]^3 + e^{-6G_{se}}[t]^4 = c_0. \tag{14}$$

Setting c_0 to $u_0 e^{-G_{mc}+\alpha} = u_0 e^{-2G_{se}+\varepsilon+\alpha}$, an example of equilibrium concentrations as a function of ε is shown in Fig. 12(c). Small ε results in low concentrations of the completed square: initial growth needs to be significantly forward-biased to complete assembly, so that intermediate assembly concentrations can reach a high enough level to allow favorable hierarchical attachment. For larger, more complex hierarchical systems, this problem worsens: larger assemblies would require initial concentrations to be even further from equilibrium. At some point, in contrast to the powered model, the assembly of sufficiently large assemblies at $\tau = 2$ would require initial monomer attachments take place at $\tau = 1$. For kTAM systems, seeded growth can reduce this issue of tile depletion, but for 2HAM systems, as intermediate assembly concentrations must be able to approach concentrations for favorable $\tau = 2$ growth, the same approach is not possible. From a kinetic perspective, the problem is even worse, as assembly–assembly interactions will initially involve two very low concentration terms in a bimolecular reaction rather than one low concentration assembly and one high concentration monomer, making assembly–assembly interactions in the initial moments of growth extremely slow.

This particular system in the unpowered model is also uniquely-addressed, rather than being actually algorithmic, which minimizes the depletion of each monomer. For a system that instead makes use of an individual tile t_i n times in an assembly A, it is easily seen that in an unpowered model, that assembly must have a concentration $[A] \leq [t_i]/n$ so that total tile concentrations are conserved. Two sufficiently large assemblies of this kind, in a situation of near-equilibrium $\tau = 2$ growth for monomers, would always remain at concentrations making hierarchical attachment unfavorable, as they would be constrained to significantly lower concentrations than the monomers.

These examples illustrate that 2HAM systems, as studied in theoretical computer science, cannot easily be moved to more physical models and retain similar behavior. The difficulty, in both the powered and unpowered cases, is that the concept of a globally-applicable threshold τ for binding breaks down when

applied not just to monomer–assembly interactions, but also to assembly–assembly interactions in a hierarchical model. Whether assembly–assembly binding is favorable cannot be approximated by considering only bond strengths: the concentrations of each assembly must be considered separately. This is part of the approximation in irreversible models, but in monomer-addition models, tile concentrations can be considered equal, at least initially, while in a hierarchical model, assembly concentrations cannot be considered in the same way while retaining physical relevance. Assemblies could be considered as each having distinct, time-varying τs, which would start at ∞, but could end up higher or lower than the initial τ for tiles.

Annealing could be one possibility to address concentrations in hierarchical assembly, by gradually making bonds stronger over time as temperature is decreased. However, since simple annealing would have a uniform effect on a system over time, it would not be enough to allow 2HAM systems in general to assemble in a more physical model. Assemblies of significantly different sizes, for example, could end up at very different concentrations, and interact favorably at different temperatures, or tiles could be depleted creating some smaller assemblies before hierarchical assembly at lower temperatures created new assemblies requiring those tiles for continued growth.

Whether there is a non-trivial subset of systems which assemble hierarchically in both the 2HAM and more physically-relevant models is an open question; the lack of a way to approximate a globally-applicable τ significantly limits the classes of 2HAM-designed systems that can be clearly seen to behave similarly in more physically-relevant models, and thus limits the model's physical relevance. Yet hierarchical self-assembly, in general, is a worthwhile area of research. Theoretical work on the 2HAM has shown several important differences in the computational power of hierarchical assembly in comparison to monomer-addition assembly.[12] From another perspective, supramolecular chemists have successfully designed hierarchically-assembling structures of a few components,[80,81] and fascinatingly complex hierarchical self-assembly is seen throughout biology,[19,82,83] often assembling without annealing. A common element in many of these systems is the presence of bonds of different types and strengths at different levels of hierarchical assembly. Particularly in biological systems, there is also often significant restructuring and chaperoning involved, some bonds can become active only when subassemblies are completed, and gene expression feedback can allow monomer concentrations to be controlled by the assembly process. None of these phenomena can be effectively accounted for in the 2HAM, but they could provide possible routes to a more physically-relevant abstract model of hierarchical tile self-assembly. Indeed, the fundamental importance of physical factors not accounted for in the 2HAM makes the design of physically-robust hierarchically assembling algorithmic systems a fascinating and open problem.

7 Conclusions

Through the development of these models and techniques, complex algorithmic self-assembly with DNA tiles has gradually moved toward experimental practicality. Error rates have fallen significantly with the use of proofreading tile systems and constant-temperature growth, while seeded systems with barriers to nucleation have allowed for strong control of spurious nucleation rates. Experimental work is now close to a point where some of algorithmic assembly's theoretical potential can be realized, with the construction of large tile systems that are able to perform comparatively complex computations in their growth.

From an experimental standpoint, systems of 1000 unique SST tiles have been successfully used to build uniquely-addressed structures.[25] Meanwhile, systems of around 22 DX tiles and 35 unique ends, have been used to construct algorithmic structures of thousands of tiles with decreasing error rates.[15] Whether the formidable sequence space of SST tiles can be used for complex algorithmic assembly or DX tile systems can be scaled to allow for more complex computation remains to be seen. The ability of tile systems to algorithmically construct finite terminal assemblies that stop growing and remain stable also presents a challenge that is currently being researched.[15]

At a larger scale, almost all of the advances in physical understanding of tile assembly have been made through models that consider aspects in isolation. Experimental constructions have made use of proofreading, zig-zag ribbons for nucleation control, and constant-temperature growth with varying tile concentrations all at once.[31,62] The implications of doing so, and the interactions between the methods, are unclear: in the zig-zag tile systems of Barish *et al.*[31] and later work, for example, several proofreading blocks assemble partially in one direction and are then completed in the other direction of ribbon growth, rather than assembling individually as considered in theoretical work on proofreading. Theoretical work that unifies the numerous physical aspects of tile assembly could not only provide insight into the interactions between methods, but also potentially motivate novel techniques that could solve several problems at once.

Acknowledgements

This work was partially supported by National Science Foundation awards 0832824, 1162589, and 1317694.

References

1 K. G. Libbrecht, The physics of snow crystals, *Rep. Prog. Phys.*, 2005, **68**, 855.
2 D. Vavylonis, Q. Yang and B. O'Shaughnessy, Actin polymerization kinetics, cap structure, and fluctuations, *Proc. Natl. Acad. Sci. U. S. A.*, 2005, **102**, 8543–8548.
3 J. S. Lindsey, Self-assembly in synthetic routes to molecular devices. Biological principles and chemical perspectives: a review, *New J. Chem.*, 1991, **15**, 153–180.
4 B. Grünbaum and G. C. Shephard, *Tilings and Patterns*, W. H. Freeman & Co., New York, NY, USA, 1986.

5 C. Janot, *Quasicrystals: a primer*, Oxford Univ. Press, Oxford, UK, 1994.
6 M. Senechal, *Quasicrystals and geometry*, Cambridge Univ. Press, Cambridge, UK, 1996.
7 H. Wang, An unsolvable problem on dominoes, *Harvard Computation Laboratory, Technical Report BL30 (II-15)*, 1962.
8 R. Berger, The undecidability of the domino problem, *Memoirs of the AMS*, 1966, vol. 66, pp. 1–72.
9 W. Hanf, Nonrecursive tilings of the plane I, *J. Symbolic Logic*, 1974, **39**, 283–285.
10 D. Myers, Nonrecursive tilings of the plane II, *J. Symbolic Logic*, 1974, **39**, 286–294.
11 D. Doty, Theory of algorithmic self-assembly, *Commun. ACM*, 2012, **55**, 78–88.
12 M. J. Patitz, An introduction to tile-based self-assembly and a survey of recent results, *Nat. Comput.*, 2004, **13**, 195–224.
13 C. G. Evans, R. F. Hariadi and E. Winfree, Direct atomic force microscopy observation of DNA tile crystal growth at the single-molecule level, *J. Am. Chem. Soc.*, 2012, **134**, 10485–10492.
14 B. Wei, M. Dai and P. Yin, Complex shapes self-assembled from single stranded DNA tiles, *Nature*, 2012, **485**, 623–626.
15 C. G. Evans, PhD thesis, *Crystals that count! Physical principles and experimental investigations of DNA tile self-assembly*, California Institute of Technology, 2014.
16 A. L. Mackay, Generalised crystallography, *Izvj. Jugosl. Centr. Krist.*, 1975, **10**, 15–36.
17 J.-M. Lehn, Toward self-organization and complex matter, *Science*, 2002, **295**, 2400–2403.
18 E. Nogales and H.-W. Wang, Structural mechanisms underlying nucleotide-dependent self-assembly of tubulin and its relatives, *Curr. Opin. Struct. Biol.*, 2006, **16**, 221–229.
19 P. G. Leiman, S. Kanamaru, V. V. Mesyanzhinov, F. Arisaka and M. G. Rossmann, Structure and morphogenesis of bacteriophage T4, *CMLS, Cell. Mol. Life Sci.*, 2003, **60**, 2356–2370.
20 C. Lin, Y. Liu, S. Rinker and H. Yan, DNA tile based self-assembly: Building complex nanoarchitectures, *ChemPhysChem*, 2006, **7**, 1641–1647.
21 T. J. Fu and N. C. Seeman, DNA double-crossover molecules, *Biochemistry*, 1993, **32**, 3211–3220.
22 E. Winfree, F. Liu, L. A. Wenzler and N. C. Seeman, Design and self assembly of two-dimensional DNA crystals, *Nature*, 1998, **394**, 539–544.
23 P. Yin, R. F. Hariadi, S. Sahu, H. M. T. Choi, S. H. Park, T. H. LaBean and J. H. Reif, Programming DNA tube circumferences, *Science*, 2008, **321**, 824–826.
24 C. G. Evans and E. Winfree, DNA sticky end design and assignment for robust algorithmic self-assembly, in *DNA Computing and Molecular Programming (DNA)*, ed. D. Soloveichik and B. Yurke, Lecture Notes in Computer Science (LNCS), Springer, Berlin, 2013, vol. 8141, pp. 61–75.
25 Y. Ke, L. L. Ong, W. M. Shih and P. Yin, Three-dimensional structures self-assembled from DNA bricks, *Science*, 2012, **338**, 1177–1183.
26 E. Winfree, *Simulations of computing by self-assembly*, tech. rep. CaltechCSTR:1998.22, Pasadena, CA, 1998.
27 R. Schulman and E. Winfree, Programmable control of nucleation for algorithmic self-assembly, *SIAM J. Comput.*, 2010, **39**, 1581–1616.
28 P. W. K. Rothemund, N. Papadakis and E. Winfree, Algorithmic self-assembly of DNA sierpinski triangles, *PLoS Biol.*, 2004, **2**, e424.
29 R. D. Barish, P. W. K. Rothemund and E. Winfree, Two computational primitives for algorithmic self-assembly: copying and counting, *Nano Lett.*, 2005, **5**, 2586–2592.
30 K. Fujibayashi, R. Hariadi, S. H. Park, E. Winfree and S. Murata, Toward reliable algorithmic self-assembly of DNA tiles: a fixed-width cellular automaton pattern, *Nano Lett.*, 2008, **8**, 1791–1797.
31 R. D. Barish, R. Schulman, P. W. K. Rothemund and E. Winfree, An information-bearing seed for nucleating algorithmic self-assembly, *Proc. Natl. Acad. Sci. U. S. A.*, 2009, **106**, 6054–6059.
32 L. Adleman, Q. Cheng, A. Goel, M.-D. Huang, D. Kempe, P. M. de Espanés and P. W. K. Rothemund, Combinatorial optimization problems in self-assembly, *Symposium on the Theory of Computing (STOC)*, ACM, Montreal, Quebec, Canada, 2002, pp. 23–32.
33 P. W. K. Rothemund and E. Winfree, The program-size complexity of self-assembled squares, *Symposium on the Theory of Computing (STOC)*, ACM, Portland, Oregon, USA, 2000, pp. 459–468.
34 L. Adleman, Q. Cheng, A. Goel and M.-D. Huang, Running time and program size for self-assembled squares, *Symposium on the Theory of Computing (STOC)*, ACM, Hersonissos, Greece, 2001, pp. 740–748.
35 D. Soloveichik and E. Winfree, Complexity of self-assembled shapes, *SIAM J. Comput.*, 2007, **36**, 1544–1569.
36 D. Doty, M. J. Patitz and S. M. Summers, Limitations of self-assembly at temperature 1, *Theor. Comput. Sci.*, 2011, **412**, 145–158.
37 A. M. Mohammed and R. Schulman, Directing self-assembly of DNA nanotubes using programmable seeds, *Nano Lett.*, 2013, **13**, 4006–4013.
38 R. Schulman and E. Winfree, Synthesis of crystals with a programmable kinetic barrier to nucleation, *Proc. Natl. Acad. Sci. U. S. A.*, 2007, **104**, 15236–15241.
39 H.-L. Chen, R. Schulman, A. Goel and E. Winfree, Reducing facet nucleation during algorithmic self-assembly, *Nano Lett.*, 2007, **7**, 2913–2919.
40 S. Woo and P. W. Rothemund, Self-assembly of two-dimensional DNA origami lattices using cation-controlled surface diffusion, *Nat. Commun.*, 2014, **5**, 4889.
41 S. Hamada and S. Murata, Substrate-assisted assembly of interconnected single-duplex DNA nanostructures, *Angew. Chem., Int. Ed.*, 2009, **48**, 6820–6823.
42 S. Hamada and S. Murata, Theoretical model of substrate-assisted self-assembly of DNA nanostructures, *RSC Adv.*, 2012, **2**, 7406–7412.
43 H.-L. Chen, A. Goel and C. Luhrs, Dimension augmentation and combinatorial criteria for efficient error-resistant DNA self-assembly, *Symposium on Discrete Algorithms (SODA)*, SIAM, San Francisco, California, 2008, pp. 409–418.

44 J. C. Mitchell, J. R. Harris, J. Malo, J. Bath and A. J. Turberfield, Self-assembly of chiral DNA nanotubes, *J. Am. Chem. Soc.*, 2004, **126**, 16342–16343.
45 K. Fujibayashi and S. Murata, Precise simulation model for DNA tile self-assembly, *IEEE Trans. Nanotechnol.*, 2009, **8**, 361–368.
46 R. Schulman, C. Wright and E. Winfree, Increasing redundancy exponentially reduces error rates during algorithmic self-assembly, *ACS Nano*, 2015, **9**, 5760–5771.
47 M. J. Patitz, Simulation of self-assembly in the abstract tile assembly model with ISU TAS, 2011, arXiv:1101.5151 [cs.MS].
48 *The Xgrow Simulator*, http://dna.caltech.edu/Xgrow/.
49 H.-L. Chen and A. Goel, Error free self-assembly using error prone tiles, *DNA Computing (DNA)*, Lecture Notes in Computer Science (LNCS), Springer, 2005, vol. 3384, pp. 702–707.
50 D. Soloveichik and E. Winfree, Complexity of compact proofreading for self-assembled patterns, *DNA Computing (DNA)*, Lecture Notes in Computer Science (LNCS), 2006, vol. 3892, pp. 125–135.
51 J. SantaLucia and D. Hicks, The thermodynamics of DNA structural motifs, *Annu. Rev. Biophys. Biomol. Struct.*, 2004, **33**, 415–440.
52 P. Gàcs, Reliable cellular automata with self-organization, *J. Stat. Phys.*, 2001, **103**, 45–267.
53 D. Lubrich, S. J. Green and A. J. Turberfield, Kinetically controlled self-assembly of DNA oligomers, *J. Am. Chem. Soc.*, 2009, **131**, 2422–2423.
54 R. M. Dirks and N. A. Pierce, Triggered amplification by hybridization chain reaction, *Proc. Natl. Acad. Sci. U. S. A.*, 2004, **101**, 15275–15278.
55 J. E. Padilla, M. J. Patitz, R. Pena, R. T. Schweller, N. C. Seeman, R. Sheline, S. M. Summers and X. Zhong, Asynchronous signal passing for tile self-assembly: fuel efficient computation and efficient assembly of shapes, *Unconventional Computation and Natural Computation (UCNC)*, Lecture Notes in Computer Science (LNCS), Springer, 2013, vol. 7956, pp. 174–185.
56 N. Jonoska and D. Karpenko, Active tile self-assembly, part 1: universality at temperature 1, *Int. J. Found. Comput. Sci.*, 2014, **25**, 141–163.
57 K. Fujibayashi, D. Zhang, E. Winfree and S. Murata, Error suppression mechanisms for DNA tile self-assembly and their simulation, English, *Nat. Comput.*, 2009, **8**, 589–612.
58 U. Majumder, T. H. LaBean and J. H. Reif, Activatable tiles: compact, robust programmable assembly and other applications, *DNA Computing (DNA)*, Lecture Notes in Computer Science (LNCS), Springer, 2008, vol. 4848, pp. 15–25.
59 J. E. Padilla, R. Sha, M. Kristiansen, J. Chen, N. Jonoska and N. C. Seeman, A signal-passing DNA-strand-exchange mechanism for active self-assembly of DNA nanostructures, *Angew. Chem., Int. Ed.*, 2015, **54**, 5939–5942.
60 E. Winfree and R. Bekbolatov, Proofreading tile sets: error correction for algorithmic self-assembly, *DNA Computing (DNA)*, Lecture Notes in Computer Science (LNCS), Springer, 2004, vol. 2943, pp. 126–144.
61 D. Soloveichik, M. Cook and E. Winfree, Combining self-healing and proofreading in self-assembly, *Nat. Comput.*, 2008, **7**, 203–218.
62 R. Schulman, B. Yurke and E. Winfree, Robust self-replication of combinatorial information *via* crystal growth and scission, *Proc. Natl. Acad. Sci. U. S. A.*, 2012, **109**, 6405–6410.
63 J. H. Reif, S. Sahu and P. Yin, Compact error-resilient computational DNA tiling assemblies, *DNA Computing (DNA)*, Lecture Notes in Computer Science (LNCS), Springer, 2005, vol. 3384, pp. 293–307.
64 D. W. Oxtoby, Homogeneous nucleation: theory and experiment, *J. Phys.: Condens. Matter*, 1992, **4**, 7627–7650.
65 D. Kashchiev and G. M. van Rosmalen, Review: Nucleation in solutions revisited, *Cryst. Res. Technol.*, 2003, **38**, 555–574.
66 V. I. Kalikmanov, Multi-component nucleation, *Nucleation Theory*, Springer, Netherlands, Dordrecht, 2013, pp. 239–251.
67 *Nucleation*, ed. A. Zettlemoyer, Marcel Dekker, New York, 1969.
68 R. Schulman, S. Lee, N. Papadakis and E. Winfree, One dimensional boundaries for DNA tile self-assembly, in *DNA Computing (DNA)*, Madison, WI, USA, June 1–3, 2003, ed. J. Chen and J. Reif, Lecture Notes in Computer Science (LNCS), Springer Berlin Heidelberg, 2004, vol. 2943, pp. 108–125.
69 H.-L. Chen and M.-Y. Kao, Optimizing tile concentrations to minimize errors and time for DNA tile self-assembly systems, *DNA Computing and Molecular Programming (DNA)*, Lecture Notes in Computer Science (LNCS), Springer, 2011, vol. 6518, pp. 13–24.
70 R. P. Sear, Quantitative studies of crystal nucleation at constant supersaturation: experimental data and models, *CrystEngComm*, 2014, **16**, 6506–6522.
71 T. Witten and L. Sander, Diffusion-limited aggregation, a kinetic critical phenomenon, *Phys. Rev. Lett.*, 1981, **47**, 1400–1403.
72 J. P. Hirth and J. Lothe, *Theory of Dislocations*, Wiley, New York, 1982.
73 R. F. Hariadi, E. Winfree and B. Yurke, Determining hydrodynamic forces in bursting bubbles using DNA nanotube mechanics, *Proc. Natl. Acad. Sci. U. S. A.*, 2015, **112**, E6086–E6095, [eqn (13) of main text].
74 J. Zenk, C. Tuntivate and R. Schulman, Kinetics and thermodynamics of Watson-Crick base pairing driven DNA origami dimerization, *J. Am. Chem. Soc.*, 2016, **138**, 3346–3354.
75 H.-L. Chen and D. Doty, Parallelism and time in hierarchical self-assembly, Symposium on Discrete Algorithms (SODA), SIAM, Kyoto, Japan, 2012, pp. 1163–1182.
76 M. Patitz, R. Schweller and S. Summers, Exact shapes and turing universality at temperature 1 with a single negative glue, in *DNA Computing and Molecular Programming (DNA)*, ed. L. Cardelli and W. Shih, Lecture Notes in Computer Science (LNCS), Springer, Berlin, Heidelberg, 2011, vol. 6937, pp. 175–189.
77 J. Reif, S. Sahu and P. Yin, Complexity of graph self-assembly in accretive systems and self-destructible systems, in *DNA Computing (DNA)*, ed. A. Carbone and N. A. Pierce, Lecture Notes in Computer Science (LNCS), Springer, Berlin, Heidelberg, 2006, vol. 3892, pp. 257–274.

78 D. T. Gillespie, Exact stochastic simulation of coupled chemical reactions, *J. Phys. Chem.*, 1977, **81**, 2340–2361.
79 F. Horn and R. Jackson, General mass action kinetics, *Arch. Ration. Mech. Anal.*, 1972, **47**, 81–116.
80 J. A. A. W. Elemans, A. E. Rowan and R. J. M. Nolte, Mastering molecular matter. Supramolecular architectures by hierarchical self-assembly, *J. Mater. Chem.*, 2003, **13**, 2661–2670.
81 O. Ikkala and G. ten Brinke, Hierarchical self-assembly in polymeric complexes: towards functional materials, *Chem. Commun.*, 2004, 2131–2137.
82 F. F. V. Chevance and K. T. Hughes, Coordinating assembly of a bacterial macromolecular machine, *Nat. Rev. Microbiol.*, 2008, **6**, 455–465.
83 P. Aldridge and K. T. Hughes, Regulation of flagellar assembly, *Curr. Opin. Microbiol.*, 2002, **5**, 160–165.

