

Thermodynamic Binding Networks

David Doty^{1*}, Trent A. Rogers^{2**}, David Soloveichik^{3***}, Chris Thachuk^{4†},
and Damien Woods^{5‡}

¹ University of California, Davis, ²University of Arkansas, ³University of Texas at Austin, ³California Institute of Technology, ⁵Inria

Abstract. Strand displacement and tile assembly systems are designed to follow prescribed kinetic rules (i.e., exhibit a specific time-evolution). However, the expected behavior in the limit of infinite time—known as thermodynamic equilibrium—is often incompatible with the desired computation. Basic physical chemistry implicates this inconsistency as a source of unavoidable error. Can the thermodynamic equilibrium be made consistent with the desired computational pathway? In order to formally study this question, we introduce a new model of molecular computing in which computation is driven by the thermodynamic driving forces of enthalpy and entropy. To ensure greatest generality we do not assume that there are any constraints imposed by geometry and treat monomers as unstructured collections of binding sites. In this model we design Boolean AND/OR formulas, as well as a self-assembling binary counter, where the thermodynamically favored states are exactly the desired final output configurations. Though inspired by DNA nanotechnology, the model is sufficiently general to apply to a wide variety of chemical systems.

1 Introduction

Most of the models of computing that have come to prominence in molecular programming are essentially kinetic. For example, models of DNA strand displacement cascades and algorithmic tile assembly formalize desired interaction rules followed by certain chemical systems over time [9, 13]. Basing molecular computation on kinetics is not surprising given that computation itself is ordinarily viewed as a *process*. However, unlike electronic computation, where

* doty@ucdavis.edu. Supported by NSF grant CCF-1619343.

** tar003@email.uark.edu Supported by the NSF Graduate Research Fellowship Program under Grant No. DGE-1450079, NSF Grant CAREER-1553166, and NSF Grant CCF-1422152.

*** david.soloveichik@utexas.edu. Supported by NSF grants CCF-1618895 and CCF-1652824.

† thachuk@caltech.edu. Supported by NSF grant CCF-1317694.

‡ damien.woods@inria.fr. Part of this work was carried out at California Institute of Technology. Supported by Inria (France) as well as National Science Foundation (USA) grants CCF-1219274, CCF-1162589, CCF-1317694.

thermodynamics holds little sway, chemical systems operate in a Brownian environment [2]. If the desired output happens to be a meta-stable configuration, then thermodynamic driving forces will inexorably drive the system toward error. For example, *leak* in most strand displacement systems occurs because the thermodynamic equilibrium of a strand displacement cascade favors incorrect over the correct output, or does not discriminate between the two [12]. In DNA tile assembly, we typically must find and exploit kinetic barriers to unseeded growth to enforce that growth happens only from seed assemblies, otherwise thermodynamically favored assemblies will quickly form that are not the intended self-assembly program execution from the seed/input [1, 11].

We introduce the Thermodynamic Binding Networks (TBN) model, where information processing is due entirely to the thermodynamic tradeoff between *entropy* and *enthalpy*, and not any particular reaction pathway. In most experimental systems considered in DNA nanotechnology, thermodynamic favorability is determined by a tradeoff between: (1) the number of base pairs formed or broken (all else being equal, a state with more base pairs bound is more favorable); (2) the number of separate complexes (all else being equal, a state with more free complexes is more favorable). We use the terms enthalpy and entropy to describe (1) and (2) respectively (although this use does not perfectly align with their physical definitions, see Section 2). Intuitively, the entropic benefit of configurations with more separate complexes is due to additional microstates, each describing the independent three-dimensional positions of each complex. Although the general case of a quantitative trade-off between enthalpy and entropy is complex, we develop an elegant formulation based on the limiting case in which enthalpy is infinitely more favorable than entropy. Intuitively, this limit corresponds to increasing the strength of binding, while diluting (increasing the volume), such that the ratio of binding to unbinding rate goes to infinity. Systems studied in molecular programming can in principle be engineered to arbitrarily approach this limit. Indeed, this is the regime previously studied in the context of leak reduction for strand displacement cascades [12]. Figure 1 shows a simple TBN, which can exist in 9 possible binding configurations. The favored (stable) configuration is the one that, among the maximally bound ones (bottom row), maximizes the number of separate complexes (bottom right).

As a central choice in seeking a general theory, we *dispense with geometry*: formally, we treat monomers simply as multisets of binding sites (domains). Viewed in the context of strand displacement, this abstracts away secondary structure (the order of domains on a strand), allowing us to represent arbitrary molecular arrangements such as pseudoknots [4], and handle non-local error modes such as spurious remote toeholds [5]. In the context of tile self-assembly, we consider configurations in which binding does not follow the typical regular lattice structure. Since the TBN model does not rely on geometric constraints to enforce correct behavior, showing that specific undesired behavior is prevented by enthalpy and entropy alone leads to a stronger guarantee. Thus, for example proving leaklessness in this model would imply that even if pseudoknots, or other typically disallowed structures form, we would still have little leak. Indeed, by

casting aside the vagaries of DNA biophysics (e.g., persistence length, number of bases per turn, sequence dependence on binding strength, etc.), our aim is to develop a general theory of programmable systems based on molecular bonds, a theory that will apply to bonds based on other substrates such as proteins, base stacking, or electric charge.

After introducing the TBN model in Section 2, we give results on Boolean circuit-based and self-assembly-based computation. In Section 3 we show how to construct AND and OR gates where the thermodynamically favored configurations encode the output. We develop provable guarantees on the entropic penalty that must be overcome to produce an incorrect 1 output, showing how the logic gates can be designed to make the penalty arbitrarily large. Although completely modular reasoning seems particularly tough in this model, we develop a proof technique based on logically excising domains to handle the composition of Boolean gates—specifically trees of AND gates. Further work is needed to generalize these results to arbitrary circuits.

In Section 4 we look at self-assembly, beginning with questions about large assemblies. On the one hand we exhibit a class of TBNs with thermodynamically stable assemblies (with simple ‘tree’ connectivity) of size exponential in the number of constituent monomer types. On the other hand, we show that this bound is essentially tight by giving an exponential size upper bound on the size of stable assemblies in general. These self-assembly results, along with the binary counter result below, tell us that *monomer-efficient* self-assembly is indeed possible within this model, but that (somewhat surprisingly for a model that favors enthalpy infinitely over entropy) super-exponential size polymers are necessarily unstable, even if they are self-assemblable in kinetic-based models.

For clarity of thought in separating the computational power of thermodynamics and kinetics, throughout much of this paper we do not identify any particular kinetic pathway leading to the desired TBN stable state. Of course real-world physical systems do not operate at thermodynamic equilibrium, and might take longer than the lifetime of the universe to get there. Thus, for such ‘kinetically trapped’ systems, encoding desired output in thermodynamic equilibrium is not enough by itself. To address this, in Section E we give a kinetically and thermodynamically favoured binary counter that assembles in both the abstract Tile Assembly Model and the TBN model. Similarly, the strand displacement AND gate from ref. [12] can be shown to compute correctly in the TBN model [3]. Nonetheless, more work is needed to come up with TBN schemes that have fast kinetic pathways, in addition to the provable thermodynamic guarantees.

2 Model

Let $\mathbb{N}, \mathbb{Z}, \mathbb{Z}^+$ denote the set of nonnegative integers, integers, and positive integers, respectively. A key type of object in our definitions is a multiset, which we define in a few different ways as convenient. Let \mathcal{A} be a finite set. We can define a multiset over \mathcal{A} using the standard set notion, e.g., $\mathbf{c} = \{a, a, c\}$, where $a, c \in \mathcal{A}$. Formally, we view multiset \mathbf{c} as a vector assigning counts to \mathcal{A} . Let-

ting $\mathbb{N}^{\mathcal{A}}$ denote the set of functions $f : \mathcal{A} \rightarrow \mathbb{N}$, we have $\mathbf{c} \in \mathbb{N}^{\mathcal{A}}$. We index entries by elements of $a \in \mathcal{A}$, calling $\mathbf{c}(a) \in \mathbb{N}$ the *count of a in \mathbf{c}* . Fixing some arbitrary ordering on the elements of $\mathcal{A} = \{a_1, a_2, \dots, a_k\}$, we may equivalently view \mathbf{c} as an element of \mathbb{N}^k , where for $i \in \{1, 2, \dots, k\}$, $\mathbf{c}(i)$ denotes $\mathbf{c}(a_i)$. Let $\|\mathbf{c}\| = \|\mathbf{c}\|_1 = \sum_{a \in \mathcal{A}} \mathbf{c}(a)$ denote the *size* of \mathbf{c} . For any vector or matrix \mathbf{c} , let $\text{amax}(\mathbf{c})$ denote the largest absolute value of any component of \mathbf{c} .

We model molecular bonds with precise binding specificity abstractly as binding “domains”, designed to bind only to other, specific binding domains. Formally, consider a finite set \mathcal{D} of *primary domain types*. Each primary domain type $a \in \mathcal{D}$ is mapped to a *complementary domain type* (a.k.a., *codomain type*) denoted a^* . Let $\mathcal{D}^* = \{a^* \mid a \in \mathcal{D}\}$ denote the set of codomain types of \mathcal{D} . The mapping is assumed 1-1, so $|\mathcal{D}^*| = |\mathcal{D}|$. We assume that a domain of primary type $a \in \mathcal{D}$ binds only to its corresponding complementary type $a^* \in \mathcal{D}^*$, and vice versa.¹ The set $\mathcal{D} \cup \mathcal{D}^*$ is the set of *domain types*.

We assume a finite set \mathcal{M} of *monomer types*, where a monomer type $\mathbf{m} \in \mathbb{N}^{\mathcal{D} \cup \mathcal{D}^*}$ is a non-empty multiset of domain types, e.g., $\mathbf{m} = \{a, b, b, c^*, a^*\}$ with $a, b, c \in \mathcal{D}$ being primary domain types. A *thermodynamic binding network* (TBN) is a pair $\mathcal{T} = (\mathcal{D}, \mathcal{M})$ consisting of a finite set \mathcal{D} of primary domain types and a finite set $\mathcal{M} \subset \mathbb{N}^{\mathcal{D} \cup \mathcal{D}^*}$ of monomer types. A *monomer collection* $\vec{\mathbf{c}} \in \mathbb{N}^{\mathcal{M}}$ of \mathcal{T} is multiset of monomer types; intuitively, $\vec{\mathbf{c}}$ indicates how many of each monomer type from \mathcal{M} there are, but not how they are bound.²

Since one monomer collection usually contains more than one copy of the same domain type, we use the term *domain* to refer to each copy separately.³ We similarly reserve the term *monomer* to refer to a particular instance of a monomer type if a monomer collection has multiple copies of the same monomer type.

A single monomer collection $\vec{\mathbf{c}}$ can take on different configurations depending on how domains in monomers are bound to each other. To formally model configurations, we first need the notion of a bond assignment. Let (U, V, E) be the bipartite graph describing all possible bonds, where U is the multiset of all primary domains in all monomers in $\vec{\mathbf{c}}$, V is the multiset of all codomains in all monomers in $\vec{\mathbf{c}}$, and E is the set of edges between primary domains and their complementary codomains $\{\{u, v\} \mid u \in U, v \in V, v = u^*\}$. A *bond assignment* M is a matching⁴ on (U, V, E) . Then, a *configuration* α of monomer collection $\vec{\mathbf{c}}$ is the (multi)graph $(U \cup V, E_M)$, where the edges E_M describe both the association

¹ That is, we assume *like-unlike* binding such as that found in DNA Watson-Crick base-pairing, as opposed to *like-like* binding such as hydrophobic molecules with an affinity for each other in aqueous solution, or base stacking between the blunt ends of DNA helices [6, 14]. It is not clear the extent to which this choice affects the computational power of our model.

² Because a monomer collection is a multiset of monomer types, each of which is itself a multiset, we distinguish them typographically with an arrow.

³ For instance, the monomer collection shown in Fig. 1 has 2 domains of type a , 2 domains of type b , and 1 domain of type a^* and b^* each.

⁴ A matching of a graph is a subset of edges that share no vertices in common. In our case this enforces that a domain is bound to at most one other domain.




Fig. 1: An example TBN $\mathcal{T} = (\mathcal{D}, \mathcal{M})$. $\mathcal{D} = \{a, b\}$ and $\mathcal{M} = \{\mathbf{m}_1, \mathbf{m}_2, \mathbf{m}_3, \mathbf{m}_4\}$, where monomers $\mathbf{m}_1 = \{a, b\}$, $\mathbf{m}_2 = \{a^*, b^*\}$, $\mathbf{m}_3 = \{a\}$, and $\mathbf{m}_4 = \{b\}$. Note that the order of domains does not matter (thus, $\{a, b\} = \{b, a\}$). There are nine distinct configurations for the monomer collection $\vec{\mathbf{c}} = \{\mathbf{m}_1, \mathbf{m}_2, \mathbf{m}_3, \mathbf{m}_4\}$ consisting of a single copy of each of these monomers. The five in the top row are *not saturated* meaning that they do not maximize the number of bound domains, whereas the four configurations in the bottom row are all *saturated*. In addition to being saturated, the configuration in the bottom right is *stable* as it maximizes the number of separate complexes (3) among all saturated configurations (the other saturated configurations have 2).

of domains within the same monomer, and the bonding due to M . Specifically, for each pair of domains $d_i, d_j \in \mathcal{D} \cup \mathcal{D}^*$ that are part of the same monomer in $\vec{\mathbf{c}}$, let $\{d_i, d_j\} \in E_M$, calling this a *monomer edge*, and for each edge $\{d_i, d_i^*\}$ in the bond assignment M , let $\{d_i, d_i^*\} \in E_M$, calling this a *binding edge*. Let $[\vec{\mathbf{c}}]$ be the set of all configurations of a monomer collection $\vec{\mathbf{c}}$. We say the size of a configuration, written $|\alpha|$, is simply the number of monomers in it.

Another graph that will be useful in describing the connectivity of the monomers, independent of which exact domains are bound, is the *monomer binding graph* $G_\alpha = (V_\alpha, E_\alpha)$, which is obtained by contracting each monomer edge of α . In other words, V_α is the set of monomers in α , with an edge between monomers that share at least one pair of bound domains.

Which configurations are thermodynamically favored over others depends on two properties of a configuration: its bond count and entropy. The *enthalpy* $H(\alpha)$ of a configuration is the number⁵ of binding edges (i.e., the cardinality of the matching M). The *entropy* $S(\alpha)$ of a configuration is the number of connected

⁵ We are assuming bonds are of equal strength (although the definition can be naturally generalized to bonds of different strength).




Fig. 2: **Basic AND gate:** Monomers $i_1 = \{a, b\}$ and $i_2 = \{c, d\}$ represent the input, $o = \{e, f\}$ represents the output, while the remainder are intermediate monomers to implement the logic relating the input to the output. If either or both inputs are missing, then the only stable configuration has the present input monomers free (unbound) and the output monomer o not free (bound). If both input monomers are present, then there are two stable configurations: one with inputs free (and o not free), or the one depicted with o free and both inputs bound.

components of α .⁶ Each connected component is called a *polymer*.⁷ Note that a polymer is itself a configuration, but of a smaller monomer collection $\vec{c}' \leq \vec{c}$. As with all configurations, the size of a polymer is the number of monomers in it.

Intuitively, configurations with higher enthalpy $H(\alpha)$ (more bonds formed) and higher entropy $S(\alpha)$ (more separate complexes) are thermodynamically favored. What happens if there is a conflict between the two? One can imagine capturing a tradeoff between enthalpy and entropy by some linear combination of $H(\alpha)$ and $S(\alpha)$. In DNA nanotechnology applications, the tradeoff can be controlled by increasing the number of nucleotides constituting a binding domain (increasing the weight on $H(\alpha)$), or by decreasing concentration (increasing the weight on $S(\alpha)$).⁸

⁶ Our use of the terms “enthalpy” and “entropy”, and notation H and S is meant to evoke the corresponding physical notions. Note, however, that there are other contributions to physical entropy besides the number of separate complexes. Indeed, the free energy contribution of forming additional bonds typically contains substantial enthalpic and entropic parts.

⁷ We are generalizing the convention for the word “polymer” in the chemistry literature. We have no requirement that a polymer be linear, nor that it consist of repeated subunits. We chose “polymer” rather than “complex” to better contrast with “monomer”.

⁸ In typical DNA nanotechnology applications, the Gibbs free energy $\Delta G(\alpha)$ of a configuration α can be estimated as follows. Bonds correspond to domains of length l bases, and forming each base pair is favorable by ΔG_{bp}^o . Thus, the contribution of

In the rest of this paper, we study the particularly interesting limiting case in which enthalpy is *infinitely* more favorable than entropy.⁹ We say a configuration α is *saturated* if it has no pair of domains d and d^* that are both unbound; this is equivalent to stating that α has maximal bonding among all configurations in $[\vec{\mathbf{c}}]$. We say a configuration $\alpha \in [\vec{\mathbf{c}}]$ is *stable* (aka thermodynamically favored) if it is saturated and maximizes the entropy among all saturated configurations, i.e., every saturated configuration $\alpha' \in [\vec{\mathbf{c}}]$ obeys $S(\alpha') \leq S(\alpha)$. Let $[\vec{\mathbf{c}}]_{\square}$ denote the set of stable configurations of monomer collection $\vec{\mathbf{c}}$. See Fig. 1 for an example thermodynamic binding network that has a single stable configuration. We note that, consistent with our model, in DNA strand displacement cascades “long” domains are assumed to always be paired, and systems can be effectively driven by the formation of more separate complexes [16].

3 Thermodynamic Boolean formulas

Fig. 2 shows an example of a TBN that performs AND computation, based on the CRN strand displacement gate from ref. [12]. Realized as a strand displacement system, it has a kinetic pathway taking the untriggered (left) to the triggered (right) configuration. The inputs are specified by the presence (logical value 1) or absence (logical value 0) of the input monomers \mathbf{i}_1 and \mathbf{i}_2 . The output convention followed is the following. The output is 1 if and only if *some* stable configuration has the output monomer \mathbf{o} unbound to any other monomer (free). This can be termed the *weak* output convention. Alternatively, in the *strong* output convention, output 1 implies *every* stable configuration has the output monomer \mathbf{o} free, and output 0 implies *every* stable configuration has the output monomer \mathbf{o} bound to some other monomer. More complex AND gate designs are compatible with the strong output convention (not shown).

Note that even the weak output convention, coupled with a kinetic pathway releasing the output given the correct inputs, can be used to argue that: (1) if the correct inputs are present the output will be produced (via kinetic argument),

$H(\alpha)$ to $\Delta G(\alpha)$ is $(\Delta G_{\text{bp}}^{\circ} \cdot l)H(\alpha)$. At 1 M, the free energy penalty due to decreasing the number of separate complexes by 1 is $\Delta G_{\text{assoc}}^{\circ}$. At effective concentration C M, this penalty increases to $\Delta G_{\text{assoc}}^{\circ} + RT \ln(1/C)$. As the point of zero free energy, we take the configuration with no bonds, and all monomers separate. Thus, the contribution of $S(\alpha)$ to $\Delta G(\alpha)$ is $(\Delta G_{\text{assoc}}^{\circ} + RT \ln(1/C))(|\alpha| - S(\alpha))$, where $|\alpha|$ is the total number of monomers. To summarize,

$$\Delta G(\alpha) = (\Delta G_{\text{bp}}^{\circ} \cdot l)H(\alpha) + (\Delta G_{\text{assoc}}^{\circ} + RT \ln(1/C))(|\alpha| - S(\alpha)).$$

Note that, as expected, this is a linear combination of $H(\alpha)$ and $S(\alpha)$, and that increasing the length of domains l weighs $H(\alpha)$ more heavily, while decreasing the concentration C weighs $S(\alpha)$ more heavily. Typically $G_{\text{bp}}^{\circ} \approx -1.5$ kcal/mol, and $G_{\text{assoc}}^{\circ} \approx 1.96$ kcal/mol [10].

⁹ Note that the other limiting case, where entropy is infinitely more favorable, is degenerate: the most favorable configuration in that case always has every monomer unconnected to any other.

(2) if the correct inputs are not present then ultimately little output will be free (thermodynamic argument). In the context of strand displacement cascades, TBNs can explore arbitrary structures (pseudoknots, remote toeholds, etc) since we do not impose any ordering on domains in a monomer, nor any geometry. This strengthens the conclusion of (2), showing that arbitrary (even unknown) kinetic pathways must lead to a thermodynamic equilibrium with little output.

While individual AND gates can be proven correct with respect to the above output conventions (e.g., through the SAT solver of ref. [3]), it remains to be shown that these components can be safely composed into arbitrary Boolean circuits. Note that the input and output monomers have orthogonal binding sites. This is important for composing AND gates, where the output of one acts as an input to another. As is typical for strand displacement logic, OR gates can be trivially created when multiple AND gates have the same output. Dual-rail AND/OR circuits are sufficient to compute arbitrary Boolean functions without explicit NOT gates. Nonetheless it is not obvious that the input convention (complete presence or absence of input monomers) matches the output convention (weak or strong). It is also not clear how statements about the stable configurations of the whole circuit can be made based on the stable configurations of the individual modules.

We now show that correct composition can be proven in certain cases. Although we believe that the gate shown in Fig. 2 is composable, the argument below relies on a different construction. We further consider a restricted case of AND gate formulas (trees).

An important concept in the argument below is the notion of “distance to stability”. This refers to the difference between the entropy of the stable configurations and the largest entropy of a saturated configuration with incorrect output. The larger the distance to stability, the larger the entropy penalty to incorrectly producing the output. Unlike the simple AND gate from Fig. 2, the constructions below can be instantiated to achieve arbitrary desired distance to stability (by increasing the redundancy parameter n).

Many open questions remain. Can our techniques be generalized to arbitrary circuits, rather than just trees of AND gates? Can we prove these results for logic gates that have a corresponding kinetic pathway (like the AND gates in Fig. 2 which can be instantiated as strand displacement systems)? Finally, in our Boolean gate constructions, we assume that the monomer collection has exactly one copy of certain monomers. It remains open whether these schemes still work if there are many copies of all monomers.

3.1 Translator cascades

We begin with the simplest of circuits, *translator cascades* ($x_1 \rightarrow x_2 \rightarrow \dots \rightarrow x_{k+1}$), which simply propagate signal through k layers when the input signal x_1 is present. Logically a translator gate is simply a repeater gate. The input is the presence or absence of the input monomer consisting of n copies of domain x_1 . Our analysis below implies that if and only if the input is present, there is a stable configuration with n copies of x_{k+1} domain in the same polymer. The




Fig. 3: A cascade of k translator gates discussed in Section 3.1, with redundancy parameter $n = 3$. We say that a configuration of a formula has output 1 if the terminator monomer $\{x_{k+1}, \dots, x_{k+1}\}$ is free, and has output 0 otherwise. Redundancy parameter n specifies the number of copies of monomers and domains as shown.

terminator gadget converts this output to the weak output convention defined above (whether or not the monomer consisting of n copies of domain x_{k+1} is free). The following Lemma shows that we can exactly compute the distance from stability of a translator cascade shown in Fig. 3. Besides being a “warm-up” for AND gate cascades, the Lemma is used in the proof of Theorem 2.

Observation 1. *The intended configuration α of a monomer collection representing a depth k , redundancy n translator cascade, without input, and with output 0, is saturated and has $S(\alpha) = nk + 1$. (See Fig. 3.)*

Lemma 1. *If γ is a saturated configuration of a monomer collection representing a depth k , redundancy n translator cascade, without input, and with output 1, then $S(\gamma) = n(k - 1) + 2$.*

The proof of Lemma 1 appears in Appendix A. Taken together, Observation 1 and Lemma 1 imply that the redundancy parameter (n) guarantees the distance to stability ($n - 1$) for a translator cascade of any length.

3.2 Trees of AND gates

In this section we motivate how Boolean logic gates can be composed such that the overall circuit has a guaranteed distance to stability, relative to a redundancy parameter n . Specifically, we start with the AND gate design of Fig. 4, and we give a concrete argument for a tree of these AND gates (e.g., Fig. 5).

Theorem 2. *Consider a TBN for AND gates, with redundancy n , composed into a tree of depth k . If at least one of the inputs is not present, the distance to stability for any saturated configurations with output 1 is at least $n - 2k - 1$.*

Proof. Let γ be any saturated configuration of the TBN with output 1. Consider the missing input and define the *leak path* to be the linear sequence of AND gates from the missing input to and including the terminator gadget. For




Fig. 4: AND gates used in Section 3.2, with redundancy parameter $n = 4$. Two saturated configurations are shown: γ_0 is the intended configuration corresponding to output of 0. γ_1 is the intended configuration corresponding to output of 1. Input domains are x_i y_i , and output domains are x_{i+1} . The output is considered to be 1 in any configuration where all n output domains are in the same polymer, 0 otherwise. Dashed boxes represent that any domain type appearing inside of a box does have a complement appearing outside of the box.




Fig. 5: Shown highlighted is a leak path through a tree of AND gates from a missing input (“0”) to erroneous output (“1”).

convenience we imagine relabelling all the domains in the leak path indexed by the position of the AND gate in the leak path. For example, Fig. 5 highlights the leak path through the tree from a missing input (“0”) to erroneous output (“1”). Specifically, the domain names as shown in Fig. 4 appear in the i th AND gate (for $1 \leq i \leq k$), where x_{k+1} feeds into the terminator gadget. Domains y_i connect the leak path to the rest of the tree.

Definition 1. Given a configuration α of a monomer collection \vec{c} , we say we excise a domain d if we create a new configuration α' by removing the node corresponding to d and all incident edges. (Note that α' is a configuration of a monomer collection of a different TBN.)

Manipulation 1. Excise all domains of type y_i and codomains of type y_i^* on monomers of the leak path involved in fan-in, $1 \leq i \leq k$, yielding the new configuration γ' . Note that if domain y_i is on a monomer other than the leak path, then it is not excised.

The leak path in γ' now has no domains in common with the rest of the tree (and thus no bonds). Let γ'_L be the subconfiguration of the leak path, and let γ'_R be the subconfiguration of the rest of the system. (Note $\gamma' = \gamma'_L \cup \gamma'_R$.)

Observation 3. *Given a saturated configuration α , if you excise all domains or codomains of a particular type (or both its domains and codomains) yielding α' , then α' is saturated.*

By Observation 3 γ' is saturated since for every domain type y_i and codomain type y_i^* , every instance of y_i^* is excised; $1 \leq i \leq k$. This implies γ'_L and γ'_R are also saturated.

Manipulation 2. Excise all domains of type a_i and b_i and all codomains of type a_i^* and b_i^* in γ'_L , $1 \leq i \leq k$, yielding the new configuration γ''_L . By Observation 3, γ''_L is saturated.

Claim 4. $S(\gamma') \geq S(\gamma)$.

Proof of the claim. Entropy can only be decreased via excision if an entire monomer is excised. Since Manipulation 1 only excised domain and codomain types from the set $\mathcal{D}' = \bigcup_{i=1}^k \{y_i, y_i^*\}$, and those domain types only appear on monomers which also have domain instances with types not in \mathcal{D}' , then no entire monomer was excised. ■

Claim 5. $S(\gamma''_L) \geq S(\gamma'_L) - 3k$.

Proof of the claim. For every layer i , $1 \leq i \leq k$, there are 3 monomers that only contain domain and codomain types in the set $\{a_i, b_i, a_i^*, b_i^*\}$. Therefore, γ''_L contains at most 3 fewer monomers than γ'_L , for each of the k layers. ■

Claim 6. $S(\gamma''_L) = n(k-1) + 2$.

Proof of the claim. Recognize that γ''_L is a saturated configuration of a monomer collection representing a depth k , redundancy n translator cascade, without input, and with output 1. The claim follows by Lemma 1. ■

Claim 7. $S(\gamma) \leq n(k-1) + 2 + S(\gamma'_R) + 3k$.

Proof of the claim.

$$\begin{aligned}
S(\gamma) &\leq S(\gamma') && \text{by Claim 4} \\
&= S(\gamma'_L) + S(\gamma'_R) \\
&\leq S(\gamma''_L) + S(\gamma'_R) + 3k && \text{by Claim 5} \\
&\leq n(k-1) + 2 + S(\gamma'_R) + 3k && \text{by Claim 6}
\end{aligned}$$

■

Now, take the monomers from the leak path in γ , and configure them into the “untriggered configuration” (see Fig. 4, left), yielding subconfiguration β . Let $\alpha = \beta \cup \gamma'_R$. Note that β is saturated, and therefore α is a saturated configuration of the entire tree (*i.e.*, the same TBN as γ).

Observation 8. $S(\alpha) = S(\gamma'_R) + k(n+1) + 1$.

Finally, consider the entropy gap between α and γ .

$$\begin{aligned}
S(\alpha) - S(\gamma) &\geq S(\gamma'_R) + k(n+1) + 1 - S(\gamma) && \text{by Observation 8} \\
&\geq S(\gamma'_R) + k(n+1) + 1 \\
&\quad - (n(k-1) + 2 + S(\gamma'_R) + 3k) && \text{by Claim 7} \\
&= n - 2k - 1
\end{aligned}$$

Therefore, there exists a saturated configuration with output 0 over the same TBN as γ , but with entropy at least $n - 2k - 1$ larger, thus establishing the theorem. \square

Theorem 2 seems to suggest that in order to maintain the bound on distance to stability for incorrect computation, the redundancy parameter n should increase to compensate for an increase in circuit depth k . However, a more sophisticated argument shows that manipulations 1 and 2 can decrease entropy by at most $k + 1$. Following the above argument, the distance to stability is found to be $n - 2$. This is optimal because a single AND gate with redundancy $n = 2$ can be shown to have no entropy gap between output 0 and output 1 configurations.

4 Thermodynamic self-assembly: Assembling large polymers

TBNs can not only exhibit Boolean circuit computation, but they can also be thought of as a model of self-assembly. Here we begin to explore this connection by asking a basic question motivated by the abstract Tile Assembly Model (aTAM) [13]: how many different monomer *types* are required to assemble a large polymer?

Favoring enthalpy infinitely over entropy, on its face, *appears* to encourage large polymers. Perhaps we can imagine designing a single TBN \mathcal{T} that can assemble arbitrarily large polymers where for each $n \in \mathbb{N}$, \mathcal{T} has a stable polymer α composed of at least n monomers. In this section we show that this is impossible: every TBN $\mathcal{T} = (\mathcal{D}, \mathcal{M})$ has stable polymers of size at most exponential in the number of domain types $|\mathcal{D}|$ and monomer types $|\mathcal{M}|$ (Theorem 13). The proof shows that any polymer ρ larger than the bound can be partitioned into at least two saturated (maximally bound) polymers, which implies that ρ is not stable. Fig. 6 gives an example. We also show that this upper bound is essentially tight by constructing a family of systems with exponentially large stable polymers (Theorem 9). Taken together, the exponential lower bound of Theorem 9 and




Fig. 6: A polymer ρ composed of several copies of four monomer types, which is not stable since it can be broken into several smaller polymers (bottom panel) such that all domains are bound.




Fig. 7: An example of a TBN from Theorem 9 for $n = 4$ and $k = 2$.

upper bound of Theorem 13 give a relatively tight bound on the maximum size achievable for stable TBN polymers.

Is it possible to construct algorithmically *interesting* TBN polymers that are stable? In the full version of this paper, we show that a typical binary counter construction from the aTAM model is not stable, but can be modified to become stable in our model. Importantly, this TBN binary counter demonstrates that in principle algorithmically complex assemblies could have effective assembly pathways (aTAM) as well as be thermodynamically stable (TBN).

4.1 Superlative trees: TBNs with exponentially large stable polymers

The next theorem shows that there are stable polymers that are exponentially larger than the number of domain types and monomer types required to assemble them.

Theorem 9. *For every $n, k \in \mathbb{Z}^+$, there is a TBN $\mathcal{T} = (\mathcal{D}, \mathcal{M})$ with $|\mathcal{D}| = n - 1$ and $|\mathcal{M}| = n$, having a stable polymer of size $\frac{k^n - 1}{k - 1}$.*

Proof. An example of \mathcal{T} for $n = 4$ and $k = 2$ is shown in Fig. 7. Let $\mathcal{D} = \{d_1, \dots, d_n\}$ and $\mathcal{M} = \{\mathbf{m}_1, \dots, \mathbf{m}_n\}$, where, for each $j \in \{2, \dots, n - 1\}$, $\mathbf{m}_j = \{d_{j-1}^*, k \cdot d_j\}$ (i.e., 1 copy of d_{j-1}^* and k copies of d_j), $\mathbf{m}_1 = \{k \cdot d_1\}$, and $\mathbf{m}_n = \{d_{n-1}^*\}$. Define $\vec{\mathbf{c}} \in \mathbb{N}^{\mathcal{M}}$ by $\vec{\mathbf{c}}(\mathbf{m}_j) = k^{j-1}$ for $j \in \{1, \dots, n\}$. Then $\|\vec{\mathbf{c}}\| = \sum_{j=1}^n k^{j-1} = \frac{k^n - 1}{k - 1}$. Observe that $[\vec{\mathbf{c}}]$ has a unique (up to isomorphism) saturated configuration α (which is therefore stable), described by a complete k -ary tree: level $j \in \{1, \dots, n - 1\}$ of the tree is composed of k^{j-1} copies of \mathbf{m}_j , each bound to k children of type \mathbf{m}_{j+1} in level $j + 1$. \square

The remainder of Section 4 is devoted to proving that no stable polymer ρ can have size *more* than exponential in $|\mathcal{D}|$ and $|\mathcal{M}|$. We note that this is easy to prove in the special case if the monomer binding graph G_ρ is a tree, such as the TBN in Theorem 9 in Appendix C.

4.2 A linear algebra framework

We prove Theorem 13, the main result of Section 4, by viewing TBNs from a linear algebra perspective. Let $\mathcal{T} = (\mathcal{D}, \mathcal{M})$ be a TBN, with $\mathcal{D} = \{d_1, \dots, d_d\}$ and $\mathcal{M} = \{\mathbf{m}_1, \dots, \mathbf{m}_m\}$. For a matrix \mathbf{A} , let $\mathbf{A}(i, j)$ denote the entry in the i 'th row and j 'th column. Define the $d \times m$ *positive monomer matrix* $\mathbf{M}_{\mathcal{T}}^+$ of \mathcal{T} by $\mathbf{M}_{\mathcal{T}}^+(i, j) = \mathbf{m}_j(d_i)$. Define the $d \times m$ *negative monomer matrix* $\mathbf{M}_{\mathcal{T}}^-$ of \mathcal{T} by $\mathbf{M}_{\mathcal{T}}^-(i, j) = \mathbf{m}_j(d_i^*)$. Define the $d \times m$ *monomer matrix* $\mathbf{M}_{\mathcal{T}}$ of \mathcal{T} to be $\mathbf{M}_{\mathcal{T}} = \mathbf{M}_{\mathcal{T}}^+ - \mathbf{M}_{\mathcal{T}}^-$. Note that $\mathbf{M}_{\mathcal{T}}^+$ and $\mathbf{M}_{\mathcal{T}}^-$ are matrices over \mathbb{N} , but $\mathbf{M}_{\mathcal{T}}$ is over \mathbb{Z} .

The rows of the monomer matrix $\mathbf{M}_{\mathcal{T}}$ correspond to domain types and the columns correspond to monomer types. The mapping from a TBN \mathcal{T} to a monomer matrix $\mathbf{M}_{\mathcal{T}}$ is not 1-1: $\mathbf{M}_{\mathcal{T}}(i, j)$ is the number of d_i domains minus the number of d_i^* domains in monomer type \mathbf{m}_j , which would be the same, for instance, for monomer types $\mathbf{m}_1 = \{d_1, d_3\}$ and $\mathbf{m}_2 = \{d_1, d_1^*, d_3\}$. Let $\vec{\mathbf{c}}$ be a monomer collection and let $\mathbf{d} = \mathbf{M}_{\mathcal{T}} \vec{\mathbf{c}} \in \mathbb{N}^d$; for $i \in \{1, \dots, d\}$, $\mathbf{d}(i)$ is the number of d_i domains minus the number of d_i^* domains in the whole monomer collection $\vec{\mathbf{c}}$.

Let $\alpha \in [\vec{\mathbf{c}}]$ be saturated; α can only have a domain d_i unbound if all copies of its complement d_i^* are bound, and vice versa. If $\mathbf{d}(i) > 0$, in α there is an excess of d_i domains, and all d_i^* domains are bound. If $\mathbf{d}(i) < 0$, in α there is an excess of d_i^* domains, and all d_i domains are bound. This leads to the following observation.

Observation 10. *Let $\mathcal{T} = (\mathcal{D}, \mathcal{M})$ be a TBN and $\vec{\mathbf{c}} \in \mathbb{N}^{\mathcal{M}}$ a monomer collection. Let $\mathbf{d} = \mathbf{M}_{\mathcal{T}} \vec{\mathbf{c}}$. Then for every configuration $\alpha \in [\vec{\mathbf{c}}]$, α is saturated if and only if, for all $i \in \{1, \dots, d\}$, if $\mathbf{d}(i) \geq 0$ (respectively, if $\mathbf{d}(i) \leq 0$), then $\mathbf{d}(i)$ is the number of unbound d_i (resp., d_i^*) domains in α .*

Let $\mathcal{T} = (\mathcal{D}, \mathcal{M})$ and $\mathcal{T}' = (\mathcal{D}, \mathcal{M}')$ be TBNs with the same set of domain types. Then we call \mathcal{T}' a *relabeling* of \mathcal{T} if there exists a subset $D \subseteq \mathcal{D}$ such that \mathcal{M}' can be obtained from \mathcal{M} by starring any instance of $d_i \in D$ in \mathcal{M} and unstarring any instance of d_i^* in \mathcal{M} . Since this corresponds to negating the i 'th row of $\mathbf{M}_{\mathcal{T}}$, which negates the i 'th entry of the vector $\mathbf{d} = \mathbf{M}_{\mathcal{T}} \vec{\mathbf{c}}$, this gives the following observation.

Observation 11. *Let $\mathcal{T} = (\mathcal{D}, \mathcal{M})$ be a TBN and $\vec{\mathbf{c}} \in \mathbb{N}^{\mathcal{M}}$ a monomer collection. There exists a relabeling \mathcal{T}' of \mathcal{T} so that $\mathbf{M}_{\mathcal{T}'} \vec{\mathbf{c}} \geq 0$.*

Combining Observations 10 and 11 results in the following observation, which essentially states that for any given monomer collection $\vec{\mathbf{c}}$, we may assume without loss of generality that domains unbound in saturated configurations $\alpha \in [\vec{\mathbf{c}}]$ are all primary domain types.

Observation 12. *Let $\mathcal{T} = (\mathcal{D}, \mathcal{M})$ be a TBN and $\vec{\mathbf{c}} \in \mathbb{N}^{\mathcal{M}}$ a monomer collection. There exists a relabeling \mathcal{T}' of \mathcal{T} so that, letting $\mathbf{d} = \mathbf{M}_{\mathcal{T}'} \vec{\mathbf{c}}$, for all configurations $\alpha \in [\vec{\mathbf{c}}]$, α is saturated if and only if, for all $i \in \{1, \dots, d\}$, $\mathbf{d}(i) \in \mathbb{N}$ is the number of unbound primary domains of type $d_i \in \mathcal{D}$ in α .*

The following lemma is a key technical tool for showing that a polymer is not stable (or equivalently that a stable configuration has entropy greater than 1 and therefore cannot be a single polymer). It generalizes the idea shown in Fig. 6 that if one can find a monomer subcollection $\vec{\mathbf{c}}_1$ in a larger collection $\vec{\mathbf{c}}$, and $\vec{\mathbf{c}}_1$ has a saturated configuration with *no* bonds left unbound, then one can create a saturated configuration $\gamma \in [\vec{\mathbf{c}}]$ with no bonds between $\vec{\mathbf{c}}_1$ and $\vec{\mathbf{c}} - \vec{\mathbf{c}}_1$. (Thus γ has at least two polymers.)

More generally, given a monomer collection $\vec{\mathbf{c}}$ with at least as many d_i as d_i^* domains (under appropriate relabeling this holds for each i by Observation 11), if we can partition $\vec{\mathbf{c}}$ into subcollections $\vec{\mathbf{c}}_1$ and $\vec{\mathbf{c}}_2$, and each of them *also* has at least as many d_i as d_i^* domains for each $i \in \{1, \dots, d\}$, then every stable configuration $\alpha \in [\vec{\mathbf{c}}]_{\square}$ has at least two polymers, since there is a saturated configuration of $\vec{\mathbf{c}}$ in which there are no bonds between $\vec{\mathbf{c}}_1$ and $\vec{\mathbf{c}}_2$.¹⁰

Lemma 2. *Let $\mathcal{T} = (\mathcal{D}, \mathcal{M})$ be a TBN, let $\vec{\mathbf{c}} \in \mathbb{N}^{\mathcal{M}}$ be a monomer collection of \mathcal{T} such that $\mathbf{M}_{\mathcal{T}}\vec{\mathbf{c}} \geq \mathbf{0}$, and let $\alpha \in [\vec{\mathbf{c}}]_{\square}$ be a stable configuration. If there exist nonempty subcollections $\vec{\mathbf{c}}_1, \vec{\mathbf{c}}_2 \in \mathbb{N}^{\mathcal{M}}$ where 1) $\vec{\mathbf{c}}_1 + \vec{\mathbf{c}}_2 = \vec{\mathbf{c}}$ and 2) $\mathbf{M}_{\mathcal{T}}\vec{\mathbf{c}}_1 \geq \mathbf{0}$ and $\mathbf{M}_{\mathcal{T}}\vec{\mathbf{c}}_2 \geq \mathbf{0}$, then $S(\alpha) > 1$.*

The proof of Lemma 2 appears in Appendix B.

4.3 Exponential upper bound on polymer size

We now show a converse to Theorem 9, namely Theorem 13, showing that stable polymers have size at most exponential in the number of domain and monomer types. The proof of Theorem 13 closely follows Papadimitriou’s proof that integer programming is contained in NP [7]. That proof shows, for any linear system $\mathbf{A}\mathbf{x} = \mathbf{b}$, where \mathbf{A} is a given $n \times m$ integer matrix, $\mathbf{b} \in \mathbb{Z}^n$ is a given integer vector, and \mathbf{x} represents the m unknowns, that if the system has a solution $\mathbf{x} \in \mathbb{N}^m$, then it has a “small” solution $\mathbf{x}' \in \mathbb{N}^m$. “Small” means that $\text{amax}(\mathbf{x}')$ is at most exponential in $n + m + \text{amax}(\mathbf{A}) + \text{amax}(\mathbf{b})$. The technique of [7] proceeds by showing that any sufficiently large solution $\mathbf{x} \in \mathbb{N}^m \setminus \{\mathbf{0}\}$ can be split into two vectors $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{N}^m \setminus \{\mathbf{0}\}$ such that $\mathbf{x}_1 + \mathbf{x}_2 = \mathbf{x}$, where $\mathbf{A}\mathbf{x}_1 = \mathbf{0}$, so \mathbf{x}_2 is also a solution: $\mathbf{A}\mathbf{x}_2 = \mathbf{A}(\mathbf{x} - \mathbf{x}_1) = \mathbf{A}\mathbf{x} - \mathbf{A}\mathbf{x}_1 = \mathbf{A}\mathbf{x} = \mathbf{b}$. This is useful because \mathbf{x}_1 and \mathbf{x}_2 satisfy the hypothesis of Lemma 2, which tells us that all stable configurations $\alpha \in [\mathbf{x}]$ obey $S(\alpha) > 1$, so any single-polymer configuration of \mathbf{x} is not stable.

We include the full proof for three reasons: 1) self-containment, 2) it requires a bit of care to convert our inequality $\mathbf{A}\mathbf{x} \geq \mathbf{0}$ into an equality as needed for the

¹⁰ Observations 10, 11, and 12 are not really *necessary* for our technique, but simplify the description of the conditions under which $\vec{\mathbf{c}}_1$ and $\vec{\mathbf{c}}_2$ would be saturated: specifically, that if $\mathbf{d} = \mathbf{M}_{\mathcal{T}}\vec{\mathbf{c}}$ is in the nonnegative orthant, then so are $\mathbf{d}_1 = \mathbf{M}_{\mathcal{T}}\vec{\mathbf{c}}_1$ and $\mathbf{d}_2 = \mathbf{M}_{\mathcal{T}}\vec{\mathbf{c}}_2$. If we did not use relabeling (thus could not guarantee that \mathbf{d} is in the nonnegative orthant) then the requisite condition to apply Lemma 2 would be that \mathbf{d} , \mathbf{d}_1 , and \mathbf{d}_2 all occupy the same orthant; i.e., for all $i \in \{1, \dots, d\}$, if any of $\mathbf{d}(i)$, $\mathbf{d}_1(i)$, or $\mathbf{d}_2(i)$ are negative, then the other two are not positive.

technique,¹¹ and 3) although the proof of [7] is sufficiently detailed to prove our theorem, the statement of the theorem in [7] hides the details about splitting the vector, which are crucial to obtaining our result.

We require the following discrete variant of Farkas' Lemma, also proven in [7].

Lemma 3 ([7]). *Let $a, d, l \in \mathbb{Z}^+$, $\mathbf{v}_1, \dots, \mathbf{v}_l \in \{0, \pm 1, \dots, \pm a\}^d$, and $K = (ad)^{d+1}$. Then exactly one of the following statements holds:*

1. *There exist l integers $n_1, \dots, n_l \in \{0, 1, \dots, K\}$, not all 0, such that $\sum_{j=1}^l n_j \mathbf{v}_j = \mathbf{0}$.*
2. *There exists a vector $\mathbf{h} \in \{0, \pm 1, \dots, \pm K\}^d$ such that, for all $j \in \{1, \dots, l\}$, $\mathbf{h}^\top \cdot \mathbf{v}_j \geq 1$.*

Intuitively, statement (1) of Lemma 3 states that the vectors can be added to get $\mathbf{0}$ (they are “directions of balanced forces” [7]). This is false if and only if statement (2) holds: the vectors all lie on one side of some hyperplane, whose orthogonal vector \mathbf{h} would then have positive dot product with each of the vectors \mathbf{v}_j (thus adding any of them would move positively in the direction \mathbf{h} and could never cancel to get $\mathbf{0}$).

Intuitively, Theorem 13 states that the size of polymers in stable configurations is upper bounded by a function which is exponential in d . We prove this by first defining a constant K which is exponential in d . If each of the m individual monomer counts is less than K , then we are done since no polymer in the configuration can have size bigger than mK . If some of the monomer counts are greater than K (call these *large-count monomers*), we consider two cases.

For the first case, we consider the scenario where the vectors which describe the monomer types with large monomer counts are such that they can “balance” each other out with relatively small linear combination coefficients. If this is the case, then we can make a saturated subconfiguration which has at least one polymer using these small linear combination coefficients and large-count monomer types since the domains and codomains completely “balance” each other out. We can then use the rest of the counts of the configuration to make another saturated subconfiguration which has at least one polymer. This is shown mathematically by applying Lemma 3 to show that the monomer counts in the polymer can be split to find a configuration consisting of two separate saturated polymers. This means that there is a saturated configuration that has at least two polymers which contradicts the assumption α is a single stable polymer.

If there exist no such linear combination to “balance out” out the vectors describing the large-count monomers, then Lemma 3 tells us all of these vectors lie on the same side of some hyperplane. In this case, we show that counts of

¹¹ In particular, the proof of [7] upper bounds the size of \mathbf{x} in terms of the entries of both \mathbf{A} and \mathbf{b} . However, the naïve way to solve a linear inequality $\mathbf{A}\mathbf{x} \geq \mathbf{0}$ using an equality, by introducing slack variables \mathbf{b} and asking for solutions $\mathbf{x} \in \mathbb{N}^m$, $\mathbf{b} \in \mathbb{N}^n$ such that $\mathbf{A}\mathbf{x} = \mathbf{b}$, allows for the possibility that $\|\mathbf{b}\|$ is very large compared to $\|\mathbf{A}\|$, in which case upper bounding $\|\mathbf{x}\|$ in terms of both \mathbf{A} and \mathbf{b} does not help to bound $\|\mathbf{x}\|$ in terms of \mathbf{A} alone.

the small-count monomers play a role in bounding the counts of the large-count monomers. Intuitively, if all of the vectors describing the large-count monomers lie on the same side of some hyperplane, they are missing domains and codomains which will allow them to bind together. The domains and codomains they need in order to bind together, then must be found on the small-count monomer. Consequently, this means the size of polymers will be bound by the counts of small-count monomers (which is exponential in K).

The proof of the following theorem is in Appendix D.

Theorem 13. *Let $\mathcal{T} = (\mathcal{D}, \mathcal{M})$ be a TBN with $d = |\mathcal{D}|$ and $m = |\mathcal{M}|$. Let $a = \max_{\mathbf{m} \in \mathcal{M}, d_i \in \mathcal{D} \cup \mathcal{D}^*} \mathbf{m}(d_i)$ be the maximum count of any domain in any monomer. Then all polymers of every stable configuration α of \mathcal{T} have size at most $2(m + d)(ad)^{2d+3}$.*

References

1. Robert D Barish, Rebecca Schulman, Paul WK Rothmund, and Erik Winfree. An information-bearing seed for nucleating algorithmic self-assembly. *Proceedings of the National Academy of Sciences*, 106(15):6054–6059, 2009.
2. Charles H Bennett. The thermodynamics of computation—a review. *International Journal of Theoretical Physics*, 21(12):905–940, 1982.
3. Keenan Breik, Lakshmi Prakash, Chris Thachuk, Marijn Heule, and David Soloveichik. Computing properties of stable configurations of thermodynamic binding networks. *In preparation*, 2017.
4. Robert M Dirks, Justin S Bois, Joseph M Schaeffer, Erik Winfree, and Niles A Pierce. Thermodynamic analysis of interacting nucleic acid strands. *SIAM review*, 49(1):65–88, 2007.
5. Anthony J Genot, David Yu Zhang, Jonathan Bath, and Andrew J Turberfield. Remote toehold: a mechanism for flexible control of dna hybridization kinetics. *Journal of the American Chemical Society*, 133(7):2177–2182, 2011.
6. Thomas Gerling, Klaus F. Wagenbauer, Andrea M. Neuner, and Hendrik Dietz. Dynamic DNA devices and assemblies formed by shape-complementary, non-base pairing 3D components. *Science*, 347(6229):1446–1452, 2015.
7. Christos H Papadimitriou. On the complexity of integer programming. *Journal of the ACM (JACM)*, 28(4):765–768, 1981.
8. Matthew J. Patitz. An introduction to tile-based self-assembly and a survey of recent results. *Natural Computing*, 13(2):195–224, 2014.
9. Andrew Phillips and Luca Cardelli. A programming language for composable DNA circuits. *Journal of the Royal Society Interface*, 6(Suppl 4):S419–S436, 2009.
10. John SantaLucia Jr and Donald Hicks. The thermodynamics of DNA structural motifs. *Annu. Rev. Biophys. Biomol. Struct.*, 33:415–440, 2004.
11. Rebecca Schulman and Erik Winfree. Programmable control of nucleation for algorithmic self-assembly. *SIAM Journal on Computing*, 39(4):1581–1616, 2009.
12. Chris Thachuk, Erik Winfree, and David Soloveichik. Leakless DNA strand displacement systems. In *Proceedings of DNA Computing and Molecular Programming 21*, pages 133–153. Springer, 2015.
13. Erik Winfree. *Algorithmic self-assembly of DNA*. PhD thesis, California Institute of Technology, 1998.

14. Sungwook Woo and Paul W. K. Rothemund. Programmable molecular recognition based on the geometry of DNA nanostructures. *Nature Chemistry*, 3:620–627, 2011.
15. Peng Yin, Rizal F. Hariadi, Sudeer Sahu, Harry M. T. Choi, Sung Ha Park, Thomas H. LaBean, and John H. Reif. Programming DNA tube circumferences. *Science*, 321:824–826, 2008.
16. David Yu Zhang, Andrew J. Turberfield, Bernard Yurke, and Erik Winfree. Engineering entropy-driven reactions and networks catalyzed by DNA. *Science*, 318(5853):1121–1125, 2007.

A Proof of Lemma 1

Proof. By structural induction on cascade depth k . Consider as a minimal element a saturated configuration γ_1 having output 1, for a translator cascade of depth $k = 1$, redundancy n , and without input. By assumption, the terminator monomer $\{x_2, \dots, x_2\}$ is free. To saturate the n codomains of type x_2^* , the $\{x_2^*, \dots, x_2^*\}$ monomer and the n $\{x_1^*, x_2\}$ monomers must be in the same polymer. To saturate the n codomains of type x_1^* , the n $\{x_1\}$ monomers must also be in the same polymer containing the n $\{x_1^*, x_2\}$ monomers. There are therefore two polymers containing the following monomers: (1) the terminator monomer $\{x_2, \dots, x_2\}$, and (2) every other monomer. Thus, $S(\gamma_1) = 2$.

Assume that if γ_k is a saturated configuration having output 1 for a translator cascade of depth k , redundancy n , and without input, then $S(\gamma_k) = n(k-1) + 2$.

Consider a saturated configuration γ_{k+1} having output 1, for a translator cascade of depth $k+1$, redundancy n , and without input. Let $L_i = n \times \{x_i^*, x_{i+1}\} \cup n \times \{x_i\}$ be the gate monomers of each layer i ; $1 \leq i \leq k+1$. We first modify γ_{k+1} into a saturated configuration γ'_{k+1} with output 1, such that there are no bonds between monomers in L_1 and L_2 . The only possible bonds between monomers in L_1 and L_2 is between a monomer $\{x_1^*, x_2\} \in L_1$ and a monomer $\{x_2^*, x_3\} \in L_2$. Let p be the number of bonds between these two types of monomers. If $p = 0$, we are done. Otherwise, we note that there must be p free $\{x_2\}$ monomers in γ_{k+1} . Let γ'_{k+1} be the configuration where the p bonds between L_1 and L_2 monomers in γ_{k+1} are replaced by p new bonds between $\{x_2\}$ and $\{x_2^*, x_3\}$ monomers, both from L_2 . Thus γ'_{k+1} remains saturated and can be partitioned into two saturated sub-configurations: γ'_{L_1} containing the monomers from L_1 and γ'_R containing the remainder. Since there are no bonds between the sub-configurations then $S(\gamma'_{k+1}) = S(\gamma'_{L_1}) + S(\gamma'_R)$.

First, note that the $2n$ monomers from L_1 can only form a single saturated configuration containing n polymers, with each polymer being an $\{x_1\}$ monomer bound to a $\{x_1^*, x_2\}$ monomer. Thus, $S(\gamma'_{L_1}) = n$. Second, note that γ'_R is a saturated configuration having output 1 for a translator cascade of depth k , redundancy n , and without input. Thus, $S(\gamma'_R) = n(k-1) + 2$ by the inductive assumption. Therefore, $S(\gamma'_{k+1}) = nk + 2$ establishing the claim. \square

B Proof of Lemma 2

Proof. Let $\mathbf{d} = \mathbf{M}_{\mathcal{T}} \vec{\mathbf{c}}$, $\mathbf{d}_1 = \mathbf{M}_{\mathcal{T}} \vec{\mathbf{c}}_1$, and $\mathbf{d}_2 = \mathbf{M}_{\mathcal{T}} \vec{\mathbf{c}}_2$. For $i \in \{1, \dots, d\}$, $\mathbf{d}(i) \geq 0$ is the number of unbound domains of type $d_i \in \mathcal{D}$ in every saturated configuration $\alpha \in [\vec{\mathbf{c}}]$ (and no d_i^* domains are unbound in α). Let $\alpha_1 \in [\vec{\mathbf{c}}_1]$ and $\alpha_2 \in [\vec{\mathbf{c}}_2]$ be saturated configurations. Define the configuration $\alpha \in [\vec{\mathbf{c}}]$ by $\alpha = \alpha_1 \cup \alpha_2$; note that $S(\alpha) > 1$ since there are no bonds between α_1 and α_2 . Excess domains of each type in α are given by the vector $\mathbf{d}_1 + \mathbf{d}_2 = \mathbf{M}_{\mathcal{T}} \vec{\mathbf{c}}_1 + \mathbf{M}_{\mathcal{T}} \vec{\mathbf{c}}_2 = \mathbf{M}_{\mathcal{T}} (\vec{\mathbf{c}}_1 + \vec{\mathbf{c}}_2) = \mathbf{M}_{\mathcal{T}} \vec{\mathbf{c}} = \mathbf{d}$. Thus α is saturated by Observation 10. Since $S(\alpha) > 1$, every stable configuration of $[\vec{\mathbf{c}}]$ has at least two polymers as well. \square




Fig. 8: **left:** Example showing why a path (dashed red edges) in an acyclic monomer binding graph that repeats the same primary domain/codomain pair *in the same order* implies the configuration is not stable. This configuration is similar to Fig. 7, but by extending the depth of the tree without creating a new domain type, we are forced to reuse the same domain type three times along a path, so two occurrences as one moves along the path (marked with black arrows) have the same primary/codomain ordering. **right:** Exchanging the bonds between these occurrences (dashed green edges) disconnects the graph—increasing entropy—while keeping the configuration saturated.

C Upper bound on the size of a stable polymer if it is a tree

Proposition 1. *Let $\mathcal{T} = (\mathcal{D}, \mathcal{M})$ be a TBN, $d = |\mathcal{D}|$, and $l = \max_{\mathbf{m} \in \mathcal{M}} \|\mathbf{m}\|_1$ be the maximum number of domains in any monomer type. Then any stable polymer ρ of \mathcal{T} such that G_ρ is acyclic has at most $1 + l \cdot \frac{(l-1)^d - l + 1}{l^2 - 3l + 2}$ monomers.*

Proof. Fig. 8 shows an example. A path in G_ρ traverses an edge by either moving from a monomer with a codomain d_i^* to a monomer with a primary domain d_i , or vice versa. Suppose there is a simple path in G_ρ of length $2d + 1$. Then by the pigeonhole principle, some domain d_i is traversed twice with the same primary/codomain ordering. Then by binding the first instance of d_i to the second instance of d_i^* and the second instance of d_i to the first instance of d_i^* , we introduce a cycle. However, since the number of edges has not changed, the new graph must be disconnected, i.e., split into at least two polymers. But since this new configuration is saturated, the original could not have been stable. Therefore the diameter of any acyclic stable monomer binding graph is at most $2d$.

Note that l upper bounds the degree of G_ρ . Any degree- l acyclic graph of diameter $\leq 2d$ has at most $1 + l \cdot \sum_{j=1}^{d-1} (l-1)^{j-1} = 1 + l \cdot \frac{(l-1)^d - l + 1}{l^2 - 3l + 2}$ nodes: the worst case is a rooted complete $(l-1)$ -ary tree of depth d (so diameter $2d$), whose root is the only node that can have up to l children since it has no parent. \square

D Proof of Theorem 13

Proof. It suffices to prove the theorem in the special case where α is itself a single stable polymer. Let $\vec{\mathbf{c}}$ be the monomer collection of α . By Observation 11 we can assume without loss of generality that $\mathbf{M}_\mathcal{T} \vec{\mathbf{c}} \geq \mathbf{0}$.

To turn this inequality into an equality, we introduce slack variables. Let \mathbf{I}_d be the $d \times d$ identity matrix. Let $\mathbf{M}' = [\mathbf{M}_{\mathcal{T}} | -\mathbf{I}_d]$ ($\mathbf{M}_{\mathcal{T}}$ concatenated horizontally with $-\mathbf{I}_d$); note \mathbf{M}' has dimension $d \times (m + d)$. Let $\mathbf{d} = \mathbf{M}_{\mathcal{T}} \vec{\mathbf{c}} \in \mathbb{N}^d$; by Observation 12 $\mathbf{d}(i)$ is the number of unbound primary domains $d_i \in \mathcal{D}$ in every saturated configuration in $[\vec{\mathbf{c}}]$ (and no such configuration has any unbound codomains). Concatenating $\vec{\mathbf{c}}$ with \mathbf{d} to obtain $\vec{\mathbf{c}}' \in \mathbb{N}^{m+d}$, we have $\mathbf{M}' \vec{\mathbf{c}}' = \mathbf{0}$. Let $\mathbf{m}'_1, \dots, \mathbf{m}'_{m+d} \in \mathbb{N}^d$ denote the columns of \mathbf{M}' .

Intuitively, \mathbf{M}' can be thought of as the monomer matrix for a modified TBN that has all the monomer types of \mathcal{M} , and in addition, for each $i \in \{1, \dots, d\}$, also has monomer type $\mathbf{m}'_{m+i} = \{d_i^*\}$. For each $i \in \{1, \dots, d\}$, $\vec{\mathbf{c}}'(m+i)$ represents how many monomers of type \mathbf{m}'_{m+i} should be added to a saturated configuration of $[\vec{\mathbf{c}}]$ to bind to every unbound d_i primary domain, so that the resulting configuration has no unbound domains (i.e., $\mathbf{M}' \vec{\mathbf{c}}' = \mathbf{0}$).

Let $K = (ad)^{d+1}$. If $\max(\vec{\mathbf{c}}) < K$ (i.e., all monomers have count at most K), then $\|\vec{\mathbf{c}}\| < mK = m(ad)^{d+1} < 2(m+d)(ad)^{2d+3}$, and we are done. Otherwise, let $J = \{j_1, j_2, \dots, j_l\} = \{j \in \{1, \dots, m+d\} \mid \vec{\mathbf{c}}'(j) \geq K\}$ denote the set of indices of components of $\vec{\mathbf{c}}'$ that are at least K (we call these *large-count monomer types* below), and let $\bar{J} = \{1, \dots, m+d\} \setminus J$ be all the other indices. Let $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_l$ denote the corresponding columns of \mathbf{M}' , i.e., $\mathbf{v}_i = \mathbf{m}'_{j_i}$.

We now consider two cases.

1. There exist integers n_1, \dots, n_l between 0 and K , not all 0, such that $\sum_{j=1}^l n_j \mathbf{v}_j = \mathbf{0}$. Define $\vec{\mathbf{c}}'_1 \in \mathbb{N}^{m+d}$ by $\vec{\mathbf{c}}'_1(j_i) = n_i$ for $i \in \{1, \dots, l\}$ and $\vec{\mathbf{c}}'_1(j) = 0$ for all $j \in \bar{J}$. In other words, $\vec{\mathbf{c}}'_1$ has n_i copies of the i 'th large-count monomer type and none of the others. Define $\vec{\mathbf{c}}'_2 = \vec{\mathbf{c}}' - \vec{\mathbf{c}}'_1$. Note that by the way $\vec{\mathbf{c}}'_1$ was constructed $\vec{\mathbf{c}}'_2 \geq \mathbf{0}$. We have $\mathbf{M}' \vec{\mathbf{c}}'_1 = \sum_{j=1}^l n_j \mathbf{v}_j = \mathbf{0}$ by definition. From this it follows that $\mathbf{M}' \vec{\mathbf{c}}'_2 = \mathbf{M}'(\vec{\mathbf{c}}' - \vec{\mathbf{c}}'_1) = \mathbf{M}' \vec{\mathbf{c}}' - \mathbf{M}' \vec{\mathbf{c}}'_1 = \mathbf{0}$. Define $\vec{\mathbf{c}}_1 = (\vec{\mathbf{c}}'_1(1), \vec{\mathbf{c}}'_1(2), \dots, \vec{\mathbf{c}}'_1(m))$ and $\vec{\mathbf{c}}_2 = (\vec{\mathbf{c}}'_2(1), \vec{\mathbf{c}}'_2(2), \dots, \vec{\mathbf{c}}'_2(m))$. Since $\vec{\mathbf{c}}'_1 + \vec{\mathbf{c}}'_2 = \vec{\mathbf{c}}'$, we have $\vec{\mathbf{c}}_1 + \vec{\mathbf{c}}_2 = \vec{\mathbf{c}}$. The last d columns of \mathbf{M}' , i.e., $\mathbf{m}'_{m+1}, \dots, \mathbf{m}'_{m+d}$, are all negative unit vectors, and $\vec{\mathbf{c}}'(m+1), \dots, \vec{\mathbf{c}}'(m+d) \geq 0$. Thus, since $\mathbf{M}' \vec{\mathbf{c}}'_1 = \mathbf{0}$, we have that $\mathbf{M}_{\mathcal{T}} \vec{\mathbf{c}}_1 \geq \mathbf{0}$. A similar argument establishes that $\mathbf{M}_{\mathcal{T}} \vec{\mathbf{c}}_2 \geq \mathbf{0}$. Since $\vec{\mathbf{c}}_1, \vec{\mathbf{c}}_2 \in \mathbb{N}^{\mathcal{M}}$ and $\vec{\mathbf{c}}_1 + \vec{\mathbf{c}}_2 = \vec{\mathbf{c}}$, by Lemma 2, $S(\alpha) > 1$, a contradiction since we assumed α is a single polymer.
2. No such integers exist. We show that this implies that the large-count monomer types can't be *too* large. By Lemma 3 there is a vector $\mathbf{h} \in \{0, \pm 1, \dots, \pm K\}^d$ such that $\mathbf{h}^\top \cdot \mathbf{v}_j \geq 1$ for all $j \in \{1, \dots, l\}$; put another way, $\mathbf{h}^\top \cdot \mathbf{m}'_j \geq 1$ for all $j \in J$. Taking the dot product of \mathbf{h}^\top with both sides of the equation $\mathbf{M}' \vec{\mathbf{c}}' = \mathbf{0}$ gives $\mathbf{h}^\top \cdot \mathbf{M}' \vec{\mathbf{c}}' = \mathbf{h}^\top \cdot \mathbf{0} = 0$, i.e., $\sum_{j=1}^{m+d} (\mathbf{h}^\top \cdot \mathbf{m}'_j) \vec{\mathbf{c}}'(j) = 0$. Moving the sum elements corresponding to \bar{J} to the other side of this equation,

$$\sum_{j \in J} (\mathbf{h}^\top \cdot \mathbf{m}'_j) \vec{\mathbf{c}}'(j) = - \sum_{j \in \bar{J}} (\mathbf{h}^\top \cdot \mathbf{m}'_j) \vec{\mathbf{c}}'(j)$$

Since $1 \leq \mathbf{h}^\top \cdot \mathbf{m}'_j$ for all $j \in J$,

$$\sum_{j \in J} \vec{\mathbf{c}}'(j) \leq \sum_{j \in J} (\mathbf{h}^\top \cdot \mathbf{m}'_j) \vec{\mathbf{c}}'(j) = - \sum_{j \in \bar{J}} (\mathbf{h}^\top \cdot \mathbf{m}'_j) \vec{\mathbf{c}}'(j) < (m+d)dK^2a,$$

where the last inequality follows from the fact that there are $|\bar{J}| < m+d$ entries in the sum, d terms in each dot product $\mathbf{h}^\top \cdot \mathbf{m}'_j$, $0 \leq \vec{\mathbf{c}}'(j) < K$ for all $j \in \bar{J}$, $\text{amax}(\mathbf{h}) \leq K$, and $\text{amax}(\mathbf{M}') \leq a$. Also note $\sum_{j \in \bar{J}} \vec{\mathbf{c}}'(j) < (m+d)K$. To complete the proof, observe that since $\vec{\mathbf{c}}'$ is $\vec{\mathbf{c}}$ with extra concatenated nonnegative components,

$$\begin{aligned} \|\vec{\mathbf{c}}\| &\leq \|\vec{\mathbf{c}}'\| = \sum_{j \in J} \vec{\mathbf{c}}'(j) + \sum_{j \in \bar{J}} \vec{\mathbf{c}}'(j) < (m+d)dK^2a + (m+d)K \\ &< 2(m+d)dK^2a = 2(m+d)(ad)^{2d+3}. \quad \square \end{aligned}$$

E Self-assembly: kinetics versus thermodynamics

In this section we compare our thermodynamic model (TBNs) with a well-studied kinetic model of self-assembly called the abstract Tile Assembly Model (aTAM). We first show that it is not necessarily the case that the output (terminal) assembly of an aTAM system is thermodynamically stable when interpreted as a polymer in the TBN model. Specifically we show this for an aTAM binary counter system. Then, we propose a new aTAM system which assembles a binary counter that is also stable polymer in the TBN model. This shows that it is possible, at least for a binary counter, to have systems that are tile/monomer efficient and that self-assemble to an intended target structure in both a kinetic and thermodynamic sense.

E.1 Interpreting aTAM assemblies as TBN polymers

The fundamental components in the aTAM are non-rotatable unit-sized 2D squares with a single “glue” on each of their four sides. A glue consists of a label (which is usually a string or color) and a non-negative integer value called the strength. The glues of two tiles can bind provided that their labels and strengths match. Growth in the aTAM occurs on the integer square lattice of points (\mathbb{Z}^2) and begins with a connected initial configuration of tiles, called a seed, and then proceeds in an asynchronous manner with one tile attaching at a time. A tile can attach to the existing assembly provided that the combined strength of all the glues with which it binds is greater than or equal to a parameter of the system which we call the temperature. An assembly has an associated binding graph: an undirected graph on \mathbb{Z}^2 with nodes corresponding to tiles in the assembly and edges corresponding to bound glues. An assembly to which no more tiles can attach is said to be terminal, Fig. 9(a) shows an example. Formally, an aTAM system is defined to be a triple (T, σ, τ) where T is a finite

set of tiles, $\sigma \in T$ is the seed tile and $\tau \in \mathbb{Z}^+$ is the temperature. Comprehensive informal and formal definitions of the aTAM can be found in [8].

We can interpret any aTAM assembly as a TBN polymer in the following way. Each aTAM tile is mapped to a TBN monomer and each aTAM glue on a tile side is mapped to a TBN domain in a monomer. The binding graph of an aTAM assembly is mapped to a corresponding TBN configuration (which, due geometry, is a graph and not a multigraph). Notion in the aTAM tiles having a square shape, there being a seed tile, temperature, and binding occurring on the integer lattice have no corresponding notions in the TBN.

E.2 A tile-based binary counter that is not a thermodynamically stable polymer

In this subsection we show that it is not necessarily the case that the terminal assembly of an aTAM system is thermodynamically stable when interpreted as a polymer in the TBN model. To get an immediate intuitive sense of a difference in the two models simply observe that a monomer in the TBN model, can “self-saturate” meaning that if it contains two domains which are complementary, then those two domains can bind to each other as opposed to binding to domains of other monomers.¹² Even worse, in a TBN system there is no geometric requirement on the arrangement of bonds, so polymers can form into arbitrary networks. However, in the aTAM square tiles can never self-bind, nor can they bind to tiles that are not adjacent on the 2D square lattice.

Part (a) of Fig. 9 shows the unique terminal assembly of an aTAM system which assembles a binary counter that counts to 16. We briefly describe the system here. The system has a tile set T which is the union of all tiles shown in Fig. 9(a), a seed tile σ which is the bottom leftmost tile, and works at temperature 2. Strength 1 and 2 glues are respectively denoted as thin and thick coloured lines between adjacent tiles. The tile set T contains four “hard-coded” tiles with strength 2 glues so that the leftmost (or “first”) column of the counter can grow directly from the seed at temperature 2. Notice in the assembly that the strength 2 glues which initiate the growth of new columns alternate between occurring on the east of the bottom tile of the column and occurring on the east of some other tile in the column. When they occur on the bottom, growth continues upwards with a carry c being propagated, causing bits to be flipped from west to east, until a 0 is met. If the strength 2 glues initiates a column from elsewhere, there are tile types in T so that the values of glues are simply copied to the east except for the least significant (i.e. bottom) bit which is incremented from 0 to 1. Growth of the counter is halted when the red strength 2 glue is exposed which causes a column of tiles to bind which do not have any output glues.

We can interpret the assembly shown in Fig. 9(a) as a TBN polymer α by viewing the square tiles as monomers and their glues as domains, as described in Section E.1. Observe that the TBN configuration β in Fig. 9(b) has

¹² Experimentally, for example, such a situation can occur when using the DNA single stranded tile motif [15].




Fig. 9: **(a)** A binary counter in the abstract Tile Assembly Model (aTAM) that self-assembles a $k \times (2^k + 1)$ rectangle from $k + O(1)$ square tiles, where $k = 4$. Thick lines between tiles are “strength 2” domains and thin lines are “strength 1” domains. Colors of glues between tiles help to indicate “positional specificity”. (For aTAM glues we ignore all “*” annotations in the figure.) This assembly can be interpreted as a TBN polymer (where tiles become monomers, and glues become domains). However, the resulting TBN polymer is not stable, since **(b)** shows another stable TBN configuration using the same set of monomers, but with greater entropy since it is composed of two polymers (i.e. monomers can “self-saturate”). **(c)** An aTAM counter using $O(k)$ tiles, that is also a stable polymer when interpreted in the TBN model. Compared to part (a) and (b), in (c) each monomer encode its height i (i.e., the significance of the relevant bit), by having domains and codomains subscripted with i or $i - 1$. In the proof of its stability, this polymer is denoted γ .

the same multiset of monomers as α . Furthermore, one of the tile-monomers of β is self-saturated (each of its domains is bound to another domain on the same monomer). Finally, observe that $H(\alpha) = H(\beta)$ (both have the same number of bound domains) but $S(\alpha) < S(\beta)$ (β has one more polymer than α). Hence α is not a stable configuration. This yields the straightforward conclusion that a terminal assembly producible in the aTAM is not necessarily a thermodynamically stable polymer when interpreted in the TBN model.

E.3 A thermodynamically stable binary counter

Part (c) of Fig. 9 shows a binary counter that is a terminal assembly producible by an aTAM system. Let γ be the interpretation of that assembly as a TBN polymer. This section gives a proof of Lemma 4, but first we give a little intuition.

To create the tile assembly system that builds the terminal assembly in Fig. 9(c), the tile assembly system in Fig. 9(a) is modified so that the labels of the glues encode the natural number “height” at which they appear. Intuitively we can already see an obvious difference between β and γ since in γ the “top” and “bottom” domain of each monomer is distinct, hence it is not possible for any monomer in γ to bind only to itself; it must bind to other monomers in a stable configuration. However, this observation alone is not sufficient to show that the single-polymer configuration γ in Fig. 9(c) is stable. In the remainder of this section we argue that that is indeed the case.

Lemma 4. *γ is a stable polymer in the TBN model.*

We define the *seed column* to be the set of monomers in the leftmost column of γ and the *end column* to be the set of monomers in the rightmost column. We also enumerate rows of γ from bottom to top starting with 1 for the bottom row. We refer to any monomer that contains a subscripted c or c^* glue in row $i \in \{1, 2, 3, 4\}$ of γ (counting from the bottom) as a c^i -monomer. In addition, if a c^i -monomer contains a c_{i-1}^* domain, we refer to it as a c_n^i -monomer, otherwise we refer to it as a c_s^i -monomer. Furthermore, we refer to monomers in row i which contain subscripted n domains as n^i -monomers.

To see that the configuration containing exactly γ is thermodynamically stable, we will show that all stable configurations consist of one polymer with all domains bound.

Let C be *any* stable configuration consisting of the monomers shown Fig. 9(c). Let δ be the unique polymer in C which contains the *seed monomer* (i.e. the bottom-left tile of Fig. 9(c) interpreted as a monomer). We will show that δ must in fact contain all monomers in the polymer shown in Fig. 9(c), which in turn implies Lemma 4.

Let row i of δ be the set of all monomers in δ that appear on row i of γ . Likewise left column i of δ be the set of all monomers in δ that appear on column i of γ .

The following observation states that δ cannot contain any partial columns.

Claim 14. *The rows of δ all contain the same number of monomers.*

Proof. γ does not contain any unbound domains, therefore by stability neither does δ . If two rows of δ each had a different number of monomers then this would imply that some domain of δ is unbound, which is impossible. \square

The next observation also follows from the fact that δ cannot contain any unbound domains.

Claim 15. *The polymer δ contains the seed column and end column.*

Proof. First, note that since δ contains the seed monomer by assumption, it must contain the entire seed column; otherwise some domain would be unbound because the domains in the seed column are unique and do not appear on any other monomers. To see that δ must contain the end column we examine the bottom row of δ . We have already established that δ must contain the seed column. Now, observe that the bottom monomer of the seed row contains a horizontal starred domain and all the inner monomers (the monomers which do not appear in the seed or end columns) contain exactly one starred horizontal domain and one unstarred horizontal domain. Thus, if the bottom monomer in the end column was not in δ , there would be one extra unstarred horizontal domain in the bottom row of δ which means δ would contain unbound domains. And, for the same reason δ must contain the entire seed column, it must contain the entire end column. \square

Since δ contains the seed and end columns and the top row of δ cannot contain any unbound domains, we have the following:

Observation 16. *The polymer δ contains both c^4 monomers.*

For $i \in \{1, 2, 3, 4\}$, we denote the number of c^i monomers in δ by $|c^i|$. Likewise, we denote the number of c_n^i and c_s^i monomers in δ by $|c_n^i|$ and $|c_s^i|$ respectively.

Claim 17. *Let $i \in \{2, 3\}$. Then $|c_n^i| = |c_s^i|$.*

Proof. To see this, first note that the number of subscripted 0 domains must equal the number of subscripted 0^* domains in row i so that no domains are left unbound. Denote the number of monomers which contain a subscripted n domain in row i by $|n^i|$. Observe that monomers on row i that have both subscripted n^* and n domains each have an equal number of subscripted 0 and 0^* domains. Also, notice that for each i the two monomers on row i of the seed and end columns together (i.e. unioned) have an equal number of subscripted 0 and 0^* domains. Now observe that c_s^i monomers have an extra subscripted 0 domain and c_n^i -monomers have an extra subscripted 0^* domain. Also, we note that a row of monomers in row i is made up completely of monomers from the seed and end columns, c_n^i -monomers, c_s^i -monomers, and monomers that each contain a pair of subscripted n, n^* domains. Hence, in order for row i to have an equal number of subscripted 0 and 0^* domains, the number of c_n^i and c_s^i monomers must be equal. \square

Claim 18. For $i > 2$, $|c^i| = \frac{|c^{i-1}|}{2}$ and $|c^2| = |c^1|$.

Proof. In order for all domains in δ to be bound, row i must contain just enough c -monomers to bind with all the c_n^{i-1} -monomers. It immediately follows from this that $|c^2| = |c^1|$ since the c -monomers in row 1 each contain the c_0^* domain which in turn is bound to some c_0 domain which in turn are only found on row 2 in c^2 -monomers.

Let $i > 2$. Claim 17 states that the number of c_n^{i-1} -monomers and c_s^{i-1} -monomers is equal. Hence exactly half of the c^{i-1} monomers, in fact the set of c_n^{i-1} -monomers, contain the c_{i-2}^* domain that that can bind to a c^i monomer. Also, all c^i monomers have a c_{i-2} domain that binds to a c^{i-1} monomer. Thus, the number of c -monomers in row i is exactly half the number c -monomers in row $i - 1$. In other words, row i contains exactly $\frac{|c^{i-1}|}{2}$ of c^i -monomers when $i > 2$. \square

The next claim allows us to determine the number of columns contained in δ based on the number of c^1 -monomers contained in δ .¹³

Claim 19. The number of columns in δ is $2 + 2|c^1| - 1$.

Proof. To see why Claim 19 holds, note that Claim 15 tells us that δ contains both the seed and end columns. In addition, it follows from Claim 14 that for each c^1 -monomer in δ , δ must contain a full column. In order for δ not to have any exposed domains there must be an n^1 -monomer between each c^1 -monomer. Once again, applying Claim 14, for each of these n^1 -monomers, δ must contain a full column. Consequently, it follows that δ contains $2 + 2|c^1| - 1$ columns. \square

We are now ready to conclude this section with the proof of its main claim, Lemma 4.

Proof of Lemma 4. It follows from Observation 16 that δ contains both c^4 -monomers. Claim 18 then implies that row 3 has 4 c^3 -monomers. Applying Claim 18 again, we have that δ contains 8 c^2 -monomers. By the kind of reasoning used in the proof of Claim 18 once more, we have that δ contains 8 c^1 -monomers. Using Claim 19, we have that δ contains 17 columns and thus contains all the monomers of γ (shown in Fig. 9(c)). Consequently, any stable configuration has exactly one polymer, and since γ is saturated, it is stable. \square

¹³ Although we are not yet saying whether the number of c^1 -monomers in δ is the same as the number of c^1 -monomers in γ