Flipping Tiles: Concentration Independent Coin Flips in Tile Self-Assembly

Cameron T. Chalk *, Bin Fu **, Alejandro Huerta *, Mario A. Maldonado *, Eric Martinez *, Robert T. Schweller *, and Tim Wylie

University of Texas Rio Grande Valley, Edinburg, TX 78539-2999
Email: {cameron.chalk01,bin.fu,alejandro.huerta02, eric.m.martinez02,robert.schweller,timothy.wylie}@utrgv.edu

Abstract. In this paper we introduce the *robust coin flip* problem in which one must design an abstract tile assembly system (aTAM system) whose terminal assemblies can be partitioned such that the final assembly lies within either partition with exactly probability 1/2, regardless of what relative concentration assignment is given to the tile types of the system. We show that robust coin flipping is possible within the aTAM, and that such systems can guarantee a worst case $\mathcal{O}(1)$ space usage. As an application, we then combine our coin-flip system with the result of Chandran, Gopalkrishnan, and Reif [3] to show that for any positive integer n, there exists a $\mathcal{O}(\log n)$ tile system that assembles a constant-width linear assembly of expected length n that works for all concentration assignments. We accompany our primary construction with variants that show trade-offs in space complexity, initial seed size, temperature, tile complexity, bias, and extensibility, and also prove some negative results. Further, we consider the harder scenario in which tile concentrations change arbitrarily at each assembly step and show that while this is not solvable in the aTAM, this version of the problem can be solved by more exotic tile assembly models from the literature.

1 Introduction

Self-assembly is the process by which local interactivity among unorganized, autonomous units results in their amalgamation into compounds. One of the premiere models for studying the theoretical possibilities of self-assembly is the *abstract tile assembly model* (aTAM) [22] in which system monomers are 4-sided Wang tiles that attach to a growing seed assembly whenever matching glues present a sufficient bonding strength. The motivation for studying the aTAM stems from the feasibility of a nanoscale DNA implementation [12], along with the universal computational power of the model [19], which permits many features including *algorithmic* self-assembly of general shapes [20], and more [8,17].

A promising new direction in self-assembly is the consideration of *randomized* self-assembly systems. In randomized self-assembly (a.k.a. nondeterministic self-assembly), assembly growth is dictated by nondeterministic, competing assembly

* Research supported in part by National Science Foundation Grant CCF-1117672.
** Research supported by National Science Foundation Early Career Award 0845376.

paths yielding a probability distribution on a set of final, terminal assemblies. By careful design of tile-sets and the relative concentration distributions of these tiles, a number of new functionalities and efficiencies have been achieved that are provably impossible without this non-determinism. For example, by precisely setting the concentration values of a generic set of tile species, arbitrarily complex strings of bits can be *programmed* into the system to achieve a specific shape with high probability [9,15]. Alternately, if the concentration of the system is assumed to be fixed at a uniform distribution, randomization still provides for efficient expected growth of linear assemblies [3] and low-error computation at temperature-1 [6]. Even in the case where concentrations are unknown, randomized self-assembly can build certain classes of shapes without error in a provably more efficient manner than without randomization [2].

Motivated by the power of randomized self-assembly, along with the potential for even greater future impact, we focus on the development of the most fundamental randomization primitive: the *robust* generation of a uniform random bit. In particular, we introduce the problem of self-assembling a uniformly random bit within $\mathcal{O}(1)$ space that is guaranteed to work for all possible concentration distributions. We define a tile system to be a *coin flip* system, with respect to some tile concentration distribution, if the terminal assemblies of the system can be partitioned such that each partition has exactly probability 1/2 of assembling. We say a system is a *robust coin flip* system if such a partition exists that guarantees 1/2 probability for all possible tile concentration distributions. By designing systems that flip a fair coin for all possible (adversarially chosen) concentration distributions, we achieve an intrinsically fair coin-flipping system that is robust to the experimental realities of imprecise quantity measurements. Such intrinsically fair systems may further allow for increased scalability of randomized self-assembly systems in scenarios where exact concentrations of species are either unknown or intractable to predict at successive assembly stages.

Our results. Our primary result is an aTAM construction that constitutes a robust fair coin flip system which completes in a guaranteed $\mathcal{O}(1)$ space. We apply our robust coin-flip construction to the result of Chandran, Gopalkrishnan, and Reif [3] to show that for any positive integer n, there exists a $\mathcal{O}(\log n)$ tile system that assembles a constant width-4 linear assembly of expected length n that works for all concentration assignments. We accompany this result with a proof that such concentration independent assembly of width-1 assemblies is not possible with fewer than n tile types. We further accompany our main coin-flip construction with variant constructions that provide trade-offs among standard aTAM metrics such as space, tile complexity, and temperature, as well as new metrics such as coin bias, and the *extensibility* of the system, which is the maximum number of distinct locations a single assembly of the system can add a tile. We show that 1-extensible systems, while computationally universal, cannot robustly coin-flip in bounded space without incurring a bias, but can robustly coin-flip in bounded expected space. We also consider the more extreme model in which concentrations may change adversarially at each assembly step. We show that the aTAM cannot robustly coin flip in bounded space within this model,

SUMMARY OF POSITIVE COIN FLIP RESULTS

Robust Coin Flip in the aTAM					
Space	Bias	τ	$\|\sigma\|$	k-ext	Theorem
$\mathcal{O}(1)$	-	1	7	2	1
$\mathcal{O}(1)$	-	2	1	2	2
unbounded	-	2	1	1	4
c	$<p^{(c/2)+1}$	2	1	1	5

Unstable Concentrations Robust Coin Flip					
Model	Space	Bias	τ	$\|\sigma\|$	Theorem
neg-aTAM	$\mathcal{O}(1)$	-	1	2	9
neg-hTAM	$\mathcal{O}(1)$	-	1	1	9
polyTAM	$\mathcal{O}(1)$	-	2	3	9
GTAM	$\mathcal{O}(1)$	-	1	2	9

Table 1: τ represents the temperature of the system, $|\sigma|$ represents the number of tiles in the seed assembly, and k-ext denotes the extensibility of the system. p represents the largest disparity in relative tile concentration between any pair of tile types in the system for a given concentration distribution.

but a number of more exotic extensions of the aTAM from the literature are able to robustly coin flip in $\mathcal{O}(1)$ space. We summarize our results in Table 1. The problem of self-assembling random bits has been considered before [11], but their technique, and in fact almost all randomized techniques to date, do not work when arbitrary concentrations are considered.

2 Definitions and Model

2.1 Tiles, Assemblies, and Tile Systems

Consider some alphabet of glue types Π. A tile is a unit square with 4 edges each assigned some glue type from Π. Further, each glue type $g \in \Pi$ has some non-negative integer strength $str(g)$. Each tile may be assigned a finite length string label, e.g., "black","white","0", or "1". Further, for simplicity, we assume each tile center is located at a pixel $p = (p1, p2) \in \mathbb{Z}^2$. For a given tile t, we denote the tile center of t as its position. As notation, we denote the set of all tiles that constitute all translations of the tiles in a set T as the set T^*. An *assembly* is a set of tiles each assigned unique coordinates in $\mathbb{Z}^2$. For a given assembly α, define the *bond graph* G_α to be the weighted graph in which each element of α is a vertex, and each edge weight between tiles is $str(g)$ if the tiles share an overlapping glue g, and 0 otherwise. An assembly α is said to be *τ-stable* for a positive integer τ if the bond graph G_α has min-cut at least τ, and *τ-unstable* otherwise. A *tile system* is an ordered triple $\Gamma = (T, \sigma, \tau)$ where T is a set of tiles called the *tile set* (we refer to elements of T as tile types), σ is an assembly called the *seed* and τ is a positive integer called the *temperature*. When considering a tile a that is some translation of an element of a tile set T, we will use the term *tile type* of a to reference the element of T that a is a translation of. Assembly proceeds by growing from assembly σ by any sequence of single tile attachments from T so long as each tile attachment connects with strength at least τ. Formally, we define what can be built in this fashion as the set of producible assemblies:

Definition 1 (Producibility). *For a given tile system $\Gamma = (T, \sigma, \tau)$, the set of* ***producible assemblies*** *for system Γ, $\texttt{PROD}_\Gamma$, is defined recursively:*

- *(Base) $\sigma \in \texttt{PROD}_\Gamma$*
- *(Recursion) For any $A \in \texttt{PROD}_\Gamma$ and $b \in T^*$ such that $C = A \cup \{b\}$ is τ-stable, then $C \in \texttt{PROD}_\Gamma$.*

As additional notation, we say $A \rightarrow_1^\Gamma B$ if A may grow into B through a single tile attachment, and we say $A \rightarrow^\Gamma B$ if A can grow into B through 0 or more tile attachments. An **assembly sequence** for a tile system Γ is a sequence (finite or infinite) $\overrightarrow{\alpha} = \langle \alpha_1, \alpha_2, \dots \rangle$ in which $\alpha_1 = \sigma$, each α_{i+1} is a single-tile extension of α_i, and each α_i is τ-stable. The **frontier** of an assembly α, written as $F(\alpha, \Gamma)$, is a partial function that maps an assembly α and a tile system Γ to a set of tiles $\{t \in T^* | \alpha \cup \{t\} \in \texttt{PROD}_\Gamma \wedge t \notin \alpha\}$. We further define $\texttt{TERM}_\Gamma$ to be the subset of $\texttt{PROD}_\Gamma$ consisting only of assemblies for which no further tile in T may attach.

Definition 2 (Finiteness and Space). *For a given tile assembly system $\Gamma = (T, \sigma, \tau)$, we say Γ is* ***finite*** *iff $\forall \sigma \in \texttt{PROD}_\Gamma, \exists \alpha \in \texttt{TERM}_\Gamma : \sigma \rightarrow^\Gamma \alpha$. That is, each producible assembly has a growth path ending in a finite, terminal assembly. If Γ is not finite, we say it is* ***infinite****. Define the* ***space of an assembly*** *α as $|\alpha|$. Let the* ***space of a tile assembly system*** *be defined as the $\max\limits_{\alpha \in \texttt{TERM}_\Gamma} |\alpha|$ iff Γ is finite. If Γ is infinite, let* ***space*** *remain undefined. Note that a finite system may have infinite/unbounded space.*

Definition 3 (Extensibility). *Consider a tile assembly system $\Gamma = (T, \sigma, \tau)$, and assembly $\alpha \in \texttt{PROD}_\Gamma$. We denote the set of all locations at which a tile may stably attach to α as L_α. More formally, $L_\alpha = \{p_t | t \in F(\alpha, \Gamma)\}$. We say a tile system Γ is k-**extensible** iff $\forall \alpha \in \texttt{PROD}_\Gamma, |L_\alpha| \leq k$. Informally, a tile assembly system is k-extensible iff at any point in the assembly process, the assembly can only grow in at most k locations.*

2.2 Probability in Tile Assembly

We use the definition of probabilistic assembly presented in [1,3,6,15,9]. Let P be a function denoting a **concentration distribution** over a tileset T representing the concentrations of each tile type with the restrictions $\forall t \in T, P(t) > 0$ and $\sum\limits_{t \in T} P(t) = 1$. For a tile t, we sometimes refer to $P(t)$ as the **concentration** of t. Using a concentration distribution, we can consider probabilities for certain events in the system. To study probabilistic assembly, we can consider the assembly process as a Markov chain where each producible assembly is a state and transitions occur with non-zero probability from assembly A to each B whenever $A \rightarrow_1^\Gamma B$. For each B that satisfies $A \rightarrow_1^\Gamma B$, let $t_{A \rightarrow B}$ denote the tile in T whose translation is added to A to get B. The transition probability from A to B is defined to be

$$TRANS(A,B) = \frac{P(t_{A\to B})}{\sum_{\{C|A\to_1^{\Gamma} C\}} P(t_{A\to C})} \tag{1}$$

The probability that a tile system Γ terminally assembles an assembly A is defined to be the probability that the Markov chain ends in state A. For each $A \in \texttt{TERM}_\Gamma$, let $\texttt{PROB}^P_{\Gamma\to A}$ denote the probability that Γ terminally assembles A with respect to concentration distribution P.

Definition 4 (Expected Space). *For a given finite tile system $\Gamma = (T, \sigma, \tau)$, let the expected space of Γ relative to a concentration distribution P be defined as*

$$\texttt{EXPECTEDSPACE}_\Gamma = \sum_{\alpha \in \texttt{TERM}_\Gamma} |\alpha| \cdot \texttt{PROB}^P_{\Gamma\to\alpha} \tag{2}$$

Definition 5 (Coin Flipping). *We consider a finite tile system Γ a* ***coin flip tile system with bias*** *b with respect to a concentration distribution P for some $b \in \mathbb{R}$ iff the set of terminal assemblies in $\texttt{PROD}_\Gamma$ is partitionable into two sets X and Y such that $\left|\sum_{x\in X} \texttt{PROB}^P_{\Gamma\to x} - \sum_{y\in Y} \texttt{PROB}^P_{\Gamma\to y}\right| \le 2b$. A* ***fair coin flip tile system*** *is a coin flip tile system with bias 0. We consider a finite tile system Γ a* ***robust coin flip tile system with bias*** *b iff the set of terminal assemblies in $\texttt{PROD}_\Gamma$ is partitionable into two sets X and Y such that $\left|\sum_{x\in X} \texttt{PROB}^C_{\Gamma\to x} - \sum_{y\in Y} \texttt{PROB}^C_{\Gamma\to y}\right| \le 2b$ for all concentration distributions C. A* ***robust fair coin flip tile system*** *is a robust coin flip tile system with bias 0.*

3 Robust Fair Coin Flipping in the aTAM

In this section we show systems capable of robust fair coin flips in the aTAM. Figure 1 shows a simple fair coin flip aTAM system for the uniform concentration distribution. To solve this problem for arbitrary concentration distributions, more involved techniques are required.

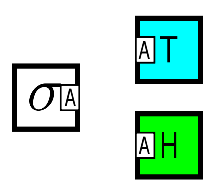

Fig. 1: A non-robust fair coin flip for the uniform concentration distribution.

Theorem 1. *There exists a $\mathcal{O}(1)$ space 2-extensible robust fair coin flip tile system $\Gamma = (T, \sigma, 1)$ in the aTAM with $|\sigma| = 7$.*

Proof. To show this we present a tile system $\Gamma = (T, \sigma, 1)$ in which two terminal states exist and are equiprobable for all concentration distributions P. $|T| = 9$ and σ contains 7 tiles. The system terminates nondeterministically and contains either 2 h tiles and 1 t tile or 2 t tiles and 1 h tile. The system leverages any

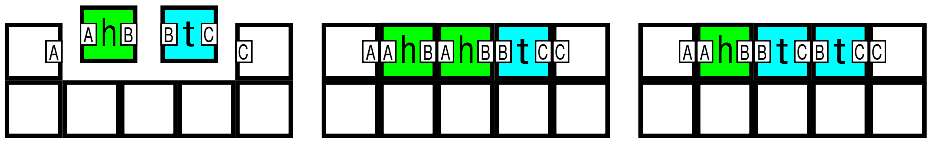

Fig. 2: Shown are the σ, h, and t tiles on the left, and the terminal states of the assembly system representing heads and tails. A, B and C glues are strength 1. Non-matching glues have 0 strength.

difference in tile concentrations between h and t by ensuring that placement of a t tile increases the probability of terminating in an assembly containing $2h$ tiles and vice versa. A graphical representation of σ, the h and t tiles, and terminal states of the assembly system is shown in Figure 2. Without loss of generality, assume the leftmost bottom tile in σ sits at position $(0, 0)$. We will refer to each producible assembly sans σ by the labels of the tiles in positions $(1, 1), (2, 1)$ and $(3, 1)$ as such: $\_\_t, h\_\_, \_ht$ and so forth. We now show that $\texttt{PROB}^P_{\Gamma \to hht} = \frac{1}{2}$ for all concentration distributions P. Let c_h be the concentration of the tile labeled h and c_t be the concentration of the tile labeled t, then

$$\begin{aligned}
\texttt{PROB}^P_{\Gamma \to hht} &= TRANS(\sigma, __t) \cdot TRANS(__t, _ht) \cdot TRANS(_ht, hht) \\
&\quad + TRANS(\sigma, __t) \cdot TRANS(__t, h_t) \cdot TRANS(h_t, hht) \\
&\quad + TRANS(\sigma, h__) \cdot TRANS(h__, h_t) \cdot TRANS(h_t, hht) \\
&= \frac{c_t}{c_t + c_h} \cdot \frac{c_h}{c_h + c_h} \cdot \frac{c_h}{c_h} + \frac{c_t}{c_t + c_h} \cdot \frac{c_h}{c_h + c_h} \cdot \frac{c_h}{c_t + c_h} \\
&\quad + \frac{c_h}{c_t + c_h} \cdot \frac{c_t}{c_t + c_t} \cdot \frac{c_h}{c_t + c_h} \\
&= \frac{{c_t}^2 + 2c_t c_h + {c_h}^2}{2{c_t}^2 + 4c_t c_h + 2{c_h}^2} = \frac{1}{2}.
\end{aligned}$$

□

3.1 Extension to a Single-Seed

A common constraint in the aTAM is that σ contains only one tile. Thus, no seed structure must be formed prior to the self-assembly process. The construction shown in Figure 3 addresses this constraint and works in a similar fashion as the construction in Theorem 1. Note that this system requires $\tau = 2$.

Theorem 2. *There exists a $\mathcal{O}(1)$ space 2-extensible robust fair coin flip tile system $\Gamma = (T, \sigma, 2)$ in the aTAM with $|\sigma| = 1$.*

Proof. Our tile set is shown in Figure 3. Without loss of generality, assume σ sits at position $(0, 0)$. Until the tile labeled S (see Figure 3) is placed, the assembly process is deterministic. Upon attachment of S, cooperative binding locations allow the attachment of tiles h and t nondeterministically. We denote the assemblies following the placement of S similarly to the proof of Theorem 1. We refer to assemblies containing tile S by the labels of tiles in positions $(1, -1), (1, 0)$ and $(2, 0)$ as $\_\_t, \_\_h, \_ht$ and so forth. Reflecting the analysis shown in Theorem 1, we have $\texttt{PROB}^P_{\Gamma \to hht} = .5$ for all concentration distributions P, which implies $\texttt{PROB}^P_{\Gamma \to htt} = .5$ as there are two terminal assemblies. □

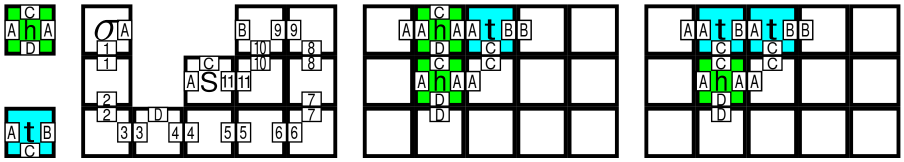

Fig. 3: T is shown. Our seed, labeled σ, begins a deterministic attachment process ending with the placement of the tile labeled S. Glues labeled $\{1, 2, 3, \ldots, 11\}$ are of strength 2. Glues labeled $\{A, B, C, D\}$ are of strength 1, ensuring that the nondeterministic attachments of tiles h and t do not begin until the cooperative binding locations are opened by placement of the tile labeled S. The nondeterministic sequence of attachments following the placement of S is similar to that of Theorem 1.

3.2 1-Extensible Coin Flipping

The previous sections showcase 2-extensible solutions to the robust fair coin flip problem. A natural question follows: is there a 1-extensible solution? Theorem 3 shows that there is no $\mathcal{O}(1)$ space solution in the aTAM. Using algorithms based on John von Neumann's randomness extractor [21] we can achieve an unbounded space robust fair coin flip system (Theorem 4) as well as a $\mathcal{O}(1)$ space construction which incurs a small bias (Theorem 5).

Theorem 3. *There does not exist a $\mathcal{O}(1)$ space 1-extensible robust fair coin flip tile system in the aTAM.*

Proof. We prove this by contradiction. Assume that there exists a $\mathcal{O}(1)$ space 1-extensible robust fair coin flip aTAM tile system $\Gamma = (T, \sigma, \tau)$. We now specify a concentration distribution for m tiles in T that contradicts this claim. Assume that Γ generates assemblies of size at most h. Consider a series of phases $p_1, \ldots, p_n$ such that p_{i+1} is derived from p_i by the attachment of the tile in the frontier of p_i with the largest concentration. Select a parameter $t = 10mn^3$, and let $c_1 = 1$ and $c_{i+1} = tc_i$ for $i = 1, \ldots, m-1$. Let the concentration for each $t_i \in T$ be $\frac{c_i}{c_1+c_2+\cdots+c_m}$.

For each assembly p_i, let $q_{i_1}, \ldots, q_{i_u}$ be the set of tile types in the frontier of p_i listed in increasing order by their concentrations. Let c_{i_u} denote the concentration of tile type q_{i_u}. With probability $\frac{c_{i_u}}{c_{i_1}+\cdots+c_{i_u}}$, tile type q_{i_u} is attached. We have

$$\frac{c_{i_u}}{c_{i_1}+\cdots+c_{i_u}} \geq \frac{1}{\frac{(u-1)c_{i_{u-1}}}{c_{i_u}}+1} \tag{3}$$

$$\geq \frac{1}{\frac{(u-1)}{t}+1} \geq \frac{1}{\frac{m}{t}+1} \tag{4}$$

$$\geq \frac{1}{\frac{1}{10n^3}+1}. \tag{5}$$

Algorithm 1 Unbounded

```
procedure UnboundedFCFE(h, t)
    coin = {heads, tails}
    pdist = {h, t}
    repeat
        flip_1 ← flip(coin, pdist)
        flip_2 ← flip(coin, pdist)
    until flip_1 ≠ flip_2
    return flip_2
end procedure
```

Algorithm 2 Bounded

```
procedure BoundedFCFE(h, t, k)
    coin = {heads, tails}
    pdist = {h, t}
    round ← 1
    while round ≤ k do
        flip_1 ← flip(coin, pdist)
        flip_2 ← flip(coin, pdist)
        if flip_1 ≠ flip_2 then
            return flip_2
        end if
        round ← round + 1
    end while
    return flip(coin, h, t)
end procedure
```

Therefore, with probability at least

$$\left(\frac{1}{\frac{1}{10n^3}+1}\right)^n \geq \left(\frac{1}{\frac{1}{10n^3}+1}\right)^{10n^3\cdot\frac{1}{10n^2}} \tag{6}$$

$$\geq \left(\frac{1}{e}\right)^{\frac{1}{10n^2}} > 0.6 \tag{7}$$

we follow the sequence $p_1, \ldots, p_n$ to generate an assembly. This is a contradiction. Note that we use the facts that $(1+\frac{1}{x})^x$ is an increasing function for all real $x > 1$, and $\lim_{x\to+\infty}(1+\frac{1}{x})^x = e \approx 2.17828$. □

In response to Theorem 3, we give a 1-extensible aTAM system capable of robust fair coin flips in unbounded space in Theorem 4. In 1951, John von Neumann gave a simple method for extracting a fair coin from a biased one [21]. We show two algorithms based on the Von Neumann extractor. Algorithm 1 uses an unbounded number of *rounds* to extract a fair coin flip. We use Algorithm 1 to show that a *fair coin flip extractor* can be implemented in the aTAM to achieve an unbounded space, 1-extensible, robust coin flip tile system. We extend this method in Algorithm 2 to create a *bounded fair coin flip extractor* by adding a parameter k which controls the maximum number of rounds allowed. This is a *bounded coin flip extractor* that is implemented in the aTAM and achieves $\mathcal{O}(1)$ space, is 1-extensible, and is a robust coin flip tile system with bounded bias.

We now describe our 1-extensible aTAM tile system that implements Algorithm 1. In Algorithm 1, a coin is a set of cardinality 2 with possible values *heads* and *tails*. *flip* is a function that selects and returns a *heads* or *tails* value based on the probabilities h and t, where $h, t \in (0, 1)$ and $h + t = 1$. In our construction, calls to the *flip* function are carried out by a non-deterministic competition for attachment between a *0* tile and a *1* tile. Aside from calls to the *flip* function, the rest of the algorithm can be implemented by deterministic

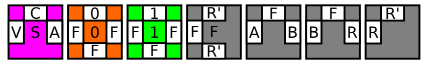

(a) Tile set that makes two nondeterministic *flips* corresponding to the two calls to the *flip* function in 1.

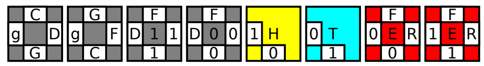

(b) Tile set that checks the result of the two flips and possibly starts another round if a fair bit has not been achieved. A *HEADS* or *TAILS* tile is placed if a fair bit has been achieved.

Fig. 4: The tile labeled S is the seed of the tile assembly system and the temperature is 2. The strength of the glues are as follows: str(0)=1, str(1)=1, str(A)=2, str(B)=2, str(C)=1, str(D)=1, str(F)=1, str(G)=2, str(R)=2, and str(R')=2.

tile placements. Figure 4 gives the tile set used in the construction. Consider all tiles labeled *H* as *HEADS* tiles and all tiles labeled *T* as *TAILS* tiles where their placement implies the returning of heads and tails, respectively. Consider all tiles labeled *E* as *ERR* tiles. The set of tiles in Figure 4(a) starts the process and makes two non-deterministic placements of a *1* tile or a *0* tile. The set of tiles in Figure 4(b) checks the result of the two flips. If the order of the flips, starting from the left, is *10*, it outputs a *HEADS* tile. If the order of the flips is *01*, it outputs a *TAILS* tile. Otherwise, it outputs an *ERR* tile, which starts another loop. Figure 5 shows examples of assemblies that can grow in Round 1 of the algorithm. This construction yields Theorem 4. The full analysis of this construction is omitted in this version due to space.

Theorem 4. *There exists a 1-extensible, robust coin flip tile system in the aTAM. The tile system achieves $\mathcal{O}(1/pq)$ expected space, where p and q denote the relative concentrations of the two tiles with the largest difference in concentration for a given concentration distribution.*

We now extend Algorithm 1 by adding a parameter k, which controls the maximum number of rounds allowed (Algorithm 2). This *bounded fair coin flip extractor* can be implemented in the aTAM to achieve a $\mathcal{O}(1)$ space, 1-extensible, robust coin flip tile system with bounded bias. The bounded k-rounds can be controlled by the implementation of a 1-extensible version the the aTAM counter construction from [5] for a desired base, leading to a tradeoff in bias, space, and tile complexity. We state the primary tradeoff in Theorem 5 between space and bias, and omit the tradeoff in tile complexity in this version, as well as construction details and analysis.

Theorem 5. *There exists a c space 1-extensible robust coin flip tile system in the aTAM with bias less than $p^{(c/2)+1}$, where p denotes the larger relative concentration of the pair of tiles with the largest difference in concentration for a given concentration distribution.*

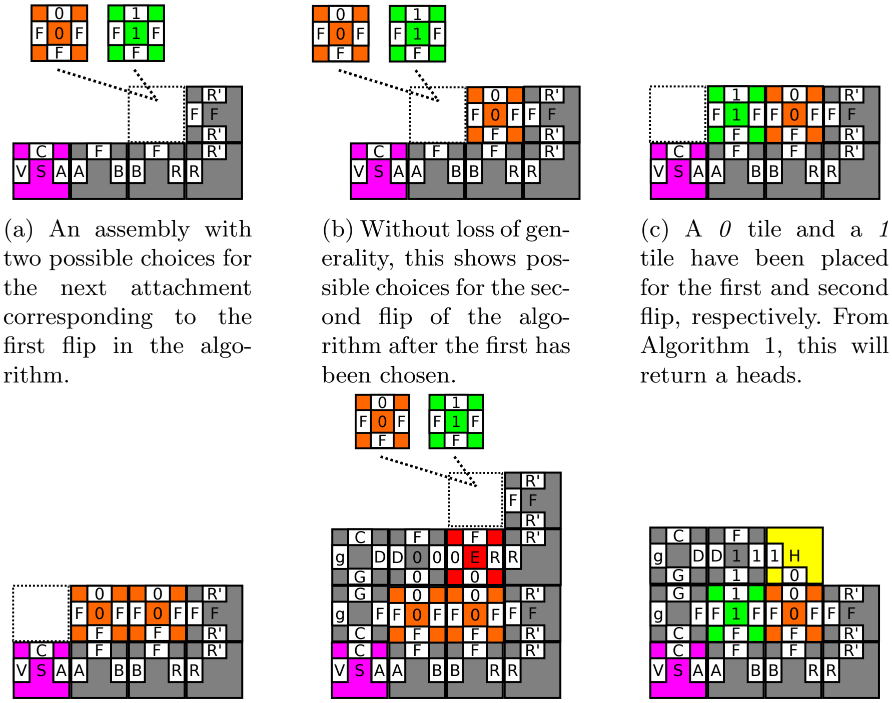

(a) An assembly with two possible choices for the next attachment corresponding to the first flip in the algorithm.

(b) Without loss of generality, this shows possible choices for the second flip of the algorithm after the first has been chosen.

(c) A *0* tile and a *1* tile have been placed for the first and second flip, respectively. From Algorithm 1, this will return a heads.

(d) Two *0 tiles* were placed for the first two flips. From Algorithm 1, the system must perform another round.

(e) An assembly where the first round of the algorithm failed to generate a bit and proceeds to start a new round.

(f) An assembly where the first round of the algorithm was a valid flip and it generates a heads.

Fig. 5: A sample of producible assemblies for Round 1

4 Robust Simulation of Randomized Linear Assemblies

As an application of the primitive shown in Theorem 2, we show that a class of randomized linear aTAM tile assembly systems can be simulated in a concentration robust manner with a minor scale factor.

We first briefly describe a scale (m, n)-simulation of a given tile system, based on the block replacement schemes of [4]. Consider an aTAM system $\Gamma = (T, \sigma, \tau)$ and a proposed simulator system $\Gamma' = (T', \sigma', \tau')$. Now consider the mapping from $\texttt{TERM}_{\Gamma}$ to $\texttt{TERM}_{\Gamma'}$ obtained by replacing each tile in an assembly $A \in \texttt{TERM}_{\Gamma}$ with a rectangular $m \times n$ block of tiles over U, according to some fixed $m \times n$ block mapping R. If there exists such a mapping M from $\texttt{TERM}_{\Gamma}$ to $\texttt{TERM}_{\Gamma'}$ that is bijective, then we say that Γ' simulates the production of Γ at scale factor (m, n). Further, we say that Γ *robustly simulates* Γ' for concentration distribution P if for all terminal assemblies $A \in \texttt{TERM}_{\Gamma}$, $\texttt{PROB}^{P}_{\Gamma \to A} = \texttt{PROB}^{C}_{\Gamma' \to M(A)}$ for all concentration distributions C over T', i.e., Γ' produces terminal assemblies with probability independent of concentration assignment, and with exactly the same probability distribution as the concentration dependent system it simulates.

We now define a class of linear assembly systems for which we can construct robust, concentration independent simulations.

Definition 6 (Unidirectional two-choice linear assembly systems). *A tile system Γ is a* unidirectional two-choice linear assembly system *iff:*

1. *Γ is 1-extensible,*
2. *$\forall \alpha \in \mathtt{PROD}_\Gamma, |F(\alpha, \Gamma)| \leq 2$,*
3. *$\forall \beta \in \mathtt{PROD}_\Gamma, \beta$ is a $1 \times n$ line for some $n \in \mathbb{N}$.*

Theorem 6. *For any unidirectional two-choice linear assembly system $\Gamma = (T, \sigma, \tau)$ in the aTAM, there is an aTAM system $\Gamma_s = (T', \sigma', \tau')$ that robustly simulates Γ for the uniform concentration distribution at scale factor 5×4; further, $|T'| = c|T|$ for some constant c.*

Proof. Let $\Gamma = (T, \sigma, \tau)$ be a unidirectional two-choice linear assembly system. Define an *undecided assembly* to be any assembly $\alpha \in \mathtt{PROD}_\Gamma$ such that $|F(\alpha, \Gamma)| = 2$. For each undecided assembly, we will construct a gadget utilizing the technique in Theorem 2. We call the two tiles of an undecided assembly's frontier h and t. Consider $\alpha_h = \alpha \cup h$ and $\alpha_t = \alpha \cup t$. We simulate Γ in reference to a uniform concentration distribution, so α transitions to α_h with probability .5 and to α_t with probability .5. Figure 6 shows an example of utilizing a 5×4 gadget in Γ_s to simulate the transition from α to α_h or α_t. By application of Theorem 2, the gadget will grow into one of two possible states with probability .5 for any concentration distribution. By chaining the gadgets together we can robustly simulate the nondeterministic attachments in Γ. Each tile is simulated by a 5×4 block of tiles, therefore $|T'| = c|T|$ for some constant c. □

As a corollary to Theorem 6, we can create a tile system to build an expected length n assembly for all concentration distributions with $\mathcal{O}(\log n)$ tile complexity. First, we will prove that there is no aTAM tile system which generates linear (width-1) assemblies of expected length n for all concentration distributions ([3] showed that this is possible for the uniform concentration distribution).

Theorem 7. *There is no aTAM tile system to generate an assembly of width-1 and expected length n for all concentration distributions with less than n tile complexity.*

Proof. Towards a contradiction, assume a self-assembly system can generate a linear assembly with expected length n and uses at most $k < n$ tiles. There is at least one assembly S that is of length at least n. Let $S = t_1 \cdots t_{i-1} t_i \cdots t_m t_i ...$, where $t_i \cdots t_m t_i$ is the first cycle that appears in S since there are less than n tiles. We define the concentration of the types of tiles as follows:

Let $c_1 = 1$, $c_j = c_{j-1}/n^{100}$ for $j = 1, ..., k$. The concentration of each type t_i is $\frac{c_i}{c_1 + c_2 + \cdots + c_k}$. Therefore, with probability at least $(\frac{1}{1+\frac{1}{n^{99}}})^n$, the assembly

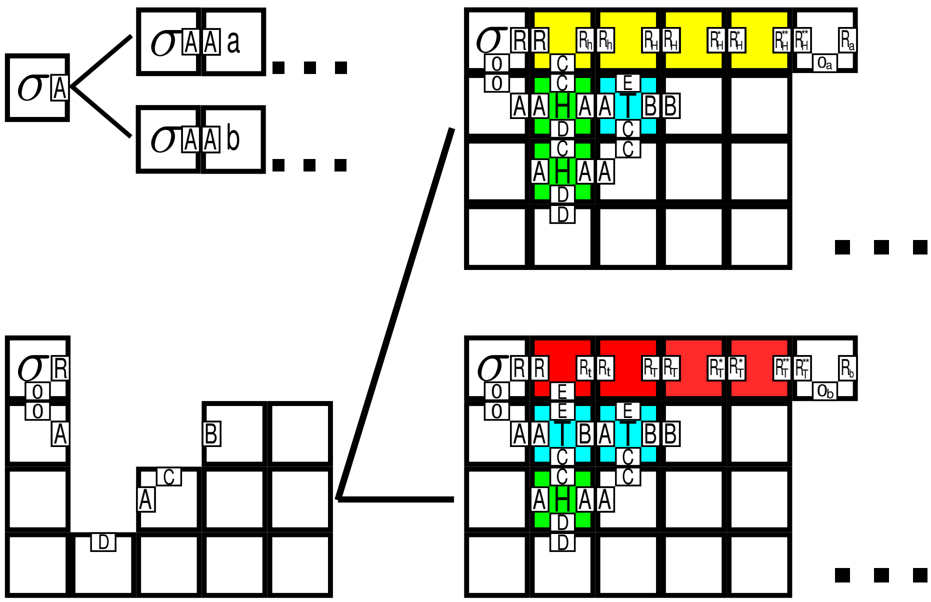

Fig. 6: A simulation of one non-deterministic linear tile attachment. Each non-determinstic attachment will require a 5×4 robust coin flip gadget shown in Fig. 3. The assembly may continue after simulating a non-determinstic attachment by building another 5×4 robust coin flip gadget, building a deterministic 5×4 block, or terminating.

$t_1 \cdots t_{i-1} t_i \cdots t_m t_i$, or one at least as long, will be generated. With probability at least $(\frac{1}{1+\frac{1}{n^{99}}})^{n^3} > 0.9$, an assembly at least as long as $t_1 \cdots t_{i-1}(t_i \cdots t_m t_i)^{n^2} \ldots$ will be generated, which has length at least n^2. This contradicts the assumption that the expected length is n. □

We now contrast the width-1 impossibility result of Theorem 7 with a result showing that width-4 linear assemblies do allow for efficient growth to expected length n in a concentration independent manner. To achieve this, we apply Theorem 6 to the unidirectional two-choice linear assembly system presented in [3], which yields the following result.

Corollary 1. *There exists an aTAM tile system $\Gamma = (T, \sigma, \tau)$ which terminates in a width-4 expected length n assembly for all concentration distributions. $|T| = \mathcal{O}(\log n)$.*

Proof. Let m be $\left\lfloor \frac{n}{5} \right\rfloor$. Consider $\Gamma = (T, \sigma, \tau)$ to be a robust simulation at scale factor 5×4 of a unidirectional two-choice linear assembly system that terminates in an expected length m linear assembly using $O(\log m)$ tile types. Note that such a unidirectional two-choice linear assembly system exists as shown in [3] and can be robustly simulated as shown by Theorem 6. If $5m = n$, then Γ terminates in an expected length n assembly with width-4; otherwise, we add $n \mod 5$ length deterministically. Since our scale factor is constant, $|T| = O(\log n)$. □

5 Robust Fair Coins with Unstable Concentrations

As an extension to the idea of concentration independent solutions outlined in this paper, we consider an adversarial model wherein the concentration distribution of tiles changes during each stage of the assembly process; in other words, the concentrations are unstable.

Definition 7 (Unstable Concentrations Robust Fair Coin Flip). *Let an* unstable concentration distribution P *be a function mapping* $z \in \mathbb{Z}^+$ *to concentration distributions over a tile set* T*. Let* P_i *denote* $P(i)$*. For each* B *that satisfies* $A \to_1^\Gamma B$*, let* $t_{A\to B}$ *denote the tile in* T *whose translation is added to* A *to get* B*. The transition probability from* A *to* B *is defined to be*

$$TRANS(A,B) = \frac{P_{|A|}(t_{A\to B})}{\sum_{\{C|A\to_1^\Gamma C\}} P_{|A|}(t_{A\to C})} \tag{8}$$

We consider a finite tile system Γ *an* ***unstable concentrations robust fair coin flip*** *iff the set of terminal assemblies in* $\texttt{PROD}_\Gamma$ *is partitionable into two sets* X *and* Y *such that* $\sum\limits_{x\in X} \texttt{PROB}^C_{\Gamma\to x} = \sum\limits_{y\in Y} \texttt{PROB}^C_{\Gamma\to y}$ *for all unstable concentration distributions* C*.*

We now prove that there is no unstable concentration robust fair coin flip system in the aTAM. First, we state and prove a lemma that will be useful in our proof.

Lemma 1. *For any producible assembly* $A \in \texttt{PROD}_\Gamma$ *and any tile type* $t \in T$*, there exists another assembly* A^* *such that for any sequence of assemblies* $\langle A_0 = A, A_1, A_2, \ldots, A_h\rangle$ *where* A_{i+1} *is derived from* A_i *by attaching a tile of type* t $(i = 0, 1, 2, \cdots, h-1)$*, and tile type* t *cannot be attached to* A_h*, then* $A_h = A^*$*.*

Proof. Let A^* be the least-sized producible assembly such that $A^* \setminus A$ contains only tiles of type t and the frontier of A^* contains no tiles of type t. We will show that A can only grow A^* if only allowed to attach tile type t.

Towards a contradiction, assume there exists a sequence of assemblies from A such that $A_h \neq A^*$. If A_h is some subassembly of A^*, note that we may still attach tiles of type t to reach A^*, implying that A_h does not fit the specified requirements. Otherwise, let A_n be the first assembly in the sequence which contains a tile not in A^*. Consider A_{n-1}. There is no tile of type t attachable to A_{n-1} such that the tile is not in A^*. If there were, that tile of type t would be attachable to A^*, contradicting the definition of A^*. Therefore no such A_n can exist, implying that A_h must be A^*.

Theorem 8. *There does not exist a* $\mathcal{O}(1)$ *space unstable concentrations robust fair coin flip tile system in the aTAM.*

Proof. Towards a contradiction, assume that a space-n solution does exist.

As the assembly process proceeds, the key point to consider is when the current assembly enters a state in which multiple distinct positions may attach a tile. In such a case select one type t of all attachable tiles, and increase its concentration to ensure, with high probability, that assembly proceeds by attaching only tiles of type t up until there is no position to attach type t tiles. Such a type t is called a *dominate* type. Let the concentration of the dominate tile type t be $(1-\frac{1}{100n^2})$. For each step i, let t_i denote the dominate type of concentration $(1-\frac{1}{100n^2})$.

When there is more than one position to attach the same type of tile t, we are assured by Lemma 1 that a unique assembly will result after repeatedly placing tiles of type t (in any order) until placement of t is no longer an option.

Given this setup, we have that at each step i, the assembly does not grow with a dominate type with probability at most $\frac{1}{10n^2}$. With probability at most $\frac{1}{10n}$, there is a step i among n steps that the assembly does not grow with the dominate type.

Therefore, there is a terminal assembly that will be generated with probability at least 0.9. This is a contradiction. □

Motivated by the impossibility of robust coin flipping in the aTAM under unstable concentrations, we now consider some established extensions of the aTAM from the literature. In particular, we show that robust coin flipping with unstable concentrations is possible within the aTAM with negative glues [22,10,18], the polyTAM [13], the hexTAM [7] with negative glues, and the GTAM [14].

Theorem 9. *There exists a $\mathcal{O}(1)$ space unstable concentration robust fair coin-flip tile system in the aTAM with negative glues, polyTAM, hexTAM with negative glues, and the GTAM.*

Proof. Consider a tile assembly system $\Gamma = (T, \sigma, \tau)$ with 3 producible assemblies: σ, a terminal assembly *heads*, and a terminal assembly *tails*. Further, $\sigma \rightarrow_1^{\Gamma}$ *heads* and $\sigma \rightarrow_1^{\Gamma}$ *tails*. Let $t_{\sigma \rightarrow heads}$ and $t_{\sigma \rightarrow tails}$ be the same tile c, then $TRANS(\sigma, heads) = TRANS(\sigma, tails) = \frac{P(c)}{2P(c)} = \frac{1}{2}$. Systems which meet these characteristics within the mentioned models can be seen in Figure 7.

□

6 Conclusions and Future Work

In this paper we have introduced the problem of designing robust, fair coin flipping systems. Generating such coin flips is fundamental for the implementation of randomized self-assembly algorithms. By incorporating concentration independent robustness into the design of such systems, we directly address the practical issue of limited control over species concentrations. Our goal in this work is to provide a stepping stone for the creation of general, robust randomized self-assembly systems. As evidence towards the feasibility of this goal, we have shown how our gadgets can be applied to convert a large class of linear systems into equivalent systems with the concentration robustness property. A

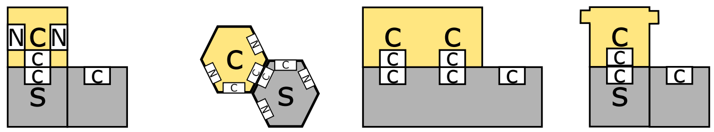

(a) (b) (c) (d)

Fig. 7: The terminal assemblies representing "heads" in some alternate models. C is a strength-τ glue and N is a strength-(-1) glue in (a) the aTAM tile system and (b) the hexTAM tile system. (c) C is a strength-1 glue in a $\tau = 2$ polyTAM tile system. (d) C is a strength-1 glue in a $\tau = 1$ GTAM tile system. The abutting geometry does not allow two C tiles to attach.

more general open problem is as follows: given a general tile system, is it possible to convert the system to an approximately equivalent system that is concentration robust? If possible, how efficiently can this be accomplished in terms of scale factor and approximation factor?

Another direction for future work is the consideration of generalizations of the coin flip problem. Our partition definition for coin flip systems extends naturally to distributions with more than two outcomes, as well as non-uniform distributions. What general probability distributions can be assembled in $O(1)$ space, and with what efficiency? We have also introduced the online variant of concentration robustness in which species concentrations may change at each step of the self-assembly process. We have shown that when such changes are completely arbitrary, coin flipping is not possible in the aTAM. A relaxed version of this robustness constraint could permit concentration changes to be bounded by some fixed rate. In such a model, how close to a fair flip can a system guarantee in terms of the given rate bound? As an additional relaxation, one could consider the problem in which an initial concentration assignment may be *approximately* set by the system designer, thereby modeling the limited precision an experimenter can obtain with a pipette.

A final line of future work focusses on applying randomization in self-assembly to computing functions. The parallelization within the abstract tile assembly model allows for substantially faster arithmetic than what is possible in non-parallel computational models [16]. Can randomization be applied to solve these problems even faster? Moreover, there are a number of potentially interesting problems that might be helped by randomization, such as primality testing, sorting, or a general simulation of randomized boolean circuits.

References

1. Florent Becker, Ivan Rapaport, and Eric Rémila, *Self-assembling classes of shapes with a minimum number of tiles, and in optimal time*, Foundations of Software Technology and Theoretical Computer Science (FSTTCS), 2006, pp. 45–56.
2. Nathaniel Bryans, Ehsan Chiniforooshan, David Doty, Lila Kari, and Shinnosuke Seki, *The power of nondeterminism in self-assembly*, Proc. of the 22nd ACM-SIAM Sym. on Discrete Algorithms, SODA'11, SIAM, 2011, pp. 590–602.
3. Harish Chandran, Nikhil Gopalkrishnan, and John Reif, *Tile complexity of linear assemblies*, SIAM Journal on Computing **41** (2012), no. 4, 1051–1073.

4. Ho-Lin Chen and Ashish Goel, *Error free self-assembly using error prone tiles*, DNA Computing, LNCS, vol. 3384, 2005, pp. 62–75.
5. Qi Cheng, Gagan Aggarwal, Michael H. Goldwasser, Ming-Yang Kao, Robert T. Schweller, and Pablo Moisset de Espanés, *Complexities for generalized models of self-assembly*, SIAM Journal on Computing **34** (2005), 1493–1515.
6. Matthew Cook, Yunhui Fu, and Robert T. Schweller, *Temperature 1 self-assembly: Deterministic assembly in 3D and probabilistic assembly in 2D*, Proc. of the 22nd ACM-SIAM Sym. on Discrete Algorithms, SODA'11, 2011, pp. 570–589.
7. ErikD. Demaine, MartinL. Demaine, SndorP. Fekete, MatthewJ. Patitz, RobertT. Schweller, Andrew Winslow, and Damien Woods, *One tile to rule them all: Simulating any tile assembly system with a single universal tile*, Automata, Languages, and Programming, LNCS, vol. 8572, 2014, pp. 368–379.
8. D. Doty, *Theory of algorithmic self-assembly*, Communications of the ACM **55** (2012), no. 12, 78–88.
9. David Doty, *Randomized self-assembly for exact shapes*, SIAM Journal on Computing **39** (2010), no. 8, 3521–3552.
10. David Doty, Lila Kari, and Benoît Masson, *Negative interactions in irreversible self-assembly*, Proc. of The 16th Inter. Meeting on DNA Computing and Molecular Programming, Lecture Notes in Computer Science, Springer, 2010, pp. 37–48.
11. David Doty, JackH. Lutz, MatthewJ. Patitz, ScottM. Summers, and Damien Woods, *Random number selection in self-assembly*, Unconventional Computation, Lecture Notes in Computer Science, vol. 5715, 2009, pp. 143–157.
12. Constantine Evans, *Crystals that count! physical principles and experimental investigations of dna tile self-assembly*, Ph.D. thesis, California Inst. of Tech., 2014.
13. Sándor P. Fekete, Jacob Hendricks, Matthew J. Patitz, Trent A. Rogers, and Robert T. Schweller, *Universal computation with arbitrary polyomino tiles in non-cooperative self-assembly*, Proc. of the 25th ACM-SIAM Sym. on Discrete Algorithms, SODA'15, SIAM, 2015, pp. 148–167.
14. Bin Fu, MatthewJ. Patitz, RobertT. Schweller, and Robert Sheline, *Self-assembly with geometric tiles*, Automata, Languages, and Programming, LNCS, vol. 7391, 2012, pp. 714–725.
15. Ming-Yang Kao and Robert T. Schweller, *Randomized self-assembly for approximate shapes*, Inter. Coll. on Automata, Languages, and Programming, Lecture Notes in Computer Science, vol. 5125, 2008, pp. 370–384.
16. Alexandra Keenan, Robert Schweller, Michael Sherman, and Xingsi Zhong, *Fast arithmetic in algorithmic self-assembly*, Unconventional Computation and Natural Computation, LNCS, vol. 8553, 2014, pp. 242–253.
17. Matthew J. Patitz, *An introduction to tile-based self-assembly and a survey of recent results*, Natural Computing **13** (2014), no. 2, 195–224.
18. MatthewJ. Patitz, RobertT. Schweller, and ScottM. Summers, *Exact shapes and turing universality at temperature 1 with a single negative glue*, DNA Computing and Molecular Programming, LNCS, vol. 6937, 2011, pp. 175–189.
19. Paul W. K. Rothemund and Erik Winfree, *The program-size complexity of self-assembled squares (extended abstract)*, Proc. of the 32nd ACM Sym. on Theory of Computing, STOC'00, 2000, pp. 459–468.
20. David Soloveichik and Erik Winfree, *Complexity of self-assembled shapes*, SIAM Journal on Computing **36** (2007), no. 6, 1544–1569.
21. John von Neumann, *Various Techniques Used in Connection with Random Digits*, Journal of Research of the National Bureau of Standards **12** (1951), 36–38.
22. Erik Winfree, *Algorithmic self-assembly of DNA*, Ph.D. thesis, California Institute of Technology, June 1998.