Reachability Bounds for Chemical Reaction Networks and Strand Displacement Systems

Anne Condon[1], Bonnie Kirkpatrick[1], and Ján Maňuch[1,2]

[1] Department of Computer Science, University of British Columbia, Vancouver, British Columbia
[2] Department of Mathematics, Simon Fraser University, Burnaby, British Columbia

Abstract. Chemical reaction networks (CRNs) and DNA strand displacement systems (DSDs) are widely-studied and useful models of molecular programming. However, in order for some DSDs in the literature to behave in an expected manner, the initial number of copies of some reagents is required to be fixed. In this paper we show that, when multiple copies of all initial molecules are present, general types of CRNs and DSDs fail to work correctly if the length of the shortest sequence of reactions needed to produce any given molecule exceeds a threshold that grows polynomially with attributes of the system.

1 Introduction

DNA strand displacement systems (DSDs) [10,13] and chemical reaction networks (CRNs) [8,7] are important molecular programming models. DSDs provide sophisticated molecular realizations of logic circuits and even artificial neurons [4,5], while CRNs elegantly express chemical programs that can then be translated into DSDs [8,9]. CRNs and thus DSDs can in principle simulate Turing-general models of computation [6,3], and DSDs can be energy efficient [11,6,9,14]. It is also possible in principle to recycle molecules in DSDs by running reversible reactions or displacements in both forwards and reverse directions, so that t steps of the system use just $O(\log t)$ molecules [1].

However, correct behavior of some published DSDs [3,1] requires that an exact numbers of some reactants are present initially, and it is currently impractical to obtain the exact numbers in a wet lab. We previously considered the conditions for a class of CRNs to work correctly when multiple copies of all initial molecules are present and showed that the length of the shortest trace (sequence of reactions) needed to "reach", i.e., produce, any given molecule is bounded by a polynomial function of some attributes of a CRN in this class [1]. This reachability upper bound reveals important limits of molecular programs that fall in the class covered by our result: we cannot write such programs that run correctly in a closed chemical system and for which the number of steps (reactions) of the program is sufficiently large relative to the volume of initial reagents.[1]

[1] The volume is the physical volume of all the molecules. It can be approximated by the number of all the types of reagents in the initial configurations.

D. Stefanovic and A. Turberfield (Eds.): DNA 18, LNCS 7433, pp. 43–57, 2012.

In this work we provide two new reachability upper bounds that significantly extend our earlier work. The first new theorem applies to *tagged* CRNs which, as we explain below, are important because they can be translated into DSDs of comparable volume that can simulate the CRNs traces. The second new theorem applies to a broader class of DSDs than does the translated version of our first result. In the rest of this introduction we motivate our results in more detail. Sections 2 and 3 provide technical details of both theorems; additional details are in the full version of the paper. We list some open questions in Section 4.

New result for chemical reaction networks (CRNs). Figure 1 illustrates a CRN of the type to which our new result applies (a formal definition of CRN is in Section 2). Each reaction r is reversible, and has unique *tag* species τ_r^+ and τ_r^- on its left and right sides respectively. We explain later why we focus on tagged CRNs, and also explain why we ignore reaction rate constants in our example and results. When a single copy of each species in the set $\{A, B, E, \tau_1^+, \tau_2^+, \tau_3^+, \tau_4^+\}$ is initially present, it takes six reaction steps to produce the product H, and to do so, reaction 1 must run in the forwards direction, then later run backwards, then forwards again. However, if another copy of A and B are present initially then H can be generated with just four reactions. The behavior of the system with two copies does not mirror its behavior with one copy; in this sense it is incorrect. While for this simple example it might not seem important how many steps are needed to produce a particular product, it is critically important in contexts where the product is the result of a computation and an erroneous result could be produced as a result of cross-talk, or short-circuiting of multiple copies of the intended computation.

$$
\begin{aligned}
&(1)\ \tau_1^+ + A + B \rightleftharpoons C + D + \tau_1^- \\
&(2)\ \tau_2^+ + C + D + E \rightleftharpoons C + D + F + \tau_2^- \\
&(3)\ \tau_3^+ + A + B + F \rightleftharpoons A + B + G + \tau_3^- \\
&(4)\ \tau_4^+ + C + D + G \rightleftharpoons C + D + H + \tau_4^-
\end{aligned}
$$

Fig. 1. Example of a simple tagged chemical reaction network (CRN)

In this paper, our notion of correctness is that of *copy tolerance* [1]. We say that a CRN $\mathbf{C}$ is x-copy-tolerant if the length of the shortest trace that produces any species s in $\mathbf{C}$ and in $\mathbf{C}^{(x)}$ is the same, where $\mathbf{C}^{(x)}$ is the CRN with the same reactions as $\mathbf{C}$ but with x initial copies of each initial molecule of $\mathbf{C}$. A system is copy-tolerant if it is x-copy-tolerant for all x. The CRN of Figure 1 is not 2-copy-tolerant. Copy-tolerance is a weak notion of correctness; if a CRN $\mathbf{C}$ is not 2-copy tolerant then, for example, $\mathbf{C}$ also fails to satisfy the stronger requirement that each possible trace of $\mathbf{C}$ in the 2-copy setting is an interleaving of two possible traces in the single copy setting. We chose to work with a weak notion of correctness because it makes our results stronger, i.e., they apply also to notions of correctness that are stronger than copy-tolerance.

Our first reachability upper bound, Theorem 2, shows that in order for a tagged CRN **C** to be copy-tolerant, the number of steps needed for **C** to produce any given species must be suitably bounded. The bound is a polynomial function of the volume and other attributes of **C**.

We prove our result for *tagged* CRNs—CRNs with a unique species on the left and right side of each reaction (see Figure 1)— for two reasons. First, the tags make it possible for us to prove strong results. The second reason stems from the fact that our ultimate goal is to prove limits on the power of DSDs, which can be realized with DNA strands, rather than for CRNs which are a useful theoretical abstraction. When translating an untagged CRN to a DSD, two sets of auxiliary DNA strand complexes, so-called transformer, are introduced per reaction of the CRN, one set for each side of the reaction. Each set of transformers includes unique strands that do not otherwise appear in the DSD. The CRN tag species represent the sets of transformer DNA strands. Put another way, to translate an untagged CRN to a DSD using current methods, it is necessary to first add tags to the CRN and then map the tags to the sets of transformer species. Thus, by proving a reachability upper bound for a tagged CRN, we are obtaining a result for the DSD realization of the corresponding untagged CRN. The result would apply also to other realizations of CRNs, perhaps even using molecules other than DNA, in which transformer molecules are needed in the realization. Our earlier result [1] did not apply to general tagged CRNs.

Unlike the example of Figure 1, chemical reactions have associated kinetic rate constants that, along with species counts, determine reaction propensities [8,7]. In particular, a CRN behaves stochastically if multiple reactions are applicable to the molecules available at one or more points in the sequence of reactions. However, in examples such as the stack machine of Qian et al. [3] and the Gray code counter of Condon et al. [1], correctness of the CRN does not depend on the relative propensities of applicable reactions (although efficiency of the CRN does). Since our results are expressed in terms of number of reactions rather than reaction propensities, they apply to stochastic CRNs. We can interpret our reachability result as a hitting time in the stochastic context where a hitting time is the minimum number of reactions required to reach a goal state from a initial state.

New result for strand displacement systems (DSDs). The second main contribution of this paper is a limit on the types of DSDs that are correct in multi-copy settings. In strand displacement (Figure 2), an initially unbound "signal" strand I binds to a "template" T, causing another signal strand O that was initially bound to T to become unbound. DSDs are collections of strands that can change configurations via successive strand displacements in a pre-programmed fashion [13,14,2]; we provide a formal definition later.

Our first result on tagged CRNs implies a reachability upper bound for DSD realizations of CRNs, but says nothing about DSDs more generally. In Theorem 3 we elucidate this simple upper bound which is obtained by applying the CRN result to limited types of DSDs, those whose signal strands consist of exactly two domains: a toehold and a long-domain. However, since the signal types are

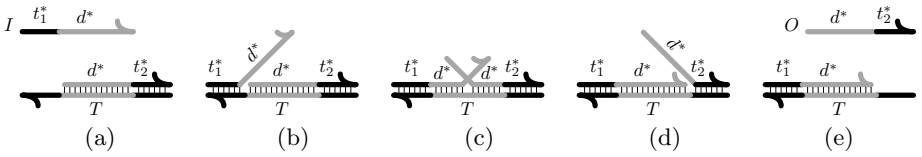

Fig. 2. Strand displacement. (a) An unbound DNA strand I, with a short toehold (dark line) and long-domain (lighter line), plus a duplex consisting of a template strand T and a third strand O that is bound to T. (b) I binds to T via its toehold. (c) Through a process of branch migration, the long-domain of I becomes bound to T, displacing bonds of O. (d) O is bound to T by only a toehold. (e) The toehold bonds break, making O unbound.

limited, this result does not apply to general DSDs. This is because, while tagged CRNs can be translated to DSDs having parameters such as the volume and the number of types of reactants polynomial in the volume of the CRN [9], it is not clear whether the converse is true. To see why, consider signal strands that have three domains: a toehold and two long-domains such that they each start with the same long-domain d^* and toehold t^*, and end with a distinct long-domain. Assume there are δ different types of these signal strands where δ is the number of long-domains on the template we will consider. Note for the DSD template having δ long-domains, over the course of several displacements, there are factorially many different configurations—ways in which signal strands are bound to the template. Figure 3 provides a simple example where any permutation of the signal species could bind to the template. Now, we want to create a tagged CRN that is equivalent to this DSD. Such a tagged CRN in which each template configuration is a distinct species would thus have the number of distinct species and reactions factorial in the volume (number of toehold and long-domains) of the DSD. Since each reaction in the tagged CRN requires a unique tag which needs to be present in the initial configuration, the overall volume of the tagged CRN would be also factorial in the volume of the DSD. It is not clear how else to translate such a DSD to a (tagged) CRN of comparable volume.

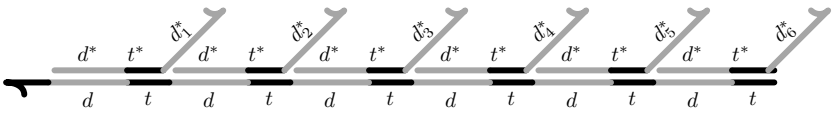

Fig. 3. A template with 6 long-domains and $6! = 720$ possible configurations. Dark lines are toeholds and the lighter ones are long-domains. The template contains $\delta = 6$ toehold-long-domain-toehold blocks. In each block, any one of the signal species may be bound. Thus the number of possible configurations of this template is $\delta! = 6!$.

Can "long" computations be correctly performed by DSDs, even in the presence of many copies? Our second reachability upper bounds for DSDs, Theorems 4 and 5, answer this in the negative, showing that, if sufficiently many

copies are present, then any unbound DNA strand that can be produced (i.e., reached) by a sequence of strand displacements can always be reached within a number of displacements that grows at most polynomially in the volume of the single-copy DSD. Thus, for example, we cannot write DSD programs that run correctly in the multi-copy setting and for which the minimum number of displacements needed to produce some given signal strand is exponential in the initial volume.

As further motivation, we describe another application of our DSD reachability bound. The CRN of Figure 4 describes a traditional 3-bit binary counter. Initially, three species, namely $0_3, 0_2$ and 0_1 represent the bits 0 at each index of the counter. Exactly one reaction can advance the counter from each value (all in the forward direction), until the counter reaches $1_3 1_2 1_1$. For the n-bit generalization of this counter, the number of species is just $2n$ (two species per bit) while the number of steps is 2^n. Thus the volume is logarithmic in the number of steps. Another very nice feature of this CRN is that it works correctly even if multiple copies of the initial species are present, not only in the sense of being copy-tolerant but also in the sense that the trace of the multi-copy system is an interleaving of traces of the single-copy system, even in the presence of cross-talk (details omitted).

$$
\begin{array}{ll}
(1)\ \mathbf{0}_1 & \rightleftharpoons \mathbf{1}_1 \\
(2)\ \mathbf{0}_2 + 1_1 & \rightleftharpoons \mathbf{1}_2 + 0_1 \\
(3)\ \mathbf{0}_3 + 1_2 + 1_1 & \rightleftharpoons \mathbf{1}_3 + 0_2 + 0_1
\end{array}
$$

Fig. 4. Binary counter CRN

However, if tags are added to the counter in order that it can be translated to a DSD using tags as discussed previously, the volume of species for the DSD realization of the counter becomes exponential in n. This is because reaction (1) is executed in the forward direction 2^{n-1} times and is never executed in the reverse direction; thus 2^{n-1} copies of the tag on the left side of reaction (1) must be present initially. Is there an alternative (tag-less) DSD realization of the n-bit CRN binary counter whose volume grows polynomially in n? Our DSD result implies that there is no such realization. If there were, then our reachability upper bound implies that in the multi-copy setting the bit 1_n could be produced in a polynomial number of steps. But since we know that it takes 2^{n-1} steps to produce 1_n even in the multi-copy setting, we have a contradiction.

2 Reachability Upper Bound for CRNs

In this section we first provide formal definitions of tagged CRNs. We then provide our main technical result, and conclude with a restatement of this result to obtain our reachability upper bound theorem for copy-tolerant CRNs.

2.1 Definition of Tagged CRNs

Notation. If $\mathcal{S}$ is a multiset, we will denote the set of distinct elements in $\mathcal{S}$ as $[\![\mathcal{S}]\!]$. If S is a set and k is a positive integer, then $k \cdot S$ denotes the multiset containing k copies of each element in S. Similarly, if $\mathcal{S}$ is a multiset, then $k \cdot \mathcal{S}$ denotes the union of k copies of $\mathcal{S}$. The set operations on multisets are defined in a usual way. In addition, we define the intersection $\mathcal{S} \cap T$ of a multiset $\mathcal{S}$ and a set T as $\mathcal{S} \cap (|\mathcal{S}| \cdot T)$, i.e., $\mathcal{S} \cap T$ contains only elements in $[\![\mathcal{S}]\!] \cap T$, and for each $x \in [\![\mathcal{S}]\!] \cap T$, the number of copies of x in $\mathcal{S} \cap T$ is the same as the number of copies of x in $\mathcal{S}$.

Definition 1 (Tagged CRN). *A* tagged chemical reaction network *is a tuple* $\mathbf{C} = \langle S, T, R, \mathcal{S}_0, \mathcal{T}_0 \rangle$ *with variables defined as follows:*

- S *is a set of signal species and* T *is the set of tag species, and* $S \cap T = \emptyset$.
- R *is a set of reversible or irreversible reactions, where each* $r \in R$ *is an ordered pair* $(\mathcal{I}_r, \mathcal{P}_r)$ *of multisets of signal and tag molecules such that* $\mathcal{I}_r \cap T = \{\tau_r^+\}$ *and* $\mathcal{P}_r \cap T = \{\tau_r^-\}$. *Intuitively, a reaction* $r = (\mathcal{I}_r, \mathcal{P}_r)$ *either consumes the molecules in* $\mathcal{I}_r$ *and produces the molecules* $\mathcal{P}_r$*, or, if the reaction is reversible, it can also consume* $\mathcal{P}_r$ *and produce* $\mathcal{I}_r$*. In the first case, we say that the reaction was applied in the* forward *direction and denote it as* $+r$*, in the second case in the* backward *direction and denote it as* $-r$*. The symbols* $+r$ *and* $-r$ *will be called* oriented *reactions.*
- $\mathcal{S}_0$ *is a multiset of signal molecules and* $\mathcal{T}_0$ *is a multiset of tag molecules present initially at time-step zero. The* volume *of CRN* $\mathbf{C}$ *is the number of molecules in* $\mathcal{S}_0 \cup \mathcal{T}_0$.

Tags limit the number of times a reaction can be applied in the same direction without being applied in the reverse direction. For example, if r is a reversible reaction and $\mathcal{T}_0$ contains only one copy of τ_r^+ and no copies of τ_r^-, then in any valid trace, the oriented occurrences of r has to alternate, starting with $+r$. If r is an irreversible reaction and $\mathcal{T}_0$ contains x copies of τ_r^+, then in any valid trace, there are at most x occurrences of $+r$ (and no occurrences of $-r$). Limiting the number of tags forces a system to recycle molecules in long traces.

Definition 2. *Consider a tagged CRN system* $\mathbf{C} = \langle S, T, R, \mathcal{S}_0, \mathcal{T}_0 \rangle$*. Define the* bandwidth *of signal species* s *as the maximum number of occurrences of* s *on the left side* $\mathcal{I}_r$ *(respectively, left* $\mathcal{I}_r$ *or right side* $\mathcal{P}_r$*) of any irreversible (respectively, reversible) reaction* $r \in R$*. Define the* maximum bandwidth $b_{\mathbf{C}}$ *(respectively,* total bandwidth $B_{\mathbf{C}}$*) of* $\mathbf{C}$ *as the maximum (respectively, the sum) of bandwidth over all signal species in* S*. Similarly, the* proper bandwidth *of signal species* s*, the* maximum proper bandwidth $\tilde{b}_{\mathbf{C}}$ *and the* total proper bandwidth $\tilde{B}_{\mathbf{C}}$ *are defined analogously but using* $\mathcal{I}_r \setminus \mathcal{I}_r \cap \mathcal{P}_r$ *instead of* $\mathcal{I}_r$ *and* $\mathcal{P}_r \setminus \mathcal{I}_r \cap \mathcal{P}_r$ *instead of* $\mathcal{P}_r$*. For any reversible reaction* $r \in R$*, let* t_r *be the maximum of the number of occurrences of* τ_r^+ *or* τ_r^- *in* $\mathcal{T}_0$*; and for any irreversible reaction* $r \in R$*, let* t_r *be the number of occurrences of* τ_r^+ *in* $\mathcal{T}_0$*. Let* $T_{\mathbf{C}}$ *be the sum of* t_r*'s over all reactions* $r \in R$*. We define the* x-copy *of* $\mathbf{C}$*, for* $x \in \mathbb{Z}^+$*, as the CRN* $\langle S, T, R, x \cdot \mathcal{S}_0, x \cdot \mathcal{T}_0 \rangle$.

Let $\rho = r_1, r_2, ..., r_m$ be a sequence of oriented reactions where $r_i \in R$ for all i. For reaction r if $\text{sign}(r) = +$, let $\mathcal{A}_r = \mathcal{I}_r$ and $\mathcal{B}_r = \mathcal{P}_r$ where as if $\text{sign}(r) = -$, let $\mathcal{A}_r = \mathcal{P}_r$ and $\mathcal{B}_r = \mathcal{I}_r$. The configuration of the system at each step i is defined as $(\mathcal{S}_i, \mathcal{T}_i)$ where $\mathcal{S}_i = (\mathcal{S}_{i-1} \setminus (\mathcal{A}_r \cap S)) \cup (\mathcal{B}_r \cap S)$ and, similarly, $\mathcal{T}_i = (\mathcal{T}_{i-1} \setminus (\mathcal{A}_r \cap T)) \cup (\mathcal{B}_r \cap T)$. A reaction sequence ρ is *valid* if $\mathcal{A}_r \cap S \subseteq \mathcal{S}_{i-1}$ and $\mathcal{A}_r \cap T \subseteq \mathcal{T}_{i-1}$ for all i, meaning that for each molecule in $\mathcal{A}_r$ there must be one in $\mathcal{S}_{i-1} \cup \mathcal{T}_{i-1}$ to remove. A *trace* is a valid reaction sequence.

2.2 The Main Upper Bound

Our main upper bound, Theorem 1, shows that in the multi-copy setting, any product of a tagged CRN can be produced within a number of reactions that is bounded by a function of the number of signal species, the bandwidth, and the number of tags of the CRN.

Theorem 1. *Let* $\mathbf{C} = \langle S, T, R, S_0, \mathcal{T}_0 \rangle$ *be a tagged CRN and let* $s_{end} \in S$*. If some trace of* $\mathbf{C}$ *produces* s_{end}*, then in a* $(|S| - |[\![\mathcal{S}_0]\!]|)(b_{\mathbf{C}} + \tilde{b}_{\mathbf{C}}(T_{\mathbf{C}} - 1)) \leq |S| b_{\mathbf{C}} T_{\mathbf{C}}$*-copy CRN of* $\mathbf{C}$*,* s_{end} *can be produced in at most* $(|S| - |[\![\mathcal{S}_0]\!]|)(b_{\mathbf{C}} + \tilde{b}_{\mathbf{C}}(T_{\mathbf{C}} - 1))T_{\mathbf{C}} \leq (|S| - 1) b_{\mathbf{C}} T_{\mathbf{C}}^2$ *steps.*

Proof. Let $\rho = r_1, r_2, \ldots, r_m$ be a valid sequence of oriented reactions in a single-copy system producing s_{end} starting from the initial set S_0. Consider any prefix of this sequence, say $\rho_i = r_1, \ldots, r_i$. Construct a new sequence ρ_i' by canceling all pairs $+r, -r$ for any reaction $r \in R$. It does not matter how these pairs are formed. Let S' be the set of signal molecules appearing on the left hand side of reactions in ρ_i'. Now, let us see what happens if we apply this sequence on the initial set $\mathcal{S}_0 \cup \mathcal{T}_0 \cup k \cdot S'$, where k is sufficiently large so that the reaction sequence is valid. We can make the following observations:

(1) The final number of copies of each signal species is the same as if we would apply ρ_i on $\mathcal{S}_0 \cup \mathcal{T}_0 \cup k \cdot S'$.
(2) For each reaction $r \in R$, ρ_i' contains either only forward or only backward occurrences of r (or no occurrences), and their number is limited by the number t_r of corresponding tags in $\mathcal{T}_0$. As a consequence, the length of ρ_i' is at most $T_{\mathbf{C}}$.
(3) Consider a signal molecule $s \in S'$. Since each reaction in ρ_i' removes at most $b_{\mathbf{C}}$ copies of s and the length of ρ_i' is at most $T_{\mathbf{C}}$, before each reaction in ρ_i', there are at least $k - \tilde{b}_{\mathbf{C}}(T_{\mathbf{C}} - 1)$ copies of s.
(4) Hence, it follows that if we set $k = b_{\mathbf{C}} + \tilde{b}_{\mathbf{C}}(T_{\mathbf{C}} - 1)$, then before each reaction in ρ_i', there are at least $b_{\mathbf{C}}$ copies of any signal in S', and hence, the reaction sequence is valid. Note that this is true even if we randomly permute reactions in ρ_i'.

For each signal s appearing in the single-copy trace and not appearing in the initial set $\mathcal{S}_0$, let $r_{\text{index}(s)}$ be the first reaction in ρ which produces a copy (or more) of s. Let $s_1, ..., s_n$ be the sequence of all signals not in $\mathcal{S}_0$ ordered by their indices, i.e., $\text{index}(s_1) \leq \text{index}(s_2) \leq \cdots \leq \text{index}(s_n)$. Furthermore, without loss

of generality we can assume $s_n = s_{\text{end}}$. Let $S_i = \{s_1, ..., s_i\}$. We can make one additional observation:

(5) For each s_i, the left side of each reaction in $\rho'_{\text{index}(s_i)}$ contains only signals in $[\![\mathcal{S}_0]\!] \cup S_{i-1}$. By (4), if we start in a configuration which contains the multiset of signals and tags $\mathcal{S}_0 \cup \mathcal{T}_0 \cup (b_{\mathbf{C}} + \tilde{b}_{\mathbf{C}}(T_{\mathbf{C}} - 1)) \cdot ([\![\mathcal{S}_0]\!] \cup S_{i-1})$, $\rho'_{\text{index}(s_i)}$ is a trace producing a copy of s_i.

Construction

(S1) Start with the initial set containing $b_{\mathbf{C}} + \tilde{b}_{\mathbf{C}}(T_{\mathbf{C}} - 1)$ copies of $[\![\mathcal{S}_0]\!]$ and the empty sequence of reactions.
(S2) For each $i = 1, .., n$: add $b_{\mathbf{C}} + \tilde{b}_{\mathbf{C}}(T_{\mathbf{C}} - 1)$ copies of $\mathcal{S}_0 \cup \mathcal{T}_0$ to the initial set and append $b_{\mathbf{C}} + \tilde{b}_{\mathbf{C}}(T_{\mathbf{C}} - 1)$ times sequence $\rho'_{\text{index}(s_i)}$ to the constructed sequence of reactions.

Claim 1. *After each step i in (S2), the constructed sequence is valid and the final configuration contains* $b_{\mathbf{C}} + \tilde{b}_{\mathbf{C}}(T_{\mathbf{C}} - 1)$ *copies of each signal in* $[\![\mathcal{S}_0]\!] \cup S_i$.

Proof. Proof by induction: *Base case:* For $i = 0$, after (S1), we have $b_{\mathbf{C}} + \tilde{b}_{\mathbf{C}}(T_{\mathbf{C}} - 1)$ copies of each signal in $[\![\mathcal{S}_0]\!]$ and the empty sequence of reactions is valid. *Induction step:* Inductive assumption: before step i, we have $b_{\mathbf{C}} + \tilde{b}_{\mathbf{C}}(T_{\mathbf{C}} - 1)$ copies of each signal in $[\![S_0]\!] \cup S_{i-1}$ and the sequence constructed so far is valid. By (5), if we add a copy of $\mathcal{S}_0$ and run the reaction sequence $\rho'_{\text{index}(s_i)}$ on the current configuration, the trace is valid. By (1), this newly added part (a copy of $\mathcal{S}_0$ and reactions in $\rho'_{\text{index}(s_i)}$) will not decrease the number of any signal. Finally, $\rho'_{\text{index}(s_i)}$ must contain the last reaction of $\rho_{\text{index}(s_i)}$, i.e., $r_{\text{index}(s_i)}$ which produces at least one copy of s_i. If we repeat this $b_{\mathbf{C}} + \tilde{b}_{\mathbf{C}}(T_{\mathbf{C}} - 1)$ times, we will still have at least $b_{\mathbf{C}} + \tilde{b}_{\mathbf{C}}(T_{\mathbf{C}} - 1)$ copies of signals in $[\![\mathcal{S}_0]\!] \cup S_{i-1}$ plus $b_{\mathbf{C}} + \tilde{b}_{\mathbf{C}}(T_{\mathbf{C}} - 1)$ copies of s_i. □

The bound: The construction uses $(n+1)(b_{\mathbf{C}} + \tilde{b}_{\mathbf{C}}(T_{\mathbf{C}} - 1))$ copies of $\mathcal{S}_0$, $n(b_{\mathbf{C}} + \tilde{b}_{\mathbf{C}}(T_{\mathbf{C}} - 1))$ copies of $\mathcal{T}_0$ and repeats $n(b_{\mathbf{C}} + \tilde{b}_{\mathbf{C}}(T_{\mathbf{C}} - 1))$ times the trace $\rho'_{\text{some index}}$. By (2), the length of each $\rho'_{\text{some index}}$ trace is at most $T_{\mathbf{C}}$, hence the total length of the constructed sequence is at most $n(b_{\mathbf{C}} + \tilde{b}_{\mathbf{C}}(T_{\mathbf{C}} - 1))T_{\mathbf{C}}$. Furthermore, n can be bounded by $|S| - |[\![\mathcal{S}_0]\!]|$. □

Finally, we restate Theorem 1 for copy-tolerant CRNs.

Theorem 2. *If a tagged CRN* $\mathbf{C} = \langle S, T, R, S_0, \mathcal{T}_0 \rangle$ *is* $|S|b_{\mathbf{C}}T_{\mathbf{C}}$*-copy-tolerant and* s_{end} *can be produced in* $\mathbf{C}$*, then the length of the shortest trace of* $\mathbf{C}$ *that produces* s_{end} *is at most* $(|S| - 1)b_{\mathbf{C}}T_{\mathbf{C}}^2$.

A natural question is whether we could improve the bound in condition (3) of the proof of Theorem 1. The following examples shows that it is not possible in general.

Example 1. Assume that ρ contains exactly an even number, T, of oriented reactions $+r_1, \ldots, +r_T$ designed as follows. First for every partition π of ρ into two sets ρ_1^π and ρ_2^π of same size, we introduce a new signal s_π. Let $\mathit{\Pi}$ be the set of all such partitions. Next, we define reactions $r_1, \ldots, r_T$ in such a way that each of these signals is either an input or a product of each reactions:

$$\mathcal{I}_{r_i} = \{s_\pi;\ r_i \in R_1^\pi, \pi \in \mathit{\Pi}\}\,,$$
$$\mathcal{P}_{r_i} = \{s_\pi;\ r_i \in R_2^\pi, \pi \in \mathit{\Pi}\}\,.$$

Note that after all reactions in ρ are applied, the number of copies of any of the signals s_π is not changed, since there is exactly $T/2$ reactions in ρ adding one copy of s_π and $T/2$ reactions removing one copy of s_π.

Now, we show that for any permutation of the reactions in ρ, there is a signal molecule with $k - T/2$ copies when the first $T/2$ reactions in this order are applied, and hence, k in (3) has to be set to at least $T/2$. Consider the partition π_0 of ρ into the first and the second $T/2$ reactions of this order. Then the signal s_{π_0} appears in the input set of the first $T/2$ reactions, and thus, the number of copies of s_{π_0} is $k - T/2$ after applying the first $T/2$ reactions.

3 Reachability Upper Bound for DSDs

In this section we first define the type of DSD to which our results apply, along with related notation needed for our results. We then provide our main upper bound, and conclude with a restatement of this result to obtain our reachability upper bound theorem for copy-tolerant DSDs.

3.1 Definition of DSDs

A *basic DNA strand displacement system (DSD)* is a pair $\mathit{\Delta} = (\mathcal{S}, \mathcal{C}_{init})$ of strands and initial configuration (secondary structure) for those strands, plus allowable *positional displacements*, defined as follows.

- $\mathcal{S}$ is a finite multiset of *strands*; $\mathcal{S}$ may contain many strands of a given type. Strands are composed of subsequences of finite strings of symbols, called *domains*. Domains are partitioned into two groups: *toeholds* and *long-domains*. Corresponding to each domain x is a complementary domain x^*; x is a toehold if and only if x^* is. The strands are partitioned into two groups: *signals* or *templates*. There is no bound on the number of toeholds and long-domains of a template or a signal. A *regional interval* is a sequence of domains beginning and ending with a toehold that alternates between toehold and long-domains. Each *template* strand is a concatenation of one or more regional intervals.

  We say that the DSD $\mathit{\Delta}$ has *simple signals*, if each signal in $\mathcal{S}$ is composed of exactly one toehold and one long-domain.

– $\mathcal{C}_{init}$ is an initial configuration, where a *configuration* is a secondary structure formed by the strands of $\mathcal{S}$ where domains can bind to their complements. Moreover, each signal strand is either unbound or is bound to a template strand by a single toehold and a single long domain that is adjacent to that toehold and each regional interval of its template must have exactly one open toehold. There are no intra-template bonds or intra-signal bonds. Note that this implies that configurations are pseudoknot-free and contain no hairpin loops. The *volume* of DSD Δ is the number of nucleotides, taken over all strands in the initial configuration $\mathcal{C}_{init}$.

Starting with the initial configuration, DSDs can progress through a sequence of configurations via positional strand displacements (PDs). PDs can move the open toehold of the regional interval to the right or to the left. A PD moving the open toehold to the right is specified by a positive number k, a template strand T and a signal strand called the invader, say of type I, see Figure 2(a), where we can now assume that only positions $k-1, k, k+1$ of template T are shown. The template should have at least $k+1$ domains. The domain d at position k of the template should be a long-domain and should be preceded at position $k-1$ by a toehold, say t. For the displacement to be applicable to a given configuration $\mathcal{C}$, it must be that in $\mathcal{C}$ an additional signal strand, which we refer to as the releasee, is bound to d at position k and to a toehold at position $k+1$ of the template T, and the toehold at position $k-1$ is unbound (open). The invader is unbound in $\mathcal{C}$ and contains the substring t^*d^*.

A displacement models the following steps in Figure 2(b,c,d), when toeholds and long-domains are actual DNA sequences. First, toehold t^* of the invader binds to the toehold t of the template at position $k-1$. Then a branch migration ensues, whereby domain d^* of the invader binds to d at position k of the template and the releasee is no longer bound at this position. Finally, if it exists, the bond between the releasee and the toehold at position $k+1$ is broken. Thus in the resulting configuration $\mathcal{C}'$, substring t^*d^* of the invader is bound to td on the template at positions $k-1$ and k and the releasee is unbound, see Figure 2(e).

Formally a *positional displacement (PD)* of DSD Δ is a tuple of the form (I, T, k, z), where I is a signal strand type, T is a template strand, k is a positive integer and $z \in \{L, R\}$. PD (I, T, k, z) is *applicable* to a configuration $\mathcal{C}$ if the following conditions hold:

1. Strand T has at least $k+1$ domains and the kth domain, say d, must be a long-domain. Also a strand of type O, called the releasee, is bound to the kth domain of T.
2. In the configuration $\mathcal{C}$, a strand of type I is unbound.
3. If $z = R$ the following conditions hold (conditions for $z = L$ are symmetric):
   (a) The $(k-1)$st domain of T must exist and be a toehold, say t.
   (b) A strand of type I must contain substring t^*d^*. (If $z = L$, it must contain d^*t^*.)
   (c) The releasee must also be bound to a toehold at position $k+1$ of T. No other domains of the releasee are bound.

(d) In the configuration $\mathcal{C}$ the toehold at position $k-1$ of strand T is unbound. We call this toehold the *input toehold* of PD (I, T, k, z).

The PD must *release* exactly one strand of type O. Suppose that PD (I, T, k, z) is applicable to $\mathcal{C}$. Let $\mathcal{C}'$ be obtained from $\mathcal{C}$ by removing the bonds between T and the releasee and by adding bonds either between any substring t^*d^* of an unbound strand of type I of $\mathcal{C}$ and the domains td at positions $k-1$ and k of T if $z = L$, or between any substring d^*t^* of I and the substring dt at positions k and $k+1$ of T if $z = R$. Then we say that (I, T, k, z) induces $\mathcal{C}'$ from $\mathcal{C}$. We say that a signal is *simple* if the whole string for the signal consists of t^*d^* or d^*t^*. A DSD is *simple* if all the signals in the DSD are simple. This definition excludes *cooperativity* where two invading strands release a single releasee or one invading strand releases two releasees, because, by definition, every PD must be initiated by one invader and release exactly one releasee.

A sequence of PDs $\rho = p_1, p_2, \ldots p_{|\rho|}$ is *valid* with respect to $\mathcal{C}_{init}$ if there is a sequence $\mathcal{C}_1, \mathcal{C}_2, \ldots \mathcal{C}_{|\rho|+1}$ of configurations of Δ with $\mathcal{C}_1 = \mathcal{C}_{init}$ such that for all $i, 1 \leq i \leq |\rho|$, p_i is applicable to $\mathcal{C}_i$ and induces $\mathcal{C}_{i+1}$ from $\mathcal{C}_i$. When $\mathcal{C}_{init}$ is clear from the context, we simply say that ρ is valid. A valid sequence *produces* a strand $s \in \mathcal{S}$ if in $\mathcal{C}_{|\rho|+1}$, the strand s is unbound. Let Invaders(ρ) be the set of types of invaders of ρ. Let Unbound$(\rho, \mathcal{C}_{init})$ be the set of types of unbound signals in $\mathcal{C}_{|\rho|+1}$ and Unbound(ρ) the set of types of unbound signals in $\mathcal{C}_1 \cup \cdots \cup \mathcal{C}_{|\rho|+1}$.

Let $\rho = p_1, p_2, ..., p_{|\rho|}$ be a sequence of PDs. The *regional subsequence* $\rho(T[u, v])$ is the subsequence of ρ whose PDs $p_i = (I_i, T_i, k_i, z_i)$ have positions k_i inside $T[u, v]$.

3.2 The Upper Bounds

First, we use the fact that a DSD with simple strands can be simulated by a tagged CRN with volume that is polynomial in the volume of DSD, and thus we can use the bound in Theorem 1 to obtain the following result. If $\Delta = (\mathcal{S}, \mathcal{C}_{init})$ is a DSD, we define $\Delta^{(x)}$ to be the DSD $(x \cdot \mathcal{S}, x \cdot \mathcal{C}_{init})$.

Theorem 3. *Let Δ be a DSD with simple signals. Let B be the number of types of initially bound signal strands and D be the total number of long-domains of all templates. If Δ can produce s_{end}, then $\Delta^{(2D(2D+B))}$ can produce s_{end} via a sequence of at most $4D^2(2D+B)$ PDs.*

As shown in Figure 3, this strategy will not work in the case of general signal strands. Instead of simulating a DSD by a tagged CRN, in Theorem 4, we will prove a bound for general (i.e., not with simple signals) DSDs directly, reusing some ideas of the proof for tagged CRNs.

Let Δ be a DSD. Roughly, our goal is to show that if there is a valid sequence of PDs that produces a given signal s_{end} from $\mathcal{C}_{init}$, then in a DSD with many copies of $\mathcal{C}_{init}$ there is a valid sequence of PDs that produces s_{end} in a number of steps that is bounded by a polynomial in the volume of Δ. We will build up to the statement and proof through a series of definitions and claims. Our

polynomial bound will be a function of two attributes of Δ: the number B of types of signal strands that are all bound (i.e., every copy is bound) in $\mathcal{C}_{init}$ and the total number D of long-domains of all templates in $\mathcal{C}_{init}$.

Let $\alpha = p_1, p_2, ..., p_{|\alpha|}$ be a valid sequence of PDs that produces s_{end}. For each type s of signal strand that is Unbound$(\alpha) \setminus \mathcal{C}_{init}$, let index$(s)$ be the index of the first PD of α that releases s. Let $s_1, \ldots, s_B$ be the sequence of all such signals ordered by their indexes, i.e., index$(s_1) <$ index$(s_2) < \ldots <$ index(s_B). Let $S_i = \{s_1, \ldots, s_i\}$. We assume without loss of generality that there are B such types and also that $s_B = s_{end}$. Let $\alpha_i = p_1, p_2, \ldots, p_{\text{index}(s_i)}$.

Let $T[u, v]$ be a regional interval, $d = (v - u)/2$ the number of long-domains in $T[u, v]$, and let $\alpha_i(T[u, v]) = p_1, p_2, ..., p_{|\alpha_i(T[u,v])|}$, where $p_j = (I_j, T_j, k_j, z_j)$ for every $j = 1, \ldots, |\alpha_i(T[u, v])|$. We construct a subsequence $\beta_i(T[u, v])$ of the PDs in $\alpha_i(T[u, v])$. The PDs in this subsequence will be of two types, *marked* and *connector*.

Markers. Mark the first PD p_1 of $\alpha_i(T[u, v])$, and then mark the *last* PD of $\alpha_i(T[u, v])$ to bind to each long-domain in the regional interval $T[u, v]$. Let $p_{m_1}, \ldots, p_{m_{d+1}}$ be the subsequence of marked PDs ($1 = m_1 < m_2 < \cdots < m_{d+1}$). It is easy to see that the sequence of PD positions, $k_{m_2}, \ldots, k_{m_d}$, consists of two interleaved monotonic subsequences: $U = u + 1, u + 3, \ldots, \ell - 2$ and $V = v - 1, v - 3, \ldots, \ell + 2$, where ℓ is the long-domain position of the last PD in $\alpha_i(T[u, v])$. Furthermore, the marked PDs with the long-domains in the first subsequence have direction R and in the second subsequence direction L.

Connector sequences. Now, we must connect the marked PDs by introducing *connector sequences* of PDs between each consecutive pair of marked PDs with the goal being for each subsequent PD to use the toehold opened by the previous PD.

Let $\bar{z}$ indicate the opposite direction from z. For the connector sequence connecting p_{m_1} and p_{m_2}, select as a connector the *first* PD in $\alpha_i(T[u, v])$ with direction $\bar{z}_{m_2}$ that binds to each long-domain of $T[u, v]$ between positions k_{m_1} and k_{m_2} inclusive. It is easy to see that either all selected connector PDs are before p_{m_2} in the sequence $\alpha_i(T[u, v])$, or $p_{m_1} = p_{m_2}$ and the connector sequence consists of the same PD. In the second case, p_{m_1} is the only PD in $\alpha_i(T[u, v])$ with the long-domain position $u + 1$ or $v - 1$.

Consider $j = 2, \ldots, s$. Each PD of the connector sequence connecting p_{m_j} to $p_{m_{j+1}}$ will be between p_{m_j} and $p_{m_{j+1}}$ in the sequence $\alpha_i(T[u, v])$. We will consider two cases.

1. If $z_{m_j} = z_{m_{j+1}}$, then no connector PDs are needed.
2. If $z_{m_j} \neq z_{m_{j+1}}$, then we select the connectors as follows. In the subsequence between reactions p_{m_j} and $p_{m_{j+1}}$, choose as a connector the *first* PD that binds to each position between k_{m_j} and $k_{m_{j+1}}$, excluding position k_{m_j} and including position $k_{m_{j+1}}$. Note that each PD in this connector sequence must have direction z_{m_j}.

The construction is illustrated in Figure 5. The sequence $\beta_i(T[u, v])$ contains all the marked PDs and all the connector PDs, with distinct indices. Note that

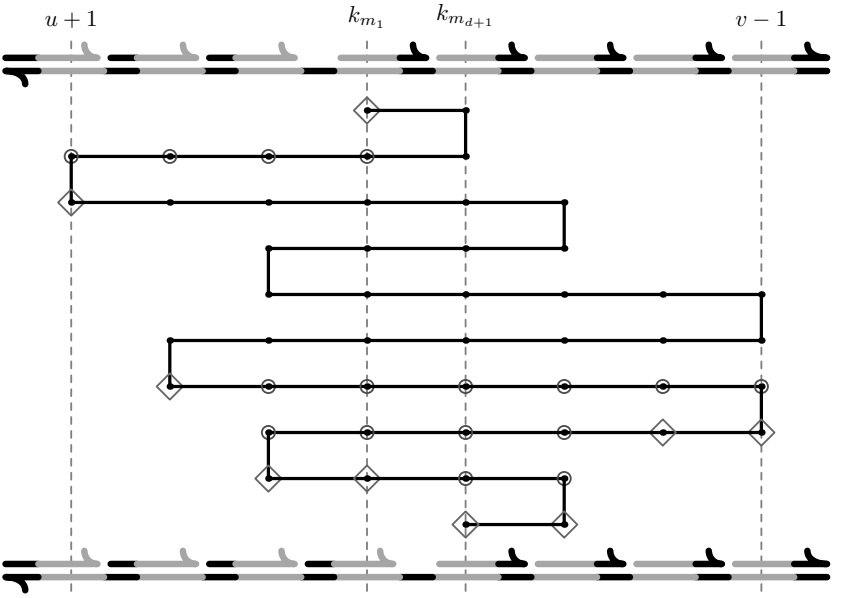

Fig. 5. An example of construction of the $\beta_i(T[u,v])$ subsequence. At the top is the form of the initial configuration of the regional interval and at the bottom the final configuration. Each dot represents a PD of regional subsequence $\alpha_i(T[u,v])$, each diamond a marked PD and each circle a connector PD. The sequence of PDs $\beta_i(T[u,v])$ is then a subsequence of $\alpha_i(T[u,v])$ which contains only the marked and connector PDs.

this is a subsequence of $\alpha_i(T[u,v])$ since for every $j = 1,\ldots,s$, the connector sequence connecting p_{m_j} to $p_{m_{j+1}}$ contains only PDs between between p_{m_j} and $p_{m_{j+1}}$.

We next provide a sequence of claims that we use to prove our main result. All proofs can be found in the full version of the paper.

Claim 2. *Each PD, p_{m_j} for $j \geq 2$ in sequence $\beta_i(T[u,v])$ can use the toehold opened by the previous PD in the sequence. The PD p_{m_1} can use the initially open toehold.*

Claim 3. *The length of $\beta_i(T[u,v])$ is at most $(d+1)(d+2)/2$, where $d = (v-u)/2$ is the number of long-domains in $T[u,v]$.*

Claim 4. *The length of β_i is at most $(D+1)(D+2)/2$ and thus $|\text{Invaders}(\beta_i)| \leq (D+1)(D+2)/2$. Also, $\text{Invaders}(\beta_i)$ contains only types of unbound strands of $\mathcal{C}_{init}$ or strand types in $S_{i-1} = \{s_1,\ldots,s_{i-1}\}$.*

Claim 5. *β_i is valid with respect to*

$$\mathcal{C}_{init} \cup (D+1)(D+2)/2 \cdot (\mathcal{C}_{init} \cup S_{i-1}).$$

Moreover,

$$(D+1)(D+2)/2 \cdot (\mathcal{C}_{init} \cup S_{i-1}) \subseteq \text{Unbound}(\beta_i, \mathcal{C}_{init} \cup (D+1)(D+2)/2 \cdot (\mathcal{C}_{init} \cup S_{i-1})).$$

Claim 6. *Let $\beta_i^{(D+1)(D+2)/2}$ denote the sequence β_i concatenated $(D+1)(D+2)/2$ times, modified just so that the PDs of each copy refer to templates of different copies of $(D+1)(D+2)/2 \cdot \mathcal{C}_{init}$. Then $\beta_i^{(D+1)(D+2)/2}$ is valid with respect to the configuration*

$$(D+1)(D+2)/2 \cdot \mathcal{C}_{init} \cup (D+1)(D+2)/2 \cdot (\mathcal{C}_{init} \cup S_{i-1}).$$

Moreover,

$$(D+1)(D+2)/2 \cdot (\mathcal{C}_{init} \cup S_i) \subseteq \text{Unbound}(\beta_i^{(D+1)(D+2)/2}, \\ (D+1)(D+2)/2 \cdot \mathcal{C}_{init} \cup (D+1)(D+2)/2 \cdot (\mathcal{C}_{init} \cup S_{i-1})).$$

The proof of our main technical result, Theorem 4, follows from the preceding claim.

Theorem 4. *Let Δ be a DSD with B types of initially bound signal strands and let D be the total number of long-domains of all templates. If Δ can produce s_{end}, then $\Delta^{((D+1)(D+2)(B+1)/2)}$ can produce s_{end} via a sequence of at most $(D+1)^2(D+2)^2B/4$ PDs.*

Finally, we restate Theorem 4 for copy-tolerant DSDs. We say that a DSD is x-copy-tolerant if the length of the shortest PD sequence that produces any signal strand s in Δ and in $\Delta^{(x)}$ is the same. A DSD is copy-tolerant if it is x-copy-tolerant for all x.

Theorem 5. *Let Δ be a DSD with B types of initially bound signal strands and let D be the total number of long-domains of all templates. If Δ can produce s_{end} and Δ is $(D+1)(D+2)(B+1)/2$-copy tolerant, then Δ can produce s_{end} via a sequence of at most $(D+1)^2(D+2)^2B/4$ PDs.*

4 Open Questions

There are many open questions about the potential for CRNs and DSDs to be correct in the multi-copy setting. First, can our reachability upper bound results be strengthened? There are two possible ways to strengthen our result for CRNs (Theorem 2): either by reducing the length of the shortest computation needed to produce s_{end} or to show that the system is not x-copy tolerant for some $x < |S|b_C T_C$. Similarly, there are two ways to strengthen the reachability upper bounds for DSDs.

Also, can our result on DSDs be extended to DSDs with more complex primitives, such as cooperative strand displacement [12] or irreversible reactions? What if long-domains can form intra-molecular bonds, e.g., forming hairpins, in addition to inter-molecular bonds?

This paper considers only reachability bounds, i.e., bounds on the number of reactions (steps) needed to reach (produce) a given product. However, real chemical reaction networks behave stochastically, with rates that depend on

relative quantities of species. It is plausible that the lack of robustness implied by our theorems, i.e., errors that occur in the multi-copy setting in CRNs that fail to satisfy the conditions of the theorem, would be very unlikely to occur in some CRNs and thus would not be an issue in a real system. Analyses of robustness of CRNs under stochastic assumptions, perhaps computing expected hitting times, would help us better understand the degree to which robustness issues are a problem.

References

1. Condon, A., Hu, A.J., Maňuch, J., Thachuk, C.: Less haste, less waste: On recycling and its limits in strand displacement systems. J R Soc. Interface (2012)
2. Cardelli, L.: Two-domain DNA strand displacement. In: Proc. of Developments in Computational Models (DCM 2010). Electronic Proceedings in Theoretical Computer Science, vol. 26, pp. 47–61 (2010)
3. Qian, L., Soloveichik, D., Winfree, E.: Efficient Turing-Universal Computation with DNA Polymers. In: Sakakibara, Y., Mi, Y. (eds.) DNA 16 2010. LNCS, vol. 6518, pp. 123–140. Springer, Heidelberg (2011)
4. Qian, L., Winfree, E.: Scaling up digital circuit computation with DNA strand displacement cascades. Science 332, 1196–1201 (2011)
5. Qian, L., Winfree, E., Bruck, J.: Neural network computation with DNA strand displacement cascades. Nature 475, 368–372 (2011)
6. Seelig, G., Soloveichik, D., Zhang, D.Y., Winfree, E.: Enzyme-free nucleic acid logic circuits. Science 314(5805), 1585–1588 (2006)
7. Soloveichik, D.: Robust stochastic chemical reaction networks and bounded tau-leaping. J. Comput. Biol. 16(3), 501–522 (2009)
8. Soloveichik, D., Cook, M., Winfree, E., Bruck, J.: Computation with finite stochastic chemical reaction networks. Nat. Comp. 7, 615–633 (2008)
9. Soloveichik, D., Seelig, G., Winfree, E.: DNA as a universal substrate for chemical kinetics. Proc. Nat. Acad. Sci. USA 107(12), 5393–5398 (2010)
10. Yurke, B., Mills, A.P.: Using DNA to power nanostructures. Genet. Program. Evolvable Mach. 4(2), 111–122 (2003)
11. Yurke, B., Turberfield, A.J., Mills Jr., A.P., Simmel, F.C., Neumann, J.L.: A DNA-fuelled molecular machine made of DNA. Nature 406, 605–608 (2000)
12. Zhang, D.Y.: Cooperative hybridization of oligonucleotides. J. Am. Chem. Soc. 133, 1077–1086 (2011)
13. Zhang, D.Y., Turberfield, A.J., Yurke, B., Winfree, E.: Engineering entropy-driven reactions and networks catalyzed by DNA. Science 318, 1121–1125 (2007)
14. Zhang, D.Y., Seelig, G.: Dynamic DNA nanotechnology using strand displacement reactions. Nature Chemistry 3, 103–113 (2011)