Sampling on Lattices with Free Boundary Conditions Using Randomized Extensions

Sarah Cannon*† Dana Randall*‡

Abstract

Many statistical physics models are defined on an infinite lattice by taking appropriate limits of finite lattice regions. A key consideration is which boundary to use when taking those limits, since the boundary can have significant influence on physical and computational properties of the limit. We consider configurations on finite regions with free or partially free boundary conditions and show that by randomly extending the boundary by a few layers, choosing among only a constant number of allowable extensions, we can generalize the arguments used in the fixed boundary setting to infer bounds on the mixing time for free boundaries. We demonstrate this principled approach using randomized extensions for 3-colorings of regions of $\mathbb{Z}^2$ and lozenge tilings of regions of the triangular lattice, building on arguments for the fixed boundary cases due to Luby et al. [14]. Our approach yields an efficient algorithm for sampling free boundary 3-colorings of regions with one reflex corner, the first result to efficiently sample free boundary 3-colorings of any nonconvex region. We also consider self-reducibility of free boundary 3-colorings of rectangles, and show our algorithm can be used to approximately count the number of free-boundary 3-colorings of a rectangle.

1 Introduction

Sampling proper k-colorings uniformly has been the focus of much research. Given a graph $G = (V, E)$ and an integer k, a k-coloring is an assignment of colors $[k] = \{1, \ldots, k\}$ to the vertices so that all pairs of neighboring vertices have distinct colors. A natural Markov chain known as Glauber dynamics starts with any valid coloring, chooses $(v, c) \in V \times [k]$ uniformly, and recolors v with color c if this yields a valid coloring.

Glauber dynamics have been extensively studied, primarily in the case when k is large compared to Δ, the maximum degree of G. If $k \geq \Delta + 2$ the chain is known to connect the state space. Vigoda [22] showed that the chain is rapidly mixing, or converging rapidly to its stationary distribution, as is necessary for efficient sampling, when $k \geq 11\Delta/6$. For graphs with large girth and large Δ, this degree constraint can be reduced [10, 12, 18]. Bubley et al. [2] showed that the chain is rapidly mixing when $k \geq 5$ whenever $\Delta = 3$ or when $k \geq 7$, $\Delta = 4$ and G is triangle free, notable because this includes the Cartesian lattice $\mathbb{Z}^2$. See [6] for a survey on efficiently sampling k-colorings.

While Glauber dynamics have proven more challenging to analyze when the number of colors is small, a notable exception is the interesting case of 3-coloring on finite regions on $\mathbb{Z}^2$. Colorings of the Cartesian lattice are of interest in statistical physics to study phase transitions of the k-state anti-ferromagnetic Potts model, a basic model of anti-ferromagnetism. In this model, vertices are assigned "spins" from $\{1, \ldots, k\}$; in the anti-ferromagnetic case, neighbors have a preference for unequal spins, and at zero temperature this preference becomes required, mapping the valid configurations precisely to the set of proper k-colorings, each occurring with equal probability. While much less has been shown for Glauber dynamics when k is small, 3-colorings on $\mathbb{Z}^2$ are known to map bijectively to Eulerian orientations of the graph, well-known in the statistical physics community as the "ice model." This structure allows more analysis, and Luby et al. [14] showed that a related Markov chain based on "tower moves" that can at once update a linear collection of sites is rapidly mixing on any simply connected region of $\mathbb{Z}^2$, provided the boundary is fixed in advance (known as *fixed boundary conditions*). Randall and Tetali [20] subsequently showed fast convergence of Glauber dynamics itself could be inferred from the comparison method of Diaconis and Saloff-Coste [4]. Unfortunately this analysis does not extend to higher dimensions, and it is expected by physicists that the chain will converge slowly in high dimensions, perhaps even dimensions 3 or 4. Recently this has been verified by Peled and Galvin et al. in $\mathbb{Z}^d$ for sufficiently

*College of Computing, Georgia Institute of Technology, Atlanta, GA 30332-0765

†sarah.cannon@gatech.edu. Supported in part by a Clare Boothe Luce Graduate Fellowship, NSF DGE-1148903, and a Simons Award for Graduate Students in Theoretical Computer Science.

‡randall@cc.gatech.edu. Supported in part by NSF grants CCF-1526900 and CNS-1544090.

high dimensions d [7, 8, 19].

There remain basic questions of convergence of Glauber dynamics, even in the case of sampling 3-colorings in $\mathbb{Z}^2$, particularly in the context of *free boundary conditions*. This question is interesting computationally because for some models the convergence rates are known to depend significantly on the types of boundaries (see, e.g., [1, 16, 17]). Goldberg, Martin and Paterson [9] extended the Markov chain studied by Luby et al. to the case of free boundary conditions on *rectangular* subregions of $\mathbb{Z}^2$, but their argument does not seem to extend to other simply connected regions. While rectangular regions are of the most significance in physics, the restriction to this class of graphs precludes "L-shaped" regions that are necessary for the self-reducibility that allows us to approximately count using a well-known reduction between sampling and approximate counting due to Jerrum, Valiant and Vazirani [11]. Moreover, the proof of Goldberg et al. [9] identifies a set of weights for linear towers near or intersecting the boundary, and shows that these weights allow the coupling of proof of Luby et al. [14] to extend to rectangular regions with free boundary conditions. There is no explanation for why their weights work or why a similar approach apparently fails for more general lattice regions.

Similarly, lozenge tilings on the triangular lattice are easier for fixed boundaries and remain challenging for free boundaries. A *lozenge tiling* is a covering of a lattice region with rhombus shaped lozenges, each covering exactly two adjacent triangles, so that every triangle is covered by a unique lozenge (thus it is dual to perfect matchings on hexagonal lattice regions). This 2-dimensional tiling problem has received a lot of attention because of its remarkable properties: each lozenge tiling seems to jump out of the page to form a 3-dimensional structure comprised of supported boxes, and indeed this "height function" has allowed many deep mathematical discoveries, most notably the "Arctic circle theorem" [3].

Fixed boundaries are most natural for lozenge tilings, where all lozenges are required to stay within the boundary of the region, but there is also some interest in the free boundary case when they may overlap the boundary arbitrarily. Glauber dynamics identify a window of 6 triangles arranged in a hexagon and rotate the lozenges by 60° if the hexagon contains exactly 3 lozenges; this move is equivalent to adding or removing a box in the height function representation. Luby et al. [14] show a related tower chain mixes in polynomial time for lozenge tilings in any simply connected region with fixed boundary conditions, and the argument of Randall and Tetali again shows Glauber dynamics also converge quickly. The case of lozenge tilings with free boundary conditions was studied by Martin and Randall [15], using a correspondence between tilings and non-intersecting lattice paths, dynamic programming, and an approach based on determinants. This allows us to approximately count the number of configurations on hexagonal regions with free boundaries, but since the method does not seem to easily generalize to other regions, we again cannot use self-reducibility to construct an algorithm for efficiently sampling.

In this paper, we study Glauber dynamics for 3-colorings on subregions of $\mathbb{Z}^2$ and lozenge tilings on subregions of the triangular lattice with mixtures of free and fixed boundaries. We say the fixed boundary cells are "compatible" if the heights of all the fixed cells are determined uniquely from the height of any one in the height function representation. This includes all boundaries where the set of free cells and the set of fixed cells are each continuous, including the cases when the boundary cells are all fixed or all free. Note that fixed boundary cells must be compatible for Glauber dynamics to connect the state spaces; see Figure 6 for examples of compatible and non-compatible boundaries.

We show Glauber dynamics converge quickly for 3-colorings and lozenge tilings on any simply connected region with compatible fixed boundary conditions, provided all "reflex" boundary cells are fixed. By reflex cells we mean those that make a left turn on the Cartesian lattice when transversing the boundary in the clockwise direction, or angles greater than 60° on on the triangular lattice.

A significant special case is 3-colorings on rectangular regions of $\mathbb{Z}^2$, as well as lozenge tilings on triangular regions of the triangular lattice, each with free boundary conditions. Moreover, if first we flip a coin to determine the color of the single reflex cell in any "L-shaped" region of $\mathbb{Z}^2$, and then run Glauber dynamics on the remaining cells, our proof shows this chain will converge quickly; this is the first result to efficiently sample uniformly from free boundary 3-colorings of any non-convex region. Another important class for which this shows fast convergence is L-shaped regions with hybrid boundaries, where the fixed piece is contiguous and includes the unique reflex cell. This is precisely the case required to establish self-reduciblity of the sampling problem, and hence we can use sampling to approximately count the number of 3-colorings of any rectangular region with free boundary conditions efficiently for the first time. Likewise, a similar argument establishes self-reducibility for lozenge tilings on triangular regions, again giving the first reduction between counting and sampling.

Our proofs follow a similar strategy as Gold-

berg et al. [9]; however, rather than checking whether there are weights for towers near or crossing the boundary that suffice for the coupling, we *derive* these probabilities using *randomized extensions* and a type of reduction to regions with fixed boundaries. We embed the region with free boundaries in a larger region with fixed boundaries by extending the boundary by a few layers. This extension is randomized so that every free cell on the boundary can be extended to allow moves that can either increase or decrease the height. Naively, this can be accomplished by allowing all exponentially many extensions of the boundary, but this is too many to be practical. Somewhat surprisingly, we show that randomly choosing from merely a *constant number* of extensions suffices. We then extend the coupling arguments from [14] to bound the convergence time of the chain on regions with free boundaries. A comparison argument allows us to infer that Glauber dynamics is efficient as well. Handling regions with a reflex corner is more challenging, so we require all but at most one of the reflex cells on the boundary be fixed.

An interesting feature of this new approach is that our proofs rely only on these reductions to larger regions with fixed boundary conditions, and do not require advance knowledge of the probabilities of tower moves near or crossing the original boundary. Instead these probabilities can be determined from the proof, and for 3-colorings on rectangular regions of $\mathbb{Z}^2$ with free boundaries, the probabilities we derive exactly coincide with those found by [9], as they must. Moreover, our approach of deriving these probabilities allows us to generalize the arguments to lozenge tilings and regions with hybrid free and fixed boundary conditions. We believe that the method of randomized extensions will hold for other problems for which self-reducibility requires a mixture of free and fixed boundary cells.

2 Background: Tower chains and mixing times

It has proven difficult to construct direct proofs bounding the convergence times of Glauber dynamics on various planar lattice structures such as 3-colorings and lozenge tilings. Luby et al. [14] instead introduced "tower-based Markov chains" that sample from these structures efficiently by updating multiple locations at once. Randall and Tetali [20] subsequently used the comparison method of Diaconis and Saloff-Coste [4] to show that the single-site Glauber dynamics are also rapidly mixing. We carefully desribe these chains and this analysis, as this will be relevant in what follows.

Three-colorings. We begin by considering 3-colorings. Let R be an $m \times n$ rectangle, subdivided into mn cells of unit width and height. A *3-coloring* of R assigns to each cell one of three colors, denoted 0, 1, and 2, such that no two cells that are horizontally or vertically adjacent have the same color; see Figure 1(a). A *tower* in a 3-coloring is set of contiguous cells in a row or column whose values can all be incremented or decremented by one (modulo 3) to obtain another valid 3-coloring; applying this change to all cells in a tower simultaneously is a *tower move*. The *height* of a tower is the number of cells whose color changes in a tower move. We specifically consider towers of height $h \geq 2$ and refer to towers of height 1 as *flips*. For any tower of height at least two, exactly one of the tower's extremal cells must have all four neighbors with the same color. Let that cell be called the *end* of the tower, and the other the *start* of the tower. The *start color* of the towers i the color the start cell is changed to by the tower move. We will denote the boundary of R by ∂R. A tower *abutting* a boundary of R has its start or end cell next to ∂R, and a tower *adjacent to* a boundary of R has all cells next to ∂R. Figure 1(a) shows towers that are abutting a boundary (top right), adjacent to a boundary (bottom), and both (top left). In Section 3.3 we consider 3-colorings of regions that are not rectangles, but the conditions we impose ensure these remain the only types of towers we need to consider.

There is a well-known bijection between 3-colorings and Eulerian orientations of regions of the two-dimensional Cartesian lattice. Luby et al. [14] define a Markov chain in the context of Eulerian orientations on lattice regions where the orientations of boundary edges are fixed and are part of the input. This chain, $\mathcal{M}_{eul}$, can be used to efficiently sample Eulerian orientations, and therefore 3-colorings. Each Eulerian orientation decomposes naturally into two *routings*, where a routing is a collection of edge-disjoint lattice paths leading up and to the right. Specifically, all directed edges pointing right or up in the Eulerian orientation belong to one routing, while edges pointing down or left belong to the other. Markov chain $\mathcal{M}_{eul}$ only incorporates tower moves from one of these two routings, but it can just as easily include moves for both. This corresponds more naturally to colorings. We let $\widehat{\mathcal{M}}_{eul}$ denote the Markov chain that uniformly chooses one of the two routings (right/up or down/left), and then makes a move of $\mathcal{M}_{eul}$ in that routing. This corresponds to all tower moves on 3-colorings that change the color of a linear set of cells in one move. It is known that $\widehat{\mathcal{M}}_{eul}$ is rapidly mixing, following immediately from the analysis of $\mathcal{M}_{eul}$ by Luby et al.

Lozenge tilings. Tower chains for lozenge tilings are defined similarly. Let R be any subregion of the triangular lattice in the plane. A *lozenge tiling* of R is a matching in which every triangle in R is

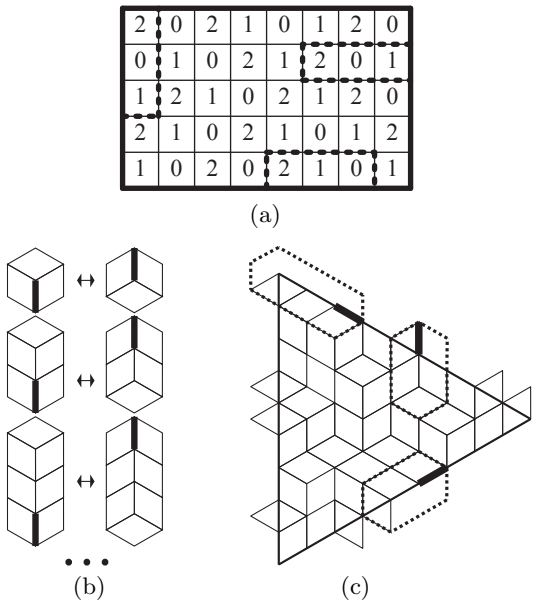

Figure 1: (a) A 3-coloring of a rectangle R. (b) Lozenge tower moves, of heights 1, 2, and 3. (c) A free boundary lozenge tiling of triangle E.

matched with another triangle with which it shares an edge. In a *fixed boundary* tiling all paired triangles must be inside R, while in a *free boundary* tiling triangles may paired across the boundary of R. Luby et al. also presented a tower Markov chain $\mathcal{M}_{loz}$ that efficiently samples fixed boundary lozenge tilings of a finite, simply connected region of the triangular lattice. $\mathcal{M}_{loz}$ specifically considers only towers in the vertical direction. We let $\widehat{\mathcal{M}}_{loz}$ denote a version of their chain which includes towers in all directions and at each iteration uniformly chooses one of the three principle directions. Rapid mixing of $\widehat{\mathcal{M}}_{loz}$ is a simple extension of work in [14].

We focus on lozenge tilings of equilateral triangle shaped regions with free boundary conditions, denoted E. A *tower* in a lozenge tiling of E is an arrangement of lozenges as outlined in Figure 1(b), which shows towers of heights 1, 2, and 3, or some rotation of this configuration. A *tower move* collectively rotates by 180° the lozenges in the tower. Edge e of the triangular lattice shown in bold *starts* the tower move. As shown in Figure 1(c), free boundary lozenge tilings similarly have towers *abutting* a boundary (top right), *adjacent* to a boundary (bottom), or both (top left). Tower moves in these boundary cases are the natural extension of interior tower moves, as are edges starting a tower, also shown in bold. A tower whose outline is six triangles forming a hexagon is a *flip*; the *height* of a tower is the minimum number of flips needed to move from one tower configuration to the other.

Markov chains and mixing times. Our intention is to bound the mixing time of certain Markov chains, i.e., the time it takes to nearly converge to equilibrium. The time a Markov chain $\mathcal{M}$ takes to converge to its stationary distribution π is measured in terms of the distance between π and $\mathcal{P}^t$, the distribution at time t. Let $\mathcal{P}^t(x,y)$ be the t-step transition probability and Ω be the state space. The *mixing time* of $\mathcal{M}$ is

$$\tau(\epsilon) = \min\{t : \|\mathcal{P}^{t'}, \pi\|_{tv} \leq \epsilon, \ \ \forall\ t' \geq t\},$$

where $\|\mathcal{P}^t, \pi\|_{tv} = \max_{x\in\Omega} \frac{1}{2} \sum_{y\in\Omega} |\mathcal{P}^t(x,y) - \pi(y)|$ is the *total variation distance* at time t. As is standard practice, we assume $\epsilon = 1/4$ and consider mixing time $\tau = \tau(1/4)$. We say $\mathcal{M}$ is *rapidly mixing* if τ is bounded above by a polynomial in n and *slowly mixing* if it is bounded below by an exponential in n. (For more on standard definitions regarding Markov chains, see [13] or [21].)

A *coupling* of a Markov chain $\mathcal{M}$ is a joint Markov process $(\mathcal{A}, \mathcal{B})$ on $\Omega \times \Omega$ such that each of the marginals $\mathcal{A}$ and $\mathcal{B}$ is a faithful copy of $\mathcal{M}$ and, once the two coordinates coalesce, they move in unison: if $A_t = B_t$, then $A_{t+1} = B_{t+1}$. *Path coupling* arguments are a convenient way of bounding the mixing time of a Markov chain by considering only a subset U of the joint state space $\Omega \times \Omega$ of a coupling. To show that $\mathcal{M}$ is rapidly mixing, we consider an appropriate metric Φ on Ω and prove that the two marginal chains, if in a joint configuration in subset U, get no farther away in expectation after one iteration. The version due to Dyer and Greenhill [5] is stated below.

THEOREM 2.1. ([5]) *Let $\Phi : \Omega \times \Omega \to \mathbb{Z}$ be a metric which takes values in $[0, B]$, where $\Phi(\sigma, \tau) = 0$ iff $\sigma = \tau$. Let $\mathcal{M}$ be an ergodic Markov chain on Ω and let $(\mathcal{A}, \mathcal{B})$ be a coupling of $\mathcal{M}$, with $\Phi_t := \Phi(A_t, B_t)$ and $\Delta\Phi_t := \Phi_{t+1} - \Phi_t$. Let $U \subseteq \Omega \times \Omega$ be such that for all $(A_t, B_t) \in \Omega \times \Omega$, there exists a path $X_t = Z_0, Z_1, ..., Z_r = Y_t$ such that $(Z_i, Z_{i+1}) \in U$ for $0 \leq i < r$ and*

$$\sum_{i=0}^{r-1} \Phi(Z_i, Z_{i+1}) = \Phi(X_t, Y_t).$$

Suppose that whenever $(A_t, B_t) \in U$, the coupling satisfies

$$\mathbb{E}[\Delta\Phi_t \mid A_t, B_t] \leq 0.$$

Additionally, assume there exists $\alpha > 0$ such that for all t such that $\Phi_t \neq 0$,

$$\mathbb{P}(\Delta\Phi_t \neq 0) > \alpha.$$

Then the mixing time of $\mathcal{M}$ satisfies

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

3 Three-colorings of subregions of $\mathbb{Z}^2$

Our interest will be studying a tower chain $\mathcal{M}_C$ for 3-colorings of regions with a hybrid of free and fixed boundary conditions.[1]

3.1 Rectangles with one free boundary.

We first demonstrate our approach in the simplest setting of an $m \times n$ rectangle R where the boundary cells on three sides of R are fixed (including all four corners), while the colors on the last side are allowed to change. Without loss of generality, assume R's right side has free boundary colors. See Figure 2(a), where cells with fixed colors are shaded. We wish to sample from such 3-colorings of R, and introduce an approach to do so.

Random extensions. Our strategy will be to extend the boundary in a carefully defined way. Starting from a rectangle R with one free side, we create a rectangle R' that extends three units beyond R to the right. Given a coloring χ of R, we extend it to a coloring χ' of R' as follows. First, within R, $\chi' = \chi$. Along the right side of R, the colors in the additional columns will each be a copy of the values of χ in R's rightmost column C, denoted $\chi(C)$, with the values incremented or decremented by 1 (mod 3). With probability $1/2$, the colors of the first column to the right of R are $\chi(C)+1 (\text{mod } 3)$, the colors of the second column are $\chi(C)-1 = \chi(C)+2 (\text{mod } 3)$, and the colors of the third column are $\chi(C)+1 (\text{mod } 3)$. The colors in these columns can be seen as $\chi(C)$ incremented, then incremented again, and then decremented; we will refer to such a configuration as "up-up-down," or UUD for short. With the remaining probability $1/2$, the columns right of C will have a "down-down-up" (DDU) configuration, consisting of, left to right, $\chi(C)-1 (\text{mod } 3)$, $\chi(C)+1 (\text{mod } 3)$, and $\chi(C)-1 (\text{mod } 3)$. We then treat the colors along the boundary of this new region R' as fixed. See Figure 2(b) and (c).

The Markov chain $\mathcal{M}_C$. We now describe a Markov chain, defined using these randomized extensions, that samples from the desired distribution. Recall R is $m \times n$; we label the lower left cell of R as $(0,0)$, the lower right cell as $(m-1,0)$, the upper left cell as $(0,n-1)$, and the upper right cell as $(m-1,n-1)$. Rectangle R' then has coordinates from 0 to $m+2$ in the x direction and 0 to $n-1$ in the y direction.

Starting at any initial 3-coloring χ_0 of R, iterate:

- Choose, uniformly at random, $r \in \{u, d\}$.
- Extend coloring χ_i of R to a coloring χ_i' of R': if $r = u$, the extension has an UUD configuration; if $r = d$, the extension is DDU.
- Starting from coloring χ_i' of rectangle R', apply a single iteration of Markov chain $\widehat{\mathcal{M}}_{eul}$ with a new notion of the height of a tower:
  - Choose, uniformly at random,
    $$(x, y, c, p) \in \{1, 2, ..., m+1\} \times \{1, 2, ..., n-2\} \times \{0, 1, 2\} \times (0, 1).$$
  - If $\chi_i'(x,y) \neq c$ and no neighbors of (x,y) have color c in χ_i', then recolor (x,y) with color c.
  - Else, if (x,y) is the start of a tower of color c whose intersection with R has size $h > 0$, make this tower move if $p < 1/2h$.
- Let χ_{i+1} be the resulting coloring of R.

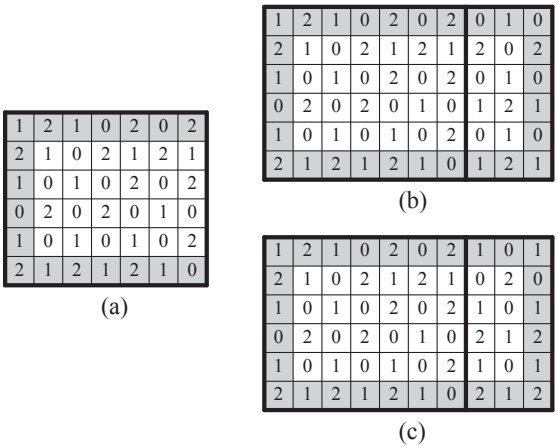

Figure 2: (a) A 3-coloring of rectangle R; shaded cells have fixed colors. (b) R's UUD extension. (c) R's DDU extension.

Ergodicity and reversibility of $\mathcal{M}_C$. It is not immediately obvious that the Markov chain $\mathcal{M}_C$ is ergodic or reversible. We will show that the specific extensions we defined are sufficient to make the chain converge on R to the uniform distribution. Connectivity of the state space follows from an argument based on height functions, similar to connectivity proofs for fixed boundary 3-colorings [14] and free boundary 3-colorings of rectangles [9]. Moreover, the self-loops imply the chain is aperiodic, so $\mathcal{M}_C$ is ergodic. Showing $\mathcal{M}_C$ is reversible, which we need in order to show samples are generated from the uniform distribution, is more subtle. To do this, we enumerate all types of moves and determine the probabilities with which they occur. There are five cases to consider: interior flips, interior towers, boundary flips, towers abutting the boundary, and towers adjacent to the boundary. Because there is only one free boundary, there are no towers that fall into both the last two cases. We first consider

[1] The definition of $\mathcal{M}_C$, and more precisely the random extensions, will change slightly according to context, as detailed for each case.

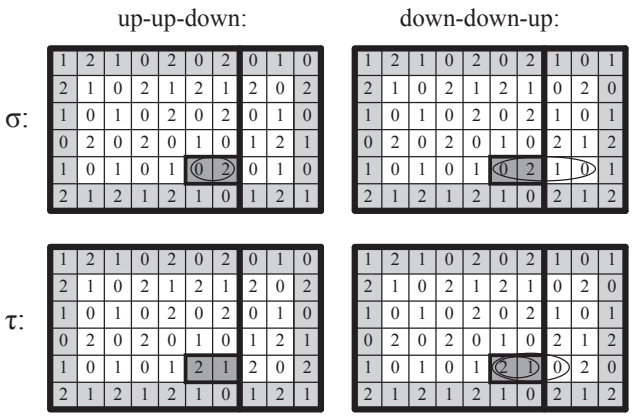

Figure 3: Colorings σ and τ differ by a tower of height 2 (dark grey) abutting R's free boundary. The circled moves transition τ to σ and vice versa.

towers abutting the boundary as a detailed illustrative case, and then continue to analyze the remaining four cases. We define $q := 1/(3(m+1)(n-2))$ to be the probability of picking a given color and given location in R' (excluding its fixed boundary).

Tower abutting the free boundary, height $h > 1$. Consider two 3-colorings σ and τ of R, differing by a tower of height $h \geq 2$ abutting R's free right boundary, as in Figure 3 where $h = 2$. Without loss of generality, let the tower's leftmost cell be $(\overline{x}, \overline{y})$, where $\sigma(\overline{x}, \overline{y}) = 0$ and $\tau(\overline{x}, \overline{y}) = 2$, and suppose the tower's colors decrease (mod 3) from left to right. From coloring σ, the transitions yielding coloring τ are $(u, \overline{x}, \overline{y}, 2, p)$ for $p < 1/2h$ and $(d, \overline{x}, \overline{y}, 2, p)$ for $p < 1/2h$. These towers may go beyond the boundary of R, but due to the nature of our extensions they will always terminate before the fixed boundary of R' and so will never be rejected. Additionally, although their height in R' may be larger than h, their intersection with R will always be of height exactly h. From τ, the transitions yielding σ are $(d, m-1, \overline{y}, c_1, p)$ and $(d, m, \overline{y}, c_2, p)$, each for $p < 1/2h$. There is exactly one value of c_1 and exactly one value of c_2 which will yield a non-stationary transition; in Figure 3, $c_1 = 2$ and $c_2 = 1$. Again, one of these towers has height larger than h in R', but its intersection with R is of height h and this determines the probability with which the move occurs. It follows that

$$\mathbb{P}(\sigma, \tau) = 2 \cdot \frac{1}{2 \cdot (m+1)(n-2) \cdot 3 \cdot 2h} = \frac{q}{2h} = \mathbb{P}(\tau, \sigma).$$

Interior flip. Consider two 3-colorings σ and τ of R, differing at a single location $(\overline{x}, \overline{y})$ that is not on R's free right boundary. Suppose, without loss of generality, $\sigma(\overline{x}, \overline{y}) = 0$ and $\tau(\overline{x}, \overline{y}) = 1$, meaning all neighbors of location $(\overline{x}, \overline{y})$ have color 2. It follows that, for any values of f and any values of p, the move $(f, \overline{x}, \overline{y}, 1, p)$ transitions σ to τ and the move $(f, \overline{x}, \overline{y}, 0, p)$ transitions τ to σ. Thus,

$$\mathbb{P}(\sigma, \tau) = \frac{1}{(m+1)(n-2) \cdot 3} = q = \mathbb{P}(\tau, \sigma).$$

Interior tower, height $h > 1$. Consider two 3-colorings σ and τ of R, differing by a tower of height $h > 1$ that is not abutting or adjacent to the free boundary of R. Without loss of generality, suppose this tower stretches in the horizontal direction from location $(\overline{x}, \overline{y})$ to location $(\overline{x} + h - 1, \overline{y})$, with $\sigma(\overline{x}, \overline{y}) = 0$, $\tau(\overline{x}, \overline{y}) = 1$, $\sigma(\overline{x} + h - 1, \overline{y}) = c_1$ and $\tau(\overline{x} + h - 1, \overline{y}) = c_1 + 1 (\text{mod } 3)$. It follows that, for any values of f, the move $(f, \overline{x}, \overline{y}, 1, p)$ is a transition from σ to τ if $p < 1/2h$ and the move $(f, \overline{x} + h - 1, \overline{y}, c_1, p)$ is a transition from σ to τ if $p < 1/2h$. These are the only transitions between σ and τ. Thus,

$$\mathbb{P}(\sigma, \tau) = \frac{1}{(m+1)(n-2) \cdot 3 \cdot 2h} = \frac{q}{2h} = \mathbb{P}(\tau, \sigma).$$

Free boundary flip. Consider two 3-colorings σ and τ of R, differing at a single location $(\overline{x}, \overline{y})$, where $\overline{x} = m - 1$. Without loss of generality, suppose $\sigma(m-1, \overline{y}) = 1$ and $\tau(m-1, \overline{y}) = 2$. There are three moves that can transform σ to τ: a tower whose intersection with R is height 1, $(u, m-1, \overline{y}, 2, p)$ for $p < 1/2$; a single flip $(d, m-1, \overline{y}, 2, p)$ for any p; and another tower whose intersection with R is of height 1, $(d, m, \overline{y}, 1, p)$ for $p < 1/2$. Similarly, there exist three moves that can transform τ to σ: $(d, m-1, \overline{y}, 1, p)$ for $p < 1/2$, $(u, m-1, \overline{y}, 1, p)$ for any p, and $(u, m, \overline{y}, 2, p)$ for $p < 1/2$. This yields

$$\mathbb{P}(\sigma, \tau) = \frac{2}{2 \cdot (m+1)(n-2) \cdot 3 \cdot 2} + \frac{1}{2 \cdot (m+1)(n-2) \cdot 3} = \frac{1}{(m+1)(n-2) \cdot 3} = q = \mathbb{P}(\tau, \sigma).$$

Tower adjacent to the free boundary, height $h > 1$. Lastly, consider two 3-colorings σ and τ of R, differing by a tower of height $h > 1$ that is adjacent to R's free right boundary. Without loss of generality, suppose this tower stretches from $(m-1, \overline{y})$ to $(m-1, \overline{y}+h-1)$, with $\sigma(m-1, \overline{y}) = 1$, $\tau(m-1, \overline{y}) = 2$, $\sigma(m-1, \overline{y}+h-1) = c_1$ and $\tau(m-1, \overline{y}+h-1) = c_1 + 1 (\text{mod } 3)$. There is one move that can transform σ to τ, namely a tower move of height h selected by $(d, m-1, \overline{y}, 2, p)$ with $p < 1/2h$. There is also one move that can transform τ to σ, specifically $(u, m-1, \overline{y}+h-1, c_1, p)$, for any $p < 1/2h$. It follows that

$$\mathbb{P}(\sigma, \tau) = \frac{1}{2 \cdot (m+4)(n+4) \cdot 3 \cdot 2h} = \frac{q}{4h} = \mathbb{P}(\tau, \sigma).$$

Type of move between σ and τ	$\mathbb{P}(\sigma,\tau)$ $= \mathbb{P}(\tau,\sigma)$
Interior flip	q
Boundary flip	q
Interior tower, height $h > 1$	$\frac{q}{2h}$
Tower abutting the boundary, height $h > 1$	$\frac{q}{2h}$
Tower adjacent to the boundary, height $h > 1$	$\frac{q}{4h}$

Table 1: Types of moves and the probabilities with which they occur in $\mathcal{M}_C$, where $q := \frac{1}{3(m+1)(n-2)}$.

This completes our case analysis, proving $\mathcal{M}_C$ is indeed reversible for rectangles with one free boundary. A summary of the probabilities with which these moves occur can be found in Table 1.

Rapid mixing of $\mathcal{M}_C$. Our final concern is the convergence time. We use a path coupling argument to show $\mathcal{M}_C$ is rapidly mixing and can efficiently sample 3-colorings with one free boundary. Consider a joint process $(\mathcal{A}, \mathcal{B})$ on $\Omega \times \Omega$, where each of $\mathcal{A}$ and $\mathcal{B}$ is a copy of Markov chain $\mathcal{M}_C$. Let A_t and B_t, respectively, be their marginal distributions at iteration t. We couple by making the same choice of (r, x, y, c, p) for both $\mathcal{A}$ and $\mathcal{B}$ at each iteration. The distance Φ between two 3-colorings of R is the number of flips needed to transform one 3-coloring into the other. We consider the case where A_t and B_t differ by a single flip, and we analyze the change in distance $\Delta\Phi_t := \Phi(A_{t+1}, B_{t+1}) - \Phi(A_t, B_t)$.

LEMMA 3.1. *Suppose A_t and B_t differ by one flip, that is, $\Phi(A_t, B_t) = 1$. After one iteration of the joint process $(\mathcal{A}, \mathcal{B})$,*

$$\mathbb{E}[\Delta\phi_t \mid A_t, B_t] \leq 0.$$

Proof. Let $(\overline{x}, \overline{y})$ be the cell where A_t and B_t differ. There are two cases to consider, where $(\overline{x}, \overline{y}) \in int(R)$ or $(\overline{x}, \overline{y})$ is on the free boundary of R. Cell $(\overline{x}, \overline{y})$ could be adjacent to fixed-color cells, but $\mathcal{A}$ and $\mathcal{B}$ are most likely to move apart when it is not so we suppose this is the case. For $(\overline{x}, \overline{y}) \in int(R)$, that $\mathcal{A}$ and $\mathcal{B}$ move no farther apart in expectation follows from the path coupling argument for fixed boundary 3-colorings of Luby et al. [14]; we do not reproduce those details here.

Suppose $(\overline{x}, \overline{y})$ is on the free right boundary of R, i.e. $\overline{x} = m - 1$. Without loss of generality, suppose $A_t(m-1, \overline{y}) = 0$ and $B_t(m-1, \overline{y}) = 1$. We analyze the moves of the coupling that may affect the distance between $\mathcal{A}$ and $\mathcal{B}$. Moves not in the same row or column as $(m-1, \overline{y})$ or an adjacent row or column have the same effect in A_t and B_t, and contribute 0 to the expected change in distance. Interior towers going right in rows $\overline{y} - 1$ or $\overline{y} + 1$ will end at or before column $m - 1$, because $(m-1, \overline{y})$'s left, up, and down neighbors must all have the same color, 2. Interior towers in column $m - 2$ coming up or down towards row $\overline{y}$ will similarly end before they reach row $\overline{y}$. Towers moves contained in column m have no effect once any coloring of R' is restricted back to R, and so cannot affect the distance between A_t and B_t.

Suppose the extension is such that there is a tower in column $m - 1$ coming towards location $(m-1, \overline{y})$, starting at $(m-1, y')$ where $y' > \overline{y} + 1$, and that this tower reaches $(m-1, \overline{y}+1)$. It follows that $(m-1, \overline{y}+1)$'s left and right neighbors must have the same color, either 0 or 1. In one of A_t and B_t (specifically, the one where $(m-1, \overline{y})$'s color matches the color of $(m-1, \overline{y}+1)$'s neighbors), this tower will end at $(m-1, \overline{y}+1)$, and in the other it will end at $(m-1, \overline{y})$. In the first of these cases it has height $y' - \overline{y}$, and in the second it has height $y' - \overline{y} + 1$. For $p < 1/(2(y' - \overline{y} + 1))$, this tower move will occur in both A_t and B_t, resulting in a change in distance of -1 as $A_{t+1} = B_{t+1}$. For $1/(2(y' - \overline{y} + 1)) < p < 1/(2(y' - \overline{y}))$, only the tower move ending at $(m-1, \overline{y}+1)$ will occur, resulting in a change in distance of $+(y' - \overline{y})$. Altogether, this means that as the desired values c and r that yield this tower occur with probability 1/6, conditioning on the proper choice of (x, y) by this iteration of the Markov chain,

$$\begin{aligned}
&\mathbb{E}[\Delta\Phi_t \mid (x,y) = (m-1, y')] \\
&\quad = \frac{1}{6}\left(\frac{1}{2(y' - \overline{y} + 1)}(-1)\right. \\
&\qquad\qquad \left. + \left(\frac{1}{2(y' - \overline{y})} - \frac{1}{2(y' - \overline{y} + 1)}\right)(y' - \overline{y})\right) \\
&\quad = \frac{1}{12}\left(\frac{-1}{y' - \overline{y} + 1} + 1 - \frac{y' - \overline{y}}{y' - \overline{y} + 1}\right) \\
&\quad = \frac{1}{12}\left(1 - \frac{1 + y' - \overline{y}}{y' - \overline{y} + 1}\right) = 0.
\end{aligned}$$

The expected change in distance due to such a move is 0. The same analysis holds for tower moves in column $m-1$ starting at $(m-1, y')$ where $y' < \overline{y} - 1$ that reach neighbor $(m-1, \overline{y}-1)$ of $(m-1, \overline{y})$. The expected change in distance due to such move is also 0.

Analysis is similar for a tower starting at location $(x', \overline{y})$, where $x' < m - 2$, and coming right to location $(m-2, \overline{y})$. Towers originating here in A_t and B_t will differ in their height's intersection with R by one, although one of these towers may continue into the extension. As above, in one of A_t and B_t, this tower will end at $(m-2, y)$. In the other, where the tower ends depends on the frame. For one value of r, the tower will end at $(m-1, y)$, while for the other value of r, the tower will end at $(m+1, y)$. In either case,

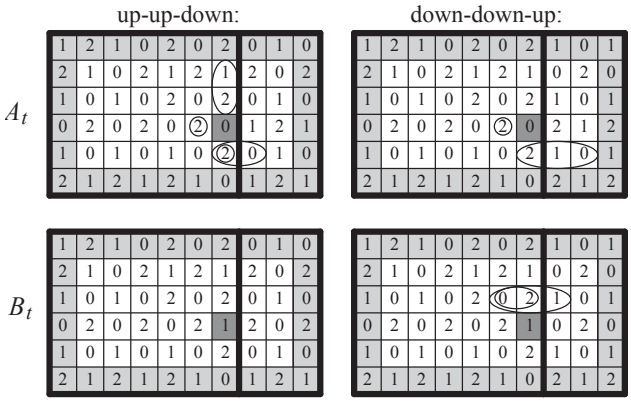

Figure 4: Two colorings A_t and B_t differing by a single flip. Moves increasing $\Phi(A_t, B_t)$ are circled.

the size of the tower intersected with R, the notion of height used in $\mathcal{M}_C$, is the same and equal to $m - x'$. In all, this means that as the desired value of c that starts this tower move occurs with probability $1/3$, regardless of the value of r we have

$$\begin{aligned}
&\mathbb{E}[\Delta\Phi_t \mid (x,y) = (m-1, y')] \\
&\quad = \frac{1}{3}\left(\frac{1}{2(m-x')}(-1)\right. \\
&\quad\quad \left. + \left(\frac{1}{2(m-x'-1)} - \frac{1}{2(m-x')}\right)(m-x'-1)\right) \\
&\quad = \frac{1}{6}\left(\frac{-1}{m-x'} + 1 - \frac{m-x-1}{m-x'}\right) \\
&\quad = \frac{1}{6}\left(1 - \frac{m-x'}{m-x'}\right) = 0.
\end{aligned}$$

It remains to consider moves beginning at $(m-1, \overline{y})$, its neighbors, and in the random extension in rows $\overline{y}$, $\overline{y}-1$, and $\overline{y}+1$. We split these up into three cases: Moves beginning in row $\overline{y}-1$, moves beginning in row $\overline{y}$, and moves beginning in row $\overline{y}+1$.

Consider moves beginning in row $\overline{y}-1$; the analysis for moves beginning in row $\overline{y}+1$ will be symmetric. We differentiate into two cases, whether $(m-1, \overline{y}-2)$ and $(m-2, \overline{y}-1)$ have the same colors or different colors. First, suppose they have the same color, and without loss of generality suppose it is color 0; see Figure 4. If $(m-1, \overline{y}-1)$ is selected by the coupling, a move occurs only if $c = 1$. For $r = u$, in A_t a flip will occur and in B_t a tower move of height 2 in the upward direction will occur if $p < 1/4$. If both moves occur, the change in distance is -1, while if only the flip occurs, the change in distance is $+1$. Altogether, this yields

$$\begin{aligned}
&\mathbb{E}[\Delta\Phi_t \mid (r,x,y,c) = (u, m-1, \overline{y}-1, 1)] \\
&\qquad = \frac{1}{4}(-1) + \left(1 - \frac{1}{4}\right)(+1) = \frac{1}{2}.
\end{aligned}$$

In the same case for $r = d$, in A_t there is a tower whose intersection with R is height 1 while no move occurs in B_t, increasing the distance by 1.

$$\mathbb{E}[\Delta\Phi_t \mid (r,x,y,c) = (d, m-1, \overline{y}-1, 1)] = +\frac{1}{2}.$$

Finally, we must examine moves for $(x,y) = (m, \overline{y}-1)$ and $(x,y) = (m+1, \overline{y}-1)$. Inspection shows the second of these produces no moves in A_t or B_t. For $(x,y) = (m, \overline{y}-1)$, if $r = d$ there are no moves in either, while if $r = u$ then there is a tower in A_t whose intersection with R is height 1 while no such move occurs in B_t, increasing the distance by one.

$$\mathbb{E}[\Delta\Phi_t \mid (r,x,y,c) = (u, m, \overline{y}-1, 2)] = +\frac{1}{2}.$$

Altogether, noting that each extension occurs with probability $1/2$, each color is selected with probability $1/3$, and the correct value of $x \in \{m-1, m, m+1\}$ is selected with probability $1/3$, it follows that in the case where $(m-1, \overline{y}-2)$ and $(m-2, \overline{y}-1)$ have the same color,

$$\begin{aligned}
&\mathbb{E}[\Delta\Phi_t \mid (x,y) \in \{m-1, m, m+1\} \times \{\overline{y}-1\}] \\
&\qquad = \frac{1}{18} \cdot \left(\frac{1}{2} + \frac{1}{2} + \frac{1}{2}\right) = \frac{1}{12}.
\end{aligned}$$

Next, consider the case where $(m-1, \overline{y}-2)$ and $(m-2, \overline{y}-1)$ have different colors, where without loss of generality $(m-1, \overline{y}-2)$ has color 1 and $(m-2, \overline{y}-1)$ has color 0. Note a vertically reflected version of this case can be seen in row $\overline{y}+1$ in Figure 4. In A_t for value $r = d$, there are no nonstationary choices of $(x,y) \in \{m-1, m, m+1\} \times \{\overline{y}-1\}$, and the same holds for B_t and $r = u$. Meanwhile the moves $(d, m-1, \overline{y}-1, 1, p)$ and $(d, m, \overline{y}-1, 2, p)$ both begin left-going towers in B_t. These tower moves may end at a fixed boundary of R and be rejected, but in the worst case they do not so we assume this happens. Both occur if $p < 1/2h$ and increase the distance by h. In A_t the move $(u, m-1, \overline{y}-1, 0, p)$ begins a down-going tower that in the worst case is not rejected and is of some height h'; the move occurs if $p < 1/2h'$, and increases the distance by h'. In total, noting that each extension occurs with probability $1/2$, each color is selected with probability $1/3$, and the correct value for $x \in \{m-1, m, m+1\}$ is selected with probability

1/3, we see that in the case where $(m-1, \overline{y}-2)$ and $(m-2, \overline{y}-1)$ have different colors,

$$\mathbb{E}[\Delta\Phi_t \mid (x,y) \in \{m-1, m, m+1\} \times \{\overline{y}-1\}] \leq \frac{1}{18} \cdot \left(\frac{1}{2h}(h) + \frac{1}{2h}(h) + \frac{1}{2h'}(h') \right) = \frac{1}{12}.$$

We have shown that no matter the colors of $(m-1, \overline{y}-2)$ and $(m-2, \overline{y}-1)$, it is true that

$$(*)\ \mathbb{E}[\Delta\Phi_t \mid (x,y) \in \{m-1, m, m+1\} \times \{\overline{y}-1\}] \leq \frac{1}{12}.$$

The same analysis hold for moves occurring in row $\overline{y}+1$:

$$(*)\ \mathbb{E}[\Delta\Phi_t \mid (x,y) \in \{m-1, m, m+1\} \times \{\overline{y}+1\}] \leq \frac{1}{12}.$$

It remains to consider moves in row $\overline{y}$ where $x \in \{m-2, m-1, m, m+1\}$. Consider a move beginning at $(m-2, \overline{y})$. If this cell's three neighbors other than $(m-1, \overline{y})$ are all the same, suppose without loss of generality they are equal to 0; see Figure 4. Then for a selection of $(m-2, \overline{y}, 1)$, in A_t a flip occurs, and in B_t a tower whose intersection with R is height 2 going right occurs if $p < 1/2h$, regardless of the value of r. If both moves occur the change in distance is -1, while if only the flip occurs the change in distance in 1:

$$\mathbb{E}[\Delta\Phi_t \mid (x,y,c) = (m-2, \overline{y}, 1)] = \frac{1}{4}(-1) + \left(1 - \frac{1}{4}\right)(+1) = \frac{1}{2}.$$

If $(m-2, \overline{y})$'s three neighbors other than $(m-1, y)$ are not all the same, then two must be the same and the third different; we suppose two are color 0, and the third is color 1. In A_t, cell $(m-2, y)$'s fourth neighbor is also color 0, and so (in the worst case) there is a tower of height h beginning at $(m-2, \overline{y})$ going in the direction of its neighbor with color 1. Meanwhile in B_t, cell $(m-2, \overline{y})$ has two neighbors with color 1 and two neighbors with color 0, so no move occurs. It follows that also in this case,

$$\mathbb{E}[\Delta\Phi_t \mid (x,y,c) = (m-2, \overline{y}, 1)] \leq \frac{1}{2h}(h) = \frac{1}{2}.$$

In both cases, noting that the value of c is selected with probability 1/3, it follows that

$$(*) \qquad \mathbb{E}[\Delta\Phi_t \mid (x,y) = (m-2, \overline{y})] \leq \frac{1}{3} \cdot \frac{1}{2} = \frac{1}{6}.$$

Finally, we consider moves in row $\overline{y}$ beginning in columns $m-1$, m, and $m+1$. Inspection shows all such moves will coalesce the two chains, and the nonstationary moves that do so are precisely those discussed in the analysis of the reversibility of boundary flips in the previous section. Across both A_t and B_t and both extensions, there are two flips and four towers whose intersection with R is of height 1 that coalesce the two chains, each occurring for a different (f, x, c) triplet; each such triplet occurs with probability 1/18.

$$(*) \qquad \mathbb{E}[\Delta\Phi_t \mid (x,y) \in \{m-1, m, m+1\} \times \{\overline{y}\}] = \frac{1}{18}\left(2(1) + 4\left(\frac{1}{2}\right)\right)(-1) = -\frac{2}{9}.$$

For each case above, each choice of (x, y) occurs with probability $\frac{1}{(m+1)(n-2)}$. Conditioning on possible values of (x, y) and using the above results labeled with $(*)$ that summarize all nonzero contributions to the expected change in distance,

$$\mathbb{E}[\Delta\Phi_t] \leq \frac{1}{(m+1)(n-2)}\left(3 \cdot \frac{1}{12} + 3 \cdot \frac{1}{12} + 1 \cdot \frac{1}{6} - 3 \cdot \frac{2}{9}\right) = 0.$$

We have shown that at iteration t, Markov chains $\mathcal{A}$ and $\mathcal{B}$ get no further apart in expectation according to distance function Φ, concluding the proof. □

THEOREM 3.1. *On $m \times n$ rectangles with one free boundary, $\mathcal{M}_C$ is rapidly mixing.*

Proof. We apply the path coupling theorem (Theorem 2.1). $\mathcal{M}_C$ is ergodic. Metric Φ takes on values in $[0, mn^2 + m^2 n]$ and satisfies the stated path condition, with U being the set of all pairs of colorings differing by a single flip. Whenever A_t, B_t differ by a single flip, $\mathbb{E}[\Delta\Phi_t \mid A_t, B_t] \leq 0$ (Lemma 3.1). It remains to check α; it follows from the analysis in Section 6.B of [9], using a height function argument and noting that the height value at every fixed boundary cell will be the same, that $\mathbb{P}(\Delta\Phi_t \neq 0) > 1/(24(m+1)(n-2)(m+n))$. We've now verified that all the conditions of Theorem 2.1 hold, so the mixing time of $\mathcal{M}_C$ satisfies

$$\tau(\varepsilon) \leq \left\lceil \frac{eB^2}{\alpha} \right\rceil \lceil \log(\varepsilon^{-1}) \rceil = O\left((m^3 n^6 + m^6 n^3) \log\left(\varepsilon^{-1}\right)\right).$$

□

COROLLARY 3.1. *Glauber dynamics on rectangle R with one free boundary is rapidly mixing.*

Proof. This corollary follows directly from work of Randall and Tetali [20] comparing the mixing time of fixed-boundary chains using the comparison method of Diaconis and Saloff-Coste [4]. □

(a)

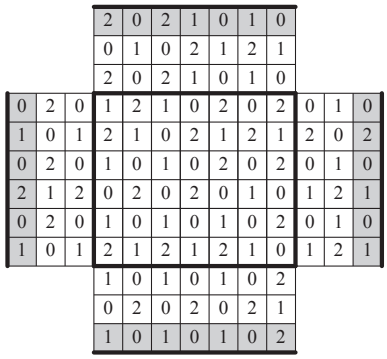

(b)

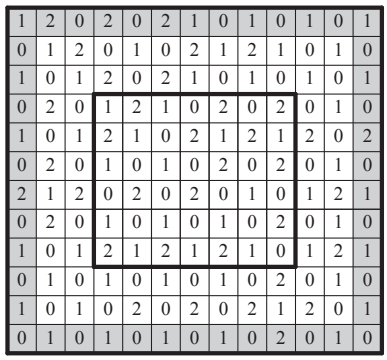

Figure 5: (a) A free boundary 3-coloring and the sides of one of its 16 extensions. (b) The corners of the extension.

3.2 Rectangles with free boundary conditions.

We are now prepared to tackle rectangles where all four sides have free boundary conditions. We extend each side randomly as before. A key insight was to extend each side independently, leading to 16 possible random extensions. An important consideration is that the corners must be carefully completed so $\mathcal{M}_C$ is reversible and mixes rapidly. We show this in what follows.

Random extensions. On the right side of R, extend R with coloring χ to R' with coloring χ' as above. Rectangle R is also extended three cells along all remaining sides in the same way; independently, the values of coloring χ on the boundary row or column are repeated, in order moving away from R, incremented, decremented, incremented (mod 3) with probability 1/2 (UUD), and decremented, incremented, decremented (mod 3) with the remaining probability 1/2 (DDU). This results in 16 possible configurations for the sides of the random extension, each occurring with equal probability. One is shown in Figure 5(a).

Given χ' on the sides of the extension, corners are completed deterministically. If two sides meeting at a corner have opposite UUD and DDU configurations, four colors within the corner are uniquely determined because χ' must be a valid 3-coloring. In other cases, the corner cell closest to R should be given the unique color different from the color in the closest corner of R that still yields a valid 3-coloring. In both cases, the remaining colors in the corners will not affect any Markov chain moves intersecting R, so can be completed arbitrarily; we choose to complete them canonically by repeatedly giving each cell with at least two colored neighbors the lowest admissible color. See Figure 5(b).

Markov chain $\mathcal{M}_C$. In the case of four free boundaries, $\mathcal{M}_C$ is nearly identical to the previous section. One of sixteen possible extensions is selected uniformly at random, and then a move of $\widehat{M}_{eul}$ is performed on the random extension.

Suppose R is $m \times n$; as before, we label the lower left cell of R as $(0,0)$, the lower right cell as $(m-1,0)$, the upper left cell as $(0,n-1)$, and the upper right cell as $(m-1,n-1)$. Rectangle R' then has coordinates ranging from -3 to $m+2$ in the x direction and -3 to $n+2$ in the y direction.

Starting at any initial 3-coloring χ_0 of R, iterate:

- Choose, uniformly at random, a four character string $f \in \{u,d\}^4$.
- Extend coloring χ_i of R to a coloring χ_i' of R'. If the first character of f is u, the top of the random extension has an UUD configuration; if it is d, the extension is DDU. The same holds for the left, bottom, and right sides of R and the second, third, and fourth characters of f, respectively.
- Starting from coloring χ_i' of rectangle R', apply a single iteration of Markov chain $\widehat{\mathcal{M}}_{eul}$:
  - Choose, uniformly at random,
    $$(x,y,c,p) \in \{-2,-1,...,m+1\} \times \{-2,-1,...,n+1\} \times \{0,1,2\} \times (0,1).$$
  - If $(x,y) \in R$, $\chi_i'(x,y) \neq c$, and no neighbors of (x,y) have color c in χ_i', then recolor (x,y) with color c.
  - Else, if (x,y) is the start of a tower of color c whose intersection with R has size $h > 0$, make the tower move if $p < \frac{1}{2h}$.
- Let χ_{i+1} be the resulting coloring of the subrectangle R of R'.

Ergodicity and reversibility of $\mathcal{M}_C$. Markov chain $\mathcal{M}_C$ includes all moves of irreducible Markov chain $\widetilde{M}$ of [9], which acts on the same state space as $\mathcal{M}_C$, so $\mathcal{M}_C$ is irreducible. As it has self loops, $\mathcal{M}_C$ is also aperiodic and thus ergodic. That $\mathcal{M}_C$ is reversible on free boundary rectangles follows from the arguments for reversibility of $\mathcal{M}_C$ on rectangles with one free and three fixed boundaries in Section 3.1. The main difference is the value q, the probability of picking any given location and color, which is now $q = 1/(3(m+4)(n+4))$. However, there are also two additional cases

for the types of moves that two 3-colorings could differ by: colorings σ and τ could differ by flips and towers at corners of R.

Corner flip. Consider two 3-colorings σ and τ of R, differing at a single location that is one of the four corners of R. Suppose, without loss of generality, this occurs at the bottom left corner $(0,0)$, $\sigma(0,0) = 1$, and $\tau(0,0) = 2$. Cells $(0,1)$ and $(1,0)$ must then have color 0 in both tilings. The first and last characters of f will have no effect on any moves between σ and τ, and for any given values of these characters (labeled as $*$'s below) there are five moves that will transform σ to τ, specifically four towers whose intersection with R is height 1 and a single flip: $(*du*, 0, 0, 2, p)$ for $p < 1/2$; $(*\,u\,d*, 0, 0, 2, p)$ for $p < 1/2$; $(*\,d\,d*, 0, 0, 2, p)$ for any p; $(*dd*, -1, 0, 1, p)$ for $p < 1/2$; and $(*dd*, 0, -1, 1, p)$ for $p < 1/2$. Similarly, there are five moves that transform τ to σ. Altogether, it follows that

$$\begin{aligned}\mathbb{P}(\sigma,\tau) &= \frac{4}{2\cdot 2\cdot (m+4)(n+4)\cdot 3\cdot 2} \\ &\quad + \frac{1}{2\cdot 2\cdot (m+4)(n+4)\cdot 3} \\ &= \frac{3q}{4} = \mathbb{P}(\tau,\sigma).\end{aligned}$$

Tower abutting and adjacent to the free boundary. Consider two 3-colorings σ and τ of R, differing by a tower of height $h > 1$ that is both abutting and adjacent to a boundary of R (that is, it contains a corner of R). Without loss of generality, suppose this tower stretches in the horizontal direction from location $(0,0)$ to location $(h-1,0)$, the colors of the tower cells increase from left to right (mod 3), $\sigma(0,0) = 0$, $\tau(0,0) = 1$, $\sigma(h-1,0) = c_1$ and $\tau(h-1,0) = c_1 + 1 (\text{mod } 3)$. For any given first and last characters of f, there are two moves that transition σ to τ: $(*\ d\ d\ *, 0, 0, 1, p)$ for $p < 1/2h$ and $(*\ d\ d\ *, -1, 0, 0, p)$ for $p < 1/2h$. Similarly there are two moves that transform τ to σ, namely $(*\ u\ u\ *, h-1, 0, c_1, p)$ for $p < 1/2h$ and $(*\,d\,u*, h-1, 0, c_1, p)$ for $p < 1/2h$. This implies

$$\mathbb{P}(\sigma,\tau) = \frac{2}{2\cdot 2\cdot (m+4)(n+4)\cdot 3\cdot 2h} = \frac{q}{4h} = \mathbb{P}(\tau,\sigma).$$

The probabilities of all moves, including towers and flips at the corners of R, can be seen in the first two columns of Table 2. It follows that $\mathcal{M}_C$ converges to the uniform distribution over (free boundary) 3-colorings of R, as desired.

Relating $\mathcal{M}_C$ to previous Markov chains. Goldberg et al. presented a tower-based Markov chain $\widetilde{M}$ for 3-colorings of rectangles with free boundary conditions [9]. Their probabilities for the various moves enabled a path-coupling proof, but they did not provide any justification for the probabilities chosen. Here we see the transitions probabilities of $\widetilde{M}$ are identical to those we derived for $\mathcal{M}_C$, up to a small multiplicative factor that arises because R is smaller than R' and $\widetilde{M}$ has higher stationary probability. See Table 2, where $q = 1/(3(m+4)(n+4))$ and $s = 1/(12mn)$.

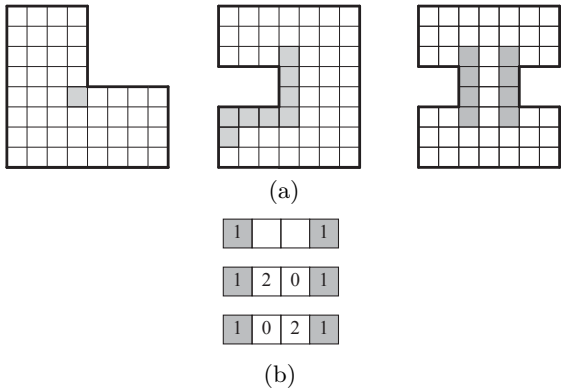

Figure 6: (a) Three regions we can sample efficiently; gray cells must have fixed colors. (b) The two possible 3-colorings are not connected by flips, so $\mathcal{M}_C$ is not irreducible.

Rapid mixing of $\mathcal{M}_C$. As $\widetilde{\mathcal{M}}$ is rapidly mixing, $\mathcal{M}_C$ is as well with just an additional constant multiplicative factor. Goldberg et al. [9] show that Glauber dynamics is also rapidly mixing using the comparison method.

3.3 Regions with mixed boundary conditions.

In Section 3.1, we examined $\mathcal{M}_C$ on a rectangle with some boundary cells fixed and some free; this is an example of *mixed boundary constraints*. For a more general statement, we only require that R be simply connected, the boundary cell at any reflex corner be fixed, and $\mathcal{M}_C$ be irreducible, that is, the state space of 3-colorings for the region and its boundary constraints must be connected by flip moves. We will discuss a necessary and sufficient condition for $\mathcal{M}_C$ to be irreducible later in this section. Some regions that satisfy our conditions, with fixed cells gray, are shown in Figure 6(a).

Random extensions. Markov chain $\mathcal{M}_C$ samples from mixed boundary 3-colorings in the same way as in the more restrictive cases in the previous sections. Given a coloring χ of a region G, we extend to G' and χ' probabilistically, with each free boundary side receiving an UUD or DDU configuration with probability 1/2; we make a fixed-boundary move on this larger region; and we restrict back to G. However, for shapes with reflex corners, this is complicated by the fact that extensions

Type of move	Probability in $\mathcal{M}_C$	Probability in $\widetilde{\mathcal{M}}$ of [9]
Interior flip	q	s
Boundary flip	q	s
Corner flip	$\frac{3q}{4}$	$\frac{3s}{4}$
Interior tower, height $h > 1$	$\frac{q}{2h}$	$\frac{s}{2h}$
Tower abutting the boundary, height $h > 1$	$\frac{q}{2h}$	$\frac{s}{2h}$
Tower adjacent to the boundary, height $h > 1$	$\frac{q}{4h}$	$\frac{s}{4h}$
Tower abutting and adjacent to the boundary, height $h > 1$	$\frac{q}{4h}$	$\frac{s}{4h}$

Table 2: Types of moves for rectangles with free boundaries, and the probability with which they occur in $\mathcal{M}_C$ and in $\widetilde{\mathcal{M}}$ from [9]. Here $q = \frac{1}{3(m+4)(n+4)}$ and $s = \frac{1}{12mn}$.

of the sides adjacent to the reflex corner will overlap. We deal with this by considering combinatorial extensions, and region G', while locally planar, doesn't have a global embedding into the plane, or even into three-dimensional space. We can define $\mathcal{M}_C$ simply by considering only the probabilities of moves restricted to the original region, which are the same as listed in Table 2. We now let towers *start* and *end* in cells adjacent to the boundary, as Goldberg et al. did in their definition of $\widetilde{\mathcal{M}}$ [9].

Markov chain $\mathcal{M}_C$. Starting at any valid 3-coloring χ_0 of G,

- Choose, uniformly at random, a cell g of G, a color $c \in \{0, 1, 2\}$, and a value $p \in (0, 1)$.
- If g is not a fixed-color boundary cell, $\chi'_i(g) \neq c$, and no neighbors of g have color c in χ'_i, then recolor g with color c with probability 1 if g is not a convex corner of G and probability 3/4 if g is a convex corner of G.
- Else, if g starts a tower with start color c and height $h \geq 2$ not containing any fixed-color boundary cells, make the tower move if $p < 1/bh$, where $b = 4$ for a boundary adjacent tower and $b = 2$ otherwise.

Despite stating this Markov chain in a way that doesn't use randomized extensions, we note it was our insights into randomized extensions that allowed us to realize sampling from mixed extensions was possible, determine what the conditions on non-convex regions must be to allow such sampling to be efficient, and decide the probabilities of different types of moves.

Conditions for Irreducibility of $\mathcal{M}_C$. As Figure 6(b) demonstrates, there are some mixed boundary conditions for which $\mathcal{M}_C$ is not irreducible. We present a necessary and sufficient condition for $\mathcal{M}_C$ with mixed boundary constraints to be irreducible.

It is well-known [9, 14] that grid 3-colorings can be mapped to height functions. One can do this by choosing any cell and fixing its height h to be its color. Then heights of remaining cells are determined according to the following rules. If cell a is immediately left of cell b and a's color is incremented (mod 3) to obtain b's color, then $h(a) + 1 = h(b)$, while if a's color is decremented (mod 3) to obtain b's color, then $h(a) - 1 = h(b)$. The same rule holds if a is immediately above b. This height function will be consistent in the sense that neighboring cells will differ in height by 1 and $h(a) \cong \chi(a) (\text{mod } 3)$. There may be many height functions depending on the chosen a starting point, but any two h and h' will satisfy $h(g) - h'(g) = 3k$ for all g and for a fixed constant k.

A mixed boundary condition for S is *height consistent* if there is a height function h^* on the fixed boundary cells of S such that for any valid 3-coloring of R, there is a mapping of this coloring to a height function h that agrees with h^* on all fixed boundary cells. We note that height consistency can depend on the values of the fixed boundary cells; the example in Figure 6(b) is not height consistent, but if one of its fixed cells had a different value it would be.

LEMMA 3.2. *Mixed boundary conditions for a region R are height consistent if and only if $\mathcal{M}_C$ is irreducible.*

Proof. This follows from the irreducibility argument of [9], with a few minor modifications to account for fixed boundary constraints. It is precisely the fact that every 3-coloring has the same height values at all fixed boundary locations that enables this proof to go through. For boundary conditions that are not height consistent, the state space Ω is not connected because no moves change the height values of fixed color boundary cells. □

Ergodicity and reversibility of $\mathcal{M}_C$. As $\mathcal{M}_C$ has self loops it is aperiodic, so if the mixed boundary conditions are height consistent then it is ergodic. If not, $\mathcal{M}_C$ remains ergodic on each connected component of its state space. Reversibility of $\mathcal{M}_C$ when it is defined as above is straightforward; moves are defined by specific probabilities that won't change.

Rapid mixing of $\mathcal{M}_C$. This follows from our previous work on $\mathcal{M}_C$.

THEOREM 3.2. *If $\mathcal{M}_C$ is irreducible on a region R of area A with mixed boundary conditions such that every reflex cell is fixed, then it is rapidly mixing.*

Proof. We apply the path coupling theorem (Theorem 2.1) using the same distance function as above, first verifying the hypotheses hold. $\mathcal{M}_C$ is ergodic. Metric Φ takes on values in $[0, A^2]$, and satisfies the stated path condition with U being the set of all pairs of colorings differing by a single flip. If A_t and B_t differ on a single flip, locally it will look like an interior or boundary flip as considered in the analysis of Section 3.1, or a corner flip; this latter case is analyzed in [9]. Possible additional adjacent cells will be fixed, but this will only lead to fewer moves that increase the distance between $\mathcal{A}$ and $\mathcal{B}$. In all cases, we have $\mathbb{E}[\Delta\Phi_t | A_t, B_t] \leq 0$. Bounding the variance of $\Delta\Phi_t$ away from 0 simply requires height consistent mixed boundary conditions, which we have by Lemma 3.2. One can show, adapting the analysis in Section 6.B of [9], that $\mathbb{P}(\Delta\Phi_t \neq 0) \geq \Omega(1/A^2)$. By Theorem 2.1, the mixing time of $\mathcal{M}_C$ satisfies $\tau(\varepsilon) \leq O(A^6 \log(\varepsilon^{-1}))$. □

COROLLARY 3.2. *Whenever Glauber dynamics is irreducible, it is rapidly mixing for simply-connected regions with mixed boundary conditions where all reflex cells are fixed.*

Proof. This follows directly from the comparison argument in [20] and is very similar to the comparison argument given for free-boundary 3-colorings in [9]. □

3.4 Efficiently sampling non-convex regions.

As a corollary to our above result, we present the first method to efficiently sample from any non-convex region. Consider any L-shaped region L, as in the leftmost part of Figure 6(a). Fix the color at L's single reflex vertex uniformly at random. Now, in polynomial time, $\mathcal{M}_C$ can generate a uniformly random 3-coloring of L with this mixed boundary condition. This method results in a uniformly random 3-coloring of L.

4 Lozenge tilings on the triangular lattice

Next, we demonstrate randomized extensions are a general approach that can be applied to additional structures: lozenge tilings with free and mixed boundaries.

4.1 Tilings with free boundaries.

Let E be an $n \times n \times n$ triangle in the triangular lattice. Let σ be any free boundary lozenge tiling of E, where lozenges are allowed to cross the boundary of E but the interior of E must be tiled completely. Of note, E has no fixed boundary lozenge tilings, but exponentially many free boundary lozenge tilings. We efficiently sample free boundary lozenge tilings of E with a Markov chain based on the same principled approach from the previous sections: we build a random extension around E, fix its boundary, do a single move of the tower Markov chain $\widehat{\mathcal{M}}_{loz}$ for fixed-boundary lozenge tilings, and restrict the result back to E.

Random extensions. Given a free-boundary lozenge tiling σ of E, we randomly extend to form a tiling σ' of a region E'. Where random extensions for 3-colorings were formed with "up-up-down" or "down-down-up" configurations, here random extensions are formed with"left-left-right" (LLR) or "right-right-left" (RRL) arrangements of lozenges. However, the extensions of different sides of the triangle are not independent as they are in the 3-coloring case.

Consider any edge e of the triangular lattice that is contained in the boundary of E, and suppose there is no lozenge of σ crossing the boundary of E at e. We define e's *left shadow* as the six faces of the triangle lattice shaded in Figure 7(a), rotated appropriately, and e's right shadow symmetrically as in Figure 7(b). To form a random extension of E, we will choose e's left shadow with probability 1/2 and e's right shadow with probability 1/2, where each shadow is tiled as in Figure 7. For any edge of the triangular lattice on the boundary of E that bisects a lozenge of tiling σ, e's right and left shadows also consist of six faces of the triangular lattice, but slightly different faces; see Figure 7(c) and (d) for these shadows and their tilings. Shadows for edges on other boundaries of E are rotations of this construction. In both cases above, the left shadow of edge e is tiled by two lozenges "shifted" left followed by a lozenge "shifted" right, a configuration we will call "left-left-right" (LLR). Similarly we define a "right-right-left" configuration (RRL).

To construct an extension E' tiled by σ' for E and σ, we extend all edges of the triangular lattice in the boundary of E via their left shadow with probability 1/2 and via their right shadow with the remaining probability 1/2. These extensions are not independent, and all edges are extended in the same direction. An LLR configuration and RRL configuration for one side of E are depicted in (a) and (b), respectively, of Figure 8. Once the sides of the random extension are generated, the corners can be completed deterministically as in Figure 8(b) in the case of an LLR extension, and with a reflection of that tiling for an RRL extension. We call the union of all these tiles E', and fix its boundary. Unlike with 3-colorings, the extension surrounding E is not always the same shape, although it will be very sim-

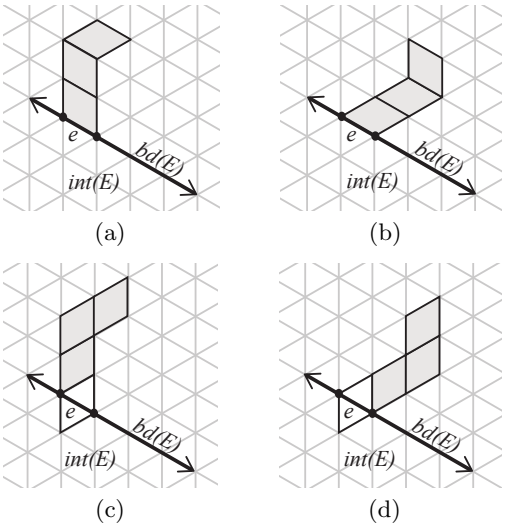

Figure 7: Edge $e \subseteq \partial E$ not crossed by a lozenge, (a) its left shadow, and (b) its right shadow. Edge $e \subseteq \partial E$ crossed by a lozenge, (c) its left shadow, and (d) its right shadow.

ilar. A key fact is that the set of edges of the triangular lattice that have at least one endpoint in the interior of E' is always the same, regardless of the tiling σ of E that was used to generate E' and independent of which extension is chosen. We will use q to denote the number of such edges.

Markov chain $\mathcal{M}_L$. Starting at any free-boundary lozenge tiling σ_0 of E, iterate:

- Choose, uniformly at random, $f \in \{l, r\}$.
- If $f = l$, extend tiling σ_i of E to a tiling σ_i' of a region E' via its LLR extension; if $f = r$, extend tiling σ_i of E via its RRL extension.
- Starting from tiling σ_i' of region E', apply a single iteration of $\widehat{\mathcal{M}}_{loz}$:
  - Choose, uniformly at random, an edge e of the triangular lattice with at least one endpoint in the interior of E'.
  - Choose, uniformly at random, $p \in \{0, 1\}$.
  - If e starts a tower whose intersection with E is of height $h > 0$, make the tower move if $p < 1/2h$.
- Let σ_{i+1} be the resulting free-boundary tiling of E.

We briefly clarify what we mean by the height of a tower's intersection with E. A single tower move can be seen as the result of a series of rotations of three lozenges comprising a hexagon H; such rotations are shown at the top of Figure 1(b). Our previous definition of the height of a tower is precisely the number of such moves that collectively comprise the tower move. For a tower in tiling σ_i' of E', we say the height of its intersection with E is the number of the moves comprising the tower that have a nontrivial intersection with E.

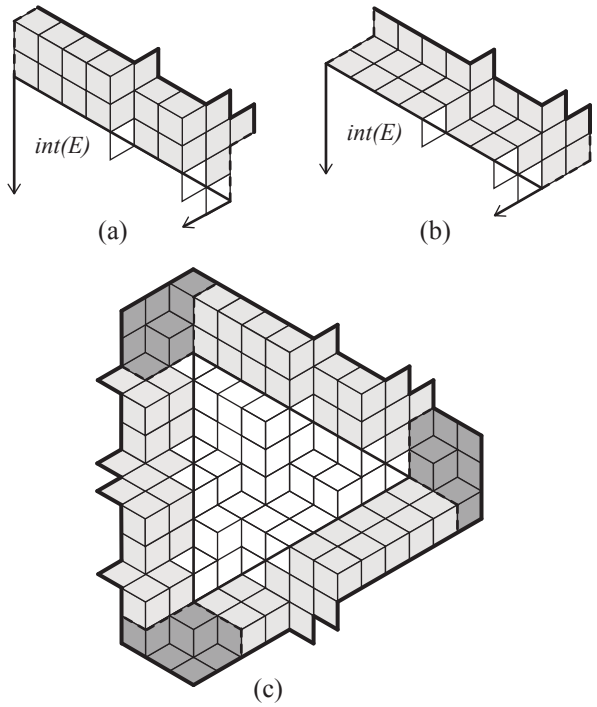

Figure 8: (a) An LLR extension, (b) an RRL extension, and (c) an LLR extension, corners completed deterministically.

Ergodicity and reversibility of $\mathcal{M}_L$. To show $\mathcal{M}_L$ is irreducible, we note it is possible to transform any free boundary lozenge tiling into a tiling entirely with lozenges of the same orientation. As there are also self-loops, $\mathcal{M}_L$ is ergodic. It remains to demonstrate that $\mathcal{M}_L$ is reversible, which will imply it converges to the uniform distribution over free boundary lozenge tilings of E, as desired.

We consider two tilings σ and τ differing by a single move of $\mathcal{M}_L$. There are several cases to consider, based on the type of move that separates σ and τ. In every case below, the tiling in the neighborhood of the difference, where σ and τ agree, is irrelevant.

Interior flip. Consider two tilings σ and τ differing by a single flip (a tower of height 1) entirely contained in E. Regardless of f, in σ there are three edges that start height 1 tower moves yielding τ, and each is selected with probability $1/q$. If selected, each tower move occurs with probability $1/2$. The same is true of moves transforming τ into σ. The probability of moving between two tilings differing by a single interior flip is

$$\mathbb{P}(\sigma, \tau) = 3 \cdot \frac{1}{q} \cdot \frac{1}{2} = \frac{3}{2q} = \mathbb{P}(\tau, \sigma).$$

Interior tower, height $h > 1$. Consider two tilings σ and τ differing by a tower of height $h > 1$ entirely contained in E. Regardless of f, in both σ and τ there is one edge that starts a tower move yielding the

Type of move	$\mathbb{P}(\sigma,\tau) = \mathbb{P}(\tau,\sigma)$ in $\mathcal{M}_L$
Interior flip	$\frac{3}{2q}$
Boundary flip	$\frac{5}{4q}$
Corner flip	$\frac{3}{4q}$
Interior tower, height $h > 1$	$\frac{1}{2qh}$
Tower abutting the boundary, height $h > 1$	$\frac{1}{2qh}$
Tower adjacent to the boundary, height $h > 1$	$\frac{1}{4qh}$
Tower abutting and adjacent to the boundary, height $h > 1$	$\frac{1}{4qh}$

Table 3: Probabilities of different types of moves for $\mathcal{M}_L$ on a triangle E. Value q is the number of edges of the triangular lattice with at least one endpoint in $int(E')$.

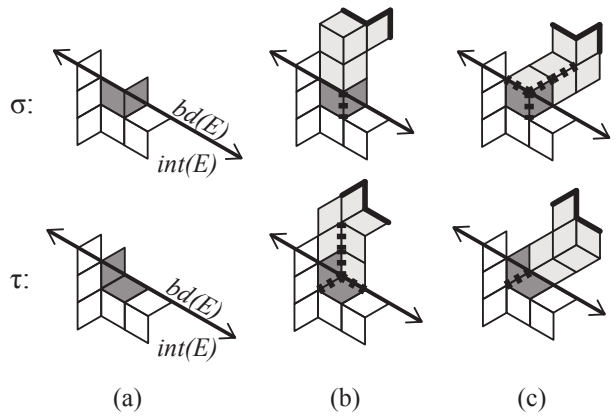

Figure 9: (a) Tilings differing by a boundary flip. (b) Their LLR extensions. (c) Their RRL extensions. Dashed edges start moves transitioning between the two.

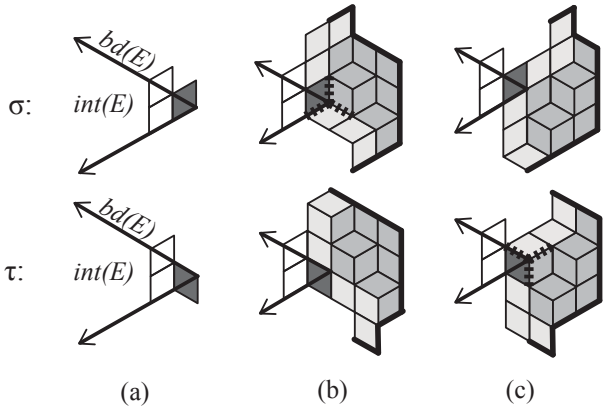

Figure 10: (a) Tilings differing by a corner flip. (b) Their LLR extensions. (c) Their RRL extensions. Dashed edges start moves transitioning between the two.

other. This tower move occurs with probability $1/2h$ if the correct edge is selected. The probability of moving between two tilings differing by an interior tower is

$$\mathbb{P}(\sigma,\tau) = \frac{1}{q}\cdot\frac{1}{2h} = \frac{1}{2qh} = \mathbb{P}(\tau,\sigma).$$

Boundary flip. Consider two tilings σ and τ differing by a single flip along the boundary of E, not at a corner. In both σ and τ, for one extension (LLR or RRL) there is one edge whose selection yields a move between the two if $p < 1/2$, while for the other there are four; see Figure 9. All of these edges start towers whose intersection with E is height 1, although these towers might have larger height in E'. In total, since the probability of selecting the correct extension and edge is $1/2q$, the probability of moving between these two tilings is

$$\mathbb{P}(\sigma,\tau) = 5\cdot\frac{1}{2}\cdot\frac{1}{2q} = \frac{5}{4q} = \mathbb{P}(\tau,\sigma).$$

Corner flip. Consider two tilings σ and τ differing by a single flip at a corner of E; see Figure 10. In both σ and τ, for one extension (LLR or RRL) there are three edge selections which yield a move between the two, while for the other there are none. All of these edges start towers of height 1, and so each move occurs with probability $1/2$. In total, the probability of moving between these two tilings is

$$\mathbb{P}(\sigma,\tau) = 3\cdot\frac{1}{2}\cdot\frac{1}{2q} = \frac{3}{4q} = \mathbb{P}(\tau,\sigma).$$

Tower abutting the boundary, height $h > 1$. Across the two possible extensions (LLR or RRL), there will be two edge selections yielding this move, each starting a tower whose intersection with E is height h; see Figure 11(a), where $h = 2$. The probability of making this move is then

$$\mathbb{P}(\sigma,\tau) = 2\cdot\frac{1}{2h}\cdot\frac{1}{2q} = \frac{1}{2qh} = \mathbb{P}(\tau,\sigma).$$

Tower adjacent to the boundary, height $h > 1$. Across the two possible extensions (LLR or RRL), there will be one edge selection yielding this move, and it starts a tower of height h; see Figure 11(b), where

$h = 2$. The probability of making this move is then

$$\mathbb{P}(\sigma, \tau) = \frac{1}{2h} \cdot \frac{1}{2q} = \frac{1}{4qh} = \mathbb{P}(\tau, \sigma).$$

Tower abutting and adjacent to the boundary, height $h > 1$. Across the two possible extensions (LLR or RRL), there will be one edge selection yielding this move, and it starts a tower of height h; see Figure 11(c), where $h = 3$. The probability of making this move is then

$$\mathbb{P}(\sigma, \tau) = \frac{1}{2h} \cdot \frac{1}{2q} = \frac{1}{4qh} = \mathbb{P}(\tau, \sigma).$$

We also briefly note that tower moves for towers abutting two boundaries, whether adjacent to a boundary or not, occur with the same probabilities as towers abutting one boundary. A summary of the probabilities with which different types of moves occur can be found in Table 3, where q is equal to the number of edges of the triangular lattice with a least one endpoint in $int(E')$. Recall q is the same for all random extensions of E, regardless of tiling σ of E. While these move probabilities arise naturally from randomized extensions, it would be difficult to correctly guess these specific values without using computational techniques.

Rapid mixing of $\mathcal{M}_L$. We use a path coupling argument. Consider a joint process $(\mathcal{A}, \mathcal{B})$ on $\Omega \times \Omega$, where each of $\mathcal{A}$ and $\mathcal{B}$ is a copy of $\mathcal{M}_L$. We faithfully couple by making the same choice of (f, e, p) for both $\mathcal{A}$ and $\mathcal{B}$ at each iteration. The marginal distributions of these two coupled chains at iteration t are denoted by A_t and B_t, respectively. We define distance Φ between two free boundary lozenge tilings of E to be the minimum number of flips needed to transform one tilling to the other. We say $\Phi_t := \Phi(A_t, B_t)$, and let $\Delta\Phi_t = \Phi_{t-1} - \Phi_t$. The first crucial step in the path coupling argument is to demonstrate that whenever A_t and B_t differ by one flip, in expectation $\mathcal{A}$ and $\mathcal{B}$ get no farther apart after one iteration of the coupling. That is, we show $\mathbb{E}[\Delta\Phi_t] \leq 0$.

We begin with a structural lemma about free boundary lozenge tilings and our random extensions. For any given tower, we define the *center line* of the tower to be the edge e that starts the tower, extended to reach the boundary of the tower. We will use ∂E to denote the boundary of triangle E. We say that a tower *crosses ∂E at vertex v* if v is the first vertex of ∂E that the tower's center line contains, when it is traversed starting with e. Let the *state* of edge $a \in \partial E$ denote whether or not it is crossed by a lozenge or not; we will talk about a and b having the same state or different states. Consider the following lemma, where we use the expression *outside of E* to mean an edge whose interior is entirely disjoint from $E \cup \partial E$.

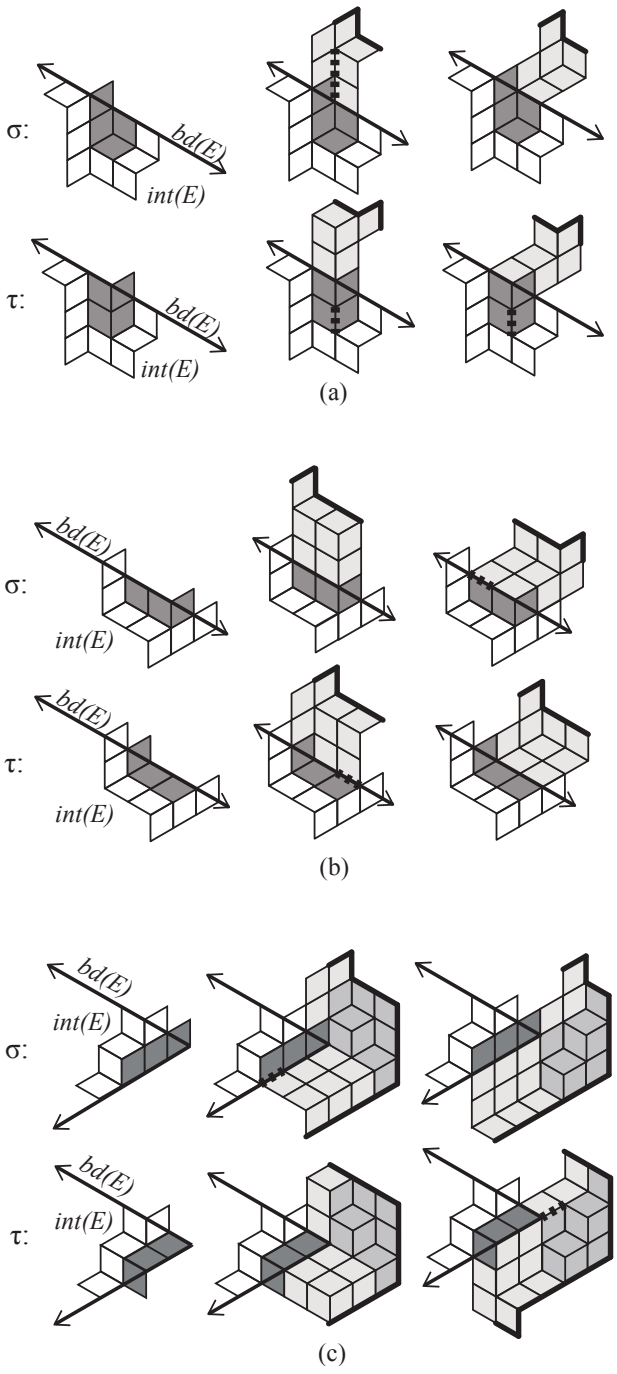

Figure 11: (a) Tilings differing by a tower abutting the boundary. Center: Their LLR extensions. Right: Their RRL extensions. (b) Tilings differing by a tower adjacent to the boundary. Center: Their LLR extensions. Right: Their RRL extensions. (c) Tilings differing by a tower abutting and adjacent to the boundary. Center: Their LLR extensions. Right: Their RRL extensions. Dashed edges start towers that transition σ to τ or vice versa.

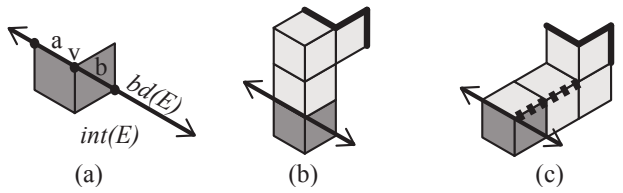

Figure 12: (a) A sample tiling. (b) An LLR extension. (c) An RRL extension. Dashed edges start the only towers that begin outside E and cross ∂E at v.

LEMMA 4.1. *Let σ be any lozenge tiling of E. Let σ^L and σ^R denote σ's LLR and RRL extensions. Let v be any vertex in ∂E that is not a corner, and let a and b be the two edges in ∂E adjacent to v. If a and b have different states, then across both σ^R and σ^L there are at most two towers started by an edge outside of E that cross ∂E at v. If a and b have the same state, there are no towers started by edges outside E that cross ∂E at v.*

Proof. We first show if a and b have the same state, there can be no tower that crosses ∂E at v. Any vertex v contained in the interior of the center line of a tower must be incident to at least two lozenges with a 120° angle at v, and these lozenges must be consecutive around v. If a and b are both crossed by a lozenge of σ, v cannot possibly be on the center line of any tower; the lozenges crossing a and b are the only two incident on v that have a 120° angle at v, and they are not consecutive around v. If a and b are both not crossed by a lozenge of σ then it is possible for them to be on the center line of some tower, but only if that tower runs along ∂E. In this case, v cannot possibly be the first vertex of the tower's center line in ∂E. This concludes our proof of the first statement.

Suppose, without loss of generality, b is clockwise from a, a is not crossed by a lozenge of σ, and b is crossed by a lozenge of σ; see Figure 12(a). In σ_R, there are two edges outside of E that start a tower whose center line crosses ∂E at v, shown in Figure 12(c) where these two edges are dashed. In σ^L, there are no such edges; see Figure 12(b). While in σ^L there is a tower whose center line crosses ∂E at v, it is started by an edge inside E, not an edge outside E. □

LEMMA 4.2. *Suppose A_t and B_t differ by one flip of a hexagon H, that is, $\Phi(A_t, B_t) = 1$. After one iteration of the joint process $(\mathcal{A}, \mathcal{B})$,*

$$\mathbb{E}[\Delta\Phi_t \mid A_t, B_t] \leq 0.$$

Proof. First, let e' be an edge on the boundary of H, in the interior of E. Edge e' will start a tower of height $h \geq 1$ in exactly one of A_t and B_t, depending on the orientation of the lozenge outside H that is also incident to e'. This tower move leads to an increase in distance of h and occurs with probability $\frac{1}{2h}$. For each such e',

$$\mathbb{E}[\Delta\Phi_t \mid e = e'] = \frac{1}{2h}h = \frac{1}{2}.$$

As was established by Luby et al. [14], such towers are the only moves originating in E that can increase the distance between $\mathcal{A}$ and $\mathcal{B}$ in expectation. Any other moves either are disjoint from H and have the same effect in A_t and B_t or are neutral moves. Neutral moves coalesce A_t and B_t with some probability and increase the distance between them with some probability, but the expected change in distance between A_t and B_t can be calculated to be 0 (we also saw neutral moves in the context of 3-colorings). For moves originating at an edge e in the boundary of E outside of H, the same can be shown to be true in each extension, and thus it is true in the average of the two extensions, as desired.

We now analyze in detail the five possible cases (up to rotations) for the location of H, shown in Figure 13. Recall q is the number of edges that have at least one endpoint inside E', and this number is independent of the tiling σ used to generate E'.

<u>*Case 1:*</u> If H is contained in E and not adjacent to the boundary of E, there are six moves that increase the distance between A_t and B_t. Specifically, these are the six towers started by the edges of the boundary of H, each occurring in exactly one of A_t and B_t with probability $1/2qh$ and leading to an expected increase in distance of h. The total contributions of each of these moves to $\Delta\Phi_t$ is then $(h)1/2qh = 1/2q$. There are also six moves that decrease the distance, specifically the six edges of the triangular lattice in the interior of H, each starting a tower of height 1 in exactly one of A_t and B_t. Each of these moves occurs with probability $1/2q$ and decreases the distance by 1. Altogether,

$$\mathbb{E}[\Delta\phi_t] = 6 \cdot \left(\frac{1}{2q}\right) + 6 \cdot \left(\frac{1}{2q} \cdot (-1)\right) = 0.$$

Any moves originating in the extension in this case will be neutral in expectation, as the random extension is the same for both A_t and B_t.

<u>*Case 2:*</u> If H is contained in E and adjacent to the boundary of E, then in each extension, locally H and its neighborhood are as in Case 1, above. Averaging across both extensions, we find

$$\mathbb{E}[\Delta\phi_t] = \frac{1}{2}(0) + \frac{1}{2}(0) = 0.$$

<u>*Case 3:*</u> Suppose H is bisected by one boundary of E and not adjacent to any other boundary of E.

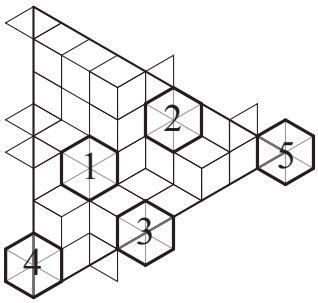

Figure 13: If A_t and B_t differ by the rotation of three lozenges in a single hexagon H, there are five cases to consider for the location of H.

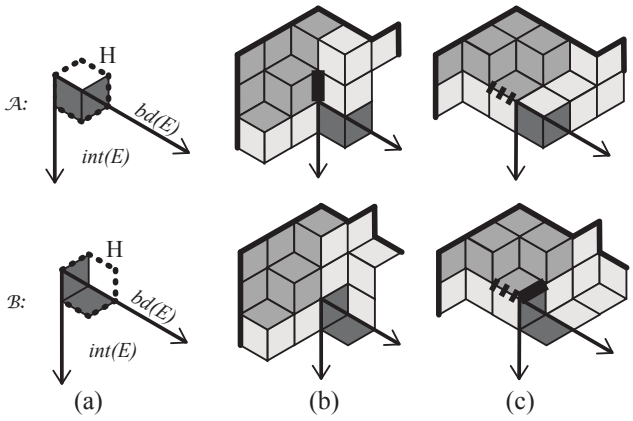

Figure 14: (a) A_t and B_t differ on hexagon H. (b) An LLR extension. (c) An RRL extension. The dashed edge starts a neutral move, and the bold edges start moves increasing Φ_t.

We first consider moves that coalesce A_t and B_t. In the boundary flip discussion in the previous subsection, we showed that $\mathbb{P}(A_t, B_t) = \mathbb{P}(B_t, A_t) = 5/4q$. This includes both tower moves started by the four edges in $int(H) \cap (E \cup \partial E)$ and tower moves started in the extension near H. Inspection shows these probabilities are disjoint, as moves transforming A_t to B_t are started by different edges than moves transforming B_t to A_t. Each of these transitions decreases the distance between A_t and B_t by one, and so contributes $-5/4q$ to the expected change in distance.

We now consider transitions that cause A_t and B_t to move farther apart. Here there are three boundary edges of H inside E. As above, for each of these three edges, in exactly one of A_t and B_t the edge starts a tower move that occurs with probability $1/2qh$ and increases the distance between A_t and B_t by h. These tower moves occur regardless of which extension σ^L or σ^R is chosen, although this choice may cause the height of a tower in E' to vary. Each edge contributes an expected $1/2q$ to $\Delta\Phi_t$. Again, all other moves started by edges of $E \cup \partial E$ contribute 0 to the expected change in distance.

To complete our examination of moves that might change the distance between A_t and B_t, we consider towers that begin outside E but intersect E. Towers that cross ∂E at H's center point will decrease the distance between A_t and B_t, and have been included in the coalescence probabilities above. Towers that cross ∂E away from H will have the same effect in A_t and B_t and contribute 0 to $\Delta\Phi_t$. Let v be one of the two vertices in $\partial E \cap \partial H$. In exactly one of A_t and B_t, the two edges in ∂E adjacent to v are in the same state. In the other of A_t and B_t, the two edges are in different states. Suppose, without loss of generality, they are in different states in A_t. By Lemma 4.1, across both σ^R and σ^L there are two towers crossing ∂E at v in A_t, while there are no such moves in B_t. Each tower move potentially increases the distance by h, the height of its intersection with E, and occurs with probability $1/2qh \cdot 1/2$. The latter term is present because the correct extension must be chosen for these moves to occur. Each such move increases the distance by $1/4q$ in expectation, and there are four such moves, two for each vertex of $\partial E \cap \partial H$. In total, we see that

$$\mathbb{E}[\Delta\Phi_t] \leq 2 \cdot \frac{-5}{4q} + 3 \cdot \frac{1}{2q} + 4 \cdot \frac{1}{4q} = 0.$$

Case 4: We note here that the analysis is very similar to Case 3, despite the seeming complexity. Again by the boundary flip discussion in the previous subsection, $\mathbb{P}(A_t, B_t) = \mathbb{P}(B_t, A_t) = 5/4q$. These probabilities are disjoint, because moves transforming A_t to B_t are started by different edges than moves transforming B_t to A_t. Each of these transitions decreases the distance between A_t and B_t by one, and so contributes $-5/4q$ to the expected change in distance.

The two edges of ∂H inside E both contribute $1/2q$ to the expected change in distance. If w is the vertex of H farthest from the corner, by Lemma 4.1, in one of A_t and B_t there are two towers that start outside E and cross the boundary at w and in the other there are no such towers. Each of these contributes $1/4q$. Edge $e' \in \partial H \cap \partial E$, in each of the extensions σ^L and σ^R, locally looks like an edge of H inside E. As in Case 2, in each extension, e' starts a tower in exactly one of A_t and B_t, and on average contributes $1/2q$ to the expected change in distance. All other moves started by edges of $E \cup \partial E$ contribute 0 to the expected change in distance. It only remains to consider towers that first cross the boundary of E at corner vertex v.

We can see that there are three possible edges in the extension whose selection yields a tower crossing ∂E at v in at least one of A_t and B_t. All are adjacent to v; see

Figure 14. One of these edges, parallel to the side of E that bisects H and dashed in Figure 14, starts a tower of height $h = 1$ in one of A_t or B_t and a tower of height $h = 2$ in the other. Selecting this edge is a *neutral* move, contributing 0 to the change in distance. The tower moves started by this edge occur in both A_t and B_t, decreasing the distance by one, with probability 1/4q. With probability $1/2q - 1/4q = 1/4q$, only the height 1 tower moves occurs, increasing the distance by one. In total, the contribution to the change in distance from this edge is 0. The two remaining towers crossing ∂E at v each occur in only one extension, with probability $1/2qh \cdot 1/2$, and increase the distance by h, contributing $1/4q$ each. In total, we have

$$\mathbb{E}[\Delta\Phi_t] \leq 2 \cdot \frac{-5}{4q} + 2 \cdot \frac{1}{2q} + 2 \cdot \frac{1}{4q} + \frac{1}{2q} + 2 \cdot \frac{1}{4q} = 0.$$

Case 5: We first consider moves that may increase the distance between A_t and B_t. The single edge of ∂H inside E contributes an expected $1/2q$ to the change in distance as above. For each of the two vertices in $\partial H \cap \partial E$, by Lemma 4.1 and the argument above there are two towers originating in the extension of E that cross ∂E at that vertex. Each of these towers, of which there are four total, two crossing ∂E at each vertex, contributes an expected $1/4q$ to the change in distance.

The only remaining moves that could affect the distance between A_t and B_t are precisely those that coalesce the two chains. As demonstrated in the corner flip discussion of the previous subsection, the probability that A_t becomes B_t is $3/4q$, as is the probability that B_t becomes A_t. The moves transforming A_t into B_t are started by different edges than the moves transforming B_t into A_t, so these probabilities are disjoint. All of these moves change the distance between A_t and B_t by -1. In total, we have

$$\mathbb{E}[\Delta\Phi_t] \leq 1 \cdot \frac{1}{2q} + 4 \cdot \frac{1}{4q} + (-1) \cdot \frac{3}{4q} + (-1) \cdot \frac{1}{3q} = 0.$$

We have seen that for every case of A_t and B_t at distance 1, $\mathbb{E}[\Delta\Phi_t] \leq 0$, as desired. □

THEOREM 4.1. *$\mathcal{M}_L$ is rapidly mixing and efficiently samples free boundary lozenge tilings of E. Its mixing time satisfies $O(n^8 \log(\varepsilon^{-1}))$, where n is the length of one side of E.*

Proof. We apply the path coupling theorem (Theorem 2.1). $\mathcal{M}_L$ is ergodic. Metric Φ takes values in $[0, n^3]$, which follows from adapting the fixed-boundary analysis of [14]. Metric Φ satisfies the stated path condition, with U the pairs of tilings differing by a flip. Whenever A_t and B_t differ by a flip, we showed that $\mathbb{E}[\Delta\Phi_t \mid A_t, B_t] \leq 0$ (Lemma 4.2). Again adapting work of [14], $\mathbb{P}[(\Delta\Phi_t \neq 0] \geq 1/2qh$, as one can always find a move that decreases Φ in expectation, where $q = O(n^2)$ and $h = O(n)$. All of the necessary conditions are met, so by Theorem 2.1, the mixing time of $\mathcal{M}_L$ satisfies

$$\tau(\varepsilon) \leq \left\lceil \frac{eB^2}{\alpha} \right\rceil \lceil \log(\varepsilon^{-1}) \rceil = \left\lceil \frac{e(n^3)^2}{1/4q} \right\rceil \lceil \log(\varepsilon^{-1}) \rceil$$
$$= O(n^9 \log(\varepsilon^{-1})). \qquad \square$$

The results in this section can be generalized to efficiently sample mixed-boundary lozenge tilings of simply connected regions R. We require any corners of angle other than 60° be contained in the fixed boundary. Additionally, we require $\mathcal{M}_L$ be ergodic, that is, the fixed boundary is such that flip moves are sufficient to move from any lozenge tiling to any other.

5 Approximate counting

We have thus far focused only on sampling. Seminal work of Jerrum, Valiant and Vazirani [11] shows how efficient algorithms for sampling can be used to construct efficient algorithms to approximately count, provided the underlying problem is *self-reducible*. Algorithms that require regions to have a specific shape, such as a rectangle in $\mathbb{Z}^2$ or a triangle in the triangular lattice, typically fail to be self-reducible. The reduction requires incrementally fixing the value of a cell (for colorings choosing the most likely color, for example). This requires us to draw samples from nearly square (or triangular) regions where the colors (resp., tiles) on some of the boundary cells have been fixed. Figure 15 gives a typical step of the reduction from sampling to approximate counting. Our mixed boundary Markov chains allow us to sample from exactly those regions that appear in the reduction from sampling to counting. Any reflex angles on the boundary will always be fixed, meaning Theorems 3.2 and 4.1 apply and $\mathcal{M}_C$ and $\mathcal{M}_L$ can be used to sample iteratively. Now both problems are self-reducible as required, and we use the reduction to efficiently count. While there are other known methods to efficiently count lozenge tilings on certain lattice regions with free boundary conditions [15], this new algorithm applies to a broader class of regions. In addition, this approach based on random extensions yields the first efficient algorithm for approximately counting 3-colorings on classes of lattice regions with free boundaries.

References

[1] N. Bhatnagar, S. Greenberg, and D. Randall. The effects of boundary conditions on mixing rates of Markov chains. *Approximation, Randomization and Combinatorial Optimization Algorithms and Techniques*, pages 280–291, 2006.

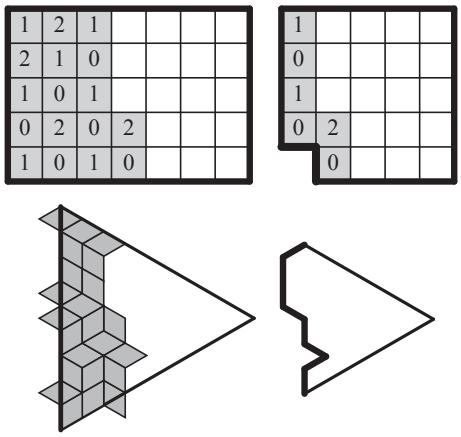

Figure 15: Left: A partial free boundary 3-coloring and lozenge tiling. Right: The mixed boundary regions we sample from in one step in the approximate counting process.

[2] R. Bubley, M. Dyer, C. Greenhill, and M. Jerrum. Approximately counting colorings of small degree graphs. *SIAM Journal of Computing*, 29:387–400, 1999.

[3] H. Cohn, M. Larsen, and J. Propp. The shape of a typical boxed plane partition. *New York Journal of Mathematics*, 4:137–165, 1998.

[4] P. Diaconis and L. Saloff-Coste. Comparison theorems for reversible Markov chains. *The Annals of Applied Probability*, 3:696–730, 1993.

[5] M. Dyer and C. Greenhill. A more rapidly mixing Markov chain for graph colorings. *Random Structures & Algorithms*, 13(3-4):285–317, 1998.

[6] A. Frieze and E. Vigoda. A survey on the use of Markov chains to randomly sample colourings. *Combinatorics, Complexity and Chance*, 2007.

[7] D. Galvin, J. Kahn, D. Randall, and G. Sorkin. Phase co-existence in the 3-coloring model on $\mathbb{Z}^2$. *http://arxiv.org/abs/1210.4232*, 2012.

[8] D. Galvin and D. Randall. Torpid mixing of local Markov chains on 3-colorings of the discrete torus. *18th Symposium on Discrete Algorithms (SODA)*, 2007.

[9] L. A. Goldberg, R. Martin, and M. Paterson. Random sampling of 3-colorings in $\mathbb{Z}^2$. In *Random Structures and Algorithms*, volume 24, pages 279–302, 2004.

[10] T.P. Hayes and E. Vigoda. A non-Markovian coupling for randomly sampling colorings. In *Foundations of Computer Science, 2003. Proceedings. 44th Annual IEEE Symposium on*, pages 618–627, Oct 2003.

[11] M. Jerrum, L. Valiant, and V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoretical Computer Science*, 43:169–188, 1986.

[12] L.C. Lau and M. Molloy. Randomly colouring graphs with girth five and large maximum degree. In JosR. Correa, Alejandro Hevia, and Marcos Kiwi, editors, *LATIN 2006: Theoretical Informatics*, volume 3887 of *Lecture Notes in Computer Science*, pages 665–676. Springer Berlin Heidelberg, 2006.

[13] D. Levin, Y. Peres, and E. Wilmer. *Markov chains and mixing times*. Providence, R.I. American Mathematical Society, 2009. With a chapter on coupling from the past by J.G. Propp and D.B. Wilson.

[14] M. Luby, D. Randall, and A. Sinclair. Markov chain algorithms for planar lattice structures. *SIAM Journal on Computing*, 31:167–192, 2001.

[15] R. Martin and D. Randall. Pfaffian algorithms for sampling routings on regions with free boundary conditions. *3rd International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM) in Lecture Notes in Computer Science*, 1671:257–268, 1999.

[16] F. Martinelli, A. Sinclair, and D. Weitz. Fast mixing for independent sets, colorings and other models on trees. *Proc. 15th ACM/SIAM Symposium on Disccrete Algorithms*, pages 449–458, 2004.

[17] F. Martinelli, A. Sinclair, and D. Weitz. The Ising model on trees: Boundary conditions and mixing time. *Comm. Mathematical Physics*, 250:301334, 2004.

[18] M. Molloy. The Glauber dynamics on colourings of a graph with high girth and maximum degree. In *Proceedings of the Thirty-fourth Annual ACM Symposium on Theory of Computing*, STOC '02, pages 91–98, New York, NY, USA, 2002. ACM.

[19] R. Peled. High-dimensional Lipschitz functions are typically flat. *To appear in Annals of Probability*, 2015.

[20] D. Randall and P. Tetali. Analyzing Glauber dynamics by comparison of Markov chains. *J. Mathematical Physics*, 41:1598–1615, 2000.

[21] A. Sinclair. *Algorithms for Random Generation & Counting: a Markov Chain Approach*. Birkhäuser, Boston, 1993.

[22] E. Vigoda. Improved bounds for sampling colorings. *J. Mathematical Physics*, 41(3):1555–1569, 2000.