
Two-dimensional tile displacement can simulate cellular automata

Erik Winfree^{1,2,3*} and Lulu Qian^{1,2,3*}

¹*Bioengineering*, ²*Computer Science*, ³*Computation and Neural Systems*

California Institute of Technology, Pasadena, CA 91125, USA

**e-mail: winfree@caltech.edu, luluqian@caltech.edu*

Tile displacement is a newly-recognized mechanism in DNA nanotechnology that exploits principles analogous to toehold-mediated strand displacement but within the context of self-assembled DNA origami tile arrays. Here, we formulate an abstract model of tile displacement for the simplest case: individual assemblies interacting with monomer tiles in solution. We give several constructions for programmable computation by tile displacement, from circuits to cellular automata, that vary in how they use energy (or not) to drive the system forward (or not), how much space and how many tile types they require, and whether their computational power is limited to PTIME or PSPACE with respect to the size of the system. In particular, we show that tile displacement systems are Turing universal and can simulate arbitrary two-dimensional synchronous block cellular automata, where each transition rule for updating the state of a 2×2 neighborhood is implemented by just a single tile.

Keywords: DNA origami, tile displacement, cellular automata, reversible computation

1. INTRODUCTION

A guiding principle in theoretical computer science has been “mechanism-to-model” exploration of connections between physical implementation and computational capabilities. For example, what can be computed by systems of AND gates and OR gates is strictly less than what can be computed by systems of NOR gates, which in turn is less than what can be computed by finite state machines coupled with an unbounded memory tape [1]. Likewise, molecular programming theory aims to understand how fundamental molecular mechanism can be used to build systems, and how the choice of mechanism determines the range of what can be built. An example would be the self-assembly of molecular structures by programmable cooperative binding, which can reliably grow structures that cannot grow reliably via non-cooperative binding [2]. When a new molecular mechanism is discovered, it is of interest to understand the nature – the limitations and capabilities – of systems that exploit that mechanism. Doing so entails formulation of an abstract model that captures the essential features of the mechanism, which can then be rigorously analyzed.

Since its invention two decades ago [4], toehold-mediated DNA strand displacement has been a central mechanism for programming dynamical function in DNA nanotechnology [5, 6]. As shown in figure 1A, a stable complex of two strands can be reconfigured such that an invading strand replaces the original partner via a branch migration process. The single-stranded portion of the original complex – known as the toehold because that is where the invading strand initiates contact – is critical

for the displacement: an invader that does not match and bind to the toehold may be a million-fold slower to perform the displacement, and thus neglected as “leak”. In abstract models that consider networks of more complex (but still non-pseudoknotted) DNA molecules that interact in solution using toehold-mediated strand displacement reactions – including both the irreversible mechanism shown here and a reversible variant known as “toehold exchange” – have been shown capable of simulating arbitrary formal chemical reaction network dynamics and even Turing-universal computation [7, 8, 9, 10, 11, 12, 13]. However, a limitation of these results is that they are intrinsically distributed computations, where state is encoded within a collection of molecules in solution, and therefore a single test tube can perform only one computation at a time. More complex molecular mechanisms, such as the hypothetical polymer-modifying enzymes envisioned by Bennett [14], are in principle capable of performing independent Turing-universal computation in parallel in the same test tube.

The mechanism of “tile displacement”, shown in figure 1B, was recently discovered during investigations into why the self-assembly of DNA origami tiles [15, 16, 17] failed to become kinetically trapped in undesired intermediates that the naive theory predicted [3]. There is a strong analogy to toehold-mediated strand displacement. Beyond using components that are two orders of magnitude larger than the individual strands involved in strand displacement, the tile displacement mechanism has several distinct features. (1) Nucleotides are on multiple helices that are oriented orthogonally to the axis of branch

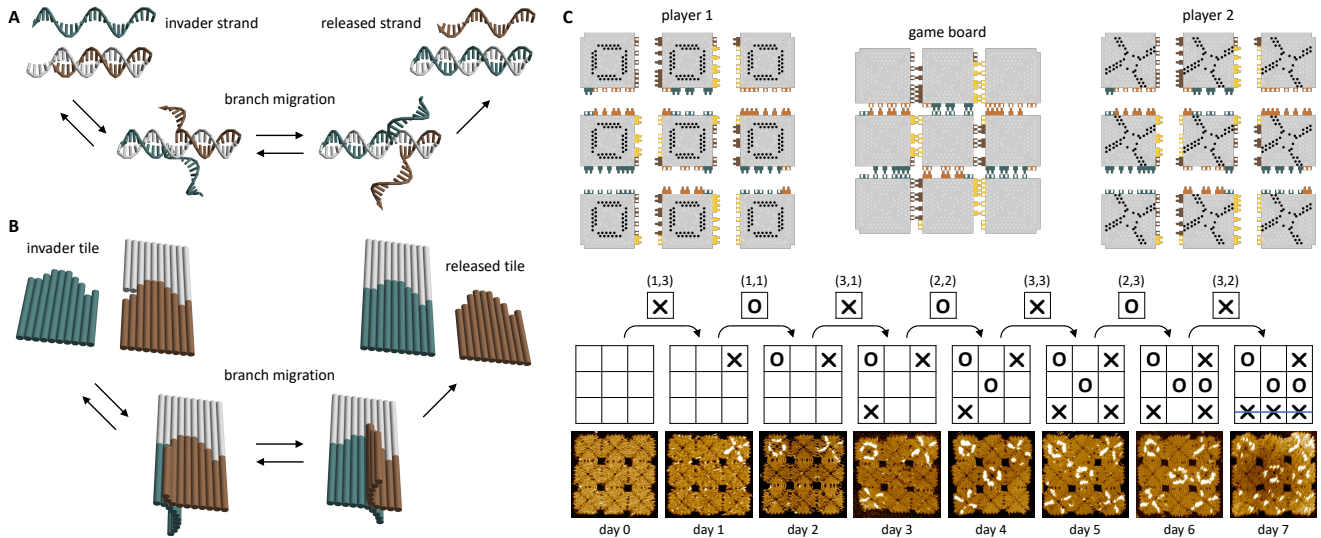


Figure 1: **(A)** Strand displacement mechanism. For scale, the DNA molecules are roughly 2 nm in diameter and 7 nm long. **(B)** Tile displacement mechanism. These hypothetical 10-helix DNA origami tiles are smaller than the 22-helix square tiles from [3]. **(C)** A tic-tac-toe game implemented using tile displacement (adapted from reference [3]). Scale bar for atomic force microscopy (AFM) images is 100 nm.

migration, rather than being on a single helix that is parallel to (identical to) the axis of branch migration. (2) Tile-tile binding may be due to multiple helix-end stacking bonds [18] in addition to (or instead of) being due to base-pair formation. (3) Toehold and branch migration domain specificity can be encoded both by tile geometry and by the sequences in multiple very short (1 or 2 nt) sticky ends, rather than being exclusively encoded by nucleotide sequences within a single helix. (4) The released tile will be less flexible than a single-stranded oligonucleotide. (5) Rather than having just one “side” and initiating displacement via a single toehold, tiles may have many (e.g. four) sides and may initiate displacement via cooperative action of multiple toeholds, as highlighted by the replacement of the central tile of a 3×3 tic-tac-toe game board [3] shown in figure 1C. Despite these differences, it remains that tile displacement is highly sensitive to toehold and branch migration sequences, such that the kinetics of tile displacement without a matching toehold may be orders of magnitude slower and similarly negligible as “leak”. Indeed, systems of interacting tile monomers and tile assemblies were shown to be reconfigurable by toehold-mediated tile displacement [3], and the same or similar constructs ought to be sufficient to implement more complex information-processing networks following, for example, the seesaw motif for circuits [19] or the two-domain scheme for formal chemical reaction network dynamics [10].

Here we are interested in whether the tile displacement mechanism enabled new ways of programming dynamical behaviors, beyond simply replicating strand displacement on a larger scale. The ability to perform displacement within a two-dimensional array being an especially novel feature of tile displacement, we ask whether – unlike ex-

isting strand displacement constructions – reconfiguration of a single tile assembly in a constant soup of monomer tiles might be sufficient for substantial computation, in which case parallel computation could be achieved with each tile assembly performing an independent computation. We present three results. First, with a feedforward Boolean circuit laid out on the initial array, there is a tile set that, via displacement, propagates signals along wires and executes the specified logic. This system is powered by the energy of toehold formation; the final state is in an energy minimum and cannot be reused. Second, as a simplification and generalization of the first construction, any one-dimensional cellular automaton can be directly translated into a set of tiles such that a wave of tile displacement converts an assembly, initially empty but for the input, into the space-time history of the cellular automaton. This system is powered by a concentration difference between the monomer tiles that are invading over those that are displaced. The above two constructions displace each tile in the original array at most once, using energy that is linear in the area used. The third construction addresses whether iterated computation can be performed in-place, which requires replacing the tile at a given location an unbounded number of times. Remarkably, using locally reversible asynchronous tile displacement, we can simulate arbitrary synchronous block cellular automata that use the 2×2 Margolus neighborhood, including his globally reversible Billiard Ball Model that is known to be Turing universal by simulation of infinite or finite recurrent Boolean circuits [20, 21, 22]. A key issue is how to bias the computation forward; we show that it is enough to include a large empty part of the array into which entropy is injected.

This paper does not aim for novel advances in molecu-

lar programming that make technical applications closer to reality. We are sharing these observations mainly because we find them to be beautiful and surprising. Tile displacement may indeed be useful for reconfiguration of adaptive molecular systems, but for most implementation goals that are merely computational, there are more direct and more reliable ways to achieve them using other mechanisms in DNA nanotechnology. However, it is remarkable that a molecular mechanism accidentally discovered in the laboratory gives rise to a theoretical model with such natural and direct connections to an esoteric but well-studied model of computation that arose in the study of the fundamental physics and ballistic motion. We hope you will see through our imperfect figures and clumsy explanations to see the poetry within the concepts [23].

2. TILE DISPLACEMENT MODEL

The abstract model developed in this work, which we call the Single-Assembly Tile Displacement (SATiDi) model, defines the behavior of a single tile assembly within a sea of monomer tiles. There are a number of assumptions that must hold in order for the model to be experimentally plausible, while also allowing its definition to be fairly clean.

Singularity. The concentration of multi-tile assemblies is sufficiently low (e.g. there is exactly one) that they do not interact with each other.

Monomers. Binding between two monomer tiles is sufficiently weak (at the given temperature and concentrations) that any dimers are fleeting and their presence can be neglected.

Stability. Tiles within an assembly (e.g. with four neighbors) are sufficiently strongly attached that they will not dissociate; tiles on the boundaries and corners, with only two or three neighbors, may have special binding interactions that make them as stable as the internal tiles.

No growth. With a single-side attachment being unstable for dimers, similarly new tiles may not attach by a single side to a facet of a multi-tile assembly. When the assembly is rectangular, as will be exclusively considered here, that means the number of tiles in an assembly will never change.


Full toeholds. For consistent tile displacement kinetics, we require that the displacement process on each side has its own mediating toehold, so a tile that is bound to four neighbors will be displaced by a tile that forms a toehold on each of the four sides. While fewer toeholds may be sufficient for displacement, it is all too plausible that their kinetics would be irregular; our simulator will issue a warning whenever such a displacement possibility is encountered.

Energetics. Tile displacement reactions must be either energetically neutral or energetically downhill, i.e. the number of toeholds formed is either the same as or more than the number of toeholds broken.

Uniform design. Each side of every tile will consist of a first toehold domain, a branch migration domain, and a second toehold domain. We will assume that the branch migration domains are distinct on the north, east, west, and south such that they force tile to maintain a specific orientation (although non-oriented versions of the model could be formulated when non-oriented tiles are desirable). Because branch migration domains cannot be changed by tile displacement, they will not be formally represented or accounted for in the model.

The model is illustrated in figure 2A, where both a valid neutral tile displacement and a valid downhill tile displacement are shown. Neutral displacement can be thought of as generalizing the “toehold exchange” mechanism from strand displacement [24]: formation of mediating toehold ensures fast kinetics, while dissociation of prior toeholds both ensures that the reaction is energetically neutral and opens up those toeholds for use in subsequent steps, as shown in figure 2B. Toehold exchange in tile displacement was demonstrated experimentally [3], although not in the exact geometric configuration required here; tuning of toehold strength (e.g. via temperature) would be required to ensure that the dissociation step (which may involve breaking four toeholds simultaneously) is sufficiently fast while still being effective for mediating the reaction.

Formally, a SATiDi system is defined by (1) a finite set of square tile types \mathcal{S} , each of which specifies an ordered pair of bond types (toeholds) for each of the four sides, (2) a bond strength function for each bond type b , $E_b > 0$, (3) a concentration for each tile type i , c_i , and (4) a standard tile displacement rate constant k . The associated set of assemblies \mathcal{A} consists of finite arrays of tile types (or *empty*). Given a specific assembly, we say that a specific toehold on a specific tile is *closed* if the corresponding toehold on the neighboring tile has the same (i.e. matching) bond type (i.e. they form a bond), while we say that it is *open* otherwise. The *bond energy* $E(A)$ of an assembly A is the sum $\sum -E_b$ over all closed toeholds in the assembly, while the *free energy* $G(A)$ of the assembly is its bond energy plus the sum $\sum \ln c_i/c_0$ over all tiles in the assembly, where c_0 is the reference concentration (e.g. 1 M). Given these, we associate a formal chemical reaction network (CRN) with reactions



where A is an assembly with tile t_j at some position x , A' is the same assembly but with t_i instead at that same position x , and t_i is a *valid displacement*: on all sides where t_j has a neighbor, t_i forms a matching bond with (at




Figure 2: (A) Abstract tile displacement model. (B) Irreversible and reversible strand displacement. (C) DNA origami tile implementation. (D) Single-stranded tile implementation.

least one) open toehold, and the total number of matching bonds increases or stays the same (i.e. the assembly’s bond energy decreases or stays the same). When all reactions are reversible, which implies that the bond energy of the assembly never changes, the CRN satisfies detailed balance with respect to the assembly bond energy, with monomer tiles having zero energy.

We consider standard stochastic kinetics according to Gillespie simulation with chemostatted constant monomer tile concentration [25, 26]. For an initial state containing a single assembly, this results in a finite continuous-time Markov chain (CTMC) where the set of states are all assemblies reachable via tile displacement reactions, and the transition $A \rightarrow A'$ involving invading tile t_i , as above, will have rate $k \times c_i$. If all reactions are reversible, this CTMC will satisfy detailed balance with respect to the assembly free energy, such that the equilibrium probability of assembly A is

$$p(A) = \frac{1}{Z} e^{-G(A)} \quad \text{with} \quad Z = \sum_{A'} e^{-G(A')}$$

where the partition function sum Z is taken with respect to all assemblies reachable by tile displacement.

A tile displacement system simulation is considered *unreliable* if at any time there is an energetically neutral or downhill tile replacement that does not form at least one new toehold with each neighboring tile. In this case, the simulation issues a *warning*, as illustrated in figure 2A.

We briefly consider possible experimental implementations of single-assembly tile displacement systems. Figure 2C shows the motivating DNA origami tile scheme, using a geometrically-symmetrical tile design modeled after those used in several prior experimental works [27, 17,

3, 28, 29]. More speculatively, in figure 2D we envision an implementation that makes use of topologically two-dimensional arrays of single-stranded tiles [30, 31], which have been shown to tolerate a wide variety of structural variations (including single-stranded regions as we would require for toeholds) and permitting strand displacement reactions that remove tiles from the array [32, 33]. However, single-strand tile reactions analogous to the four-toehold reversible tile displacement reactions required here have not been experimentally demonstrated. Regardless of whether considering DNA origami tiles or single-stranded tiles – or something else – a major obstacle to any experimental implementation would be the creation of the initial array with a desired initial pattern. One possible avenue – still difficult – would be to initially assemble a uniquely-addressed DNA origami array [17] or single-stranded tile array [31], use that array to geometrically organize the desired pattern of non-uniquely-addressed tiles needed for tile displacement behaviors, and then via photocleavable bonds or other mechanisms, remove and dispose of the uniquely-addressed array. But for now, we will assume that arbitrary initial assemblies can be synthesized.

3. WIRES, GATES, AND CIRCUITS

To get a feel for how tile displacement systems can be programmed, we begin with the most basic task: signal transmission. As shown in figure 3A, this can be accomplished using a single tile type (“wire”) that is used in the initial assembly to indicate where the wire is, plus a single tile type (“signal”) that carries the signal x . Two additional tile types (“top” and “bottom”) are used to provide neighboring tiles for the wire, as in general the

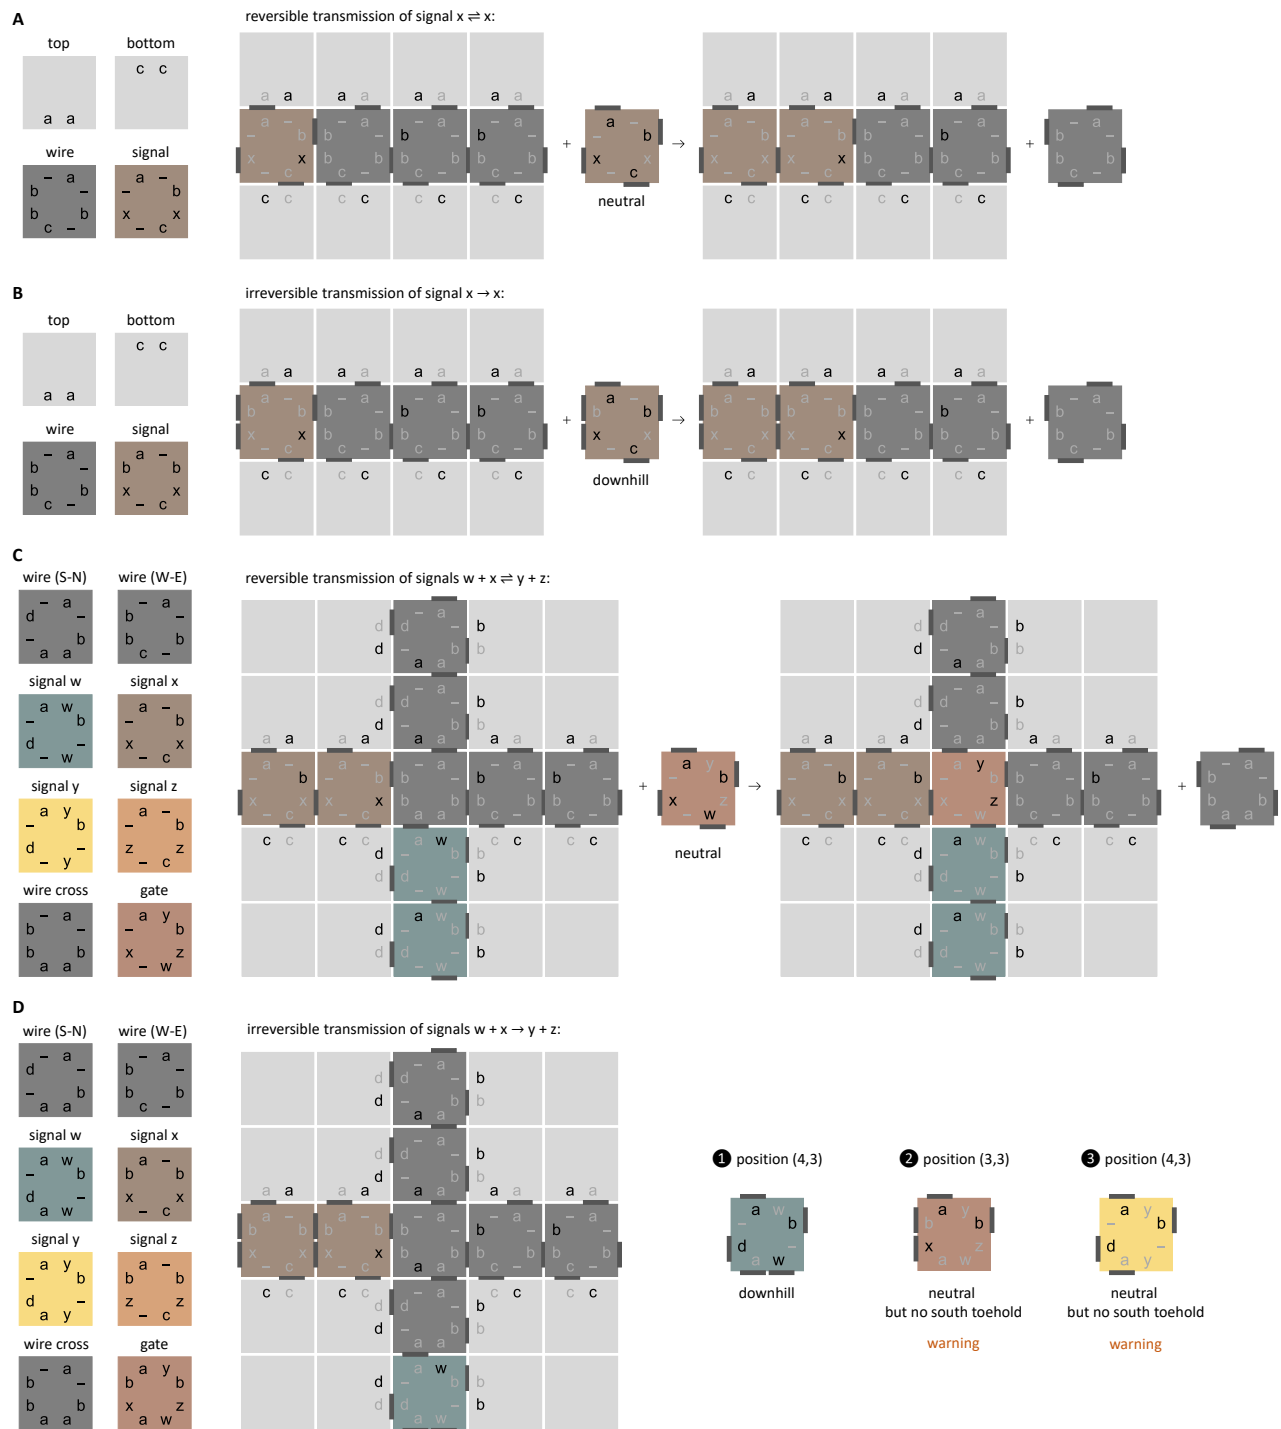


Figure 3: **(A)** Reversible wire. All toeholds are strength 1 except for toehold “-”, which is inert, i.e. strength 0. The top and bottom tiles have toeholds that are not shown, such that in the assembly the unlabeled sides are bound to each other via matching closed toeholds. In assemblies, closed toeholds are shown with light grey labels and a solid dark grey bar indicates their bond. For open toeholds, black or light grey is used to highlight relevant locations for the tile displacement reaction of interest, but have no formal meaning. In the monomer tile that is a reactant of the indicated reaction, solid dark grey bars indicate where new toehold bonds will be formed. Here and in later figures, only the forward reaction is shown for any reversible reactions (i.e. the assembly and monomer tile that are the products of the indicated reaction do not have their relevant toeholds highlighted for the backward reaction). **(B)** Irreversible wire. **(C)** Reversible wire cross. The initial assembly shown here illustrates the moment when both reversible signals arrive at the wire cross location. At this time, a reversible bond-energy neutral reaction can occur that inserts the gate tile in the central location, enabling reversible signals y and z to propagate on the output wires. **(D)** An unreliable irreversible wire cross that has two possible types of warnings.

wire will be embedded within a larger assembly. Each tile displacement reaction is neutral with respect to the bond energy, so when both the wire and signal tiles are at the same concentration, every tile displacement occurs at the same rate, and the signal transmission performs an unbiased random walk. Thus the expected time for signal transmission along a wire of length N is $O(N^2)$.

Faster signal transmission is possible if each tile displacement step is irreversible, which can be accomplished if new toehold bonds are formed such that the bond energy change is downhill. Shown in figure 3B, the wire is as before, but now the signal tile has an additional toehold. Thus, tile displacement reactions are energetically downhill, forming one net additional bond with each reaction step, and the expected time for signal transmission is now $O(N)$.

When a horizontal and a vertical wire meet, we can perform a computational step. Figure 3C shows two reversible wires, one carrying signal x and the other carrying signal w , meeting at a “wire cross” tile in the center. At this location, reversible tile displacement by a “gate tile” can effect the $w + x \rightleftharpoons y + z$ reaction. Because the initial wire cross tile has four closed toeholds, tile displacement by the gate tile must form all four new toehold bonds, and thus tile displacement here prior to arrival of both the x and w signals would be energetically unfavorable and would not occur. This gate design is robust and flexible: it is straightforward to design more powerful variants. For example, the horizontal wire can carry one of two signals, 0 or 1, the vertical wire also can carry 0 or 1, and there are now four gate tiles, one for each input combination, with output signals that effectively compute the logic function of interest. Specifically, to compute NAND and output using the same signal varieties, we would use four gate tiles that replace (w, x, y, z) respectively by $(0, 0, 1, 1)$, $(0, 1, 1, 1)$, $(1, 0, 1, 1)$, and $(1, 1, 0, 0)$.

Can we similarly perform logic gate operations using irreversible wires, thus making computation faster? Unfortunately, the above schemes no longer work in this case, as illustrated in figure 3D. The problem is that now, prior to arrival of the second signal, an energetically neutral tile displacement is possible at the gate position that simply ignores the missing input wire. Indeed, if the vertical wire is meant to be capable of carrying two signals (here w or y), then an energetically neutral tile displacement could analogously flip the signal content. Thus, this tile set and gate design is deemed unreliable and would issue warnings in our simulator. The lesson is that irreversible reactions will form an additional toehold when operating as intended, and this presents the possibility that a similar context that differs by just one open toehold, where the reaction is not intended to occur, will be energetically neutral and lead to an error.

So does this mean that linear-time binary signal transmission and circuit computation is impossible with tile displacement systems? Thankfully, no. The trick is that while the leading wavefront of signal propagation and

computation still must be reversible, in order to reliably discriminate single-toehold differences, it can be safe to irreversibly latch a decision in a context where all neighboring tile contain the same information, so differences between a 0 signal and a 1 signal by necessity involve two toeholds. Now, if the irreversible tile displacement involves the formation of one extra toehold in the correct context, in the incorrect context it would have to ignore two toeholds and thus would be uphill. This principle is illustrated in the design shown in figure 4A, where the latch tile can irreversibly insert itself into a three-tile-long segment of signal-carrying wire. Exactly where the latch tile inserts does not matter; the signal ratchets forward either way.

We are now ready to take these designs for wires and gates, and combine them to construct feedforward logic circuits that compute in time linear with the depth of the circuit (as laid out in an array). There is, however, one more problem to solve if we want to build circuits that utilize multiple types of logic gates (e.g. XOR, AND, OR, NAND, NOR, WIRECROSS, and others). When an invading gate tile (e.g. “XOR 10” shown in figure 4B) displaces the initial gate tile (e.g. “XOR”), it makes bonds with toeholds on the neighboring four tiles but not with the displaced tile itself – therefore, information about which function should be computed must be contained in the neighboring tiles, and not just the gate tiles. We achieve this goal by using a gate-specific toehold in the initial gate tile, which directs the incorporation of a translator tile in the final position of each wire, as shown in figure 4B. Now the translator tile contains information about which logic function should be computed. Thus, an arbitrary number of gate types may coexist in the same system.

Simulations of two feedforward circuit computations are shown in figure 5. Using just a systolic array of XOR gates, a collection of parity outputs (involving different subsets of the inputs) are produced, incidentally creating a Sierpinski triangle pattern within the completed wires and gates. The second circuit makes use of four gate types: a NOR gate that produces the same signal on both output wires, a NAND/XOR gate that produces NAND to the north and XOR to the east, a WIRECROSS that sends its south input to the north and its west input to the east, and a WIREPASS that sends its south input to the east and its west input to the north. The positions and identities of the gates are laid out in the initial tile array. Computation of an $N \times N$ circuit will take expected time $O(N)$.

The tile system with eastward and northward latching binary wires and the five types of gate functions discussed above consists of 66 tile types altogether. For a circuit that can be laid out effectively in this format, an area of $O(N^2)$ tiles can support N^2 gates. Arbitrary feedforward circuits with N gates can be implemented in $O(N^2)$ area using a standard crossbar array architecture (for example see [34]) and a new FANOUT gate that copies one

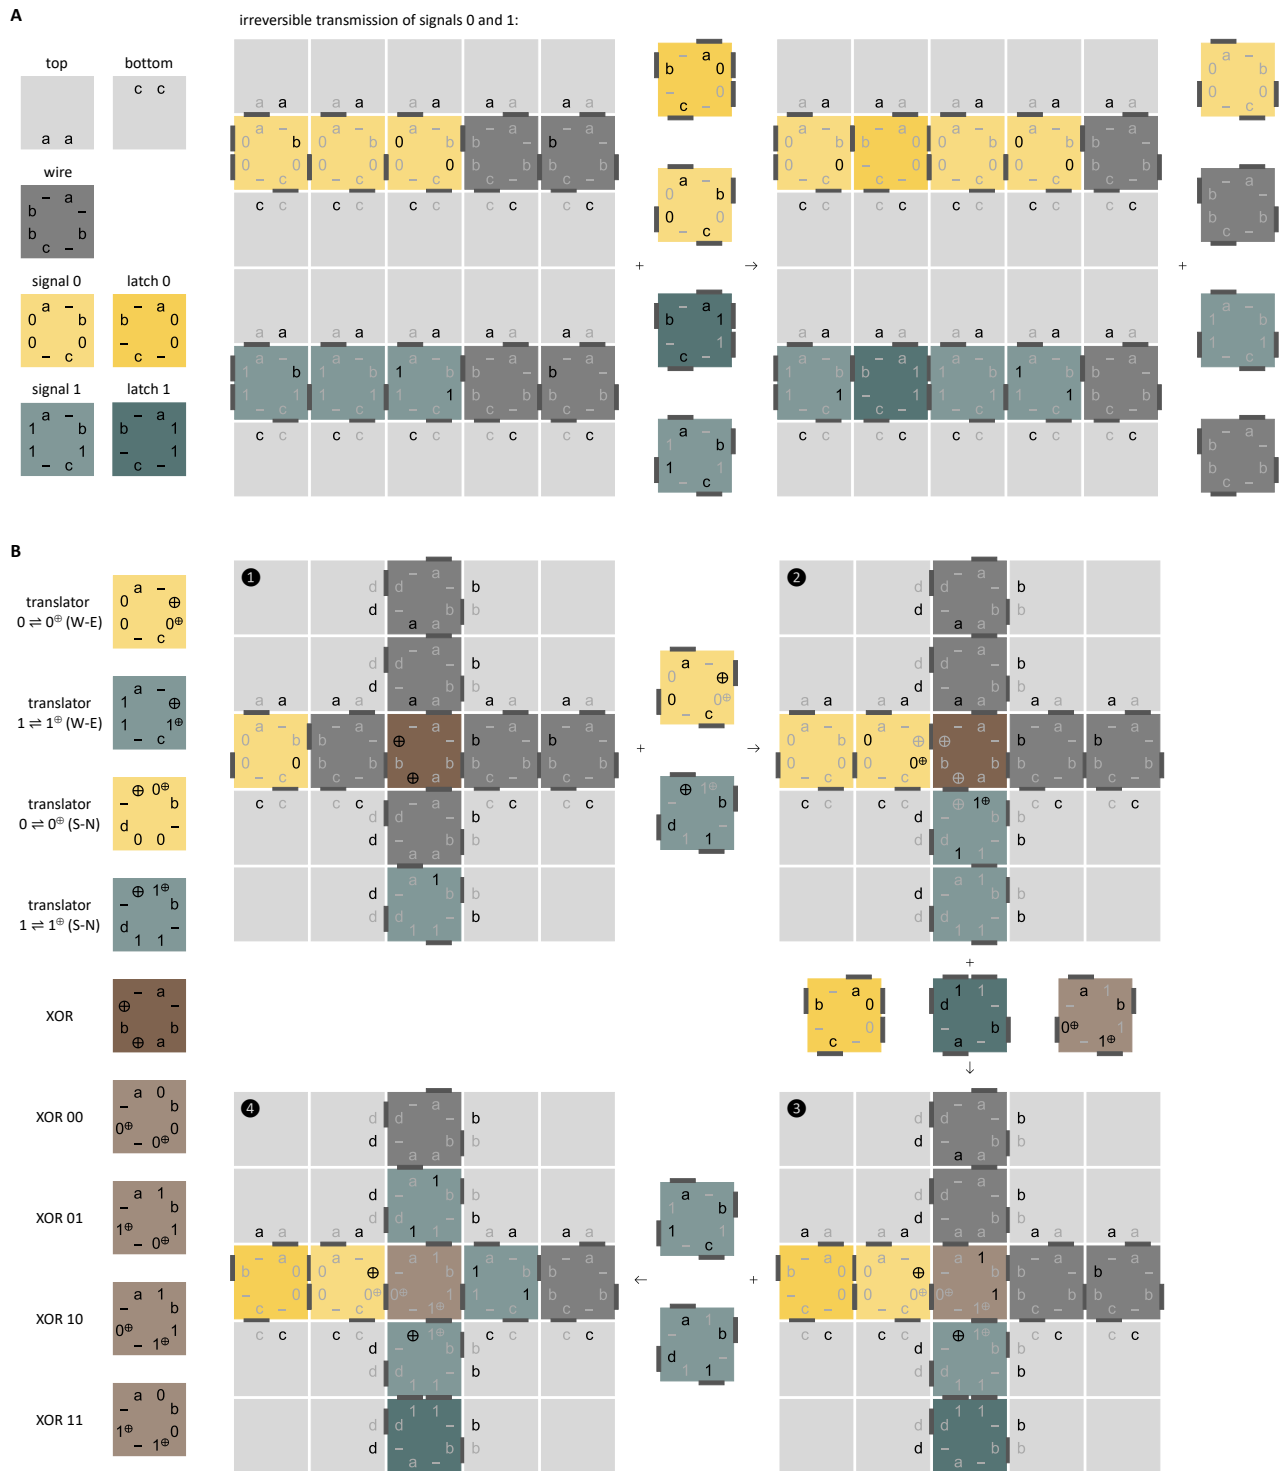


Figure 4: **(A)** Irreversible wire with no warnings. The next effect of two separate, independent tile displacement steps is shown for each wire. Importantly, no unreliable tile displacement reactions are possible. **(B)** An XOR gate. The gate itself is entirely reversible; latch steps in the input and output wires are sufficient for ensuring net progress within a circuit.

input and ignores the other (using 75 tile types if the new gate is just added, or 57 tile types if the redundant NAND/XOR and XOR gates are removed).

Can we do better than just feedforward circuits? It is clear from inspection that our latching binary-signal wires

can transmit information in either direction, depending on where the signal first arrives, and it is straightforward to implement gates that receive inputs from any two sides and produce outputs on the two other sides, so we can arrange for signals to go around in cycles. Furthermore, the

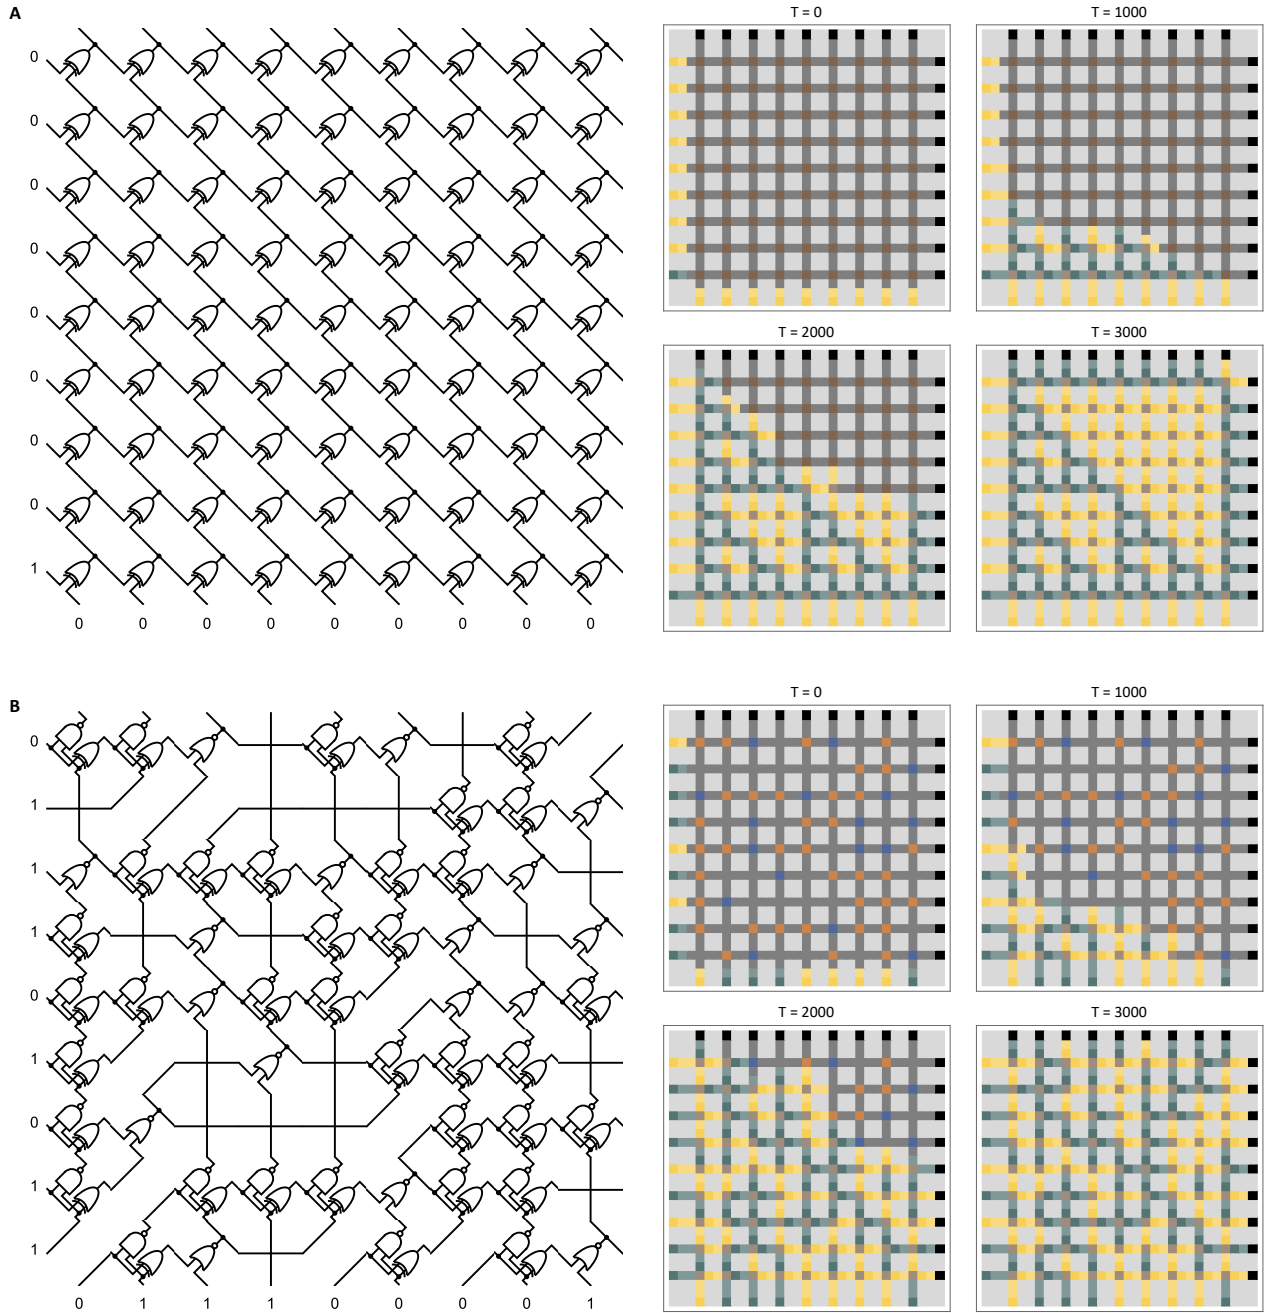


Figure 5: **(A)** A 9 by 9 array of XOR gates. Black tiles are “caps” that terminate the output wires. **(B)** A 18-input 18-output logic circuit composed of XOR, NAND, and NOR gates. Two types of wire routing are implemented with a WIRECROSS tile that sends its south input to the north and its west input to the east, and a WIREPASS tile that sends its south input to the east and its west input to the north. The simulation time T measures the number of tile displacement reactions that have occurred, rather than the real time in the Gillespie simulation.

tile displacement model in principle allows displacement to occur an arbitrary number of times in a given location. As a trivial example, the reversible wire of figure 3A can endlessly perform a random walk, back and forth forever. This raises the prospect of a tile displacement system of size N simulating a recurrent (iterated, feedback) circuit of size $O(N)$, which can perform computations of length $2^{O(N)}$ – exponentially more than what a feedforward cir-

cuit of the same size can do. This is to say, with respect to the size of the initial array and tile set, our existing construction can solve $PTIME$ problems, while a reach goal would be to solve $PSPACE$ problems like recurrent circuits can. Unfortunately, this is not compatible with the use of latching wires to ensure linear-time signal propagation: an area- N tile array initially has at most $O(N)$ open toeholds, and thus at most $O(N)$ irreversible

tile displacement steps can take place before the system comes to a standstill – or more precisely, until it must henceforth rely exclusively on reversible steps.

4. 1D CA SPACE-TIME HISTORIES

Boiling down what we learned about circuits to its reversible essence, we can re-implement the above computations using fewer tile types, more compact layouts with just one tile per logic gate, and power for driving the computation forward coming from concentration differences rather than from irreversible toehold formation.

We start by providing generalized construction for simulating the space-time history of one-dimensional block cellular automata (1D BCA) that is very similar to their simulation by algorithmic self-assembly of DNA tiles [35, 2]. The instantaneous state of a 1D BCA is just a one-dimensional array of symbols from a given alphabet \mathcal{A} , and in each time step the entire array is synchronously updated by applying a rule $(x, y) \rightarrow (f(x, y), g(x, y))$ to a partitioning of the array into pairs, where f and g are functions that define the BCA and the parity of the partition alternates on each time step. The size of the array may be infinite, finite, or expanding, with given initial state and boundary conditions (typically a finite core then periodic). Our tile displacement system construction, shown in figure 6A, makes use of $2 + 2N + N^2$ tile types for a 1D BCA with an alphabet of size N . The initial array uses 1 tile in the lower-left corner, N tiles to define input boundary conditions to be fed in at each time step from the left, N tiles to define the input boundary conditions to be fed in at each time step from the bottom, and 1 tile type filling in the remaining “blank” uncomputed region of the array. The remaining N^2 tile types encode every input/output case for the update rule. For example, a binary alphabet ($N = 2$) will result in 10 tile types (figure 6B). The n^{th} synchronous update of the 1D BCA will be encoded in the n^{th} diagonal of the tile array. Similar to the gate tiles in the circuit construction, displacement must match all four open toehold positions, else it will be energetically unfavorable. This can only happen when both the tile to the left and the tile below have already updated, thus ensuring that the computed information is based on the correct information from the preceding diagonal.

Because our model insists that any tile that can be displaced in a simulation must have a non-zero concentration as a monomer in solution, every reaction will be reversible. However, by chemostatting the blank tile at a lower concentration than the rule tiles, each displacement reaction can be biased forward by some factor $r = c_{\text{rule}}/c_{\text{blank}}$. From detailed balance of the CRN and CTMC, this ensures that the equilibrium probability of the rule-tile containing assembly is r times higher than that of the blank-tile containing assembly. Although the system will never get irreversibly locked into a final output assembly state, the complete assembly with all rule tiles in place will be r^m times more likely than an assem-

bly with m blank tiles still present, which we consider “good enough”. Note that if a final irreversible step is desired to lock in place the completed computation, this is also possible by adapting the techniques used in the circuit construction, just at the upper right corner.

Comparing the circuit construction of figure 5A to the cellular automaton space-time history construction in figure 6BC, both of which compute parallel systolic arrays of XOR gates, we see that for the same size array, the cellular automaton approach computes roughly 9 times more gates. It also uses just 10 tile types, compared to 30 for the circuit construction (if the tiles used for logic gates other than XOR are omitted).

However, our cellular automaton construction, by its very nature as a cellular automaton, receives information only in the initial 1D boundary conditions, and thus an assembly cannot specify a two-dimensional layout for the circuit that will be computed by tile displacement. A simple modification of the ideas resolves this apparent limitation: we generalize the construction to *cellular automaton transformers* whose cell update now depends both on the current state $(x, y) \in \mathcal{A}_1$ and a time-and-space-dependent input pattern $(p, q) \in \mathcal{A}_0$, as shown in figure 6D. Instead of an initial array containing uniform blank tiles, the initial array will contain a layout of “pattern” tiles that each encode the information p that the gate below it will need to read, as well as the information q that the tile to its left will need to read. If \mathcal{A}_1 is size N and \mathcal{A}_0 is size M , then there are $2N$ input tiles, M^2 pattern tiles, and N^2M^2 rule tiles. Each reversible tile displacement reaction now must match *four* variable pieces of information, in two pattern toeholds and two state toeholds. As shown in figure 6EF, laying out exactly the same circuit as in figure 5B now requires 9 times less space, uses just 39 tile types ($N = 2$ state bit values plus a terminator, $M = 5$ logic functions, but not all combinations are needed) instead of 57, and, with concentration bias again, computes significantly faster.

Both these constructions exhibit strong similarities to computation via algorithmic growth during self-assembly of tiles – in the first case, 2D tiles growing a 2D structure from a 1D boundary [35, 2], and in the second case, 3D tiles growing an additional layer on top of a patterned 2D initial assembly [36]. A significant difference is that rather than growing in size, the tile displacement system always remains the same size; rather than each tile attachment requiring new bond energy to counteract the lost entropy due to localization of the tile, the tile displacement system remains neutral with respect to bond energy because each incoming tile is balanced by an outgoing tile. Thus, rather than finding suitable operating conditions by balancing temperature (controlling the bond energies) against tile concentrations (which simultaneously affect the kinetics), in tile displacement we balance concentration against concentration (which permits similar bias at different speeds and temperatures). These benefits reflect similar observations about the increased robust-




Figure 6: **(A)** General case implementation of 1D block cellular automaton. Here a and b , written in roman font, denote specific toeholds. In contrast, x , y , f and g are variables and thus shown in italics. There will be a separate rule tile for each possible pair $x, y \in \mathcal{A}$, with f and g being dependent on x and y , and similarly for the input tiles. **(B)** An example 1D block cellular automaton that computes the same function as the circuit shown in Fig. 5A. **(C)** Simulation snapshots. **(D)** General case implementation of 1D block cellular automaton transformer. **(E)** An example 1D block cellular automaton transformer that computes the same function as the circuit shown in Fig. 5B. **(F)** Simulation snapshots.

ness of strand displacement and toehold exchange compared to direct hybridization of complementary oligonucleotides [37, 38]. Seen more generally, tile displacement systems involve reconfiguration of a constant-sized assembly via local propagation of information, which is reminiscent of the distinction between crystal growth from monomers in dilute solution (the case generally assumed in algorithmic self-assembly of DNA tiles) versus crystallization from the melt (wherein the initial state is a disorganized constant-density liquid of monomers, within which crystalline order locally propagates during crystal growth).

Have we identified new concepts for tile displacement systems that allow us to perform more computation in a limited space? Powering computation forward via concentration bias in reversible reactions has given rise to compact constructions that naturally avoid the unreliability warnings that plagued our initial wire and circuit constructions, but the computational power still remains PTIME. One way of looking at this is that the free energy of the assembly, $G(A)$ decreases every time a higher-concentration tile replaces a lower-concentration tile, yet the minimum (most favorable) free energy occurs if all tiles in the array are highest-concentration tiles. That is to say, the free energy is bounded below, and if each forward computational step is biased by a minimum amount, there are a bounded number of such steps that can occur before the computation is done. The situation is not so different from the limitation we encountered when powering computation by new toehold formation in irreversible displacement steps. Is this limitation to PTIME a feature of tile displacement systems in general, or is it particular to the lack of imagination in the constructions we have presented so far?

5. 2D CA IN-PLACE EXECUTION

We can get some ideas from the notion of a cellular automaton transformer, which reads a 2D pattern as a wave of activity passes over it, leaving a new pattern in its wake. Suppose that the new pattern can be read by a second wave, corresponding to a second cellular automaton transformer using a new set of rule tiles. For example, the initial pattern might use toehold alphabet \mathcal{A}_0 , the first cellular automaton transformer uses states in alphabet \mathcal{A}_1 and writes a new pattern using alphabet \mathcal{A}_2 by utilizing the two locations that, in figure 6D, have useless inert “-” toeholds. Then, the second cellular automaton transformer can read \mathcal{A}_2 , store its transient state in \mathcal{A}_3 , and write a third pattern using \mathcal{A}_4 . To drive the computation forward, the first transformer’s rule tiles should have a higher concentration than the pattern tiles, and the second transformer’s rule tiles should have a higher concentration than the first transformer’s rule tiles. This idea could be extended to K waves, each with its own set of rule tiles. This would improve upon the previous constructions, in which each location in the array experiences just net one forward tile displacement step – at

that location, either one has the initial tile, or the final tile. Whereas, in an implementation of a multiple-wave cellular automaton transformer, each location would go through a sequence of changes, one for each wave. In a sense, we achieve K -fold more computation within the same assembly area. This is somewhat analogous to freezing cellular automata, which are restricted to change a cell’s state a limited number of times [39].

There are two problems here, as you have probably already noticed. First, if the concentration ratio from wave to wave is r , then a K -wave computation requires a ratio of r^K between the lowest-concentration tiles and the highest-concentration tiles. That quickly becomes impractical, and theoretically unappealing. Second, each wave requires a new set of tiles – yet for PSPACE computations we would require an exponential number of tile updates and thus a comparable number of waves. So this idea doesn’t get us where we want to go.

To keep a constant number of tile types while allowing an unbounded number of tile displacement steps per site, perhaps we could have a small number K of waves, but have wave K output its new pattern using alphabet \mathcal{A}_0 so that the tiles of wave 1 can read it – thus allowing iterated computation, such as binary counters and perhaps universal space-bounded algorithms. This is indeed the essence of the construction we’ll arrive at, but it comes at a cost: for wave 1 tiles to displace wave K tiles, they cannot be at a lower concentration, which basically implies that all rule tiles must be at the same concentration, and we have no concentration bias pushing the computation forward. (This conclusion is not specific to periodic waves of cellular automata transformers; it follows in general that if we want to implement a computation that may update a given site an unknown and unbounded number of times, then every tile type may at some point be an incoming tile and at other times be the outgoing tile, so the concentrations of all rule tiles must be equal.) If we have already accepted that our designs should exclusively use bond-energy neutral tile displacement, then in fact the bond energy and free energy of our assembly will remain constant over time – we are truly dealing with reversible computation. Thankfully, reversible computation is by no means impossible [14, 40].

Our approach will be to exhibit a surprisingly natural correspondence between certain tile displacement systems and the well-studied class of two-dimensional block cellular automata (2D BCA) that arose in the study of reversible computation by discrete models of ballistic physical dynamics [20, 21, 22]. The 2D BCA model is a natural generalization of the 1D BCA discussed above: rather than partitioning a 1D array into pairs of cells that get synchronously rewritten with alternating partition parity on alternate time steps, we now partition a 2D array into 2×2 blocks of cells that get synchronously rewritten with alternating partition parity on alternating time steps (compare figure 7A with figure 7E). The formalism

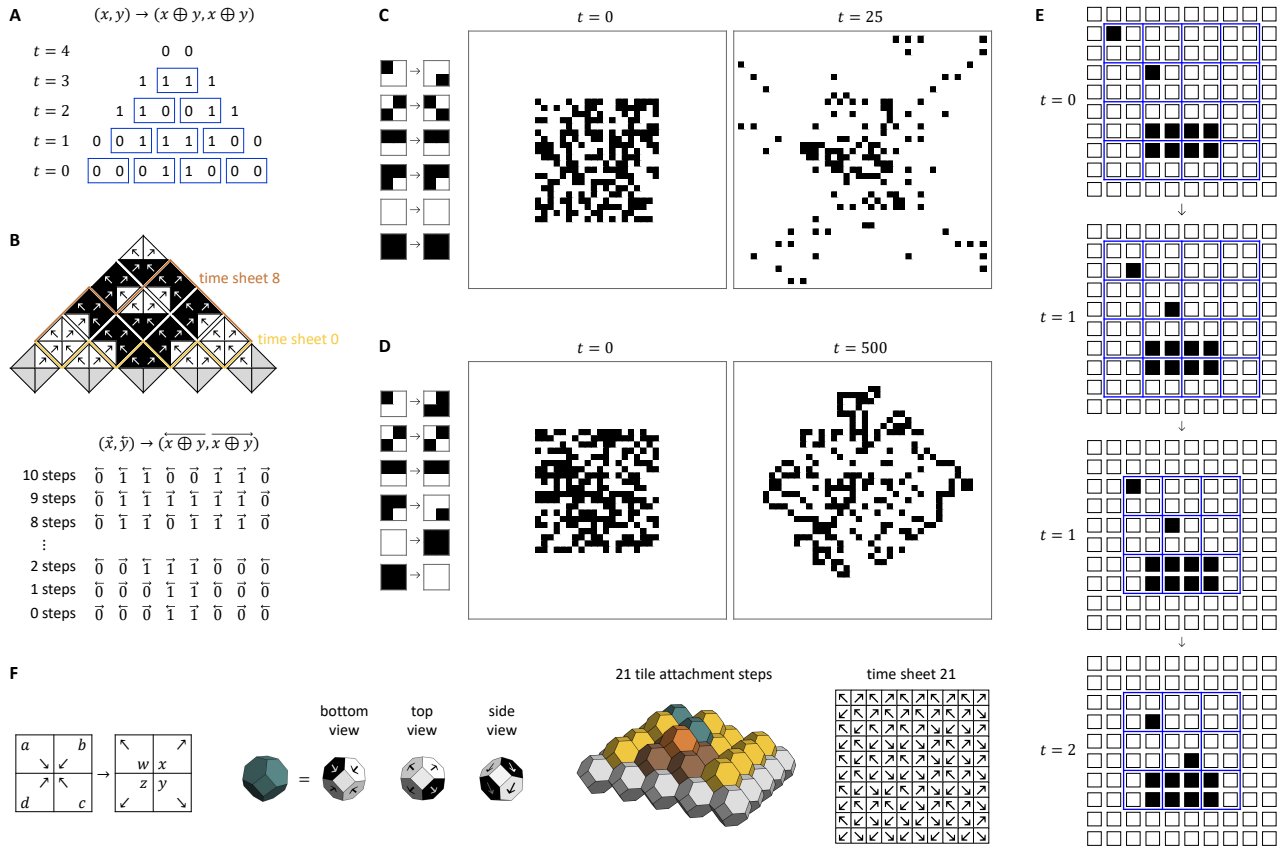


Figure 7: (A) Execution of a synchronous 1D block cellular automaton. (B) Asynchronous 2D tile self-assembly that simulates the computation in (A). (C) and (D) Simulations of two example 2D block cellular automata: Billiard Ball Model (C) and Critters (D). (E) Execution of a synchronous 2D block cellular automaton. (F) Asynchronous 3D tile self-assembly that simulates the computation in (E).

allows the rewrite rules to be arbitrary functions

$$f\left(\begin{bmatrix} a & b \\ d & c \end{bmatrix}\right) = \begin{bmatrix} w & x \\ z & y \end{bmatrix}$$

but if the rewrite function is a bijection, then the 2D BCA is logically reversible in the sense that iterating with f^{-1} instead of with f will bring the simulation backwards in time. The most famous 2D BCA rule, the Billiard Ball Model (BBM), is logically reversible, rotationally and mirror symmetric, conserves the total number of 1s, can directly simulate reversible circuits, and with an infinite periodic initial state can simulate universal Turing machines [20]. Example simulations of two binary-state 2D BCA, the BBM and “Critters”, are shown in figure 7CD. With larger alphabets, 2D BCA can simulate arbitrary classical cellular automata and Turing machines, either of the irreversible or reversible variety. (Generalizations to using blocks larger than 2×2 is also natural, but will not be considered here.)

There are three obstacles to implementing arbitrary 2D BCA as tile displacement systems, and we will solve them all. The first is that tile displacement reactions are asynchronous (occurring at random locations and in random orders) while 2D BCA require synchronous updates of the

entire array (and fail utterly if the same update function is applied asynchronously with no other modifications). The second is that the mechanics of tile displacement must be designed to avoid irreversible steps that close too many toeholds at once. And the third obstacle is that with exclusively reversible reactions and no concentration bias, there must be some other way to drive the system forward if we don’t want to wait forever.

For the first challenge, we adapt prior methods for imbuing asynchronous cellular automata with locally synchronizing mechanisms [41, 42, 43]. The specific approach used here generalizes the approach used for simulation of 1D cellular block automata space-time histories in the previous section. Figure 7A gives an example of a 1D BCA, with boxes highlighting the partitioning into pairs with alternate parity on each synchronous time step. Figure 7B shows the same computation interpreted as 2D tile self-assembly where, starting from the 5 tiles at the bottom that encode the 8 input bits as well as their partitioning, rule tiles attach whenever they can match two sides of existing tiles in the assembly, thus asynchronously growing the space-time history. We have augmented the tiles with arrows that point to where incoming tiles could attach; thus, in the initial assembly of 5 tiles, the sites

where tile can attach are exactly those locations where arrows are pointing inward toward the incoming tile. A cut through the assembly’s space-time diagram corresponds a particular moment during the asynchronous self-assembly process – we show a cut after 0 tile additions (yellow) and another after 8 tile additions (orange). We call these “time sheets” because at different horizontal (x) positions, they are at different heights (t), and thus reading out the binary (black/white) states along a time sheet path correspond to states at different time steps of the underlying synchronous cellular automaton. Nonetheless, the time-sheet state information, augmented with the relevant arrows, is all that is needed to correctly complete the computation using an asynchronous update rule that executes only when arrows point toward each other, otherwise leaving the cells untouched. This process exactly mimics the self-assembly of the deterministic space-time history, despite its non-deterministic order of execution.

There is an exactly analogous arrow-augmented asynchronous update rule for 2D BCA. Rather than square tiles, we now have truncated octahedra as “tiles”, but the self-assembling structure is again a space-time history of the correct synchronous cellular automaton computation. Tiles may attach when they match four hexagonal faces of existing tiles in the assembly. (The small square faces are inert.) Again, if we imagine arrows orthogonal to the hexagonal faces of tiles, pointing out of the tile, then valid sites for attachment of a new tile correspond exactly to situations where all four arrows on the matching faces are pointing toward each other. The growth front for a give stage of assembly again corresponds to a (now two dimensional) time sheet, and we can write out the states of each exposed hexagonal face in a two dimensional array along with the orientation of its corresponding arrow. It is now a simple observation that the asynchronous addition of a tile corresponds exactly to an asynchronous update of a 2×2 block with four inward-pointing arrows, resulting in updates of the four cells and reversing all four arrows. Another way of thinking of it is that after a block asynchronously updates, it will not be able to update again until all four overlapping 2×2 blocks have first updated and flipped the arrows back. Thus, the arrow-augmented asynchronous updating corresponds exactly to synchronous parallel updating with alternating-parity partitioning into blocks.

The second challenge is to implement this type of asynchronous block cellular automaton updating rule using tile displacement. Our construction, shown in figure 8, introduces additional complexities due to the fact that all tile displacement reactions are physically reversible, even if the 2D BCA logic update rules are irreversible, combined with the need to ensure that when one toehold is closed, the neighboring toehold must be opened – thus we must be able to guarantee a mismatch. Tripling the cell state alphabet by adding α , β , γ markers solves both problems. For each 2D BCA update case, we make three tiles, one inputting α -symbols and outputting β -

symbols, another inputting β -symbols and outputting γ -symbols, and the third inputting γ -symbols and outputting α -symbols. When an $\alpha \rightarrow \beta$ tile inserts into the array, that simulates a forward-time asynchronous update. The swapping of which toehold is open and which is closed reflects the flipping orientation of arrows in the asynchronous cellular automaton; we can read the arrows from a tile array by looking at the open toeholds and drawing the arrow from α to β , from β to γ , or from γ to α . Boundary conditions for finite arrays must also be handled, using the same principles.

Each side of a tile encodes the state of a specific cell in the 2D BCA (at a particular time mod 3, as per α - β - γ of the open toehold), and thus the grid of simulated BCA cells is oriented at a 45° angle relative to the array of tiles. State being encoded on the sides of tiles also facilitates that each tile displacement step corresponds to an update of a whole 2×2 block, and the fact that two tiles share the same side location reflects that each cell in a 2D BCA can be updated either by an odd-parity block or an even-parity block.

The final challenge concerns how to drive the computation forward. Let us first consider reversible 2D BCA rules. In this case, after any forward tile displacement step, there is exactly one monomer tile type that can reverse the reaction: the tile that was just displaced. What this means is that the full state space of the tile displacement system’s CTMC is essentially linear; though fat and fuzzy, it has the same thickness both arbitrarily far into the future and arbitrarily far into the past. The thickness has to do with all the possible contours of the time sheet for a given average time. Thus we can say that the state space of the tile displacement system consists exclusively of correct reachable states of the computation; for a reversible 2D BCA simulating a compact recurrent circuit for solving a PSPACE problem, the tile displacement systems’s state space will also be exponentially long and will reach the same correct conclusion. Stochastic Gillespie simulation of the tile displacement CRN will result in an unbiased random walk back and forth along this fuzzy-linear state space. (Every assembly in this reachable state space has the same energy.) However, unlike a standard reversible Turing machine with Brownian dynamics [14], whose state spaces is strictly a linear graph so the expected random walk hitting time for reaching the end of an T step computation is $O(T^2)$, the time sheet diffuses much more slowly. As a rough estimate for an $N \times N$ tile array that requires N^2 forward updates to move the time sheet 1 net synchronous update step into the future under ideal circumstances, the same N^2 updates if half forward and half backward will be expected to net move the time sheet N steps either forward or backward, which corresponds to just $1/N$ equivalent synchronous update steps. This being just a polynomial inefficiency, perhaps we should not be too concerned.

More interesting is what happens if the 2D BCA rules are irreversible. This means there are multiple cases for




Figure 8: **(A)** General case implementation of 2D block cellular automaton. **(B)** Example updates in the Billiard Ball Model.

the 2×2 block input that map to the same output. Therefore the state space for the tile displacement system will be exponentially branched in the backwards-in-time direction (as pictured by Bennett in figure 10 of his review paper [14], but thicker and fuzzier). Consequently Brownian dynamics will tend to be entropically biased toward where there are more states, and the system will

run backwards. Can this entropic driving force be used to encourage a system to perform a desired computation by designing a system whose *reverse* dynamics are what we want? Attempting to do so would be risky, and probably futile, because the 2D block update rule being irreversible means that there are some states that have no local predecessor, and backward progress will get stuck as such

local configurations are encountered.

A better way to exploit an entropic driving force is to have nondeterministic, stochastic forward update rules added to an otherwise-reversible system. For example, consider a tile displacement simulation of the BBM model, with boundary tiles designed to implement a reflecting boundary (as they must for the system to remain reversible). If we design a special boundary tile that can either reflect a ball or (in the forward direction) produce a new ball out of nothing, then we obtain a new system that is still entirely reversible in the sense that there exists (at least one) possible applicable block update in all circumstances, so the system cannot get stuck either in the forward direction or the reverse direction. With uniform tile concentrations, all assembly states will still be isoenergetic. But started with an empty $N \times N$ array, forward updates of the special tile will about half the time produce a new ball, which will entropically drive the system to a density such that forward production of balls is balanced by the reverse reaction, the absorption of balls into the special tile. At this point, which will be $O(N^2)$ synchronous time steps in the future, the time sheet will stop advancing on average. Another way of looking at it is that with all reactions being neutral, equilibrium will reflect equipartition among all reachable states, and the combinatorially greatest number of states will have a number of ball near the optimal density – so, that’s what we are likely to observe. And the only way to get there is to run the time sheet forward enough to emit that number of balls.

A gas-filled BBM simulation is not of great use by itself, but we can make use of it by also placing a circuit in the array, and drawing a 2-cell-thick wall around it. In the BBM model, balls bounce off walls, and walls are stable. Thus, despite random stochastic gas entering the areas of the array outside the box, the circuit will remain perfectly isolated from the gas. But due to the time-sheet coupling enforced by the asynchronous arrow rules, the time sheet that is being driven forward by the expansion of the gas will simultaneously drive the circuit forward.

Unfortunately, for an array of area N^2 , we will only drive the computation forward by $O(N^2)$ steps – this is no better than the PTIME computational power of the original circuit construction. Essentially, in a small confined space, our circuit “heats up” and stops working. To run it for a long time, we need a larger space into which we can release the simulated heat. For example, if we are willing to entertain a half-infinite-plane array for tile displacement, we can draw a BBM wall down the middle, release gas on one side, and let the other side simulate an interesting recurrent circuit. Now, although the array is infinite (or very large) in direct proportion to how much computation we want to do, we can say that we have confined the interesting part of the computation – the circuit itself – to a very small area relative to the potentially exponentially long computation. This isn’t PSPACE computation in terms of the size of the

array, but rather in terms of the size of the part of the array that we care about. Similar constructions can be used to drive forward computation not just for other reversible 2D BCA rules, but even for irreversible rules: the cellular automaton alphabet can be expanded to encode an inert “wall”, time sheets within the walled region and outside of it remain coupled, and sufficient entropy must be generated by stochastic rules outside the wall, to be dissipated into a sufficiently larger area.

6. DISCUSSION

Tile displacement within arrays of square DNA origami tiles was discovered accidentally [3]. While some aspects of the formal model, such as the four-sided generalization of toehold exchange, were invented for mathematical elegance rather than detailed realism, they are not too far flung from what has been experimentally demonstrated and characterized. So it is quite delightful that within the design space for tile displacement systems, we find natural implementations for feedforward circuits and one-dimensional cellular automata that compute in linear time, powered by irreversible toehold formation or by concentration gradients. Even more delightful is that attempts to squeeze out more computational power per area seemed almost inevitably to lead us consider physical constraints such as energy, reversibility, and asynchrony – which in turn lead to classical two-dimensional cellular automata models that arose in early studies of the physics of computation [14, 20]. Our strongest result (despite weak time efficiency) is that a tile displacement array of size N can reversibly simulate a recurrent reversible circuit (via the Billiard Ball Model cellular automaton) for an arbitrary number of steps. In other words, the reachability question for tile displacement is PSPACE complete – a result strongly reminiscent of Thachuk & Condon’s beautiful PSPACE-hardness result for CRNs and DSDs [44].

We believe that there remains a lot undiscovered within the tile displacement design space. For example, while our constructions showed that the asynchronous tile displacement model can simulate synchronous cellular automata, the needed flipping-arrow mechanism for local synchronization seems almost built-in to the tile displacement model in the form of open and closed toeholds for toehold exchange, and it’s not obvious how to directly simulate asynchronous cellular automaton models such as reversible surface CRNs [45, 34, 43]. We might also ask whether using information within the branch migration domains rather than just in toeholds – or whether having even more toeholds and branch migration domains on a tile’s sides – could have advantages either theoretically or experimentally. Further, the most interesting systems demonstrated experimentally in the initial work on tile displacement [3] involved systems of interacting multi-tile arrays, rather than a single array and a monomer. How do our single-assembly results fit into that larger picture? Finally, might tile displacement systems be combined with

other molecular mechanisms to solve our problems driving the computation forward – for example, an oscillator [46] that periodically activates and deactivates the α , β , and γ monomer tiles in sequence.

When first discovered, the tile displacement mechanism seemed most closely related to strand displacement mechanisms, only two dimensional. However, as we investigated the capabilities of single-assembly tile displacement, many parallels to passive tile assembly [2] became prominent. Tile displacement systems appear to combine the principles of DNA strand displacement and self-assembly in different ways than hairpin-based programmable self-assembly [47], signal-passing tile self-assembly [48, 49], CRN-controlled tile assembly [50, 51], and other models we are aware of. Comparing the benefits, drawbacks, and relationships between these models may help uncover a more unified way of thinking about programmable molecular systems.

And even if tile displacement systems, as explored theoretically here, never become useful experimentally, we hope that it was interesting and perhaps inspiring to look long and deep at a simple mechanism until intricate patterns emerge.

ACKNOWLEDGEMENTS

The authors thank William Poole and Ho-Lin Chen for useful discussions. This work was partially supported by NSF awards 2008589 and 1813550.

- [1] John E Savage. *Models of computation*, volume 136. Addison-Wesley Reading, MA, 1998.
- [2] David Doty. Theory of algorithmic self-assembly. *Communications of the ACM*, 55(12):78–88, 2012.
- [3] Philip Petersen, Grigory Tikhomirov, and Lulu Qian. Information-based autonomous reconfiguration in systems of interacting DNA nanostructures. *Nature Communications*, 9(1):1–10, 2018.
- [4] Bernard Yurke, Andrew J Turberfield, Allen P Mills, Friedrich C Simmel, and Jennifer L Neumann. A DNA-fuelled molecular machine made of DNA. *Nature*, 406(6796):605–608, 2000.
- [5] Nadrian C Seeman. *Structural DNA nanotechnology*. Cambridge University Press, 2015.
- [6] David Yu Zhang and Georg Seelig. Dynamic DNA nanotechnology using strand-displacement reactions. *Nature Chemistry*, 3(2):103–113, 2011.
- [7] Andrew Phillips and Luca Cardelli. A programming language for composable DNA circuits. *Journal of the Royal Society Interface*, 6(suppl.4):S419–S436, 2009.
- [8] Stefan Badelt, Casey Grun, Karthik V Sarma, Brian Wolfe, Seung Woo Shin, and Erik Winfree. A domain-level DNA strand displacement reaction enumerator allowing arbitrary non-pseudoknotted secondary structures. *Journal of the Royal Society Interface*, 17(167):20190866, 2020.
- [9] 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.
- [10] Luca Cardelli. Two-domain DNA strand displacement. *Mathematical Structures in Computer Science*, 23(2):247–271, 2013.
- [11] Lulu Qian, David Soloveichik, and Erik Winfree. Efficient Turing-universal computation with DNA polymers. In *DNA Computing and Molecular Programming (LNCS volume 6518)*, pages 123–140. Springer, 2010.
- [12] Matthew R Lakin and Andrew Phillips. Modelling, simulating and verifying Turing-powerful strand displacement systems. In *DNA Computing and Molecular Programming (LNCS volume 6937)*, pages 130–144. Springer, 2011.
- [13] Wataru Yahiro and Masami Hagiya. Implementation of Turing machine using DNA strand displacement. In *Theory and Practice of Natural Computing (LNCS volume 10071)*, pages 161–172. Springer, 2016.
- [14] Charles H Bennett. The thermodynamics of computation—a review. *International Journal of Theoretical Physics*, 21(12):905–940, 1982.
- [15] Paul WK Rothmund. Folding DNA to create nanoscale shapes and patterns. *Nature*, 440(7082):297–302, 2006.
- [16] Wenyan Liu, Hong Zhong, Risheng Wang, and Nadrian C Seeman. Crystalline two-dimensional DNA-origami arrays. *Angewandte Chemie International Edition*, 50(1):264–267, 2011.
- [17] Grigory Tikhomirov, Philip Petersen, and Lulu Qian. Fractal assembly of micrometre-scale DNA origami arrays with arbitrary patterns. *Nature*, 552(7683):67–71, 2017.
- [18] Sungwook Woo and Paul WK Rothmund. Programmable molecular recognition based on the geometry of DNA nanostructures. *Nature Chemistry*, 3(8):620–627, 2011.
- [19] Lulu Qian and Erik Winfree. A simple DNA gate motif for synthesizing large-scale circuits. *Journal of the Royal Society Interface*, 8(62):1281–1297, 2011.
- [20] Norman Margolus. Physics-like models of computation. *Physica D: Nonlinear Phenomena*, 10(1-2):81–95, 1984.
- [21] Tommaso Toffoli and Norman Margolus. *Cellular automata machines: a new environment for modeling*. MIT press, 1987.
- [22] Norman H Margolus. Crystalline computation. In *Feynman and computation: exploring the limits of computers*, pages 267–305. Perseus Books, 1999.
- [23] Arthur Taylor Winfree. The scientist as poet. *The Cornell Engineer*, 1964.
- [24] David Yu Zhang and Erik Winfree. Control of DNA strand displacement kinetics using toehold exchange. *Journal of the American Chemical Society*, 131(47):17303–17314, 2009.
- [25] Daniel T Gillespie et al. Stochastic simulation of chemical kinetics. *Annual Review of Physical Chemistry*, 58(1):35–55, 2007.
- [26] Tim Schmiedl and Udo Seifert. Stochastic thermodynamics of chemical reaction networks. *The Journal of Chemical Physics*, 126(4):044101, 2007.
- [27] Grigory Tikhomirov, Philip Petersen, and Lulu Qian. Programmable disorder in random DNA tilings. *Nature Nanotechnology*, 12(3):251–259, 2017.
- [28] Ali Aghebat Rafat, Sandra Sagredo, Melissa Thalhammer, and Friedrich C Simmel. Barcoded DNA origami structures for multiplexed optimization and enrichment of DNA-based protein-binding cavities. *Nature Chemistry*, 12(9):852–859, 2020.
- [29] Shubham Mishra, Soyoun Park, Tomoko Emura, Hidaka Kumi, Hiroshi Sugiyama, and Masayuki Endo. Photocontrolled DNA origami assembly by using two photoswitches. *Chemistry—A European Journal*, 27(2):778–784, 2021.
- [30] Peng Yin, Rizal F Hariadi, Sudheer Sahu, Harry MT Choi, Sung Ha Park, Thomas H LaBean, and John H Reif. Programming DNA tube circumferences. *Science*, 321(5890):824–826, 2008.

- [31] Bryan Wei, Mingjie Dai, and Peng Yin. Complex shapes self-assembled from single-stranded DNA tiles. *Nature*, 485(7400):623–626, 2012.
- [32] Bryan Wei, Mingjie Dai, Cameron Myhrvold, Yonggang Ke, Ralf Jungmann, and Peng Yin. Design space for complex DNA structures. *Journal of the American Chemical Society*, 135(48):18080–18088, 2013.
- [33] Bryan Wei, Luvena L Ong, Jeffrey Chen, Alexander S Jaffe, and Peng Yin. Complex reconfiguration of DNA nanostructures. *Angewandte Chemie International Edition*, 53(29):7475–7479, 2014.
- [34] Tatiana Brailovskaya, Gokul Gowri, Sean Yu, and Erik Winfree. Reversible computation using swap reactions on a surface. In *DNA Computing and Molecular Programming (LNCS volume 11648)*, pages 174–196. Springer, 2019.
- [35] Erik Winfree. On the computational power of DNA annealing and ligation. In *DNA-Based Computers (DIMACS volume 27)*, pages 199–221. American Mathematical Society, 1996.
- [36] Minqi Lin, Jin Xu, Dafang Zhang, Zhihua Chen, Xuncai Zhang, Zhen Cheng, Yufang Huang, and Yanbiao Li. 3D DNA self-assembly model for graph vertex coloring. *Journal of Computational and Theoretical Nanoscience*, 7(1):246–253, 2010.
- [37] David Yu Zhang, Sherry Xi Chen, and Peng Yin. Optimizing the specificity of nucleic acid hybridization. *Nature Chemistry*, 4(3):208–214, 2012.
- [38] Weiyang Tang, Weiye Zhong, Yun Tan, Guan A Wang, Feng Li, and Yizhen Liu. DNA strand displacement reaction: a powerful tool for discriminating single nucleotide variants. *DNA Nanotechnology*, pages 377–406, 2020.
- [39] Cameron Chalk, Austin Luchsinger, Eric Martinez, Robert Schweller, Andrew Winslow, and Tim Wylie. Freezing simulates non-freezing tile automata. In *DNA Computing and Molecular Programming (LNCS volume 11145)*, pages 155–172. Springer, 2018.
- [40] Kenichi Morita. *Theory of reversible computing*. Springer, 2017.
- [41] Jia Lee, Ferdinand Peper, Susumu Adachi, Kenichi Morita, and Shinro Mashiko. Reversible computation in asynchronous cellular automata. In *Unconventional Models of Computation (LNCS 2509)*, pages 220–229. Springer, 2002.
- [42] Jia Lee, Ferdinand Peper, Kenji Leibnitz, and Ping Gu. Characterization of random fluctuation-based computation in cellular automata. *Information Sciences*, 352:150–166, 2016.
- [43] Samuel Clamons, Lulu Qian, and Erik Winfree. Programming and simulating chemical reaction networks on a surface. *Journal of the Royal Society Interface*, 17(166):20190790, 2020.
- [44] Chris Thachuk and Anne Condon. Space and energy efficient computation with DNA strand displacement systems. In *DNA Computing and Molecular Programming (LNCS volume 7433)*, pages 135–149. Springer, 2012.
- [45] Lulu Qian and Erik Winfree. Parallel and scalable computation and spatial dynamics with DNA-based chemical reaction networks on a surface. In *DNA Computing and Molecular Programming (LNCS volume 8727)*, pages 114–131. Springer, 2014.
- [46] Niranjan Srinivas, James Parkin, Georg Seelig, Erik Winfree, and David Soloveichik. Enzyme-free nucleic acid dynamical systems. *Science*, 358(6369):eaal2052, 2017.
- [47] Peng Yin, Harry MT Choi, Colby R Calvert, and Niles A Pierce. Programming biomolecular self-assembly pathways. *Nature*, 451(7176):318–322, 2008.
- [48] Jennifer E Padilla, Matthew J Patitz, Robert T Schweller, Nadrian C Seeman, Scott M Summers, and Kingsi Zhong. Asynchronous signal passing for tile self-assembly: Fuel efficient computation and efficient assembly of shapes. *International Journal of Foundations of Computer Science*, 25(04):459–488, 2014.
- [49] Jennifer E Padilla, Ruojie Sha, Martin Kristiansen, Junghuei Chen, Natasha Jonoska, and Nadrian C Seeman. A signal-passing DNA-strand-exchange mechanism for active self-assembly of DNA nanostructures. *Angewandte Chemie International Edition*, 54(20):5939–5942, 2015.
- [50] David Yu Zhang, Rizal F Hariadi, Harry MT Choi, and Erik Winfree. Integrating DNA strand-displacement circuitry with DNA tile self-assembly. *Nature Communications*, 4(1):1–10, 2013.
- [51] Nicholas Schiefer and Erik Winfree. Universal computation and optimal construction in the chemical reaction network-controlled tile assembly model. In *DNA Computing and Molecular Programming (LNCS volume 9211)*, pages 34–54. Springer, 2015.