Crystals that Count!

Physical Principles and Experimental Investigations of DNA Tile Self-Assembly

Thesis by
Constantine Glen Evans

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy

California Institute of Technology
Pasadena, California
2014
(Defended May 27, 2014)

for Glen, who should have been able to call me an idiot while reading this thesis

αἰαῖ ταὶ μαλάχαι μέν, ἐπὰν κατὰ κᾶπον ὄλωνται,
ἠδὲ τὰ χλωρὰ σέλινα τό ῾τ εὐθαλὲς οὖλον ἄνηθον,
ὕστερον αὖ ζώοντι καὶ εἰς ἔτος ἄλλο φύοντι:
ἄμμες δ᾽ οἱ μεγάλοι καὶ καρτεροί, οἱ σοφοὶ ἄνδρες,
ὁππότε πρᾶτα θάνωμες, ἀνάκοοι ἐν χθονὶ κοίλᾳ
εὕδομες εὖ μάλα μακρὸν ἀτέρμονα νήγρετον ὕπνον.
καὶ σὺ μὲν ὦν σιγᾷ πεπυκασμένος ἔσσεαι ἐν γᾷ,
ταῖς Νύμφαισι δ᾽ ἔδοξεν ἀεὶ τὸν βάτραχον ᾄδειν.
ταῖς δ᾽ ἐγὼ οὐ φθονέοιμι: τὸ γὰρ μέλος οὐ καλὸν ᾄδει.
Lament for Bion 99-107

Acknowledgments

When I arrived at Caltech, it was as a young theorist interested in string theory. As my fascination with that field quickly faded, I found myself lost and directionless. It was by chance that I found my advisor Erik Winfree's lab while searching for something that was both interesting and not string theory.

In the years that followed, Erik continually supported me through numerous personal and academic ordeals, at the same time patiently assisting me in learning to do actual research in a field I was hitherto entirely unfamiliar with. It has only been through his patience and guidance that I have survived to finish this.

The lab, too, has always been a great support, from Rebecca Schulman and Rizal Hariadi, who helped to start off my research and teach me entirely new topics, to Paul Rothemund, Nadine Dabby, Joseph Schaeffer, Damien Woods, David Doty, Sungwook Woo, Niranjan Srinivas, Lulu Qian, Ashwin Gopinath, Lucinda Acosta, and the many others who have made research and life here so interesting.

Outside of research, I have deeply appreciated my friends at Caltech who have made my time as a graduate student here enjoyable, particularly Catherine Beni, Camden Jansen, Daniel Thai, Daniel Leighton, Christian Griset, and too many others to name; our antics and nonsense were as much sources of delight to me as they were sources of annoyance to others. I thank Megan Wiseley for helping me to become a somewhat stable student here, and Nicole De Silva for motivating me to complete this thesis. I must acknowledge, too, the frequent advice of my sister Marina Karastamatis, which has always been sharp and accurate and which I have often regretted ignoring.

For countless reasons, I am eternally indebted to my mother, Joyce Evans, who has devoted unimaginable amounts of time and effort to my many endeavors and has been at once a caring mother and stalwart research and business partner. She has, astoundingly, put up with all of my antics and bumblings through the years, whether at UCSD, Egea, or especially Caltech, despite always having so many other worries. If, when she was a graduate student, I taught her "the meaning of patience" [18], then I am sure that as a graduate student myself I have redefined it for her.

Lastly, this thesis is in memory of my uncle, Glen Evans. His illustrious research career as part of the Human Genome Project and a pioneer of synthetic biology, his brilliance, his eccentricity and perhaps also his irreverence were always inspirations to me. It was from him, growing up, that I first learned how to pipette, and my (still rather limited) knowledge of molecular biology and lab technique largely stems from my working with him at

Egea Biosciences and elsewhere. Watching him struggle with his illness, still supporting my pursuits and our many shared interests, from the time I first started at Caltech until his death in 2010, was devastating. Were he still alive today, I am sure that this thesis would not have escaped his sharp yet insightful criticism; I can only hope that in my research I can someday accomplish some fraction of what he did.

Abstract

Algorithmic DNA tiles systems are fascinating. From a theoretical perspective, they can result in simple systems that assemble themselves into beautiful, complex structures through fundamental interactions and logical rules. As an experimental technique, they provide a promising method for programmably assembling complex, precise crystals that can grow to considerable size while retaining nanoscale resolution. In the journey from theoretical abstractions to experimental demonstrations, however, lie numerous challenges and complications.

In this thesis, to examine these challenges, we consider the physical principles behind DNA tile self-assembly. We survey recent progress in experimental algorithmic self-assembly, and explain the simple physical models behind this progress. Using direct observation of individual tile attachments and detachments with an atomic force microscope, we test some of the fundamental assumptions of the widely-used kinetic Tile Assembly Model, obtaining results that fit the model to within error. We then depart from the simplest form of that model, examining the effects of DNA sticky end sequence energetics on tile system behavior. We develop theoretical models, sequence assignment algorithms, and a software package, StickyDesign, for sticky end sequence design.

As a demonstration of a specific tile system, we design a binary counting ribbon that can accurately count from a programmable starting value and stop growing after overflowing, resulting in a single system that can construct ribbons of precise and programmable length. In the process of designing the system, we explain numerous considerations that provide insight into more general tile system design, particularly with regards to tile concentrations, facet nucleation, the construction of finite assemblies, and design beyond the abstract Tile Assembly Model.

Finally, we present our crystals that count: experimental results with our binary counting system that represent a significant improvement in the accuracy of experimental algorithmic self-assembly, including crystals that count perfectly with 5 bits from 0 to 31. We show some preliminary experimental results on the construction of our capping system to stop growth after counters overflow, and offer some speculation on potential future directions of the field.

Contents

Thesis Outline

Algorithmic tile assembly systems are fascinating. In theory, from very simple sets of tiles, they can grow into beautiful, complex structures as a consequence of fundamental interactions. The work behind this thesis is motivated by relatively simple questions: with DNA tiles, can we actually build such structures? And in doing so, what complexities lie between theoretical abstractions and physical realities?

In the last chapter of this thesis, Chapter 5, some of our earliest experimental results are included: one particularly beautiful binary counter crystal we found in a sample in 2010. Our journey to both reproduce those results reliably, and better understand the physical principles of tile assembly, led us on a meandering path in various directions before finally returning us to our original goals. Rather than follow our meanderings, this thesis instead explains our research as a path from the abstract and general to the specific and concrete, building up a foundation for the construction of the experimental systems it eventually presents.

In Chapter 1, we begin by surveying the physical principles of DNA self-assembly as they are currently understood in the field. Algorithmic DNA tile systems have been studied for some time, but despite extensive theoretical research, numerous challenges have arisen in experimental implementations, and these have only started to be overcome in recent experiments. We explain various simple physical models of tile assembly, and how they have offered insight into experimental results and motivated improvements in tile system design and growth. We also look forward into areas with the potential for future research. This chapter is adapted from a review paper with Erik Winfree currently under revision for publication in Accounts of Chemical Research.

That simple physical models can be intuitively developed to explain a process, however, does not mean that they are correct, even if they qualitatively match experimental results. In Chapter 2, we quantitatively examine several assumptions of the widely used kinetic Tile Assembly Model (kTAM) by directly observing tiles with an atomic force microscope as they attach to and detach from crystals on a mica surface near equilibrium. With the resulting "movies," we were able to gather statistics of tile attachment and detachment rates, which fit the kTAM's rate assumptions within experimental error; at the same time, our experiments demonstrated the viability of AFM movies as a technique for directly examining DNA tile systems during growth. This work was published in the Journal of the American Chemical Society in 2012. The original experimental design and experiments were envisioned and run by Rizal Hariadi and Erik Winfree. I ran further experiments, and developed the image interpretation process and Bayesian methods used to analyze the data.

Now more confident in our models, and with a better understanding of design and experimental methods, we can start considering details that haven't been analyzed yet. One of these areas is in the energetics of tile detachment. For simplicity, the kTAM assumes all bonds have equal strength. In reality, however, DNA hybridization is highly sequence dependent. In Chapter 3, we present our theoretical and simulation-based research into how variations in sequence energetics affect error rates in tile assembly. We show that both uniformity of sticky end energies and cross-talk between different sticky ends are important considerations in tile assembly, and then develop algorithms and a software package to design sticky end sequences. We also examine how concentration variations might allow compensation for lack of sequence uniformity, and how systems can be analyzed to determine sensitivity of different sticky end pairs to cross-talk. This work was published in the proceedings of the 19th DNA Computing and Molecular Programming conference in 2013. Arising initially out of simple thoughts about more accurately simulating tile systems and designing new sticky ends for the binary counter, it unexpectedly became an increasingly complex topic that was studied in numerous conversations between me and Erik Winfree. I developed much of the theoretical basis of the results presented in the chapter, ran simulations, and constructed the algorithms and software.

At last, we can return to the binary counter system. In Chapter 4, we examine the binary counter as an example of a specific tile system design. With theoretical models, simulations, and qualitative experimental observations, we consider numerous aspects of tile system design through this single example, particularly with regards to the effects of tile concentrations and facet nucleation. We also consider the surprisingly challenging design of tile systems that reliably build finite assemblies in the kTAM, develop a "capping" mechanism to do so with the binary counter, and take a brief sojourn into the "clever" design of a system that functions only in the kTAM.

In Chapter 5, we conclude with experimental results in algorithmic self-assembly. We present the original crystal that beguiled us with its accuracy (and uniqueness), along with numerous other crystals that have now been grown with similar error rates. We also examine preliminary results with a capping mechanism to produce finite ribbons of programmable length with the binary counter system. The counting crystals in this chapter represent a significant improvement over earlier works on algorithmic self-assembly, primarily achieved through improvements in how tile systems are grown and the better understanding of the binary counter system detailed in Chapter 4. The binary counter system in the chapter was derived from the work of Rebecca Schulman on zig-zag copying ribbons, who also developed many of the experimental techniques we would later use [64]. I ran experiments and simulations, developed the theoretical considerations in Chapter 4, modified the improved zig-zag copy ribbon design of Schulman et al to incorporate counting, and designed the capping system.

Chapter 1

Physical Principles for DNA Tile Self-Assembly

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. These 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 assemblies using the same tiles, based on inputs to the assembly process. This chapter considers challenges that have arisen in the experimental implementation of tile assembly using DNA tiles, and how those challenges are being overcome in more recent experiments, some of which have involved dozens of tile types assembling into complex crystals of several thousand tiles. We explain examples of simple physical models of DNA tile assembly phenomena that have offered insight into experimental results and motivated improvements in tile system design and growth, and identify numerous areas with the potential for future research that could assist with experimental progress.

1.1 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 [37, 75, 40]. The simplicity and power of such self-assembly is striking: unlike many synthetic structures that are built individually by some outside process, 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 already made seminal contributions to the understanding of periodic crystalline structures, such as the classification of all two-dimensional and three-dimensional symmetry groups [25]. Tiling theory also uncovered more complex, and often beautiful, possibilities, such as the existence of aperiodic tilings related to quasicrystals [31, 65]. Perhaps more surprising, a direct correspondence between

Manuscript submitted and under revision for publication as Constantine G Evans and Erik Winfree, "Physical Principles for DNA Tile Self-Assembly."

tilings and Turing machines [76, 5] 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) [27, 44]. 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 also must 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 [14, 53]; 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, 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 periodically to create an object of unbounded size, as well as analogous examples from biology such as microtubules that self-assemble from $\alpha-$ and $\beta-$tubulin. 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 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 icosohedral head, whiskers, tail, baseplate, and tail fibers) 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 [39]. DNA tiles, which can have a variety of structures, attach to each other by short, single-stranded "sticky end" regions of typically 5 to 10 nucleotides. The complementary binding of DNA allows binding rules to be programmed by setting the sticky end sequences of each tile type, allowing for flexible, 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 nt (like those found on double-crossover (DX) tiles [19, 81]) can allow 30 to 60 unique sticky end sequences, and longer sticky ends (like those found on single-stranded tiles (SSTs) [82]) can allow for considerably more,

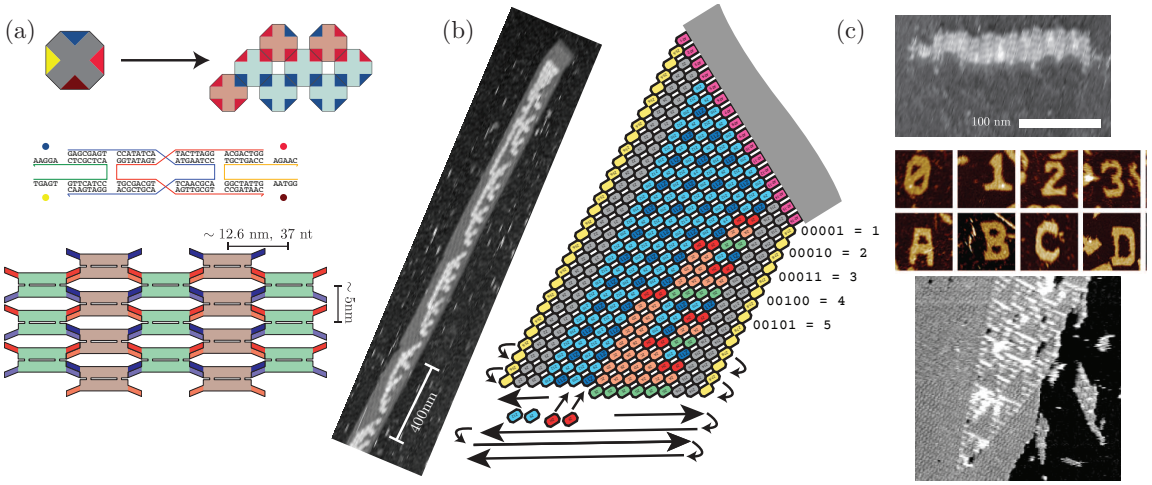

Figure 1.1: Examples of DNA tile self-assembly. (a) shows an abstract representation of a tile and the structure of one double-crossover (DX) tile [63], along with lattices representations [16]. (b) shows a schematic and atomic force microscopy (AFM) image of a zig-zag binary counting system (Chapter 5). (c) shows other examples of assembly: from top to bottom, a periodic crystal [16], 150nm images of crystals formed from uniquely-addressed tiles [77], and an algorithmic Sierpinski pattern initiated by a single error in a sea of zeros [22].

according to plausible sequence design criteria [17]. DNA tiles have been used in a wide range of large periodic structures [39] and structures with uniquely-addressed tiles that have ranged up to a thousand unique tiles in three dimensions [33, 77].

Experimental implementation of algorithmic self-assembly, however, has posed several challenges. Of primary importance, growth errors can occur when tiles attach in locations where they have imperfect matches, and spurious nucleation can occur, resulting in undesired structures [79, 61, 57, 3, 22]. Nonetheless, considerable progress has been made in overcoming these challenges, and significantly more complex structures have been successfully built [22, 4]. This progress has largely stemmed not from improvements in experimental technique, but from a better understanding of the physics behind DNA tile assembly, and simple models of growth that have motivated fundamental changes to tile system design and protocols.

In this review, we consider several examples of major improvements in experimental algorithmic self-assembly, and show how these were derived from an understanding of the physical principles behind the self-assembly process.

1.2 Basic Models

Two complementary models of tile assembly are widely studied, both from Winfree [79]. 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.

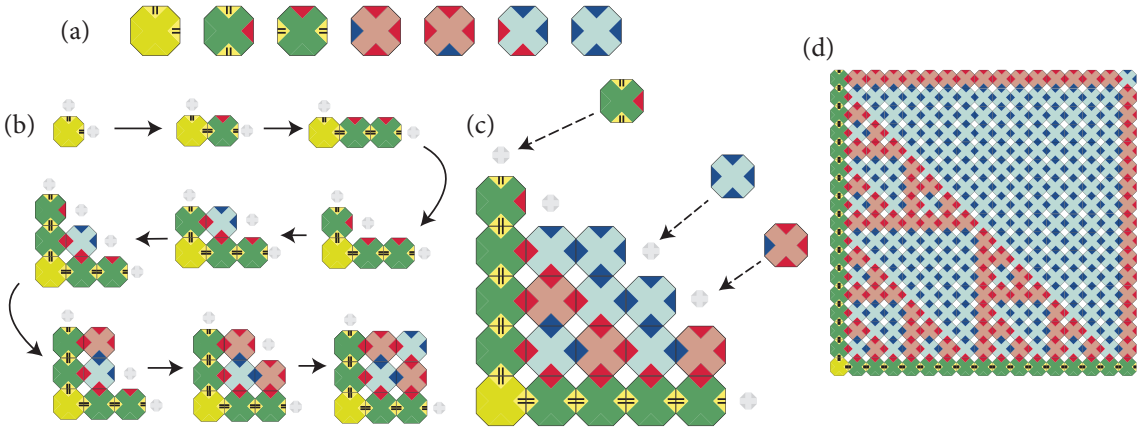

Figure 1.2: A $\tau = 2$ XOR/Sierpinski pattern tile system in the abstract Tile Assembly Model (aTAM). (a) shows the 7 tiles in the set: a yellow seed tile with double-strength ends (denoted by two lines), two boundary tiles, and four rule tiles. (b) shows a possible series of attachments, with gray tiles indicating potential attachment sites. (c) shows the possible tiles that can attach to a small assembly, while (d) shows a larger assembly.

In the aTAM, tile systems consist of tiles with four sticky ends of integer *strengths*: usually 1 ("single-strength" or "weak") or 2 ("double-strength" or "strong"). Growth begins from a designated *seed tile*; tiles then attach asynchronously, one at a time from all possible attachments, if their sticky ends match ends on adjacent tiles in the assembly with strengths that sum to at least the system "temperature" or "threshold" τ. Once attached, tiles never detach. Figure 1.2 illustrates this process for a $\tau = 2$ XOR system that grows a Sierpinski triangle pattern: from an initial seed tile, tiles attaching by double-strength bonds create a V-shaped edge. The center is filled in by tiles attaching by two single-strength bonds: one from the tile left of the site site, and one from the tile below.

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 [79, 1]. Squares [58, 2] and arbitrary shapes [68] can be efficiently assembled by relatively few tile types, and numerous other properties have been shown [14, 53]. These results rely on cooperativity: with $\tau = 1$, at least for certain classes of tile systems, significant computation can't be performed or is much more difficult [15].

The aTAM is useful for considering the theoretical potential of tile assembly, but is not clearly relevant to 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 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, although we will discuss a mass-action version of the kTAM later. To simplify the analysis, 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 single 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 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, 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 doesn't 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'll assume that the standard free energies of correct bonds can simply be added, with no cooperative effects, and that mismatched bonds don't hinder binding.

5. Bonds are assumed to have identical standard free energies, $\Delta G^\circ_{\text{se}}$, or some integer multiple higher strength bonds. A tile that would attach by a total bond strength of b (e.g. b single-strength bonds) in the aTAM would result in a standard free energy change of $b\Delta G^\circ_{\text{se}}$.

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

$$r_f = k_f[c] \qquad\qquad r_b = k_f e^{b\Delta G^\circ_{\text{se}}/RT+\alpha} \tag{1.1}$$

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, and α is a constant unitless free energy change from other factors, such as the loss of rotational entropy during binding.

By defining $\hat{k}_f \equiv k_f e^{\alpha}$; a sign-reversed, dimensionless free energy $G_{se} \equiv -\Delta G^\circ_{\text{se}}/RT$; and a dimensionless free energy analogue G_{mc}, where $[c] = e^{-G_{mc}+\alpha}$, corresponding to the free tile concentration, we can construct the symmetric rate equations

$$r_f = \hat{k}_f e^{-G_{mc}} \qquad\qquad r_b = \hat{k}_f e^{-bG_{se}} \tag{1.2}$$

that are the core of the kTAM [79].

The kTAM employs assumptions that are not valid in general, but experiments have given some insight into their validity. Crystal growth near equilibrium has been observed directly on a surface, showing single-tile attachments and detachments with rates fitting the kTAM within error [16]. Results from ribbon and nanotube growth largely fit predictions and simulations in the kTAM [43, 63], and many experimental results qualitatively fit the model [3, 10, 57]. Lattice defects and evidence of assembly interactions have been observed, however, motivating designs and conditions to avoid them [16, 22], 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 [22, 43].

To more accurately model particular situations, numerous extensions to the kTAM have been studied. These have included models that account for DNA hybridization energetics and unequal concentrations [17, 20], and mass-action and fixed-volume stochastic models [61, 72] that allow analysis of nucleation and crystal scission. Other aggregate interactions, however, and lattice defects remain open areas for future kinetic models.

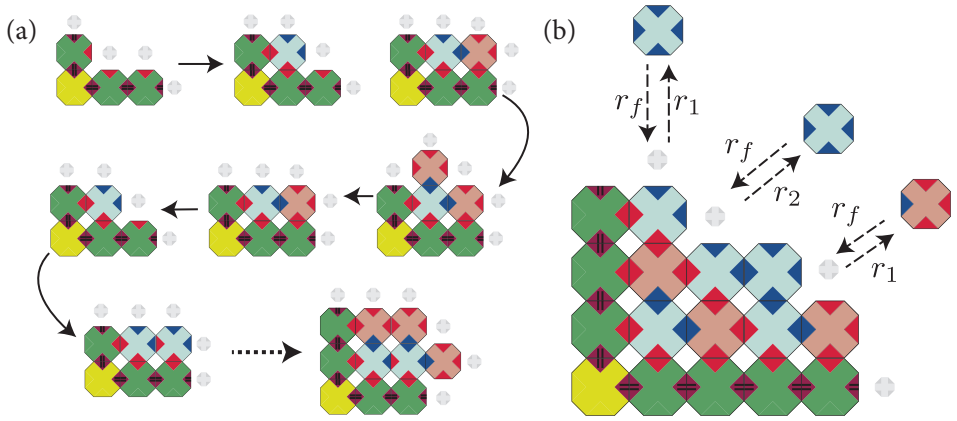

Figure 1.3: The Sierpinski system in the kinetic Tile Assembly Model (kTAM). (a) shows one possible path of attachments and detachments, disregarding the time between each change. (b) shows attachment and detachment rates for various tiles and lattice sites.

Simulations of the kTAM often employ the Xgrow or ISU TAS simulators, which support both aTAM and kTAM simulations, and have some support for kTAM extensions [52, 72, 20].

1.3 Errors and Proofreading

Connections between the aTAM and kTAM

That aTAM tile systems will behave similarly in the kTAM is not immediately apparent. Tiles 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 equilibrium. To examine *growth* in the kTAM, we define a ϵ, where $G_{mc} = 2G_{se} - \epsilon$. This corresponds to the supersaturation of tiles in solution, and the extent to which tile attachment is faster than tile detachment. If 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 close to equilibrium, 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 $r_b/r_f = e^{(2-b)G_{se}-\epsilon}$. 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 $r_2/r_f = e^{-\epsilon} < 1$, resulting in a random walk that will be slightly biased forward. So 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 [79].

This equivalence is only in the limit of slow growth, where $r_f \to 0$. For faster growth, with smaller values of G_{mc}, 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

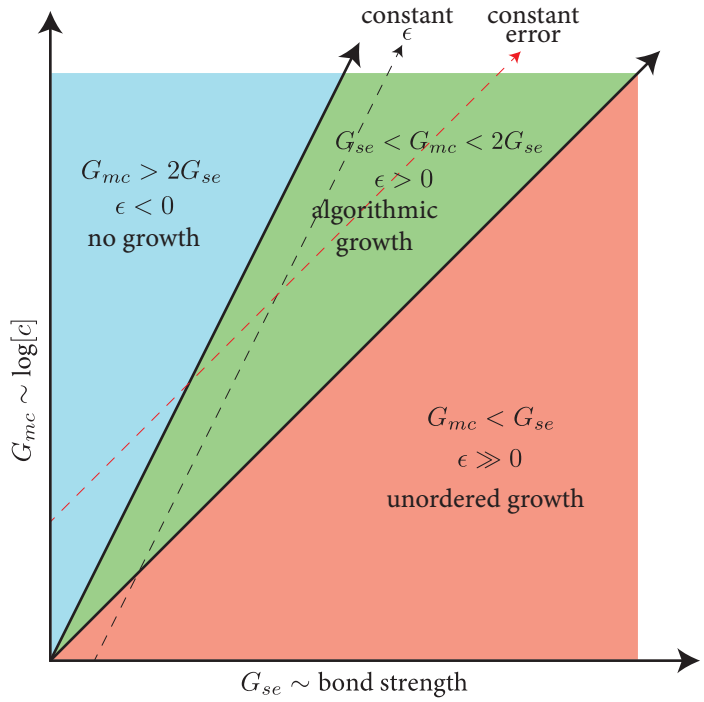

Figure 1.4: Phase diagram for G_{mc} ("**m**onomer **c**oncentration") and G_{se} ("**s**ticky **e**nd strength") in the kTAM with a $\tau = 2$ system. For $G_{mc} > 2G_{se}$, detachment rates are higher than attachment rates of tiles attached total strength 2, and there is no growth. For $G_{mc} < G_{se}$, 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.

detaches, as in Figure 1.5, 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 Figure 1.1(c). Figure 1.4 shows a phase diagram for growth with different values of G_{se} and G_{mc}.

Growth errors present a significant obstacle for algorithmic growth at practical speeds, and were found in simulations and experimental results [57]. Fortunately, growth errors in the kTAM can be simply modeled, and through clever tile system design, error rates can be drastically reduced.

We can construct the *kinetic trapping model* by formalizing our intuitive reasoning [79]. 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 Figure 1.5.

In a filled state, there is a possibility that a further tile can attach in an adjacent site by two bonds, and growth

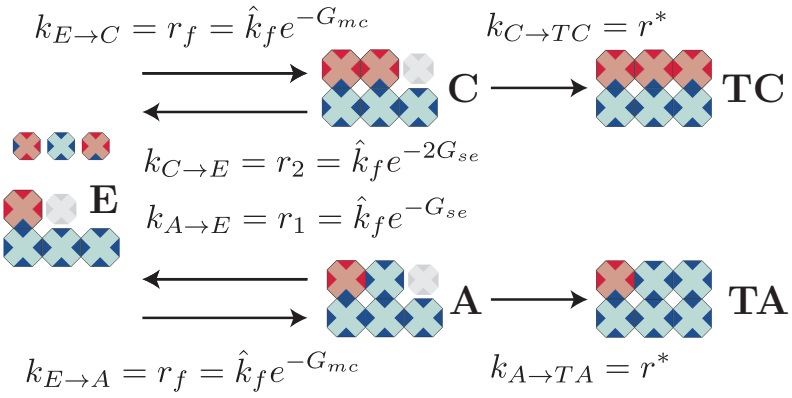

Figure 1.5: Transition rates in the kinetic trapping model, as described in the text.

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—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. The rate at which growth will continue in this fashion can be approximated as the growth rate for two-bond attachments $r^* = r_f - r_2$. This results in transition rates, with the addition of "trapped" states for a correct tile (TC) or almost-correct tile (TA), of

$$\dot{P}(t) = \begin{array}{c} \\ E \\ C \\ A \\ TC \\ TA \end{array} \begin{array}{c} \begin{array}{ccccc} E & C & A & TC & TA \end{array} \\ \begin{pmatrix} -2r_f & r_2 & r_1 & 0 & 0 \\ r_f & -r_2 - r^* & 0 & 0 & 0 \\ mr_f & 0 & -r_1 - r^* & 0 & 0 \\ 0 & r^* & 0 & 0 & 0 \\ 0 & 0 & r^* & 0 & 0 \end{pmatrix} \end{array} P(t) \ . \tag{1.3}$$

where $P(t)$ is a probability vector of the site being in a state at time t. The per-site error rate for growth can be approximated by considering these transitions as a matrix flow problem, and comparing the probability that a correct tile becomes trapped (TC) versus an erroneous tile (TA). If we consider growth near equilibrium, where $G_{mc} = 2G_{se} - \epsilon$ and ϵ small, this results in a per-tile error rate

$$P_{\text{error}} \approx m e^{-G_{se}+\epsilon} \ . \tag{1.4}$$

As expected, in the limit of large G_{mc}, and thus large G_{se}, this error rate goes to 0: by slowing down growth and lowering concentrations, growth errors can be reduced arbitrarily. Our growth rate is $r^* \propto e^{-2G_{se}}(e^{\epsilon} - 1)$, so for small ϵ, $r^* \propto P_{\text{error}}^2$. 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 [7, 67]! In order to build error-free, complex structures using algorithmic self-assembly at high yields, extremely long assembly times would be required; to allow for viable, complex algorithmic self-assembly, some other way to reduce error rates is necessary.

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

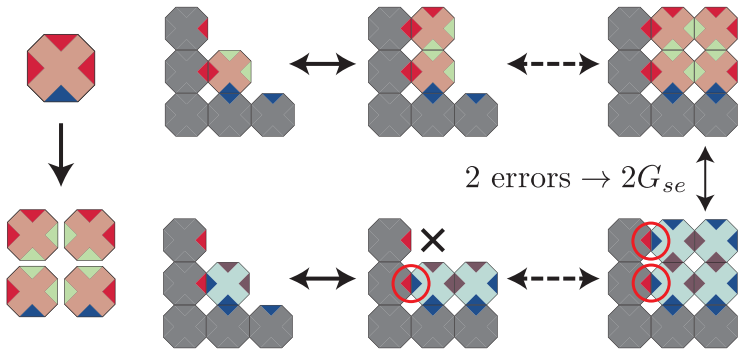

Figure 1.6: The uniform proofreading construction of Winfree and Bekbolatov [80]. (a) shows the transformation of a single tile into a 2×2 block of tiles, with unique internal bonds that for simplicity are not distinguished here; bonds along each edge are also not distinguished. (b) shows the attachment of a correct vs. incorrect tile in a block; for the incorrect block to form, two erroneous attachments must take place.

studied for cellular automata [23]. 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 [41, 13]. The behaviors and potential of several abstract designs have been studied [49, 32, 21, 42], 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 [48].

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 [80] 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 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 Figure 1.6.

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, as $\epsilon \to 0$, 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

$$P_{error} \approx m e^{-KG_{se}} . \tag{1.5}$$

Proofreading thus significantly reduces growth errors as modeled by the kinetic trapping model, and in simulations of a Sierpinski system, Winfree and Bekbolatov found that 2×2 proofreading reduced errors as

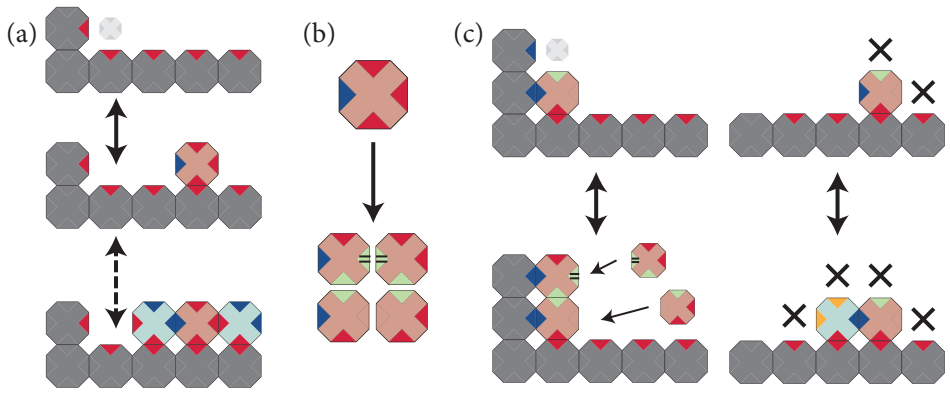

Figure 1.7: Facet nucleation and the snaked proofreading construction of Chen and Goel [7]. (a) shows a facet, along with one possible facet nucleation that can then allow for further growth. (b) shows the construction of a snaked proofreading block (null-strength bonds are denoted by the absence of a colored triangle), while (c) shows the attachment of a correctly-placed tile (left) vs. a tile that attaches along a facet.

predicted [80]. This exponential improvement of error rate is gained, however, at the cost of a K-fold increase in size of the assembled object and thus a K-fold increase in the time required for self-assembly (assuming the same tile concentrations with and without proofreading). Furthermore, with 3×3 and 4×4 proofreading, simulations do not show as much improvement as theory would suggest.

These results for larger proofreading blocks are caused by an additional form of error that can arise. In a lattice site where there is only one adjacent tile with a strength-one bond, for example, along the facet seen in Figure 1.7, *no* tile can attach in the aTAM. A tile attaching here by one bond could allow further attachments along this new, extended facet that would be out of the normal growth order of the system, thus allowing errors described as *facet nucleation errors* [7].

As there are no correct tiles that can attach in sites allowing facet nucleation, the error rate will depend on how common those 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. As facet nucleation, for $\tau = 2$, will occur only by a single weak bond, the resulting rate is $r_{\text{facet}} \propto e^{-3G_{se}}$. 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 $P_{\text{error}} = r_{\text{facet}}/r^* \sim e^{-G_{se}}$. This worst-case error probability is the same as the growth error probability without proofreading, limiting the effectiveness of only preventing growth errors.

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 [7]; in contrast, the "uniform" proofreading of Winfree and Bekbolatov uses all strength-one bonds. Double-strength and null-strength bonds are arranged such that all the tiles on one edge of the block area cannot attach after a single tile attaches by only one weak bond without a further attachment on another edge of the system, 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 [7, 10]. In demonstrating this, Chen and Goel also provide a proof of snaked proofreading's error rate reduction at arbitrary scale. In later research, Soloveichik and Winfree provide an alternative proof [66].

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 assembly [7].

The basic mechanism of snaked proofreading has been tested experimentally, showing a reduction in facet nucleation rates along a long seed structure [10], 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 [3, 4], and more recent works, including those presented in Chapter 5, have reduced error rates even further [64].

However, neither uniform nor snaked proofreading have been fully implemented in a tile system, nor have the larger proofreading blocks that may be possible with SST tiles 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 [56, 67], while proofreading in three dimensions may have the potential to be simpler and more effective [8].

1.4 Nucleation

We have so far only considered individual crystals, nucleated from a chosen 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, tile interactions 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, spurious nucleation is particularly problematic. Rothemund et al [57], for example, saw unseeded assemblies with little resemblance to Sierpinski patterns, and Fujibayashi et al [22] found numerous thin ribbon structures.

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 [62, 61]. 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 [22], 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 $[At] = [A][t](k_f/k_r)$, where k_f and k_r are the forward and reverse rates for the reaction. k_f will be the same as the forward rate constant in the kTAM, while we will use the detachment rate of the kTAM for k_r. This results in

$$[At] = [A][t]e^{-b\Delta G^{\circ}_{\text{se}}/RT - \alpha} \tag{1.6}$$

If we assume all tiles are at equal concentrations, and consider the concentrations of smaller assemblies recursively, it is easily seen that the equilibrium concentration of an assembly A is

$$[A] = [t]^N e^{-B\Delta G^\circ_{se}/RT - N\alpha} = e^{BG_{se} - NG_{mc}} \equiv e^{-G(A)} \tag{1.7}$$

where N is the total number of tiles in A, B is the total strength of all the bonds between tiles in A, and $G(A)$ is the total (unitless) free energy of A (with negative values being more favorable).

For simplicity, we have assumed that tile concentrations are equal and remain constant, though Equation 1.6 does not rely on this, and Equation 1.7 is easily modified for closed systems. For this form of open system, where tiles are "added" to replace depleted tiles, the free energy, and thus concentration, of increasingly larger assemblies is unbounded, but useful results can be obtained for finite assemblies. In practice, when starting near equilibrium with sufficiently low seed concentrations and high tile concentrations, tile concentrations can be made to remain mostly constant, as discussed later in this review.

As an interesting aside, growth error rates can be approximated by this model. If a correct assembly has an equilibrium concentration $[A]$, an assembly $[A']$ that is the same except for a single erroneous rather than correct tile in a location will have, at minimum, one mismatched single-strength bond in place of a correct single-strength bond. This difference in the total bond strength B will result in an equilibrium concentration $[A'] = [A]e^{-G_{se}}$, in line with the kinetic trapping model probability of an error occurring at $\epsilon = 0$. Similarly, if K growth errors are required, the incorrect assembly will have a total bond strength that is at least K less than the correct assembly, and an equilibrium concentration of $[A'] = [A]e^{-KG_{se}}$. These error rates do not depend upon ϵ, but ϵ must be small for the system to approach equilibrium at all.

Unseeded nucleation of crystals can be modeled by *critical nuclei*, 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 [61], 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 Figure 1.8. For $\epsilon \leq 0$, growing larger squares will always be unfavorable; $G(A)$ will always increase. For small $\epsilon > 0$, 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. So long as ϵ is small enough, little nucleation should take place without some desired seed structure to initiate it.

For the Sierpinski system in Figure 1.2, 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.

One system that allows the formation of ribbons with low spurious nucleation is the zig-zag ribbon system, illustrated in Figure 1.9 [61]. 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

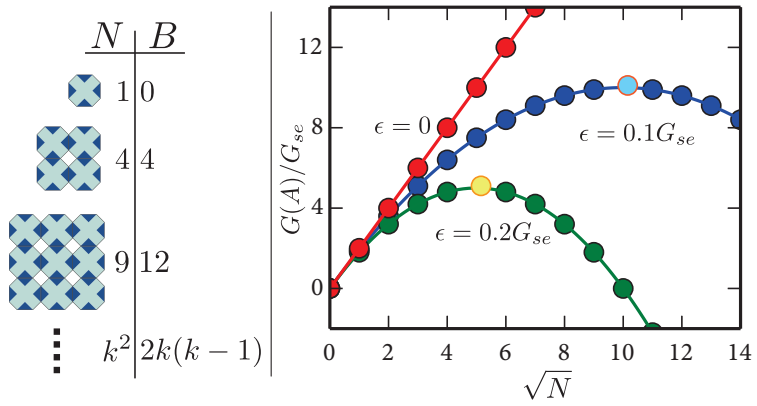

Figure 1.8: Critical nuclei for uniform squares: for smaller ϵ, squares of increasing size must form before the number of bonds vs number of tiles added makes further growth favorable, for $\epsilon = 0$, growth is never favorable. The highest $G(A)$, lowest concentration squares here are critical nuclei that bound nucleation of larger crystals.

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 [61]. Thus, by designing zig-zag systems of increasing width, spurious nucleation can be reduced arbitrarily.

Experimentally, seeding growth can be done relatively simply by using a structure of uniquely-addressed tiles to reliably form a large seed structure. Often, however, substitute structures built using other methods, such as DNA origami or assembly PCR, are used along with adapters that allow DNA tiles to attach; this method has been used for one-dimensional seeds [57, 3], as well as rectangular and cylindrical seeds for ribbons and nanotubes [4, 43].

Spurious nucleation of zig-zag ribbons has been measured experimentally, despite difficulties measuring very low nucleation rates, with results indicating a noticeable but less-than-predicted dependence on width [63]. Nucleation by origami seeds at width 4, however, was 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 [43]. These forms of nucleation control have been incorporated into some experiments [4, 64] that have qualitatively shown far fewer spuriously nucleated assemblies than those seen in experiments with seeded growth but a smaller barrier to unseeded nucleation [22].

1.5 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

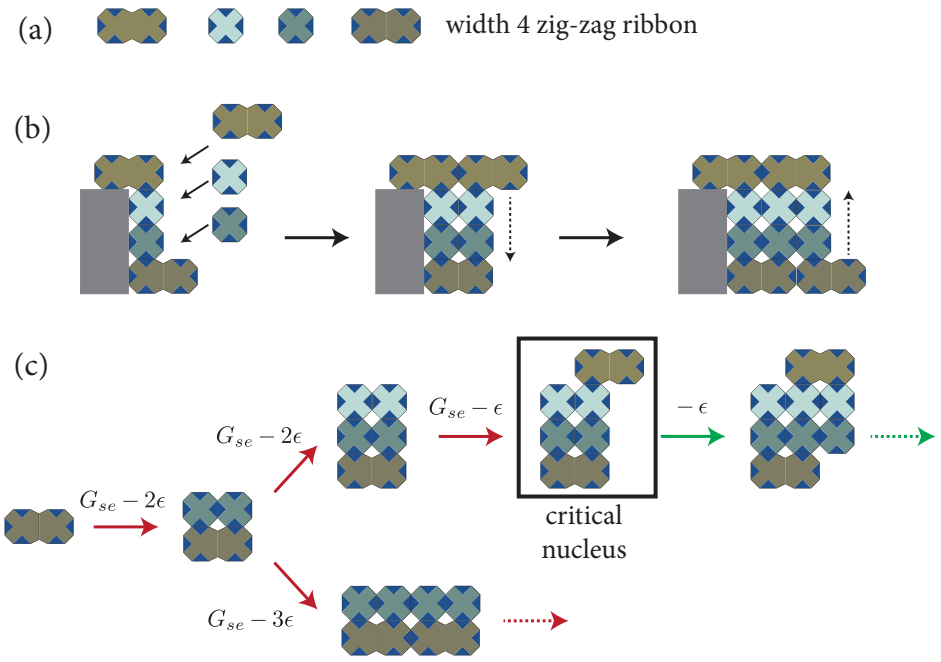

Figure 1.9: The zig-zag tile system [61]. (a) shows a width 4 tile set, which can be expanded to arbitrary width by adding tiles in the center; all top and bottom bonds are unique. Seeded growth (b) constructs a ribbon with rows that grow in alternate directions and double tiles that nucleate each subsequent row. Unseeded growth (b) 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.

of free tiles in solution will be depleted, increasing G_{mc} and moving the system closer to the 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 [81]. So long as this temperature decrease is slow enough, the system will remain near equilibrium, allowing for optimal growth.

This technique has been used in numerous experiments, and is the standard method for both non-algorithmic and algorithmic DNA tile systems [77, 4, 20, 57]. 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 vastly 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.

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 choices of concentrations remains an open problem: there has been some theoretical work [9, 17] and concentration choices based on simulations (Chapter 4), but concentrations used in experiments, which have often not been equal [4], 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 [64]. 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 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 that earlier work using ordinary annealing. For example, error rates for bit-copying ribbons were reduced from 0.26% [4] to 0.034% [64] per 2×2 block, and similar improvements for a binary counting ribbon are presented in Chapter 5.

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 [64]. 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.

1.6 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 [33]. 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, as shown in Chapter 5. 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, and is discussed in Chapter 4.

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 [4, 64]. The implications of doing so, and the interactions between the methods, are unclear: in the zig-zag tile systems of Barish et al 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.

Chapter 2

Direct AFM Observation of DNA Tile Self-Assembly

While the theoretical implications of models of DNA tile self-assembly have been extensively researched, and such models have been used to design DNA tile systems for use in experiments, there has been little research testing the fundamental assumptions of those models. In this chapter, we use direct observation of individual tile attachments and detachments of two DNA tile systems on a mica surface imaged with an atomic force microscope (AFM) to compile statistics of tile attachments and detachments. We show that these statistics fit the widely-used kinetic Tile Assembly Model, and demonstrate AFM movies as a viable technique for directly investigating DNA tile systems during growth rather than after assembly.

2.1 Introduction

Macromolecular self-assembly is common in nature as a mechanism for the construction of complex structures: the construction of microtubules from tubulin in the cytoskeleton is one notable example. Analogous self-assembling systems have been designed using a variety of often biology-derived "building blocks" to construct large, complex structures with nano-scale resolution suitable for a diverse range of applications [74, 45, 47].

Crystals formed from DNA tiles are particularly interesting in their ability to self-assemble based on designed, single-stranded DNA "sticky ends" that attach to complementary ends on other tiles [81]. This allows crystals to be designed that have complex, programmable structure despite consisting of only a relatively small number of different tile types [3, 4, 22, 11].

In this regard, the tiles are similar to Wang tiles [76], which were used as a basis for the abstract Tile Assembly Model (aTAM) of DNA tile assembly [79]. In this model, tiles fill empty lattice spaces if they can attach by at least a certain number of correct bonds, usually two. While the model ignores the kinetics of tile attachment and ignores detachment entirely, it does serve as a useful basis for designing the logic behind both simple and complex tile sets. In an attempt to make a model useful for understanding experimental results, the kinetic Tile

Reproduced with permission from Constantine G Evans, Rizal F Hariadi, and Erik Winfree. "Direct atomic force microscopy observation of DNA tile crystal growth at the single-molecule level." In: *J. Am. Chem. Soc.* 134 (2012), pp. 10485–10492. DOI: 10.1021/ja301026z.

Assembly Model (kTAM) was devised [79, 57]. In this model, free tiles in solution attach to empty binding sites at a rate determined only by their concentrations, and detach at a rate determined only by the number of matching sticky-end bindings holding them to the lattice; in other words, an on-rate and off-rate per tile of

$$k_{\text{on}} = k_f[c] \qquad\qquad k_{\text{off}} = k_f e^{-b\Delta G^\circ_{\text{se}}/RT+\alpha} \tag{2.1}$$

where k_f is the forward rate constant, $[c]$ is the concentration of free tiles of the type of interest, b is the number of correct bonds holding a particular tile in place, $\Delta G^\circ_{\text{se}}$ is the standard free energy of a single bond, and α is a constant factor accounting for other binding energies. In order to remove α, parameters $\hat{k}_f \equiv k_f e^{\alpha}$ and dimensionless "free energy" G_{mc} where $[c] = e^{-G_{mc}+\alpha}$ are used; these, along with a dimensionless energy $G_{se} \equiv \Delta G^\circ_{\text{se}}/RT$ result in

$$k_{\text{on}} = \hat{k}_f e^{-G_{mc}} \qquad\qquad k_{\text{off}} = \hat{k}_f e^{-bG_{se}} \tag{2.2}$$

where both G_{mc} and G_{se} are usually positive numbers.

2.2 EXPERIMENTAL METHODS

System Description

The systems described by the kinetic model are typically implemented with DNA tile structures similar to those shown in Figure 2. For example, in these experiments, we used two previously-published tile systems that had two different molecular structures. The first was R00/S00, a system of two tiles that form a two-dimensional lattice [57]. The second was NAoMI-B [28], a system with a single tile type that assembles in solution into three-dimensional tubes, but in our experiment formed into 2-dimensional lattices owing to association with the mica surface [59]. The R00/S00 tiles have a DAO-E molecular structure: two crossovers between helices (**D**ouble), **A**ntiparallel strand orientations (a strand in one helix has the opposite orientation after crossing over to the other helix), an **O**dd number of half-turns between crossover points in a single tile, and an **E**ven number of half-turns between crossovers in adjacent tiles. NAoMI-B has a DAO-O structure, which is similar except with an odd number of half-turns between crossovers in adjacent tiles (see Figure 2(a)) [59, 19, 81]. Both of these systems have similar properties, using tiles with four binding sites comprising five-base (R00/S00) or six-base (NAoMI-B) single-stranded sticky ends that attach to complementary sticky ends on other tiles.

The kinetic model has been widely used as a quantitative model to investigate tile systems [80, 7, 62], and with some extensions (sequence-dependent $\Delta G^\circ_{\text{se}}$s, varying concentrations) has been used in numerous computational simulations of crystal growth [20, 57, 10, 63]. Algorithmic growth behavior suggested by the kTAM and design techniques like proofreading [80] that have been based on the model have translated well into experimental findings [4, 22, 10], but no investigation has been done directly into the model's mechanistic assumptions.

In particular, the kinetic model makes the following simplifying assumptions, or idealizations [79]:

1. Free monomer concentrations remain constant during growth.
2. Crystals are perfectly rectilinear, two-dimensional, and free of defects, with single tiles that attach and detach one at a time.

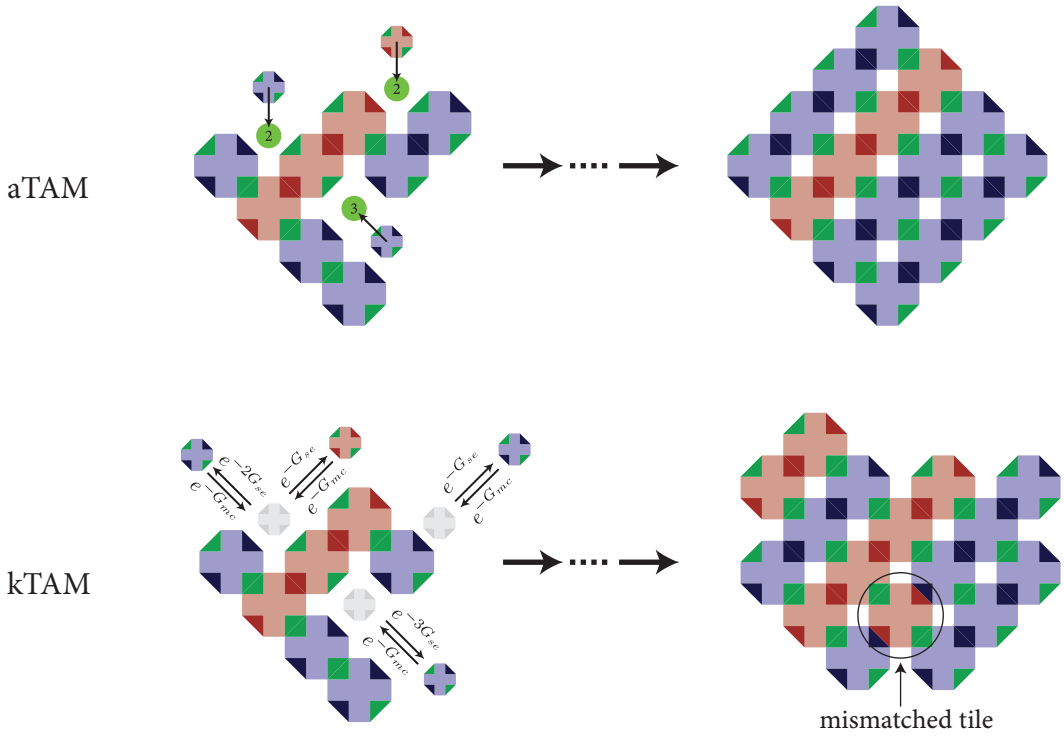

Figure 2.1: Illustrations of the abstract (aTAM) and kinetic (kTAM) Tile Assembly Models. In the aTAM, tiles attach only in locations with at least a certain number of correct bonds (here, two), and never detach. In the kTAM, tiles attach with rates dependent only on concentration, and detach based on the number of correct bonds; a few examples of attachment and detachment rates for given crystals are shown, omitting the common factor $\hat{k}_f$. The right side shows the same crystals at a later time; in the aTAM, this is the unique and final result of tile attachments, while in the kTAM it is one possible state that can continue to grow or shrink. In the kTAM, a crystal can also contain mismatched tiles; a mismatched tile is highlighted by a black circle in the bottom right crystal.

3. Attachment rates are constant and equal regardless of the number of correct or incorrect bonds for a tile at the attachment site, as shown in equation 2.2.

4. Detachment rates are exponentially dependent on (a) the number of correct attachments the tile has to the crystal, as shown in equation 2.2, and (b) on nothing else.

Assumption 1, while not true for growth in general, can be approximated by using seeded nucleation with a small enough concentration of seeds in comparison to tiles and a short enough period of sufficiently slow growth [4, 63, 62]. Assumption 2 is beyond the scope of this paper; it is worth mentioning that we were able to obtain data without interference from lattice defects or multiple attachments, though violations of the assumptions have been seen [3, 57, 22]. Assumptions 3 and 4 are of primary interest to us here, as there are a number of physical effects that could cause experiments to deviate from the assumptions. It is possible, for example, that electrostatic or steric effects in the molecular structure of the lattice could impact attachment and detachment rates in different ways depending upon the location of the binding sites in the lattice or other factors. Another, specific example is that tiles attaching by two bonds in both of our tile systems can have two

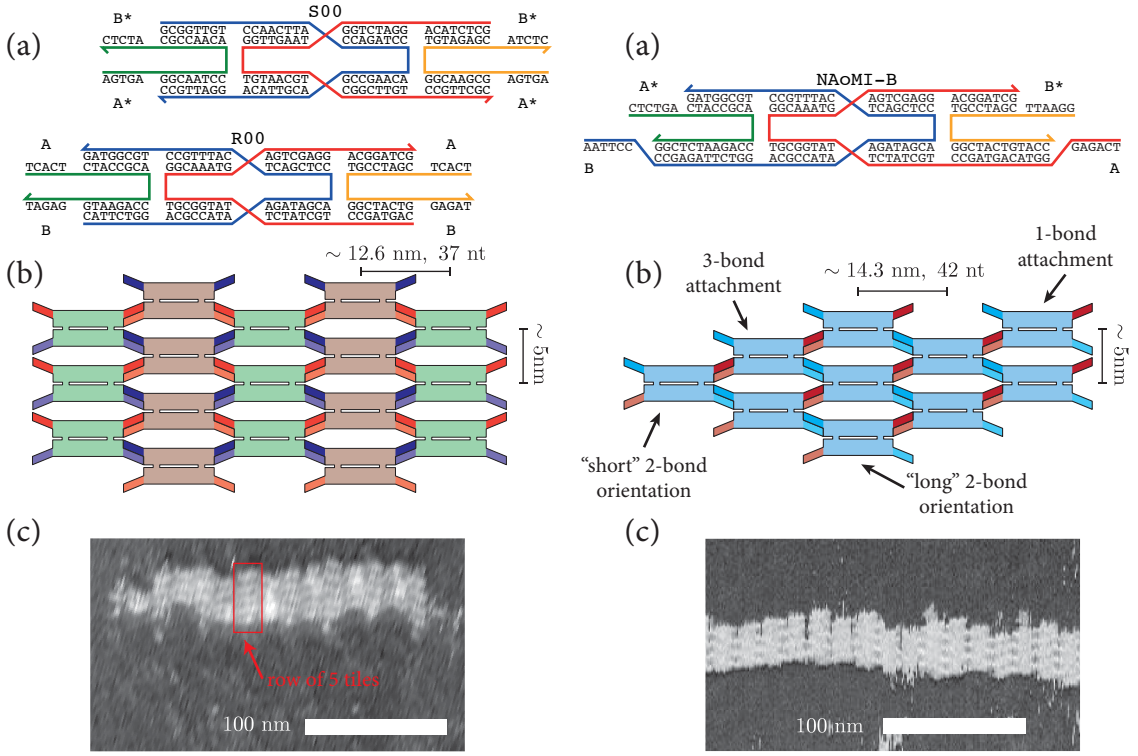

Figure 2.2: DNA tile structures and their resultant lattices for R00/S00 (left side of figure) and NAoMI-B (right side) systems: (a) shows sequence designs of tiles: A, B, ... represent symbolic sticky ends, with * denoting complements and barbs indicating 3' ends. (b) shows example lattice structures resulting from perfect (no-mismatch) growth, along with examples of the "short" and "long" orientations of tiles attached by two bonds, and (c) shows example AFM images of lattices.

different orientations of attachment to the lattice (Figure 2(b)); entropic or electrostatic effects could cause these orientations to have different rates. Any of these effects could violate 3 and 4(b). Tiles with multiple bonds to the lattice could also attach and detach via processes that might violate the simple additive assumption of 4(a), such as detachment and attachment of a single bond in a tile with four bonds.

Atomic Force Microscopy

In our experiments, we examined assumptions 3 and 4 by using fluid tapping-mode atomic force microscopy (AFM) to create movies of individual tiles attaching and detaching in lattices. The lattices were in a fluid environment with free monomers, but were weakly bound to mica for imaging while growing [51]; while most systems have previously been grown with lattices in solution, there have been systems successfully grown on surfaces in a similar fashion [26, 70, 35]. With a system in growth-biased conditions, we were able to create movies of individual crystals growing rapidly on the surface (figure 2.3(a) and supplementary movie 3); however, attachment of individual tiles took place too quickly for us to distinguish single-tile changes, and in any case, there would have been too few detachments to obtain sufficient statistics. By instead imaging near an equilibrium that we reached via adjustment of temperature, tile concentration, and other parameters (discussed later), we were able to observe distinct, single-tile attachments and detachments on single crystals, and compile statistics of these changes based on number of bonds and tile orientation.

Throughout the experiment, we used TAE/Mg^{2+} buffer (40 mM Tris, 20 mM acetic acid, 1 mM EDTA, 12.5 mM Mg acetate), and PAGE-purified strands (Integrated DNA Technologies). For R00/S00, we annealed mixes of strands (25 nM each) in buffer for each of the two individual tiles separately with a temperature decrease from 90 °C to 20 °C at 1 °C/min, starting with 90 °C incubation for 5 min; crystals could only grow when these were later mixed. For NAoMI-B, all strands were mixed together at 1 μM per strand and annealed with a 70 min linear temperature decrease from 90 °C to 20 °C; this created nanotubes that were then diluted to a 50 nM per tile concentration. As in previous work [59], nearly all DNA tile nanotubes break open and unfold during imaging, so that they appear as thin strips of two-dimensional tile crystals. Imaging was performed on a MultiMode AFM (Bruker) with a Nanoscope IIIa controller, with temperature held at 33.4 °C by an in-stage heater (R00/S00) or at 32 °C (NAoMI-B) by a space heater heating the entire room. Around 40 μL buffer was deposited on freshly cleaved mica; this was allowed to equilibrate in temperature in the AFM and was then imaged without any tiles. Around 5 μL of each annealed tile mix (two for R00/S00, one for NAoMI-B) was usually added to this buffer while still imaging. While surface binding effects, evaporation, and our AFM's open fluid cell design made it difficult to know the effective tile concentrations in the sample chamber during AFM imaging, relative concentrations were adjusted such that the crystals were near equilibrium, simplifying analysis, and were kept equal across different tile types for R00/S00. Imaging speed was kept as fast as possible while still allowing individual tile resolution, which for our AFM resulted in a speed of around 80 sec/frame at a scale of around 600 nm with a resolution of 512 pixels per scan line and 512 scan lines per image. Images of five crystals from three movies were used: one long, large movie with two crystals for R00/S00, and two smaller movies, one with two crystals and one with a single crystal, for NAoMI-B.

Our ability to observe individual tile attachments and detachments as they occurred was made possible by a number of coincidental advantages found by our AFM use in prior experiments with DNA tiles. In fluid-tapping mode, free tiles were able to remain in our buffer solution rather than binding to mica, allowing us to achieve a suitable concentration of free tiles in solution for growth. The ability to reliably image in the 30 to 40 °C range allowed us to maintain a temperature near equilibrium. Adjusting buffer salt concentrations would have allowed us to ensure that reasonably-sized lattices would bind to mica while monomers remained in solution, taking advantage of salt modulation of mica-DNA binding [6, 51]; in practice, we found that our buffer's salt concentrations were already acceptable. By combining control of temperature and tile concentrations, we were able to tune attachment and detachment rates such that individual events could be resolved at our imaging speed. We additionally used nanoAnalytics analog Q-control (Asylum Research) and tuning of AFM parameters in order to minimize any damage to the lattices during imaging.

2.3 Analysis

To analyze the resulting movies, as AFM images tend to vary considerably in contrast, distortion, and sample location from frame to frame, we manually translated visible crystals into lattice structures by counting the number of tiles in each visible "row" of tiles, and recording the offset of this row from adjoining rows. Lattices from subsequent frames were then automatically aligned, and combined with estimated in-frame position data taken from a combination of average lattice spacing and crystal position and angles. The position and timing of events at each lattice position were derived from this data. Occasionally, tiles adjacent to each other would both attach or both detach between the same frames, creating ambiguous events, but this didn't affect the results significantly, and was compensated for by assuming that the most probable order of events was what actually

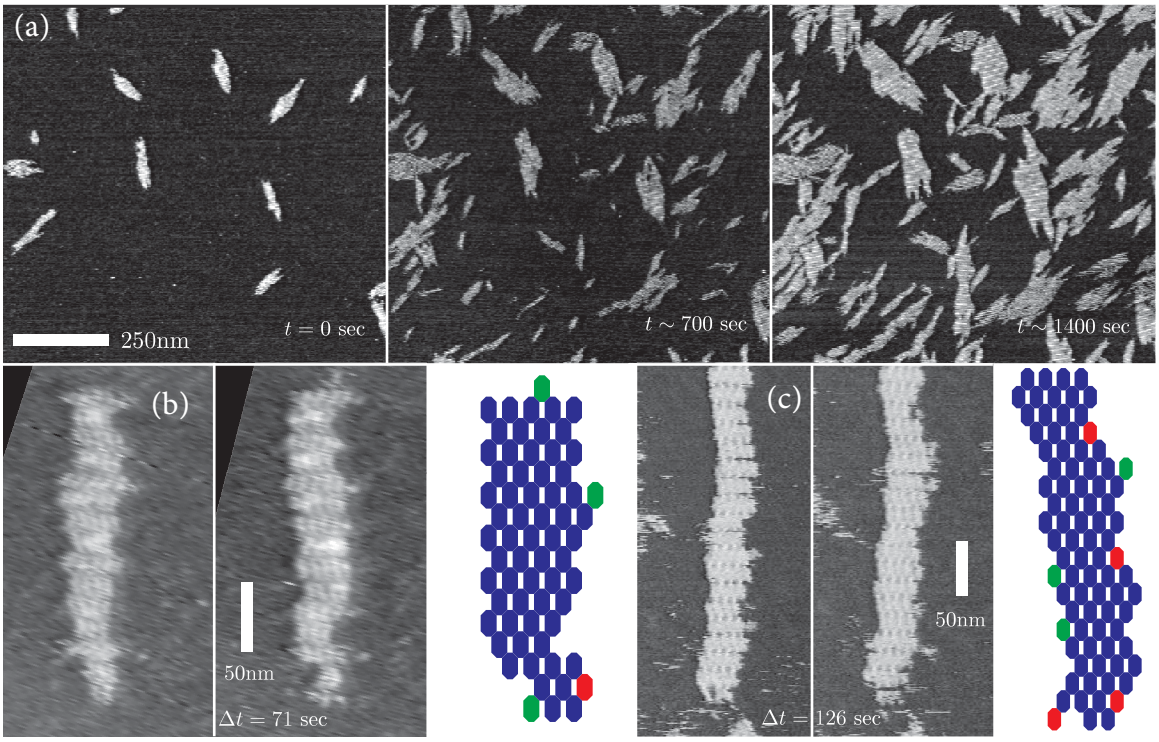

Figure 2.3: AFM images of R00/S00 and NAoMI-B systems: (a) shows several frames from a movie of quickly-growing R00/S00 crystals. (b) with R00/S00 and (c) with NAoMI-B show adjacent frames in a near equilibrium movie (supplementary movies 1 and 2) and their interpretation: blue tiles are present in both frames, while green tiles attached and red tiles detached from one frame to the next.

took place: that detaching tiles attached by the least number of bonds in the older frame detached first, and that attaching tiles attached by the most number of bonds attached first.

From this data, consisting of 115 identified attachment and detachment events for R00/S00 and 96 for NAoMI-B (not including tiles that stayed attached or lattice sites that remained empty), our challenge is to separately infer the on-rate and off-rate for each type of crystal site, and then together (globally) to extract G_{mc}, G_{se} and $\hat{k}_f$. The agreement (or lack of it) between the separate and global fits will allow us to assess assumptions 3 and 4.

Each position in a lattice can be in two states: filled (F) or empty (E). From one frame to the next, we can therefore have a filled position remained filled (FF), an empty position remain empty (EE), a filled position become empty (FE), or an empty position become filled (EF). Each pair of measurements of each position in the lattice from one frame to the next constitutes a Bernoulli trial i with a probability p_i of remaining the same (FF, EE) and a probability $1 - p_i$ of changing (FE, EF). Depending upon the initial state, these probabilities are dependent upon the time between the two measurements Δt_i (which can take into account dropped frames, position within the frame, and AFM irregularities, rather than just frame rate), and an attachment or detachment rate λ^c_θ, where c is either EF or FE (determined by the initial state of the position) and θ denotes the presumed rate-affecting parameters of the set of locations we're considering: the number of bonds surrounding the position, and, when accounting for it, the orientation of two-bond positions. As we assume these rates are constant for given parameters, the probability of a trial i over a time Δt_i resulting in no change is $p_i = e^{-\lambda^{c_i}_{\theta_i} \Delta t_i}$.

For a set of trials with the same θ and starting state (empty or full), we can use Bayesian inference to provide the probability of a rate λ_θ^c given our data. The *a posteriori* probability of our data—a set of trials $NC_{\theta,c}$ where a location remained unchanged (FF or EE), and a set of trials $C_{\theta,c}$ where a location changed (FE or EF)—with an assumed rate λ_θ^c is simply the product of the individual Bernoulli distribution probabilities:

$$p(\text{data}|\lambda_\theta^c) = \prod_{i \in NC_{\theta,c}} e^{-\lambda_\theta^c \Delta t_i} \prod_{i \in C_{\theta,c}} \left(1 - e^{-\lambda_\theta^c \Delta t_i}\right) \tag{2.3}$$

Bayesian inference would therefore assert that the probability of a rate λ_θ^c given our data is

$$p(\lambda_\theta^c|\text{data}) = \frac{p(\lambda_\theta^c)}{p(\text{data})} \left[\prod_{i \in NC_{\theta,c}} e^{-\lambda_\theta^c \Delta t_i} \prod_{i \in C_{\theta,c}} \left(1 - e^{-\lambda_\theta^c \Delta t_i}\right) \right] \tag{2.4}$$

Using the uniform prior $p(\lambda_\theta^c)$, both our prior and our likelihood $p(\text{data})$ are independent of λ_θ^c, and so can be considered a constant normalization factor to our probability. Thus our inferred rate of attachment or detachment given our set of data is the λ_θ^c that maximizes $p(\text{data}|\lambda_\theta^c)$ from equation 2.3. Our probabilities are given by normalizing $p(\text{data}|\lambda_\theta^c)$ over a reasonable range of λ_θ^c values (such that the *a posteriori* distribution drops to approximately zero on the boundaries) and using it as a probability density function.

This method provides us with the rates of attachment and detachment per lattice position for a given set of parameters. kTAM's assumptions 3 and 4 provide that these rates should be related by equation 2.2; in other words, for some $\hat{k}_f$, G_{mc} and G_{se}, we should have

$$\lambda_\theta^{EF} = \hat{k}_f e^{-G_{mc}} \qquad\qquad \lambda_\theta^{FE} = \hat{k}_f e^{-b(\theta) G_{se}} \tag{2.5}$$

Using Bayesian inference again, we can obtain G_{se}, G_{mc} and $\hat{k}_f$ for a given data set of detachments with *varying* numbers of bonds by using G_{se}, G_{mc}, and $\hat{k}_f$ as parameters, using equation 2.5 to derive detachment rates for different θ parameters, and including all data in the sets NC and C:

$$p(G_{se}, G_{mc}, \hat{k}_f|\text{data}) = \frac{p(G_{se}, G_{mc}, \hat{k}_f)}{p(\text{data})} \left[\prod_{i \in NC} e^{-\Delta t_i \lambda_{\theta_i}^{c_i}(G_{se}, G_{mc}, \hat{k}_f)} \prod_{i \in C} \left(1 - e^{-\Delta t_i \lambda_{\theta_i}^{c_i}(G_{se}, G_{mc}, \hat{k}_f)}\right) \right] \tag{2.6}$$

As we can again use a uniform prior, the probability density for G_{se}, G_{mc} and $\hat{k}_f$ can be derived by ignoring $p(\Delta G^\circ_{\text{se}}, G_{mc}, \hat{k}_f)/p(\text{data})$ and normalizing $p(G_{se}, G_{mc}, \hat{k}_f|\text{data})$ over a sufficiently large parameter space.

In order to verify that our analysis produced accurate results when presented with data fitting the kinetic model, the entire analysis, with the exception of the manual data entry, was applied to simulated data taken from the Xgrow kTAM simulator [72]. Parameters were chosen to match conditions in our experimental results, the amount of data was chosen to be similar, and the output was modified to mimic the line-scanning of an AFM. The results (Figure 4(b,d)) suggested our analysis methods were sound in the face of finite scanning rates and multiple events.

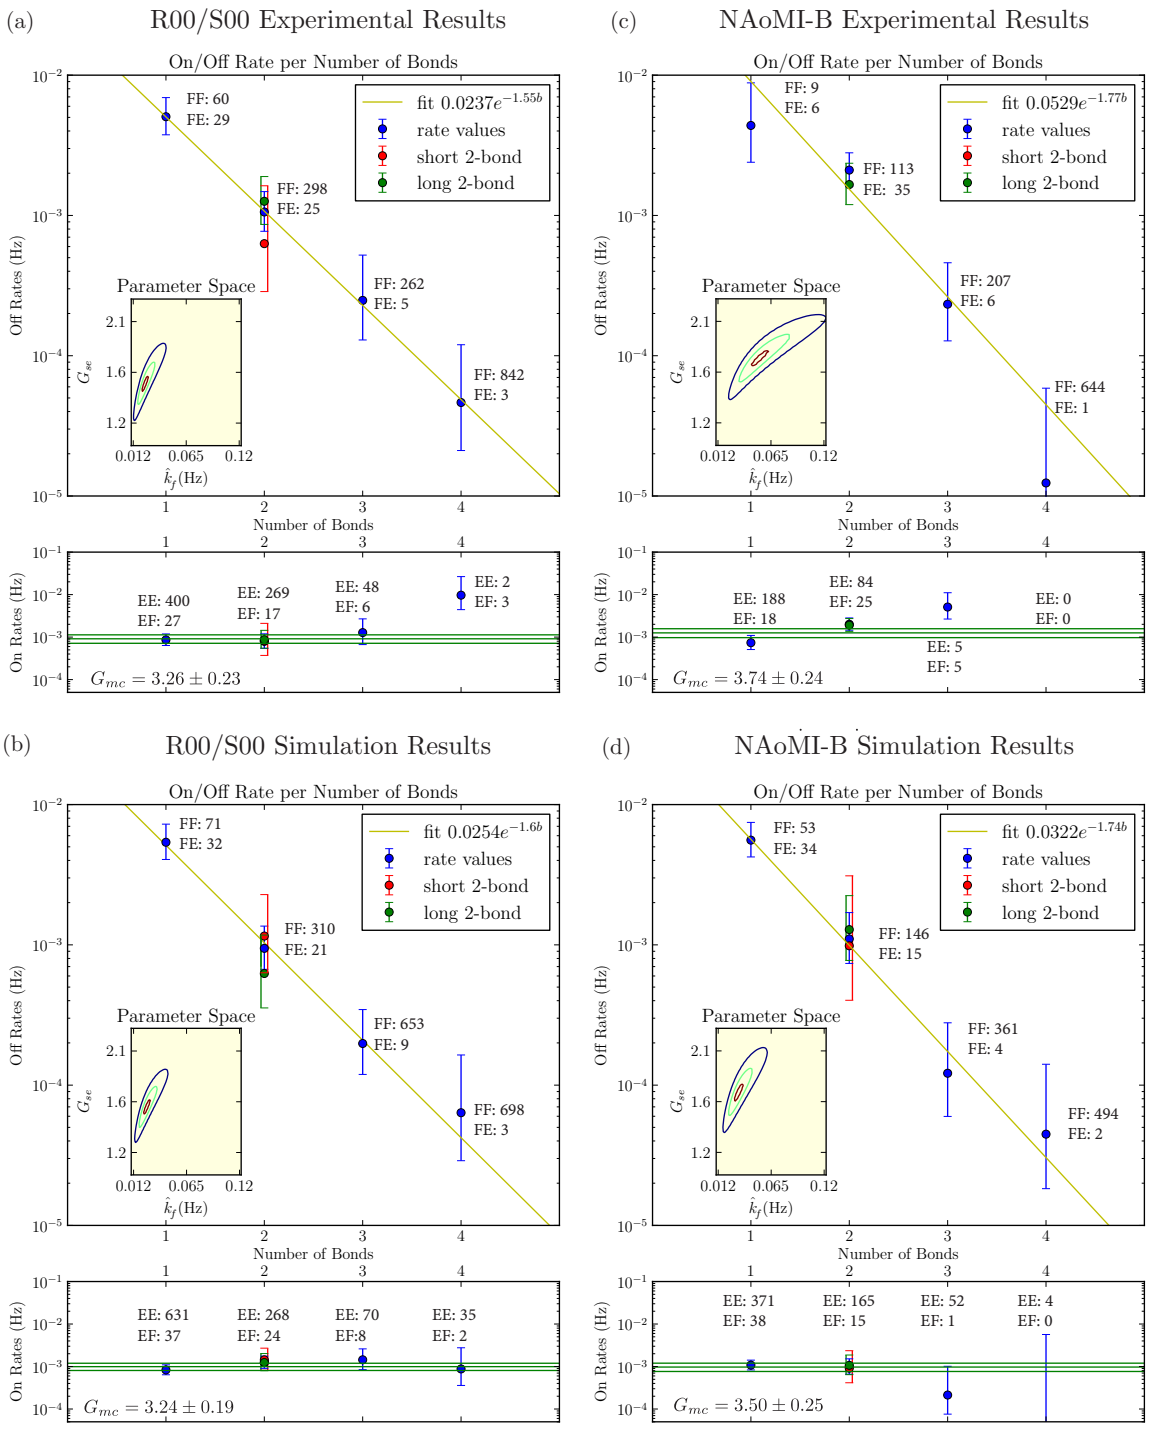

Figure 2.4: Plots of detachment and attachment rates for R00/S00 (left top) and NAoMI-B (right top) lattices. Bottom figures show simulation data with similar parameters. Insets in detachment rate plots show activation rate/G_{se} parameter space, with lines representing $p = 0.1,\ 0.5$ and 0.9. Lines in attachment rate plots show mean and 90% confidence interval for G_{mc}.

2.4 Results and Discussion

Our results, shown in figure 2.4, were compiled separately for R00/S00 and NAoMI-B. Ignoring the possibility of incorrect bonds and lattice defects in the system (assumption 2), and assuming constant monomer concentrations (assumption 1), our results largely suggest that assumptions 3 (equal attachment rates) and 4(a) (detachment rates exponentially dependent on the number of bonds) of the kTAM are valid within our experimental error. Additionally, while observing relatively large error bounds owing to the smaller amount of data for each orientation, we found that tile orientation did not alter detachment rate beyond experimental error, and thus did not violate assumption 4(b) (detachment rates dependent upon nothing but number of bonds).

As for the consistency of the quantitative results, it is worth remembering that NAoMI-B uses 6 nt sticky ends, while R00/S00 uses 5 nt sticky ends. Using the rough approximation that G_{se} scales linearly with the number of nucleotides, this suggests that the NAoMI-B G_{se} should be around 1.2 times that of the R00/S00 G_{se}. Our results give a ratio of around 1.14 ± 0.15, in line with this approximated expectation.

There were several key differences between our experimentally observed data and simulated data, many of which likely originate from differences in the relative numbers of events observed and ultimately in overall crystal "shape." In particular, data from simulation, while resulting in similar detachment rates, had significantly more empty lattice locations with three or four connecting bonds ($b = 3$ or $b = 4$). These sites are the result of a crystal that is, subjectively, more "spindly" and less compact than those we found experimentally, with long "arms" of tiles. Whether "spindly" crystals actually form is a question that would require a different experiment, but regardless of their formation, it is quite likely that crystals of that shape would not cleanly bind to mica, or would produce AFM images unsuitable for analysis; our experiment therefore inadvertently selected against crystals of this form, and instead selected for relatively solid, compact crystals that were easily imaged. In a related shape issue, the fact that imaged NAoMI-B crystals originated from nanotubes in solution meant that all observed crystals were long, narrow ribbons, resulting in no data for short-orientation two-bond attachments, as our data never included the ends of the crystals.

Our analysis of attachment and detachment rates involved only statistics of detachment and attachment for single tiles. The rates derived from such an analysis would ideally be independent of crystal shape, with only the size of the error bounds differing owing to differences in the number of events. In some cases, however, the significantly smaller number of events likely resulted in anomalous observations having a large effect on attachment and detachment rates. For R00/S00, the attachment rate for $b = 4$ is calculated from only 5 events (2 empty-to-empty and 3 empty-to-full), and is significantly higher than expected by kTAM. For NAoMI-B, the attachment rate for $b = 3$ is similarly calculated by only 10 events, and is also significantly higher than expected. By comparison, $b = 3$ for R00/S00 involved 54 events, and was in line with expectations.

With such events made rare by crystal shape, the chance of the rates being affected by errors in image interpretation, anomalous imaging results, or erroneous assumptions about the order of tile attachment is significantly increased. We expect that this accounts for the discrepancies seen in attachment rates, as all unexpected attachment rates have very few recorded events.

In addition to the directly obtainable results, we can also give rough estimates of other parameters. Bayesian inference suggested $\hat{k}_f$ values of 0.024 /M/sec and 0.053 /M/sec for R00/S00 and NAoMI-B, respectively. If we assume what should be a reasonable order-of-magnitude estimation of a 10 nM tile concentration, this

results via equations 1 and 2 in an $\alpha \sim -15$ and $k_f \sim 10^6$ /M/sec. Such a k_f value is on the same order of the $6 \cdot 10^5$ /M/sec used by Winfree [79], but the α is very different than the $+\ln(20)$ used there for simulations of growth in solution [79]. $RT\alpha$ can be interpreted as a free energy adjustment that remains constant for all attached tiles, as opposed to the $\Delta G^\circ_{\text{se}}$ which is scaled with the number of bonds. While the positive α for growth in solution can be interpreted as due to the loss of rotational entropy when a tile binds to the crystal, negative α corresponds to an energetically favorable contribution to attachment in all cases, as could be expected for growth on mica where all tiles involved bind to the mica surface [51].

In comparing our results to other experiments, it is necessary to convert standard free energies to our unitless $G_{se} = \Delta G^\circ_{\text{se}}/RT$ (equations 1 and 2), and in many cases make a change in sign per our convention. The G_{se} we calculate from the statistics is significantly lower than that found by experiments in solution rather than on a mica surface; this corresponds to a significantly "weaker" binding energy contribution for individual sticky ends. In particular, Schulman et al [63] measured an in-solution $\Delta G^\circ_{\text{2se}} = -9.43 \pm 0.21$ kcal/mol at 37 °C for two bonds, corresponding to $G_{se} = 7.7 \pm 0.17$. Nangreave et al [46], using a different form of tile with 6 nt sticky ends, found a $\Delta G^\circ_{\text{se}}$ on the order of -3 to -6 kcal/mol depending upon the form of attachment, resulting in a G_{se} of between 6 and 10. By comparison, our rates of $G_{se} = 1.55$ and $G_{se} = 1.77$ are extremely low. Indeed, a simple order of magnitude comparison of kTAM-expected detachment rates derived from the 7.7 value of Schulman would suggest that, were the $\Delta G^\circ_{\text{se}}$ in our on-mica conditions even close to so large, our entire experiment, and viewing growth with any level of detail at all, would be infeasible: the difference in rates between even tiles attached by two and three bonds would be on the order of $e^{-2 \cdot 7.7}/e^{-3 \cdot 7.7} \sim 2200$.

There are a number of possible factors involving DNA binding with mica and DNA hybridization near a mica surface which could explain this discrepancy. A difference in binding for ssDNA vs dsDNA to mica, for example, could result in a modified G_{se} for sticky ends, as two free, single-stranded sticky ends could have a stronger binding to mica than the bound, double-stranded combination of the two ends. Also, while DNA hybridization near a mica surface has not been widely studied, it is reasonable to expect that there could be significant differences from hybridization in solution: at the very least, the "salt bridge" of divalent cations (in our case, Mg^{2+}) between the DNA and mica would result in a very different and likely non-homogenous ionic environment [51].

2.5 Conclusions

With this work, we've shown that we can make single-tile-resolution AFM movies of single DNA tile system crystals at a sufficient frame rate to provide a consistent view of the steps involved in assembly. While we've used these movies to verify some key assumptions of the kTAM, at least for growth on a surface, the technique raises the prospect of examining growth on a surface in general at a level that has not previously been possible. For example, AFM movies could be used to examine the growth of algorithmic tile systems, possibly allowing insight into the mechanisms of both errors and error-reducing methods like proofreading and snaked proofreading.

It's worth noting that if the G_{se} for tile detachment on a mica surface is indeed significantly lower than in solution, as we have found, it may significantly increase error rates for algorithmic tile systems grown on mica surfaces, with the kinetic trapping model [79] suggesting a possible error rate of upwards of 30 to 40% for $G_{se} = 1.6$ and Xgrow simulations suggesting an error rate for the binary counter system from Barish et al [4] too high to be measured. However, for the examination of the mechanisms behind errors this could end up being a benefit, with the higher error rate providing significantly more data on what might otherwise be errors

in a relatively accurate system. As algorithmic systems can often be reliant on perfect, error-free assembly, the ability to examine the systems, especially those that involve proofreading schemes, in an environment that significantly increases error rate could be vital in understanding how to further reduce errors.

Acknowledgment

We thank Christina Wright for prototype software for analyzing AFM movies, and Shaun Lee for discussions regarding AFM movies of DNA tiles. We acknowledge National Science Foundation (NSF) awards 0832824 ("The Molecular Programming Project") and 0622254.

Supporting Information

Several AFM movies of non-equilibrium and equilibrium crystal growth are available as supporting information:

- Supplementary Movie 1: equilibrium conditions with R00/S00. 25 nM each of R00/S00 were annealed separately, mixed in the AFM sample stage buffer during imaging, and imaged at $33.4\,^{\circ}$C; the movie has a $\sim 560 \times 525$ nM field of view for 1088 sec at a frame rate of ~ 73 sec/frame.
- Supplementary Movie 2: equilibrium conditions with NAoMI-B. 50 nM of NAoMI-B was annealed and imaged at $32\,^{\circ}$C; the movie has a $\sim 430 \times 425$ nM field of view for 1012 sec at a frame rate of ~ 126 sec/frame.
- Supplementary Movie 3: growth-biased conditions for R00/S00. 100 nM each of R00 and S00-2J (S00 with added hairpins for labeling purposes creating a striped lattice, also known as $\hat{\mathrm{B}}$ in Winfree et al's Figure 4abc and supplementary materials [81]) were annealed separately, mixed in the AFM sample stage buffer during imaging, and imaged at $32\,^{\circ}$C. The movie has a ~ 2 μM field of view at a frame rate of ~ 80 sec/frame.
- Supplementary Movie 4: growth-biased conditions for NAoMI-B. ~ 200 nM of NAoMI-B was annealed from $90\,^{\circ}$C to $40\,^{\circ}$C at $1\,^{\circ}$C/min, and was then added directly to unheated buffer in the AFM and imaged at room temperature. The movie has a ~ 4 μM field of view at ~ 84 sec/frame. The concentration of tiles was doubled during the two frames exhibiting extensive line artifacts.

For presentation purposes, equilibrium movie image contrast levels were manually adjusted, with black and white being set at the bottom of histogram peaks for surface and tile heights, respectively, and gray being set midway between the two peaks. Images were then manually shifted (but not stretched) to adjust for AFM drift. Growth-biased movies underwent the same process with the addition of manual stretching of images to adjust for AFM distortion.

This material is available free of charge via the Internet at http://pubs.acs.org

Chapter 3

DNA Sticky End Design and Assignment

A major challenge in practical DNA tile self-assembly is the minimization of errors. Using the kinetic Tile Assembly Model, a theoretical model of self-assembly, it has been shown that errors can be reduced through abstract tile set design. In this chapter, we instead investigate the effects of "sticky end" sequence choices in systems using the kinetic model along with the nearest-neighbor model of DNA interactions. We show that both the sticky end sequences present in a system and their positions in the system can significantly affect error rates, and propose algorithms for sequence design and assignment.

3.1 Introduction

Self-assembly of DNA tiles is a promising technique for the assembly of complex nanoscale structures. Assembly of tiles can be programmed by designing short complementary single-stranded DNA "sticky ends." While assembly using unique tile types or simple lattices is often studied [77, 50], algorithmic growth, where small sets with few tile types can form complex assemblies, is particularly powerful theoretically, and has been studied extensively through the abstract Tile Assembly Model (aTAM) [79, 14, 53].

A number of different designs for tile structure are used for assembly [77, 57, 50]. As an example, the DAO-E tile design (Figure 3.1(a)) consists of two helices connected by two crossovers, with four 5 nucleotide (nt) sticky ends, one at each end of each helix. Experimentally, conditions are usually used such that tiles will favorably attach by two bonds between sticky-end regions, adding cooperativity to binding. In the abstract Tile Assembly Model, this is modelled by individual tiles attaching to edges of the current assembly when they can make at least two correct bonds to adjacent tiles ($T = 2$), and never detaching once attached.

The Pascal mod 3 (PM3) system shown in Figure 3.1(b) is a simple example. The tiles implement addition modulo 3, akin to Pascal's triangle. Tiles attach by their two lower-left ends, and then provide ends for future tiles to attach that sum the logical values of the two "input" ends. Growth proceeds to the upper-right, controlled by a V-shaped seed of tiles that attach by strength-2 bonds and provide edges of logical 1s.

Reproduced with permission from Constantine G. Evans and Erik Winfree. "DNA Sticky End Design and Assignment for Robust Algorithmic Self-assembly". In: *DNA 19*. Vol. 8141. LNCS. Springer, 2013, pp. 61–75. DOI: 10.1007/978-3-319-01928-4_5.

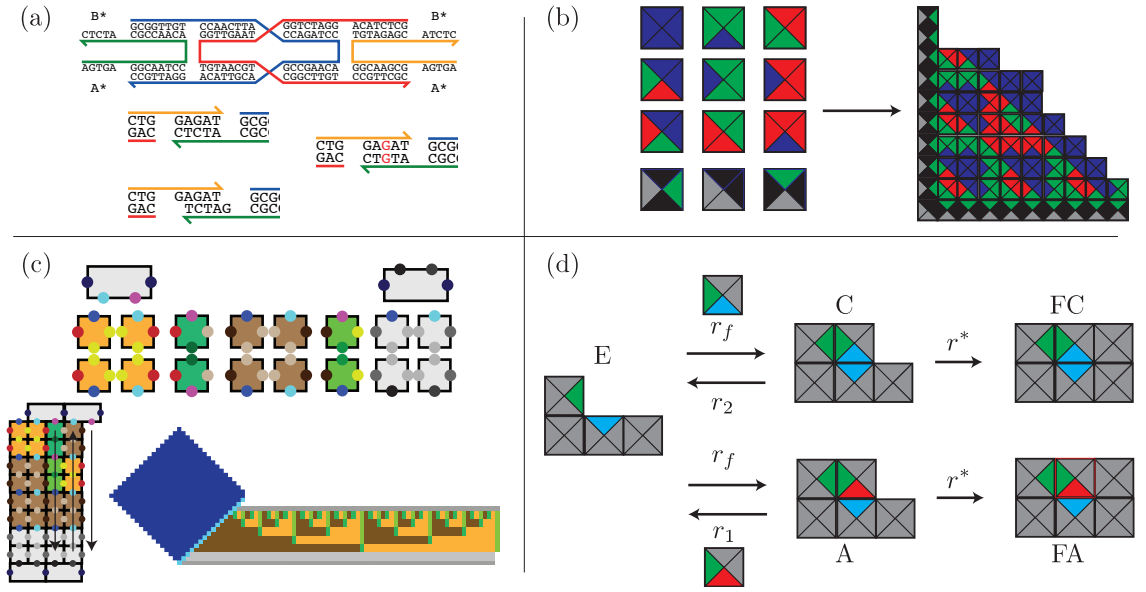

Figure 3.1: Tile systems, structures and the kinetic trapping model. (a) shows an example DAO-E tile structure [57], along with examples of complementary and partially mismatched sticky end attachments. (b) shows the Pascal mod 3 tile system along with a potential perfect assembly. Blue, green and red correspond to ends with logical values 0, 1, and 2, respectively, while black indicates double-strength bonds of the V-shaped seed. (c) shows the tiles (top) in the Barish counter system, along with an illustration of zig-zag ribbon growth (left) and an Xgrow simulation of growth from an origami seed (blue), where each pixel represents one tile. Orange and brown tiles indicate tiles with logical values of 1 and 0, respectively, while gray tiles are boundary and nucleation barrier tiles, and incrementing tiles are green. (d) illustrates the states and transition rates in the kinetic trapping model of growth errors.

A more sophisticated example, the counter system from Barish et al [4], is shown in Figure 3.1(c). In this system, a ribbon of tiles grows from a large seed structure of DNA origami. Rows of tiles grow in a zig-zag fashion, with each new row being started by a double tile that is equivalent to two permanently-attached single tiles. On "downward" rows tiles increment a bit string with two tiles per bit from the previous row, while on "upward" rows, corresponding tiles copy the newly-incremented row. These tiles implement a binary counter starting from whatever bit string was specified on the original origami seed and incrementing every two rows of tiles.

In examining algorithmic growth of experimental systems, the kinetic Tile Assembly Model provides better physical relevance [79]. Tiles are assumed to be in solution at a particular concentration, which is usually assumed to be constant. Tiles attach to empty lattice sites at a rate r_f dependent only on their concentration, and detach at a rate r_b ($b = 1, 2, \ldots$) dependent upon the number of correct "sticky end" attachments they have to the assembly:

$$r_f = \hat{k}e^{-G_{mc}} \qquad r_b = \hat{k}e^{-bG_{se}}\,. \tag{3.1}$$

Here G_{mc} is a dimensionless free energy analogue related to tile concentration by $[c] = e^{-G_{mc}+\alpha}$, G_{se} is the sign-reversed dimensionless free energy of a single bond, b is the number of correct bonds, and $\hat{k}$ is an adjusted forward rate constant $\hat{k} \equiv k_f e^{\alpha}$, where k_f is the usual second-order mass action rate constant for tile

attachment, typically $k_f = 10^6$ /M/s. This model has been used for numerous theoretical and computational simulation studies of algorithmic tile assembly [80, 10, 20, 14, 53], and has fit well with experimental findings both qualitatively and quantitatively [16, 22].

As the kinetic model allows any tile to attach regardless of correctness, it is challenging to design tile systems that exhibit algorithmic behavior while keeping erroneous growth low enough to obtain high yields of correct assemblies. Growth errors in the kinetic model are well studied, and often modelled by the kinetic trapping model. The model considers tiles attaching and detaching at a single lattice location, while having a rate for an attached tile to become "frozen" in place by further growth. This rate, $r^* = \hat{k}e^{-G_{mc}} - \hat{k}e^{-2G_{se}}$, is related to the overall growth rate of the system [79]. As tiles that attach without any correct bonds ("doubly-mismatched" tiles) will detach very quickly, to first approximation, the only states that need to be considered are empty (E), correct tile (C), and "almost correct tile" (A)—a tile that is attached by one correct bond—along with frozen states for correct and almost correct tiles (FC and FA). These states are described in Figure 3.1(d).

Numerous techniques have been studied to reduce such error rates, especially "proofreading" transformations that transform individual tiles into multiple tile blocks or sets of tiles [80, 10, 7, 56]. These techniques have been shown to significantly reduce error rates both in simulation and experimentally [22, 10, 4]. Such techniques rely on changing tile systems at an abstract level, and reduce error rates of even ideal systems. However, in implementing the abstract logic of a tile system in actual DNA tiles, design complexities cause the system's kinetics to deviate from the default kTAM parameters. In particular, the single-stranded "sticky ends" that implement the abstract ends must be chosen from a finite sequence space to be both as uniform in binding energy and as orthogonal as possible. Deviations here can introduce further errors [20].

3.2 Theoretical Model

In the kTAM, G_{se} and G_{mc} are by default considered to be constant and independent of both tiles and sticky ends. A more detailed model cannot assume this. G_{mc} is dependent upon tile concentration: the value may be different for each tile type, and may change as free tiles are depleted by attachment. However, as experimental techniques exist to keep tile concentrations approximately constant throughout assembly [64], we will assume a time-invariant (but possibly tile type dependent) G_{mc}.

G_{se}, on the other hand, will depend upon the bonds between sticky ends. Ends with different sequences will have different free energies for binding to their complements, and some ends may be able to partially bind to ends that are only partially complementary (Figure 3.1(a)). This results in a G_{se}^{ij} for each pair of sticky ends (i, j). In the default kTAM, all non-diagonal terms will be zero, and all diagonal terms will be equal. G_{se}^{ij} can thus be defined in terms of deviations from a reference G_{se}:

$$G_{se}^{ii} = G_{se} + \delta_i \qquad\qquad G_{se}^{ij} = s_{ij} G_{se} \text{ for } i \neq j \,. \tag{3.2}$$

Non-uniform sticky ends, with non-zero δ_i, will affect the detachment rate of correct and almost-correct tile attachments, while spurious non-orthogonal binding strengths s_{ij} will only decrease detachment rates for almost-correct and doubly-mismatched tile attachments. In the following theoretical analysis, the much lower likelihood doubly-mismatched interactions are ignored. For simulations, done with the Xgrow kTAM simulator [72], these interactions are taken into account when there is non-orthogonal binding.

Uniformity

Non-uniform sticky end energies have been simulated previously [20], but have not been studied analytically. In the kTAM, the growth rate of an assembly depends on the difference between on and off rates [79], which we approximate for a uniform system as $r^* = \hat{k}e^{-G_{mc}} - \hat{k}e^{-2G_{se}}$.

For a system with non-uniform energies, a tile attaching by two i bonds will have

$$r^* = \hat{k}e^{-G_{mc}} - \hat{k}e^{-2G_{se}-2\delta_i} = \hat{k}e^{-2G_{se}}\left(e^{\epsilon} - e^{-2\delta_i}\right)$$

where we define $\epsilon \equiv 2G_{se} - G_{mc}$, a measure of supersaturation: for an ideal system, $\epsilon = 0$ results in unbiased growth, whereas $\epsilon > 0$ results in forward growth and $\epsilon < 0$ causes crystals to shrink. As can be seen, the growth rate will depend on the δ_i's of the bonds in the growth region. With $\delta_i < -\frac{1}{2}\epsilon$ (negative δ corresponds to weaker binding), growth in a region won't be favorable.

In the worst case, where tiles attaching by two bonds with the smallest δ_i form a sufficiently large region, growth can only be ensured if $\epsilon > -2 \min\{\delta_i\}$, and error rates can be approximated by the kTAM with this minimum ϵ value. The kinetic trapping model in the default kTAM results in an error rate $P_{error} \approx me^{-G_{se}+\epsilon}$ for m possible incorrect tile attachments [79], so the worst-case error rate for a given $\delta_{\min} \equiv \min\{\delta_i\}$ would be

$$P_{\text{error}} \approx me^{-G_{se}-2\delta_{\min}}. \tag{3.3}$$

Figure 3.2(a) shows simulations of the PM3 system with ϵ adjusted along the lines of our worst-case growth requirements. For positive deviations, where most ends remain at the same strength, assembly time is largely unchanged, while the error rate increases. For negative (weaker bond) deviations, where ϵ is adjusted, the error rate rises per Eq. 3.3, while the assembly time decreases sharply as most tiles attach with the same G_{se}^{ii} but are at a higher concentration.

While this method to adjust tile concentrations ensures crystal growth, it may not obtain the optimal trade-off between growth rate and error rate. This trade-off has been addressed for perfect sticky ends [9, 30], but is more complicated with imperfect sticky ends and complex tile sets. Rather than simply adjusting all concentrations uniformly, the assumption can be made, which is not necessarily optimal, that error rates for a complex tile set can be reduced by ensuring that the overall growth rate remains uniform throughout the crystal. This can be achieved by modifying the concentrations of tiles to modify their G_{mc} values such that the r^* for each tile type is the same. Figure 3.2(b) shows simulations of this form of concentration-adjustment with the PM3 system. As expected, assembly time remains almost completely unchanged across a large range of deviations. Meanwhile, negative deviations do not significantly increase error rates, and positive deviations increase error rates in a manner similar to Figure 3.2(a).

Orthogonality

Unlike non-uniformity, the kinetic trapping model for growth errors can be easily extended to account for non-orthogonality. Assuming $s_{ij} \ll 1$, growth errors will be primarily caused by almost-correct tiles attaching by one correct and one incorrect bond, as in the ideal case. A uniform incorrect bond strength of s, and m possible almost-correct tiles for a given lattice site, then gives the following rates of change between the different states shown in Figure 3.1(d):

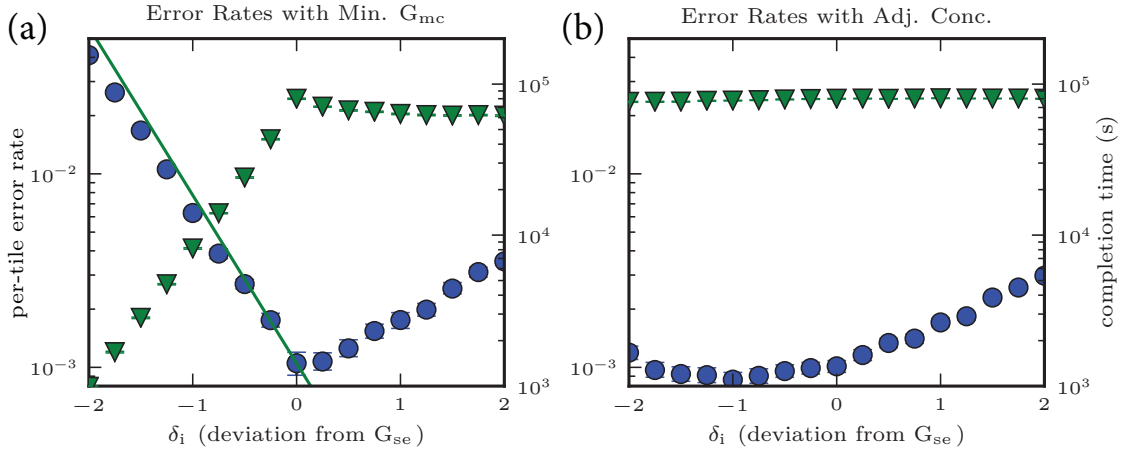

Figure 3.2: Error rates for Pascal mod 3 systems with non-uniform end interactions simulated in Xgrow. In both (a) and (b), single sticky ends have been changed so that $G_{se}^{ii} = G_{se} + \delta_i$, while all others have remained at G_{se}. In (a), the ϵ for the system has been uniformly changed to always allow forward-growth by two of the weakest bond types by setting G_{mc}. In (b), the tiles with deviating ends have had their concentration adjusted so that all tiles have the same growth rate $r^* = \hat{k}e^{-G_{mc}^n} - \hat{k}e^{-G_{se}^{ii}-G_{se}^{jj}}$, where tile type n attaches using sticky end types i and j. Blue circles show error rates; green triangles show the time taken to construct an 8000 tile assembly; the line in (a) shows Eq. 3.3. For these simulations, we set base parameters of $G_{se} = 10$ and $G_{mc} = 19.2$.

$$\dot{P}(t) = \begin{array}{c} \\ E \\ C \\ A \\ FC \\ FA \end{array} \begin{array}{c} \begin{array}{ccccc} E & C & A & FC & FA \end{array} \\ \begin{pmatrix} -2r_f & r_2 & r_{(1+s)} & 0 & 0 \\ r_f & -r_2 - r^* & 0 & 0 & 0 \\ mr_f & 0 & -r_{(1+s)} - r^* & 0 & 0 \\ 0 & r^* & 0 & 0 & 0 \\ 0 & 0 & r^* & 0 & 0 \end{pmatrix} \end{array} P(t) \ . \tag{3.4}$$

Here $P(t)$ is a vector of probabilities at time t that the site will be in a state $[E, C, A, FC, FA]$. The steady state of this is not useful, as any combination of FC and FA will be a steady state. Instead, the eventual probability of being in FA after starting only in state E at $t = 0$ will provide an error rate per additional tile in an assembly. This can be treated as a flow problem, where we consider the differential accumulation into FC and FA from E, as in Winfree [79]. From this, the probability of an almost-correct tile being trapped in place is:

$$P_{error} = \frac{m}{m + \frac{r_f + r_{1+s}}{r_f + r_2}} \approx \frac{1}{1 + \frac{1}{m}e^{(1-s)G_{se} - \epsilon}} \approx me^{(s-1)G_{se}+\epsilon} \ . \tag{3.5}$$

While tile systems will have a different number of possible almost-correct tiles for different lattice sites, making this result less applicable, the PM3 system has an equal number for every possible lattice site. Figure 3.3(a) shows error rates in simulations with interactions between single pairs of ends and for a uniform non-orthogonal interaction energy between every pair. In both cases, error rates largely follow Eq. 3.5.

Sticky End Sensitivity

When non-orthogonal sticky end interactions are not uniform, the degree of their influence on error rates may depend on which tile types they appear on and the logical interactions within the tile set. In systems where a tile never has the opportunity to attach with strength $1 + s_{ij}$, interactions between i and j may be less relevant,

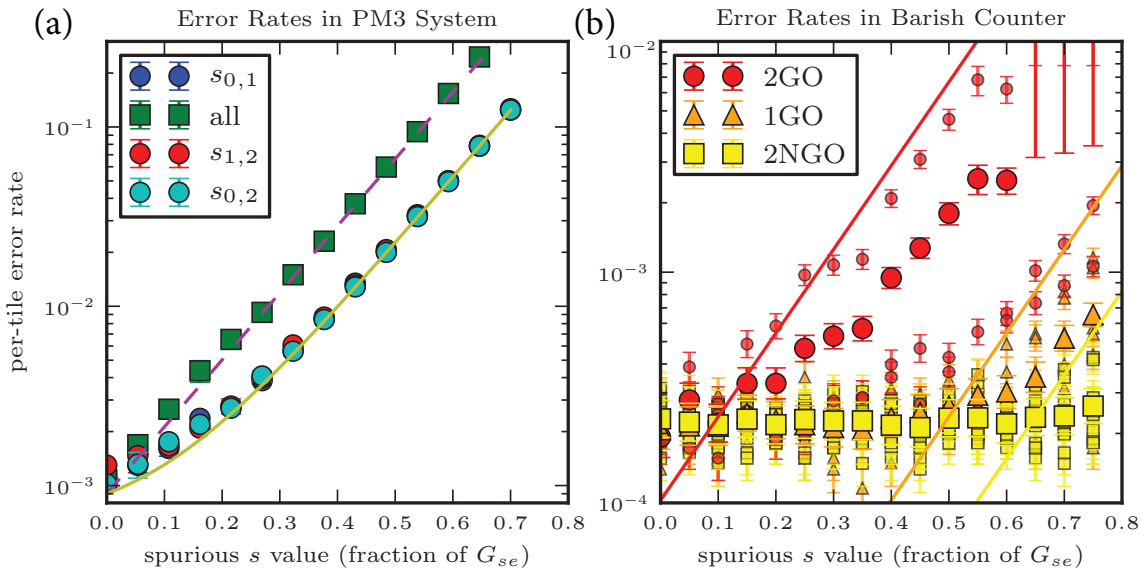

Figure 3.3: Error rates with non-orthogonal interactions. (a) shows interactions for the PM3 system; circles and solid lines show simulated and theoretical error rates, respectively, with single pairs interacting. Squares and dashed lines show error rates for a uniform non-orthogonal interaction between every pair. (b) shows error rates for sensitive single non-orthogonal pairs in the Barish counter system, along with lines showing $e^{-(s-\sigma)G_{se}}$ for various values of σ chosen to roughly follow the worst pairs of each sensitivity. Small dots represent individual pairs, while large dots show averages for sensitivity classes. For (a) $G_{se} = 10$ and $G_{mc} = 19.2$; for (b) $G_{se} = 8.35$ and $G_{mc} = 17.8$.

whereas other pairs of ends in the system may allow tiles to erroneously attach during correct growth and be simply locked in place by continued growth. For example, Figure 3.3(b) shows error rates for the Barish counter system when non-orthogonal interactions are introduced between single pairs of sticky ends. These pairs have been organized into sets (1NGO, 2NGO, 1GO, and 2GO) based on a model described below of how interactions between them may affect the tile system. As can be seen, this model has some success in predicting the impact different pairs will have on error rates.

We start by assuming that all attachments in growth occur with single tiles attaching by exactly two correct strength-1 bonds. Assuming that each tile in the system can have its ends labelled as inputs or outputs, and that every growth site has a unique tile that can attach by inputs, all lattice locations possible in the system will eventually be filled by a specific tile. Rather than looking at lattice sites that actually appear in correct growth, which would require simulation, we can combinatorially investigate all possible local neighborhoods that might appear, and conservatively examine them for possible problems. For example, whether there exists a tile that can attach with strength $1 + s_{ij}$ can be approximated by whether there are two tiles that share a common input bond on one side but not the other, so that when one tile incorrectly attaches where the other could attach correctly, it forms a strength 1 bond for the common bond and a strength s_{ij} bond for the mismatch (as in Figure 3.4(a)).

We describe end pairs where such tiles exist as being in the set of "first-order sensitive" end pairs. If the sides of the tiles are inputs for at least one tile type, and thus the tiles can attach in normal forward growth, the end pair is in the set of first-order growth oriented sensitive (1GO) pairs, whereas without consideration of input

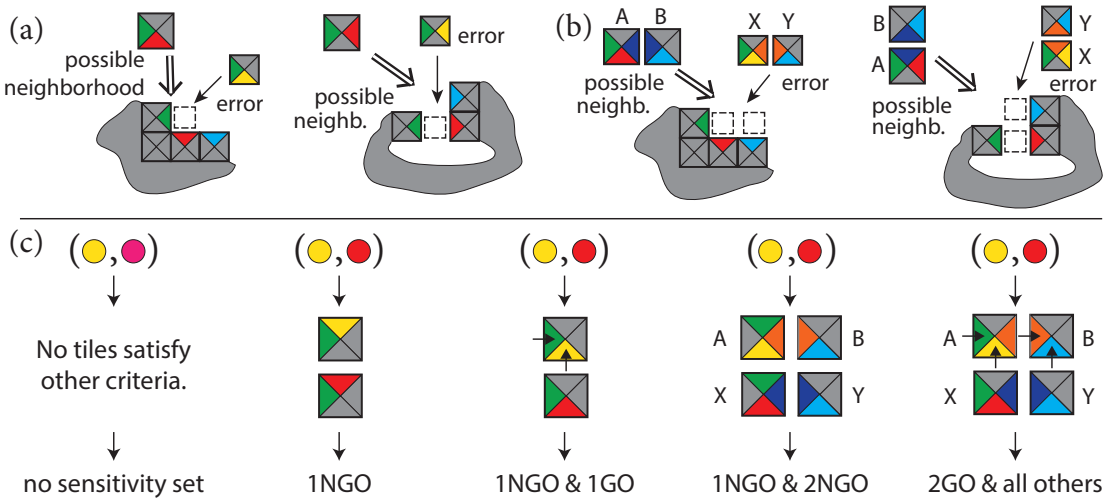

Figure 3.4: Illustration of end pair sensitivity sets. For simplicity, all left and bottom sides are considered inputs. (a) shows, for given tiles, examples of possible local neighborhoods they could attach to and tiles that could erroneously attach via first-order sensitivity. (b) shows, for given pairs of tiles A and B, examples of local neighborhoods the pair could attach to in sequence, and a pair of tiles X and Y that could erroneously attach via second-order sensitivity. (c) shows examples of tiles satisfying various criteria for the shown end pairs to be in different sensitivity sets; arrows show examples of required input sides for growth-oriented sets.

and output sides, the end pair is in the set of first-order non-growth-oriented sensitive (1NGO) pairs. End pairs (i, j) that are in 1NGO but not 1GO have tiles that can attach with strength $1 + s_{ij}$ only during growth after an error or at sites where there is no correct tile.

While end pairs in these sets have tiles that allow the first, erroneous tile attachment in the kinetic trapping model, the model also requires that a second tile be able attach by two correct bonds to the erroneous tile and adjacent tiles to trap the error in place. This is also not necessarily possible: an incorrect attachment could result in there being no adjacent correct attachment, and designing systems where this is the case is in fact the goal of proofreading systems [80].

Thus we can devise "second-order sensitive" sets of end pairs that allow this second, correct tile attachment, and are therefore expected to be more likely to cause errors. Consider a pair of tiles A and X with a common bond on one side but not the other, satisfying the criteria for a first-order sensitive pair. Whether a further tile can attach with strength 2 can be approximated by whether there is some second pair of tiles, B and Y, that can each attach to some third side of their respective original tiles, and also share a common bond on another side. In a plausible local neighborhood where A and B could attach correctly in sequence, it is possible for X to first attach erroneously, with strength $1 + s_{ij}$ (in the location where A could have bound), then for Y to attach with strength 2 (where B could have bound after A) owing to the second common bond, as in Figure 3.4(b).

As with first-order sensitivity, if the common and differing sides of the first pair of tiles are inputs, and sides of the second pair of tiles that are shared or attach to the first pair are also inputs, then the end pair involved is in the set of second-order growth oriented sensitive (2GO) pairs, whereas without consideration of inputs, the pair is in the set of second-order non-growth-oriented sensitive (2NGO) pairs.

These sets can be summarized more formally as follows, while examples of satisfying tiles are shown in

Figure 3.4(c):

- An end pair (i, j) is in the set of first-order sensitive end pairs if there exist at least two tiles in the tile system where both tiles share a common end k on one side, and on some other side, one tile has end i and the other has end j. If at least one of the two tiles has k and either i or j as inputs, then the end pair is in 1GO and 1NGO, otherwise, it is only in 1NGO.
- To determine if a first-order sensitive end pair (i, j) is in the set of second-order sensitive end pairs, consider a pair of tiles that satisfy the first-order criteria, and additional pairs of tiles that can attach to the first pair by bonds l and m (possibly the same) on a third side. If there exist a pair of these additional tiles that also share a common bond n, then the end pair is second-order sensitive. If at least one of the first tiles has k and either i or j as inputs, and one of the additional tiles attaching to it has n and either l or m as an input, then the end pair is in 2GO and 2NGO, otherwise, it is only in 2NGO.

Note that this analysis is done without determining what assemblies and thus what local neighborhoods actually form, so the combinations of inputs being considered might never appear during the growth of a correct assembly. As such, it is conceivable that, for example, an end pair could be in 2GO without ever having an effect in correct growth of an assembly. While this is a significant limitation, determining if a combination of inputs ever occurs, or if two tiles are ever assembled adjacent to each other, is in general undecidable by reduction to the Halting problem [78]. Furthermore, our current software treats all bonds as strength-1, and all tiles as single tiles, with double tiles being represented by a pair of single tiles with a fake bond that is then excluded from the sets; whilst the set definitions could be extended to account for double tiles and strength-2 bonds, we have not yet investigated the complexities involved.

Also, while pairs may be in either or both of 1GO or 2NGO, in all systems we have considered, all pairs in 1GO have also been in 2NGO, and there have been no pairs that are only in 1NGO. End pairs that aren't in *any* of these sets, and can be described as "zeroth-order," should have interactions between them that have a negligible effect on error rates in the kinetic trapping model.

Very rough theoretical estimates of the contributions that sensitive end pairs will have on a system can be obtained by considering the number of tiles that need to attach incorrectly. For pairs in 2GO, as only the initial tile will need to attach incorrectly before it can be locked in place by a correct attachment, the probability of an error every time such a situation occurs is $\sim e^{(s-1)Gse}$. For those in 1GO but not 2GO, since there is no correct attachment after the first tile attaches incorrectly, at least one further incorrect attachment will be required, giving a probability of error $\sim e^{(s-2)Gse}$ or lower. For pairs only in 2NGO or 1NGO, the probability that the first tile can attach incorrectly will depend upon the likelihood that growth is proceeding in an incorrect direction, which in turn will depend upon numerous factors, but will usually require at least one previous incorrect attachment, giving another factor of $\sim e^{-Gse}$ on top of their GO counterparts.

For the Barish counter, there are 342 pairs of ends (helix direction prevents around half the ends from attaching to the other half). Of these, 22 are 2NGO, 9 are both 1GO and 2NGO, and 3 are also 2GO. Figure 3.3(b) shows error rates for increasing values of s_{ij} where one pair has its value increased and all other spurious pairs are left with $s_{ij} = 0$. Each pair has been classified by its "worst" set. As can be seen, 2NGO pairs have little impact on error rates beyond those seen in the ideal kTAM, 1GO pairs start to have an effect after around $s_{ij} > 0.4$, and 2GO pairs are the most sensitive. In the case of the three 2GO pairs in the Barish counter, two cause errors

that prevent correct growth in the next row without an additional error, explaining the significant difference between the most sensitive 2GO pair and the two less sensitive pairs.

3.3 Sequence Design and Assignment

Sequence Design

DNA sequence set design for molecular computation is a widely-studied problem. Different applications necessitate different constraints and approaches: longer sequences with less stringent requirements can be constrained with combinatorial methods like Hamming distance [34], while work on sequences with more stringent requirements have used thermodynamic constraints [73]. However, the basic goal shared throughout most of these algorithms is to find the largest set of DNA sequences that hybridize to their complements significantly better than to any other sequences in the set, or to find a set of a certain size with the best possible "quality"; in this the problem is similar to the maximum independent set problem, which is NP complete [12, 54].

For sticky ends, the sequence lengths required, especially the 5 to 6 nt ends of DAO-E tiles, are shorter, and provide a smaller sequence space, than most other work has considered, with a few exceptions that have largely generated very small sets [73]. Using the end pair sensitivity model, we can reduce errors from non-orthogonal interactions by changing the assignment of sequences to abstract ends, as described later. However, we have no corresponding model to allow us to compensate for non-uniform energies.

The goal for our sequence design, therefore, is to find a requested number of sequences that (a) have non-orthogonal interactions less than a set constraint, and (b) have binding energies (melting temperatures) as uniform as possible given the orthogonality constraints. This contrasts with many sequence design algorithms, where a minimum melting temperature is of primary importance [71], and from algorithms that simply constrain melting temperatures to be within set constraints [73], in that our algorithm chooses a sequence with the *closest* melting temperature at each step.

As the lengths of sticky end sequences are short, complex secondary structure is limited, and thus our algorithm uses an approximation of minimum free energy (MFE) for thermodynamic calculations. Similar to the "h-measure" used in Phan et al [54], the algorithm considers hybridization between two sequences with every possible offset, and uses the nearest-neighbor interaction data from SantaLucia et al [60], including values for symmetric loops, dangles, single-base mismatched pairs, and coaxial stacking with core sequences. Furthermore, for DAO-E tiles, core helix bases adjacent to the sticky ends affect energetics, and need to be designed alongside the sticky end sequences.

Our algorithm works as follows, for length L sticky ends.

1. Generate a set of all possible available sequences A that fit user requirements. With adjacent bases considered, this could be as many as 4^{L+2} sequences.
2. Calculate end-complement binding energies $G_{se}^{ii'}$ for all sequences in A, and (to speed up computation) remove any sequence that falls outside a user-specified range around the median $G_{se}^{ii'}$ of all sequences initially in A, which we call $\overline{G_{se}}$.
3. For each sequence needed:

a) Randomly choose a sequence i from all sequences in A that are closest to $\overline{G_{se}}$, and add this to the set of chosen sequences C.

b) Calculate the G_{se}^{ij} between i and every remaining sequence j in A, and remove all sequences from A with a G_{se}^{ij} greater than a user-specified value.

4. Stop when either A is empty, or a sufficient number of sequences have been generated.

$\overline{G_{se}}$ is chosen as the desired ideal G_{se} in order to ensure a large number of sequences with similar G_{se}^{ii}s will be available; for 5 nt ends, the desired value is $G_{se} \cdot RT = 8.35$ kcal/mol at 37 °C. By adjusting parameters, the maximum number of sequences that can be chosen can be changed, as shown in Table 3.1; running the algorithm repeatedly will also find different numbers of sequences.

Sets chosen by this algorithm are guaranteed to have all ends interact less than a set amount $s_{ij} < s_{\text{desired}}$ with ends other than their complements, and to deviate from the desired correct interaction by less than a set amount $|\delta_i| < \delta_{\text{desired}}$, though when generating sets of a fixed size the largest δ_is will often be much smaller, as the software selects for the smallest δ_i values possible.

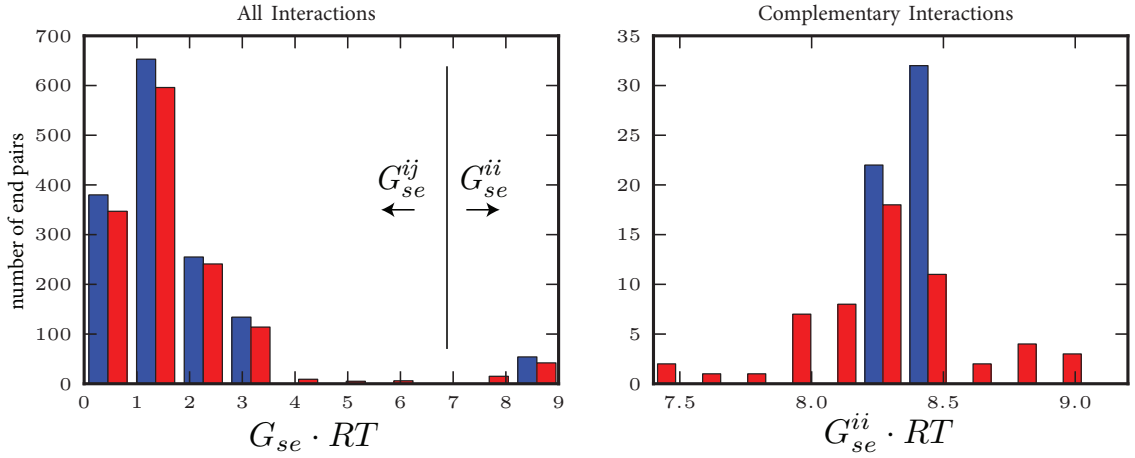

Figure 3.5: Histograms of end pair interactions with the original Barish sequences (red) and newly designed sequences (blue). (a) shows all end pairs; (b) shows a zoomed-in area containing all end-complement pairs. All energies were calculated using the energy model in our sequence designer at 37 °C.

Length (nt)	$\overline{G_{se}} \cdot RT$	$\max(s_{ij})$	# found	std(δ_i)	max δ_i
5	8.354	0.2	5	$0.04G_{se}$	$0.1G_{se}$
5	8.354	0.4	21	$0.01G_{se}$	$0.038G_{se}$
5	8.354	0.5	40	$0.01G_{se}$	$0.036G_{se}$
6	9.818	0.4	29	$0.004G_{se}$	$0.015G_{se}$
10	15.454	0.4	183	$0.01G_{se}$	$0.05G_{se}$

Table 3.1: Examples of the number of sticky ends found by our designer for varying user-specified parameters (bold). For lengths 5 and 6, examples are the best out of 100 runs, while for length 10, the example is a single run.

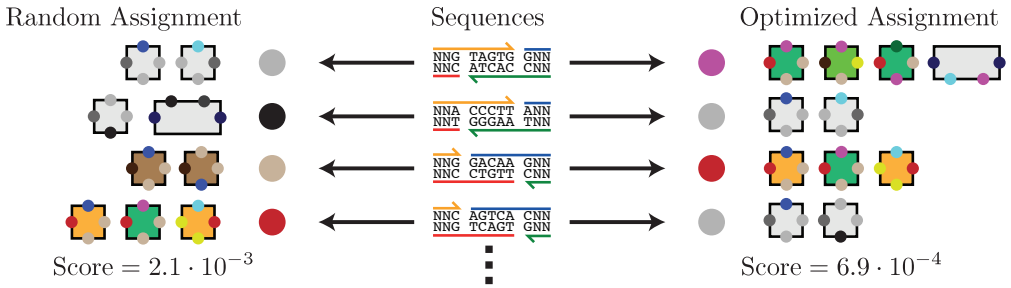

Figure 3.6: Illustration of end assignment for the Barish counter set with new sequences. For conciseness, only a portion of the ends are shown.

Figure 3.5 shows a comparison between end pair interactions in the original Barish counter system and new sequences designed with our sequence design software. As can be seen, our software prevents large non-orthogonal interactions of 4 kcal/mol $< G_{se}^{ij} \cdot RT < 6$ kcal/mol, but does not significantly reduce interactions with $G_{se}^{ij} \cdot RT < 4$ kcal/mol. However, for complementary interactions, our software is able to find a significantly more uniform set of ends.

The practical value of this designer depends on the accuracy of the underlying energy model, of course, but the same algorithm can be used with different energy models as understanding of sticky end energetics is improved. The algorithm, with some energy model modifications, may also be of use in other areas of DNA computation where very short sequences with very similar melting temperatures and low non-orthogonal interactions are needed, such as toehold regions in strand displacement systems. However, it does not consider a number of factors important for actual strand displacement regions, and starts to become computationally intractable for sequences longer than 10 or 11 nt.

Sequence Assignment

The sequence designer is able to find sets of ends with very similar complementary interactions, and low non-orthogonal interactions. However, by ensuring that sequences are assigned to ends in a system such that end pairs with higher sensitivity have lower interactions, errors can further be reduced, and perhaps more importantly, the chance that a poor choice of sequences is made for a critical pair of ends can be minimized.

We assigned ends using a simulated annealing algorithm that used, as a score, the sum of rough error estimates for each end pair (see Figure 3.4):

$$S(\text{assignment}) = \sum_{i,j \in 2GO} e^{-(s_{ij}-1.1)G_{se}} + \sum_{i,j \in 1GO \text{ and } \notin 2GO} e^{-(s_{ij}-1.5)G_{se}} + \sum_{i,j \in 2NGO \text{ and } \notin 1GO} e^{-(s_{ij}-1.65)G_{se}} + \sum_{i,j \in 1NGO \text{ and } \notin 2NGO} e^{-(s_{ij}-2)G_{se}} . \tag{3.6}$$

We call the resulting assignment 'optimized', although of course it is not guaranteed to be a global optimum. Offset values in the exponents were set by rough estimates of the worst errors for different classes in the simulations shown in Figure 3.4, and terms here for 2GO, 1GO and 2NGO are shown by solid lines in that figure. For 1NGO, the -2 parameter is chosen simply to be lower than other classes, as no system we have

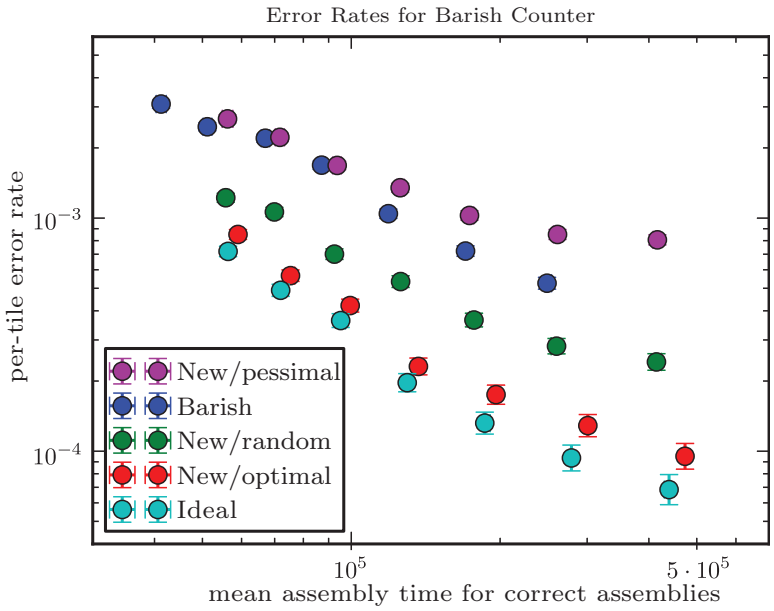

Figure 3.7: Error rates for the Barish counter system with different sticky end sequences. Error rates are calculated from the percentage of correct assemblies formed of size 673. G_{se} values are calculated from ends, or are uniformly $G_{se} \cdot RT = 8.35$ kcal/mol in the ideal case. G_{mc} values were varied between 17.6 and 17.9. 1000 simulations were run for each G_{mc} value.

examined has end pairs that are only 1NGO. Since the sequence designer chooses adjacent bases as well as sticky end sequences, sequences can be consistently assigned to ends on all tiles, as in Figure 3.6. The sequences and tiles for the Barish counter cannot be assigned in the same way, as different tiles with the same sticky end types often have different adjacent base pairs, modifying their interactions. Furthermore, as the sequence assignment algorithm only considers non-orthogonal interactions, results on a system with significant non-uniformity will likely be inconsistent.

Figure 3.7 shows simulated error rates and assembly time for counters using sequences from Barish et al [4], sequences designed by our sequence designer and randomly assigned, and the same designed sequences assigned by our simulated annealing algorithm to both minimize and maximize the score in Eq. 3.6, along with error rates and assembly time for the system under ideal kTAM conditions. For a range of G_{mc} values and resultant assembly times, there is at least a 3-fold improvement in error rate between new sequences that are pessimally and optimally assigned by our scoring function, with increasing improvement as the assembly rate, and thus ideal error rate, decreases. For optimally assigned sequences, error remains close to the ideal error rate. The original sequences and assignment for the Barish counter perform slightly better than the pessimally assigned new sequences.

3.4 Conclusions and Discussion

These methods of sticky end design and assignment serve two purposes: firstly, to design experimental systems with error rates as close to the ideal kTAM as possible, and secondly, to reduce the chance that a poor choice of sequences, or even a poor assignment of sequences to tiles, might significantly impact experimental results. The

methods should be relevant for most types of DNA tiles, and most tile systems with deterministic algorithmic behavior. Our software for these algorithms is available online [69].

The simulation results here, and the methods themselves, are reliant on the accuracy of the energy model used. While some research has been done on sticky-end energetics [46, 55, 36, 16], usually for individual pairs of tiles, it is not known how well nearest-neighbor models of DNA energetics apply to sticky ends on DNA tiles in lattices. Different tile structures may also require slightly different models, especially with regard to coaxial stacking with base pairs adjacent to the sticky ends.

It is possible that extending end sensitivity definitions to higher orders, considering more than two tile attachments, may be a useful area of investigation, especially when considering tile systems making use of similarly higher order proofreading. Indeed, proofreading can counteract at a more fundamental level some of the same errors that arise from non-orthogonal interactions. The effects of non-uniform sticky end energies, however, may still significantly impact proofreading sets, and remain a potentially fruitful area of research beyond our simplistic modeling and concentration adjustment technique.

Chapter 4

Design Considerations for Self-Assembling Binary Counters

4.1 Introduction

Algorithmic self-assembly has clear and extensive theoretical promise as a method of constructing nanoscale structures. Most theoretical work, however, has been focused on its computational power, or on general considerations in the abstract and kinetic Tile Assembly Models (aTAM and kTAM), rather than the design and implementation of specific systems. Similarly, although some experiments have been done on full algorithmic systems on their own, experiments and simulations have often been meant primarily to analyze specific theoretical ideas. Such general approaches are useful in providing widely applicable results that may be applied to specific systems. They have also allowed for research into the general behavior of complex algorithmic systems before experimental implementations of similar complexity have been feasible.

In this and the next chapter, however, we instead consider the design, theoretical details, simulation and experimental implementation of a specific algorithmic tile system: a ribbon that implements a binary counter and can halt further growth, or "cap," when the counter overflows. As the starting value of the counter can be changed by using a different seed, the same system can be used to construct ribbons of differing lengths. The system thus approaches one of the theoretical ideals of algorithmic self-assembly: a simple algorithmic system that can programmably construct a variety of finite structures based on differing inputs.

Rather than considering a single aspect of general tile assembly systems, our focus is on multiple aspects of a single system. By integrating together theory, simulation and experiments, we will try both to move toward a working experimental implementation of a specific system, and to derive more general and perhaps unexpected insights.

In this chapter, we consider several issues in the design of a capping, binary counting ribbon that assembles accurately and stops growing reliably. This will lead to an examination of how varying tile concentrations affect growth errors and facet nucleation, and a short analysis of challenges in constructing stable finite shapes in

the kTAM. Considerations of facet nucleation, capping, and growth order will also motivate a second system, the "spurious counter," that relies on specific features of the kTAM to allow growth in ways that would not be possible in the aTAM.

4.2 BINARY COUNTING

The capping binary counter is a useful simple example of algorithmic self-assembly: at the most abstract level, the logic of a binary counter in the aTAM can be implemented in only four tiles [58]. As mechanisms for computing length and addressing rows in larger assemblies, counters are widely employed in theoretical constructions [58, 2, 24]. Even by itself, a binary counter system of fixed width, with a programmable starting value and growth that stops when the counter overflows, would allow a simple illustration of algorithmic self-assembly's computational power: a single tile system that could be used, with different starting values, to create structures of precise, programmable length.

As seen in Figure 4.1(a), the basic logic of a binary counter consists of four elements: in each (vertical) row, "zero" and "one" tiles copy bits from the previous row when they attach via a "copy" sticky end, while corresponding zero and one "increment" tiles, attaching to an "increment" sticky end, perform the incrementing of the bit string in each row. Starting from the boundary at the low-bit (top in diagrams) side of the bit string, each one-valued tile from the previous row results in the attachment of a zero-increment tile, flipping the bit and carrying the increment sticky end to the next location in the row, until a zero-valued tile in the previous row is reached. At this point, a one-increment tile attaches, and further locations in the row are copied rather than incremented. To create a longer fixed-width ribbon, we simply need an initial seed structure that specifies our starting values (and thus dictates ribbon width as well), and some form of boundary tiles, perhaps assembled by double bonds, to initiate new rows. A form of this basic binary counter was implemented in the earlier experiments of Barish et al [3].

Such a system, however, will continue counting so long as as new rows can be initiated: after the counter overflows, counting will continue from all zeros. In fact, it isn't necessary for each row to complete before the next begins. This poses a problem for stopping growth. An overflow occurs when all bits in a row are 1, and thus all tiles in the next row are zero-increment tiles, with the increment sticky end reaching the high-bit side of the ribbon. Detecting an overflow must therefore take place at the high-bit side, and in order to stop growth, this detection must take place before the next row has a chance to grow. A mechanism is thus needed for preventing each row from growing until the first has completed.

Since each row is initiated by a boundary on the low-bit side, controlling the growth of this boundary controls row initiation. The zig-zag ribbon mechanism, also useful for controlling spurious crystal nucleation [62, 61], presents one possibility. As described in Section 1.4, zig-zag ribbons use "double-tile" boundaries on each side that reverse the direction of growth and initiate the next row. While counting cannot be performed starting from the high-bit side, the system can be modified such that rows growing from high to low only copy the bits from the previous row. Incrementing only takes place on rows that grow from low to high, but each row cannot start growth correctly until the previous has formed and the appropriate boundary has attached. Detecting and stopping growth after an overflow, in the aTAM, is simple: by not having a high-bit boundary that attaches to an increment sticky end, an overflow will result in no correct attachment being possible, so growth will stop. In the kTAM, facet nucleation and other errors from insufficient attachment will pose complications for finite growth, which will be discussed in Section 4.4.

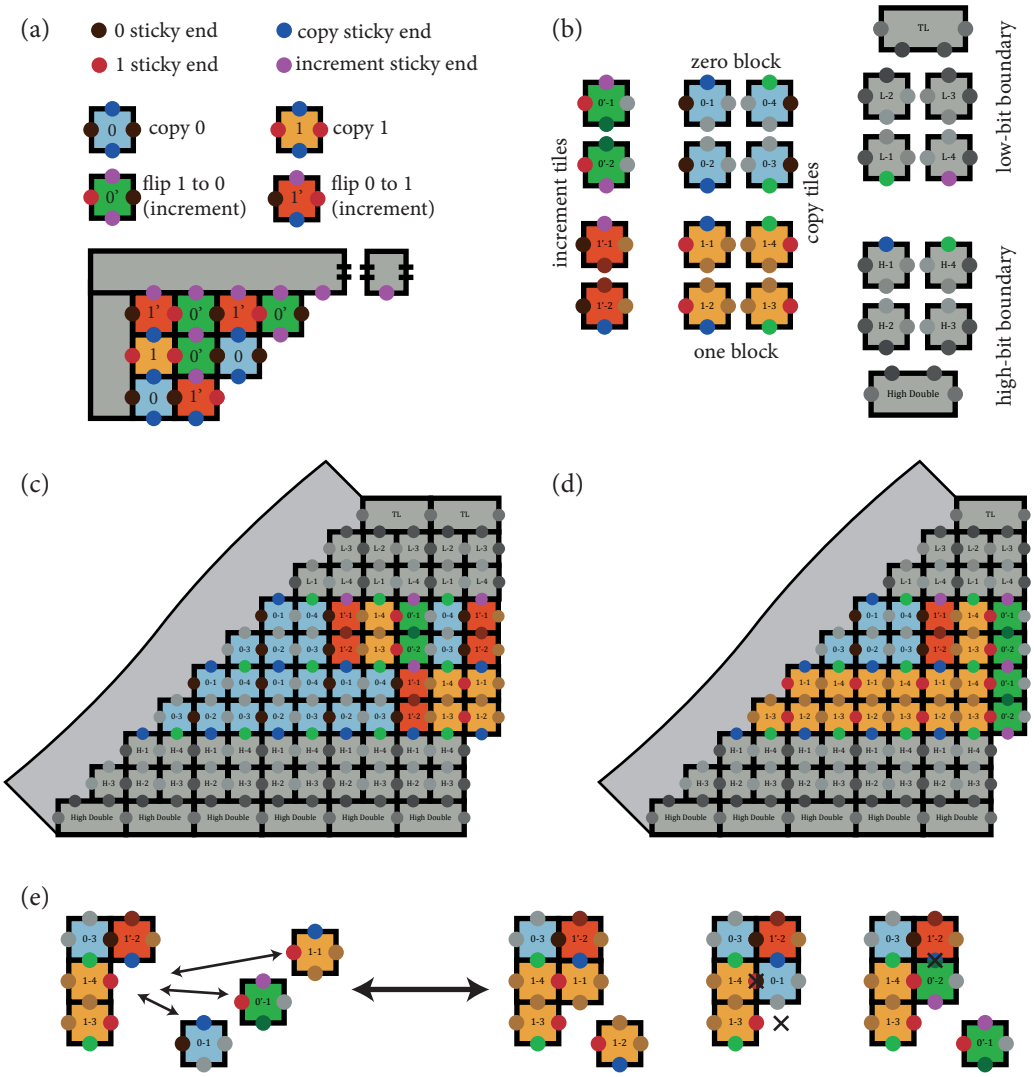

Figure 4.1: Binary counter design and logic. (a) shows the basic logic of binary counting with DNA tiles: four tiles, two for each value, are used, with one copying the bit from the previous row, while the other, starting from an "increment" sticky end on top, flips bits from the previous row until a zero is reached, incrementing the value of the row by one. Growth proceeds in this example from left to right, along a boundary of increment sticky ends. (b) shows the zig-zag counting tile system used in this and the next chapter. Zero and one tiles have been changed into 2×2 blocks; each bond type within blocks is unique (except for incrementing tiles attaching to copying blocks), but are shown as they same color. Growth as a set-width ribbon is achieved by double tiles on edges that initiate each row in a zig-zag fashion, counting in the top-to-bottom (low-to-high value) direction and copying in the bottom-to-top direction, while nucleation barrier tiles (L and H) reduce spurious nucleation in solution. The system grows off origami (as in (c)), which also sets the counter size and starting value, until reaching an overflow state where no correct tile can attach (d). While the system proofreads copying in a 2×1 fashion, incrementing versus copying tiles are not proofread, as shown in (e). Note that two-bit counters are for illustration purposes; in simulations and experiments, five-bit counters are used.

When designing this system for experimental implementation, it is useful to incorporate proofreading and control of spurious nucleation [61, 80]. As shown in Figure 4.1(b-d) (for very narrow ribbons), nucleation control is implemented by adding two boundary blocks of width 2. Full 2×2 proofreading would require that each tile be replaced by a 2×2 block, making each row into two rows, but the zig-zag system has *already* necessitated doubling the number of rows, essentially turning each tile into a 1×2 block. To minimize increased complexity, proofreading is instead partially incorporated, with each tile being transformed into a 2×1 block. The result, shown in Figure 4.1(d), is that the value of each bit in the previous row is proofread, but the increment or copy state of the row is not.

When combined with origami seeding [4], this binary counter design represents a complete system for constructing programmable-length ribbons in the aTAM. In simulations in the kTAM, it will count with reasonable accuracy. Experimentally, sequences for the tile structures and sticky ends can be designed, and counting ribbons can be grown: the system as shown in Figure 4.1(b-c), with a five-bit seed, is the system used for much of the counting results in Chapter 5. The system is based on a combination of the similar binary counting ribbon in Barish et al [4] with the tile and nucleation barrier design of the copying ribbon in Schulman et al [64], both successful experimental systems.

When moving from the aTAM to the kTAM or experiments, however, errors become important factors in growth that need to be considered. Parts of the system that have particularly high potential for errors must be examined; in particular, since no *correct* attachment can occur when the counter has overflowed, but there are *incorrect* attachments that can, stopping growth is a significant matter that will be discussed in Section 4.4. First, however, since there are more choices in system design in the kTAM—particularly, in the concentrations of tiles—we will examine how these may affect error rates.

4.3 EFFECTS OF TILE CONCENTRATIONS ON ERROR RATES

Experimental results in tile assembly have tended to result in higher error rates than theoretical models. This is not surprising: some sources of error, like lattice defects and defective tiles, are not modelled in the kTAM, and experimental uncertainties in temperature control, strand quality and pipetting mean that there will be some variation in growth temperature and in tile concentrations. This last point is of particular interest: in an experimental sample, the concentration of each tile may vary, thus changing the G^t_{mc} and equilibrium for different tiles t. In Chapter 3, similar variation in sticky end binding energies that changed the G_{se} for each pair of sticky ends was found to have significant effects on growth in simulation. It would not seem unusual for concentration variations to result in similar effects. While variations in energetics are a matter of design, concentration variations are a matter of experimental technique; knowing how significantly variations impact growth rates would give insight into how robust systems can be to experimental variations.

From another perspective, intentionally changing the concentrations of different tiles provides tile-level control of growth rates in different parts of the system. This has been used, in theory and simulation, to compensate for variations in energetics [17]. Experiments, meanwhile, have used intuitive concentration changes for various reasons, especially as a mechanism of reducing spurious nucleation by lowering concentrations of tiles that can easily form critical nuclei, and increasing and reducing the growth speed of different system components [4]. Examining how error rates are affected by concentration variations could also provide insight into optimal concentration choices, and how concentration changes to address experimental details might impact growth.

In general, having unequal tile concentrations moves the system away from the basic kTAM, where all tile types attach to lattice sites at the same rate, and instead results in each tile type t attaching at a different rate $r_f^t = k_f[c_t] = \hat{k}e^{G_{mc}^t}$. This will in turn result in a tile-specific G_{mc}^t for the tile. With large enough concentration changes, G_{mc}^t could be changed enough that growth behavior could be substantially changed: tiles with very high concentrations might be able to favorably attach by only one bond and would attach and detach very rapidly, while tiles with very low concentrations might require three bonds and would attach and detach at a vastly slower time scale, despite all bonds being the same strength. Such behavior could exhibit interesting features, but this has not been explored. We can model small changes in concentration, however, as perturbations of G_{mc}^t away from some desired G_{mc} that don't take us far away from equilibrium, so that growth by two bond attachments is still slightly favorable, and the primary change is the attachment rate of different tiles. The effect this has on error rates can be split into two parts: the change in growth errors of correct and incorrect tiles to a valid growth site, and the changes in facet nucleation that can cause backward growth.

Concentration and Growth Errors

To analyze growth rate effects, we'll consider a simple situation of two tiles. Each tile is able to attach correctly, by two bonds, to a particular growth site, and is able to attach almost-correctly, by one bond, to the correct growth site for the other tile. We'll assume that the first tile, tile A, is at an unmodified concentration with $G_{mc}^A = G_{mc} = 2G_{se} - \epsilon$, while for tile B we will modify the concentration so that $G_{mc}^B = G_{mc} - \delta$, where $\delta \equiv \ln c_B/c_A$ is positive for a raised concentration, and negative for a lowered concentration.

While error rates for each growth site could be determined in the kinetic trapping model, the energetics-based approximation can quickly provide results. For the correct site for tile A, attachment of A would result in a total free energy change of $\Delta G = G_{mc} - 2G_{se} = -\epsilon$, while incorrect attachment of B would result in $\Delta G = G_{mc} - \delta - G_{se} = G_{se} - \epsilon - \delta$; the difference provides our approximate error rate $P_{\text{err}}^A = e^{-G_{se}+\delta}$. The same process for A incorrectly attaching in the correct site for B results in $P_{\text{err}}^B = e^{-G_{se}-\delta}$.

Intuitively, these error rates mean that, when the concentration of a tile B is lowered, error rates for that tile attaching incorrectly elsewhere decrease, as attachment rates are lowered, while error rates for other tiles attaching to sites where B should attach are increased, as their attachment rates are comparatively higher. When an entire system is considered, these two effects both change the error rate. If a system has M sites where A should attach, and N sites where B should attach, then the probability of all A and B tiles attaching perfectly can be approximated as

$$P_{\text{perfect}} = \left(1 - e^{-G_{se}+\delta}\right)^M \left(1 - e^{-G_{se}-\delta}\right)^N \tag{4.1}$$

When δ is varied, this probability will change, as shown in Figure 4.2. As error rates will depend upon the number of locations of each type, M and N affect the optimal δ and error rates; intuitively, having a lower concentration of the tile that should appear less frequently will reduce overall error rates, despite raising them for its correspondingly less frequent attachment sites. The effect of any concentration variations on growth error rates is limited, however, by two factors: firstly, as δ is only logarithmically dependent upon concentration, concentration changes don't show the exponential error increase seen with energetics changes. Secondly, since changing δ increases error rates for one lattice site while reducing it for another, these two changes cancel out each other to some extent, and it is only the difference in the two effects that remains.

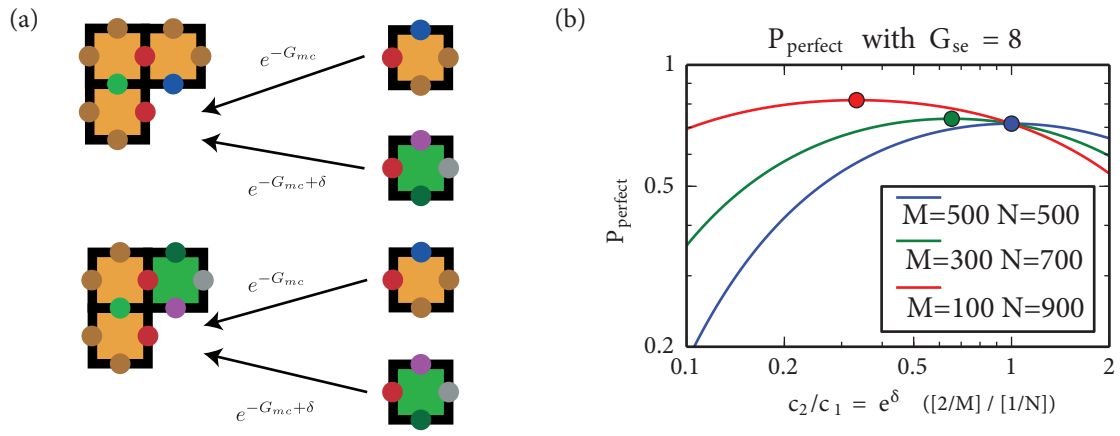

Figure 4.2: Growth errors with concentration variance. Considering two tiles, with one (green) at a varied concentration, (a) shows how each tile can attach by one bond in the location where the other would correctly attach, with two different rates. (b) shows the theoretical probability of having no errors occur in a system that should have M (first number) yellow tile locations and N (second number) green tile locations, given a varying ratio of concentrations.

Thus, as seen, the probability of perfect attachments changes relatively slowly despite significant changes in concentration; even when tiles are used equally in a correct assembly, and one tile has a concentration that is only 20% that of the other, yield still remains more than half the optimal yield. Small variations in concentration have little effect, and while concentrations can be optimized for assemblies with different numbers of different tile types by finding a δ that maximizes P_{perfect}, the increase in yield over 1000 tiles is relatively small, in turn implying a very small change in error rates. Jang et al also considered the effects of choosing optimized concentrations, and found only moderate improvements in simulated error rates [29], while Chen and Kao observed only slightly more substantial improvements in simulations despite using tile systems with significant differences in tile frequency [9].

While this result suggests that clever choices of concentration cannot significantly reduce rates of growth errors, it is actually considerably reassuring. From an experimental protocol perspective, the result suggests that experimental variability in tile concentrations will not severely impact growth error rates, and that precise control of concentrations may not be vital for accurate growth. From a design perspective, the result gives some freedom to use modified concentrations to address other considerations, without being concerned about severe impacts on growth error rates.

Concentration, Facet Nucleation and Growth Order

Since the growth rates of different parts of a system will be affected by the concentrations of the tiles involved, and thus the rates of facet nucleation are impacted, it is reasonable to think that even if concentration variations do not significantly affect growth errors, they may affect facet nucleation and growth order, potentially resulting in significant sources of errors. The effects of facet nucleation on actual error rates are difficult to determine in general: different assemblies will have varying numbers of sites where facet nucleation can occur as they grow, and even if facet nucleation does occur, with subsequent backward growth, there is the possibility that the correct tiles will still attach, just not in the correct order.

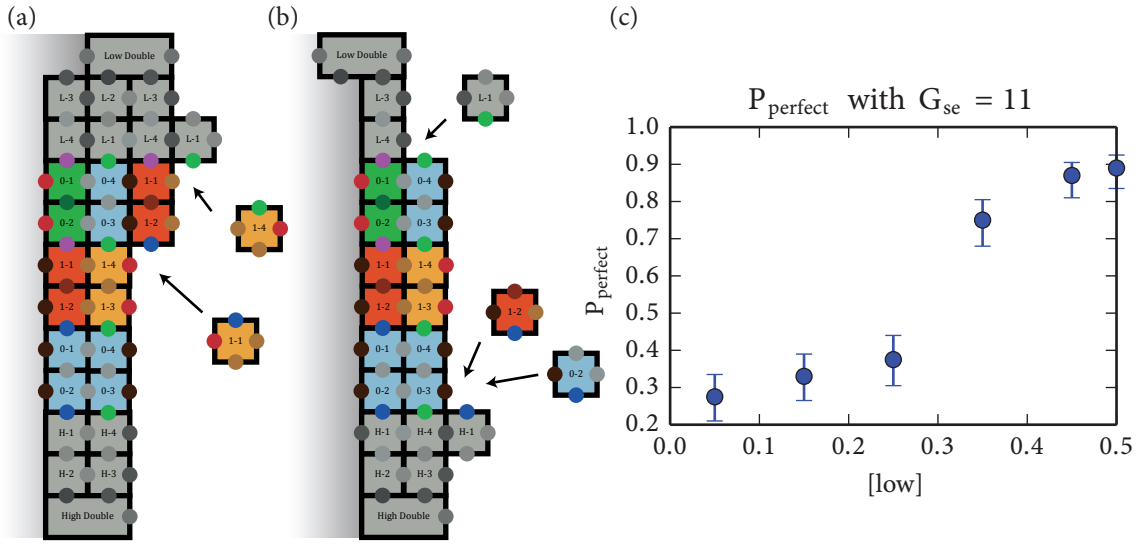

Figure 4.3: Backwards growth from facet nucleation in two different directions. If facet nucleation occurs on the low-value boundary (a), the copying row still attaches correctly despite growing backward. If it instead occurs on the high-value boundary (b), growth of the counting row is not deterministic and may be incorrect. In simulations (200 per point) of the system with [high] = 0.25[middle] and $G_{mc} = 21$, (c) shows the probability of constructing an error-free counter of approximately 30 increment rows.

For the specific case of the binary counter system, differences in sequence energetics initially motivated analysis of how changing the relative concentrations of the nucleation barrier tiles, already at lowered concentrations to reduce spurious crystal nucleation, might improve yield. In qualitative laboratory experiments, increasing the concentration of the "low-bit" barrier over the "high-bit" barrier tiles resulted in significantly better-formed ribbons, and simulations showed significantly lower error rates. In fact, simulations showed lower error rates even without considering sequence energetics, suggesting that some other, more fundamental phenomenon was involved.

The reason for these results lies in the logic of the binary counter, and in facet nucleation that does not result in errors. For both nucleation barriers, facet nucleation of further nucleation barrier tiles will not immediately result in an error: all sticky ends in both blocks are unique, and so the only tiles that can attach to sites by one bond are the same tiles that would attach correctly in the same locations. Once attached, however, growth of a row in the ribbon can proceed out of order, with either two sides of the same row growing competitively, or two rows growing simultaneously in the same direction as shown in Figure 4.3(a-b).

In zig-zag binary counter growth, the actual incrementing of the binary counter takes place in a low-bit side to high-bit side direction; growth in the opposite direction is non-deterministic, as incrementing or copying tiles can attach to the same locations (Figure 4.3(b)). Without further facet nucleation or growth errors, on the other hand, the *copying* rows, despite normally attaching from the high-bit side to low-bit side, will attach correctly regardless of growth direction. Thus facet nucleation on the low-bit side will be unlikely to result in logical errors in growth, while facet nucleation on the high-bit side will have a high probability of causing errors.

Since at equal concentrations, each boundary will be present for approximately equal time in a correctly assembling system, and facet nucleation is equally likely on either side, the relative rates of facet nucleation will

depend upon the relative concentrations of the high and low boundaries. Furthermore, if backward growth of a row has already started, it is less likely that facet nucleation on the other side will have a chance to cause incorrect attachments. Thus, since facet nucleation from the low-bit side won't result in counting errors, a higher concentration of the low-bit side tiles would be expected to result in lower actual error rates and more accurate counting.

Figure 4.3(c) shows simulations with varying concentrations of the low boundary tiles that illustrate this. The concentration of high boundary tiles is 25% that of normal counting tiles in all simulations. With concentrations of low boundary tiles below 25%, the probability of assembling a correct counter of approximately 20 rows is relatively low, while above 25%, the probability sharply increases.

This optimization, however, works by making more likely an "error" (facet nucleation of the low-bit boundary) in growth order that doesn't result in actual errors being incorporated into the assembly. Increasingly higher concentrations of the low boundary make this backward growth increasingly likely, and move away from the correct, zig-zag growth of the system. Since counting doesn't depend upon correct zig-zag growth, high concentrations of the low-boundary can significantly lower error rates. If the system were modified in some way that *did* depend upon correct zig-zag growth, then this optimization would be limited.

4.4 Finite Assemblies and Capping

One process that requires correct zig-zag growth of the counter, unfortunately, is stopping growth to produce finite assemblies. Overflows take place on the high-bit side of the counter, and must be detected there. This detection must have a global effect, stopping growth throughout the system, and thus the order of growth does need to be controlled so that facet nucleation on the low-bit side doesn't take place before growth is stopped.

In the aTAM, finite assemblies arise naturally out of the rules behind tile attachment. When an assembly reaches a state where no further correct attachments can take place, growth simply stops. In the case of the counter system, for example, growth continues until an overflow occurs, proceeding to the high-bit boundary with an increment sticky end instead of a copy sticky end. Since there is no high-bit boundary tile that can attach to an increment sticky end, growth can't continue, and the counter stops.

From a kTAM perspective, however, this assembly is highly prone to further incorrect "cancerous" growth. A growth site that doesn't allow any correct attachment in the aTAM can still be filled with an insufficiently attaching tile; as there is no correct attachment at the site, this will take place at the same rate as facet nucleation. Normal facet nucleation can also take place along the entire edge of the final row. Since there is no further correct growth, the probability that the assembly will remain correct goes to zero over time: if an insufficient attachment can take place, it eventually will.

If there are *no* locations where a tile can attach by any strength, however, such facet nucleation won't be able to take place. This motivates capping: structures with null bonds (experimentally implemented as hairpin ends) along exposed sides that "cap" an assembled system by covering all the remaining sticky ends that could allow for cancerous facet nucleation. A simple design can be seen in Figure 4.5(a), along with traces of assembly size versus time are shown for 20 simulations. These traces show a problem with this design: while some counters will never have the cap attach, or will have errors that will result in incorrect lengths, even counters with caps that do attach correctly, and remain at the correct size for some time, will eventually start growing again. At the end of the simulations shown, only 3 of the 20 assemblies remain capped.

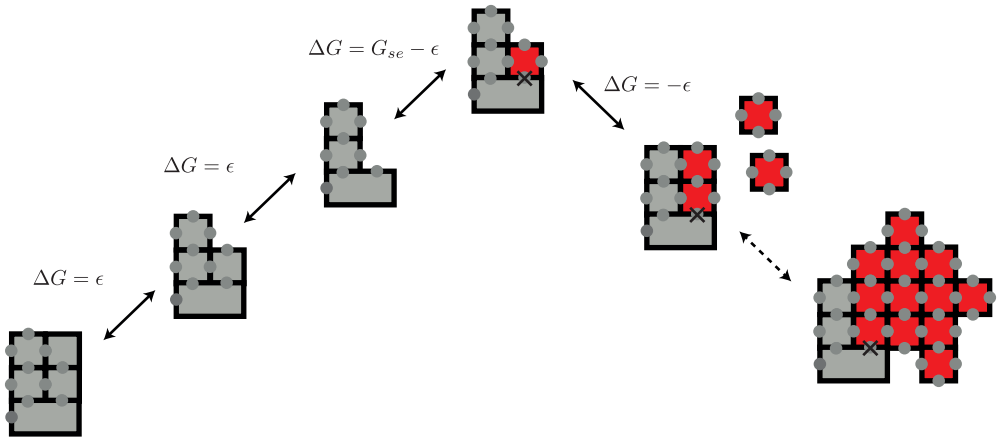

Figure 4.4: Energy diagram of cap failure. Each cap tile detachment, for tiles attached by two bonds, results in a free energy increase of ϵ. In the worst case, once the cap is detached, a single error, with a free energy increase of $G_{se} - \epsilon$, can allow unbounded favorable growth (with free energy changes of $-\epsilon$ per step).

Watching simulations, the reason for this failure is clear, and is related to the fundamentals of the kTAM. While further growth cannot take place when the cap is attached, the detachment of tiles, and the resulting random walk or "breathing" of growth, especially for tiles attached by two bonds, means that the cap will not always be attached. Once a facet nucleation event takes place while the cap is partially or fully detached, cancerous growth can start, resulting in capping failure. Intuitively, no matter how large or stable a cap, there will thus always be some rate for the cap to fail: even if the entire assembly must detach for an error to take place in the initial seeding.

Whether this challenge in constructing stable finite algorithmic assemblies is fundamental to the kTAM is unclear. *Uniquely-addressed* tile systems have no such problem: there is no rate for cancerous growth, as once an assembly has completed, detachments can only be filled with correct tiles. Beyond the kTAM, experimental works have also demonstrated uniquely-addressed finite assemblies [33, 77]. It is an open question as to whether an algorithmic system could be designed to have similar properties, or whether algorithmic behavior necessarily results in the potential for erroneous, cancerous growth.

A more reasonable goal in designing a system to construct finite assemblies, however, is to have a cap with a failure rate low enough that it will, on average, survive for significantly longer than the average time for the structure to assemble. This is similar to spurious nucleation: whereas for spurious nucleation we are concerned with the rate of formation for assemblies that are critical nuclei, for cap failure, we are concerned with the rate that critical nuclei for cancerous growth form from some series of attachments to and detachments from the final assembly, which acts as a type of seed.

As with spurious crystal nucleation, this failure rate can be related to the favorability of failure pathways. One example is shown in Figure 4.4. A fully assembled cap, with each tile attached by two bonds, can disassemble (with $G_{mc} = G_{se} - \epsilon$) by unfavorable steps of free energy change $\Delta G = \epsilon$. Once fully detached, some error must take place for cancerous growth to occur, making the pathway more unfavorable, but then cancerous growth can occur by favorable steps. In many cases, since the detachment steps become more favorable closer to equilibrium, the energy barrier for cap failure can *decrease* closer to equilibrium, where ϵ is smaller, in contrast

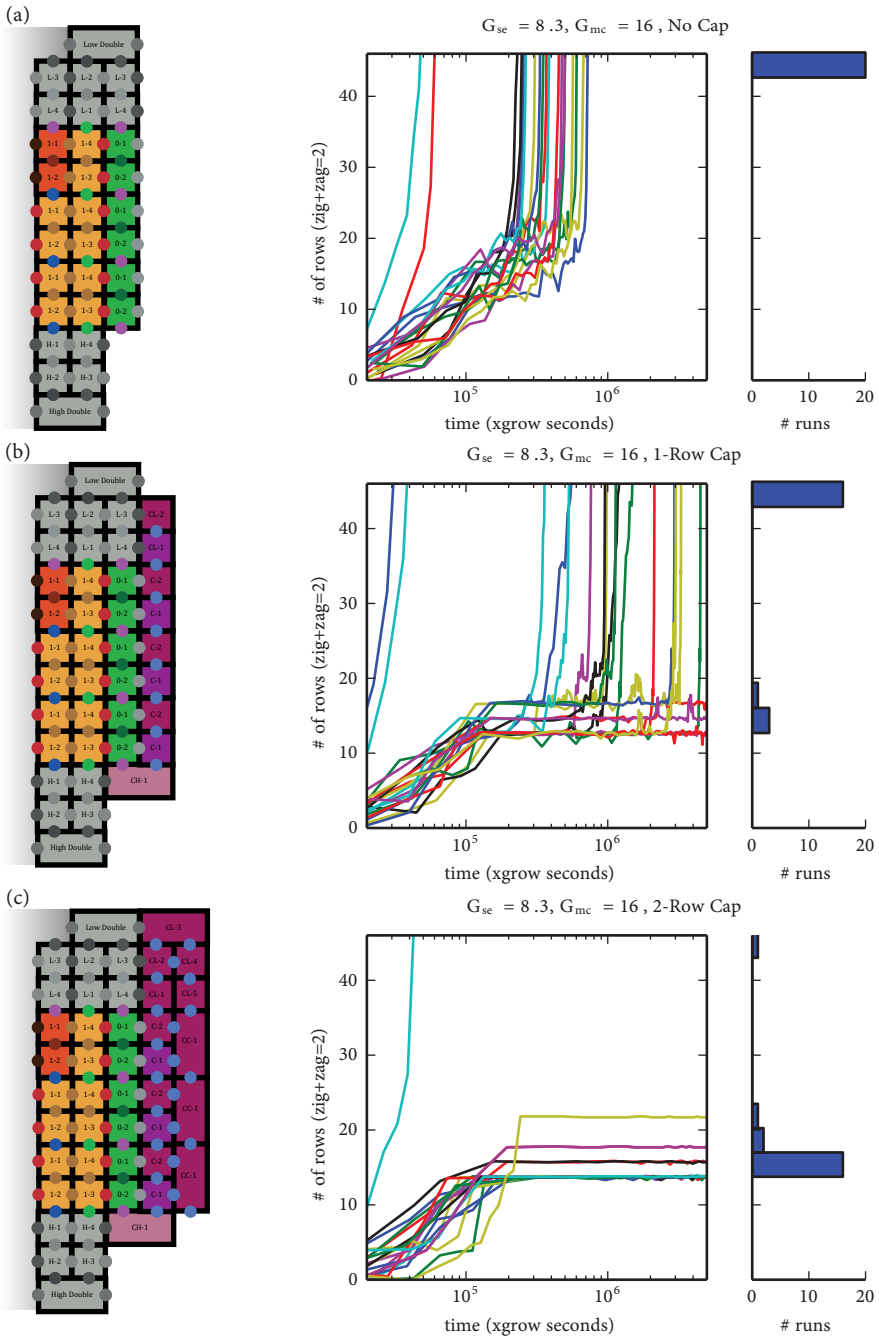

Figure 4.5: Two cap designs. (a) shows the counter with no cap. As seen in traces of counter size over time in 20 simulations that overflow at 14 rows, growth does not stop despite the lack of any aTAM-correct attachments. (b) shows a cap design that covers most sticky ends, not allowing facet nucleation that could result in cancerous growth. However, as facet nucleation can occur when the cap is partially or fully detached, the cap often fails in simulation, either before attaching or after being attached for some time. (c) shows a two-row cap design, where the second row consists of double tiles that attach by three correct bonds. While this cap does not always attach correctly, and errors can also occur in counting itself, the cap largely stays attached once attached.

to growth errors, facet nucleation, and spurious nucleation.

There are thus two ways to decrease the rate of cap failure: the barrier to cap detachment can be increased, or the barrier to cancerous growth can be increased. Proofreading could be used to reduce facet nucleation errors and growth errors that could result in cancerous growth, but it would have to be used for *all* locations where such errors could occur. Adding more tiles to the cap could help, but as each tile attaching by two bonds would add ϵ to the barrier, the effect would still diminish when approaching equilibrium. Individual tile detachment steps could be made less favorable by having tiles attach by three bonds rather than two, but this would then make erroneous, premature cap attachment problematic, as tiles would be very unlikely to fall off once attached.

One compromise is to use a cap that initially attaches weakly, allowing correction of premature attachment, and then has a series of strong, three-bond tile attachments that strongly lock the cap in place. Even a single such attachment will fall off at a rate slower than two-tile detachment by a factor of $e^{-G_{se}}$; thus multiple attachments at reasonable values of G_{se} will result in a cap that will fail, after full attachment, very slowly.

Such a system is shown in Figure 4.5(c). In simulations with the same parameters as the first cap design, some counters still never have the cap attach correctly; some also have errors in counting, possibly caused by cap tiles interfering with growth, that result in the cap stopping growth at the wrong length. However, the majority of counters stop correctly, and once capped, remain capped for the duration of the simulations.

4.5 The Spurious Counter System

As seen in the previous section, the zig-zag binary counter system can be used to make finite assemblies through the use of a capping construction. It is disappointing, however, that the cap prevents us from using significant concentration variations to reduce facet nucleation. In our experimental results, most errors seen seem to be from facet nucleation.

Full snaked proofreading could significantly ameliorate this problem [7]. For a zig-zag system, however, this would require that each row be transformed into at least two rows that grow in parallel in the same directions, making the system considerably more complex.

In considering an alternative system design, we examined the possibilities of designing systems that took advantage of specifics of the kTAM in their fundamental logic. While the aTAM approximates the behavior of the kTAM in the limit of long assembly time and low concentrations, and the abstract nature of the growth makes intuitive design of complex systems comparatively simple, the kTAM can allow behavior, especially probabilistic and rate-based behavior, that isn't possible in the aTAM. In a manner similar to a programmer writing assembly to take advantage of architecture-specific features in the middle of a program, kTAM-specific tile system designs can be made.

As an example, unlike the aTAM, where attachment takes place in steps, and detachment never occurs, in the kTAM there is an inherent hierarchy of time scales, separated by factors of $e^{G_{se}}$, for detachment of tiles attached by different numbers of bonds. This hierarchy results in a corresponding difference of time scale for growth through tiles that attach by two bonds, and "facet nucleation" caused by a tile that attaches to a facet by one bond and allows further tiles to attach. Near equilibrium, normal growth will proceed at a rate $\propto e^{-G_{se}}$, while facet nucleation will occur at a given site at a rate $\propto e^{-2G_{se}}$ [7].

In previous systems, facet nucleation has been considered an error, and constructions like snaked proofreading have been developed to reduce its rate by at least a factor of $e^{G_{se}}$ [7]. The selective use, however, of snaked

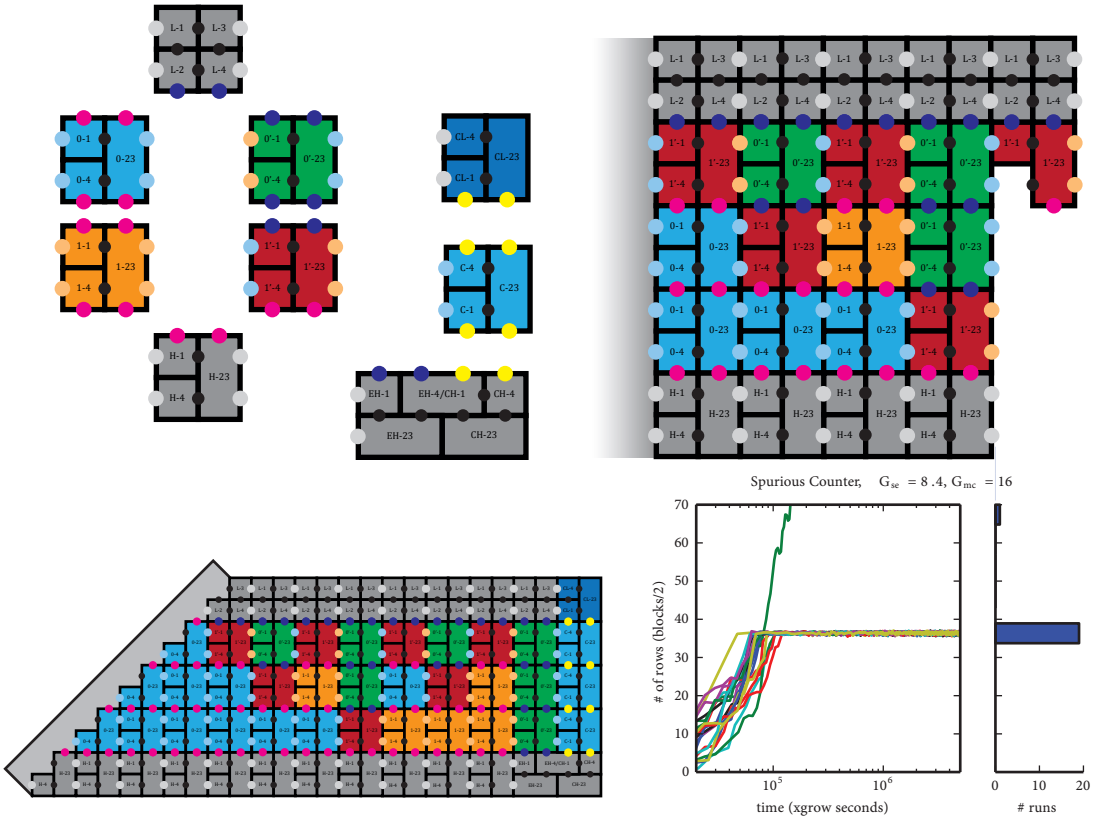

Figure 4.6: The spurious counter system. The tile system uses six 2×2 blocks, plus three other blocks for capping, as shown in (a). Unique bonds are indicated by small black dots; tiles edges without dots are null bonds that do not contribute to binding. Rows are nucleated by facet nucleation in the low-value boundary, which does not incorporate snaked proofreading, and proceeds downward by snaked growth as in (b). The cap attaches quickly, in comparison to row nucleation, after an overflow condition is reached. This counter, for simulations of a *six*-bit system starting at 101110, assemblies quickly and with high accuracy. Row nucleation and cap attachment can be regulated by changing the concentrations of the tiles involved; in simulations, a $10\times$ concentration of low boundary tiles and $2\times$ concentration of cap tiles was used relative to all other tiles.

proofreading allows us considerable control over where facet nucleation will likely occur. Facet nucleation is only an error in that it allows growth out of correct order, and attachment by less than a sufficient number of bonds for aTAM growth. With intentional facet nucleation, desired growth can be achieved with a system that will not function in the aTAM.

We make use of controlled facet nucleation to design the "spurious counter" system, illustrated in Figure 4.6(a-c). While similar in appearance to the zig-zag binary counter, growth does not proceed in a zig-zag fashion. The low-bit side boundary uses uniform proofreading that does not reduce facet nucleation, while all other blocks use snaked proofreading that significantly reduces it. From an aTAM perspective, the system should not work at all: once a row has completed, there are no further aTAM-correct attachments: the *only* way for the next row to start is with a facet nucleation event.

Since the low-bit boundary does not have reduced facet nucleation, it is very likely that facet nucleation will occur there compared with other blocks along the edge of a completed row. Since the bonds within the boundary are unique, attachment by one bond is still deterministic, and will result in a correct block from

which the next row can grow. Facet nucleation, however, will still take place at a rate of approximately $e^{-2G_{se}}$, compared with the system's average growth rate of $e^{-G_{se}}$. Thus, for high enough G_{se}, normal growth will be fast enough that a row will be very likely to complete before facet nucleation occurs to start the next row.

This use of two different growth rates—one for normal growth and one for intentional facet nucleation—allows a control over growth order with high probability that would be impossible in the aTAM. Capping in this system can be achieved in much the same way that capping is implemented in the zig-zag binary counter. A capping row that is nucleated with normal growth will tend to attach much faster than a new row will be nucleated through intentional facet nucleation.

Since facet nucleation is controlled via snaked proofreading in areas where it isn't intended to occur, the system can have significantly reduced error rates. As shown in Figure 4.6(d), simulations of wider and longer counter ribbons than the simulated zig-zag counters from Figure 4.5(b) result in very high yields under comparable conditions.

The construction of this system is one example of a design that takes advantage of specifics of the kTAM to allow growth that would not be possible in the aTAM. Whether there are other advantages of kTAM-based design, and whether there may be a modular and robust way of incorporating such features into tile system design, is an open question.

4.6 Conclusions and Discussion

In this chapter, we have considered the design of a specific tile system—a binary counter—from the aTAM to the kTAM level (experimental details are considered in the next chapter). Beyond serving as an illustration of tile system design, several considerations are of more general relevance, and could warrant further research.

The result that concentration variations have relatively little effect on growth errors is reassuring: it suggests that small experimental variations in concentrations will be unlikely to drastically increase errors. Higher tolerances such as these translate into cheaper experiments that can be performed with less skill and expertise. Facet nucleation, however, still looms as a complex challenge that is affected by concentration. It can result in complex, non-obvious behavior that is extremely dependent upon specifics of the tile systems being considered.

This perhaps reiterates the importance of snaked proofreading, or some other form of facet nucleation control. Beyond its use only in areas that are obvious candidates for proofreading, where errors can immediately arise from growth errors or facet nucleation, our results demonstrate the usefulness of snaked proofreading everywhere in a system. In the counter without a cap, apart from overflows, the low and high nucleation barrier tiles have no real algorithmic behavior or choice in tile attachments, and facet nucleation can only result in correct tiles attaching. Yet the backward growth resulting from those attachments can result in errors that may actually be far from the initial facet nucleation.

The spurious counter system is one example of the usefulness of snaked proofreading. The cap design doesn't incorporate the "strong" attachment of tiles by three bonds used in the cap for the zig-zag counter. Since proofreading is used throughout, the barrier to cancerous growth is high enough that the cap still works. Similarly, the facet nucleation complexities of the zig-zag counter can largely be ignored for the spurious counter, where undesired facet nucleation during growth is unlikely.

Our capping construction works in simulation, and preliminary results in Chapter 5 suggest that it likely forms correctly in experiments. Capping, however, and the construction of finite assemblies in the kTAM,

seems to be a major, interesting topic for further research. Finite assemblies are widely studied in the aTAM, but stopping growth is trivial. The kTAM makes the reliable construction of finite assemblies far more complex and challenging.

The mechanisms used in the spurious counter are a short foray into kTAM-based tile system design, and point to the possibility of growth behavior that would be impossible in the aTAM. Design *beyond* the aTAM in this way, rather than constructions to move kTAM system behavior closer to the aTAM, is largely unstudied. It is possible that such designs, though more difficult to understand, could result in more efficient, simpler systems with better experimental viability, much like processor-specific assembly in high-performance computing. Research into robust, perhaps modular design making use of the kTAM, beyond the clever trick we used here, could be rewarding.

Chapter 5

Toward a Capping, Self-Assembling DNA Tile Binary Counter

5.1 Introduction

Self-assembly of DNA tiles, as seen in previous chapters, has significant theoretical potential as a technique for programmably constructing nanoscale structures. The ability to have a small system of tiles that can, depending upon programmable input, form into a wide variety of different structures with high yield and precision is very appealing.

Experimentally, however, tile assembly of finite structures has largely been limited to tile systems with uniquely-addressed tiles [77, 33]. While the results of these experiments have been impressive, such constructions are limited in scale: each additional tile in the final assembly requires another unique tile type, and more unique sticky ends, while the maximum numbers of both are limited by sequence space considerations [17]. Constructing different structures requires different sets of tiles.

Utilizing algorithmic self-assembly would allow small sets that could construct very large, programmable structures. Experiments with algorithmic systems, however, have largely been designed to grow unbounded assemblies, which are then examined after growth for some period of time [22, 57, 3, 4]. Control of nucleation and error rates have been gradually diminishing challenges, and increasingly complex assemblies have become feasible.

In this chapter, we use an experimental implementation of the zig-zag binary counter system from Chapter 4, with and without a capping mechanism, incorporating several advances in design and growth as described in Chapter 1 [61, 64], to work toward low error rate, high accuracy algorithmic self assembly that can result in programmable, finite structures. In theory, our system can produce fixed-width ribbons of programmable length limited only by the number of bits in the seed and with a resolution of 13 nm. Our experimental system uses an origami seed that constructs a 5-bit ribbon, limiting length to around 800 nm.

In order to grow such structures with reasonable yield, an error rate is required that is significantly lower than $1/N$ for a structure of N tiles; for our system, this results in error rate requirements of around 1.6% per tile for a counter of 2 rows, down to 0.1% for a counter of 32 rows. While experiments with bit-copying ribbons have previously found error rates of between 0.065% to 0.009% [4, 64], experiments with binary counter ribbons have found error rates of approximately 10% per tile [3] in 2005 to 1% per tile [4] in 2009, with the best crystals counting from 0 to 6 without errors. Without the inclusion of capping, we have found error rates in preliminary samples of around 0.13% per tile, with two crystals that counted perfectly from 0 to 31. Preliminary data of the binary counter system with a capping mechanism has also resulted in ribbon structures that show evidence of being capped.

Nucleation of correct-width ribbons, however, is lower than desirable, with numerous ribbons showing lattice defects in the origami-ribbon attachment that cause reduced width counters (counting still takes place correctly). Capping results suggest that the cap structure may significantly impede correct growth of the counter system, requiring further optimization. Additionally, analysis of the types of errors seen in our results suggests certain mechanisms for errors in the system that could be examined further and optimized.

5.2 DESIGN

The binary counter system used in this chapter, illustrated in Figure 5.1(b-c), is an experimental implementation of the design presented in Chapter 4. For the counter itself, 22 DAO-E tiles were used with 5 nt sticky ends in a zig-zag ribbon system nucleating from a long rectangle origami structure with 16 "adapter" tiles, derived from the systems used in Barish et al and Schulman et al [4, 63, 64]. The core and sticky end sequences for these tiles were taken from DAO-E tiles used successfully in those experiments, and rearranged to fit the logic of the system.

The two-layer capping structure shown in 5.1(d) was based on the capping considerations and mechanism discussed in Chapter 4, but for simplicity incorporates the double tile from the low-bit boundary of the counter. The core sequences of the other 9 tiles in the structure were designed using a modified version of Pepper and spuriousC that incorporated spuriousC's scoring function into a simulated annealing algorithm [38]. The additional sticky end sequences used for the cap were designed using an early version of StickyDesign [17].

In order to interpret images of grown crystals, some form of labelling was needed. While many previous experiments incorporated hairpin labels on specific tiles, there is evidence that hairpins on tiles have significant effects on energetics and growth. We instead incorporated two internal biotin-modified dTs into the core sequences of tiles we wished to label, using the biotin-streptavidin tile labelling technique developed in Schulman et al [64]. For our earlier experiments, the T0-D tile in zero blocks was labelled, while for later experiments, the T1-C tile in one blocks was labelled instead.

As discussed in Chapter 4, the concentrations of different tiles in a system can have important effects on growth order, and need to be considered when designing a system. Throughout our experiments, we varied the concentrations we used in hopes of achieving better results. In general, concentrations of zero and one tiles were 100 nM. While later analysis (in Chapter 4) showed that the change was unlikely to have significant effects on error rates, incrementing tiles were used at 70 nM owing to their lower frequency of use. Low-bit and high-bit boundary tiles were used at between 25 nM and 50 nM, with the concentration of the low-bit side between 1.4 to 2 times that of the high-bit side. The lowered concentrations were intended to lower rates of spurious

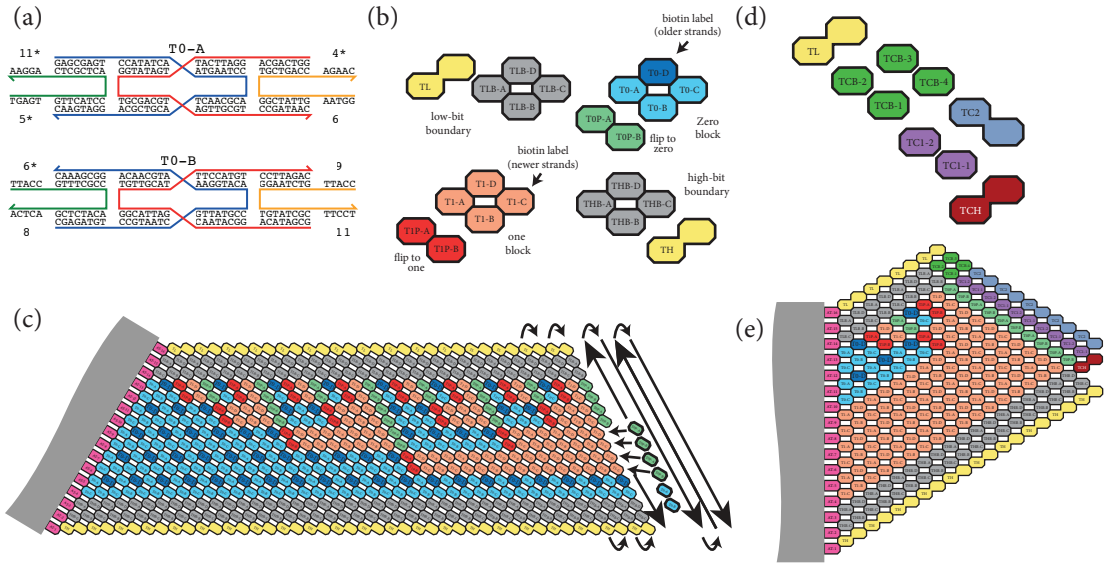

Figure 5.1: Design of the binary counter system. (a) shows the DAO-E structures of two tiles with their sequences; all tile structures are shown in Appendix A. (b) shows the tiles used for the counter system, which nucleates from an origami seed with adapter tiles to form a 16-wide ribbon (c). The two-layer cap uses the tiles in (d) to form the structure shown in (e) when the counter overflows. (e) also shows a counter ribbon that starts at a value through the use of alternate adapter strands.

nucleation, as the boundary tiles by themselves make a critical nucleus for a 6-tile-wide zig-zag ribbon, as described in Chapter 1. The excess of the low-bit boundary tiles was motivated by facet nucleation and growth order considerations discussed in Chapter 4: except when considering cap attachment, backward growth from the low-bit boundary is less likely to cause errors than backward growth from the high-bit boundary.

5.3 Experimental Methods

Experiments for this chapter were performed over a number of years, with differing samples and strands. After a successful sample in 2010, subsequent samples had significantly higher error rates, lower yields of correctly nucleated ribbons, and numerous problems with growth that were often tracked down to particular tile types; reordering component strands would often resolve these problems, while other would arise. Strands used in these experiments were a combination of numerous orders over several years, sometimes used for several other experiments; some strands were over 5–7 years old, stored in water at $-20\,^{\circ}\mathrm{C}$ and frozen and thawed an unknown number of times.

In response, for later experiments, we reordered all strands used in the system, with the exception of scaffold and staple strands for the origami seed. In doing so, we incorporated a single design change, moving the biotin label from the zero tile block to the one tile block to simplify the interpretation of counting.

In the older experiments, DNA strands (Integrated DNA Technologies) for the system were PAGE-purified. In the newer experiments, they were instead either HPLC-purified (for tile and adapter strands shorter than 60 nt) or PAGE-purified (for tile strands longer than 60 nt). In both, guard strands (discussed below) and origami staple strands were left unpurified; as they could be present in vast excess, and were not part of the

algorithmic self-assembly process, it was not expected that purity would be important. The origami scaffold used was M13mp18 (Bayou Biolabs). Concentrations used for various components are described in detail in Appendix A.

Samples of all system strands mixed together, in TAE/Mg^{2+} buffer (40 mM Tris, 20 mM acetic acid, 1 mM EDTA, 12.5 mM Mg acetate), were grown using the constant-temperature growth method described in Schulman et al and Chapter 1 [64], with gradient-capable thermocyclers (Eppendorf). Samples were first heated to $90\,^{\circ}$C, then reduced to $37\,^{\circ}$C at $1\,^{\circ}$C/min and held at this temperature for 6 hours, as an anneal for individual tiles that remained significantly above the melting temperature for ribbons. The temperature was then reduced at $0.1\,^{\circ}$C/min to a ribbon growth temperature around 32 to $33\,^{\circ}$C, depending upon the sample, and left at this temperature for between 10 to over 90 hours.

The intention of this protocol was to allow constant temperature growth near equilibrium, but settings required for melting temperature depend both on exact tile concentrations (that vary based on pipetting error and strand purity) and thermocycler temperature accuracy. As a result, instead of searching for ideal growth temperatures, we ran multiple samples with a gradient of $\pm 1\,^{\circ}$C around an approximate growth temperature.

To image, "guard strands" that were complementary to sticky-end-containing strands on particular tiles (described in Appendix A) were added in significant excess (at least $20\times$; see Appendix A) to samples, which were then left at the growth temperature for between 10 to 30 minutes. These strands served to deactivate a sufficient number of tiles in the system to prevent further growth, as described in Chapter 1. Samples were then imaged on mica at room temperature using a MultiMode AFM (Bruker) with a Nanoscope IIIa controller in fluid tapping mode. A modified buffer with 100 mM of added Na^{+} was added as needed to lower mica-DNA interactions in order to obtain cleaner images. In some experiments, a "washing" technique was also employed, where the sample was deposited on mica with the sodium-added buffer, which was then replaced several times before adding unmodified buffer; this process was meant to remove deactivated free tiles and other material that might increase background noise, while still resulting in crystals that strongly adhered to mica for optimal imaging. However, imaging in this manner was not found to be significantly better than simply imaging with some amount of added sodium buffer to an extent that the extra effort was worthwhile, and it was not used in later experiments. In all experiments, 1 µM streptavidin was used, and added after deposition of the samples on mica, sometimes after initial imaging without labels. Around 2 µL to 5 µL was usually added to samples, taking into account resulting images and background attachment of streptavidin to mica.

5.4 Results and Analysis

Counting

Samples of the counting system without a cap resulted in several long, high-accuracy ribbons, with a selection of them shown in Figure 5.2. For the older experiments, we found one ribbon that correctly nucleated and counted with only one error for 64 rows. After this point, the ribbon began to become narrower through lattice defects, and showed more errors; a possible explanation is that at this point in growth some tiles may have been depleted compared to others. In a sample with reordered strands, held at $33\,^{\circ}$C for approximately 97 hours, we found numerous ribbons that appeared to be counting correctly. Some, with errors in nucleation from origami, resulted in 4-bit or narrower ribbons, though these also counted; they are shown along with other problematic ribbons in Figure 5.3.

In total, we analyzed six 5-bit ribbons: the first correctly-nucleated ribbon found in the older sample, and the first five correctly-nucleated 5-bit ribbons from the newer sample. Excluding overflows (discussed below), these counters had a combined 264 rows, 1,380 counting blocks (not including nucleation blocks), or 8,508 tiles. With 11 errors, this resulted in an error rate per *row* of around 4.2%, per counting block of 0.80%, or per tile (with 30 tiles per increment) of 0.13%. Including only ribbons from the new sample resulted in an error rate per counting block of 0.96%. Both counter (a) from experiments with the older strands and one counter (c) from the newer experiments successfully counted from 0 to 31 with *no* errors (as a comparison, the example in Barish et al counted correctly from 0 to 6 [4]).

In all crystals, there were some locations where counting logic in adjacent rows was clearly correct, but a streptavidin label likely either did not attach or actually detached; in some cases, it was clear that the label, or even the entire tile, had been forcibly detached by the AFM tip, sometimes disappearing mid-frame. We have described these bits as labelling errors, and have not included them as actual errors for error rate purposes.

Without a cap, ribbons were free to continue growing after an overflow with only a single insufficient attachment. Thus, all ribbons seen reached an overflow and continued growing. During some overflows, the counting reset to zero, while in others, high-value ones that should correctly have been incremented instead remained ones. This behavior is not unexpected: the simplest errors resulting in continued ribbon growth after an overflow are (1) the high-value boundary attaching with one error, and continuing growth normally with an all-zero value, (2) copying one tiles attaching in place of incrementing zero tiles, allowing the high-value boundary to attach normally and continuing growth with one or more high-value ones, or (3) facet nucleation anywhere along the facet, with several different possibilities for resulting counter value. Interestingly, ribbons remained the same width after overflows, despite simulations suggesting the possibility of width changes or even cancerous growth lacking the high-value boundary. This was particularly striking in narrow, 2-bit counters in older experiments that were seen in some samples; these correctly counted and remained the same width, despite requiring an insufficient attachment every few rows in order to continue growth.

All errors other than overflows seen in analyzed counters resulted in the same logical error: after correctly incrementing in lower-value bits, a higher-value 1 bit that should have been copied from the previous row was instead incremented again, changing to a 0 and propagating the "incrementing" action to higher-value bits until reaching a 0, as would be correct for incrementing behavior. This is an expected error for the system. An incrementing zero tile block, attaching instead of a copying one, requires only one insufficient attachment, and would explain the error. Alternatively, an incrementing one tile block, attaching via facet nucleation to a zero in the previous row, would result in higher-value bits being (correctly) copied and lower-value bits being grown out of order and flipped. Finally, facet nucleation in the high-bit boundary could allow backward growth of the incrementing row. In this case, no high-value one could be flipped to a zero without further errors, but a zero could be flipped to a one by an incrementing one tile, resulting in the same backward growth error.

However, other expected errors were not seen: no instances of bits being incorrectly copied rather than incremented were found, for example, nor were there any errors that could be explained by an incrementing one tile simply attaching via growth error and continuing with normal growth, which would have resulted in a one appearing in a sea of zeros.

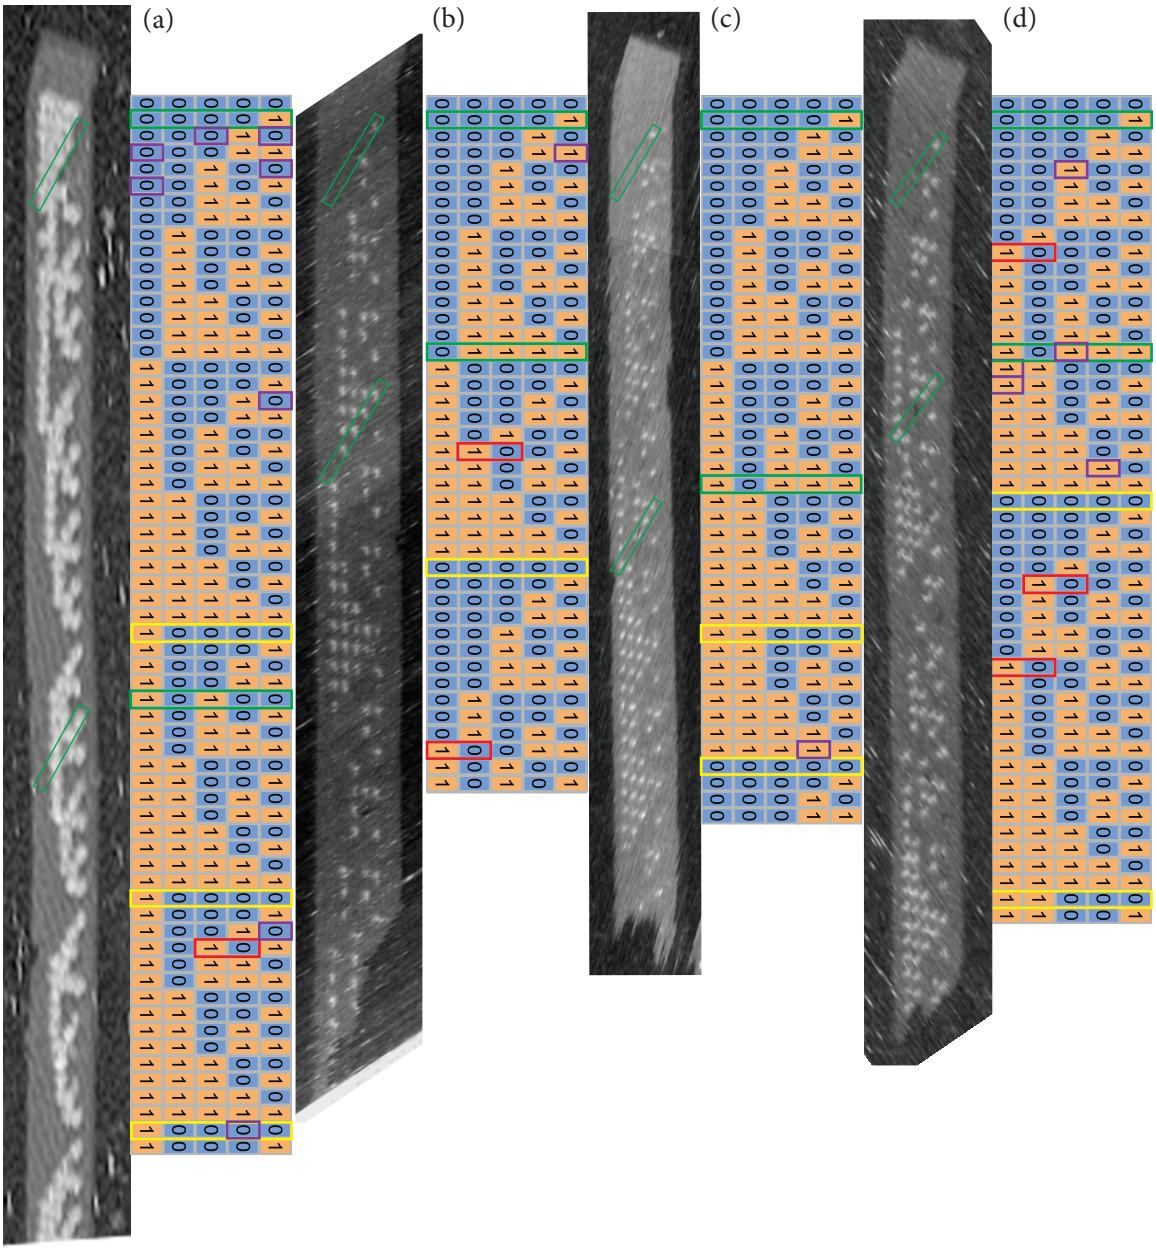

Figure 5.2: AFM images of binary counters without capping. (a) is from experiments run in 2010, and uses biotin labels on zero blocks, with older strands. (b-d) are from 2014, and use biotin labels on one blocks, with a new order of strands. (c) is a composite of two images. Red boxes show locations of counting errors; yellow boxes show overflows, and purple boxes show locations where labels were not present, but analysis of adjacent rows suggested that a labelled tile was present without an attached streptavidin, a "labelling error" that was not an actual error in growth. Green boxes illustrate corresponding rows in the images and interpretations. All ribbons are 5 bit counters, and correctly nucleated; each ribbon is approximately 80 nm wide.

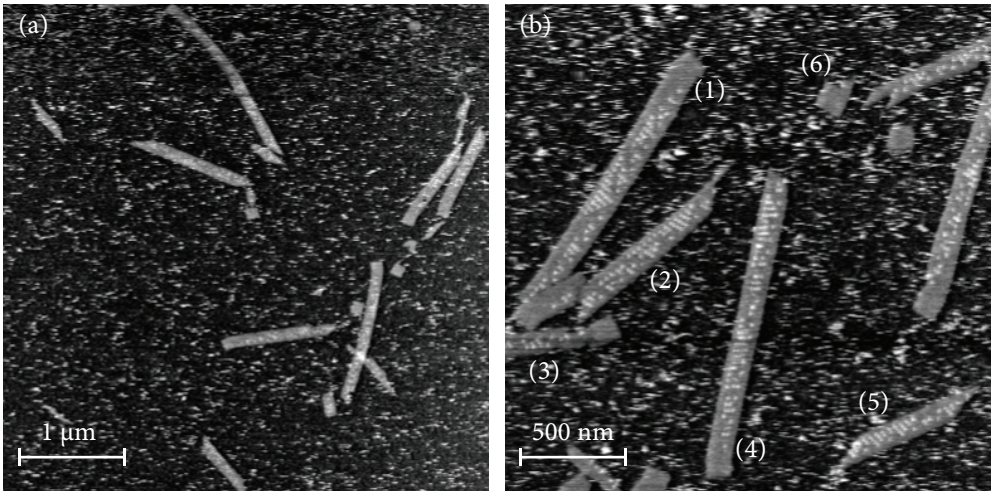

Figure 5.3: (a) shows a wide-field image of the 2014 sample, with correctly and incorrectly nucleated ribbons. (b) shows several examples of ribbons with problems: a (1) 5-bit ribbon that narrows to 4 bits, (2) broken ribbon, presumably broken during deposition, (3) poorly nucleated ribbon, (5) ribbon fragment, and (6) origami seed without any ribbon growth.

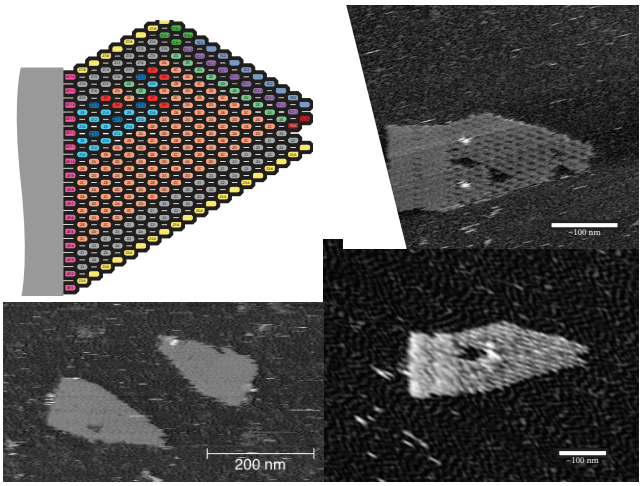

Figure 5.4: Preliminary AFM images of binary counters with evidence of capping.

Capping

In a preliminary experiment to examine cap growth, with the older set of strands, counters were grown starting from a very high value (11100) such that they would overflow very quickly. While streptavidin labelling of the results was inconsistent and problematic, several very short ribbons were seen. These ribbons, shown in Figure 5.4, all exhibited straight facets with a uniform "notched" point, also seen in the diagram of a capped ribbon. In (b) in particular, three rows of tiles after this notch are clearly visible, as expected for a correctly attached cap structure, and the structure additionally appears qualitatively similar in shape to a well-formed cap.

These results suggest that the cap structure is likely able to form correctly. However, numerous structures in the background, and the poor labelling of the structures, suggest that capping may have interfered with

counting or that there may have been continuing problems with the older set of strands. The inconsistency of results with the older strands calls into question many matters of error rates and yield when using them. Further experiments, with our newer strands, will be necessary to properly investigate capping.

5.5 Discussion

Our results are preliminary, and need to be expanded. However, they do show that accurate algorithmic self-assembly is possible at present with improved technologies. The counters seen with our newer strands were by no means elusive; they were the product of our first attempt at growth with the reordered strands, and did not require precise optimization of concentrations or temperatures. We hope that further results may be able to improve upon what we have already seen, and also result in finite, programmable structures through capping.

The types of errors that we see warrant some further analysis. Despite several counting errors being possible from the lack of full proofreading and seen in simulations, only one type has been seen experimentally. This may suggest that facet nucleation is a larger problem than might be expected. Width-change errors, seen both in our results and earlier ribbon experiments [4, 64], seem to occur through lattice defects that aren't considered in the kTAM. Furthermore, overflow behavior may be different experimentally and in the kTAM, with simulations often changing width at overflows. All of these factors may indicate areas where experiments may diverge from the kTAM. Analyzing the relative prevalence of different errors and the ways they arise in simulation and more extensive experimental results could prove insightful.

The most comparable earlier experimental results for algorithmic self-assembly come from Barish et al and Schulman et al [4, 64]. Barish et al [4], using a similar binary counting system with earlier experimental techniques, achieved an error rate of 4.1% per counting block for selected ribbons, after purification and ligation. By comparison, our unpurified, preliminary results have an error rate of 0.80% per block. Barish et al [4] also presents results for a similar ribbon that copies bits rather than incrementing, using the same older techniques. This system, which lacks the proofreading problems (described in Chapter 4) of the binary counter system and has some tolerance of facet nucleation, resulted in error rates of 0.26% per block. With techniques similar to ours, Schulman et al [64] found copy ribbon error rates of 0.034%. Thus improvements in growth methods for both systems have resulted in similar error rate improvements, while the significantly lower error rates for the copying system are likely the consequence of better proofreading and control of facet nucleation, both matters of tileset design.

These comparisons, and the types of errors seen, suggest that while the advances used to achieve our results have largely been improved experimental techniques and some new designs built around a proven base, further improvement may be limited by intrinsic design problems, especially facet nucleation. The design of a completely new binary counter system could incorporate many significant improvements, such as snaked proofreading, optimized sequence design and energetics-optimized concentrations. It is tempting to think about just how much better such a design might work, especially since the only errors in counting we have seen so far have been errors that would be greatly reduced by snaked proofreading.

As noted above, the uniform, zig-zagging proofreading of copy ribbons seems to result in more than a ten-fold reduction in error experimentally. With snaked proofreading, perhaps even another ten-fold reduction could be achieved, and use of 3×3 snaked proofreading might reduce error rates even further. Additionally, our results in Chapter 3 suggest that with improvements to sequence design and assignment, which have not yet

been implemented in any experimental system, a further five-to-ten-fold reduction in error rates might be possible. With the combination of these fundamental design improvements, and a system that is designed with careful consideration of the issues raised in this thesis, a considerable reduction in error rates might be possible: even if the reductions do not combine perfectly, or other limiting factors arise at some level (strand quality or malformed tiles, for example), it may still be possible to reduce experimental errors, with our current understanding of tile assembly, by a factor of 100 to 1000. Considering our current 0.14% per tile error rate, this could allow for high yields with structures of 70,000 to 700,000 tiles. With the latter, even a full 12-bit counter, counting to 4096, might be possible, allowing a precise ribbon of up to 0.1 mm to be constructed while keeping 13 nm resolution.

Another consideration is the potential for more complex systems, with more tile types. As discussed in Chapter 3, sequence space is a limiting factor for the number of unique sticky ends, especially for the 5 nt ends of the DAO-E tiles used in most algorithmic self-assembly experiments. Taking into account sticky end sequence orientation, and allowing for a non-orthogonal binding energy fraction of $s_{ij} \leq 0.5$, our sequence designer in Section 3.3 could provide around 80 sticky ends (40 $5'$-terminating and 40 $3'$-terminating). Even with longer sticky ends, the number of possible sticky end types may be limited when compared to the number used in complex theoretical constructions.

However, the sequence assignment algorithm in Chapter 3 should work particularly well in combination with proofreading, as sticky ends inside proofreading blocks will not typically result in sensitive sticky-end pairs, and so with more proofreading, the use of lower-quality sequences, allowing stronger non-orthogonal binding, may be possible; this would significantly enlarge the number of possible sticky ends. Also, from a theoretical perspective, even around 80 sticky ends may allow for interesting and complex algorithmic behavior; the design of tile systems that make parsimonious use of sticky ends could be an interesting area of theoretical research. It is possible that design methods could employ a larger number of tile types to compensate: 80 sticky ends types could allow for up to 4^{80} different tile types. Perhaps reuse of sticky ends, or some form of sticky-end-type reducing transformation, could be a worthwhile area of investigation.

Of more immediate experimental interest, however, remains the issue of facet nucleation and the potential of full snaked proofreading. While snaked proofreading has not yet been implemented in an experimental system beyond tests of its basic mechanism, it is extremely promising in theory and simulations [7, 10]. Is it the future of experimental algorithmic self-assembly?

Bibliography

[1] Len Adleman et al. “Combinatorial Optimization Problems in Self-assembly”. In: *STOC 2002*. Montreal, Quebec, Canada: ACM, 2002, pp. 23–32. DOI: 10.1145/509907.509913.

[2] Leonard Adleman et al. “Running Time and Program Size for Self-assembled Squares”. In: *STOC 2001*. Hersonissos, Greece: ACM, 2001, pp. 740–748. DOI: 10.1145/380752.380881.

[3] Robert D. Barish, Paul W. K. Rothemund, and Erik Winfree. “Two Computational Primitives for Algorithmic Self-Assembly: Copying and Counting”. In: *Nano Letters* 5.12 (2005), pp. 2586–2592. DOI: 10.1021/nl0520381.

[4] Robert D. Barish et al. “An information-bearing seed for nucleating algorithmic self-assembly”. In: *PNAS* (2009). DOI: 10.1073/pnas.0808736106.

[5] Robert Berger. “The undecidability of the domino problem”. In: *Memoirs of the AMS* 66 (1966), pp. 1–72.

[6] Magdalena Bezanilla et al. “Adsorption of DNA to Mica, Silylated Mica, and Minerals: Characterization by Atomic Force Microscopy”. In: *Langmuir* 11.2 (Feb. 1995), pp. 655–659.

[7] Ho-Lin Chen and Ashish Goel. “Error Free Self-assembly Using Error Prone Tiles”. In: *DNA 10*. Vol. 3384. LNCS. Springer, 2005, pp. 702–707.

[8] Ho-Lin Chen, Ashish Goel, and Chris Luhrs. “Dimension Augmentation and Combinatorial Criteria for Efficient Error-resistant DNA Self-assembly”. In: *SODA 2008*. San Francisco, California: SIAM, 2008, pp. 409–418.

[9] Ho-Lin Chen and Ming-Yang Kao. “Optimizing Tile Concentrations to Minimize Errors and Time for DNA Tile Self-assembly Systems”. In: *DNA 16*. Vol. 6518. LNCS. Springer, 2011, pp. 13–24. DOI: 10.1007/978-3-642-18305-8_2.

[10] Ho-Lin Chen et al. “Reducing facet nucleation during algorithmic self-assembly”. In: *Nano Lett.* 7 (2007), pp. 2913–2919.

[11] Arkadiusz Chworos et al. “Building Programmable Jigsaw Puzzles with RNA”. In: *Science* 306.5704 (2004), pp. 2068–2072. DOI: 10.1126/science.1104686.

[12] Russell Deaton et al. “A Software Tool for Generating Non-crosshybridizing Libraries of DNA Oligonucleotides”. In: *DNA8*. Vol. 2568. LNCS. Heidelberg: Springer, 2003, pp. 252–261.

[13] Robert M. Dirks and Niles A. Pierce. "Triggered amplification by hybridization chain reaction". In: *PNAS* 101.43 (2004), pp. 15275–15278. DOI: 10.1073/pnas.0407024101.

[14] David Doty. "Theory of Algorithmic Self-assembly". In: *Commun. ACM* 55.12 (2012), pp. 78–88. DOI: 10.1145/2380656.2380675.

[15] David Doty, Matthew J. Patitz, and Scott M. Summers. "Limitations of Self-assembly at Temperature 1". In: *Theor. Comput. Sci.* 412.1-2 (2011), pp. 145–158. DOI: 10.1016/j.tcs.2010.08.023.

[16] Constantine G Evans, Rizal F Hariadi, and Erik Winfree. "Direct atomic force microscopy observation of DNA tile crystal growth at the single-molecule level." In: *J. Am. Chem. Soc.* 134 (2012), pp. 10485–10492. DOI: 10.1021/ja301026z.

[17] Constantine G. Evans and Erik Winfree. "DNA Sticky End Design and Assignment for Robust Algorithmic Self-assembly". In: *DNA 19*. Vol. 8141. LNCS. Springer, 2013, pp. 61–75. DOI: 10.1007/978-3-319-01928-4_5.

[18] Joyce A. Evans. *Celluloid Mushroom Clouds: Hollywood and the Atomic Bomb*. Westview Press, 1998.

[19] T J Fu and Nadrian C Seeman. "DNA double-crossover molecules". In: *Biochemistry* 32 (1993), pp. 3211–3220.

[20] K. Fujibayashi and S. Murata. "Precise Simulation Model for DNA Tile Self-Assembly". In: *Nanotechnology, IEEE Transactions on* 8.3 (2009), pp. 361–368. DOI: 10.1109/TNANO.2008.2011776.

[21] Kenichi Fujibayashi et al. "Error suppression mechanisms for DNA tile self-assembly and their simulation". English. In: *Natural Computing* 8.3 (2009), pp. 589–612. DOI: 10.1007/s11047-008-9093-9.

[22] Kenichi Fujibayashi et al. "Toward Reliable Algorithmic Self-Assembly of DNA Tiles: A Fixed-Width Cellular Automaton Pattern". In: *Nano Letters* 8.7 (2008), pp. 1791–1797. DOI: 10.1021/nl0722830.

[23] Peter Gàcs. "Reliable Cellular Automata with Self-Organization". In: *Journal of Statistical Physics* 103.1-2 (2001), pp. 45–267. DOI: 10.1023/A:1004823720305.

[24] Ashish Goel and Pablo Moisset de Espanés. "Toward Minimum Size Self-Assembled Counters". In: *DNA Computing*. Ed. by Max H. Garzon and Hao Yan. Vol. 4848. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2008, pp. 46–53. DOI: 10.1007/978-3-540-77962-9_5.

[25] Branko Grünbaum and G C Shephard. *Tilings and Patterns*. New York, NY, USA: W. H. Freeman & Co., 1986.

[26] Shogo Hamada and Satoshi Murata. "Substrate-Assisted Assembly of Interconnected Single-Duplex DNA Nanostructures". In: *Angewandte Chemie* 121.37 (2009), pp. 6952–6955. DOI: 10.1002/ange.200902662.

[27] William Hanf. "Nonrecursive tilings of the plane I". In: *The Journal of Symbolic Logic* 39 (1974), pp. 283–285.

[28] Rizal F Hariadi. "Non-equilibrium dynamics of DNA nanotubes". PhD thesis. California Institute of Technology, 2011.

[29] Byunghyun Jang, Yong-Bin Kim, and Fabrizio Lombardi. "Monomer Control for Error Tolerance in DNA Self-Assembly". English. In: *Journal of Electronic Testing* 24.1-3 (2008), pp. 271–284. DOI: 10.1007/s10836-007-5016-4.

[30] Byunghyun Jang, Yongbin Kim, and Fabrizio Lombardi. "Error Tolerance of DNA Self-Assembly by Monomer Concentration Control". In: *2006 21st IEEE International Symposium on Defect and Fault Tolerance in VLSI Systems*. IEEE, 2006, pp. 89–97.

[31] Christian Janot. *Quasicrystals*. Springer, 1994.

[32] N. Jonoska and D. Karpenko. "Active Tile Self-assembly, Self-similar Structures and Recursion". In: *ArXiv e-prints* (Nov. 2012). arXiv: 1211.3085.

[33] Yonggang Ke et al. "Three-Dimensional Structures Self-Assembled from DNA Bricks". In: *Science* 338.6111 (2012), pp. 1177–1183. DOI: 10.1126/science.1227268.

[34] Alfred Kick, Martin Bönsch, and Michael Mertig. "EGNAS: an exhaustive DNA sequence design algorithm." In: *BMC Bioinformatics* 13 (2012), p. 138.

[35] Byeonghoon Kim et al. "Growth and restoration of a T-tile-based 1D DNA nanotrack". In: *Chem. Commun.* 47 (39 2011), pp. 11053–11055. DOI: 10.1039/C1CC10477B.

[36] Zhe Li et al. "Molecular Behavior of DNA Origami in Higher-Order Self-Assembly". In: *J. Am. Chem. Soc.* 132 (2013), pp. 13545–13552.

[37] Kenneth G Libbrecht. "The physics of snow crystals". In: *Reports on Progress in Physics* 68.4 (2005), p. 855.

[38] Shawn Ligocki and Erik Winfree. *Personal communication*. 2013.

[39] Chenxiang Lin et al. "DNA Tile Based Self-Assembly: Building Complex Nanoarchitectures". In: *Chem Phys Chem* 7.8 (2006), pp. 1641–1647. DOI: 10.1002/cphc.200600260.

[40] Jonathan S Lindsey. "Self-assembly in synthetic routes to molecular devices. Biological principles and chemical perspectives: a review". In: *New J. Chem* 15.2-3 (1991), pp. 153–179.

[41] Daniel Lubrich, Simon J. Green, and Andrew J. Turberfield. "Kinetically Controlled Self-Assembly of DNA Oligomers". In: *Journal of the American Chemical Society* 131.7 (2009), pp. 2422–2423. DOI: 10.1021/ja807765v.

[42] Urmi Majumder, Thomas H. LaBean, and John H. Reif. "Activatable Tiles: Compact, Robust Programmable Assembly and Other Applications". In: *DNA 13*. Vol. 4848. LNCS. Springer, 2008, pp. 15–25. DOI: 10.1007/978-3-540-77962-9_2.

[43] Abdul M. Mohammed and Rebecca Schulman. "Directing Self-Assembly of DNA Nanotubes Using Programmable Seeds". In: *Nano Letters* 13.9 (2013), pp. 4006–4013. DOI: 10.1021/nl400881w.

[44] Dale Myers. "Nonrecursive tilings of the plane II". In: *The Journal of Symbolic Logic* 39 (1974), pp. 286–294.

[45] Ki Tae Nam et al. "Virus-Enabled Synthesis and Assembly of Nanowires for Lithium Ion Battery Electrodes". In: *Science* 312.5775 (2006), pp. 885–888. DOI: 10.1126/science.1122716.

[46] Jeanette Nangreave, Hao Yan, and Yan Liu. "Studies of Thermal Stability of Multivalent DNA Hybridization in a Nanostructured System". In: *Biophysical Journal* 97.2 (2009), pp. 563–571. DOI: 10.1016/j.bpj.2009.05.013.

[47] Lorena Nasalean et al. "Controlling RNA self-assembly to form filaments". In: *Nucleic Acids Research* 34.5 (2006), pp. 1381–1392. DOI: 10.1093/nar/gkl008.

[48] Jennifer Padilla. *Personal communication*. 2013.

[49] Jennifer E. Padilla et al. "Asynchronous Signal Passing for Tile Self-assembly: Fuel Efficient Computation and Efficient Assembly of Shapes". In: *Unconventional Computation and Natural Computation*. Vol. 7956. LNCS. Springer, 2013, pp. 174–185. DOI: 10.1007/978-3-642-39074-6_17.

[50] Sung Ha Park et al. "Programmable DNA Self-Assemblies for Nanoscale Organization of Ligands and Proteins". In: *Nano Lett.* 5 (2013), pp. 729–733.

[51] David Pastré et al. "Adsorption of DNA to Mica Mediated by Divalent Counterions: A Theoretical and Experimental Study". In: *Biophys J* 85.4 (Oct. 2003), pp. 2507–2518.

[52] M. J. Patitz. "Simulation of Self-Assembly in the Abstract Tile Assembly Model with ISU TAS". In: *ArXiv* (Jan. 2011).

[53] Matthew J. Patitz. "An Introduction to Tile-Based Self-assembly". In: *Unconventional Computation and Natural Computation*. Vol. 7445. LNCS. Springer, 2012, pp. 34–62. DOI: 10.1007/978-3-642-32894-7_6.

[54] Vinhthuy Phan and Max H Garzon. "On codeword design in metric DNA spaces". In: *Nat. Comput.* 8 (2008), pp. 571–588.

[55] Andre V. Pinheiro et al. "Steric Crowding and the Kinetics of DNA Hybridization within a DNA Nanostructure System". In: *ACS Nano* 6 (2013), pp. 5521–5530.

[56] John H. Reif, Sudheer Sahu, and Peng Yin. "Compact Error-Resilient Computational DNA Tiling Assemblies". In: *DNA 10*. Vol. 3384. LNCS. Springer, 2005, pp. 293–307. DOI: 10.1007/11493785_26.

[57] Paul W. K Rothemund, Nick Papadakis, and Erik Winfree. "Algorithmic Self-Assembly of DNA Sierpinski Triangles". In: *PLoS Biol* 2.12 (2004), e424. DOI: 10.1371/journal.pbio.0020424.

[58] Paul W. K. Rothemund and Erik Winfree. "The Program-size Complexity of Self-assembled Squares (Extended Abstract)". In: *STOC 2000*. Portland, Oregon, USA: ACM, 2000, pp. 459–468. DOI: 10.1145/335305.335358.

[59] Paul W K Rothemund et al. "Design and characterization of programmable DNA nanotubes". In: *J Am Chem Soc* 126.50 (2004), pp. 16344–16352.

[60] John SantaLucia and Donald Hicks. "The Thermodynamics of DNA Structural Motifs". In: *Annu. Rev. Biophys. Biomol. Struct.* 33 (2004), pp. 415–440.

[61] R. Schulman and E. Winfree. "Programmable Control of Nucleation for Algorithmic Self-Assembly". In: *SIAM Journal on Computing* 39.4 (2010), pp. 1581–1616. DOI: 10.1137/070680266.

[62] Rebecca Schulman and Erik Winfree. "Programmable Control of Nucleation for Algorithmic Self-assembly". In: *DNA 10*. Vol. 3384. LNCS. Springer, 2005, pp. 319–328. DOI: 10.1007/11493785_28.

[63] Rebecca Schulman and Erik Winfree. "Synthesis of crystals with a programmable kinetic barrier to nucleation". In: *PNAS* 104.39 (2007), pp. 15236–15241. DOI: 10.1073/pnas.0701467104.

[64] Rebecca Schulman, Bernard Yurke, and Erik Winfree. "Robust self-replication of combinatorial information via crystal growth and scission". In: *PNAS* 109.17 (2012), pp. 6405–6410. DOI: 10.1073/pnas.1117813109.

[65] Marjorie Senechal. *Quasicrystals and geometry*. CUP Archive, 1996.

[66] David Soloveichik, Matthew Cook, and Erik Winfree. "Combining self-healing and proofreading in self-assembly". In: *Natural Computing* 7.2 (2008), pp. 203–218. DOI: 10.1007/s11047-007-9036-x.

[67] David Soloveichik and Erik Winfree. "Complexity of compact proofreading for self-assembled patterns". In: *DNA 11*. Vol. 3892. LNCS. 2006, pp. 125–135.

[68] David Soloveichik and Erik Winfree. "Complexity of Self-Assembled Shapes". In: *SIAM J. Comput.* 36.6 (Feb. 2007), pp. 1544–1569. DOI: 10.1137/S0097539704446712.

[69] *StickyDesign*. URL: http://dna.caltech.edu/StickyDesign/.

[70] Xuping Sun et al. "Surface-Mediated DNA Self-Assembly". In: *Journal of the American Chemical Society* 131.37 (2009), pp. 13248–13249. DOI: 10.1021/ja906475w.

[71] F Tanaka. "Design of nucleic acid sequences for DNA computing based on a thermodynamic approach". In: *Nucleic Acids Res.* 33 (2005), pp. 903–911.

[72] *The Xgrow Simulator*. URL: http://dna.caltech.edu/Xgrow/.

[73] Dan Tulpan et al. "Thermodynamically based DNA strand design." In: *Nucleic Acids Res.* 33 (2005), pp. 4951–4964.

[74] Sylvain Vauthey et al. "Molecular self-assembly of surfactant-like peptides to form nanotubes and nanovesicles". In: *Proceedings of the National Academy of Sciences* 99.8 (2002), pp. 5355–5360. DOI: 10.1073/pnas.072089599.

[75] Dimitrios Vavylonis, Qingbo Yang, and Ben O'Shaughnessy. "Actin polymerization kinetics, cap structure, and fluctuations." In: *PNAS* 102.24 (2005), pp. 8543–8548.

[76] Hao Wang. "An unsolvable problem on dominoes". In: *Harvard Computation Laboratory, Technical Report BL30 (II-15)* (1962).

[77] Bryan Wei, Mingjie Dai, and Peng Yin. "Complex shapes self-assembled from single-stranded DNA tiles". In: *Nature* 485.7400 (31, 2012), pp. 623–626.

[78] Erik Winfree. "On the Computational Power of DNA Annealing and Ligation". In: *DNA Computers. DIMACS Series in Discrete Mathematics and Computer Science*. AMS, 1996, pp. 199–221.

[79] Erik Winfree. *Simulations of Computing by Self-Assembly*. Tech. rep. CaltechCSTR:1998.22. Pasadena, CA, 1998.

[80] Erik Winfree and Renat Bekbolatov. "Proofreading Tile Sets: Error Correction for Algorithmic Self-Assembly". In: *DNA 9*. Vol. 2943. LNCS. Springer, 2004, pp. 126–144.

[81] Erik Winfree et al. "Design and self-assembly of two-dimensional DNA crystals". In: *Nature* 394.6693 (1998), pp. 539–544.

[82] Peng Yin et al. "Programming DNA Tube Circumferences". In: *Science* 321.5890 (2008), pp. 824–826. DOI: 10.1126/science.1157312.

Appendix A

Supplementary Material for Binary Counter Experiments

A.1 Concentrations

Sample mixing was performed in stages. Individual strands were first mixed into individual tile mixes. For single tiles with the new strands, these were 7 μM of each shorter strand (strands 1 and 4) and 6 μM of each longer strand (strands 2 and 3); the effective concentration was considered to be 6 μM. Each double tile was mixed with 3 μM of strands 1 through 5, and 4 μM of strands 1 and 6; the effective concentration was considered to be 3 μM. Earlier experiments used larger concentrations for excess strands, either 3 or 5 times the concentrations of other strands; this was not found to make a significant difference in results.

Mixes were then constructed for each two or four-tile block in the system. These were finally mixed into samples. Concentrations used were varied slightly throughout experiments, depending upon the particular mixes and strands used and upon previous results, but as examples the concentrations used for the new counter results shown were:

- 50 nM of TL and TLB tiles (low-bit boundary tiles)
- 40 nM of TH and THB tiles (high-bit boundary tiles)
- 100 nM of T0 and T1 tiles (counting tiles)
- 70 nM of T0P and T1P tiles (incrementing tiles)
- 40 nM of all adapter tiles used
- 20 nM of origami staple strands
- 5 pM of origami scaffold (M13mp18)

Concentrations used for the capping results, with the older set of strands, were:

- 35 nM of TL and TLB tiles (low-bit boundary tiles)
- 25 nM of TH and THB tiles (high-bit boundary tiles)
- 100 nM of T0 and T1 tiles (counting tiles)
- 70 nM of T0P and T1P tiles (incrementing tiles)
- 25 nM of TCH
- 100 nM of TC1 tiles
- 35 nM of TCB tiles
- 25 nM of TC2
- 50 nM of all adapter tiles used
- 50 nM of origami staple strands
- 5 pM of origami scaffold (M13mp18)

Guard strands were prepared as a 7 μM per strand mix, though as strands were not purified, and were in significant excess, differing amounts were used in differing experiments without any differences seen in results. Around 5 to 10 μL of this was usually added to 20 μL samples at growth temperature, after first heating the guard mix to the same temperature.

A.2 Tile Designs

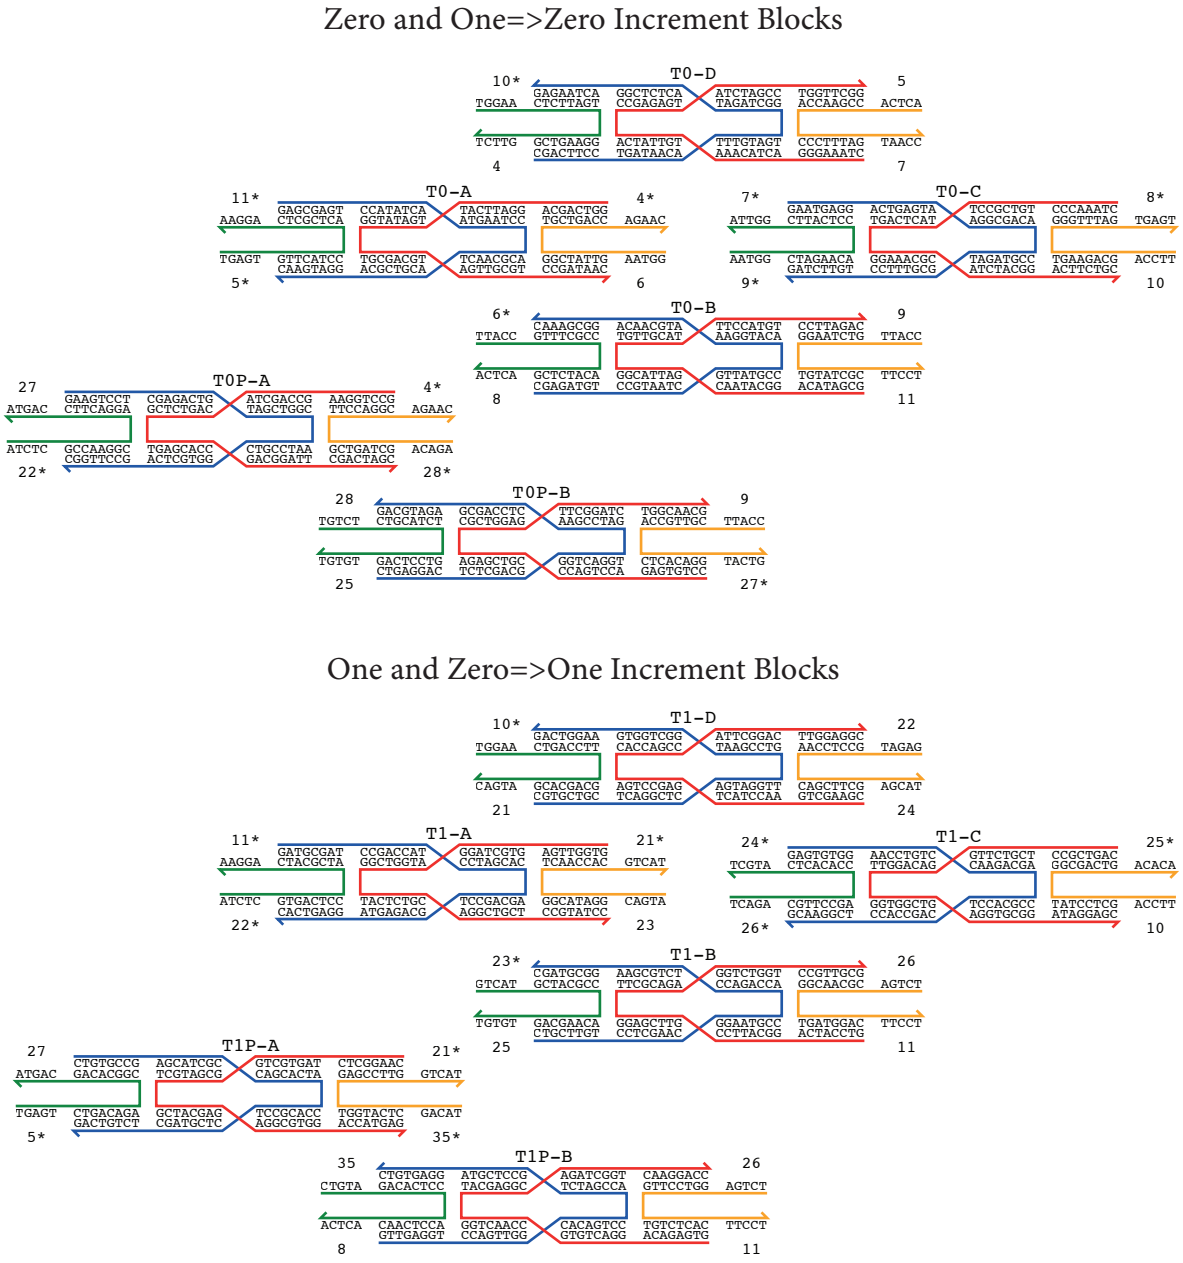

Figure A.1: Zero, one and increment tiles for the binary counter system.

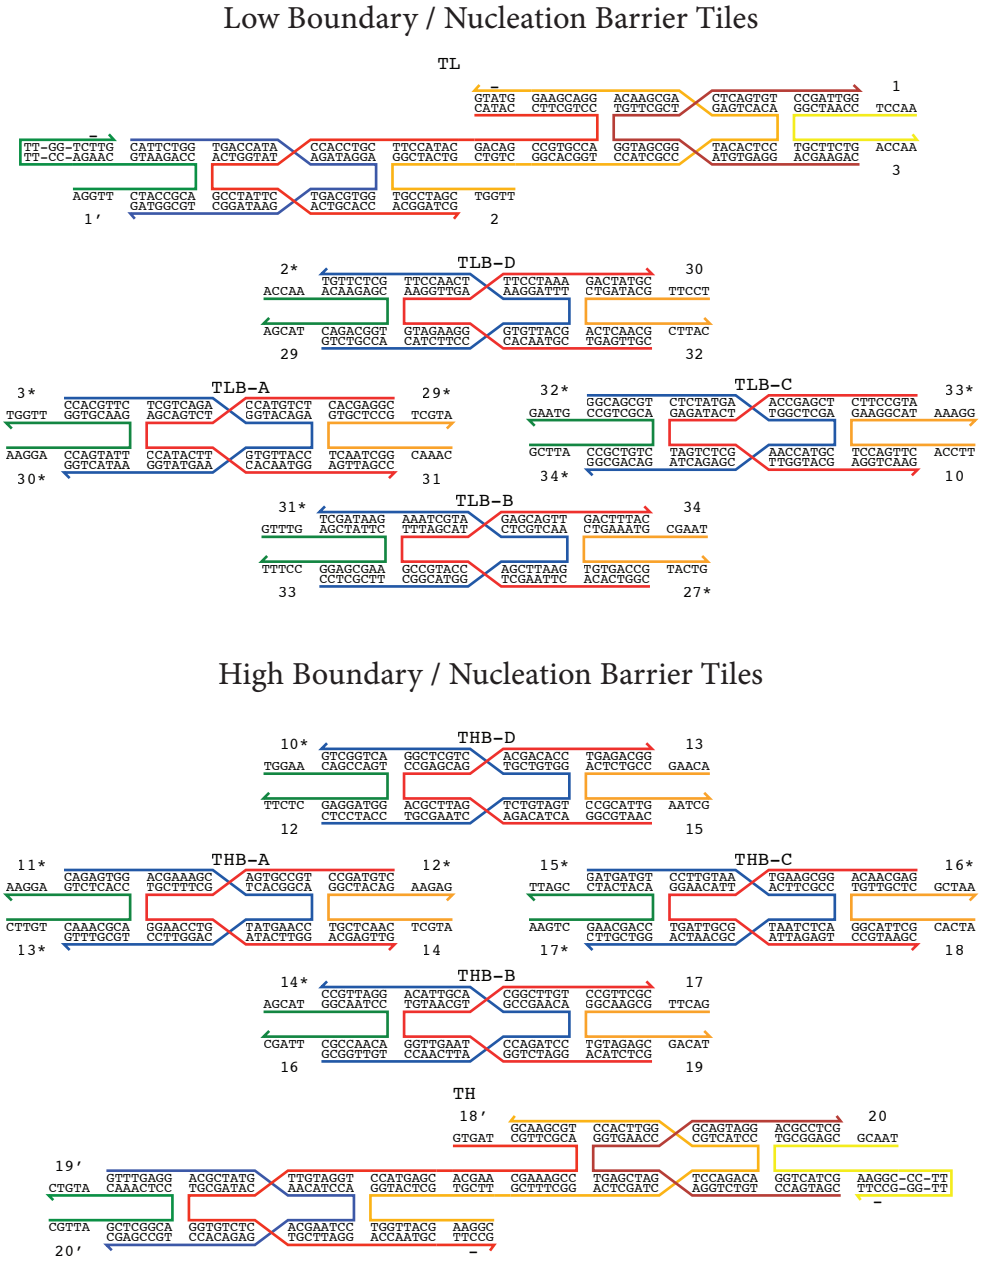

Figure A.2: Boundary and nucleation barrier tiles for the binary counter system.

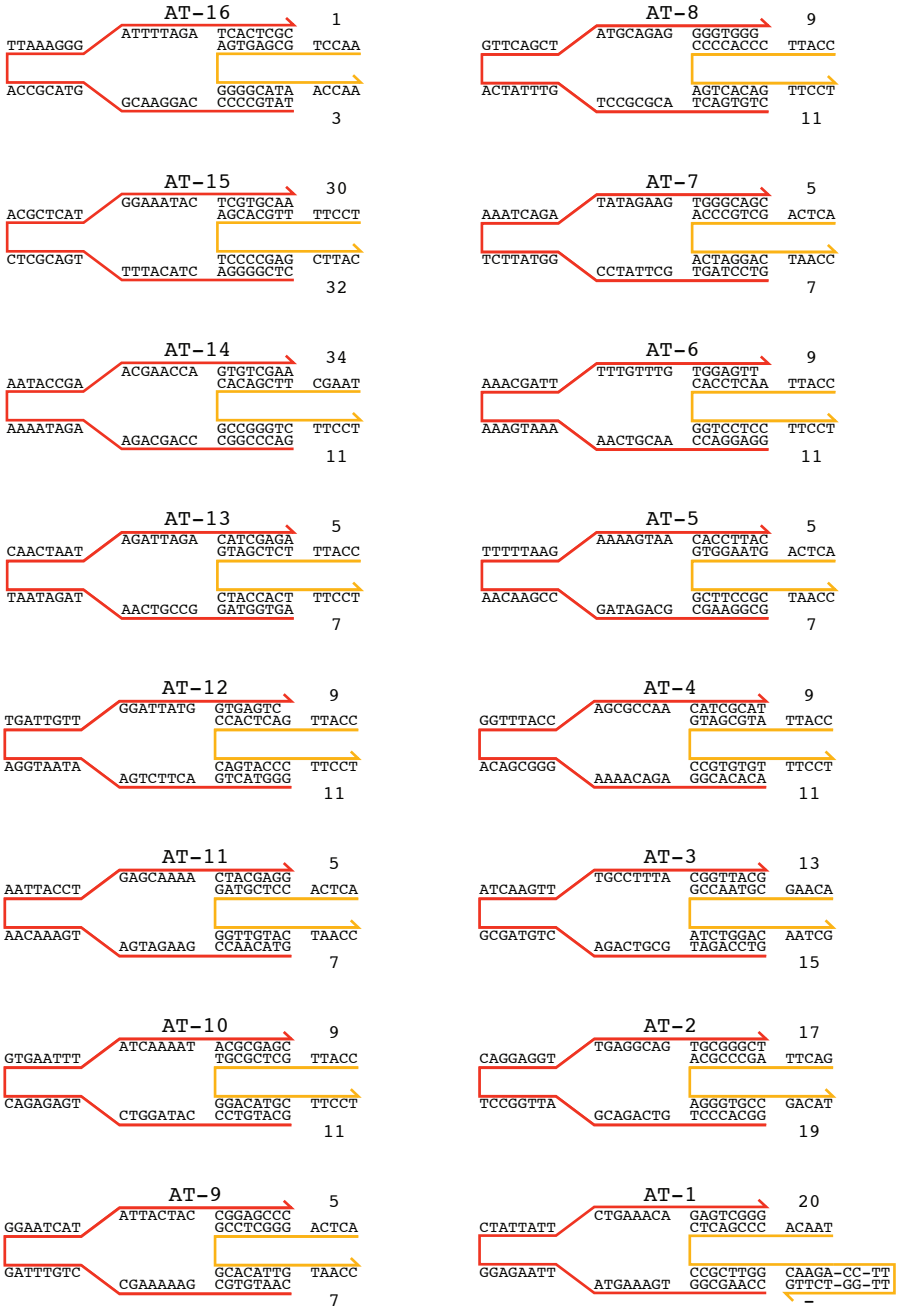

Figure A.3: Adapter tiles for the binary counter system.

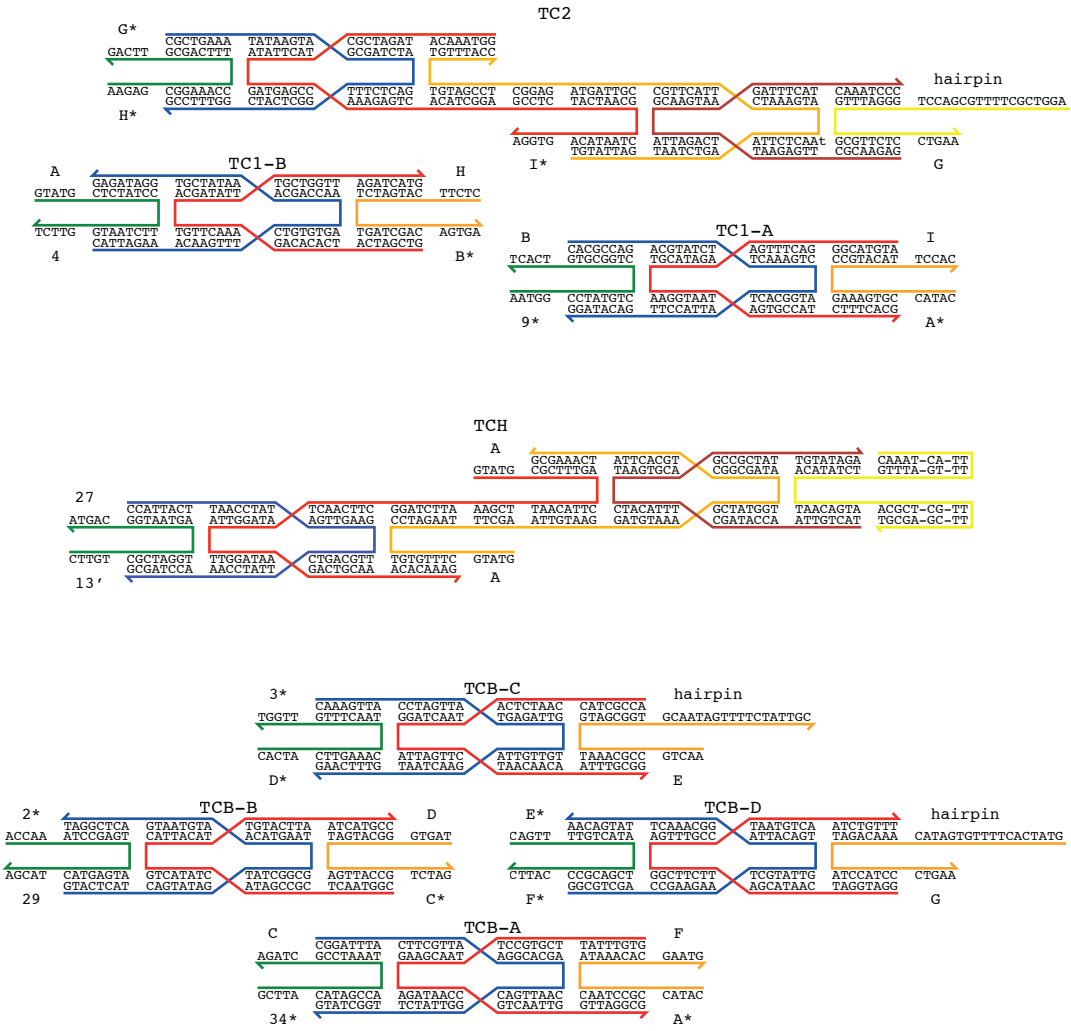

Figure A.4: Capping tiles for the binary counter system.

A.3 Other System Strands

Guard Strands

Guard strands were the Watson-Crick complements of the following strands (with sequences numbered left-to-right in the preceding diagrams): THB-D-4, THB-B-4, TH-6, T0-D-4, T0-B-4, T1-D-4, T1-B-4, TL-6, T0P-B-4, T1P-B-4, TLB-D-4, TC1-A-4, TC2-6, TCB-B-4, and TCB-D-4.