

Programming Substrate-Independent Kinetic Barriers with Thermodynamic Binding Networks

Keenan Breik¹, Cameron Chalk¹, David Doty², David Haley², and David Soloveichik¹

¹ University of Texas at Austin

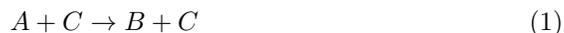
² University of California, Davis

Abstract. Engineering molecular systems that exhibit complex behavior requires the design of kinetic barriers. For example, an effective catalytic pathway must have a large barrier when the catalyst is absent. While programming such energy barriers seems to require knowledge of the specific molecular substrate, we develop a novel substrate-independent approach. We extend the recently-developed model known as thermodynamic binding networks, demonstrating programmable kinetic barriers that arise solely from the thermodynamic driving forces of bond formation and the configurational entropy of forming separate complexes. Our kinetic model makes relatively weak assumptions, which implies that energy barriers predicted by our model would exist in a wide variety of systems and conditions. We demonstrate that our model is robust by showing that several variations in its definition result in equivalent energy barriers. We apply this model to design catalytic systems with an arbitrarily large energy barrier to uncatalyzed reactions. Our results yield robust amplifiers using DNA strand displacement, a popular technology for engineering synthetic reaction pathways, and suggest design strategies for preventing undesired kinetic behavior in a variety of molecular systems.

1 Introduction

Abstract mathematical models of molecular systems, such as chemical reaction networks, have long been useful in *natural* science to study the properties of natural molecules. With recent experimental advances in synthetic biology and DNA nanotechnology [1, 3, 8, 9], such models have come to be viewed also as *programming languages* for describing the desired behavior of synthetic molecules.

We can describe a chemical program with abstract chemical reactions such as



In particular, a program may require (1) and forbid (2). But what remains hidden at this level of abstraction is a well-known chemical constraint: if (1) is possible, then (2) must also be, no matter the exact substances. Knowing this, we might try to slow (2) by ensuring B has high free energy. But then $B + C$ must also

have high free energy, so (1) slows in tandem. The only option to slow (2) but not (1) is to use a *kinetic barrier*: designing A so that, although it is possible for A to reconfigure into B , the system must traverse a higher energy (less favorable) intermediate in the absence of C .

To develop a substrate-independent approach to engineering kinetic barriers we need to rely on a universal thermodynamic property that would be relevant in a wide variety of chemical systems. We focus on the entropic penalty of association (decreasing the number of separate complexes). Intuitively, the entropic penalty is due to decreasing the number of microstates corresponding to the independent three-dimensional positions of each complex (configurational entropy). This thermodynamic penalty can be made dominant compared with other factors by decreasing the concentration.

To formalize this entropic penalty, we use the *thermodynamic binding networks* (TBN) model [5]. TBNs represent molecules as abstract *monomers* with binding sites that allow them to bind to other monomers. For a configuration γ , the TBN model defines $H(\gamma)$ as the number of bonds formed, and $S(\gamma)$ as the number of free complexes,³ and the *energy* $E(\gamma) = -wH(\gamma) - S(\gamma)$ as a (negative) weighted sum of the two.⁴ To be applicable to a wide variety of chemical systems, the TBN model does not impose geometric constraints on bonding (monomers are simply multisets of binding sets). Implementation of TBNs requires choosing a concrete physical substrate and geometric arrangement that permits the desired configurations to form.

We augment the TBN model with a notion of kinetic paths (changes in configuration) due to merging of different complexes and splitting them up (and in the full version of this paper, making, breaking, or exchanging bonds). This gives rise to a notion of *paths* of configurations, with different energies. Define

³The quantities $H(\gamma)$ and $S(\gamma)$ are meant to evoke the thermodynamic quantities of enthalpy and entropy, although the mapping is not exact. Indeed, there are other contributions to physical entropy besides the number of separate complexes, and the free energy contribution of forming additional bonds typically contains substantial enthalpic and entropic parts.

⁴In typical DNA nanotechnology applications, the Gibbs free energy $\Delta G(\gamma)$ of a configuration γ can be estimated as follows. Bonds correspond to domains of length l bases, and forming each base pair is favorable by $\Delta G_{\text{bp}}^{\circ}$. Thus, the contribution of $H(\gamma)$ to $\Delta G(\gamma)$ is $(\Delta G_{\text{bp}}^{\circ} \cdot l)H(\gamma)$. 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 $S(\gamma)$ to $\Delta G(\gamma)$ is $(\Delta G_{\text{assoc}}^{\circ} + RT \ln(1/C))(|\gamma| - S(\gamma))$, where $|\gamma|$ is the total number of monomers. To summarize,

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

Note that, as expected, this is a linear combination of $H(\gamma)$ and $S(\gamma)$, and that increasing the length of domains l weighs $H(\gamma)$ more heavily, while decreasing the concentration C weighs $S(\gamma)$ more heavily. Typically $G_{\text{bp}}^{\circ} \approx -1.5$ kcal/mol, and $G_{\text{assoc}}^{\circ} \approx 1.96$ kcal/mol [7].

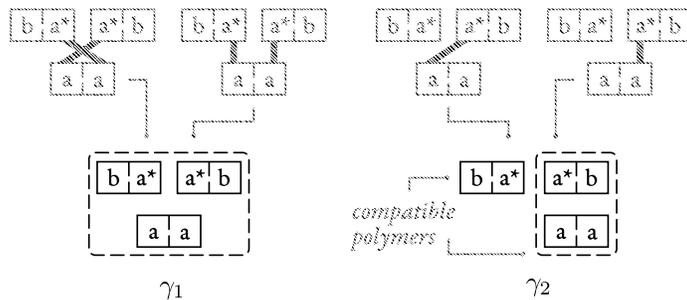


Fig. 1: Two configurations γ_1 and γ_2 of the TBN $\mathcal{T} = \{\{a, a\}, \{a^*, b\}, \{a^*, b\}\}$. Note that \mathcal{T} has 3 monomers but 2 monomer types and 6 sites but 3 site types. A dashed box indicates monomers that are part of the same polymer. A single configuration (bottom) can correspond to multiple ways of binding complementary sites (top), which are not distinguished in our model. In γ_2 the polymer on the left has exposed sites $\{b, a^*\}$ and the polymer on the right $\{a, b\}$; they are thus compatible since the exposed site a^* of the left is complementary to exposed site a of the right. Since γ_2 has compatible polymers it is not saturated, but γ_1 is.

the *height* of a path starting at γ as the maximum value of $E(\delta) - E(\gamma)$ over all configurations δ on the path. Then the kinetic energy *barrier* separating configuration δ from configuration γ is the height of the minimum-height path from γ to δ .

In Section 2 we introduce our main kinetic model. We further show that when $w \geq 1$, it is sufficient to consider only fully bonded configurations in the energy barrier analysis. In Sections 3.1 and 3.2 we develop two constructions for catalytic systems. Both constructions yield families of TBNs parametrized by a complexity parameter n such that the uncatalyzed energy barrier scales linearly with n . The catalyzed energy barrier is always 1. We show a direct DNA strand displacement implementation of one of the constructions. Finally we show an autocatalytic TBN, with an arbitrarily large energy barrier to undesired triggering, that exponentially amplifies its input signal (Section 3.2.1).

2 Kinetic model

Our kinetic models build on thermodynamic binding networks (TBN) [5]. Intuitively, we model a chemical system as a collection of molecules, each of which has a collection of binding sites, which can bind if they are complementary. Although the TBN model is more general, DNA domains can be thought of as the prototypical example of binding sites. No geometry is enforced, which allows the model to handle topologically complex structures, such as pseudoknots.

TBN (Figure 1 illustrates the concepts of the following two subsections.) Formally, a *TBN* is a multiset of monomer types. A *monomer type* is a multiset

of site types. A *site type* is a formal symbol, such as a , and has a *complementary* type, denoted a^* . We call an instance of a monomer type a *monomer* and an instance of a site type a *site*.

Configuration We may describe the configuration of a TBN at any moment in terms of which monomers are grouped into polymers. This way a *polymer* is a set of monomers, and a *configuration* is a partition of the monomers into polymers.⁵

The *exposed sites* of a polymer is the multiset of site types that would remain if one were to remove as many complementary pairs of sites as possible. Each such pair is counted as a *bond*. Note that bonds are not specified as part of a configuration, and intuitively we think of polymers as being maximally bonded. Two polymers are *compatible* if they have some complementary exposed sites. A configuration is *saturated* if no two polymers are compatible. This is equivalent to having the maximum possible number of bonds.

Notice that a polymer may have two incompatible halves. This represents spontaneous co-localization and comes with an energy penalty, as discussed later.

Path (Figure 2 illustrates the concepts of the following three subsections.) One configuration can change into another by a sequence of small steps. If γ can become δ by replacing two polymers with their union, then γ *merges* to δ and δ *splits* to γ , and we write $\gamma <^1 \delta$. We denote by \leq^1 , $<$, \leq the reflexive, transitive, and reflexive transitive closures of $<^1$. A *path* is a nonempty sequence of configurations where each merges or splits to the next. Note that there is a path between any two configurations.⁶

We could imagine smaller steps that manipulate individual bonds. But surprisingly, a bond-aware model leads to essentially equivalent kinetic barriers, which we prove in the full version of this paper. Thus keeping track of bonds is an unnecessary complication.

Energy For a configuration γ , denote by $H(\gamma)$ the number of bonds summed over all polymers. Denote by $S(\gamma)$ the number of polymers. Note that a saturated configuration has maximum $H(\gamma)$. The *energy* of γ is

$$E(\gamma) = -wH(\gamma) - S(\gamma),$$

where the bond strength w represents the benefit from gaining a bond relative to gaining a polymer. Note that $H(\gamma) \geq 0$ and $S(\gamma) > 0$, so $E(\gamma) < 0$, and that lower energy, which results from more bonds or more polymers, is more favorable. The choice to make favorability correspond to lower energy (more negative) is

⁵Note that swapping two monomers of the same type between different polymers produces a different configuration. Distinguishing different monomers of the same type allows us to equate the space of configurations with the lattice of partitions, which is a key tool in the full version of this paper.

⁶For instance, although this path is likely energetically unfavorable, we can merge all initial polymers into one, and then split into the desired end polymers.

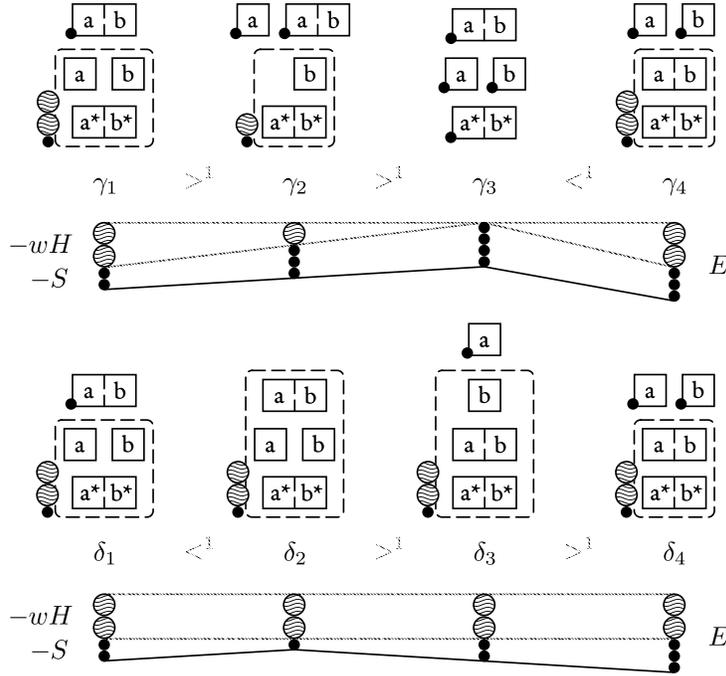


Fig. 2: A path p consisting of the configurations $\gamma_1, \gamma_2, \gamma_3, \gamma_4$ and a path p' consisting of the configurations $\delta_1, \delta_2, \delta_3, \delta_4$ of the TBN $\mathcal{T} = \{\{a\}, \{b\}, \{a, b\}, \{a^*, b^*\}\}$. The energy of each configuration is shown below it. A large wavy disc indicates energy due to a bond. A small solid disc indicates energy due to a polymer. Here bond strength $w = 2$, so a wavy disc is twice as tall as a solid disc. The height of p is $h(p) = E(\gamma_3) - E(\gamma_1) = (-4) - (-2w - 2) = 2$. The height of p' is $h(p') = E(\delta_2) - E(\delta_1) = 1$.

motivated by consistency with the standard physical chemistry notion of free energy. We call a minimum energy configuration *stable*.

Merging incompatible polymers forms no additional bonds and so is always unfavorable, since $S(\gamma)$ drops without $H(\gamma)$ rising. In contrast, when bond strength $w > 1$, merging compatible polymers is always favorable. So every stable (that is, minimum energy) configuration is saturated. This regime is typical of many real systems, and in particular, we can engineer DNA strand displacement systems [10] to have large bond strength w by increasing the length of domains (see also footnote 4).

Barrier With notions of paths and energy, we can establish the difficulty of passing from a configuration γ to another δ . The *height* $h(p)$ of a path p starting from γ is the greatest energy difference $E(\delta) - E(\gamma)$ from γ to any configuration δ along p . Notice that $h(p) \geq E(\gamma) - E(\gamma) = 0$. The *barrier* $b(\gamma, \delta)$ from γ to δ is the least height of any path from γ to δ . Notice that $b(\gamma, \delta) \geq 0$ as well.

2.1 Bounds on energy change

Merging compatible polymers or splitting into incompatible polymers changes energy in a predictable way. Splitting into incompatible polymers keeps all bonds and results in one more polymer, so overall it drops energy by 1. Similarly, merging compatible polymers results in one fewer polymer but at least one more bond. So overall it drops energy by $w - 1$ when bond strength $w \geq 1$.

To make this precise, we introduce two other partial orders on configurations. Let $\gamma \sqsubseteq \delta$ mean that γ can become δ by merges of *compatible polymers*. Let $\gamma \preceq \delta$ mean that γ can become δ by merges of *incompatible polymers*.

Claim 2.1. If $\gamma \preceq \delta$, then $E(\gamma) = E(\delta) - (S(\gamma) - S(\delta))$. If $\gamma \sqsubseteq \delta$, then $E(\delta) \leq E(\gamma) - (w - 1)(S(\gamma) - S(\delta))$.

The above claim says nothing about the general case $\gamma \leq \delta$. In order to apply the bounds, it will prove useful to decompose \leq into \preceq and \sqsubseteq . Any sequence of merges can be modified so that all merges between compatible polymers come first:

Claim 2.2. If $\gamma \leq \delta$, then some α has $\gamma \sqsubseteq \alpha \preceq \delta$.

Proof (sketch). Intuitively we want to reorder merges. In the full version of the paper, we show how to treat a merge as an object that can be applied in a context other than its original configuration, where the original polymers involved may not exist. With this machinery, the overall argument shows that if $\gamma \leq \delta$, then we can form α by starting with γ and doing as many merges as possible while preserving $\gamma \sqsubseteq \alpha \leq \delta$. No additional merge forms a bond, so $\alpha \preceq \delta$. \square

2.2 Saturated paths

We prefer to reason about saturated configurations because there are substantially fewer of them and they have special properties, which simplifies proofs. In this section we show that the barrier remains essentially the same even if we consider paths that traverse only saturated configurations. This may be surprising since breaking some bonds might seem to allow a path to bypass an otherwise large barrier.

We see an example of the special properties of saturated configurations in the following claim, which is used in later sections to show a large energy barrier in our constructions.

Claim 2.3. If γ and δ are saturated, then $b(\gamma, \delta) \geq S(\gamma) - S(\delta)$.

Proof. Consider a path p from γ to δ . Since γ and δ are saturated, $H(\gamma) = H(\delta)$, so $h(p) \geq E(\delta) - E(\gamma) = S(\gamma) - S(\delta)$. So $b(\gamma, \delta) \geq S(\gamma) - S(\delta)$. \square

Now we turn to the main result of this section. A *saturated path* is a path along which every configuration is saturated. For example, the bottom path p' in Figure 2 is saturated. If γ and δ are saturated, then let $b_{\text{sat}}(\gamma, \delta)$ denote the

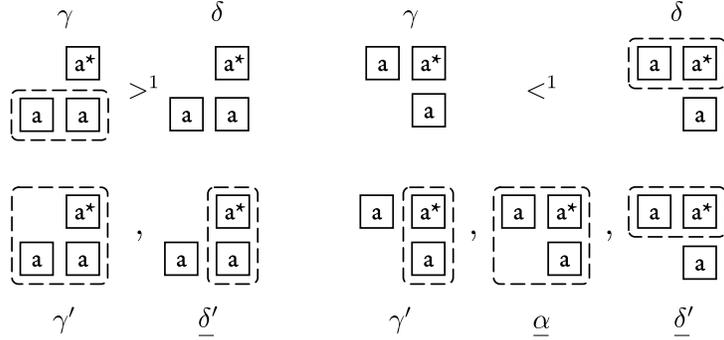


Fig. 3: An example of the two cases of the proof of Claim 2.5. An underline indicates a quantity constructed in the proof.

barrier from γ to δ when allowing only saturated paths. Since a saturated path is a path, $b_{\text{sat}}(\gamma, \delta) \geq b(\gamma, \delta)$. It turns out that if bond strength $w \geq 2$, then the reverse inequality also holds, so $b_{\text{sat}}(\gamma, \delta) = b(\gamma, \delta)$. And if $w \geq 1$, then the reverse inequality “almost” holds.

We first need a technical result proven in the full version of this paper:

Claim 2.4. If $\gamma <^1 \delta$ and $\gamma \leq \gamma'$, then some α has $\delta \leq \alpha$ and $\gamma' <^1 \alpha$.

As in the proof of Claim 2.2, the proof of the above claim relies on treating merges as objects that can be applied in different contexts. In particular, the same merge that changes γ to δ is applied to get from γ' to α .

To connect $b_{\text{sat}}(\gamma, \delta)$ to $b(\gamma, \delta)$, we first focus on a single step along a path and then extend to the full path. Let $[\gamma]$ denote the set of saturated γ' with $\gamma \preceq \gamma'$.

Claim 2.5. Let bond strength $w \geq 1$. If γ merges or splits to δ and $\gamma' \in [\gamma]$, then some $\delta' \in [\delta]$ and saturated path p' from γ' to δ' has

$$E_{\max}(p') \leq \max\{E(\gamma), E(\delta)\} + \max\{0, 2 - w\},$$

where $E_{\max}(p')$ is the maximum energy of any configuration along p' .

Proof. Let bond strength $w \geq 1$, and suppose $\gamma' \in [\gamma]$.

First consider the case where γ splits to δ (see Figure 3, left). Then $\delta <^1 \gamma$. So $\delta \leq \gamma$. By assumption $\gamma' \in [\gamma]$, so $\gamma \leq \gamma'$, so transitively $\delta \leq \gamma'$. By Claim 2.2, there is δ' with $\delta \preceq \delta' \preceq \gamma'$. By assumption γ' is saturated, and now $\delta' \preceq \gamma'$, so δ' is saturated, and so $\delta' \in [\delta]$. So let p' be the path from γ' to δ' by splits into incompatible polymers guaranteed to exist by $\delta' \preceq \gamma'$. Each such split drops energy, so the claim holds.

Next consider the case where γ merges to δ (see Figure 3, right). Then $\gamma <^1 \delta$. If $\gamma = \gamma'$, then let $\delta' = \delta$, and let p' be γ followed by δ . Otherwise $\gamma < \gamma'$. In that case let α be the configuration guaranteed by Claim 2.4. Then $\delta \leq \alpha$. So by Claim 2.2, there is δ' with $\delta \preceq \delta' \preceq \alpha$. Now by assumption γ' is saturated,

and by construction $\gamma' \leq \alpha$, so α is saturated. Next $\delta' \preceq \alpha$ means δ' is saturated, and so $\delta' \in [\delta]$.

Let p' be the concatenation of two paths p'_1 and p'_2 defined as follows. Let p'_1 be γ' followed by α . This merge can at worst result in one less polymer and no additional bonds, so $E(\alpha) \leq E(\gamma') + 1$.

Then let p'_2 be the path from α to δ' by splits into incompatible polymers guaranteed to exist by $\delta' \preceq \alpha$. Each such split drops energy, so α is the highest energy intermediate configuration. But by assumption $\gamma \leq \gamma'$ and $\gamma \neq \gamma'$, so $S(\gamma) - S(\gamma') \geq 1$. So by Claim 2.1, and assuming bond strength $w \geq 1$, we get $E(\gamma') \leq E(\gamma) - (w - 1)$. So

$$\begin{aligned} E(\alpha) &\leq E(\gamma') + 1 \\ &\leq E(\gamma) + 2 - w, \end{aligned}$$

which implies the claim. \square

To extend the result to a full path, we apply it to each configuration along the path.

Claim 2.6. For bond strength $w \geq 1$ and saturated γ and δ , we have

$$b_{\text{sat}}(\gamma, \delta) \leq b(\gamma, \delta) + \max\{0, 2 - w\}.$$

Proof. Suppose bond strength $w \geq 1$ and γ and δ are saturated, and consider a path p from γ to δ . By assumption γ is saturated, so $\gamma \in [\gamma]$. So we can apply Claim 2.5 to each configuration of p in turn to get a saturated path p' .

Let E' denote the maximum energy along p' and E denote the maximum energy along p . Then Claim 2.5 ensures $E' \leq E + \max\{0, 2 - w\}$. And both p' and p start in the same configuration with the same energy, so $h(p') \leq h(p) + \max\{0, 2 - w\}$. So $b_{\text{sat}}(\gamma, \delta) \leq b(\gamma, \delta) + \max\{0, 2 - w\}$. \square

Since $b(\gamma, \delta) \leq b_{\text{sat}}(\gamma, \delta)$, we have the following corollary of Claim 2.6.

Corollary 2.7. For bond strength $w \geq 2$ and saturated γ and δ , we have $b_{\text{sat}}(\gamma, \delta) = b(\gamma, \delta)$.

For Claim 2.6 and Corollary 2.7, bond strength $w \geq 1$ is necessary. If $w < 1$, then $b_{\text{sat}}(\gamma, \delta)$ can be larger than $b(\gamma, \delta)$ by an arbitrary amount.

3 TBNs with programmable energy barriers

We present two constructions for TBNs with equal energy (stable) “initial” and “triggered” configurations, such that the energy barrier to get from one to the other can be made arbitrarily large: both constructions are parameterized by n , with the energy barrier scaling linearly with n . Further, both constructions admit catalysts that reduce the energy barrier to 1.

The first construction (translator cycle), discussed Section 3.1, is based on a DNA strand displacement catalyst, and the progress from the initial to triggered

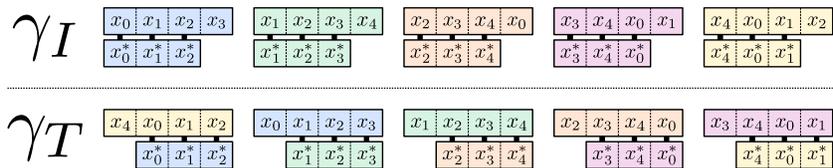


Fig. 4: The two stable configurations of a translator cycle with complex length $z = 3$ and number of complex types $c = 5$.

configurations with the catalyst can be physically implemented as a strand displacement cascade. Although this system has been previously proposed in [11], for the first time we rigorously prove an energy barrier.

The second construction (grid gate), discussed in Section 3.2, does not have an evident physical implementation (e.g., as a strand displacement system), but surpasses the translator cycle system in the following ways: (1) a proof of *copy tolerance*⁷: the energy barrier is proven in more general contexts where multiple copies of monomers are present, in any ratio, (2) *autocatalysis*: the grid gate can be modified so that the catalyst transforms the gate into a polymer that has the same excess domains as the catalyst, which can itself catalyze the transformation of additional gates (leading to exponential amplification).

Throughout both sections, we make the assumption that $w \geq 2$, so that by Corollary 2.7, it is sufficient to describe energy barriers by pathways in the saturated model; if we weaken this assumption to $w \geq 1$, then by Claim 2.6 the barrier proved is within 1 of the barrier in the unrestricted pathway model (allowing unsaturated configurations).

The constructions demonstrate that catalysts and autocatalysts with arbitrarily high energy barriers can be engineered solely by reference to the general thermodynamic driving forces of binding and formation of separate complexes, which are captured in the TBN model.

3.1 Translator cycle

Consider the TBN illustrated in Figure 4. There are two particular configurations that interest us, an initial configuration γ_I and a triggered configuration γ_T . The two configurations are stable. In the presence of a catalyst monomer $\{x_4, x_0, x_1, x_2\}$ (or an extra copy of any *top monomer*—any of the monomers with unstarred domains), a height 1 pathway exists to reach γ_T , illustrated by Figure 5. Further, this catalytic system is realizable as a DNA strand displacement cascade; more information about this connection can be found in the full version of this paper, and in [11]. In the case of many copies of each complex, since the catalyst is in fact any of the top monomers, the system may be used as an amplifier: at the end of the pathway shown in Figure 5, another monomer with domains

⁷Our proof technique for the barrier of the translator cycle does not seem sufficient to prove copy tolerance, although we conjecture the translator cycle is copy tolerant.

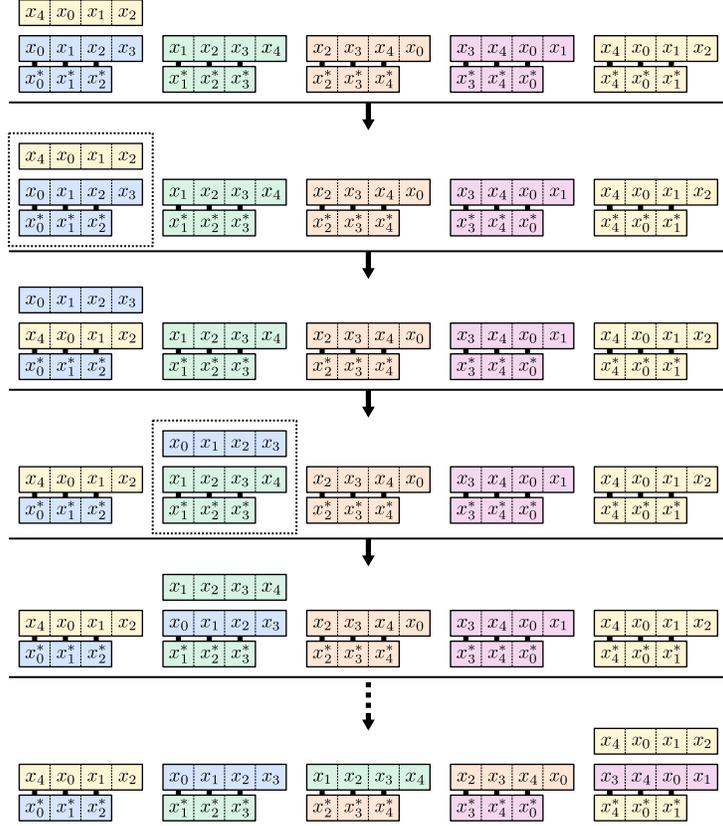


Fig. 5: A segment of the height 1 path which is possible because an extra copy of a top monomer, $\{x_4, x_0, x_1, x_2\}$, is present to act as a catalyst. If no catalyst is initially present, if one complex splits—thus reducing the number of bonds by z —the top monomer from the split complex can be used as a catalyst, resulting in a path of height $z + 1$ similar to the path shown in the figure.

$\{x_4, x_0, x_1, x_2\}$ becomes free which can catalyze another set of complexes which are in the initial configuration. Note, however, that proving the energy barrier for the translator cycle in the multi-copy setting remains open.

To program a large energy barrier, we give a formal definition for generalizing the translator cascade, parameterized by *complex length* z and *number of complex types* c . Given $z \leq c$, a (z, c) -*translator cycle* is a TBN with monomers $\mathcal{M} = \{m_i, m_i^* \mid i \in \{0, \dots, c-1\}$, where $m_i = \{x_i, x_{i+1 \pmod{c}}, \dots, x_{i+z \pmod{c}}\}$ and $m_i^* = \{x_i^*, x_{i+1 \pmod{c}}^*, \dots, x_{i+z-1 \pmod{c}}^*\}$. The *initial configuration* is $\gamma_I = \{\{m_i, m_i^*\} \mid i \in \{0, \dots, c-1\}\}$ and the *triggered configuration* is $\gamma_T = \{\{m_{i-1 \pmod{c}}, m_i^*\} \mid i \in \{0, \dots, c-1\}\}$.

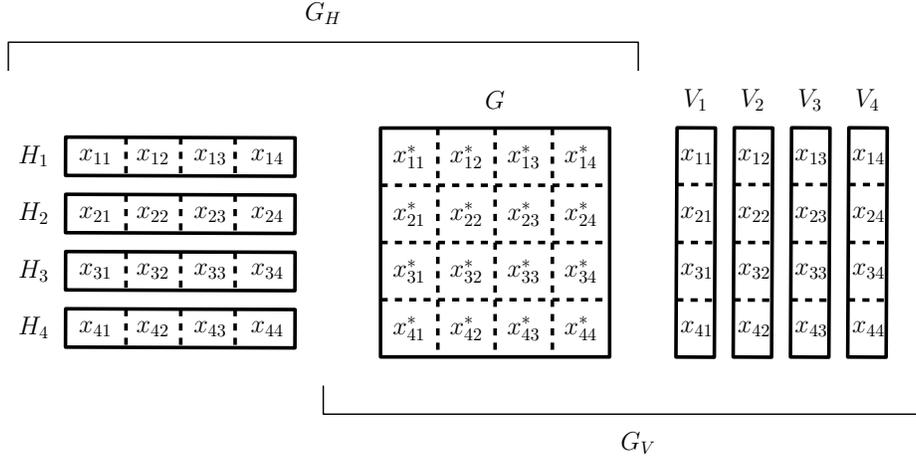


Fig. 6: The monomer types in the grid gate TBN for the case $n = 4$.

What we want is that when the catalyst is absent, there is a large energy barrier to move from γ_I to γ_T . This barrier depends on the complex length z and the number of complex types c , and can be made arbitrarily large.

We can begin a path to γ_T by splitting any complex apart, thus reducing the number of bonds by z , and then use the top monomer as a catalyst in the same way as the with-catalyst pathway shown in Figure 5. So, if z is not large, then there is a small barrier. We can also bring all complexes together while reducing the number of polymers by $c - 1$, and then split into the triggered complexes, so if c is not large, there is also a small barrier. Somewhat surprisingly, it is not sufficient to set $z = c = n$ to attain a barrier of $\Omega(n)$; a complicated path exists which has height $\Theta(\frac{c}{z})$ which is illustrated in the full version of the paper.

Note that if the cascade has $c = z^2$, then the uncatalyzed paths described above have height $\Omega(z)$. Are there other paths with smaller heights (that is, is the energy barrier $\Omega(z)$ in that case)? We prove that indeed the energy barrier is $\Omega(z)$:

Theorem 3.1. *If $z = n$ and $c = n^2$, then $b(\gamma_I, \gamma_T) = \Omega(n)$.*

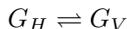
The proof appears in the full version of this paper.

3.2 Grid gate

Consider the TBN illustrated in Figure 6. We focus on two polymer types G_H and G_V depicted in the figure, and show that there is a barrier to convert G_H to G_V and vice versa. To generalize to arbitrarily high energy barriers, we parameterize the construction by n as follows. Define the following monomer types: “horizontal” $H_i := \{x_{ij}\}_{j=1}^n$ for $i \in \{1, \dots, n\}$, “vertical” $V_j := \{x_{ij}\}_{i=1}^n$ for $j \in \{1, \dots, n\}$, and “gate” $G := \{x_{ij}^*\}_{i,j=1}^n$. We fix a network \mathcal{T} which contains any number of

any of these monomer types, so long as there are enough of other monomers to completely bind all the G monomers (i.e., in saturated configurations there are no exposed starred sites).

The main result of this section is that there is an energy barrier of n to convert a G_H polymer to a G_V polymer and vice versa. In the notation of chemical reaction networks, this binding network implements the net reaction



with energy barrier n . In Section 3.2.1 we will show how this energy barrier can be reduced to 1 by a catalyst.

We state our results in increasing order of strength. First, we show an energy barrier to changing a single G_H to a single G_V (and vice versa), in the presence of exactly one of each V_j monomer (i.e., just enough to create a G_V). Let γ_H denote the configuration of the TBN depicted in Figure 6 in which G and each H_i are grouped into a single polymer G_H , and each V_j is in its own polymer. Let γ_V denote the symmetric configuration in which G and each V_j are grouped into a single polymer G_V , and each H_i is in its own polymer.

Claim 3.2. $b(\gamma_H, \gamma_V) = b(\gamma_V, \gamma_H) = n$.

Proof. Consider a saturated path p from γ_H to γ_V . Let δ denote the first configuration in the path with some H_i separate from G . Since δ is saturated but in this configuration H_i is not bound to G , all of the V_j must be bound together with G . So the configuration δ' which immediately precedes δ in the path must have all of the monomers in a single polymer. Since p is a saturated path, $H(\gamma_H) = H(\delta')$, and since p was arbitrary, $b(\gamma_H, \gamma_V) \geq h(p) \geq E(\delta') - E(\gamma_H) = S(\gamma_H) - S(\delta') = n$. The path which achieves height n proceeds by merging G_H and all the V_j 's, and then splitting off all the H_i 's, leaving G_V behind. By symmetry, $b(\gamma_H, \gamma_V) = b(\gamma_V, \gamma_H)$. \square

We can generalize this basic result to a *copy-tolerant* setting, in which we assume the starting configuration has several copies of each G_H and G_V , as well as excess copies of each H_i and V_j , showing that the energy barrier of n remains for conversion of any G_H 's into G_V 's, and vice versa.

To make our result precise, we parameterize a set of configurations of the network based upon the count of polymers G_H and G_V and the number of horizontal and vertical monomers that are unbound. For $c_{G_H}, c_{G_V} \in \mathbb{N}$, define $\Gamma(c_{G_H}, c_{G_V})$ to be the set of configurations of the network \mathcal{T} that contain precisely c_{G_H} polymers of type G_H and c_{G_V} polymers of type G_V , and in which all other monomers are in separate polymers by themselves. It turns out that the stable configurations of the network are exactly the ones in some $\Gamma(c_{G_H}, c_{G_V})$ (proof in the full version of this paper). The following theorem generalizes Claim 3.2 to the multi-copy context. Note that obtaining a similar copy-tolerant result for the translator cycle remains open.

Theorem 3.3. *Let $\gamma \in \Gamma(c_{G_H}, c_{G_V})$. If $\gamma' \in \Gamma(c_{G_H} - \Delta, c_{G_V} + \Delta)$ for some nonzero $\Delta \in \mathbb{N}$, then $b(\gamma, \gamma') = n$.*

The proof appears in the full version of this paper.

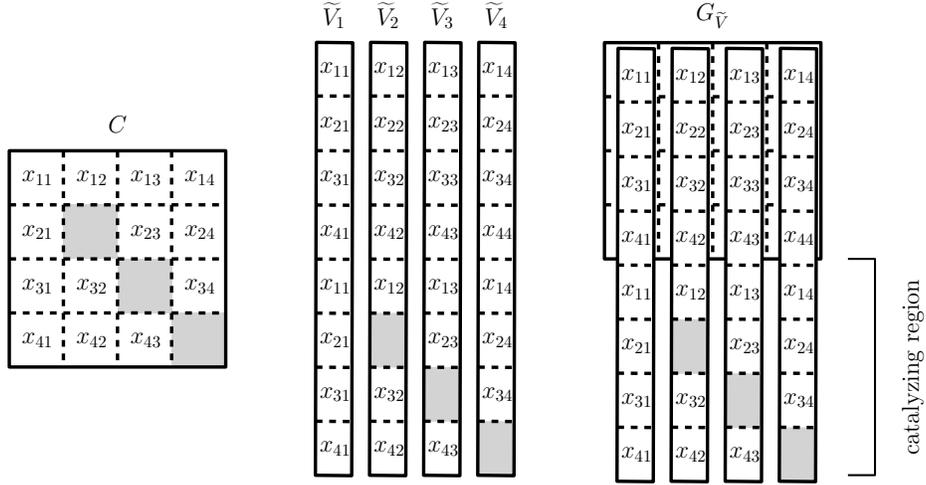
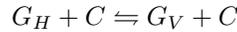


Fig. 7: Catalysts and autocatalysts in the grid gate TBN for the case $n = 4$. left: C is a single monomer that acts as a catalyst to convert between G_H and G_V . middle: Modified vertical monomers $\{\tilde{V}_j\}_{j=1}^n$ with extra sites. right: After C converts G_H to G_V with modified vertical monomers, G_V has the same excess sites as C and acts as a catalyst itself.

3.2.1 Catalysis and autocatalysis The kinetic barrier shown for the grid gate can be disrupted by the presence of new monomer types. In fact, the model admits a *catalyst* monomer C that lowers the energy barrier from n to 1, i.e., in the presence of C , a G_H can be converted into a G_V , and vice versa, with a sequence of merge-split pairs. In the notation of chemical reaction networks, this binding network implements the net reaction



with energy barrier 1 but maintains a high energy barrier for the reaction $G_H \rightleftharpoons G_V$.

For the grid gate of size $n \times n$, we define a catalyst: $C := \{x_{ij} \mid 1 \leq i, j \leq n, i \neq j\} \cup \{x_{11}\}$, illustrated in Figure 7 (left). C is a monomer containing all of the non-diagonal unstarred sites, while retaining exactly one of the diagonal sites. The mechanism by which C can transform G_H to G_V with merge-split pairs is by an alternating processes of merges and splits shown in Figure 8. Intuitively, in step i of the catalyzed reaction $G_H + C \rightarrow G_V + C$, site x_{ii}^* on G switches its association from x_{ii} on H_i (left) to its counterpart on V_i (right) by merging the evolving polymer (center) with V_i and then splitting off H_i . Only the diagonal sites need be considered in the intermediate steps since the catalyst C (blue, center, overlaid) balances the other sites of G .

Consider a network $\mathcal{T}^C := \{G, C\} \cup \{H_i\}_{i=1}^n \cup \{V_j\}_{j=1}^n$ which includes one instance of every monomer type that has been introduced, as well as the catalyst. As before, we shall be interested in net transitions between G_H and G_V , and so for

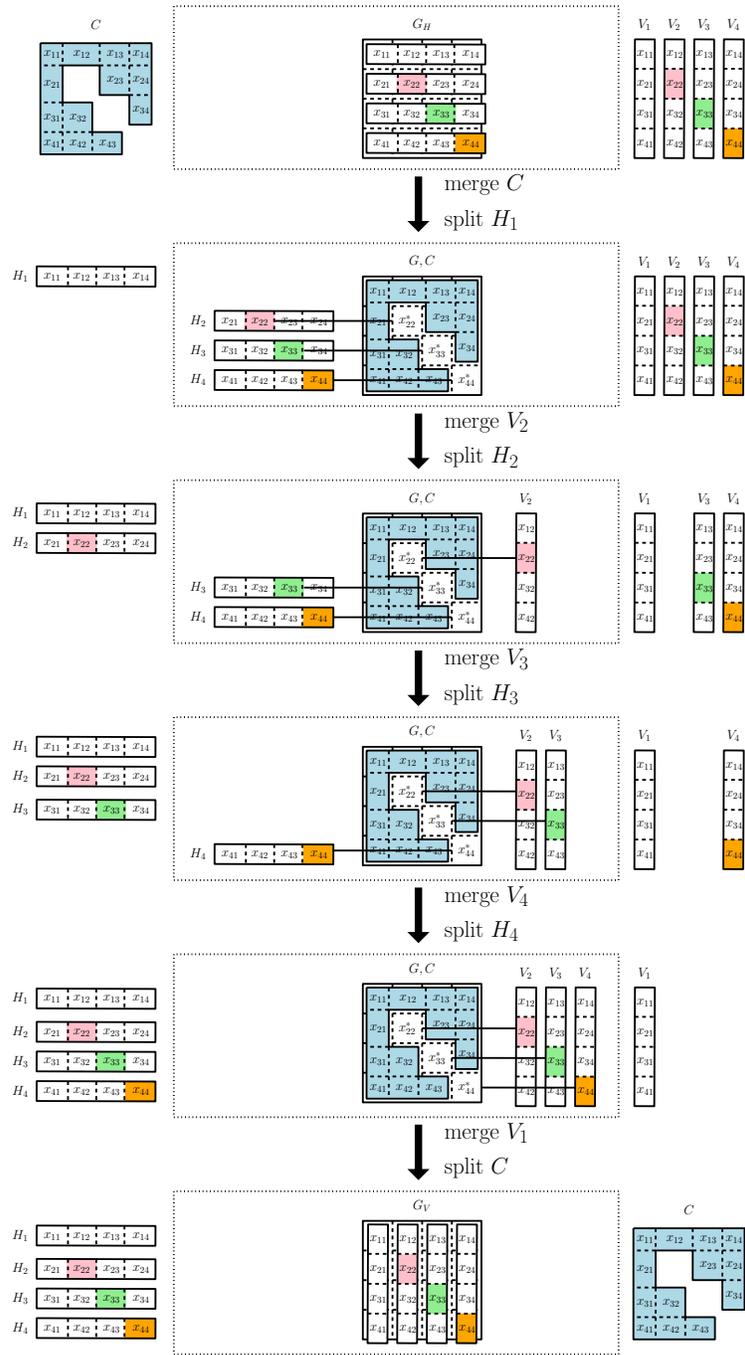


Fig. 8: Full pathway for reaction $G_H + C \rightarrow G_V + C$. In each stage, exactly one merge and one split occurs, and the center polymer remains saturated.

the network \mathcal{T}^C we define the following configurations: $\gamma_H^C := \{G_H, C\} \cup \{V_j\}_{j=1}^n$ and $\gamma_V^C := \{G_V, C\} \cup \{H_i\}_{i=1}^n$.

Theorem 3.4 states that transitions that previously had arbitrarily large energy barriers are reduced in barrier to one.

Theorem 3.4. $b(\gamma_H^C, \gamma_V^C) = b(\gamma_V^C, \gamma_H^C) = 1$.

By adding to the network additional instances of $\{G\} \cup \{H_i\}_{i=1}^n \cup \{V_j\}_{j=1}^n$, these results can be extended to the more general case of $\gamma(c_{G_H}, c_{G_V}) \cup \{C\}$ by induction.

The grid gate can also be modified to act in an *autocatalytic* manner. By modifying the vertical monomers it is possible for G_V to have a set of exposed monomers acting as a ‘‘catalyzing region’’, which has the same structure and function as C (see Figure 7, middle and right). We give the formal construction and related proofs in the full version of this paper.

4 Future work

Single-molecule applications are rare compared with multi-copy situations. Indeed, using a catalyst or autocatalyst for amplification requires many copies of the components. Although we proved the correctness of the grid gate in the multi-copy context, we could not easily generalize our proof of the correctness of the translator cycle. Since the translator cycle has a straightforward strand displacement implementation, it is important to show that the uncatalyzed energy barrier is still n in bulk applications typical of strand displacement experiments. Success in this line of research provides a promising approach to making leakless strand displacement amplifiers [10]. More generally, it is important to develop general methods for extending single-copy TBN results to multi-copy contexts.

Ideally we not only want a large energy barrier to ‘‘bad’’ configurations, but we want to avoid getting stuck in local minima that keep us from getting to the good configurations. We can define ‘‘self-stabilizing’’ TBNs with the property that from *any* configuration, the TBN can reach *some* stable configuration with a low energy barrier path. As argued in the full version of this paper, this property is true for the grid gate, but it is not true for the translator cycle. This property is worth more general exploration.

Can we use the definition of energy in TBNs to bootstrap a reasonable notion of probability of configurations or paths? For instance, in statistical thermodynamics it is common to consider the Boltzmann distribution induced by energy E , where for each configuration γ , $\Pr[\gamma] = e^{-E(\gamma)} / (\sum_{\beta} e^{-E(\beta)})$. This is the probability of seeing γ at thermodynamic equilibrium. One can also use the relative energy between two states to predict the relative rates of transition between them, which might allow defining a notion of path probability in the kinetic theory of TBNs.

A useful chemical module consists of the reaction $X + X \rightarrow Y + Y$ (or more generally converting $k > 1$ copies of X to Y), which can act as a ‘‘threshold’’ to detect whether there are at least k copies of X . Analogous to a catalytic system, implementing the above reaction while forbidding $X \rightarrow Y$ requires control of the

energy barrier, and cannot be done simply by varying the energies of X and Y . Can we construct TBNs with arbitrarily large energy barriers in this case?

What is the computational complexity of deciding whether $b(\gamma, \delta) \leq k$, given two configurations γ, δ and a threshold k ? This problem is decidable in polynomial space in the number of monomers in γ : any configuration can be written down in polynomial space, and guessing merges and splits yields a nondeterministic polynomial space algorithm (placing it in PSPACE by Savitch's Theorem). However, low-height paths could be of exponential length, and thus it is not clear that the problem is in NP, since the obvious witness is a path with height $\leq k$. Is this problem possibly PSPACE-complete?

Can kinetic barriers be proven in geometric self-assembly models such as the abstract Tile Assembly Model (aTAM) [4, 6, 12]? One approach to providing thermodynamic arguments for correctness of self-assembly systems is by showing the stable configurations of the system as a TBN are exactly the desired terminal assemblies in the aTAM. It has been shown that the class of systems which have this notion of stability is limited [2]. An alternative approach would be to relax the requirement that the set of desired assemblies correspond exactly to stable configurations, and instead argue that there is a large kinetic barrier to reach undesired assemblies even if they are stable.

Acknowledgement. DD and DH were supported by NSF grant CCF-1619343. CC, KB, and DS were supported by NSF grants CCF-1618895 and CCF-1652824.

References

1. Luca Cardelli. Strand algebras for DNA computing. *Natural Computing*, 10(1):407–428, 2011.
2. Cameron Chalk, Jacob Hendricks, Matthew J. Patitz, and Michael Sharp. Thermodynamically favorable computation via tile self-assembly. In *Proceedings of the 17th International Conference on Unconventional Computation and Natural Computation*, 2018.
3. Yuan-Jyue Chen, Neil Dalchau, Niranjan Srinivas, Andrew Phillips, Luca Cardelli, David Soloveichik, and Georg Seelig. Programmable chemical controllers made from DNA. *Nature Nanotechnology*, 8(10):755–762, 2013.
4. David Doty. Theory of algorithmic self-assembly. *Communications of the ACM*, 55(12):78–88, December 2012.
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: 23rd International Conference*, pages 249–266. Springer, 2017.
6. Matthew J. Patitz. An introduction to tile-based self-assembly and a survey of recent results. *Natural Computing*, 13(2):195–224, 2014.
7. John SantaLucia Jr and Donald Hicks. The thermodynamics of DNA structural motifs. *Annual Review of Biophysics and Biomolecular Structure*, 33:415–440, 2004.
8. David Soloveichik, Georg Seelig, and Erik Winfree. DNA as a universal substrate for chemical kinetics. *Proceedings of the National Academy of Sciences*, 107(12):5393–5398, 2010.

9. Niranjan Srinivas, James Parkin, Georg Seelig, Erik Winfree, and David Soloveichik. Enzyme-free nucleic acid dynamical systems. *Science*, 358(6369):eaal2052, 2017.
10. Chris Thachuk, Erik Winfree, and David Soloveichik. Leakless DNA strand displacement systems. In *International Workshop on DNA-Based Computers*, pages 133–153. Springer, 2015.
11. Boya Wang, Chris Thachuk, Andrew D. Ellington, Erik Winfree, and David Soloveichik. Effective design principles for leakless strand displacement systems. Submitted for publication.
12. Erik Winfree. Simulations of computing by self-assembly. Technical Report CaltechCSTR:1998.22, California Institute of Technology, 1998.