

Sampling Biased Monotonic Surfaces using Exponential Metrics*

Sam Greenberg[†]

Dana Randall[‡]

Amanda Pascoe Streib[§]

Abstract

Monotonic surfaces spanning finite regions of \mathbb{Z}^d arise in many contexts, including DNA-based self-assembly, card-shuffling and lozenge tilings. One method that has been used to uniformly generate these surfaces is a Markov chain that iteratively adds or removes a single cube below the surface during a step. We consider a biased version of the chain, where we are more likely to add a cube than to remove it, thereby favoring surfaces that are “higher” or have more cubes below it. We prove that the chain is rapidly mixing for any uniform bias in \mathbb{Z}^2 and for bias $\lambda > d$ in \mathbb{Z}^d when $d > 2$. In \mathbb{Z}^2 we match the optimal mixing time achieved by Benjamini, Berger, Hoffman, and Mossel in the context of biased card shuffling [2], but using much simpler arguments. The proofs use a geometric distance function and a variant of path coupling in order to handle distances that can be exponentially large. We also provide the first results in the case of *fluctuating bias*, where the bias can vary depending on the location of the tile. We show that the chain continues to be rapidly mixing if the biases are close to uniform, but that the chain can converge exponentially slowly in the general setting.

1 Introduction

In this paper, we are concerned with designing provably efficient algorithms for sampling from a family of discrete monotonic surfaces, where we bias the distribution to favor surfaces that are “higher.” There is a long history of sampling from various families of monotonic surfaces because many natural combinatorial problems are known to have an associated height function that can be interpreted as a piecewise-linear surface. In statistical physics, for example, domino tilings of the Cartesian lattice and lozenge tilings of the triangular lattice are natural models of diatomic molecules that have associated 3-dimensional height functions. Similarly, states of the zero temperature antiferromagnetic Potts model have a height function that maps 3-colorings in \mathbb{Z}^{d-1} to surfaces in \mathbb{Z}^d (see, for example, [13]). In each case, random sampling provides insight into the thermodynamic properties of these physical systems. We focus on surfaces that are unions of $d - 1$ dimensional faces of the Cartesian lattice \mathbb{Z}^d and discuss extensions to these other natural families at the end of the paper.

In two dimensions, monotonic surfaces, called *staircase walks*, are paths within a finite region of the lattice \mathbb{Z}^2 that step to the right or down at every edge (see Figure 1a). Markov chains for sampling staircase walks have been used to analyze card-shuffling algorithms by associating to a permutation a set of staircase walks [18, 2]. One simple Markov chain \mathcal{M}_U for sampling uniformly from the set of staircase walks, known as the “mountain / valley chain,” tries to invert a mountain

*A preliminary version of this paper appeared in the Proceedings of the 20th ACM/SIAM Symposium on Discrete Algorithms, 76–85, 2009.

[†]Department of Defense, Arlington, VA.

[‡]School of Computer Science, Georgia Institute of Technology, Atlanta GA, 30332-0765. Supported in part by NSF grants CCF-1526900, CCF-1637031 and CCF-1733812.

[§]Center for Computing Sciences, Bowie, MD 20715.

that moves to the right and then down to a valley that goes down and then to the right, or vice versa. This chain has also been studied in the context of *Dyck paths*, or staircase walks that start at $(0, n)$, end at $(n, 0)$ and do not cross below the line $x + y = n$. Dyck paths are enumerated by the n th Catalan number, and \mathcal{M}_U has proven useful for sampling from these and other Catalan structures [15]. Wilson [18] gave tight bounds on the convergence rate of the mountain / valley chain in the general case and in the case of Dyck paths by showing that in both cases it mixes in time $\Theta(n^3 \log n)$.

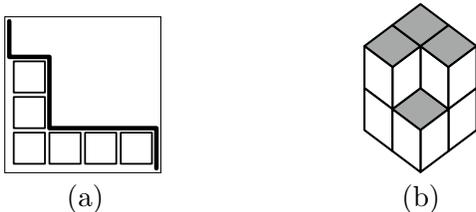


Figure 1: Monotonic surfaces in two and three dimensions.

Monotonic surfaces in \mathbb{Z}^3 correspond bijectively with *lozenge tilings* of finite regions of the triangular lattice [13]. A lozenge tiling is a covering of the lattice region with *lozenges*, or rhombuses that cover two adjacent triangular faces. Thus, a lozenge tiling is just a perfect matching (or dimer covering) in the dual graph. When we look at a two-dimensional picture of a lozenge tiling, our eyes automatically interpret the picture as a surface that is the upper envelope of a set of supported cubes in \mathbb{Z}^3 (see Figure 1b). We call this set of cubes σ a *downset*, since the set is downwardly-monotonic, and we let $M(\sigma)$ denote the monotonic surface formed by σ ; clearly these surfaces and downsets are also in bijection.

A natural Markov chain for uniformly sampling lozenge tilings tries to identify three closely packed lozenges forming a hexagon and rotates them by 180 degrees. Equivalently, we interpret this move quite naturally using the bijection with surfaces and this move corresponds to perturbing the surface locally by adding or removing a single cube, so this is precisely \mathcal{M}_U in \mathbb{Z}^3 . It is easy to see that this chain connects the state space because, starting from any configuration, we can remove cubes until we reach the “empty configuration.” This Markov chain is known to be rapidly mixing, or quickly converging to equilibrium, so it is effective for efficiently generating samples from close to the uniform distribution [13, 16, 18, 8].

There also has been interest in a *biased* version \mathcal{M}_B of this local Markov chain, where we are more likely to add unit cubes than remove them. More precisely, if σ is formed from τ by the addition of a single cube at position \bar{x} , then $\lambda_{\bar{x}} = P(\tau, \sigma)/P(\sigma, \tau)$ is the *bias at \bar{x}* . If $\lambda_{\bar{x}} > 1$ for every \bar{x} , then the stationary distribution favors configurations with more cubes.

Biased surfaces arise in nanoscience in the context of DNA-based self-assembly growth models (see, e.g., [10, 17, 19, 20]). In this setting, roughly “square” shaped tiles are constructed from strands of DNA so that each side of the tiles is single-stranded. Certain pairs of tiles are encouraged to line up and attach along edges by encoding corresponding sides with complementary sequences of base pairs. At appropriately chosen temperatures, these tiles will have a good chance of assembling according to these prescribed rules, although they also have a chance of disassociating and breaking apart. Majumder, Sahu and Reif [14] consider a DNA self-assembly model that allows the left column and bottom row of a large square to first form, and then iteratively allows tiles to associate with the large substrate if their left and bottom neighbors are already present (see Figure 1a). Likewise, tiles can disassociate if their upper and right neighbors are not present, although disassociation happens at a lower rate. The dynamics of this model are precisely captured

by the local Markov chain \mathcal{M}_B on 2-dimensional monotonic surfaces and the chain must be rapidly mixing if the substrate is to efficiently self-assemble, as required. The 3-dimensional analogue has also been used to study self-assembly, where now tiles are shaped like cubes (as in Figure 1b) and complementary sequences are used to encourage specified pairs of faces to attach. The problem of generating biased surfaces in two-dimensions was independently studied in the context of biased card shuffling, where we allow nearest-neighbor transpositions but favor putting each pair in order at each step [2].

Previous work has focused primarily on the case where the biases are *uniform*; that is, $\lambda_{\bar{x}} = \lambda$ for every \bar{x} , for which the stationary probability will be proportional to $\lambda^{|\sigma|}$, where $|\sigma|$ is the number of unit cubes lying below the surface σ . In two dimensions, the uniform bias Markov chain is equivalent to an asymmetric exclusion process, which Benjamini, Berger, Hoffman, and Mossel [2] studied in order to analyze a biased card shuffling algorithm that favors putting each pair of cards in the lexicographically correct order. They give a bound of $\Theta(n^2)$ on the mixing rate of the biased chain on $h \times w$ regions of \mathbb{Z}^2 (where $h + w = n$) for any uniform bias $\lambda > 1$ that is a constant. The bounds are optimal when $h = w = n/2$. The recent work of Levin and Peres [12] extends our current results to the case when the bias has a dependence on n , the size of the lattice. In addition, for the three dimensional variant, Caputo, Martinelli and Toninelli [8] introduced different methods to prove that the biased chain mixes in time $\tilde{O}(n^3)$ for any constant bias λ , where the \tilde{O} notation suppresses logarithmic factors. In dimensions $d > 3$ almost nothing is known, in either the biased or unbiased settings.

1.1 Our results

We make progress in several aspects of the problem of sampling biased surfaces. In two dimensions we show that the biased chain is rapidly mixing for any uniform bias on a large family of simply-connected regions, even when the bias is arbitrarily close to one. Our proof is significantly simpler than the arguments of [2], while achieving the same optimal bounds on the mixing time for square regions when the bias is constant. In fact, we get optimal bounds for all rectangular $h \times w$ regions. We also show the chain is rapidly mixing for a large family of simply-connected d -dimensional lattice regions provided the bias $\lambda \geq d^2$. The key observation underlying these results is that there is an exponential metric on the state space such that the distance between pairs of configurations is always decreasing in expectation. We then show how to modify the path coupling theorem to handle the case when the distances are exponentially large and the expected change in distance is small during moves of the coupled chain. Previously, Berger, Kenyon, Mossel, and Peres [3] also used an exponential metric in the context of Glauber dynamics on trees, appealing to the multiplicative version of the path coupling theorem originally given by Buble and Dyer [5]. Our version of the path coupling theorem makes explicit when we can use exponential metrics to bound convergence times. We believe that this new theorem is of independent interest, and it has already been used, for example, in the context of sampling lattice triangulations [7] and rectangular dissections [6].

Last, we consider “fluctuating bias,” where the rate at which we add or remove a cube depends on its location. This setting is quite natural for the self-assembly growth process where the tile at a particular location may have site-specific sequences along its bounding edges. We show that the fast convergence results still hold whenever the biases are close to uniform over locations. In fact, in two dimensions, we show that the chain mixes in time $O(n^2)$ even if the biases are not close together, as long as all of the biases are bounded away from one by a constant. However, in the general setting we may see very different behavior. We construct an example where the convergence rate requires exponential time starting at any initial configuration. In this example every move occurs with at least inverse polynomial probability and at each location it is at least as

likely to add a cube as to remove. This demonstrates that the behavior of these growth processes is quite complicated in the case of fluctuating bias.

The remainder of the paper is organized as follows. In Section 3 we review the path coupling method and introduce the modified path coupling theorem that is more appropriate when distances are exponentially large. In Section 4, we formalize the model and Markov chain and show how to bound the mixing time of the chain. In Section 5, we generalize these techniques to apply to the setting of fluctuating bias. Finally, in the last section we discuss other related problems, including sampling biased 3-colorings in \mathbb{Z}^d .

2 The Markov chain

We can now formalize our definition of monotonic surfaces and the Markov chain that makes local updates to these surfaces. Throughout the majority of the paper we focus on surfaces arising in the context of staircase walks, lozenge tilings, and higher dimensional surfaces that bound downsets of unit boxes in \mathbb{Z}^d . We discuss the generalizations to other families of monotone surfaces such as those arising from 3-colorings in \mathbb{Z}^d in the conclusions.

We start by first considering monotonic surfaces forming over simple $h \times w$ regions R in \mathbb{Z}^2 . Later we will show that we can generalize these results to a family of simply-connected regions in \mathbb{Z}^d that we call *nice*. The generalization is straightforward, but requires some careful notation, so for simplicity we postpone the details until Section 4.2. Recall that in \mathbb{Z}^2 , a *monotonic surface* in R is a path starting and ending on the boundary of R that takes precisely h steps down and w steps to the right and is composed entirely of edges with both endpoints in R . Such a path is illustrated in Figure 1a when R is a 4×4 square. Notice that any monotonic surface can be interpreted as the upper boundary of a set of unit squares (which we call a downset), where each square in the set is supported below and to the left by other squares in the set or the boundary of R . Let M be the bijection between downsets σ and their corresponding monotonic surfaces $M(\sigma)$. We let the state space Ω_{mon} be the set of all downsets of R .

Now we can describe the Markov chain \mathcal{M}_B on Ω_{mon} . For simplicity, we start by defining the *unbiased* chain \mathcal{M}_U that converges to the *uniform distribution* over monotonic paths Ω_{mon} . Assume without loss of generality that $h \leq w$. As there are typically many fewer allowable moves than the volume of R , we concentrate the selections to places where there is a possible move (in this case, to the cubes to the left or right of all the “down” edges). Start at an arbitrary downset, e.g., let $\sigma_0 = R_L$, where R_L is the empty downset, and repeat the following steps. If we are at a downset σ_t at time t , pick an edge D from among all the “down” edges. Also, pick an integer $b \in \pm 1$ uniformly at random. If $b = +1$, add the cube to the right of the edge D to create σ_{t+1} , if this is a valid downset. If $b = -1$, let σ_{t+1} be obtained from σ_t by removing the cube to the left of D if this is a valid downset. In all other cases, keep σ_t unchanged so that $\sigma_{t+1} = \sigma_t$.

For technical reasons, we analyze a slower version \mathcal{M}'_U of this chain, which we describe next. Start at an arbitrary downset, and repeat the following steps. If we are at a downset σ_t at time t , pick a diagonal D that is parallel to the vector $\bar{u}^* = (1, 1)$ and which intersects the monotonic path at a vertex v . Also, pick an integer $b \in \pm 1$ uniformly at random. If $b = +1$, add the cube above and to the right of the vertex v to create σ_{t+1} , if this is a valid downset. If $b = -1$, let σ_{t+1} be obtained from σ_t by removing the cube below and to the left of v if this is a valid downset. In all other cases, keep σ_t unchanged so that $\sigma_{t+1} = \sigma_t$.

The new Markov chain \mathcal{M}'_U does not add any valid transitions in comparison with \mathcal{M}_U , so the effect of this change is to merely add an extra self-loop probability of a factor of $(h + w)/h$ (i.e. the number of diagonals divided by the number of down edges). Thus, the mixing time of \mathcal{M}'_U is

$(h + w)/h$ times the mixing time of \mathcal{M}_U . This can be a significant difference for very asymmetric regions.

Lemma 1. *For any rectangular region R , the chains \mathcal{M}_U and \mathcal{M}'_U connect the state space Ω_{mon} .*

Proof. Let σ be any downset and let x^{max} be any cube in σ such that $\sum_i x_i^{max}$ is maximized, if it exists. We can always remove x^{max} and move to $\sigma' = \sigma \setminus x^{max}$ without violating the downset condition. Thus, from any valid downset σ we can always remove points and get to the “lowest” downset R_L . Also, such a sequence of steps can be reversed to move from R_L to any other downset ρ . \square

Since we have shown that the moves of \mathcal{M}_U and \mathcal{M}'_U each connect the state space and all valid (non-lazy) moves have the same transition probabilities, we can conclude from detailed balance that the chains converge to the uniform distribution over downsets Ω_{mon} .

We now define the *biased* Markov chain \mathcal{M}'_B by using Metropolis-Hastings transition probabilities so that we converge to the desired distribution on biased surfaces. This new chain connects the state space by the same argument as in Lemma 1. As in the uniform setting, \mathcal{M}_B is defined similarly, but chooses a down edge instead of a diagonal at each step.

The Biased Markov chain \mathcal{M}'_B

Starting at any σ_0 , iterate the following:

- Choose a diagonal D and a bit $b \in \{\pm 1\}$ as described above, with v the point of intersection of d with the path σ_t .
- If $b = +1$, add the cube above and to the right of the vertex v to create σ_{t+1} , if this is a valid downset.
- If $b = -1$, then with prob. $\frac{1}{\lambda_x}$ let σ_{t+1} be obtained from σ_t by removing the cube x below and to the left of v if this is a valid downset.
- Otherwise let $\sigma_{t+1} = \sigma_t$.

The biased Markov chains \mathcal{M}_B and \mathcal{M}'_B converge to the correct distribution on Ω_{mon} by the detailed balance condition.

3 Path coupling with exponential metrics

Path coupling is a standard technique used to bound mixing times, and although a naive application of it is not sufficient here, we will see that with some new ideas, we can make it work. One of the innovations behind our proofs is to introduce a new metric, and in some cases this metric requires a modified Path Coupling theorem. We present the background and our new theorem here.

A *coupling* of a chain \mathcal{M} is a Markov process on $\Omega \times \Omega$ such that the marginals each agree with \mathcal{M} and, once the two coordinates coalesce, they move in unison thereafter. The Coupling Lemma bounds the total variation distance by the probability that the processes have coalesced (see, for example, [1]):

Theorem 2. *For any coupling, we have $d_{tv}(P^t(x, \cdot), \pi) \leq P(X_t \neq Y_t)$.*

Definition 3. For initial states x, y let

$$T^{x,y} = \min\{t : X_t = Y_t \mid X_0 = x, Y_0 = y\},$$

and define the coupling time to be $T = \max_{x,y} \mathbb{E}[T^{x,y}]$.

The following lemma bounds the mixing time in terms of the coupling time of any coupling (see, for example, [1]):

Theorem 4. For any coupling with coupling time T , the mixing time satisfies $\tau(\epsilon) \leq \lceil Te \ln \epsilon^{-1} \rceil$.

The goal, then, is to define a good coupling and show that the coupling time is polynomially bounded. Path Coupling is a convenient way of establishing this property by only considering a subset of the joint state space $\Omega \times \Omega$.

Theorem 5. (Dyer and Greenhill [9]) Let φ be an integer valued metric defined on $\Omega \times \Omega$ which takes values in $\{0, \dots, B\}$. Let U be a subset of $\Omega \times \Omega$ such that for all $(x_t, y_t) \in \Omega \times \Omega$ there exists a path $x_t = z_0, z_1, \dots, z_r = y_t$ between x_t and y_t such that $(z_i, z_{i+1}) \in U$ for $0 \leq i < r$ and $\sum_{i=0}^{r-1} \varphi(z_i, z_{i+1}) = \varphi(x_t, y_t)$. Let \mathcal{M} be a Markov chain on Ω with transition matrix P . Consider any random function $f : \Omega \rightarrow \Omega$ such that $\mathbb{P}[f(x) = y] = P(x, y)$ for all $x, y \in \Omega$, and define a coupling of the Markov chain by $(x_t, y_t) \rightarrow (x_{t+1}, y_{t+1}) = (f(x_t), f(y_t))$. Suppose there exists $\beta \leq 1$ such that

$$\mathbb{E}[\varphi(x_{t+1}, y_{t+1})] \leq \beta \varphi(x_t, y_t),$$

for all $(x_t, y_t) \in U$.

1. If $\beta < 1$, then the mixing time satisfies

$$\tau(\epsilon) \leq \frac{\ln(B\epsilon^{-1})}{1 - \beta}.$$

2. If $\beta = 1$ (i.e., $\mathbb{E}[\Delta\varphi(x_t, y_t)] = \mathbb{E}[\varphi(x_{t+1}, y_{t+1})] - \mathbb{E}[\varphi(x_t, y_t)] \leq 0$, for all $x_t, y_t \in U$), let $\alpha > 0$ satisfy $\Pr[\varphi(x_{t+1}, y_{t+1}) \neq \varphi(x_t, y_t)] \geq \alpha$ for all t such that $x_t \neq y_t$. The mixing time of \mathcal{M} then satisfies

$$\tau(\epsilon) \leq \left\lceil \frac{eB^2}{\alpha} \right\rceil \lceil \ln \epsilon^{-1} \rceil.$$

To understand why it is difficult to use coupling to prove that \mathcal{M}'_B is rapidly mixing, we first examine the straightforward coupling of (σ_t, ρ_t) in the uniform bias case. The natural coupling simply chooses the same diagonal D and bit b to generate both σ_{t+1} and ρ_{t+1} . We first consider a natural distance metric on $\Omega_{mon} \times \Omega_{mon}$ called the Hamming distance, where $h(\sigma_t, \rho_t) = |\sigma_t \oplus \rho_t|$ (and \oplus is the symmetric difference). However, with this coupling and metric, we face difficulty with even the simplest of pairs (σ_t, ρ_t) .

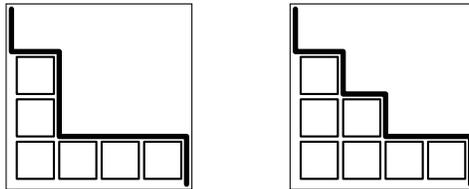


Figure 2: A pair of downsets σ_t (left) and ρ_t (right) where $\rho_t = \sigma_t \cup \{(2, 2)\}$.

Examine the pair of downsets in Figure 2. They differ on a single point, so $h(\sigma_t, \rho_t) = 1$. In order to use the coupling theorem above, the expected distance $\mathbb{E}[h(\sigma_{t+1}, \rho_{t+1})]$ must be at most $h(\sigma_t, \rho_t)$. For this pair of downsets, there are two moves that decrease that distance; if \mathcal{M}'_B chooses the diagonal $D_0 = \{(0, 0) + t\bar{u}^* : t \geq 0\}$ and either $b = +1$ or $b = -1$, then $(\sigma_{t+1}, \rho_{t+1})$ is (ρ_t, ρ_t) or (σ_t, σ_t) , respectively. In either case, the distance between σ_{t+1} and ρ_{t+1} decreases by 1. There are also two moves that *increase* the distance. If \mathcal{M}'_B chooses $D = \{(1, 0) + t\bar{u}^* : t \geq 0\}$ and $b = +1$ or $D = \{(0, 1) + t\bar{u}^* : t \geq 0\}$ and $b = +1$, then ρ_{t+1} gains a new point $((3, 2)$ or $(2, 3)$, respectively), but σ_{t+1} remains unchanged; no addition to σ_t of a vector along that diagonal leaves a valid downset. With either of these choices, the distance between σ_{t+1} and ρ_{t+1} increases by 1. If $\lambda = 1$, this is sufficient for coupling; the expected change in distance is 0. Unfortunately, for any $\lambda > 1$, the two bad moves happen with probability 1, whereas the two good moves happen with probability $1/\lambda$ and $1/\lambda$, respectively. Therefore the expected distance between the pair (σ_t, ρ_t) *increases* after one step. In higher dimensions, the situation becomes even worse. For the pair of 3 dimensional downsets in Figure 6, there are *three* moves which increase the Hamming distance and only two moves which decrease the distance. Of course, the three moves that increase the distance succeed with probability 1, but one of the two moves which decreases the distance only succeeds with probability $1/\lambda$.

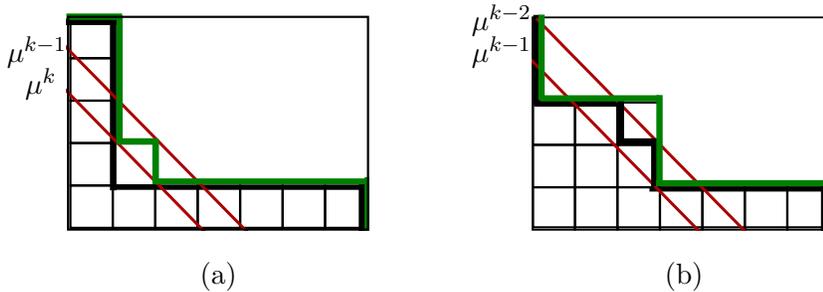


Figure 3: Path coupling with an exponential distance metric.

One promising remedy is to alter the distance metric. The bad cases described above involve two downsets that differ on some point x , where the two moves which decrease the distance involve removing x from $\sigma_t \oplus \rho_t$, while the moves that increase the distance involve adding $x + \bar{u}_i$ to $\sigma_t \oplus \rho_t$ for some i (where $\bar{u}_1 = (1, 0)$ and $\bar{u}_2 = (0, 1)$). Since the bad moves happen with greater probability than the good moves, we consider a distance metric that counts the distance between two sets that differ on x as greater than the distance between two sets that differ on $x + \bar{u}_i$. Specifically, we give a different weight to each Northwest-Southeast diagonal, with the weights smaller along higher diagonals (as in Figure 3a). This allows us to make the change in distance nonpositive in the above cases. Of course we must ensure that the difference in weight is not too great. This is because the opposite situation might happen as well, where the two bad moves involve removing $x - \bar{u}_i$ for some i , (as in Figure 3b); although this situation was not a problem for the Hamming distance metric, if we assign too much weight to those bad moves, the change in distance might be positive in this case. We find the following distance metric suffices. First, let R be the $h \times w$ rectangle in \mathbb{Z}^2 and define $\mu = \sqrt{\lambda} \geq 1$. Then for two downsets σ, ρ in R , let

$$\phi(\sigma, \rho) = \mu^{w+h} \sum_{x \in \sigma \oplus \rho} \mu^{-\|x\|_1}, \quad (1)$$

where $\|\cdot\|_1$ is the L_1 norm. Notice that all elements on each Northwest-Southeast diagonal have the same L_1 norm, and so this metric assigns a weight of μ^k for the k th diagonal from the top right,

as in Figure 4. Notice also that this definition ensures that the distance between any two downsets is either 0 or at least 1. We present the proof that this metric is decreasing in expectation at every step in Section 4.1.

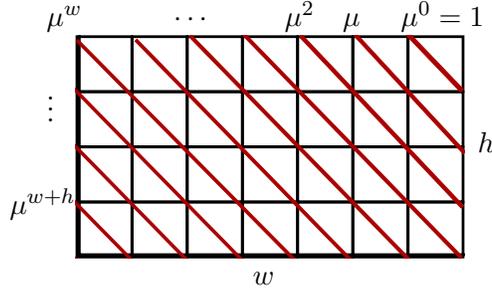


Figure 4: Exponential distance metric.

Unfortunately, this definition of the distance metric presents new problems. First, the distances might now take on non-integer values, while the Path Coupling Theorem requires integer valued metrics. In fact, if this restriction is merely removed, then the theorem is no longer true as the distances might get smaller and smaller without coalescence occurring in a polynomial number of steps. However, it is enough to add the additional condition that no pairs of configurations can have a distance within the open interval $(0, 1)$. The second, more serious concern is that the maximum distance between two configurations can be exponentially large in n . If the distance only changes by a small (polynomial) amount in each step, then we cannot expect the distance to be zero in only a polynomial number of steps. For example, for small λ we can find configurations x_t and y_t so that $\mathbb{E}[\varphi(x_{t+1}, y_{t+1})] \geq (1 - 2^{-n})\varphi(x_t, y_t)$, so the expected change is too small to apply the first part of Theorem 5. Moreover, the maximum distance B is very large, so we cannot get a good bound on the mixing time using the second part of Theorem 5 either.

The following modification of the Path Coupling Theorem allows us to handle cases when the distances can be exponentially large and the expected change in distance is small (or even zero). We show that it suffices to prove that the expected change in the absolute value of the distance is proportional to the current distance, and with this condition the mixing time is polynomially bounded. We apply this new theorem to the biased Markov chain \mathcal{M}'_B in Section 4.1.

Theorem 6. *Let ϕ be a metric defined on $\Omega \times \Omega$ which takes finitely many values in $\{0\} \cup [1, B]$. Let U be a subset of $\Omega \times \Omega$ such that for all $(X_t, Y_t) \in \Omega \times \Omega$ there exists a path $X_t = Z_0, Z_1, \dots, Z_r = Y_t$ such that $(Z_i, Z_{i+1}) \in U$ for $0 \leq i < r$ and $\sum_{i=0}^{r-1} \phi(Z_i, Z_{i+1}) = \phi(X_t, Y_t)$.*

Let \mathcal{M} be a Markov chain on Ω and let (X_t, Y_t) be a coupling of \mathcal{M} , with $\phi_t = \phi(X_t, Y_t)$. Suppose there exists $\beta \leq 1$ such that, for all $(X_t, Y_t) \in U$,

$$\mathbb{E}[\phi_{t+1}] \leq \beta \phi_t.$$

1. *If $\beta < 1$, then the mixing time satisfies*

$$\tau(\varepsilon) \leq \frac{\ln(B\varepsilon^{-1})}{1 - \beta}.$$

2. *If there exists $\kappa, \eta \in (0, 1)$ such that $\mathbb{P}[|\phi_{t+1} - \phi_t| \geq \eta \phi_t] \geq \kappa$ for all t provided that $X_t \neq Y_t$, then*

$$\tau(\varepsilon) = O\left(\frac{\ln^2 B \ln \varepsilon^{-1}}{\kappa \ln^2(1 + \eta)}\right).$$

There are two important differences between Theorem 5 and Theorem 6. The first is that Theorem 6 allows for non-integer metrics (provided that for all $X, Y \in \Omega$, $\phi(X, Y) < 1$ implies $\phi(X, Y) = 0$). This is a minor restructuring of the proof of Theorem 5 [13], and follows exactly from their method. The second is that β may equal 1 while B is exponentially large; this is the case in which both parts of Theorem 5 are insufficient to prove rapid mixing. This second case also can be shown with a slight modification of the original proof, essentially replacing the original distance $\phi(X_t, Y_t)$ with $\ln(\phi(X_t, Y_t))$. There are some technical details concerning the expectation and variance of the logarithm, but the novelty of Theorem 6 is more in the statement of the result than a new method of proof.

Note that including this case of $\beta = 1$ and exponential B requires a strong bound on the variance of ϕ_t . Without this bound on variance, Theorem 6 is not true; if $\phi_0 = 2^n$ and $\phi_{t+1} = \phi_t - 1$ for all $t \geq 1$, then clearly it will take time exponential in n for $\phi_t = 0$.

In order to prove Theorem 6 permitting exponential metrics, we define a new variable ψ , which is essentially $\ln(\phi)$. However, if we hope to prove rapid mixing by looking at $\ln(\phi)$, we need to bound the time to reach $\ln(0) = -\infty$, and the expected time could be unbounded. In particular, in order to prove rapid mixing, we need that the sequence $\{\psi_t\}$ has bounded differences. The technical fix that we make relies on the assumption that $\phi_t \notin (0, 1)$, so we need only bound the time until we reach a negative value for $\ln(\phi_t)$. Hence we define

$$\psi_t = \begin{cases} \ln(\phi_t) & \text{if } \phi_t > 0 \\ -2 \ln 2 & \text{if } \phi_t = 0 \end{cases}.$$

This means that $\psi_t \in [-2 \ln 2, \ln B]$. The particular value at zero is chosen so that if the expected distance ϕ_t is non-decreasing, then the expected value of ψ_t is non-decreasing, and that if the variance of ϕ_t is at least a linear factor, then the variance of ψ_t is at least a constant.

The following Martingale Lemma follows the proof of Lemma 6 in [13].

Lemma 7. *Given any bounded function $\phi(t)$, with $\phi(t) \leq M$ for some $M \in \mathbb{R}$ and for all $t \geq 0$, and a stopping value q , let $T = \min\{t : \phi(t) \leq q\}$. If, for all $t \geq 0$, we have $\mathbb{E}[\phi(t+1) - \phi(t)] \leq 0$ and $\mathbb{E}[(\phi(t+1) - \phi(t))^2] \geq Q$ for some $Q > 0$, then*

$$\mathbb{E}[T] \leq \frac{(M - q)^2}{Q}.$$

Proof. Define the process $Z(t) := (M - \phi(t))^2 - Qt$. Examining the expected difference between $Z(t)$ and $Z(t+1)$, we have

$$\begin{aligned} \mathbb{E}[Z(t+1) - Z(t)] &= \mathbb{E}[(M - \phi(t+1))^2 - (M - \phi(t))^2] - Q \\ &= \mathbb{E}[-2M(\phi(t+1) - \phi(t)) + \phi(t+1)^2 - \phi(t)^2] - Q \\ &= 2(\phi(t) - M)\mathbb{E}[\phi(t+1) - \phi(t)] + \mathbb{E}[(\phi(t+1) - \phi(t))^2] - Q \geq 0. \end{aligned}$$

Also, since the differences $Z(t+1) - Z(t)$ are bounded, so $\{Z(t)\}$ is a submartingale. T is a stopping time for $Z(t)$, so we may apply the Optional Stopping Theorem for submartingales to deduce that

$$\begin{aligned} \mathbb{E}[T] &\leq \frac{1}{Q} [\phi(0)(2M - \phi(0)) + q^2 - 2qM] \\ &\leq \frac{(M - q)^2}{Q}. \end{aligned}$$

□

Now we may prove the exponential metric theorem, Theorem 6.

Proof. Part (1) This case follows directly from the proof of Theorem 5, while allowing non-integer valued metrics. Since $\mathbb{E}[\phi_{t+1}] \leq \beta\phi_t$ for all t , it follows that

$$\mathbb{E}[\phi(X_t, Y_t)] = \mathbb{E}[\phi_t] \leq \beta^t \phi_0 = \beta^t \phi(X_0, Y_0) \leq \beta^t B. \quad (2)$$

Since ϕ_t is nonnegative, takes values in $\{0\} \cup [1, B]$ and is equal to zero whenever $X_t = Y_t$, we have

$$\begin{aligned} \mathbb{E}[\phi_t] &\geq \mathbb{P}[\phi_t \geq 1] \\ &= \mathbb{P}[X_t \neq Y_t], \end{aligned}$$

by Markov's inequality. Then, since by Equation 2, $\mathbb{E}[\phi_t] \leq \epsilon$ whenever $t \geq \log(B\epsilon^{-1})/\log(\beta^{-1})$, the Coupling Lemma 2 implies

$$\tau(\epsilon) \leq \ln(B\epsilon^{-1})/\ln(\beta^{-1}).$$

Since $\ln(\beta^{-1}) > 1 - \beta$, Part (1) follows.

Part (2) For part (2), we will show that ψ_t satisfies the conditions of Lemma 7, with $q = -2 \ln 2$, $M = \ln B$ and $Q = \kappa \ln^2(1 + \eta)$. Note that this proves the theorem, since Theorem 4 implies that

$$\tau(\epsilon) \leq \lceil \mathbb{E}[T] e \ln \epsilon^{-1} \rceil \leq \left\lceil \frac{(M - q)^2}{Q} e \ln \epsilon^{-1} \right\rceil = O\left(\frac{\ln^2 B \ln \epsilon^{-1}}{\kappa \ln^2(1 + \eta)}\right).$$

First we show that $\mathbb{E}[\phi_{t+1} - \phi_t] \leq 0$ implies $\mathbb{E}[\psi_{t+1} - \psi_t] \leq 0$. We may assume $\phi_t \neq 0$. Given the value of ϕ_t , let $\{r_0, r_1, r_2, \dots, r_N\}$ be the possible values for ϕ_{t+1} , each occurring with probability $\{\zeta_0, \zeta_1, \zeta_2, \dots, \zeta_N\}$. That is, $\mathbb{P}[\phi_{t+1} = r_i | \phi_t] = \zeta_i$, with $\sum_{i=0}^N \zeta_i = 1$. Assume $r_0 = 0$. Now,

$$\begin{aligned} \mathbb{E}[\psi_{t+1} | \psi_t] &= \zeta_0(-2 \ln 2) + \sum_{i=1}^N \zeta_i \ln(r_i) \\ &= -(2 \ln 2)\zeta_0 + \ln\left(\prod_{i=1}^N r_i^{\zeta_i}\right) \\ &\leq -(2 \ln 2)\zeta_0 + (1 - \zeta_0) \ln\left(\frac{\sum_{i=1}^N \zeta_i r_i}{1 - \zeta_0}\right) \\ &\leq \ln(\mathbb{E}[\phi_{t+1} | \phi_t]) - (2 \ln 2)\zeta_0 - (1 - \zeta_0) \ln(1 - \zeta_0) \\ &\leq \ln(\mathbb{E}[\phi_{t+1} | \phi_t]) \\ &\leq \ln \phi_t = \psi_t, \end{aligned}$$

where the first inequality is by the Arithmetic-Geometric Mean Inequality, and the third follows from the fact that $(1 - \zeta_0) \ln(1 - \zeta_0)/\zeta_0 \geq -2 \ln 2$ for $\zeta_0 \geq 0$.

Next we prove that if there exist constants $\kappa, \eta \in (0, 1)$ such that $\mathbb{P}[|\phi_{t+1} - \phi_t| \geq \eta\phi_t] \geq \kappa$ for $\phi_t \neq 0$, then

$$\mathbb{E}[|\psi_{t+1} - \psi_t|] \geq (\ln(1 + \eta))\kappa + \ln 2\mathbb{P}[\phi_{t+1} = 0].$$

Let $\zeta_0 = \mathbb{P}[\phi_{t+1} = 0]$. Then

$$\begin{aligned} \kappa &\leq \mathbb{P}[|\phi_{t+1} - \phi_t| \geq \eta\phi_t] \\ &= 1 \cdot \zeta_0 + \mathbb{P}[|\phi_{t+1} - \phi_t| \geq \eta\phi_t | \phi_{t+1} \neq 0](1 - \zeta_0). \end{aligned}$$

Now because, by definition, $\psi_{t+1} = \ln(\phi_{t+1})$ when $\phi_{t+1} \neq 0$, we have

$$\begin{aligned}
& \mathbb{P}[|\phi_{t+1} - \phi_t| \geq \eta\phi_t | \phi_{t+1} \neq 0] \\
&= \mathbb{P}\left[\frac{\phi_{t+1}}{\phi_t} - 1 \geq \eta | \phi_{t+1} \neq 0\right] + \mathbb{P}\left[\frac{\phi_{t+1}}{\phi_t} - 1 \leq -\eta | \phi_{t+1} \neq 0\right] \\
&= \mathbb{P}[\psi_{t+1} - \psi_t \geq \ln(1 + \eta) | \phi_{t+1} \neq 0] + \mathbb{P}[\psi_{t+1} - \psi_t \leq \ln(1 - \eta) | \phi_{t+1} \neq 0] \\
&\leq \mathbb{P}[|\psi_{t+1} - \psi_t| \geq \ln(1 + \eta) | \phi_{t+1} \neq 0].
\end{aligned}$$

Since $\phi_t \geq 1$, we have $\psi_t \geq 0$, so $|-2\ln 2 - \psi_t| \geq 2\ln 2$. Note that $m := \ln^2(1 + \eta) < \ln^2(2)$, since $\eta < 1$. This yields

$$\begin{aligned}
\mathbb{E}[(\psi_{t+1} - \psi_t)^2] &= (-2\ln 2 - \psi_t)^2 \zeta_0 + \sum_{\ell \in \Omega, \ell \neq 0} (\ln(\ell) - \ln \phi_t)^2 \mathbb{P}[\phi_{t+1} = \ell] \\
&\geq (2\ln 2)^2 \zeta_0 + m \mathbb{P}[|\psi_{t+1} - \psi_t| \geq \sqrt{m} | \phi_{t+1} \neq 0] (1 - \zeta_0) \\
&\geq (2\ln 2)^2 \zeta_0 + m \left(\frac{\kappa - \zeta_0}{1 - \zeta_0}\right) (1 - \zeta_0) \\
&= ((2\ln 2)^2 - m) \zeta_0 + m\kappa \\
&\geq m\kappa + 3\zeta_0 \ln^2 2.
\end{aligned}$$

Hence we have $Q = \kappa \ln^2(1 + \eta) < \mathbb{E}[(\psi_{t+1} - \psi_t)^2]$, as desired. □

4 Fast mixing of the uniform bias Markov chain

We start by looking at the biased Markov chain \mathcal{M}_B when the biases are uniform. We use the exponential metric introduced in Section 3 to show that \mathcal{M}_B is rapidly mixing whenever $\lambda \geq d^2$ for arbitrary dimension d , and for all $\lambda > 1$ when $d = 2$. In Section 4.3, we present a simple hitting time argument that proves the biased chain converges in polynomial time, as long as the minimum bias is at least d and the region is a d -dimensional hypercube. We conjecture that the chain is rapidly mixing for all values of a uniform bias $\lambda > 1$ in all dimensions $d \geq 2$, but do not yet have a proof for small values of λ in dimensions higher than 2.

4.1 Exponential metric for the uniform bias chain

First, we use our exponential distance metric to analyze \mathcal{M}'_B . We show that in two dimensions, the biased chain is rapidly mixing for any uniform bias, even when the bias is arbitrarily close to one. Our proof is significantly simpler than the arguments of Benjamini et al., while achieving the same optimal bounds on the mixing time for square regions when the bias is constant. In fact, our analysis extends the optimal bounds to all rectangular $h \times w$ regions of \mathbb{Z}^2 . Specifically, we prove the following theorem for rectangular regions:

Theorem 8. *Let R be a rectangular $h \times w$ region in \mathbb{Z}^2 with uniform bias $\lambda \geq 1$. Let $\chi = \lambda^{-1/2} - 1$. Then*

1. *If $\chi > 0$, then the mixing time of \mathcal{M}_B satisfies $\tau(\varepsilon) = O((hw/\chi^2)(\ln \lambda + \ln \varepsilon^{-1}))$.*
2. *If $\chi \geq 0$, then $\tau(\varepsilon) = O(hw^4 \ln^2 \lambda \ln \varepsilon^{-1})$.*

If $\lambda > 1$ is a constant, then part 1 of this theorem applies and gives a bound of $O(n)$, where $n = hw$ is the area of the region. On the other hand, for λ very close to 1, part 2 provides a polynomial bound on the mixing time. We prove similar bounds for all *nice* regions, to be defined in Section 4.2; essentially, these regions are simply-connected and have no holes. In higher dimensions, we show the chain is rapidly mixing on d -dimensional lattice regions provided the bias $\lambda \geq d^2$.

For clarity of explanation, we will first handle the simple case when R is a rectangular region in \mathbb{Z}^2 . These ideas generalize easily to more complex regions and higher dimensions, but require some extra terminology, which we define in Section 4.2.

4.1.1 Rectangular regions in two dimensions

We return to the coupling introduced in Section 3 to bound the mixing time of \mathcal{M}'_B , proving Theorem 8. The coupling of (σ_t, ρ_t) simply supplies the same diagonal D and bit b to both σ_t and ρ_t . We let U be the set of pairs of downsets that differ on a single cube. However, instead of the Hamming distance, we use the distance metric given in equation 1:

$$\phi(\sigma, \rho) = \mu^{w+h} \sum_{x \in \sigma \oplus \rho} \mu^{-\|x\|_1},$$

where $\mu = \sqrt{\lambda} \geq 1$. We will show that this distance metric satisfies non-negative contraction in ϕ_t , which is one of the requirements for Theorem 6. However, before we can prove that the distances decrease on average, we examine the moves which can *increase* the distance.

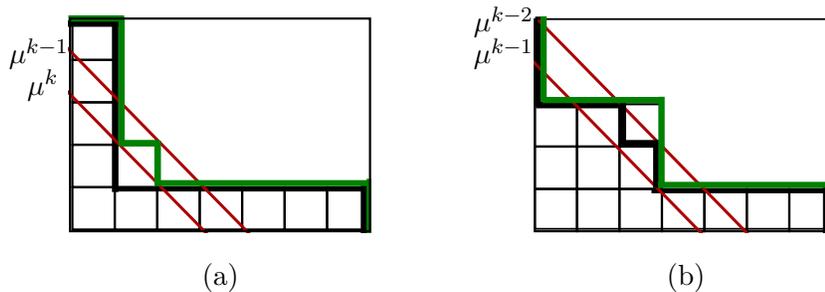


Figure 5: Path coupling with an exponential distance metric. Two cases where $\sigma_t = \rho_t \cup \{x\}$.

For a pair $(\sigma_t, \rho_t) \in U$, there are two different ways the distance can increase in $(\sigma_{t+1}, \rho_{t+1})$. If $\sigma_t = \rho_t \cup \{x\}$, we can increase the distance by attempting to add a cube v that succeeds in σ_t but fails in ρ_t , as in Figure 5a. This occurs when $v = x + \bar{u}_i$ for some i , so v is “supported” in σ_t but not ρ_t . Notice that the distance metric ϕ gives these bad moves weight that is μ times smaller than the weight of the two good moves (adding or removing x), counteracting their higher probability. The other way to increase the distance between σ_t and ρ_t is to remove a v that succeeds in ρ_t but not in σ_t . This occurs when $v = x - \bar{u}_i$ for some i , as the move creates a valid downset in ρ_t but not in σ_t , as in Figure 5b. In this case, the distance metric ϕ gives these bad moves weight that is μ times larger than the weight of the two good moves, but for small enough μ , they will still not outweigh the good moves, since the bad moves are less likely to occur than the good moves in this case. We may now prove Theorem 8.

Proof. Suppose without loss of generality that $h \leq w$. As described in Section 2, it is sufficient to analyze \mathcal{M}'_B . We will show that the distance metric ϕ defined above satisfies the conditions of the Path Coupling Theorem for Exponential metrics, Theorem 6. First we want to show that the expected change in ϕ is negative. There are at most 2 choices of (D, b) that can increase ϕ_t . We

claim that each of these has an expected increase of at most $\phi_t/\sqrt{\lambda}$. To see this, consider a move of the form $v = x + \bar{u}_i$ for some i . Then the increase in distance is $(\sqrt{\lambda})^{w+h-\|v\|_1} = \phi_t/\sqrt{\lambda}$. If the move is of the form $v = x - \bar{u}_i$ for some i , then the increase in distance is $(\sqrt{\lambda})^{-\|v\|_1} = \phi_t\sqrt{\lambda}$, and the probability of performing the move is $1/\lambda$. Therefore the expected increase is at most $\phi_t/\sqrt{\lambda}$.

There are also two choices of (D, b) that *decrease* ϕ_t ; corresponding to adding x and removing x . These each decrease ϕ_t by ϕ_t , and succeed with probability 1 and $1/\lambda$, respectively. Let $\alpha = h + w$; this is the number of choices of diagonals D . The expected change in distance satisfies

$$\begin{aligned} \mathbb{E}[\phi_{t+1} - \phi_t] &\leq \frac{1}{2\alpha} \left(2 \cdot \frac{\phi_t}{\sqrt{\lambda}} - \left(1 + \frac{1}{\lambda}\right) \phi_t \right) \\ &= \frac{\phi_t}{2\alpha} \left(\frac{2}{\sqrt{\lambda}} - 1 - \frac{1}{\lambda} \right) \\ &\leq -\frac{\phi_t \chi^2}{2\alpha}. \end{aligned}$$

Next, we check the other conditions of Theorem 6. For arbitrary $\sigma, \rho \in \Omega_{mon}$, if $x \in \sigma \oplus \rho$ for some x , then $\phi(\sigma, \rho) \geq \sqrt{\lambda}^{h+w-\|x\|_1} \geq 1$. Therefore if $\phi(\sigma, \rho) < 1$, $\phi(\sigma, \rho) = 0$. Recall that U is the set of pairs of downsets that differ on a single vector. For arbitrary $\sigma, \rho \in \Omega_{mon}$, we can connect σ to ρ by simply adding or removing the vectors in $\sigma \oplus \rho$ one by one, and $\phi(\sigma, \rho)$ is the sum of the distances. Since the volume of R is n , there are at most n possible cubes in $\sigma \oplus \rho$. Each has weight at most μ^α , so the intermediate distances each contribute at most this amount to the sum. Therefore, $\phi(\sigma, \rho) \leq n\lambda^{\alpha/2}$ for all σ, ρ . Thus, the B in Theorem 6 satisfies $B \leq n\lambda^{\alpha/2}$, and $\ln B = O(\ln n + \alpha \ln \lambda) = O(w \ln \lambda)$.

We consider two cases. If $\chi > 0$, then $\mathbb{E}[\phi_{t+1}] \leq \beta\phi_t$, where $\beta = 1 - \chi^2/\alpha$. Thus, by Theorem 6, we have that the mixing time of \mathcal{M}'_B is $\tau'(\epsilon) = O(\chi^{-2}\alpha(w \ln \lambda + \ln \epsilon^{-1})) = O(\alpha w/\chi^2(\ln \lambda + \ln \epsilon^{-1}))$. Multiplying by h/α , we get the mixing time of \mathcal{M}_B is $\tau(\epsilon) = O(hw/\chi^2(\ln \lambda + \ln \epsilon^{-1}))$, as desired. On the other hand, if $\lambda - 1 > 0$, but less than any constant, then we get a better result using the second part of Theorem 6. For any pair of σ, ρ , \mathcal{M}'_B can always add a vector v^* in their difference that maximizes $\|v\|_1$. This would change ϕ_t by an additive $\mu^{\alpha-\|v^*\|_1}$. Moreover, $\phi_t \leq n\mu^{\alpha-\|v^*\|_1}$, so the change in ϕ_t is at least ϕ_t/n . The appropriate v^* is chosen with probability at least $1/\alpha$ and the appropriate b is chosen with probability $1/2$ (and every p succeeds when adding). Therefore there is a $1/(2\alpha)$ chance of changing ϕ_t by $\phi_t/(2\alpha)$, i.e. $P(|\phi_{t+1} - \phi_t| \geq \phi_t/(2\alpha)) \geq 1/(2\alpha)$. Hence in this case, the mixing time of \mathcal{M}'_B satisfies

$$\tau'(\epsilon) = O\left(\frac{\ln^2(B)}{\frac{1}{2\alpha} \ln^2(1 + \frac{1}{\alpha})} \ln \epsilon^{-1}\right) = O(\alpha^3(w \ln \lambda)^2 \ln \epsilon^{-1}) = O(w^5 \ln^2 \lambda \ln \epsilon^{-1}),$$

and therefore the mixing time of \mathcal{M}_B is $\tau(\epsilon) = O(hw^4 \ln^2 \lambda \ln \epsilon^{-1})$. \square

4.2 More complex regions and higher dimensions

The results of Section 4.1 extend to higher dimensions and more general regions. We need to be a bit more formal, so we begin with several definitions.

Given a simply-connected region \widehat{R} in \mathbb{R}^d that is a union of unit cubes on the integer lattice, we associate to it a point set $R = \widehat{R} \cap \mathbb{Z}^d$. In this case we say that R is simply-connected. In three dimensions, a monotonic surface in R is the union of two-dimensional faces such that any cross-section along an axis-aligned plane is a two-dimensional monotonic surface. Such a surface is illustrated in Figure 1b when R is a $2 \times 2 \times 2$ region. In general, given a hypercubic region $R \subset \mathbb{Z}^d$

composed of unit hyper-cubes, a d -dimensional monotonic surface is a set of $(d - 1)$ -dimensional faces, such that any cross-section along an axis-aligned $(d - 1)$ -dimensional hyper-plane is a $(d - 1)$ -dimensional monotonic surface.

We will restrict our focus to a family of simply-connected regions that have favorable properties for the purposes of sampling monotonic surfaces. In order to define this family, we need a few preliminaries.

Definition 9. Let $\bar{u}^* = (1, 1, \dots, 1) \in \mathbb{Z}^d$. For $v \in \mathbb{Z}^d$, we define $\hat{r}(v) = \{v + k\bar{u}^* : k \in \mathbb{R}\}$, and we define the ray $r(v)$ to be the set

$$r(v) = \hat{r}(v) \cap \mathbb{Z}^d.$$

Definition 10. A d -dimensional simply-connected region $\hat{R} \subset \mathbb{R}^d$ is nice if, for all $v \in \hat{R}$, $\hat{R} \cap \hat{r}(v)$ is connected. We call its associated point set R a nice region.

A nice region \hat{R} has no holes, and in particular all monotonic surfaces in \hat{R} are the upper boundary of a subset of cubes in \hat{R} . Note that all hyper-rectangular regions are nice.

Let \bar{u}_i be the unit vector in the i th direction. Given a nice region R , we let $R_L = \{v \in R \text{ such that } v - \bar{u}^* \notin R\}$ be the lower envelope of the region.

Definition 11. Let $R \subset \mathbb{Z}^d$ be a nice region. A downset is a subset $\sigma \subseteq R$, with $R_L \subseteq \sigma$, such that for any i , if $v \in \sigma$ and $v - \bar{u}_i \in R$, then $v - \bar{u}_i \in \sigma$.

For a nice region R , we define the state space Ω_{mon} to be the set of all downsets of R . We will represent a downset by its *upper boundary*:

Definition 12. Let R be any nice region and let σ be any downset of R . We say the upper boundary of σ is $\partial(\sigma) = \{v \in \sigma \text{ such that } v + \bar{u}^* \notin \sigma\}$.

The upper boundary is a monotonic surface in bijection with the downset that defines it. It is important to notice that for any downset σ and point $v \notin \sigma$, if $\sigma \cup v$ is a valid downset, then $|\partial(\sigma)| = |\partial(\sigma \cup v)|$. This is because $\partial(\sigma \cup v) = \partial(\sigma) \cup \{v\} \setminus \{v - \bar{u}^*\}$. It follows that for any nice region R , the size of the boundary of a valid downset is fixed. This observation motivates the following definition that will be convenient when we state the mixing time of our Markov chain.

Definition 13. The span of a nice region R is $\alpha = |\partial(\sigma)|$, for any downset σ of R .

Suppose, for example, that R is an $h \times \dots \times h$ region in \mathbb{Z}^d . Then the span is $\alpha = dh^{d-1}$.

The Markov chain \mathcal{M}'_B operates as follows. Starting at an arbitrary downset, e.g., let $\sigma_0 = R_L$, where R_L is the empty downset, and repeat the following steps. If we are at a downset σ_t at time t , pick a point $v \in \partial(\sigma)$ and an integer $b \in \pm 1$ uniformly at random. If $b = +1$, let $\sigma_{t+1} = \sigma_t \cup (v + \bar{u}^*)$ if this is a valid downset. If $b = -1$, let $\sigma_{t+1} = \sigma_t \setminus \{v\}$ if this is a valid downset, however in the biased case we will do this move with the appropriate Metropolis-Hastings probabilities. In all other cases, keep σ_t unchanged so that $\sigma_{t+1} = \sigma_t$.

The Biased Markov chain \mathcal{M}'_B

Starting at any σ_0 , iterate the following:

- Choose (v, b, p) uniformly at random from $\partial(\sigma_t) \times \{+1, -1\} \times (0, 1)$.
- If $b = +1$, let $\sigma_{t+1} = \sigma \cup \{v + \bar{u}^*\}$ if it is a valid downset.

- If $b = -1$ and $p \leq \lambda_v^{-1}$, let $\sigma_{t+1} = \sigma \setminus \{v\}$ if it is a valid downset.
- Otherwise let $\sigma_{t+1} = \sigma_t$.

Observe that the chain \mathcal{M}'_B connects the state space. To see this, let σ be any downset and let v^{max} be any point in σ such that $\sum_i v_i^{max}$ is maximized. We can always remove v^{max} and move to $\sigma' = \sigma \setminus v^{max}$ without violating the downset condition. Thus, from any valid downset σ we can always remove points and get to the “lowest” downset R_L . Also, such a sequence of steps can be reversed to move from R_L to any other downset ρ . Moreover, the biased Markov chains \mathcal{M}_B and \mathcal{M}'_B converge to the correct distribution on Ω_{mon} by the detailed balance condition.

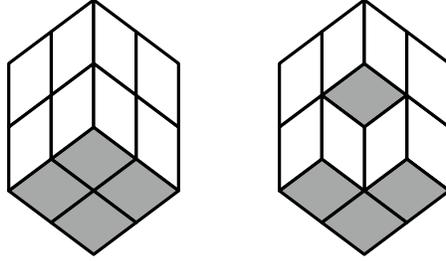


Figure 6: A pair of downsets σ_t (left) and ρ_t (right) where $\rho_t = \sigma_t \cup \{(0, 0, 0)\}$. These downsets differ on $x = (0, 0, 0)$, where \mathcal{M}_B can increase ϕ_t by adding $x + \bar{u}_i$, for any i .

The following bounds the mixing rate of the chain \mathcal{M}_B when the bias is large enough.

Theorem 14. *Let R be any nice d -dimensional region with volume n , span α , and uniform bias λ . Let $\hat{\lambda} = (d + \sqrt{d^2 - 4})/2$. Define $\chi = \hat{\lambda} - \lambda^{-1/2}$.*

1. If $\chi > 0$, then the mixing time of \mathcal{M}'_B satisfies $\tau'(\epsilon) = O(\chi^{-2}\alpha(\alpha \ln \lambda + \ln n + \ln \epsilon^{-1}))$.
2. If $\chi \geq 0$, then $\tau(\epsilon) = O(n^2\alpha(\alpha \ln \lambda + \ln n)^2 \ln \epsilon^{-1})$.

In particular, if $\lambda \geq d^2$ then $\chi^2 \geq 1/d^2$, so $\tau(\epsilon) = O(d^2\alpha(\alpha \ln \lambda + \ln n + \ln \epsilon^{-1}))$.

Note that for all nice regions, $\alpha < n$, so the mixing time of \mathcal{M}_B is always polynomially bounded for the given biases. When R is an $h \times h \times \dots \times h$ hypercube, the span of R is $\alpha = dh^{d-1}$. Hence we get $\tau'(\epsilon) = O(d^2 h^{2(d-1)} \ln \epsilon^{-1})$ whenever $1/\chi$ is a constant.

Next we define the exponential metric we need to prove Theorem 14. Let x_0 be any vector in R with maximal L_1 norm. We return to the coupling of (σ_t, ρ_t) that simply supplies the same (v^*, b, p) to both σ_t and ρ_t . Again, we let U be the set of downsets that differ on a single cube and we use the distance metric

$$\phi(\sigma, \rho) = \sum_{x \in \sigma \oplus \rho} (\sqrt{\lambda})^{\|x_0\|_1 - \|x\|_1},$$

where x_0 is any vector in R with maximal L_1 norm. We will show that this distance metric satisfies non-negative contraction in ϕ_t , which is one of the requirements for Theorem 6. First, we examine the moves which can *increase* the distance.

For a pair $(\sigma_t, \rho_t) \in U$, there are two different ways the distance can increase in $(\sigma_{t+1}, \rho_{t+1})$. If $\sigma_t = \rho_t \cup \{x\}$, we can increase the distance by attempting to add a v that succeeds in σ_t but fails in ρ_t . This occurs when $v = x + \bar{u}_i$ for some i , so v is “supported” in σ_t but not ρ_t . The other way to increase the distance between σ_t and ρ_t is to remove a v that succeeds in ρ_t but not in σ_t . This occurs when $v = x - \bar{u}_i$ for some i , as the move creates a valid downset in ρ but not in σ . The following lemma bounds the number of such increases in distance.

Lemma 15. For $\sigma_t = \rho_t \cup \{x\}$, there are at most d choices of (v^*, b) such that ϕ_t increases.

Proof. We prove the Lemma by claiming that for dimensions $i \neq j$, if \mathcal{M}_B can increase the distance by choosing $v = x + \bar{u}_i$, then it cannot increase the distance by choosing $v = x - \bar{u}_j$. This follows from a proof by contradiction: If \mathcal{M}_B can increase the distance with $x + \bar{u}_i$, then it is because $\rho_{t+1} = \rho_t \cup \{x + \bar{u}_i\}$ is a valid downset. That means $x + \bar{u}_i - \bar{u}_j \in \rho_t$. On the other hand, if \mathcal{M}_B can increase the distance with $x - \bar{u}_j$, it is because $\sigma_{t+1} = \sigma_t \setminus \{x - \bar{u}_j\}$ is a valid downset, which is only true if $x - \bar{u}_j + \bar{u}_i \notin \sigma_t$. But this contradicts the fact $\sigma_t \oplus \rho_t = \{x\}$, justifying our claim.

This implies that, to increase the distance, \mathcal{M}_B may add vectors of the form $x + \bar{u}_i$ for various dimensions i as in Figure 6, or it may remove vectors of the form $x - \bar{u}_i$ for various dimensions i , or it may add $x + \bar{u}_i$ and remove $x - \bar{u}_i$ in a single dimension i , as in Figure 7. In each of these cases, there are at most d choices of v that increase the distance. \square

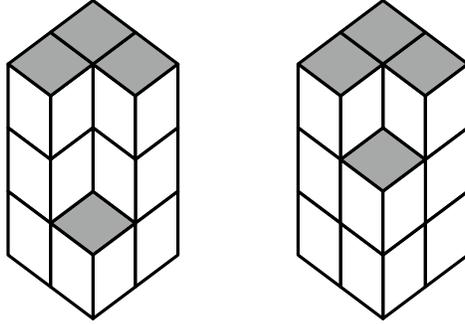


Figure 7: Downsets that differ on x , where \mathcal{M}_B increases ϕ_t by adding the vector above x or removing the vector below x .

We may now prove Theorem 14.

Proof of Theorem 14. We will show that the distance metric ϕ defined above satisfies the conditions of the Path Coupling Theorem for Exponential metrics, Theorem 6. First we want to show that the expected change in ϕ is negative. By Lemma 15, there are at most d choices of (v^*, b) that can increase ϕ_t . We claim that each of these has an expected increase of at most $\phi_t/\sqrt{\lambda}$. To see this, consider a move of the form $v = x + \bar{u}_i$ for some i . Then the increase in distance is $(\sqrt{\lambda})^{\|x_0\|_1 - \|v\|_1} = \phi_t/\sqrt{\lambda}$. If the move is of the form $v = x - \bar{u}_i$ for some i , then the increase in distance is $\lambda^{-\|v\|_1/2} = \phi_t\sqrt{\lambda}$, but the chance of choosing an appropriate p is $1/\lambda$. Therefore the expected increase is at most $\phi_t/\sqrt{\lambda}$.

There are also two choices of (v^*, b) that *decrease* ϕ_t ; corresponding to adding x and removing x . These each decrease ϕ_t by ϕ_t , and succeed with probability 1 and $1/\lambda$, respectively. Therefore the expected change in distance satisfies

$$\begin{aligned} \mathbb{E}[\phi_{t+1} - \phi_t] &\leq \frac{1}{2\alpha} \left(d \cdot \frac{\phi_t}{\sqrt{\lambda}} - \left(1 + \frac{1}{\lambda}\right) \phi_t \right) \\ &= \frac{\phi_t}{2\alpha} \left(\frac{d}{\sqrt{\lambda}} - 1 - \frac{1}{\lambda} \right) \\ &= -\frac{\phi_t}{2\alpha} \left(\frac{1}{\sqrt{\lambda}} - \frac{d + \sqrt{d^2 - 4}}{2} \right) \left(\frac{1}{\sqrt{\lambda}} - \frac{d - \sqrt{d^2 - 4}}{2} \right) \\ &\leq -\frac{\phi_t \chi^2}{2\alpha}. \end{aligned}$$

Next we check the other conditions of Theorem 6. For arbitrary $\sigma, \rho \in \Omega_{mon}$, if $x \in \sigma \oplus \rho$ for some x , then $\phi(\sigma, \rho) \geq \sqrt{\lambda}^{\|x_0\|_1 - \|x\|_1} \geq 1$. Therefore if $\phi(\sigma, \rho) < 1$, $\phi(\sigma, \rho) = 0$. Recall U is the set of pairs of downsets that differ on a single vector. For arbitrary $\sigma, \rho \in \Omega_{mon}$, we can connect σ to ρ by simply adding or removing the vectors in $\sigma \oplus \rho$ one by one, and $\phi(\sigma, \rho)$ is the sum of the distances. Since the volume of \widehat{R} is n , there are at most n possible vectors in $\sigma \oplus \rho$. Each has weight at most μ^α , so the intermediate distances each contribute at most this amount to the sum. Therefore, $\phi(\sigma, \rho) \leq n\mu^\alpha = n\lambda^{\alpha/2}$ for all σ, ρ . Thus, the B in Theorem 6 satisfies $B \leq n\lambda^{\alpha/2}$, and $\ln B = O(\ln n + \alpha \ln \lambda)$.

We consider two cases. If $\chi > 0$, then $\mathbb{E}[\phi_{t+1}] \leq \beta\phi_t$, where $\beta = 1 - \chi^2/(2\alpha)$. Thus, by Theorem 6, the mixing time of \mathcal{M}'_B satisfies $\tau'(\epsilon) = O(\chi^{-2}\alpha(\alpha \ln \lambda + \ln n + \ln \epsilon^{-1}))$. On the other hand, if $\chi \geq 0$, but less than any constant, then we use the second part of Theorem 6. For any pair of σ, ρ , \mathcal{M}_B can always add a vector v^* in their difference that maximizes $\|v\|_1$. This would change ϕ_t by $(\sqrt{\lambda})^{\alpha - \|v^*\|_1}$. As $\phi_t \leq n(\sqrt{\lambda})^{\alpha - \|v^*\|_1}$, the change in ϕ_t is at least ϕ_t/n . The appropriate v^* is chosen with probability at least $1/\alpha$ and the appropriate b is chosen with probability $1/2$ (and every p succeeds when adding). Therefore there is a $1/(2\alpha)$ chance of changing ϕ_t by $\phi_t/(2n)$, i.e. $P(|\phi_{t+1} - \phi_t| \geq \phi_t/n) \geq 1/(2\alpha)$. Hence in this case the mixing time of \mathcal{M}'_B is

$$\tau'(\epsilon) = O\left(\frac{\ln^2(B)}{\ln^2(1 + \frac{1}{n})^{\frac{1}{2\alpha}}} \ln \epsilon^{-1}\right) = O(n^2\alpha(\alpha \ln \lambda + \ln n)^2 \ln \epsilon^{-1}).$$

□

4.3 Hitting time to the maximal tiling

We now introduce a second technique that allows us to get improved bounds for the mixing rate of the uniform bias Markov chain \mathcal{M}'_B whenever $\lambda \geq d$ and the region is an $h \times h \times \dots \times h$ hypercube of volume $n = h^d$. Specifically, we prove:

Theorem 16. *Let $R \subset \mathbb{Z}^d$ be the $h \times h \times \dots \times h$ hypercube of volume $n = h^d$ and bias $\lambda \geq d$. Then the mixing time of \mathcal{M}'_B satisfies $\tau(\epsilon) = O(h^{2d-1} \ln \epsilon^{-1})$. In general, this is $o(n^2 \ln \epsilon^{-1})$, or in 2 dimensions, $O(n\sqrt{n} \ln \epsilon^{-1})$.*

The proof relies on the monotonicity of \mathcal{M}_B with respect to the trivial coupling. In other words, if (X_t, Y_t) are coupled and $X_t \subseteq Y_t$, then after one step of the coupling, $X_{t+1} \subseteq Y_{t+1}$. This implies that the coupling time is bounded by the time to hit the full cube F starting from the empty cube F' , and we can show that this will happen quickly because the distance to F is always non-increasing in expectation (see Fig. 8b).

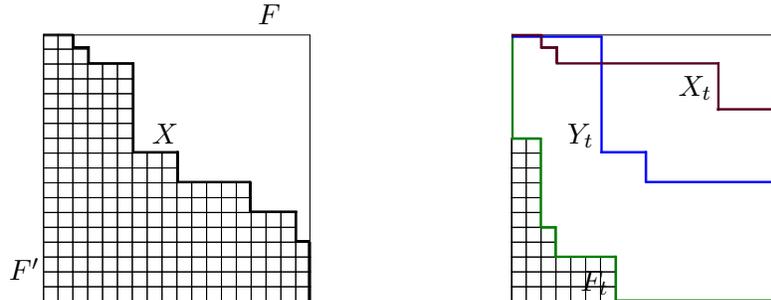


Figure 8: (a) A staircase walk with 7 peaks and 6 valleys. (b) $F_t \subseteq X_t$ and $F_t \subseteq Y_t$ for all t

We begin by proving that a monotone Markov chain with nonnegative drift will mix rapidly. Call a Markov chain \mathcal{M} on Ω *monotone* if Ω is a distributive lattice with partial order \preceq and there exists a coupling (X_t, Y_t) such that if $X_t \preceq Y_t$, then $X_{t+1} \preceq Y_{t+1}$. For $X_t \in \Omega$, let $\mathbb{P}_{\prec}(X_t) = \mathbb{P}[X_{t+1} \neq X_t \text{ and } X_{t+1} \preceq X_t]$ and let $\mathbb{P}_{\succ}(X_t) = \mathbb{P}[X_{t+1} \neq X_t \text{ and } X_t \preceq X_{t+1}]$.

Theorem 17. *Given a monotone Markov chain \mathcal{M} on Ω with maximal element F and minimal element F' , let $h(X, Y)$ be the Hamming distance between X and Y and let $H = \max_{X, Y \in \Omega} \{h(X, Y)\}$. Assume that for all $X \neq F \in \Omega$,*

$$\mathbb{P}_{\succ}(X) - \mathbb{P}_{\prec}(X) \geq \kappa \geq 0.$$

1. If $\kappa > 0$ then the mixing time is $\tau(\epsilon) \leq \lceil \frac{eH \ln H}{\kappa} \rceil \lceil \ln(\epsilon^{-1}) \rceil$.

2. If $\kappa = 0$ and for all $t \geq 0$, $\mathbb{P}[X' \neq X] \geq Q$ then the mixing time satisfies $\tau(\epsilon) \leq \lceil \frac{2eH^2}{Q} \rceil \lceil \ln(\epsilon^{-1}) \rceil$.

Proof. First we notice that the coupling time $T^{x,y} = \mathbb{E}[\min\{t : X_t = Y_t | X_0 = x, Y_0 = y\}]$ is bounded by the hitting time to reach F from F' . Define $T_F^x = \mathbb{E}[\min\{t : X_t = F | X_0 = x\}]$ to be the hitting time to reach F from x . Let $F_0 = F'$, and couple the moves of $\{F_t\}$ together with the moves of $\{X_t\}$ and $\{Y_t\}$; that is, choose the same (i, b, r) for F_t as in X_t and Y_t . Since the Markov chain is monotone with respect to the given coupling, we have $F_t \preceq X_t$ and $F_t \preceq Y_t$ for all $t \geq 0$. Thus if $F_t = F$, we also have $X_t = Y_t = F$. So the coupling time for X_t and Y_t is bounded by the hitting time of F_t to F . See Figure 8(b).

Let $\phi_t = h(X_t, F)$. For part (1), we use $\mathbb{E}[\phi_t] \leq \phi_t - \kappa\phi_{t-1}/H \leq (1 - \kappa/H)^t \phi_0$. For part (2), note that $\mathbb{E}[\phi_{t+1}^2 + \phi_t^2] \geq \phi_t^2 \geq 1$, so using Lemma 7 with $q = 0$ and $D = H$, we obtain $\mathbb{E}[T] \leq \frac{2H^2}{Q}$. Then the theorem follows from the Coupling Lemma, Theorem 4. \square

Notice that since $\lambda > 1$, we expect that for any downset $X_0 \in \Omega$, the sequence $\{X_t\}$ should approach the unique maximal element of Ω . We will show that it suffices to choose $\lambda \geq d$ to reach the full d -dimensional cube from an arbitrary position in polynomial time, thus achieving polynomial mixing time. In particular, we will show that \mathcal{M}'_B satisfies the conditions of Theorem 17.

We define a *peak* of X to be a position where we can remove a hypercube from X and a *valley* of X to be a position where we can add a hypercube to X (see Figure 8(a)). For the following lemma, define $\mathcal{V}(D)$ (and, respectively, $\mathcal{P}(D)$) to be the set of valleys (peaks) of a downset D .

Lemma 18. *For any downset $\sigma \neq F \in \Omega$,*

$$|\mathcal{P}(\sigma)| \leq (d-1)|\mathcal{V}(\sigma)| + 1. \quad (3)$$

Proof of Lemma 18. First we notice that if $\sigma \neq F$, then it has at least one valley. Furthermore, if σ has a single valley, then the number of peaks is at most $d = (d-1)|\mathcal{V}(\sigma)| + 1$. Now, assume the number of valleys is more than one and proceed as follows. Choose a valley $v = (v_1, v_2, \dots, v_d)$ that maximizes v_d . Construct σ' from σ by adding every hypercube lying above v in the d dimension; that is, $\sigma' = \sigma \cup \{v + e_d, v + 2e_d, \dots, v + (h - v_d)e_d\}$. Then $|\mathcal{P}(\sigma')| \geq |\mathcal{P}(\sigma)| - d + 1$, and $|\mathcal{V}(\sigma')| = |\mathcal{V}(\sigma)| - 1$. Hence

$$|\mathcal{P}(\sigma)| - (d-1)|\mathcal{V}(\sigma)| \leq |\mathcal{P}(\sigma')| - (d-1)(|\mathcal{V}(\sigma')|).$$

Obtain $\hat{\sigma}$ by iterating this operation until only a single valley remains; then we have

$$|\mathcal{P}(\sigma)| - (d-1)|\mathcal{V}(\sigma)| \leq |\mathcal{P}(\hat{\sigma})| - (d-1)|\mathcal{V}(\hat{\sigma})| \leq 1,$$

as desired. \square

Lemma 18 implies that if $\lambda \geq d$ then for all $\sigma \neq F \in \Omega$, $|\mathcal{P}(\sigma)| \leq (d-1)|\mathcal{V}(\sigma)| + 1 \leq d|\mathcal{V}(\sigma)|$, and so

$$\begin{aligned} \mathbb{P}_{\succ}(\sigma) - \mathbb{P}_{\prec}(\sigma) &= \frac{1}{2\alpha} \left(|\mathcal{V}(\sigma)| - \frac{1}{\lambda} |\mathcal{P}(\sigma)| \right) \\ &\geq \frac{|\mathcal{V}(\sigma)|}{2\alpha} \left(1 - \frac{d}{\lambda} \right) \geq 0. \end{aligned}$$

Hence, the Markov chain \mathcal{M}'_B has nonnegative drift towards the maximal configuration. Using Theorem 17 we can show that for $\lambda = d + \delta$ for some $\delta \geq 0$, $\tau(\varepsilon) = O(\delta^{-1}\alpha n \ln n \ln \varepsilon^{-1})$ for $\delta > 0$, and otherwise $\tau(\varepsilon) = O(\alpha n^2 \ln \varepsilon^{-1})$. However, we are able to get a better bound in Theorem 16 by introducing another metric that gives strict additive contraction for every configuration, even when $\lambda = d$.

Proof of Theorem 16.

Recall from the proof of Theorem 17 that the coupling time is bounded by the hitting time to reach F from F' . We will show that the expected time to hit F is small, using Lemma 18. Let $S_1 = \{\sigma \in \Omega : \mathcal{P}(\sigma) < d|\mathcal{V}(\sigma)\} \cup F$, and $S_2 = \{\sigma \in \Omega : \mathcal{P}(\sigma) = d|\mathcal{V}(\sigma)\}$. Define the function $\phi(\sigma) = H(F, \sigma) + I_{S_2}/(2d)$, where $H(X, Y)$ is the Hamming distance between the downsets X and Y , and I_{S_2} is the indicator function for the set S_2 . We will show that ϕ has negative additive drift towards 0.

Notice that Lemma 18 implies that if $\sigma \in S_2$ then it has exactly one valley. Moreover, if $\sigma \in S_1$ can move to S_2 in a single step, then the number of possible moves to take it to S_2 is at most 3. If σ has a single valley then it must have a valley v such that for some dimension i , $v_i = 0$; adding a cube at this valley could move σ into S_2 . If σ has three valleys then it has a single hypercube c it can remove to enter S_2 , and if σ has two valleys then it can enter S_2 by adding cubes at either of those valleys or removing a cube between them are the only ways to potentially move into S_2 . Hence, if $\sigma_t \in S_1$, then the probability that $\sigma_{t+1} \in S_2$ is at most $3/(2\alpha)$. Moreover, if $\sigma_t \in S_2$, then the probability that $\sigma_{t+1} \in S_1$ is at least $1/(2\alpha)$.

Now, conditioning on whether σ_t is in S_1 or in S_2 , we have $\mathbb{E}[\phi(\sigma_{t+1}) - \phi(\sigma_t) | \sigma_t \in S_2] \leq 0 - (2d)^{-1}(2\alpha)^{-1}$ and

$$\begin{aligned} \mathbb{E}[\phi(\sigma_{t+1}) - \phi(\sigma_t) | \sigma_t \in S_1] &\leq \frac{-|\mathcal{V}(\sigma_t)| + \lambda^{-1}|\mathcal{P}(\sigma_t)|}{2\alpha} + \frac{3}{2\alpha} \left(\frac{1}{2d} \right) \\ &\leq \frac{1}{2\alpha} (-|\mathcal{V}(\sigma_t)| + \lambda^{-1}(d-1)|\mathcal{V}(\sigma_t)|) + \frac{3}{2\alpha} \left(\frac{1}{2d} \right) \\ &\leq \frac{|\mathcal{V}(\sigma_t)|}{2\alpha} \left(-\frac{1}{d} \right) + \frac{3}{2\alpha} \left(\frac{1}{2d} \right) \\ &\leq \frac{-1}{4\alpha d}. \end{aligned}$$

Thus

$$\mathbb{E}[\phi(\sigma_t)] \leq \sigma_0 - \frac{t}{4\alpha d} \leq \left(n + \frac{1}{2d} \right) - \frac{t}{4\alpha d} \leq \varepsilon$$

whenever $t \geq 4d\alpha(n + \frac{1}{2d} - \varepsilon)$. Thus the mixing time satisfies $\tau(\varepsilon)O(\alpha n \ln(\varepsilon^{-1})) = O(h^{2d-1} \ln(\varepsilon^{-1})) = O(n^2 \ln(\varepsilon^{-1}))$ for arbitrary dimension d , but for $d = 2$, this implies $\tau(\varepsilon) = O(h^3 \ln(\varepsilon^{-1})) = O(n\sqrt{n} \ln(\varepsilon^{-1}))$. □

5 The fluctuating bias Markov chain

It turns out to be quite interesting to consider the case of *fluctuating bias* where the bias of a tile depends on its position. This situation is more realistic, particularly in the context of self-assembly. For instance, the probability of a tile lined with DNA attaching to the substrate depends on the strength of the bonds along the edges of the tile as well as the relative densities of each tile. Recall the *bias at \bar{x}* is defined as follows: if τ is formed by adding a cube at position \bar{x} to σ , then $\lambda_{\bar{x}} = P(\sigma, \tau)/P(\tau, \sigma)$ is called the *bias at \bar{x}* . The stationary probability of a configuration σ will be proportional to $\prod_{\bar{x} \in \sigma} \lambda_{\bar{x}}$.

Most of the results from Section 4 generalize to this setting, as long as we satisfy certain bounds on the amount the bias can fluctuate. When the minimum bias λ_L is large enough, we do not need any upper bound. While the upper bound might seem unnecessary even for small values of $\lambda_L > 1$, the chain can actually take exponential time to reach equilibrium if the biases vary too much (see Section 5.3).

5.1 Fast mixing with large enough minimum bias

In Section 4.3, we showed that: (i) the coupling time for the Markov chain \mathcal{M}_B was bounded by the hitting time to the maximal configuration, and (ii) the hitting time is polynomial, assuming that the bias is large enough. Clearly, (i) still holds for the fluctuating bias Markov chain \mathcal{M}_B , and the hitting time to the maximal configuration should only decrease if some of the biases are increased. Therefore, the following theorem is a simple consequence of Theorem 16:

Theorem 19. *Let R be the d -dimensional $h \times h \times \dots \times h$ hypercube with volume $h^d = n$ and fluctuating bias. Assume the minimum bias λ_L satisfies $\lambda_L \geq d$. Then the mixing time of \mathcal{M}_B satisfies $\tau(\varepsilon) = O(h^{2d-1} \ln \varepsilon^{-1}) = o(n^2 \ln \varepsilon^{-1})$.*

Moreover, we can use a similar argument, together with the uniform bias results above, to obtain the following stronger result for fluctuating bias in 2 dimensions:

Theorem 20. *Let R be a rectangular $h \times w$ region in \mathbb{Z}^2 with fluctuating bias. Suppose the minimum bias λ_L is a constant larger than 1. Then the mixing time of \mathcal{M}_B satisfies $\tau(\varepsilon) = O(wh \ln \varepsilon^{-1})$.*

Again, this is optimal up to constants.

Proof of Theorem 20. Recall from Theorem 17 that the hitting time to the maximal configuration is an upper bound on the coupling time. Thus, it suffices to show that for uniform bias λ , where $\lambda > 1$ is a constant, the hitting time to the maximal configuration is $O(wh \ln \varepsilon^{-1})$. For square regions with area n , a bound of $O(n)$ was proved by Benjamini et al. [2]. We wish to extend this to rectangular regions.

It is well-known that in 2 dimensions for uniform bias λ , where $\lambda > 1$ is a constant, the maximal configuration has constant probability in the stationary distribution. For completeness, we prove this next in the fluctuating bias setting.

Let $p(h, w; t)$ denote the number of integer partitions of t into at most h parts, each of size at most w . This is precisely the number of staircase walks in an $h \times w$ region with t squares below the curve. Consider the generating function for $p(h, w; t)$:

$$F(q) = \sum_{t=0}^{hw} p(h, w; t) q^t = \binom{h+w}{w}_q,$$

where $\binom{m}{r}_q = \prod_{i=0}^{r-1} \frac{1-q^{m-i}}{1-q^i}$ is the Gaussian binomial coefficient. Then the normalizing constant Z is equal to $F(\lambda)$. We wish to show that the weight of the highest configuration λ^{hw}/Z is at least a constant, independent of h and w .

By rearranging terms, we have

$$F(q) = \prod_{i=1}^w \frac{q^{h+i} - 1}{q^i - 1} \leq \prod_{i=1}^w \frac{q^{h+i}}{q^i - 1} = q^{hw} \prod_{i=1}^w \frac{1}{1 - q^{-i}}.$$

Let $x = 1/\lambda$. Then $F(\lambda) = q^{hw} \prod_{i=1}^w \frac{1}{1-x^i}$, and so

$$\begin{aligned} \ln(Z\lambda^{-hw}) &\leq \ln(F(\lambda)\lambda^{-hw}) = -\sum_{t=1}^w \ln(1-x^t) \\ &= \sum_{t=1}^{hw} x^t + \frac{x^{2t}}{2} + \frac{x^{3t}}{3} + \dots \\ &\leq x + 2x^2 + 3x^3 + 4x^4 + \dots \\ &= \frac{x}{(1-x)^2}. \end{aligned} \tag{4}$$

Inequality 4 follows because x^i appears in at most i terms of the sum, each with a coefficient at most 1. Therefore the maximum configuration has weight $\lambda^{hw}/Z \geq e^{-x/(1-x)^2}$, for any h and w .

By Theorem 8, we know that the mixing time of the uniform bias chain with bias λ_L is $O(wh \ln \varepsilon^{-1})$, so we expect the uniform bias chain to hit the maximal configuration in $O(wh \ln \varepsilon^{-1})$ steps.

The hitting time of the fluctuating bias Markov chain is at most the hitting time of the uniform bias chain whose bias is equal to λ_L . To see this, we will couple the two chains. We start with the uniform bias chain below the fluctuating chain, for example by setting it equal to the minimal configuration. At each step, we choose the same square to add, and whenever the fluctuating bias chain decides to remove a square, the uniform chain does too; this is possible because the bias of the fluctuating chain is at least the bias of the uniform bias chain at every square. Therefore, monotonicity is preserved during the coupling, so when the uniform bias chain hits the maximal configuration, so must the fluctuating bias chain. \square

5.2 Fast mixing when the fluctuations are bounded

We can also extend the exponential metric technique of Section 4 to handle fluctuating bias, provided the biases λ_x for $x \in R$ do not vary too much.

Theorem 21. *Let R be any nice d -dimensional region with volume n and span α , and suppose the bias at any point x satisfies $1 < \lambda_L \leq \lambda_x \leq \lambda_U$. If the maximum and minimum biases are such that*

$$\frac{d}{\sqrt{\lambda_L}} - 1 - \frac{1}{\lambda_U} \leq -\chi \tag{5}$$

for some $\chi \geq 0$, then

1. If $\chi > 0$, then the mixing time of \mathcal{M}'_B satisfies $\tau'(\varepsilon) = O(\chi^{-1}\alpha(\alpha \ln \lambda + \ln n) \ln \varepsilon^{-1})$.
2. If $\chi \geq 0$, then $\tau'(\varepsilon) = O(\alpha n^2(\alpha \ln \lambda + \ln n)^2 \ln(\varepsilon^{-1}))$.

Proof. This theorem is proved nearly identically to Theorem 14. In this case we define the distance metric as

$$\phi(\sigma, \rho) = \sum_{x \in \sigma \oplus \rho} (\sqrt{\lambda_L})^{\|x_0\|_1 - \|x\|_1}.$$

Given that \mathcal{M}'_B chose v^* and b such that ϕ_t can increase, the expected increase is at most $\phi_t \lambda_L^{-1/2}$. Indeed, if the move is of the form $v = x + \bar{u}_i$ for some i , then the increase in distance is $(\sqrt{\lambda_L})^{\|x_0\|_1 - \|v\|_1} = \phi_t \lambda_L^{-1/2}$. If the move is of the form $v = x - \bar{u}_i$ for some i , then the increase in distance is $\lambda_L^{-\|v\|_1/2} = \phi_t \sqrt{\lambda_L}$, but the chance of choosing an appropriate p is $1/\lambda_v \leq 1/\lambda_L$. Therefore the expected increase is again at most $\phi_t \lambda_L^{-1/2}$.

This implies that the expected change in distance is negative. As before, there are at most d bad moves, but now the two good moves happen with probability 1 and $1/\lambda_x \geq 1/\lambda_U$, respectively. Therefore the expected change in distance satisfies

$$\mathbb{E}_t[\phi_{t+1} - \phi_t] \leq \frac{1}{2\alpha} \left(d \cdot \frac{\phi_t}{\sqrt{\lambda_L}} - \left(1 + \frac{1}{\lambda_U}\right) \phi_t \right) \leq -\frac{\phi_t \chi}{2\alpha}.$$

The rest of the proof is identical to the proof of Theorem 14. □

5.3 Slow mixing when the fluctuations are unbounded

Let R be an $n \times n$ square region in 2 dimensions. Each $(x, y) \in R$ has a bias $\lambda_{x,y}$. In this section we show that even in this simple setting, the Markov chain \mathcal{M}'_B can be slowly mixing when the fluctuations are unbounded.

Note that we have restricted throughout this paper to the case that $\lambda_{x,y} \geq 1$ for all $(x, y) \in R$. This restriction is necessary, or the chain might not be rapidly mixing. For example, it can be shown that if $\lambda_{x,y} < 1$ when $x + y \leq n$ and $\lambda_{x,y} > 1$ when $x + y > n$, then \mathcal{M}'_B requires exponential time to converge. Indeed it will be difficult to move from a tiling that is nearly empty to one that is nearly full, even though these each occupy a constant fraction of the stationary probability.

Suppose that for all $(x, y) \in R$, $\lambda_{x,y} \geq \lambda_L > 1$. We know by Theorem 14 that for uniform bias, the Markov chain is rapidly mixing for all $\lambda > 1$ polynomially bounded away from 1 and the result is in general easier to show for larger values of λ . This leads us to expect rapid mixing in the fluctuating bias case, as long as the minimum bias λ_L is polynomially bounded away from 1. However, this is not true in general:

Theorem 22. *There exists a set of biases $\{\lambda_{x,y}\}$ on R where $\lambda_{x,y} > 1$ for all $(x, y) \in R$ and yet the mixing time of \mathcal{M}_B is $\Omega(e^{n/24})$.*

We show that if the biases below the line $x + y = n + M$ (where $M = n - \sqrt{n}$) are all close to 1 and all other biases are some very large constant ξ , then the mixing time of \mathcal{M}_B is exponentially large in n (see Figure 9). We identify sets S_1, S_2, S_3 such that $\pi(S_2)$ is exponentially smaller than both $\pi(S_1)$ and $\pi(S_3)$, which have equal weight, but to get between S_1 and S_3 , \mathcal{M}_{nn} and \mathcal{M}_T must pass through S_2 . Thus, the set S_1 forms a “bad cut” that prevents the Markov chain from mixing rapidly, regardless of the initial configuration.

To formalize these ideas, we will bound the conductance of the Markov chain. The *conductance* of an ergodic Markov chain \mathcal{M} with stationary distribution π is

$$\Phi_{\mathcal{M}} = \min_{\substack{S \subseteq \Omega \\ \pi(S) \leq 1/2}} \phi_S,$$

where $\phi_S = \phi_S^{(\mathcal{M})}$ is the *conductance* of a set $S \subset \Omega$, defined by

$$\phi_S = \frac{1}{\pi(S)} \sum_{s_1 \in S, s_2 \in \bar{S}} \pi(s_1)P(s_1, s_2).$$

Essentially, S is a bad cut if ϕ_S is exponentially small. The existence of such a set S prevents the Markov chain from mixing rapidly [11]:

Theorem 23. *For any Markov chain with conductance Φ , $\tau \geq (4\Phi)^{-1} - 1/2$.*

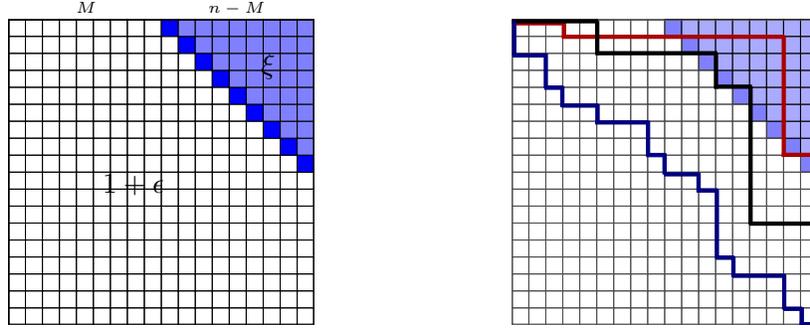


Figure 9: (a) Fluctuating bias with exponential mixing time. (b) Staircase walks in S_1, S_2 , and S_3 .

Proof of Theorem 22. We will define values $\{\lambda_{x,y}\}_{(x,y) \in R}$ such that the conductance of the Markov chain \mathcal{M}'_B is small. Let $M = n - \sqrt{n}$. For all (x, y) such that $x + y \leq n + M$, define $\lambda_{x,y} = 1 + \epsilon$, where $\epsilon = \frac{1}{4n}$. For all remaining (x, y) , let $\lambda_{x,y} = \xi$, where $\xi > 1$ will be defined later. For a staircase walk σ consisting of a sequence of steps $\sigma_i \in \{\pm 1\}$, define the *height* of σ_i as $\sum_{j \leq i} \sigma_j$, and let $\max(\sigma)$ be the maximum height of σ_i over all $1 \leq i \leq 2n$. Let S_1 be the set of configurations σ such that $\max(\sigma) < n + M$, S_2 the set of configurations such that $\max(\sigma) = n + M$, and S_3 the set of configurations such that $\max(\sigma) > n + M$.

First we will show that $\pi(S_2)$ is exponentially smaller than $\pi(S_1)$ for all values of ξ . Let $A(\sigma)$ be the number of tiles contained in σ . Then

$$\pi(S_2) = \frac{1}{Z} \sum_{\sigma \in S_2} (1 + \epsilon)^{A(\sigma)} \leq \frac{(1 + \epsilon)^{n^2 - n/2} |S_2|}{Z},$$

since there are at most $n^2 - (n - M)^2/2 = n^2 - n/2$ tiles with weight $1 + \epsilon$ in σ . By the definition of ϵ , we have $(1 + \epsilon)^{n^2 - n/2} = (1 + \frac{1}{4n})^{n^2 - n/2} \leq e^{n/4 - 1/8}$. Hence $\pi(S_2) \leq e^{n/4 - 1/8} |S_2|/Z$.

Next we will bound $|S_2 \cup S_3|$, which in turn provides an upper bound on $|S_2|$. The unbiased Markov chain is equivalent to a simple random walk $W_{2n} = X_1 + X_2 + \dots + X_{2n} = 0$, where $X_i \in \{+1, -1\}$ and where a $+1$ represents a step to the right and a -1 represents a step down. We call this random walk *tethered* since it is required to end at 0 after $2n$ steps. Compare walk W_{2n}

with the untethered simple random walk $W'_{2n} = X'_1 + X'_2 + \dots + X'_{2n}$.

$$\begin{aligned}
P\left(\max_{1 \leq t \leq 2n} W_t \geq M\right) &= P\left(\max_{1 \leq t \leq 2n} W'_t \geq M \mid W'_{2n} = 0\right) \\
&= \frac{P(\max_{1 \leq t \leq 2n} W'_t \geq M)}{P(W'_{2n} = 0)} \\
&= \frac{2^{2n}}{\binom{2n}{n}} P\left(\max_{1 \leq t \leq 2n} W'_t \geq M\right) \\
&\leq \frac{e^2 \sqrt{n}}{2\sqrt{\pi}} P\left(\max_{1 \leq t \leq 2n} W'_t \geq M\right) \\
&< 3\sqrt{n} P\left(\max_{1 \leq t \leq 2n} W'_t \geq M\right),
\end{aligned} \tag{6}$$

where 6 follows from Sterling's formula. Since the $\{X'_i\}$ are independent, we can use Chernoff bounds to see that

$$P\left(\max_{1 \leq t \leq 2n} W'_t \geq M\right) \leq 2nP(W'_{2n} \geq M) \leq 2ne^{-\frac{M^2}{2n}}.$$

Notice that $M^2/(2n) = (n - \sqrt{n})^2/(2n) = (\sqrt{n} - 1)^2/2 \geq n/3$ for $n \geq 4$. Together these show that $P(\max_{1 \leq t \leq 2n} W_t \geq M) \leq 6n^{3/2}e^{-n/3}$. In particular,

$$|S_2 \cup S_3| \leq 6 \binom{2n}{n} n^{3/2} e^{-n/3}.$$

Therefore we have

$$\begin{aligned}
\frac{\pi(S_2)}{\pi(S_1)} &\leq \frac{\frac{1}{Z} e^{n/4-1/8} |S_2 \cup S_3|}{\frac{1}{Z} \left(\binom{2n}{n} - |S_2 \cup S_3|\right)} \\
&\leq e^{n/4-1/8} \left(\frac{\binom{2n}{n}}{|S_2 \cup S_3|} - 1\right)^{-1} \\
&\leq e^{n/4-1/8} \left(\frac{\binom{2n}{n}}{6 \binom{2n}{n} n^{3/2} e^{-n/3}} - 1\right)^{-1} \\
&\leq \frac{6e^{n/4-1/8} n^{3/2}}{e^{n/3} - 6n^{3/2}} < e^{-n/24},
\end{aligned}$$

for large enough n . Therefore, $\pi(S_2)$ is exponentially smaller than $\pi(S_1)$ for every value of ξ .

Our next goal is to show that there exists a value of ξ for which $\pi(S_3) = e\pi(S_1)$, which will imply that $\pi(S_2)$ is also exponentially smaller than $\pi(S_3)$, and hence the set S_1 forms a bad cut, regardless of which state the Markov chain begins in. To find this value of ξ , we will rely on the continuity of the function $g(\xi) = Z\pi(S_3) - Z\pi(S_1)$ with respect to ξ . Notice that $Z\pi(S_1)$ is constant with respect to ξ and

$$Z\pi(S_3) = \sum_{\sigma \in S_3} (1 + \epsilon)^{A_b(\sigma)\xi} \xi^{A_a(\sigma)},$$

(where $A_b(\sigma)$ is the number of tiles below the diagonal M in σ and $A_a(\sigma)$ is the number of tiles above the diagonal M in σ) is just a polynomial in ξ . Therefore $Z\pi(S_3)$ is continuous in ξ and hence $g(\xi)$ is also continuous with respect to ξ . We will show that $g(1 + \epsilon) < 0$ and $g(4e^{2.5}) > 0$, so

by continuity we will conclude that there exists a value of ξ satisfying $1 + \epsilon < \xi < 4e^{2.5}$ for which $g(\xi) = 0$ and $Z\pi(S_3) = eZ\pi(S_1)$. Clearly this implies that for this choice of ξ , $\pi(S_3) = e\pi(S_1)$, as desired.

First, we will show that when $\xi = (1 + \epsilon)$, $Z\pi(S_3) < eZ\pi(S_1)$, so $g(1 + \epsilon) < 0$. This is easy to see, as

$$\pi(S_3) \leq \frac{1}{Z}|S_3|\gamma^{n^2} \leq \frac{1}{Z}|S_3|e \leq \frac{1}{Z}|S_1|e \leq e\pi(S_1).$$

Next we will show that $g(4e^{2.5}) > 0$. First we notice that since the maximal tiling is in S_3 ,

$$\Pi(S_3) \geq Z^{-1}(1 + \epsilon)^{n^2 - \frac{(n-M)^2}{2}} \xi^{\frac{(n-M)^2}{2}}.$$

Also,

$$\Pi(S_1) = Z^{-1} \sum_{\sigma \in S_1} (1 + \epsilon)^{A_a(\sigma)} < Z^{-1} \binom{2n}{n} (1 + \epsilon)^{n^2 - \frac{(n-M)^2}{2}}.$$

Therefore

$$e\pi(S_1)/\pi(S_3) < \frac{e \binom{2n}{n}}{\xi^{\frac{(n-M)^2}{2}}} \leq e(2e)^n \xi^{-n/2} \leq 1$$

for $n \geq 4$, since $\xi = 4e^{2.5}$. Hence $g(4e^{2.5}) = Z\pi(S_3) - eZ\pi(S_1) > 0$, as desired.

Finally, we may analyze the conductance.

$$\Phi \leq \phi_{S_1} \leq \frac{1}{\pi(S_1)} \sum_{x \in S_1} \pi(x) \sum_{y \in S_2} P(x, y) \leq \frac{1}{\pi(S_1)} \sum_{x \in S_1} \pi(x) \pi(S_2) \leq e^{-n/24}.$$

Hence, by Theorem 23, the mixing time of \mathcal{M}'_B satisfies

$$\tau \geq (4e^{-n/24})^{-1} - 1/2 \geq e^{n/24}/4 - 1/2.$$

□

6 Conclusions

In this paper, we showed that the natural local Markov chain for sampling biased monotonic surfaces is efficient for any uniform bias in two dimensions and for large enough bias in higher dimensions. We also gave the first analysis of the convergence rate when we have fluctuating bias, showing that as long as the biases do not differ by too much, the assembly is still rapid. We have made additional contributions to the study of biased card shuffling as well, corresponding to the two-dimensional case of monotone surfaces. The bound in Theorem 8 on the mixing time of \mathcal{M}_B when $d = 2$ yields a simpler proof that the nearest-neighbor transposition chain on biased permutations is rapidly mixing, using the bijection from Benjamini et al. [2]. In fact, we achieved the same optimal bounds on the mixing time. Recently our improved bounds on the mixing time of \mathcal{M}_B for rectangular regions were used to show a tighter bound on the mixing time of the nearest-neighbor transposition chain for a generalization of biased permutations arising in the context of self-organizing lists [4].

More significantly, the techniques from Section 4.2 can be extended to other applications, such as biased 3-colorings. For example, there is a well-known bijection between 3-colorings of \mathbb{Z}^2 and sets of monotonic, *edge-disjoint* paths (see, e.g. [13]). The construction generalizes to arbitrary dimension as well, forming $(d - 1)$ -dimensional monotonic surfaces that are *face-disjoint*. There

is a Markov chain \mathcal{M}_{col} arising in the context of asynchronous cellular automata which samples biased 3-colorings. Our technique can be used to show that as long as the bias satisfies $\lambda \geq 4d^2$, the mixing time of \mathcal{M}_{col} satisfies $\tau(\varepsilon) = O(n^2 \ln(\varepsilon^{-1}))$. We expect that there may be other situations where standard coupling techniques are not sufficient for which coupling with exponential metrics will be useful.

References

- [1] D. Aldous. Random walk on finite groups and rapidly mixing Markov chains. *In Seminaire de Probabilites XVII*, pages 243–297, 1983.
- [2] I. Benjamini, N. Berger, C. Hoffman, and E. Mossel. Mixing times of the biased card shuffling and the asymmetric exclusion process. *Trans. Amer. Math. Soc.*, 357:3013–3029, 2005.
- [3] N. Berger, C. Kenyon, E. Mossel, and Y. Peres. Glauber dynamics on trees and hyperbolic graphs. *Probability Theory and Related Fields*, 131:311–340, 2005.
- [4] P. Bhakta, S. Miracle, A.P. Streib, and D. Randall. Mixing times of self-organizing lists and biased permutations. *In Proceedings of the 25th ACM/SIAM Symposium on Discrete Algorithms*, SODA, pages 1–15, 2014.
- [5] R. Bublely and M. Dyer. Faster random generation of linear extensions. *Discrete Mathematics*, 201:81–88, 1999.
- [6] S. Cannon, S. Miracle, and D. Randall. Phase transitions in random dyadic tilings and rectangular dissections. *In 26th ACM/SIAM Symposium on Discrete Algorithms*, SODA, pages 1573–1589, 2015.
- [7] P. Caputo, F. Martinelli, A. Sinclair, and A. Stauffer. Dynamics of lattice triangulations on thin rectangles. *In ACM Symposium on the Theory of Computing*, STOC, 2015.
- [8] P. Caputo, F. Martinelli, and F.L. Toninelli. Mixing times of monotone surfaces and SOS interfaces: a mean curvature approach. *Comm. Math. Phys.*, 311:157–189, 2012.
- [9] M. Dyer and C. Greenhill. A more rapidly mixing markov chain for graph colorings. *Random Structures & Algorithms*, 13:285–317, 1998.
- [10] T.-J. Fu and N.C. Seeman. DNA double-crossover molecules. *Biochemistry*, 32:3211 – 3220, 1993.
- [11] M. Jerrum and A. Sinclair. Approximate counting, uniform generation and rapidly mixing markov chains. *Information and Computation*, 82:93–133, 1989.
- [12] D. Levin and Y. Peres. Mixing of the exclusion process with small bias. *J. Stat. Phys.*, 165:1035–1050, 2016.
- [13] M. Luby, D. Randall, and A.J. Sinclair. Markov chains for planar lattice structures. *SIAM Journal on Computing*, 31:167–192, 2001.
- [14] U. Majumder, S. Sahu, and J. Reif. Stochastic analysis of reversible self-assembly. *Journal of Computational and Theoretical Nanoscience*, 5:1289–1305, 2008.

- [15] L. McShine and P. Tetali. On the mixing time of the triangulation walk and other catalan structures. In Pardalos et al., editor, *DIMACS-AMS volume on Randomization Methods in Algorithm Design*, volume 43, pages 147–160, 1998.
- [16] D. Randall and P. Tetali. Analyzing glauber dynamics by comparison of markov chains. *Journal of Mathematical Physics*, 41:1598–1615, 2000.
- [17] N.C. Seeman. DNA in a material world. *Nature*, 421:427–431, 2003.
- [18] D. Wilson. Mixing times of lozenge tiling and card shuffling markov chains. *The Annals of Applied Probability*, 1:274–325, 2004.
- [19] E. Winfree. Simulations of computing by self-assembly. In *4th DIMACS Meeting on DNA Based Computers*, 1998.
- [20] E. Winfree, X. Yang, and N.C. Seeman. Universal computation via self-assembly of DNA: Some theory and experiments. *DNA Based Computers II*, 44:191–213, 1996.