

Thermodynamically Driven Signal Amplification

Joshua Petrack ✉ 

University of California–Davis, CA, USA

David Soloveichik ✉ 

University of Texas at Austin, TX, USA

David Doty ✉ 

University of California–Davis, CA, USA

Abstract

The field of chemical computation attempts to model computational behavior that arises when molecules, typically nucleic acids, are mixed together. By modeling this physical phenomenon at different levels of specificity, different operative computational behavior is observed. Thermodynamic binding networks (TBNs) is a highly abstracted model that focuses on which molecules are bound to each other in a “thermodynamically stable” sense. Stability is measured based only on how many bonds are formed and how many total complexes are in a configuration, without focusing on how molecules are binding or how they became bound. By defocusing on kinetic processes, TBNs attempt to naturally model the long-term behavior of a mixture (i.e., its thermodynamic equilibrium).

We study the problem of *signal amplification*: detecting a small quantity of some molecule and amplifying its signal to something more easily detectable. This problem has natural applications such as disease diagnosis. By focusing on thermodynamically favored outcomes, we seek to design chemical systems that perform the task of signal amplification robustly without relying on kinetic pathways that can be error prone and require highly controlled conditions (e.g., PCR amplification).

It might appear that a small change in concentrations can result in only small changes to the thermodynamic equilibrium of a molecular system. However, we show that it is possible to design a TBN that can “exponentially amplify” a signal represented by a single copy of a monomer called the *analyte*: this TBN has exactly one stable state before adding the analyte and exactly one stable state afterward, and those two states “look very different” from each other. In particular, their difference is exponential in the number of types of molecules and their sizes. The system can be programmed to any desired level of resilience to false positives and false negatives. To prove these results, we introduce new concepts to the TBN model, particularly the notions of a TBN’s entropy gap to describe how unlikely it is to be observed in an undesirable state, and feed-forward TBNs that have a strong upper bound on the number of polymers in a stable configuration.

We also show a corresponding negative result: a *doubly* exponential upper bound, meaning that there is no TBN that can amplify a signal by an amount more than doubly exponential in the number and sizes of different molecules that comprise it. We leave as an open question to close this gap by either proving an exponential upper bound, or giving a construction with a doubly-exponential difference between the stable configurations before and after the analyte is added.

Our work informs the fundamental question of how a thermodynamic equilibrium can change as a result of a small change to the system (adding a single molecule copy). While exponential amplification is traditionally viewed as inherently a non-equilibrium phenomenon, we find that in a strong sense exponential amplification can occur at thermodynamic equilibrium as well – where the “effect” (e.g., fluorescence) is exponential in types and complexity of the chemical components.

2012 ACM Subject Classification Theory of computation → Models of computation

Keywords and phrases Thermodynamic binding networks, signal amplification, integer programming

Digital Object Identifier 10.4230/LIPIcs.DNA.29.8

Related Version *Full Version*: <https://arxiv.org/abs/2307.01550>

Funding *Joshua Petrack*: NSF grants 1900931 and 1844976.

David Soloveichik: NSF grant 1901025, Sloan Foundation Research Fellowship.

David Doty: NSF grants 2211793, 1900931, and 1844976.



© Joshua Petrack, David Soloveichik, and David Doty;
licensed under Creative Commons License CC-BY 4.0

29th International Conference on DNA Computing and Molecular Programming (DNA 29).

Editors: Ho-Lin Chen and Constantine G. Evans; Article No. 8; pp. 8:1–8:22

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

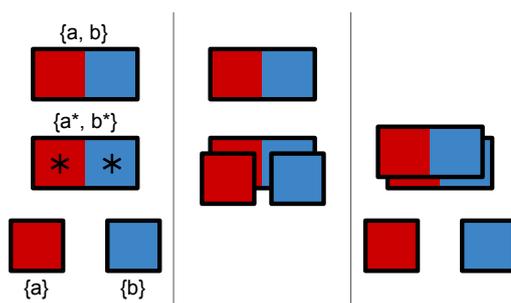
Detecting a small amount of some chemical signal, or analyte, is a fundamental problem in the field of chemical computation. The current state-of-the-art in nucleic acid signal amplification is the polymerase chain reaction (PCR)[9]. By using a thermal cycler, PCR repeatedly doubles the amount of the DNA strand that is present. One downside is the need for a PCR machine, which is expensive and whose operation can be time-consuming. The advantages of PCR are that it can reliably detect even a single copy of the analyte if enough doubling steps are taken, and it is fairly (though not perfectly) robust to incorrect results. Recent work in DNA nanotechnology achieves “signal amplification” through other kinetic processes involving pure (enzyme-free) DNA systems, such as hybridization chain reaction (HCR) [4], classification models implemented with DNA [7], hairpin assembly cascades [10], and “crisscross” DNA assembly [8].

Although highly efficacious, PCR and these other techniques essentially rely on *kinetic* control of chemical events, and the thermodynamic equilibria of these systems are not consistent with their desired output. Can we design a system so that, if the analyte is present, the thermodynamically most stable state of the system looks one way, and if the analyte is absent, the thermodynamically most stable state looks “very different” (e.g., many fluorophores have been separated from quenchers)? Besides answering a fundamental chemistry question, such a system is potentially more robust to false positives and negatives. It also can be simpler and cheaper to operate: for many systems, heating up the system and cooling it down slowly reaches the system’s thermodynamic equilibrium.

We tackle this problem of signal detection in the formal model of Thermodynamic Binding Networks (TBNs) [5, 3]. The TBN model of chemical computation ignores kinetic and geometric constraints in favor of focusing purely on configurations describing which molecules are bound to which other molecules. A TBN yields a set of stable configurations, the ways in which monomers (representing individual molecules, typically strands of DNA) are likely to be bound together in thermodynamic equilibrium. A TBN performs the task of signal amplification if its stable configurations, and thus the states in which it is likely to be observed at equilibrium, change dramatically in response to adding a single monomer. TBNs capture a notion of what signal amplification can look like for purely thermodynamic chemical systems, without access to a process like PCR that repeatedly changes the conditions of what is thermodynamically favorable.

This paper asks the question: if we add a single molecule to a pre-made solution, how much can that change the solution’s thermodynamic equilibrium? To make the question quantitative, we define a notion of distance between thermodynamic equilibria, and we consider scaling with respect to meaningful complexity parameters. First, we require an upper limit on the size of molecules in the solution and the analyte, as adding a single very large molecule can trivially affect the entire solution. Large molecules are also expensive to synthesize, and for natural signal detection the structure of the analyte is not under our control. Second, we require an upper limit on how many different types of molecules are in the solution, as it is expensive to synthesize new molecular species (though synthesizing many copies is more straightforward).

Our main result is the existence of a family of TBNs that amplify signal exponentially. In these TBNs, there are exponentially many free “reporter” monomers compared to the number of types of monomers and size of monomers. In the absence of the analyte, this TBN has a unique stable configuration in which all reporter monomers are bound. When a single copy of the analyte is added, the resulting TBN has a unique stable configuration in



■ **Figure 1** A simple thermodynamic binding network T with four monomers. Site types are differentiated by color. Bonds are shown by juxtaposing monomers so that unstarred sites cover starred sites. Left: the all-singleton configuration $\text{melt}(T)$ with four polymers. Monomers are labeled by their formal identities for reference. Middle: a configuration with two polymers. As all starred sites are covered, the configuration is saturated. Right: a configuration with three polymers. As this has the most possible polymers for a saturated configuration, it is stable.

which all reporter monomers are unbound. These TBNs are parameterized by two values: the first is the amplification factor, determining how many total reporter molecules are freed. The second is a value we call the system’s “entropy gap”, which determines how thermodynamically unfavorable a configuration of the system would need to be in order for reporters to be spuriously unbound in the absence of the analyte (false positive) or spuriously bound in its presence (false negative).

We also show a corresponding doubly exponential upper bound on the signal amplification problem in TBNs: that given any TBN, adding a single monomer can cause at most a doubly exponential change in its stable configurations. We leave as an open question to close this gap: either proving an exponential upper bound, or giving a TBN with a doubly-exponential amplification factor.

Our work can be compared to prior work on signal amplification that exhibits kinetic barriers. For example, in reference [8], a detected analyte serves as a seed initiating self-assembly of an arbitrarily long linear polymer. In the absence of the analyte, an unlikely kinetic pathway is required for spurious nucleation of the polymer to occur. However, in that system, false positive configurations are still thermodynamically favorable; if a critical nucleus is able to overcome the kinetic barrier and assemble, then growth of the infinite polymer is equally favorable as from the analyte. In contrast, in our system, there are *no* kinetic paths, however unlikely, that lead to an undesired yet thermodynamically favored configuration.

2 Definitions

2.1 General TBN Definitions

A *site type* is a formal symbol such as a , and has a *complementary* type, denoted a^* , with the interpretation that a binds to a^* (e.g., they could represent complementary DNA sequences). We also refer to site types as domain types, and sites as domains. We call a site type such as a an *unstarred* site type, and a^* a *starred* site type. A *monomer type* is a multiset of site types (e.g., a DNA strand consisting of several binding domains); for example monomer type $\vec{m} = \{a, a, a, b, c^*\}$ has three copies of site a , one of site b , and one of site c^* . A *TBN* [5, 2] is a multiset of monomer types. We call an instance of a monomer type a *monomer* and an instance of a site type a *site*.

8:4 Thermodynamically Driven Signal Amplification

We take a convention that, unless otherwise specified, TBNs are *star-limiting*: for each site type, there are always at least as many sites of the unstarred type as the starred type among all monomers. Given a TBN, this can always be enforced by renaming site types to swap unstarred and starred types, which simplifies many of the definitions below.

A *configuration* of a TBN is a partition of its monomers into submultisets called *polymers*. We say that a site type (or a site) on a polymer is *uncovered* if, among the monomers in that polymer, there are more copies of the starred version of that site type than the unstarred version (otherwise *covered*). A polymer is *self-saturated* if it has no uncovered site types. A configuration is *saturated* if all its polymers are self-saturated. A configuration α of a TBN T is *stable* if it is saturated, and no saturated configuration of T has more polymers than α . Figure 1 shows an example TBN.

An equivalent characterization of stability is in terms of merges rather than polymer counts. We say that a *merge* is the process of taking two polymers in a configuration and making a new configuration by joining them into one polymer; likewise a *split* is the process of taking one polymer in a configuration and making a new configuration by splitting it into two polymers. Maximizing the number of polymers in a saturated configuration is equivalent to minimizing the number of merges of two polymers necessary to reach a saturated configuration. To this end, some additional notation:

► **Definition 2.1.** *The distance to stability of a saturated configuration σ is the number of (splits minus merges) necessary to get from σ to a stable configuration.*

Note that this number will be the same for any path of splits and merges, as all stable configurations have the same number of polymers.

Equivalently, distance to stability is the number of polymers in a stable configuration minus the number of polymers in σ . We only consider this value for saturated configurations to ensure it is positive and because we may interpret it as a measure of how unlikely we are to observe the network in a given state under the assumption that enthalpy matters infinitely more than entropy.

The following definitions are not restricted to saturated configurations.

► **Definition 2.2.** *Given a TBN T , we say that the all-melted configuration, denoted $\text{melt}(T)$, is the configuration in which all monomers are separate.*

► **Definition 2.3.** *Given a configuration α in a TBN T , its merginess $m(\alpha)$ is the number of merges required to get from $\text{melt}(T)$ to α (or equivalently, the number of monomers in T minus the number of polymers in α).*

► **Definition 2.4.** *Given a configuration α in a TBN T , its starriness $s(\alpha)$ is the number of polymers in α which contain at least one uncovered starred site.*

We observe that α is saturated if and only if $s(\alpha) = 0$.

► **Definition 2.5.** *Given configurations α and β in a TBN T , we say $\alpha \preceq \beta$ (equivalently, $\beta \succeq \alpha$) if it is possible to reach β from α solely by splitting polymers zero or more times.*

We read $\alpha \preceq \beta$ as “ α splits to β ”. Observe that if $\alpha \succeq \beta$, then we can reach β from α in exactly $m(\beta) - m(\alpha)$ merges. In general, we may order the merges required to go from one configuration to another in whatever way allows the easiest analysis.

2.2 Comparing TBNs

We need some notion of how “different” two TBNs are, so that we can quantify how much a TBN changes after adding a single monomer.

► **Definition 2.6** (distance between configurations). *Let α and β be two configurations of a TBN, or of two TBNs using the same monomer types. We say that the distance $d(\alpha, \beta)$ between them is the L^1 distance between the vectors of their polymer counts. That is, it is the sum over all types of polymers of the difference between how many copies of that polymer are in α and in β .*

► **Definition 2.7** (distance between TBNs). *Given TBNs T and T' , let \mathcal{C} and \mathcal{C}' be their stable configurations. Define the distance between T and T' as*

$$d(T, T') = \min_{\alpha \in \mathcal{C}, \alpha' \in \mathcal{C}'} d(\alpha, \alpha'). \quad (1)$$

Note that this distance is not a metric.¹ Rather, it is a way to capture how easily we can distinguish between two TBNs; even the closest stable configurations of T and T' have distance $d(T, T')$, so we should be able to distinguish any stable configuration of one of them from all stable configurations of the other by that amount.

Note that this condition does not directly imply a stronger “experimentally verifiable” notion of distance, namely that there is some “reporter” monomer which is always bound in one TBN and always free in the other. However, the system we exhibit in this paper does also satisfy this stronger condition. We focus on the distance given here, as it is more theoretically general and our upper bound result in Section 4 apply to it.

We also need a notion of how likely we are to observe a configuration of a TBN that is not stable, in order to have a notion of the system being robust to random noise. If a TBN has one stable configuration but many other configurations that are nearly stable, we would expect to observe it in those configurations frequently, meaning that in practice we may not be able to discern what the stable configuration is as easily.

We work under the assumption that enthalpy matters infinitely more than entropy, so that we may assume that only saturated configurations need to be considered. This assumption is typical for the TBN model, and can be accomplished practically by designing binding sites to be sufficiently strong. Under this paradigm, a configuration’s distance to stability is a measure of how unlikely we are to observe it. This motivates the following definition:

► **Definition 2.8** (entropy gap). *Given a TBN T , we say that it has an entropy gap of k if, for any saturated configuration α of T , one of the following is true:*

1. α is stable.
2. There exists some stable configuration β such that $\alpha \preceq \beta$.
3. α has distance to stability at least k .

Note that by this definition, all TBNs trivially have an entropy gap of one. Note as well that stable configurations are technically also included in the second condition by choosing $\beta = \alpha$, but we list them separately for emphasis.

The second condition is necessary in this definition because any TBN necessarily has some configurations that have distance to stability one, simply by taking a stable configuration and arbitrarily merging two polymers together. These configurations are unavoidable but

¹ In particular, it fails to satisfy the triangle inequality, since T could have a stable configuration close to one of T' , so $d(T, T') = 1$, and T' could have a different stable configuration close to one of T'' , so $d(T', T'') = 1$, but T and T'' could have no close stable configurations, so $d(T, T'') > 2$.

are not likely to be problematic in a practical implementation, because a polymer in such a configuration should be able to naturally split itself without needing to interact with anything else – these configurations will never be local energy minima. Reference [2] discusses self-stabilizing TBNs in which *all* saturated configurations have this property, equivalent to an entropy gap of ∞ .

2.3 Feed-Forward TBNs

► **Definition 2.9.** *We say that a configuration α of a TBN is feed-forward if there is an ordering of its polymers such that for each domain type, all polymers with an excess of unstarred instances of that domain type occur before all polymers with an excess of starred instances of that domain type.*

We say that a TBN T is feed-forward if there is an ordering of its monomer types with this same property – that is, T is feed-forward if $\text{melt}(T)$ is feed-forward.

For example, the TBN $\{(ab), (a^*c), (b^*c^*)\}$ is feed-forward with this ordering of monomers because the a, b and c come strictly before the a^*, b^* and c^* respectively. Note that not all configurations of a feed-forward TBN are necessarily feed-forward; for instance, merging the first and third monomers in this TBN gives a non-feed-forward configuration.

An equivalent characterization can be obtained by defining a directed graph on the polymers of a configuration α and drawing an edge between any two polymers that can bind to each other, from the polymer with an excess unstarred binding site to the polymer with a matching excess starred binding site (or both directions if both are possible). The configuration α is feed-forward if and only if this graph is acyclic, and the ordering of polymers can be obtained by taking a topological ordering of its vertices.

The main benefit of considering feed-forward TBNs is that we can establish a strong lower bound on the merginess of stable configurations. If *any* TBN T has n monomers that have starred sites, it will always take at least $\frac{n}{2}$ merges to cover all those sites, because each monomer must be involved in at least one merge and any merge can at most bring a pair of them together. For instance, the non-feed-forward TBN $\{\{a, b^*\}, \{a^*, b\}\}$ can be stabilized with a single merge. In feed-forward TBNs, this bound is even stronger, as there is no way to “make progress” on covering the starred sites of two different monomers at the same time.

► **Lemma 2.10.** *If a configuration α is feed-forward, then any saturated configuration σ such that $\alpha \succeq \sigma$ satisfies $m(\sigma) - m(\alpha) \geq s(\alpha)$. That is, reaching σ from α requires at least $s(\alpha)$ additional merges.*

Intuitively, in a feed-forward configuration, the best we can possibly do is to cover all of the starred sites on one polymer at a time. We can never do better than this with a merge like merging $\{a, b^*\}$ and $\{a^*, b\}$ that would let two polymers cover all of each others’ starred sites.

Proof. Given a feed-forward configuration α , let L be the ordered list of polymers from α being feed-forward. Partition L into separate lists (keeping the ordering from L) based on which polymers are merged together in σ . That is, for each fully merged polymer $\mathbf{P} \in \sigma$ create a list $L_{\mathbf{P}}$ of the polymers from α that are merged to form \mathbf{P} , and order this list based on the ordering from L . We can order the merges to reach σ from α as follows: repeatedly (arbitrarily) pick a polymer \mathbf{P} from σ and merge all of the polymers in $L_{\mathbf{P}}$ together in order (merge the first two polymers in $L_{\mathbf{P}}$, then merge the third with the resulting polymer, and so on).

This sequence of merges gives us a sequence of configurations $\alpha = \alpha_1, \alpha_2, \dots, \alpha_\ell = \sigma$. We observe that for $1 \leq i \leq \ell - 1$, we have $s(\alpha_i) - s(\alpha_{i+1}) \leq 1$. That is, each merge can lower the starriness by at most one. We know this because each merge is merging a polymer $\mathbf{Q} \in \alpha$ with one or more other already-merged polymers from α that all come before \mathbf{Q} in L . This means \mathbf{Q} cannot cover any starred sites on any monomers it is merging with. The only way for the starriness of a configuration to decrease by more than 1 in a single merge is for the two merging polymers to cover all of each others' starred sites, so it follows that each merge in this sequence lowers starriness by at most 1. From this it follows that we need at least $s(\alpha)$ merges to get to σ , because $s(\sigma) = 0$. ◀

Letting $\alpha = \text{melt}(T)$ (note $m(\alpha) = 0$) gives the following corollary.

► **Corollary 2.11.** *Any saturated configuration σ of a feed-forward TBN T satisfies $m(\sigma) \geq s(\text{melt}(T))$.*

Because stable configurations are saturated configurations with the minimum possible merginess, this bound gives the following corollary.

► **Corollary 2.12.** *If a saturated configuration σ of a feed-forward TBN T satisfies $m(\sigma) = s(\text{melt}(T))$, then σ is stable.*

3 Signal Amplification TBN

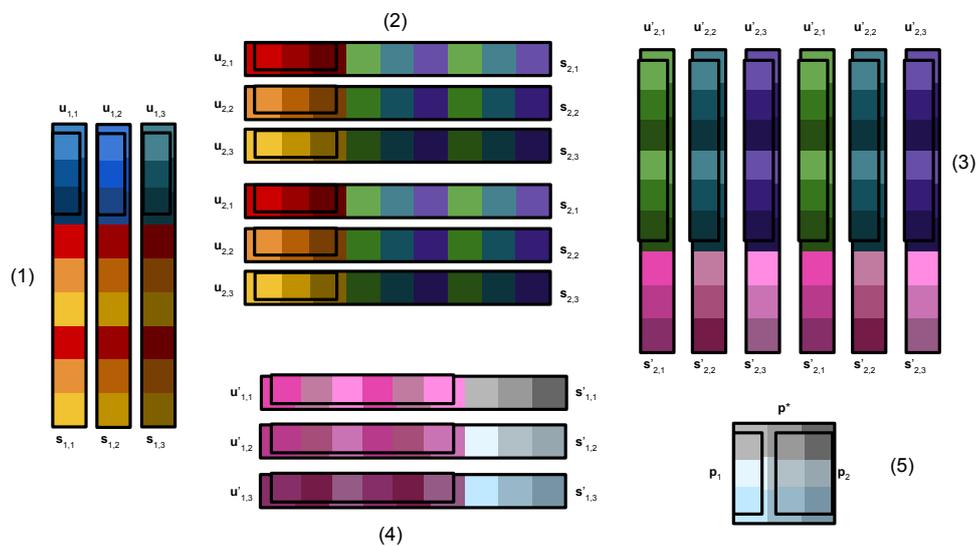
3.1 Amplification Process

In this section, we prove our main theorem. This theorem shows the existence of a TBN parameterized by two values n (the amplification factor) and k (the entropy gap). Intuitively, this TBN amplifies the signal of a single monomer by a factor of 2^n , with any configurations that give “incorrect” readings having $\Omega(k)$ distance to stability. Our proof will be constructive.

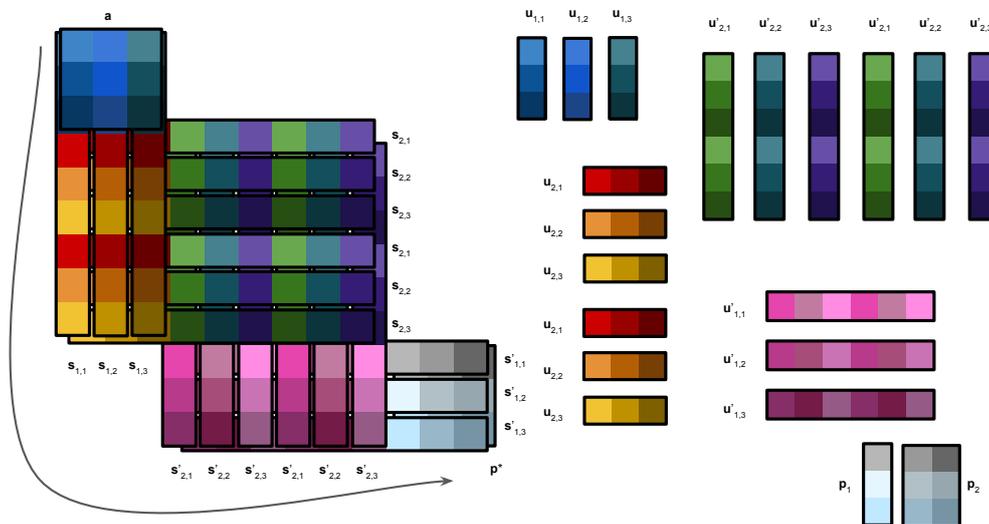
► **Theorem 3.1.** *For any integers $n \geq 1, k \geq 2$, there exists a TBN $T = T_{n,k}$ and monomer \mathbf{a} (the analyte) such that if $T^{\mathbf{a}} = T_{n,k}^{\mathbf{a}}$ is the TBN obtained by adding one copy of \mathbf{a} to $T_{n,k}$, then*

1. T and $T^{\mathbf{a}}$ each have exactly one stable configuration, denoted $\sigma_{n,k}$ and $\sigma_{n,k}^{\mathbf{a}}$ respectively, with $d(\sigma_{n,k}, \sigma_{n,k}^{\mathbf{a}}) \geq 2^n$. In particular, there are k monomer types with 2^{n-1} copies each, with all of these monomers bound in $\sigma_{n,k}$ and unbound in $\sigma_{n,k}^{\mathbf{a}}$.
2. T and $T^{\mathbf{a}}$ each have an entropy gap of $\lfloor \frac{k}{2} \rfloor - 1$.
3. T and $T^{\mathbf{a}}$ each use $\mathcal{O}(nk)$ total monomer types, $\mathcal{O}(nk^2)$ domain types, and $\mathcal{O}(k^2)$ domains per monomer.

The first condition implies that $T_{n,k}$ can detect a single copy of \mathbf{a} with programmable exponential strength - there is only one stable configuration either with or without \mathbf{a} , and they can be distinguished by an exponential number of distinct polymers. Note that this is even stronger than saying that $d(T_{n,k}, T_{n,k}^{\mathbf{a}}) \geq 2^n$, as that statement would allow each TBN to have multiple stable configurations. The second condition implies that the system has a programmable resilience to having incorrect output, because configurations other than the unique stable ones in each case are “programmably” unstable (based on k), and thus programmably unlikely to be observed. Note that throughout this paper we will use $\frac{k}{2}$ instead of $\lfloor \frac{k}{2} \rfloor$ for simplicity, as we are concerned mainly with asymptotic behavior. The third condition establishes that the system doesn't “cheat” - it doesn't obtain this amplification by either having an extremely large number of distinct monomers, or by having any single large monomers.



■ **Figure 2** The unique stable configuration $\sigma_{2,3}$ of $T_{2,3}$, with 19 polymers. All starred sites are visually “covered” by unstarred sites on another monomer. The parts of the diagram are numbered by the order that the signal from the analyte will cascade through them. Parts (1) and (2) form the “first half”, where the signal is doubled at each step. Parts (3) and (4) form the “second half”, where the signal converges so that it can get an “entropic payoff” from part (5).



■ **Figure 3** The unique stable configuration $\sigma_{2,3}^a$ of $T_{2,3}^a$. The arrow shows the conceptual order in which the analyte’s signal has been propagated, with **a** covering all $s_{1,j}$, which cover all $s_{2,j}$, which cover all $s'_{1,j}$, which finally cover p^* . This configuration has 21 polymers, 2 more than $\sigma_{2,3}$: conceptually, one of these is from adding the analyte and the other is from the analyte’s signal cascading through the layers to release P_1 and P_2 at the cost of one merge. As they have no polymers in common, $d(\sigma_{2,3}, \sigma_{2,3}^a) = 19 + 21 = 40$.

The entire TBN $T_{n,k}$ is depicted in Figures 2 and 3 with $n = 2$ and $k = 3$. The former shows the unique stable configuration before adding the analyte, and the latter shows the unique stable configuration after adding the analyte. For comparison, Figure 7 depicting the pre-analyte configuration with $n = 3$ and $k = 2$ is shown in the appendix.

We will start by constructing the first half of $T_{n,k}$ and describing “how it works”. The monomers in this first half are the driving force that allows $T_{n,k}$ and $T_{n,k}^a$ to have exponentially different stable configurations.

The first half of $T_{n,k}$ has monomer types $\mathbf{u}_{i,j}$ and $\mathbf{s}_{i,j}$ (named as those with only unstarred sites, and those with both starred and unstarred sites) for $1 \leq i \leq n, 1 \leq j \leq k$. It has domain types denoted as triples (i, j, ℓ) for $1 \leq i \leq n + 1, 1 \leq j, \ell \leq k$. Each $\mathbf{u}_{i,j}$ monomer has k different unstarred domains, one of each (i, j, ℓ) for each $1 \leq \ell \leq k$. Each $\mathbf{s}_{i,j}$ monomer has a starred copy of each domain in $\mathbf{u}_{i,j}$, and additionally has two copies of each unstarred domain $(i + 1, \ell, j)$ for each $1 \leq \ell \leq k$ (note that here the second domain type parameter varies instead of the third). For each $\mathbf{u}_{i,j}$ and $\mathbf{s}_{i,j}$ monomer, there are 2^{i-1} copies. We can conceptually break these monomers into n “layers”, each consisting of all monomers with the same value for their first parameter. The analyte we wish to detect, \mathbf{a} , is a monomer that has one copy of each unstarred domain $(1, j, \ell)$, $1 \leq j, \ell \leq k$.

Conceptually, when the analyte is absent, the most efficient way for all starred sites on each $\mathbf{s}_{i,j}$ to be covered is by the unstarred sites on a corresponding $\mathbf{u}_{i,j}$, as seen in Figure 2. Although the TBN model is purely thermodynamic, we can conceptualize that when the analyte is added, its signal can propagate “kinetically” through each layer. In the first layer, it can “displace” the k different $\mathbf{u}_{1,j}$ monomers and bind to all of the $\mathbf{s}_{1,j}$ monomers. In doing so, it brings together all the unstarred sites on all of the $\mathbf{s}_{1,j}$ monomers. Having been brought together, these sites “look like” two copies of the analyte, but with the domains from layer 2 instead of layer 1. Thus, this polymer is then able to displace *two* copies of each $\mathbf{u}_{2,j}$ from their corresponding $\mathbf{s}_{2,j}$ monomers, thus bringing all of the $\mathbf{s}_{2,j}$ together. This in turn now looks like four copies of the analyte for the domains in the third layer, and so on. Each layer allows this polymer to assimilate exponentially more $\mathbf{s}_{i,j}$, thus freeing exponentially many $\mathbf{u}_{i,j}$. Each of these displacement steps involves an equal number of splits and merges.

3.2 Convergence Process

So far, the TBN described has exactly one stable configuration before adding the analyte, and it performs the task of amplifying signal by having the potential to change its state exponentially when the analyte is added. However, there is also a stable configuration after adding the analyte in which nothing else changes, and many others in which only a small amount of change occurs. We must guarantee that the analyte’s signal “propagates” through all of the layers.

To design the system to meet this requirement, we observe that all exponentially many monomers that have been brought together must contribute to some singular change in the system that gains some entropy, to spur the signal into propagating. The typical way to accomplish this in a TBN is by having monomers that have been brought together displace a larger number of monomers from some complex at the cost of a smaller number of merges. Because the pre-analyte TBN has an entropy gap of $k - 1$ in this design so far, we can afford to give the TBN with the analyte an “entropic payoff” of $\frac{k}{2}$. When the analyte is absent, this payoff is weak enough that there will still be an entropy gap of $\frac{k}{2} - 1$; when the analyte is present, the existence of this payoff will force the signal to fully propagate, and will give the TBN with the analyte an entropy gap of $\frac{k}{2} - 1$ by making it so that any configurations in which this payoff is not achieved are also far away from stable.

Another challenge is that we cannot simply detect all our exponentially many conjoined monomers by binding them all to a single exponentially large monomer, because we need to bound the size of the largest monomer in the system. Our conceptual strategy for overcoming this is as follows: the signal will converge in much the same way as it was amplified. In the amplification step, one set of domains coming together in one layer was enough to cause two of them to come together in the next layer. In this convergence step, two sets of domains in one layer will have to converge together to activate one set in the next layer. This convergence ends in bringing together a set of binding sites that is of the same size as the analyte, which can then directly displace some monomers to gain $\frac{k}{2}$ total polymers.

We now fully define $T_{n,k}$. We start with the already described $\mathbf{u}_{i,j}$ and $\mathbf{s}_{i,j}$. To these, we first add monomer types $\mathbf{u}'_{i,j}$ and $\mathbf{s}'_{i,j}$ for $1 \leq i \leq n, 1 \leq j \leq k$. These monomers are the “converging” equivalents of $\mathbf{u}_{i,j}$ and $\mathbf{s}_{i,j}$. Conceptually, they will activate in the reverse order: two copies of each $\mathbf{s}'_{i,j}$ for $1 \leq j \leq k$, when brought together, will be able to bring together one copy of each $\mathbf{s}'_{i-1,j}$ for $1 \leq j \leq k$.

Each $\mathbf{u}'_{i,j}$ monomer has $2k$ unstarred domains: two copies each of domains $(i+1, j, \ell)'$ for each $1 \leq \ell \leq k$. Each $\mathbf{s}'_{i,j}$ has a starred copy of each of the $2k$ domains in $\mathbf{u}'_{i,j}$ and additionally has one unstarred domain $(i, \ell, j)'$ for $1 \leq \ell \leq k$ (note again here that the second domain type parameter varies instead of the third). One exception is the monomers $\mathbf{u}'_{n,j}$ and $\mathbf{s}'_{n,j}$ (the first ones to activate) which use domains $(n+1, j, \ell)$ and $(n+1, \ell, j)$ instead of $(n+1, j, \ell)'$ and $(n+1, \ell, j)'$ respectively so that they can interact with $\mathbf{s}_{n,j}$ monomers that have been brought together. Each monomer $\mathbf{u}'_{i,j}$ and $\mathbf{s}'_{i,j}$ has 2^{i-1} copies.

Finally, we add “payoff” monomers that will yield an entropic gain of $\frac{k}{2}$ when the signal from the analyte has cascaded through every layer. This choice of $\frac{k}{2}$ is arbitrary – a similar design works for any integer between 1 and k . Choosing a higher value leads to a higher entropy gap after adding the analyte and a lower entropy gap before adding it, and vice versa choosing a lower value. For simplicity of definitions we will assume k is even (though figures are shown with $k = 3$, which shows how to generalize to odd k).

We add one monomer \mathbf{p}^* , which contains the k^2 sites $(1, \ell_1, \ell_2)^{/*}$ for $1 \leq \ell_1, \ell_2 \leq k$. Note that this monomer can be replaced with k monomers of size k (in which case \mathbf{a} would be the only monomer with more than $3k$ domains), but doing so makes the proof more complex. The idea is that when all $\mathbf{s}'_{1,j}$ monomers are already together (as they can be “for free” when \mathbf{a} is present), they can cover \mathbf{p}^* in one merge; if they are apart, this requires k merges. In order to make this favorable to happen when they’re already together but unfavorable when they’re initially apart, we add another way to cover \mathbf{p}^* that takes $\frac{k}{2}$ merges. This is accomplished via monomers \mathbf{p}_j for $1 \leq j \leq \frac{k}{2}$. Each \mathbf{p}_j contains the $2k$ sites $(1, 2j-1, \ell)$ and $(1, 2j, \ell)$ for $1 \leq \ell \leq k$. We can interpret this geometrically as \mathbf{p}^* being a square, the $\mathbf{s}'_{1,j}$ covering it by rows, and the \mathbf{p}_j covering it by two columns at a time. This completes the definition of $T_{n,k}$. Recall $T_{n,k}^{\mathbf{a}}$ is $T_{n,k}$ with one added copy of \mathbf{a} .

► **Lemma 3.2.** $T_{n,k}$ has exactly one stable configuration $\sigma_{n,k}$.

► **Corollary 3.3.** $T_{n,k}$ has an entropy gap of $\frac{k}{2} - 1$.

► **Lemma 3.4.** $T_{n,k}^{\mathbf{a}}$ has exactly one stable configuration $\sigma_{n,k}^{\mathbf{a}}$, and $T_{n,k}^{\mathbf{a}}$ has an entropy gap of $\frac{k}{2} - 1$.

Proofs of these results are left to the appendix.

These results together complete the proof of Theorem 3.1: each of the more than 2^n \mathbf{u} and \mathbf{u}' monomers (which serve as reporters) are bound in $\sigma_{n,k}$ and unbound in $\sigma_{n,k}^{\mathbf{a}}$, implying their distance is more than 2^n . The largest monomer is \mathbf{a} with k^2 domains, and there are $(2n+1)k^2$ domain types and $4nk$ monomer types for the $\mathbf{s}_{i,j}$, $\mathbf{u}_{i,j}$, $\mathbf{s}'_{i,j}$, and $\mathbf{u}'_{i,j}$, plus $2 + \frac{k}{2}$ more for \mathbf{a} , \mathbf{p}^* , and \mathbf{p}_j .

3.3 Avoiding Large Polymer Formation

The TBN $T_{n,k}^a$ will, in the process of amplifying the signal of the analyte, form a single polymer of exponential size. This isn't an issue in the theoretical TBN model, but it is a practical issue because there is no way to design these monomers so that this large polymer would form.²

This can be solved by adding “translator gadgets”. These gadgets’ job is to mediate between consecutive layers. Instead of monomers from one layer directly binding to monomers from the next layer, they can split apart these translator gadgets with half of the gadget going to each layer. In exchange, the TBN will no longer have exactly one stable configuration when the analyte is present, as in the TBN model, the use of these translator gates will be purely “optional”.

We define a new TBN $\tilde{T}_{n,k}$ (as well as $\tilde{T}_{n,k}^a$, which is obtained by adding the analyte \mathbf{a}). We start with the TBN $T_{n,k}$. To assist with the amplification step, we add monomer types \mathbf{g}_i and \mathbf{g}_i^* for each $2 \leq i \leq n$. Each \mathbf{g}_i consists of one copy of each unstarred domain (i, j, ℓ) for each $1 \leq j, \ell \leq k$. Each \mathbf{g}_i^* consists of the same domains but all starred. Each of these monomers has 2^{i-1} copies. The use of these gadgets can be seen in Figure 4.

To assist with the convergence step, we add monomer types \mathbf{h}_i and \mathbf{h}_i^* for each $2 \leq i \leq n+1$. Each \mathbf{h}_i has *two* copies of each unstarred domain $(i, j, \ell)'$ for each $1 \leq j, \ell \leq k$. Each \mathbf{h}_i^* has only one copy of each of the corresponding starred domains. There are 2^{i-1} copies of each \mathbf{h}_i and 2^i copies of each \mathbf{h}_i^* . The use of these gadgets can be seen in Figure 5.

► **Theorem 3.5.** *Let $\tilde{T} = \tilde{T}_{n,k}$ and $\tilde{T}^a = \tilde{T}_{n,k}^a$ be as described. Then:*

1. \tilde{T} has exactly one stable configuration $\tilde{\sigma}_{n,k}$, and $d(\tilde{T}, \tilde{T}^a) > 2^n$.
2. \tilde{T} has an entropy gap of $\frac{k}{2}$, and \tilde{T}^a has the property that all of its configurations α that are within distance to stability $\frac{k}{2}$ satisfy $d(\tilde{\sigma}_{n,k}, \alpha) > 2^n$.
3. $\tilde{T} = \tilde{T}_{n,k}$ uses $\mathcal{O}(nk)$ total monomer types, $\mathcal{O}(nk^2)$ domain types, and $\mathcal{O}(k^2)$ domains per monomer.
4. The unique stable configuration of \tilde{T} has $\mathcal{O}(k)$ monomers in its largest polymer. There is a stable configuration of \tilde{T}^a sharing this property.

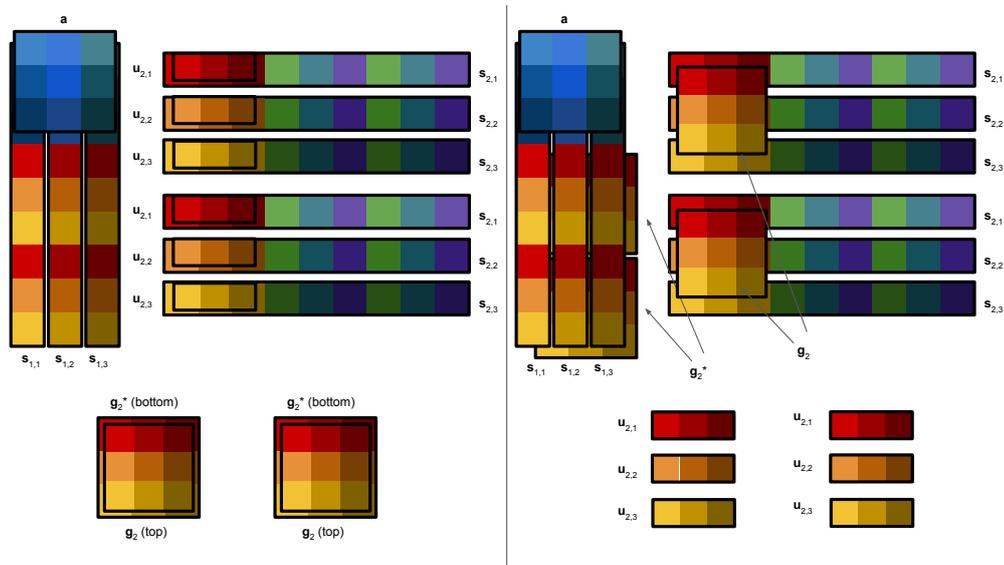
Compared to Theorem 3.1, this theorem trades away the condition that both TBNs have only a single stable configuration in exchange for the post-analyte TBN having a configuration with $\mathcal{O}(k)$ monomers per polymer, whereas the previous construction has roughly $k \cdot 2^n$ monomers in a single polymer.

The second condition is somewhat complex. This complexity’s necessity is explained by Figure 5. In that figure, if we propagate signal without using the translator gadget, we arrive at a configuration that is saturated but has only one fewer complex than a stable configuration. However, such near-stable configurations are still very different from the stable configuration of $\tilde{T}_{n,k}$, so it is still possible to distinguish the two TBNs with an amplification factor proportional to 2^n and a resilience to false positives and negatives proportional to k .

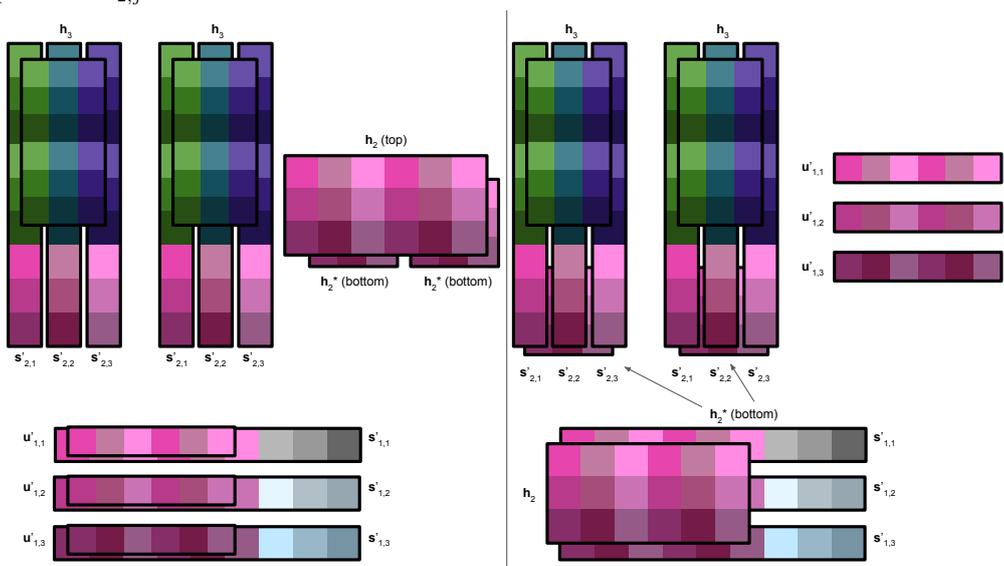
The proof of this theorem is very similar to that of Theorem 3.1, and is also left to the appendix.

² The binding graph of the monomers within this giant polymer contains many complete k -ary trees of depth n as subgraphs. If each of the nodes of this graph is a real molecule that takes up some volume, it will be impossible to embed the whole graph within 3-dimensional space as n grows.

8:12 Thermodynamically Driven Signal Amplification



■ **Figure 4** The amplifying translator gadget, before and after it is triggered to propagate the signal forward by one layer. When $s_{1,1}$, $s_{1,2}$ and $s_{1,3}$ have been brought together, instead of directly replacing all the $u_{2,j}$ monomers, they can split two $\{g_2, g_2^*\}$ complexes, and the g_2 monomers can replace the $u_{2,j}$ monomers.



■ **Figure 5** The converging translator gadget, before and after it is triggered to propagate the signal forward by one layer. When two copies each of $s'_{2,1}$, $s'_{2,2}$ and $s'_{2,3}$ have been brought together into two complexes, instead of directly replacing all the $u'_{1,j}$ monomers, they can split a $\{h_2, h_2^*, h_2^*\}$ complex, and the h_2 monomer can replace the $u'_{1,j}$ monomers.

Note that in this image, the only way to propagate the signal efficiently would be to use the translator gadget; not using it will be one unit of entropy less efficient, requiring 3 splits and 4 merges, showing that this TBN no longer has an entropy gap. If instead we hadn't used the translator gadgets in the previous layer, then all six $s'_{2,j}$ monomers in this image would be together in a single complex rather than on two separate complexes, in which case it would be equally efficient to either use this translator gadget or directly displace the $u'_{1,j}$.

4 Upper Limit on TBN Signal Amplification

In this section, we show the following theorem providing an upper bound on the distance between a TBN before and after adding a single copy of a monomer, showing that the distance is at most double-exponential in the “size” of the system:

► **Theorem 4.1.** *Let T be a TBN with d domain types, m monomer types, and at most a domains on each monomer. Let $n = \max\{d, m, a\}$. Let T' be T with one extra copy of some monomer. Then $d(T, T') \leq n^{8n^{7n^2}}$.*

Recall Definition 2.7 for the distance between TBNs. Essentially, this theorem is saying that adding a single copy of some monomer can only impact doubly exponentially many total polymers, no matter how many total copies of each monomer are in the TBN.

Our strategy for proving this theorem is to fix some ordering on polymer types, and bound the distance between the lexicographically earliest stable configuration of an arbitrary TBN under that ordering before and after adding a single copy of some monomer. To bound this distance, we cast the problem of finding stable configurations of a TBN as an integer program (IP), and use methods from the theory of integer programming value functions to give a bound on how much the solution to this IP can change given a small change in the underlying TBN.

Proof. A complete proof including all technical details can found in the full version of this paper on arxiv. Here we present only the main ideas of the proof. ◀

We first introduce a definition from [6] and some notation that was unnecessary in previous sections.

► **Definition 4.2.** *Given a (star-limiting) TBN T , the polymer basis of T , denoted $\mathcal{B}(T)$, is the set of polymers \mathbf{P} such that both of the following hold:*

- \mathbf{P} appears in some saturated configuration of a star-limiting TBN using the same monomer types as T .
- \mathbf{P} cannot be split into two or more self-saturated polymers.

The polymer basis is a useful construction because it is known to describe exactly those polymer types that may appear in stable configurations of T . It is always finite, and we will bound its size later.

Given a TBN T , let $M(T)$ denote its monomer types, and let $T(\mathbf{m})$ denote the count of monomer \mathbf{m} in T . Given a polymer \mathbf{P} and a monomer type $\mathbf{m} \in M(T)$, let $\mathbf{P}(\mathbf{m})$ represent the count of monomer \mathbf{m} in polymer \mathbf{P} .

Suppose for the rest of this section that we have a TBN T to which we wish to add a single copy of some monomer \mathbf{a} (which may or may not exist in T). Let T' be T with \mathbf{a} added.

4.1 Finding Stable Configurations via Integer Programming

Prior work [6] has shown that the problem of finding the stable configurations of a TBN can be cast as an IP. There are multiple different formulations; we will use a formulation that is better for the purpose of reasoning theoretically about TBN behavior.

Let $\{x_{\mathbf{P}} : \mathbf{P} \in \mathcal{B}(T)\}$ be variables each representing the count of polymer \mathbf{P} in a configuration of T . Then consider the following integer programming problem:

8:14 Thermodynamically Driven Signal Amplification

$$\begin{aligned}
\max \quad & \sum_{\mathbf{P} \in \mathcal{B}(T)} x_{\mathbf{P}} \\
\text{s.t.} \quad & \sum_{\mathbf{P} \in \mathcal{B}(T)} \mathbf{P}(\mathbf{m}) x_{\mathbf{P}} = T(\mathbf{m}) \quad \forall \mathbf{m} \in M(T) \\
& x_{\mathbf{P}} \in \mathbb{N} \quad \forall \mathbf{P} \in \mathcal{B}(T)
\end{aligned} \tag{2}$$

Intuitively, the linear equality constraints above express “monomer conservation”: the total count of each monomer in T should equal the total number of times it appears among all polymers.

The following was shown in [6]; for the sake of self-containment, we show it here as well

► **Proposition 4.3.** *The optimal solutions to the IP (2) correspond exactly to stable configurations of T .*

Proof. If the variables $x_{\mathbf{P}}$ form a feasible solution, then those counts of polymers are a valid configuration because they exactly use up all monomers. If the solution is optimal, then there is no saturated configuration with more polymers (as only polymers from $\mathcal{B}(T)$ can show up in stable configurations), so the configuration is stable. Conversely, if a configuration σ is stable then it can be translated into a feasible solution to the IP because it only uses polymers from $\mathcal{B}(T)$ and obeys monomer conservation. If there were a solution with a greater objective function, then this would translate to a configuration with more complexes that is still saturated (because all polymers in the polymer basis are self-saturated), contradicting the assumption of σ 's stability. ◀

We observe that adding an extra copy of some monomer to a TBN corresponds to changing the right-hand side of one of the constraints of this IP by one. Note that this is true even if we add a copy of some monomer for which there were 0 copies, as we may still include variables for polymers that contain that monomer in the former IP and simply consider there to be 0 copies of the monomer. Therefore, we are interested in sensitivity analysis of how quickly a solution to an IP can change as the right-hand sides of constraints change.

However, there is one edge case we must account for first. It is possible that T and T' have different polymer bases. This is because of the first requirement in Definition 4.2 requiring that the polymer basis respects that starred sites are limiting. If we add a single copy of a monomer, this may change which sites are limiting, if \mathbf{a} has more copies of a starred site than T had excess copies of the unstarred site. We cannot include variables for such polymers in the IP formulation without taking extra precautions, as if we do there may be optimal solutions that don't correspond to saturated configurations. Therefore, we will first account for how many copies of such a polymer T and T' may differ by:

► **Lemma 4.4.** *Suppose that some polymer \mathbf{P} is exactly one of $\mathcal{B}(T')$ and $\mathcal{B}(T)$. Then any saturated configuration of T' contains at most $|\mathbf{a}|$ copies of \mathbf{P} , where $|\mathbf{a}|$ denotes the number of sites on \mathbf{a} .*

Note that this result is slightly surprising – one natural way that one might try to design a TBN that amplifies signal is by designing the analyte so that it intentionally flips which sites are limiting. This result shows that this is an ineffective strategy: going from 5 excess copies of some site a to 5 excess copies of a^* is seemingly no more helpful in instigating a large change than going from 60 excess copies of a to 50.

Proof. If \mathbf{P} is in $\mathcal{B}(T')$ but not $\mathcal{B}(T)$, it must contain an excess of a starred site that was limiting in T , but is no longer limiting in T' . We see this because \mathbf{P} necessarily occurs in a saturated configuration of the TBN containing precisely the monomers that it is composed of; therefore, in order to not be in $\mathcal{B}(T)$, by definition of the polymer basis, it must be the case that this TBN has different limiting sites than T' .

Let a denote some such site type, so that a^* is limiting in T and a is limiting in T' , and \mathbf{P} contains an excess of a^* . Then \mathbf{a} must contain an excess of a^* , but it cannot contain more than $|\mathbf{a}|$ excess copies. Therefore, there are at most this many *total* excess copies of a^* in T' . It follows that if there are more than $|\mathbf{a}|$ copies of \mathbf{P} in a configuration of T' , then those copies of \mathbf{P} collectively have more excess copies of a^* than T' does, so some other polymer in that configuration would have to have an excess of a . This implies that such a configuration is not saturated (and therefore also cannot be stable). An identical argument shows that the same is true for polymers in $\mathcal{B}(T)$ but not $\mathcal{B}(T')$. ◀

In order to analyze and compare the two IP instances, we need them to have the same variable set. Therefore, we will include variables for all polymers from both polymer bases in both IP formulations. Let $\mathcal{B}(T, T') = \mathcal{B}(T) \cup \mathcal{B}(T')$ denote this merged polymer basis, and let $P = |\mathcal{B}(T, T')|$ denote the total number of possible polymers we must consider, or equivalently the number of variables we will have in these IPs. In each IP, we will have a constraint on each variable representing a polymer not in the relevant polymer basis, that says that that variable must equal zero.

4.2 Sensitivity Analysis

This sensitivity analysis problem of how IPs change as the right-hand sides of constraints change was studied by Blair and Jeroslow in [1]. We will not need their full theory, but we will use some of their results and methods.

In Corollary 4.7 of [1], they show that there is a constant K_3 , independent of the right-hand sides of constraints (in our case, independent of how many copies of each monomer exist) such that:

$$R_c(v) \leq G_c(v) \leq R_c(v) + K_3, \tag{3}$$

where $G_c(v)$ gives the optimal value of the objective function c of a minimization IP as a function of the vector v of right-hand sides of constraints, and $R_c(v)$ gives the optimal value of the same problem when relaxing the constraint that variables must have integer values. The objective function we've shown so far is to maximize the sum of polymer counts rather than minimize, but the same statement applies that the integer and real-valued optimal solutions differ by at most K_3 . In defining K_3 , they also show the existence of a constant M_1 such that

$$|R_c(v) - R_c(w)| \leq M_1 \|v - w\|, \tag{4}$$

where v and w are different vectors for the right-hand sides of constraints. Note that we take all norms as 1-norms. Combining these inequalities, we see that

$$|G_c(v) - G_c(w)| \leq M_1 \|v - w\| + K_3. \tag{5}$$

For example, if we want to know the difference between the total number of polymers in a stable configuration before and after adding one copy of a monomer (and if the polymer bases of T and T' are identical), then we care about increasing one element of v by 1, so our bound on this difference is $M_1 + K_3$. This statement applies to maximization and minimization problems.

4.3 From Optimal Values to Polymer Counts

For ease of analysis, we order the polymers in $\mathcal{B}(T, T')$ as follows: first we list all the polymers that are not in $\mathcal{B}(T)$, then all the polymers that are not in $\mathcal{B}(T')$, then all the polymers in $\mathcal{B}(T) \cap \mathcal{B}(T')$. We need to show that the number of copies of each individual polymer does not change too much. We do this using a technique similar to Corollary 5.10 in [1].

Let P_{tot} be the total number of polymers in a stable configuration (either before or after adding \mathbf{a} , depending on which case we are examining).

We now define a new sequence of integer programs whose optimal values give polymer counts in the lexicographically earliest stable configuration under this ordering. We do this by finding the value of each variable $x_{\mathbf{P}}$ in order. This sequence of IP problems is defined separately for both TBNs, before and after adding \mathbf{a} .

For those variables representing a polymer that is in one basis but not the other, we do not need to analyze this IP, so we simply fix such a variable's value to whatever its value is in this lexicographically earliest stable configuration, which will be 0 in one TBN and bounded by $|\mathbf{a}|$ by Lemma 4.4 in the other.

Now, to find the value of some particular variable $x_{\mathbf{Q}}$ in either of the two TBNs where $\mathbf{Q} \in \mathcal{B}(T) \cap \mathcal{B}(T')$, suppose we have already found the value $y_{\mathbf{P}}$ we wish to fix $x_{\mathbf{P}}$ to for each $\mathbf{P} < \mathbf{Q}$ under our ordering. Then we define a new IP on all the same variables as follows:

$$\begin{aligned}
 \min \quad & x_{\mathbf{Q}} \\
 \text{s.t.} \quad & \sum_{\mathbf{P} \in \mathcal{B}(T, T')} x_{\mathbf{P}} = P_{tot} \\
 & \sum_{\mathbf{P} \in \mathcal{B}(T, T')} \mathbf{P}(\mathbf{m})x_{\mathbf{P}} = T(\mathbf{m}) \quad \forall \mathbf{m} \in m(T) \\
 & x_{\mathbf{P}} = y_{\mathbf{P}} \quad \forall \mathbf{P} < \mathbf{Q} \\
 & x_{\mathbf{P}} \in \mathbb{N} \quad \forall \mathbf{P} \in \mathcal{B}(T, T')
 \end{aligned} \tag{6}$$

By construction, this IP gives us the smallest possible value that $x_{\mathbf{Q}}$ can take on in a stable configuration (as all variables must sum to P_{tot}) in which all previous $x_{\mathbf{P}}$ have fixed values. Then this process gives us a sequence of IP problems that we can sequentially compare to bound the differences between the values of the individual polymer counts in these lexicographically earliest configurations. We can repeatedly apply Equation (5) to each $x_{\mathbf{P}}$ in turn, as each variable's value before and after adding \mathbf{a} will be given by the optimal value of (6) where the only differences are in the right-hand sides of constraints.

5 Conclusion

In this paper we have defined the signal amplification problem for Thermodynamic Binding Networks, and we have demonstrated a TBN that achieves exponential signal amplification. We also showed a doubly-exponential upper bound for the problem. As TBNs model mixtures of DNA, a TBN that amplifies signal can potentially be implemented as a real system. An upper bound has implications for how effective a system designed in this way can potentially be, and shows that there are some limitations for a purely thermodynamic approach to signal detection and amplification.

One clear direction for future work is to implement such a system. This would involve creating a design that accounts for the simplifications of the TBN model. In particular, enthalpy and entropy need to be strong enough with enthalpy sufficiently stronger than

entropy. Further, the polymers formed need to be geometrically feasible. We have done some work to make this problem geometrically realizable with the inclusion of translator gadgets in Section 3.3. In principle, the polymers that are formed in this version of the system are simple enough that they should form if the DNA strands implementing them are well-designed.

Another goal would be to bridge the gap between our singly exponential amplifier and doubly exponential upper bound by either describing a TBN that can amplify signal more than exponentially, or deriving a more precise upper bound. If one wished to construct a TBN with doubly exponential amplification, an examination of our upper bound proof will show that such a TBN must have an exponentially sized polymer basis, and most likely would need to actually use an exponential amount of different polymer types in its stable configurations either with or without the analyte. Such a design seems relatively unlikely to come to fruition, and it seems more likely that our proof technique or similar techniques can be tightened in order to show a stricter upper bound. Thus, we conjecture that the true upper bound is (singly) exponential.

There are also other types of robustness that we have not discussed in this work that merit further analysis. One of these is input specificity: the question of how well the system amplifies signal if the analyte is changed slightly. Another is sensitivity to the number of copies of each component. Intuitively, our system's behavior depends on having exactly equal numbers of complementary strands within each layer; if there are too many copies of one, it may result in those excess copies spuriously propagating or blocking signal to the next layer. This issue may be intrinsic to thermodynamic signal amplifiers, or there may be some system more robust to it. Lastly, it may be experimentally useful to show that our system achieves its stable states not only in the limit of thermodynamic equilibrium, but also more practically when annealed. Some systems such as HCR are designed to reach non-equilibrium, meta-stable states when annealed. We conjecture that our system should reach equilibrium when annealed, because kinetic traps in the system are far away from being thermodynamically stable (large entropy gap). Formally studying annealing could be done by analyzing versions of the TBN model with different tradeoffs between entropy and enthalpy to model different temperatures.

References

- 1 C. E. Blair and R. G. Jeroslow. The value function of an integer program. *Mathematical Programming*, 23(1):237–273, December 1982. doi:10.1007/BF01583794.
- 2 Keenan Breik, Cameron Chalk, David Doty, David Haley, and David Soloveichik. Programming substrate-independent kinetic barriers with thermodynamic binding networks. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 18(1):283–295, 2021. doi:10.1109/TCBB.2019.2959310.
- 3 Keenan Breik, Chris Thachuk, Marijn Heule, and David Soloveichik. Computing properties of stable configurations of thermodynamic binding networks. *Theoretical Computer Science*, 785:17–29, 2019.
- 4 Harry MT Choi, Maayan Schwarzkopf, Mark E Fornace, Aneesh Acharya, Georgios Artavanis, Johannes Stegmaier, Alexandre Cunha, and Niles A Pierce. Third-generation in situ hybridization chain reaction: Multiplexed, quantitative, sensitive, versatile, robust. *Development*, 145(12):dev165753, 2018.
- 5 David Doty, Trent A. Rogers, David Soloveichik, Chris Thachuk, and Damien Woods. Thermodynamic binding networks. In Robert Brijder and Lulu Qian, editors, *DNA Computing and Molecular Programming*, pages 249–266, Cham, 2017. Springer International Publishing.

8:18 Thermodynamically Driven Signal Amplification

- 6 David Haley and David Doty. Computing properties of thermodynamic binding networks: An integer programming approach. In Matthew R. Lakin and Petr Šulc, editors, *DNA 2021: Proceedings of the 27th International Meeting on DNA Computing and Molecular Programming*, volume 205 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 2:1–2:16, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.DNA.27.2.
- 7 Randolph Lopez, Ruofan Wang, and Georg Seelig. A molecular multi-gene classifier for disease diagnostics. *Nature chemistry*, 10(7):746–754, 2018.
- 8 Dionis Mineev, Christopher M Wintersinger, Anastasia Ershova, and William M Shih. Robust nucleation control via crisscross polymerization of highly coordinated DNA slats. *Nature communications*, 12(1):1741, 2021.
- 9 Gerald Schochetman, Chin-Yih Ou, and Wanda K. Jones. Polymerase chain reaction. *The Journal of Infectious Diseases*, 158(6):1154–1157, 1988. URL: <http://www.jstor.org/stable/30137034>.
- 10 Erhu Xiong, Dongbao Yao, Andrew D. Ellington, and Sanchita Bhadra. Minimizing leakage in stacked strand exchange amplification circuits. *ACS Synthetic Biology*, 10(6):1277–1283, 2021. PMID: 34006090. doi:10.1021/acssynbio.0c00615.

A Proofs for Section 3 (Signal Amplification TBN)

► **Lemma 3.2.** $T_{n,k}$ has exactly one stable configuration $\sigma_{n,k}$.

Proof. We consider merges to get from the melted configuration to any saturated configuration. We may order these merges such that we first make all the merges necessary to cover each individual $\mathbf{s}_{i,j}$ in increasing value of i , then each individual $\mathbf{s}'_{i,j}$ in decreasing value of i . We see that at each step of this process, we may cover the monomer in question by a single merge (of its corresponding $\mathbf{u}_{i,j}$ or $\mathbf{u}'_{i,j}$). If we never merge the corresponding \mathbf{u} monomer, the only other monomers that can cover the starred sites on a given $\mathbf{s}_{i,j}$ are k different $\mathbf{s}_{i-1,\ell}$ monomers. Likewise, the only other way to cover the starred sites on a given $\mathbf{s}'_{i,j}$ is by using k different $\mathbf{s}'_{i+1,\ell}$ monomers (except for $\mathbf{s}'_{n,j}$ which needs $\mathbf{s}_{n,\ell}$ monomers).

If an \mathbf{s} or \mathbf{s}' monomer is covered in multiple different ways, we order the merges such that it is first covered by one corresponding \mathbf{u} monomer (and then ignore any other merges for now, as we are still ordering the merges to cover each \mathbf{s} monomer sequentially). We see then that if every \mathbf{s} monomer is covered by a \mathbf{u} monomer, then no \mathbf{s} monomers will be brought together during this process. Therefore, the first time in this sequence that we choose to cover an \mathbf{s} without its corresponding \mathbf{u} will require k total merges to cover that \mathbf{s} . The resulting configuration is feed-forward, so by Corollary 2.11, reaching a stable configuration requires at least one more merge per remaining \mathbf{s} monomer. This results in at least $k - 1$ extra merges compared to covering \mathbf{s} and \mathbf{s}' monomers by using \mathbf{u} and \mathbf{u}' monomers respectively.

Once all \mathbf{s} and \mathbf{s}' monomers are covered, the only other monomer with starred sites is \mathbf{p}^* , so we can make all the merges that are needed to cover it. If none of the $\mathbf{s}'_{1,j}$ monomers have been brought together, then the fewest merges it takes to cover \mathbf{p}^* is $\frac{k}{2}$, via the \mathbf{p}_j monomers. If any of them have been brought together, then it could potentially take a single merge to cover \mathbf{p}^* . However, this would have required $k - 1$ extra merges at some point during the covering of \mathbf{s} monomers, resulting in $\frac{k}{2}$ extra total merges compared to covering all \mathbf{s} monomers with \mathbf{u} monomers, then covering \mathbf{p}^* with \mathbf{p}_j monomers.

Therefore, this latter set of merges covers all starred sites in as few merges as possible, and therefore gives the unique stable configuration of $T_{n,k}$. ◀

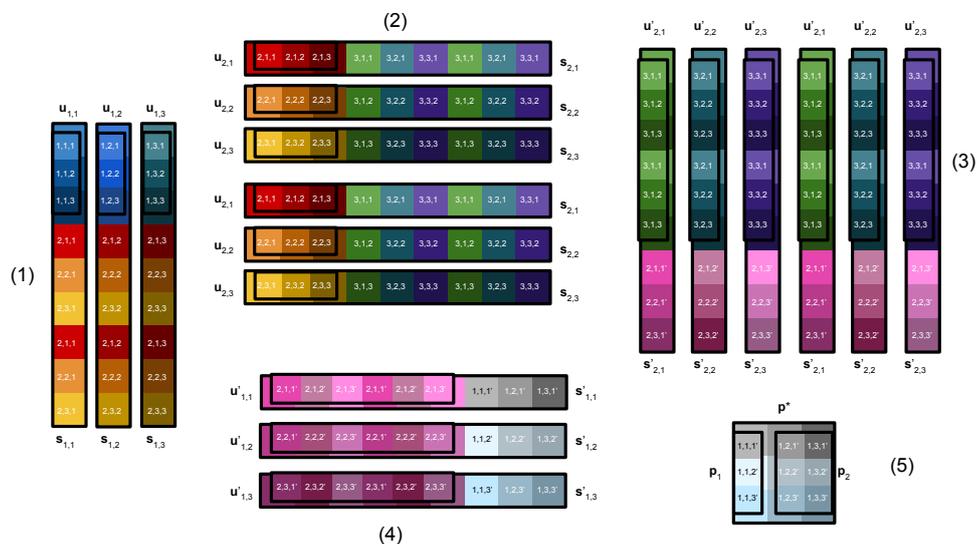
► **Corollary 3.3.** $T_{n,k}$ has an entropy gap of $\frac{k}{2} - 1$.

Proof. Recall Definition 2.8 for what we must show. Any saturated configuration that does not make all the merges in $\sigma_{n,k}$ must either have some \mathbf{s} that is not covered by its corresponding \mathbf{u} (resulting in at least $\frac{k}{2}$ extra merges, as per the above argument), or must cover \mathbf{p}^* with initially-separate $\mathbf{s}'_{1,j}$ monomers (resulting in $\frac{k}{2}$ extra merges). Thus, any such configuration has distance to stability at least $\frac{k}{2}$. Any other saturated configuration that does make all of the merges in this sequence simply makes some extra merges afterward, and therefore splits to $\sigma_{n,k}$. It follows that $T_{n,k}$ has an entropy gap of $\frac{k}{2}$ (and also of $\frac{k}{2} - 1$, for consistency in the statement of Theorem 3.1). ◀

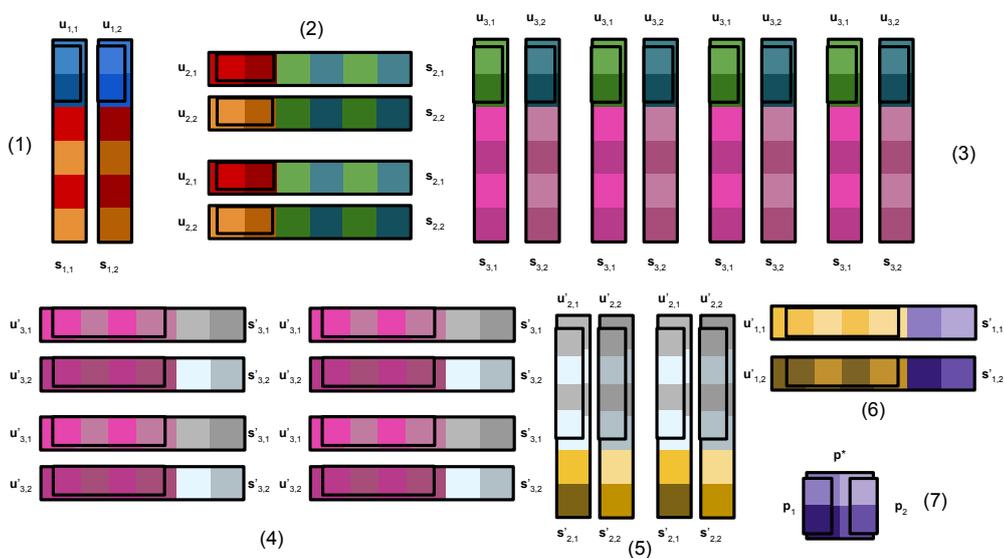
► **Lemma 3.4.** $T_{n,k}^{\mathbf{a}}$ has exactly one stable configuration $\sigma_{n,k}^{\mathbf{a}}$, and $T_{n,k}^{\mathbf{a}}$ has an entropy gap of $\frac{k}{2} - 1$.

Proof. We see that $T_{n,k}^{\mathbf{a}}$ (like $T_{n,k}$) is feed-forward (recall Definition 2.9) by first ordering \mathbf{a} along with all the $\mathbf{u}_{i,j}$, $\mathbf{u}'_{i,j}$, and \mathbf{p}_j monomers (none of which have starred sites), then all the $\mathbf{s}_{i,j}$ in increasing order of i , then all the $\mathbf{s}'_{i,j}$ in decreasing order of i , and finally \mathbf{p}^* .

Unlike $T_{n,k}$, however, we may reach a stable state by merging \mathbf{a} together with every single $\mathbf{s}_{i,j}$, every single $\mathbf{s}'_{i,j}$ and \mathbf{p}^* into a single polymer. This covers all starred sites, and requires



■ **Figure 6** A version of Figure 2 with text labels on domains, for accessibility and allowing comparison with the domains as defined in the text of the paper. This figure shows the unique stable configuration of $T_{2,3}$.



■ **Figure 7** The unique stable configuration $\sigma_{3,2}$ of $T_{3,2}$. Compared to Figure 2 (or Figure 6 above), which show $T_{2,3}$, this figure shows one more layer and a smaller entropy gap parameter. The additional layer means that there are 4 copies of each monomer in the largest parts of the figure, compared to 2 copies of each monomer in the other figures; if another layer were added, it would contain 8 copies of each monomer. The smaller entropy gap parameter manifests in this figure visually having a “2 by 2 grid” design motif, compared to the “3 by 3 grid” motif in the other figures.

exactly one merge per monomer with starred sites, so by Corollary 2.12, this configuration $\sigma_{n,k}^a$ is stable.

Now, we examine an arbitrary saturated configuration σ of $T_{n,k}^a$. We consider merges in essentially the opposite order of how they were considered when analyzing $T_{n,k}$. First, consider \mathbf{p}^* . It must be covered either by all the \mathbf{p}_j monomers, or by all the $\mathbf{s}'_{1,j}$ monomers. If we merge all the \mathbf{p}_j monomers to \mathbf{p}^* , we arrive at a configuration that is still feed-forward, but has only one fewer polymer with uncovered starred sites compared to $\text{melt}(T_{n,k}^a)$ in spite of making $\frac{k}{2}$ merges. Therefore, by Lemma 2.10, reaching a saturated configuration from this point requires at least $\frac{k}{2} - 1$ extra merges compared to $\sigma_{n,k}^a$.

Now, we may make a similar argument for all \mathbf{s} monomers in the opposite order that we considered them in Lemma 3.2. First, either we have already made $\frac{k}{2} - 1$ extra merges, or the $\mathbf{s}'_{1,j}$ monomers have all been brought together on a single polymer to cover \mathbf{p}^* . If we now make all the merges necessary to cover all starred sites on this polymer, we must do so either using all the $\mathbf{u}'_{1,j}$ or by using all the $\mathbf{s}'_{2,j}$. If we use the former, then this will require k total merges but will only reduce the count of polymers with starred sites by 1. The resulting configuration is still feed-forward, so again by Lemma 2.10 any saturated configuration we reach from this point will require at least $k - 1$ extra merges compared to σ . Otherwise, we must bring all the $\mathbf{s}_{2,j}$ monomers together to cover these sites. This does not fall victim to the same argument, because bringing these monomers with starred sites together onto the same polymer lowers the total number of polymers with uncovered starred sites. Now that they have been brought together, the same argument shows that we must either cover all the starred sites on the $\mathbf{s}'_{2,j}$ using all the $\mathbf{s}'_{3,j}$, or suffer $k - 1$ extra merges. The same argument for each layer in the converging part of the TBN also works for each layer in the amplifying part. Finally, after running through this argument we arrive at all $\mathbf{s}_{1,j}$ being brought together, which can be covered either by a single merge of \mathbf{a} or by merging the k $\mathbf{u}_{1,j}$ to it.

Overall, this shows that any saturated configuration of $T_{n,k}^a$ either makes all of the merges in $\sigma_{n,k}^a$ or it must make at least $\frac{k}{2} - 1$ extra merges. It follows that $\sigma_{n,k}^a$ is the unique stable configuration of $T_{n,k}^a$, with an entropy gap of $\frac{k}{2} - 1$ as desired. \blacktriangleleft

► **Theorem 3.5.** *Let $\tilde{T} = \tilde{T}_{n,k}$ and $\tilde{T}^a = \tilde{T}_{n,k}^a$ be as described. Then:*

1. \tilde{T} has exactly one stable configuration $\tilde{\sigma}_{n,k}$, and $d(\tilde{T}, \tilde{T}^a) > 2^n$.
2. \tilde{T} has an entropy gap of $\frac{k}{2}$, and \tilde{T}^a has the property that all of its configurations α that are within distance to stability $\frac{k}{2}$ satisfy $d(\tilde{\sigma}_{n,k}, \alpha) > 2^n$.
3. $\tilde{T} = \tilde{T}_{n,k}$ uses $\mathcal{O}(nk)$ total monomer types, $\mathcal{O}(nk^2)$ domain types, and $\mathcal{O}(k^2)$ domains per monomer.
4. The unique stable configuration of \tilde{T} has $\mathcal{O}(k)$ monomers in its largest polymer. There is a stable configuration of \tilde{T}^a sharing this property.

Proof. Recall the constructions of $\tilde{T}_{n,k}$ and $\tilde{T}_{n,k}^a$ from Section 3.3. Our argument will be very similar to that of Theorem 3.1 (i.e., the above lemmas), except we need to account for the extra monomer types.

First, consider $\tilde{T}_{n,k}$, where \mathbf{a} is absent. We wish to show that its stable configuration looks like that of $T_{n,k}$, with the added \mathbf{g} and \mathbf{h} monomers only binding to added \mathbf{g}^* and \mathbf{h}^* monomers respectively. We order the merges to get to a saturated configuration in essentially the same order as we did in analyzing $T_{n,k}$: first we will make all merges necessary to cover all $(1, j, \ell)^*$ sites, then $(2, j, \ell)^*$, and so on up to $(n + 1, j, \ell)^*$, then $(n, j, \ell)^*$, and so on. As before, at each step, we will see that we cannot make merges in any way other than those in the desired stable configuration without needing $k - 1$ extra merges for that step.

For $(i, j, \ell)^*$ sites, at each step, there is exactly one way to cover all starred sites by making one merge per monomer with these starred sites: we cover each $\mathbf{s}_{i,j}$ with a $\mathbf{u}_{i,j}$ and each \mathbf{g}_i^* with a \mathbf{g}_i . In particular, we already know from the proof of Lemma 3.2 that this is true for the $\mathbf{s}_{i,j}$ if we only use $\mathbf{s}_{i-1,j}$ and $\mathbf{u}_{i,j}$ to cover it, and that we will otherwise need to make $k - 1$ extra merges. Clearly we also cannot cover \mathbf{g}_i^* with anything other than \mathbf{g}_i without making k merges to cover it (and thus $k - 1$ extra merges), so we cannot use \mathbf{g}_i to cover $\mathbf{s}_{i,j}$.

Likewise, for $(i, j, \ell)^*$ sites, we only need to observe that each \mathbf{h}_i^* monomer can only be covered in a single merge by \mathbf{h}_i , so any other way of making merges necessarily involves $k - 1$ extra merges. So by the same argument as in Lemma 3.2 and Corollary 3.3, $\tilde{T}_{n,k}$ has exactly one stable configuration with an entropy gap of $\frac{k}{2}$. This configuration has $1 + \frac{k}{2}$ monomers in the polymer containing \mathbf{p}^* and all the \mathbf{p}_j , 3 monomers in each $\{\mathbf{h}_i, \mathbf{h}_i^*, \mathbf{h}_i^*\}$ polymer, and 2 monomers in each other polymer.

Now, consider $\tilde{T}_{n,k}^a$, where \mathbf{a} is present. If we take the stable configuration of $T_{n,k}^a$ and simply put all \mathbf{g} monomers into $\{\mathbf{g}_i, \mathbf{g}_i^*\}$ polymers, and all the \mathbf{h} monomers into $\{\mathbf{h}_i, \mathbf{h}_i^*, \mathbf{h}_i^*\}$ polymers, we have still made exactly one merge per monomer with any starred sites, so by Corollary 2.12 it is stable. If we then carry out the shifts described in Figure 4 and Figure 5, an equal number of merges and splits are made at each step, so the resulting saturated configuration is still stable. Additionally, in this configuration, the largest polymers have $k + 3$ monomers (specifically, those containing a set of $\mathbf{s}_{i,j}$ along with one copy of \mathbf{g}_i and two copies of \mathbf{g}_{i+1}^*).

All that remains to show is that all configurations of $\tilde{T}_{n,k}^a$ that are within $\frac{k}{2}$ distance to stability have exponentially many different polymers from the stable configuration of $\tilde{T}_{n,k}$. We will do this by showing that all \mathbf{u} and \mathbf{u}' monomers are free in all such configurations.

Again, this argument is very similar to the argument without the translator gadgets in Lemma 3.4. We consider merges to cover starred sites in the opposite order of the above argument for $\tilde{T}_{n,k}$. First, consider the merges necessary to cover all the $(1, j, \ell)'$ starred sites (on \mathbf{p}^*). Like before, they must be covered by either all the $\mathbf{u}'_{1,j}$ monomers or all the \mathbf{p}_j monomers, but using the latter gives a feed-forward configuration in which $\frac{k}{2} - 1$ extra merges have already been made. Thus, to be within $\frac{k}{2}$ distance to stability, we must use the $\mathbf{s}'_{1,j}$. Next, for the $(2, j, \ell)'$ starred sites, with the merges already made, there are two copies of each of these sites all together on the polymer containing all the $\mathbf{s}'_{1,j}$, and one copy of each site on each of the two \mathbf{h}_2^* monomers. If we are to merge any $\mathbf{u}'_{1,j}$ monomers to any of these in such a way that they cannot be split off without the result still being saturated, then we must merge all of the $\mathbf{u}'_{1,j}$ into one polymer. Like with the argument for $\tilde{T}_{n,k}^a$, we see that this results in $k - 1$ extra merges compared to a stable configuration. Thus, we cannot use any $\mathbf{u}'_{1,j}$, and these sites must be covered by the $\mathbf{s}'_{2,j}$ and \mathbf{h}_2 monomers.

We may do this either by using the \mathbf{h}_2 to cover both \mathbf{h}_2^* (in effect, not using the translator gadget) or by using \mathbf{h}_2 to cover all the $\mathbf{s}'_{1,j}$. The only difference in terms of the argument is that in the former case *all* of the $\mathbf{s}_{2,j}$ will be brought together in a single polymer, and in the latter case they will be split between two polymers. In the former case, it may require one extra merge to use translator gates in the next layer; however, either way, the same argument on each other layer in sequence shows that we cannot use any $\mathbf{u}'_{i,j}$ monomers without suffering $\frac{k}{2} - 1$ extra merges. Likewise, the exact same argument shows that the same thing is true of $\mathbf{u}_{i,j}$ monomers, necessitating that in any configuration that makes fewer than $\frac{k}{2} - 1$ extraneous merges, all exponentially many \mathbf{u} and \mathbf{u}' monomers must be free, as desired. ◀