

PHASE TRANSITIONS IN RANDOM DYADIC TILINGS AND RECTANGULAR DISSECTIONS*

SARAH CANNON[†], SARAH MIRACLE[‡], AND DANA RANDALL[§]

Abstract. We study *rectangular dissections* of an $n \times n$ lattice region into rectangles of area n , where $n = 2^k$ for an even integer k . We show there is a natural edge-flipping Markov chain that connects the state space. A similar edge-flipping chain is known to connect the state space when restricted to *dyadic tilings*, where each rectangle is required to have the form $R = [s2^u, (s+1)2^u] \times [t2^v, (t+1)2^v]$, where s, t, u and v are nonnegative integers. The mixing time of this Markov chain for general rectangular dissections remains open, while recent work by Cannon, Levin, and Stauffer [4] gave a polynomial upper bound on the mixing time when restricting to dyadic tilings.

We consider a weighted version of these Markov chains where, given a parameter $\lambda > 0$, we would like to generate each rectangular dissection (or dyadic tiling) σ with probability proportional to $\lambda^{|\sigma|}$, where $|\sigma|$ is the total edge length. We show there is a phase transition in the dyadic setting: when $\lambda < 1$, the edge-flipping chain mixes in time $O(n^2)$, and when $\lambda > 1$, the mixing time is $\exp(\Omega(n^2))$. The behavior for general rectangular dissections is more subtle, and even establishing ergodicity of the chain requires a careful inductive argument. As in the dyadic case, we show that the edge-flipping Markov chain for rectangular dissections requires exponential time when $\lambda > 1$. Surprisingly, the chain also requires exponential time when $\lambda < 1$, which we show using a different argument. Simulations suggest that the chain converges quickly at the isolated point $\lambda = 1$, but this case remains open.

Key words. rectangular dissection, dyadic tilings, Markov chains

AMS subject classifications. 60J10, 68W20

1. Introduction. *Rectangular dissections* arise in the study of VLSI layout [6], mapping graphs for floor layouts [18, 23], and routings and placements [26] and have long been of interest to combinatorialists [2, 22]. In each of these applications, a lattice region needs to be partitioned into rectangles whose corners lie on lattice points such that the dissection satisfies some appropriate additional constraints. For example, *equitable rectangular dissections* require that all rectangles in the partition have the same area [9] (see Figure 1). We are interested in understanding what random equitable rectangular dissections look like as well as finding efficient methods for sampling these dissections.

There has also been interest in the special case of *dyadic tilings*, or equitable rectangular dissections into dyadic rectangles. A *dyadic rectangle* is a set of the form

$$R = [s2^u, (s+1)2^u] \times [t2^v, (t+1)2^v]$$

where s, t, u and v are nonnegative integers. A dyadic tiling of the $2^k \times 2^k$ square is a set of 2^k dyadic rectangles, each of area 2^k , whose union is the full square; see Figure 1(b). Janson et al. [10] studied the asymptotics of A_k , the number of dyadic tilings of the $2^k \times 2^k$ square where $k \in \mathbf{Z}^+$. They show that every dyadic tiling must have a fault line, that is, a line bisecting the square in the vertical or horizontal direction which avoids non-trivial intersection with all rectangles in the tilings. This allows them to derive the recurrence $A_k = 2A_{k-1}^2 - A_{k-2}^4$ and show that asymptotically $A_k \sim \phi^{-1}\omega^{2^k}$, where $\phi = (1+\sqrt{5})/2 = 1.6180\dots$ is the golden ratio and $\omega = 1.84454757$ is a constant.

Although equitable partitions of lattice regions into rectangles or triangles have been extensively studied, many fundamental questions remain open. A notable exception is dissections into rectangles with area 2, commonly known as *domino tilings* or the *dimer model* from statistical physics. Researchers have discovered remarkable properties of these tilings, including striking underlying combinatorial structures [11], statistical properties of random tilings [12], and analysis showing various Markov chains for generating them are efficient [7, 15, 19].

*A preliminary version of this paper appeared in the *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 1573–1589, 2015.

[†]College of Computing, Georgia Institute of Technology, Atlanta, GA 30332-0765; sarah.cannon@gatech.edu. Supported in part by a grant from the Simons Foundation (#361047 to Sarah Cannon) and the National Science Foundation Graduate Research Fellowship Program under Grant No. NSF DGE-1650044. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

[‡]Computer and Information Sciences, University of St. Thomas, St. Paul, MN 55105; sarah.miracle@stthomas.edu.

[§]College of Computing, Georgia Institute of Technology, Atlanta, GA 30332-0765; randall@cc.gatech.edu. Supported in part by NSF Awards CCF-1526900 and 1637031.




FIG. 1. (a) An equitable rectangular dissection and (b) a dyadic tiling of the 16×16 square. Shaded rectangles are not dyadic.

Triangular dissections have been explored extensively as well, both when the vertices are in general position and when they are vertices of a planar lattice. On the Cartesian lattice \mathbf{Z}^2 , the problem becomes finding equitable (or *unimodular*) triangulations of a lattice region, where each triangle has area $1/2$. See [14] for an extensive history of work on triangulations.

Interestingly, in each of these cases, a certain “edge-flip” Markov chain has been identified that connects the state space of allowable dissections. For example, for domino tilings, the Markov chain iteratively removes a length 2 edge bordering two dominoes and replaces it with a length 2 edge in the orthogonal direction, effectively replacing two vertical dominoes with two horizontal ones, or vice versa. This chain is known to be rapidly mixing [15, 20, 25]. In the case of dyadic tilings, there is again a natural edge-flip chain that connects the set of possible configurations – if there are two neighboring rectangles in the tiling that share an edge, we can remove that edge and retiling the larger composite rectangle with the edge that bisects it in the orthogonal direction, provided the new tiling is still dyadic (see Figure 3(c),(d)). The mixing rate of this edge-flip chain was left open in [10], although the authors argue that a different nonlocal Markov chain containing additional moves does converge quickly to equilibrium. This edge-flip chain was later shown to be rapidly mixing [4].

Another edge-flip chain also connects the state space of triangulations by replacing an edge bordering two triangles with the edge connecting the triangles’ other two vertices if the quadrilateral formed by their union is convex. The edge-flip chain on triangulations of general point sets has been the subject of much interest in the computational geometry community (see, e.g., [24]). In the unweighted case the chain has only been analyzed when the points are in convex position [16, 17], in which case the triangulations are enumerated by the Catalan numbers.

Recently, Caputo et al. [5] introduced a weighted version of the lattice triangulation problem and discovered remarkable behavior. Each triangulation σ on a finite region of \mathbf{Z}^2 is assigned a weight $\lambda^{|\sigma|}$, where $\lambda > 0$ is some input parameter, with the stationary probabilities proportional to these weights. They conjecture there is a phase transition at $\lambda = 1$ and that when $\lambda < 1$ there are no long-range correlations of the triangles and Markov chains based on local edge flips converge in polynomial time, while when $\lambda > 1$ there will be large regions of aligned long-thin triangles and local Markov chains will require exponential convergence time. Their conjecture is supported by the intuition that when λ is large, triangulations with many long-thin triangles will be favored, and the geometry will force these triangles to align in the same direction. In contrast, when $\lambda < 1$, triangles with large aspect ratio will be preferred, the chain will be rapidly mixing, and there will not be any long-range order. They verify this conjecture when $\lambda > 1$ and when $\lambda < \lambda_0 < 1$ for some suitably small constant λ_0 .

1.1. Sampling Rectangular Dissections. In this paper, we study a weighted version of the equitable rectangular dissection problem and explore the mixing time of an appropriate edge-flip Markov chain. Let $n = 2^k$, for k an even integer, and let Λ_n be the $n \times n$ lattice region. We will be considering rectangular dissections of Λ_n into rectangles of area n in the dyadic and general cases. Let Ω_n be the set of dyadic tilings of Λ_n and let $\widehat{\Omega}_n$ be the set of rectangular dissections of Λ_n into rectangles of area n that are not necessarily dyadic. In the weighted setting, we are given an input parameter $\lambda > 0$ and the weight of a dyadic tiling $\sigma \in \Omega_n$ is $\pi(\sigma) = \lambda^{|\sigma|}/Z$, where $|\sigma|$ is the total length of edges in σ and $Z = \sum_{\sigma \in \Omega_n} \lambda^{|\sigma|}$ is the normalizing constant known as the partition function. Likewise, in the general dissection setting, for $\alpha \in \widehat{\Omega}_n$, we define $\widehat{\pi}(\alpha) = \lambda^{|\alpha|}/\widehat{Z}$, where $|\alpha|$ is the total length of α and \widehat{Z} is again the normalizing constant.

Let \mathcal{M}_n be the edge-flip Markov chain on Ω_n that replaces an edge bordering two rectangles with the perpendicular bisector of the combined area $2n$ rectangle, provided the resulting tiling remains dyadic (details

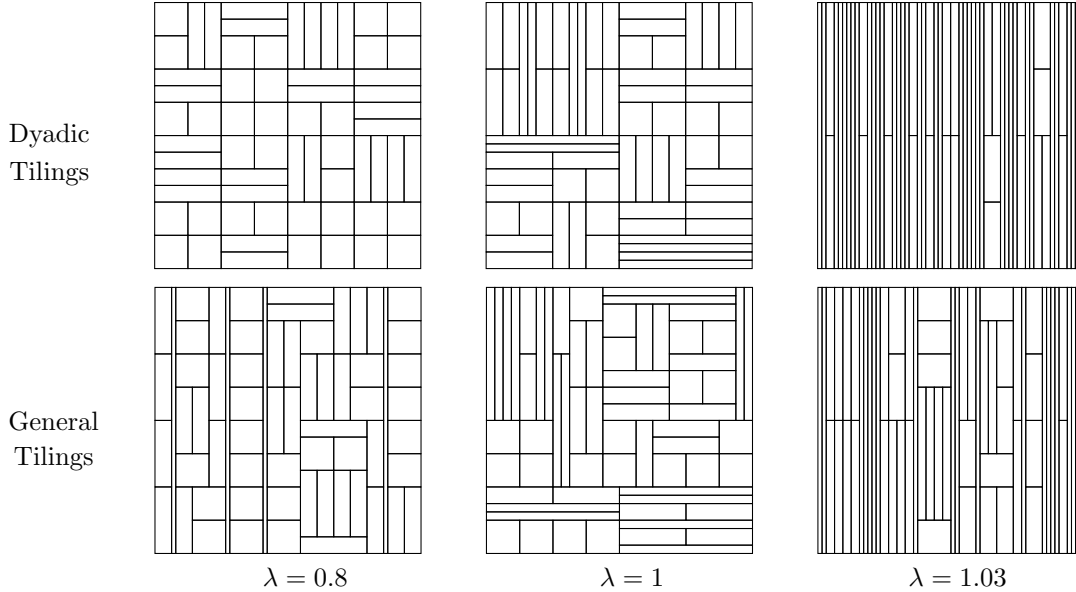


FIG. 2. \mathcal{M}_{64} and $\widehat{\mathcal{M}}_{64}$ after 1,000,000 simulated steps for various values of λ , starting with all vertical rectangles of width 1 and height 64.

are given in Section 2.). It is easy to generalize this chain to the weighted setting by modifying the transition probabilities so that the chain converges to distribution π . Likewise, we can define the natural generalization of the edge-flip chain $\widehat{\mathcal{M}}_n$ on $\widehat{\Omega}_n$ by connecting two dissections if they differ by the removal and addition of one edge. It is not obvious that this Markov chain $\widehat{\mathcal{M}}_n$ connects the state space $\widehat{\Omega}_n$, and establishing this is our first result (Theorem 7).

The remainder of the paper is concerned with the mixing times of \mathcal{M}_n and $\widehat{\mathcal{M}}_n$ as we vary the parameter λ . One might expect the same behavior for weighted rectangular dissections as in the triangulation case, namely that when λ is small we favor “balanced” rectangles (that are close to square) and we might expect the chain to be rapidly mixing, while for λ large we favor long thin rectangles, and we should expect they will mostly align vertically or horizontally. This picture is actually much more complicated in the general case, but precisely what we find in the dyadic setting. In particular, in the dyadic case we show that there is a critical point $\lambda_c = 1$, below which the chain mixes in polynomial time and above which the chain requires exponential time (Theorems 13 and 18). Such a sharp phase transition is conjectured but unproven for triangulations [5]. In subsequent work, the chain \mathcal{M}_n has been shown to converge in polynomial time also at the critical point λ_c [4]. See the top row of Figure 2 for samples generated with various values of λ for \mathcal{M}_{64} .

In the general setting the picture is more surprising. When λ is large, as in the dyadic case, the Markov chain $\widehat{\mathcal{M}}_n$ requires exponential time (Theorem 18). However, we show that the chain also requires exponential time to converge to equilibrium when λ is small (Theorem 19). Even though together these results seem to suggest that the chain will always be slow, the proofs in these two regimes (i.e., $\lambda < 1$ and $\lambda > 1$) verify that the reasons underlying the slow mixing results are quite different. When $\lambda > 1$, long thin rectangles are favored, and it will take exponential time to move from a configuration that is predominantly horizontal to one that is vertical. However, when $\lambda < 1$, balanced rectangles are favored. This is enough to dramatically speed up the mixing time in the dyadic case, but in the general setting it causes an obstacle because long thin rectangles that are well separated by many squares (or near squares) will take exponential time to disappear since their removal requires more long-thin rectangles whose creation is exponentially unlikely. Both slow mixing proofs when $\lambda > 1$ show that there is a bad cut in an equitable partition of the state space into two equal sized pieces, but the proof in the general setting when $\lambda < 1$ relies critically on a careful choice of the starting configuration. It may indeed be the case that the chain is fast if we start from the most favorable configuration consisting entirely of squares. The convergence time is unknown when $\lambda = 1$, but based on simulations we conjecture that the chain $\widehat{\mathcal{M}}_n$ converges quickly to equilibrium at this isolated point (see the bottom row of Figure 2).

We note that these results for dyadic tilings are complementary to other phase transitions discovered

in the unweighted setting. Angel et al. [1] affirmatively answered a question of Joel Spencer regarding the probability that there is a dyadic tiling if each dyadic rectangle is present with probability p , independent of the others. They show that there is a phase transition for some $p < 1$, at which point the likelihood of there not being such a tiling becomes exponentially small.

1.2. Techniques. Dyadic tilings have rich combinatorial properties that allow us to establish the presence of a phase transition in the convergence times. The proof of fast mixing of \mathcal{M}_n on dyadic tilings when $\lambda < 1$ is based on the method of exponential metrics for path coupling. Similar techniques have been used by Greenberg et al. [8] for lattice paths and Caputo et al. [5] for weighted triangulations. Both of these results rely on analysis of lattice paths while we take a more traditional approach that directly analyzes configurations of rectangles. It is worth noting that the analysis is self-contained and does not rely on computational tools to optimize the weights used in the calculations.

To show slow mixing of the Markov chains \mathcal{M}_n and $\widehat{\mathcal{M}}_n$ in the dyadic and general cases when $\lambda > 1$, we apply a standard Peierls argument. Here, a straightforward analysis suffices to show that configurations without horizontal or vertical long thin rectangles must have exponentially small weight, even after summing over all such configurations. Since we must pass through these very unlikely configurations to move from a mostly horizontal configuration to a mostly vertical one, we can conclude that the mixing time is exponential using a basic flow argument.

The proof of slow mixing for general rectangular dissections when $\lambda < 1$ is considerably more delicate. In this regime, rectangles that are close to square are preferred. We show that it will take exponential time to move from a configuration that has two well-separated long thin rectangles to one that does not have any long thin rectangles by very carefully analyzing required features of these tilings.

2. Preliminaries. We start by formalizing the problems. In the remainder of this paper, we will refer to equitable rectangular dissections instead as *tilings* in analogy to the widely used designation *dyadic tilings* to provide a uniformity of language.

Let $n = 2^k$ for some even integer k . An n -tiling is a tiling of the $[0, n] \times [0, n]$ lattice Λ_n by n axis-aligned rectangles, each of area n ; see Figure 1. We assume all rectangles are the Cartesian product of two closed intervals, $R = [x_1, x_2] \times [y_1, y_2]$, and are of dimension $2^a \times 2^b$, where $a, b \in \{0, 1, 2, \dots, k\}$ and $a + b = k$. That k is even implies n is a perfect square and there exists a “ground state” tiling consisting entirely of $\sqrt{n} \times \sqrt{n}$ squares; this is critical to the proof of Theorem 13. A tiling is *dyadic* if all rectangles are of the form $[s2^u, (s+1)2^u] \times [t2^v, (t+1)2^v]$ for some nonnegative integers s, t, u, v , that is, if all rectangles are the Cartesian product of two dyadic intervals. We will use the following lemma about dyadic intervals in some of our proofs.

LEMMA 1. *Two dyadic intervals of the same length do not overlap nontrivially.*

Proof. For any integer u , the only dyadic intervals of length 2^u are

$$[0, 2^u], [2^u, 2 \cdot 2^u], [2 \cdot 2^u, 3 \cdot 2^u], \dots [k \cdot 2^u, (k+1) \cdot 2^u], \dots$$

and none of these overlap except at their endpoints. □

2.1. The Markov Chains \mathcal{M}_n and $\widehat{\mathcal{M}}_n$. We study two related Markov chains \mathcal{M}_n and $\widehat{\mathcal{M}}_n$ whose state spaces Ω_n and $\widehat{\Omega}_n$, respectively, are all dyadic n -tilings and all n -tilings. Moves in these Markov chains consist of edge-flips, which we now define. By an *edge*, we mean a boundary between two adjacent rectangles in a tiling. Two tilings σ_1, σ_2 differ by exactly one *edge-flip* if it is possible to remove an edge in σ_1 that bisects a rectangle of area $2n$ and replace it with the bisecting edge in the perpendicular orientation to form σ_2 . For example, in Figure 3, tilings (a) and (b) differ by a single edge-flip, as do tilings (c) and (d). We say an edge e is *flippable* if it bisects a rectangle of area $2n$.

We consider biased Markov chains with a bias $\lambda \in (0, \infty)$, analogous to [5]. For a tiling σ , let $|\sigma|$ denote the sum of the lengths of all the edges in σ . First, we define the Markov chain $\widehat{\mathcal{M}}_n$ with bias λ . Note all logarithms are assumed to be base 2. Starting at any tiling σ_0 , iterate:

- Choose, uniformly at random, $(x, y, d, o, p) \in$

$$\left\{ \frac{1}{2}, \frac{3}{2}, \frac{5}{2}, \dots, \frac{2n-1}{2} \right\} \times \left\{ \frac{1}{2}, \frac{3}{2}, \frac{5}{2}, \dots, \frac{2n-1}{2} \right\} \times \{t, l, b, r\} \times \{0, 1\} \times (0, 1).$$




FIG. 3. Some tilings for $n = 16$. Non-dyadic tilings (a) and (b) differ by an edge-flip. Dyadic tilings (c) and (d) differ by an edge-flip.

Let R be the (unique) rectangle in σ_t containing (x, y) . If $d = t$, let e be the top boundary of R ; if $d = l, b$, or r , let e be the left, bottom, or right boundary of R , respectively.

- If e is a flippable edge and $\log |e| \equiv o \pmod{2}$, let σ' be the tiling obtained by flipping e to new edge e' . If $p < \lambda^{|\sigma'| - |\sigma_t|} = \lambda^{|e'| - |e|}$, then $\sigma_{t+1} = \sigma'$.
- Else, $\sigma_{t+1} = \sigma_t$.

The Markov chain \mathcal{M}_n for dyadic tilings is defined in the same way, interpreting “flippable” to mean flippable into another dyadic tiling; Figure 3 (c) and (d) shows an edge-flip between two dyadic tilings that are adjacent in Ω_n .

We note that each rectangle R of any tiling σ is of area n and so contains exactly n points in $\{\frac{1}{2}, \frac{3}{2}, \frac{5}{2}, \dots, \frac{2n-1}{2}\} \times \{\frac{1}{2}, \frac{3}{2}, \frac{5}{2}, \dots, \frac{2n-1}{2}\}$. A given flippable edge e in σ is thus selected by $2n$ different values of (x, y, d, o) , specifically, the $2n$ points (x, y) in the two rectangles e separates, each with the appropriate value of direction d and parity o . Consequently, a given flippable edge e is selected by (x, y, d, o) with probability $2n \cdot \frac{1}{n^2} \cdot \frac{1}{4} \cdot \frac{1}{2} = \frac{1}{4n} =: q$. This flip then occurs with probability $\min\{1, \lambda^{|\sigma'| - |\sigma|} = \lambda^{|e'| - |e|}\}$, according to the random value of p . These transition probabilities favor long, thin rectangles when $\lambda > 1$ and favor squares or rectangles close to square when $\lambda < 1$.

At most one of $(x, y, d, 0, p)$ and $(x, y, d, 1, p)$ results in an edge-flip; (x, y, d) selects a potentially flippable edge e in σ_t , and then an edge-flip can only occur if the length of e satisfies $\log |e| = o \pmod{2}$. This implies both \mathcal{M}_n and $\widehat{\mathcal{M}}_n$ are lazy and thus aperiodic.

In Section 2.2, below, we demonstrate that $\widehat{\mathcal{M}}_n$ and \mathcal{M}_n are irreducible and thus ergodic, so they converge to unique stationary distributions $\widehat{\pi}$ and π , respectively. By detailed balance, the distribution $\widehat{\pi}$ for general tilings can be given by $\widehat{\pi}(\sigma) = \lambda^{|\sigma|} / \widehat{Z}$, where \widehat{Z} is the normalizing constant. Similarly, for dyadic tilings, $\pi(\sigma) = \lambda^{|\sigma|} / Z$, where Z is the normalizing constant.

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

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

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, for our theorems in Section 1.1 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 .

2.2. Ergodicity of \mathcal{M}_n and $\widehat{\mathcal{M}}_n$. It remains to be shown that the moves described above connect state spaces Ω_n and $\widehat{\Omega}_n$. Connectivity for Ω_n follows from work on dyadic tilings in [10], specifically from their tree representation of a dyadic tiling. Dyadic constraints ensure rectangles exist in pairs, meaning an edge-flip is always possible for every rectangle. In particular, all $1 \times n$ and $n \times 1$ rectangles are adjacent to at least one other rectangle of the same dimensions, so can be eliminated with a single edge-flip.

However, connectivity of $\widehat{\Omega}_n$ is much less straightforward, and an interesting result in its own right. Intuitively, issues arise because rectangles in a general n -tiling do not exist in pairs and there may be many rectangles for which no edge-flip is possible; it is not even immediately evident that there is always a single valid edge-flip for any general tiling. Rectangles of height n , or alternately, rectangles of height h where there are no rectangles of larger height, may be well separated by complicated arrangements of tiles. It is not clear how to introduce another rectangle of height h next to an existing rectangle of height h so that




FIG. 4. For $n = 16$, the bold lines in (a) define an 8-region that is not linked, while in (b) they give a linked 4-region. Slicing lines intersecting the interior of each region are dashed.




FIG. 5. A linked h -region P with four slices; leftmost slice Q is white, slices at distance 1 from Q are light gray, and slices at distance 2 from Q are dark gray.

both may be eliminated, a necessary step for obtaining a tiling with no rectangles of height h or larger, for instance.

To prove ergodicity of $\widehat{\mathcal{M}}_n$, we use a double induction on h -regions, which are certain subsets of rectangles from an n -tiling in which all rectangles have height at most h . We prove there exists a sequence of edge-flips leading to a tiling of h -region P by rectangles of height h . One can repeatedly find within P an $h/2$ -tiling or an h -region of strictly smaller area than P , inductively apply a sequence of edge-flips to obtain tilings with all height $h/2$ or h rectangles, respectively, and apply a final sequence of edge-flips yielding a tiling of P by rectangles of height h . As an n -tiling is an h -region for $h = n$, there is a sequence of edge-flips connecting any two n -tilings, going through the tiling consisting entirely of $1 \times n$ rectangles.

Formally, an h -region is a simply-connected subset of rectangles from an n -tiling in which (A) all rectangles have height at most h , and (B) for all vertical segments on the boundary of the region there is a $c \in \mathbb{Z}^+$ such that the segment has length ch . Note (B) is equivalent to all horizontal segments on the boundary of the region being separated by some multiple of h . For $n = 16$, Figure 4(a) depicts an 8-region while (b) depicts a 4-region. By the interior $int(P)$ of an h -region P we mean the area occupied by the rectangles of P minus its boundary. The collection of all vertical edges on the boundary of an h -region P for which the interior of P is to the right is the *left boundary* of P , and the *right boundary* of P is defined similarly.

For any h -region P , let h_0 denote the vertical coordinate of the bottommost boundary edge of P . Call the horizontal lines at heights $h_0, h_0 + h, h_0 + 2h, \dots, h_0 + kh, \dots$ the *slicing lines* of P . By (B) in the definition of an h -region, all horizontal segments on the boundary of P are contained in some slicing line. An h -region is *linked* if every connected component of the intersection of a slicing line with $int(P)$ also intersects the interior of some rectangle in P . That is, P is linked if there are no segments of slicing lines that separate $int(P)$ into two disjoint h -regions. Figure 4 (a) is not linked as the 8-region is separated by the slicing line at height 8, while (b) is linked.

Call the connected regions of $int(P)$ separated by the slicing lines of P *slices* of P ; note that one rectangle might span two slices. Slices are shaded different colors in Figure 5. Call two slices *adjacent* if they are both incident on a common segment of a slicing line. If P is linked, then for any two adjacent slices there exists a rectangle spanning both.

For any linked h -region P , let w_0 denote the horizontal coordinate of the leftmost boundary edge of P ,

and let $w := n/h$ denote the minimum width of a rectangle in P . Note every rectangle in P has width $2^i w$, for some integer $i \geq 0$, and has width exactly w if and only if its height is h .

LEMMA 2. *Let P be a linked h -region. For any rectangle $[x_1, x_2] \times [y_1, y_2]$ in P , both x_1 and x_2 can be written in the form $w_0 + dw$, $d \in \mathbb{Z}^+$.*

Proof. Let Q be a slice of P adjacent to a leftmost boundary edge of P . All rectangles whose interior intersects Q satisfy the necessary property because all rectangles in P have width a multiple of w . We proceed by induction on the distance between some slice Q_i and Q , where two adjacent slices are at distance one; see Figure 5.

Suppose that Q_i is adjacent to some Q_{i-1} at distance $i-1$ from Q , and that the statement holds for all rectangles whose interior has a nontrivial intersection with Q_{i-1} . At least one of those rectangles R in Q_{i-1} must also have a nontrivial intersection with Q_i because P is linked. Traveling leftwards in Q_i from R , all rectangles crossed must also satisfy the desired property, as they are separated horizontally from R by some collection of rectangles, each of width a multiple of w . It follows that the left boundary edge of Q_i satisfies the desired property, and thus all rectangles in slice Q_i do. \square

For any (linked) h -region P , let S be the set of *aligned* rectangles in P , that is, all rectangles $[x_1, x_2] \times [y_1, y_2]$ in P of height h for which $y_1 - h_0$ is an integer multiple of h . Aligned rectangles are precisely those whose top and bottom boundaries are both contained in some slicing line and whose interior doesn't intersect a slicing line. All other rectangles of height h are *unaligned*.

LEMMA 3. *Let P be an h -region. Any connected component of the intersection of a slicing line with the interior of P intersects the interior of an even number (possibly 0) of rectangles of height h .*

Proof. We will prove the stronger statement that for any connected component of the intersection of $\text{int}(P)$ with the slicing line l at vertical coordinate y_l , for every $y \in \{y_l - h + 1, y_l - h + 2, \dots, y_l - 1\}$, there are an even number (possibly 0) of rectangles $R = [x_1, x_2] \times [y_1, y_2]$ of height h satisfying $y_1 = y$ that cross l .

Suppose for the sake of contradiction that the stronger statement above does not hold. We define a dual graph F on slices of P , where each vertex represents a slice of P . Two slices (vertices) are connected by an edge in F if there is a common segment l of a slicing line between them and this segment doesn't satisfy the above statement; that is, if there is some $y \in \{y_l - h + 1, \dots, y_l - 1\}$ such that there are an odd number of rectangles of height h with $y_1 = y$ that cross l . As P is simply connected F is a forest, with at least one edge by assumption. Pick some slice Q of P that corresponds to a degree one vertex of F . Let Q' be the unique adjacent slice such that segment l separating Q from Q' crosses an odd number of rectangles of height h with some common value for y_1 . Let y_l denote the vertical coordinate of this slicing line l , and suppose without loss of generality Q lies below l and Q' lies above it.

Consider the collection of rectangles $R = [x_1, x_2] \times [y_1, y_2]$ of height h crossing l , which each have $y_1 \in \{y_l - h + 1, \dots, y_l - 1\}$. Let y^* be the smallest among $y_l - h + 1, \dots, y_l - 1$ such that an odd number of rectangles of height h crossing l have $y_1 = y^*$. Because slice Q was a vertex of degree 1 in F , for each coordinate $y \in \{y_l - h + 1, \dots, y_l - 1\}$ (particularly, for $y = y^*$) there are an even number of unaligned rectangles of height h non-trivially intersecting Q and extending down into another slice satisfying $y_2 = y$. For the same reason, there are also an even number of rectangles of height h non-trivially intersecting Q satisfying $y_1 = y$ and extending upward into some slice that is not Q' .

Examine the horizontal lines l_1 and l_2 at heights $y^* + \varepsilon$ and $y^* - \varepsilon$, respectively, for some $\varepsilon < 1$; for each consider the connected component of its intersection with the interior of Q . Both must be of the same length dw , $d \in \mathbb{Z}^+$; cross the same number r of aligned rectangles of height h in Q ; and cross the same number u of unaligned rectangles of height h that don't have y_1 or y_2 equal to y^* . Note l_1 crosses an odd number of rectangles of height h with $y_1 = y^*$, an odd number extending into Q' and an even number extending into all other slices. At the same time, l_2 crosses an even number of rectangles of height h with $y_2 = y^*$ extending into other slices below Q . By looking at l_1 we conclude d is odd and by looking at l_2 we conclude d is even, a contradiction. \square

Recall S denotes the set of aligned rectangles in a linked h -region P . Define a binary coloring of all points in $P \setminus S$. For $p = (x, y) \in P \setminus S$, let $\bar{p} = (\bar{x}, y)$ be the rightmost point on the left boundary of P that is left of p . By Lemma 2, write $\bar{x} = w_0 + dw$, $d \geq 0$ an integer. Define $o(p) = 0$ if d is even and $o(p) = 1$ if d is odd. Let $r(p)$ be the number of (aligned, height h) rectangles in S that cross the segment between \bar{p} and p . Let E be the set of all points in $P \setminus S$ with $r(p) + o(p)$ even, and let O be the set of all points in $P \setminus S$




FIG. 6. An h -region P and a rectangle R crossing a slicing line illustrating the definitions $r(p)$, $o(p)$, $u(p)$, and $u^*(p)$.

with $r(p) + o(p)$ odd. We now show $r(p) + o(p)$ has constant parity on rectangles of a linked h -region.

LEMMA 4. For linked h -region P containing rectangle $R = [x_1, x_2] \times [y_1, y_2] \subseteq P \setminus S$, for all $p \in R$, $r(p) + o(p)$ is the same modulo 2.

Proof. If R does not cross a slicing line, this is trivially true as $o(p)$ and $r(p)$ are constant on the intersection of any rectangle with any slice. If R crosses slicing line l , let $p = (x, y)$ be on the left boundary of R just above l and $p' = (x', y')$ be on the left boundary of R just below l ; see Figure 6. By Lemma 2, $x = w_0 + dw$ for some $d \in \mathbb{Z}$; we now analyze the parity of d . The value of $o(p)$ as well as all rectangles of height h between p and \bar{p} affect this parity, whether aligned or not, while shorter rectangles have width at least $2w$ and do not. Let $u(p)$ be the number of unaligned rectangles of height h that cross the segment between p and \bar{p} , and let $u^*(p)$ be the number of these rectangles that also cross the component of $l \cap \text{int}(P)$ that intersects R . We note that $u(p) \equiv u^*(p) \pmod{2}$ by Lemma 3, and that $u^*(p) = u^*(p')$, provided p and p' were placed sufficiently close together. Thus $u(p) \equiv u(p') \pmod{2}$.

One can see that $x = w_0 + dw$, where $d \equiv o(p) + r(p) + u(p) \pmod{2}$. Similarly, $x' = w_0 + d'w$, where $d' \equiv o(p') + r(p') + u(p') \pmod{2}$. As $x = x'$ and $u(p) \equiv u(p') \pmod{2}$, it follows that $r(p) + o(p) \equiv r(p') + o(p') \pmod{2}$. \square

Thus E and O form a well-defined partition of the rectangles in $P \setminus S$. We now consider $h/2$ -regions within P as well as h -regions within P with strictly smaller area of two types: - ‘interior’ and ‘boundary’. Interior $h/2$ -regions or h -regions have both their left and right boundaries adjacent to rectangles of height h in S . Boundary $h/2$ -regions contain part of the boundary of P for a portion of their left or right boundary.

A h -region P has *even width* if every connected component of the intersection of any horizontal line with $\text{int}(P)$ is of length that is an even multiple of w .

LEMMA 5. Let P be a linked h -region, where no two rectangles of height h are horizontally adjacent such that there exists a valid edge-flip between them. Then at least one of the following holds:

- all connected components of E and O are boundary $h/2$ -regions
- there exists an interior $h/2$ -region
- there exists an interior h -region of strictly smaller area than P and even width

Proof. Suppose there is at least one connected component of E or of O that is not a boundary $h/2$ -region; we now proceed to find an interior $h/2$ -region or an interior h -region with strictly smaller area.

Consider all connected components of E and O that are not boundary $h/2$ -regions; that is, all components that are not simply connected, contain rectangles of height h , or are not adjacent to the boundary of P . Define a partial order on these components where $G \leq H$ if and only if G is contained within a hole in component H . Consider any minimal element G in this partial order, and suppose without loss of generality that it is a connected component of E . Any holes in G consist entirely of aligned rectangles of height h that are in S , and by hypothesis no two of these aligned rectangles are horizontally adjacent. Consider rectangle $R = [x_1, x_2] \times [y_1, y_2]$ that is part of such a hole. Choose any vertical height y^* such that the horizontal line at height y^* intersects the interior of R , as well as the interior of some rectangle R^- immediately left of R and some rectangle R^+ immediately right of R . Let p^- be on the horizontal line at height y^* inside




FIG. 7. A simply connected h -region G that is not an $h/2$ -region. Rectangles in S' are white, while region E' is shaded dark gray and region O' is light gray. The dashed line encloses h -region G' with strictly smaller area.

R^- , and let p^+ be on the same horizontal line inside R^+ . Both R^- and R^+ must be in E , implying that $r(p^-) + o(p^-)$ has the same parity as $r(p^+) + o(p^+)$; but this is a contradiction as $r(p^+) = r(p^-) + 1$ while $o(p^+) = o(p^-)$. Thus G is in fact simply connected.

Suppose G is an $h/2$ -region. Because G is a maximal connected component of E , any rectangles in $P \setminus S$ horizontally adjacent to G would also by definition be in E . Thus the left and right boundaries of G can only be adjacent to rectangles in S because, by assumption, G is not a boundary $h/2$ -region. G is an interior $h/2$ -region, and we are done.

If G is not an $h/2$ -region, then it contains rectangles of height h ; recall that as $G \subseteq E \subseteq P \setminus S$, it contains no aligned rectangles of height h . Pick any leftmost rectangle R' of height h in G , and let $S' \subseteq G$ be the set of all rectangles of height h in G for which bottom coordinate y_1 is an integer multiple of h different from the bottom of R' ; call all rectangles in S' *realigned*. For all points in $G \setminus S'$, let $r'(p)$ be the number of realigned rectangles left of p .

In analogy to Lemma 4, we show $r'(p) \pmod{2}$ is constant on any given rectangle $R \in G \setminus S'$. First, note $r'(p)$ is constant on any rectangle entirely contained in a realigned slice of G (a slice determined by horizontal lines at vertical interval ch , $c \in \mathbb{Z}$, from the bottom of R'). Suppose rectangle R crosses some line separating two realigned slices. Examine $p = (x, y)$ and $\bar{p} = (\bar{x}, \bar{y})$ on the left boundary of R just above and just below the realigned slicing line. By Lemma 2, $x = \bar{x} = x_0 + dw$, $d \in \mathbb{Z}^+$; consider the parity of d . Using the same notation as above, $o(p) = o(\bar{p})$ because realigned slicing lines are offset from the alignment of vertical boundary edges of G . If $u(p)$ denotes the number of rectangles of height h not in S' between p and the first boundary edge left of p , $u(p) = u(\bar{p})$ for the same reason. As $d = o(p) + u(p) + r'(p) \pmod{2}$ and also $d = o(\bar{p}) + u(\bar{p}) + r'(\bar{p}) \pmod{2}$, then $r(p) = r(\bar{p}) \pmod{2}$ and $r'(p) \pmod{2}$ is constant across all rectangles in $G \setminus S'$.

Partition $G \setminus S'$ into E' , rectangles for which $r'(p)$ is even, and O' , the rectangles for which $r'(p)$ is odd. Let G' be the largest connected component of O' , joined with all of its holes; see Figure 7. Note region O' always has at least one nontrivial component as any rectangles immediately to the right of R' are not in S' by hypothesis, so satisfy $r' = 1$.

The left and right boundaries of G' must be adjacent to rectangles in $S' \subseteq G$, so $G \setminus G'$ is nonempty and G' has strictly less area than $G \subseteq P$. The lengths of the vertical boundary segments of G' are integer multiples of h , and G' is simply connected as it is the union of a connected region with all of its holes. As $G' \subseteq P$, it contains no rectangles of height more than h . Thus G' is a h -region of strictly smaller area than P . We can also observe that G' has even width, as otherwise we quickly find a contradiction within E' and O' in a manner analogous to Lemma 3. \square

LEMMA 6. For any h -region P , there exists a sequence of edge-flips within P that yields a tiling of P entirely with rectangles of height h .

Proof. We proceed by a double induction, on h and on the area of P . We first note any 1-region is simply an $n \times a$ box consisting of a horizontal $n \times 1$ rectangles, that is, it is tiled with rectangles of height 1. Additionally, the smallest area h -regions consist of a single rectangle of height h , and are clearly tiled with rectangles of height h . These two examples serve as the base cases for our double induction.




FIG. 8. The edge-flip sequence used to reduce the number of rectangles of height h given an interior $h/2$ -region tiled entirely by height $h/2$ rectangles.




FIG. 9. An h -region P in which all connected components of E (dark gray) and O (light gray) are boundary $h/2$ -regions; the tiling of P after applying (1) to each; and the tiling of P by rectangles of height h resulting from edge-flips between vertically-adjacent rectangles of height $h/2$.

Let $h = 2^a$, $a \leq k$, be some height larger than one, and let P be any h -region. Suppose by induction that (1) for any $h/2$ -region, there exists a sequence of edge-flips leading to a tiling of the $h/2$ -region entirely by rectangles of height $h/2$, and (2) for all h -regions with strictly smaller area than P , there exists a sequence of edge-flips yielding a tiling consisting entirely of rectangles of height h .

If P is not linked, then we can separate P into at least two h -regions of smaller area and apply (2), so assume P is linked. We can also assume that P never contains two horizontally-adjacent rectangles of height h with a valid edge-flip between them; any such pairs can easily be eliminated with single edge-flip, creating instead two rectangles of height $h/2$, in a preprocessing step.

We now apply Lemma 5, to show that unless all connected components of E and O are boundary $h/2$ -regions, then we can always reduce the number of rectangles of height h in P .

If there exists an interior $h/2$ region, by (1), flip so that all rectangles are of height $h/2$. Recall that its boundary must consist of rectangles of height h by the definition of an interior region. For each left boundary height h rectangle R , ‘move’ it to the right via a sequence of edge-flips. Specifically, flip the horizontal edge separating the two height $h/2$ rectangles immediately to R ’s right to create two more rectangles of height h at the same alignment as R ; then, flip R ’s right edge. Repeat until there is a rectangle of height h adjacent to a right boundary rectangle of height h , at which point one final flip eliminates both; see Figure 8. After each such sequence of flips, there are two fewer rectangles of height h in P .

If there exists an interior h -region G of strictly smaller area and even width, by (2), flip edges such that the region is tiled exclusively by rectangles of height h . Because of the even width condition, at each y -coordinate, there are an even number of rectangles of height h in G ; including the rectangles of height h necessarily adjacent to the left and right boundary of G , there are still an even number of rectangles. These can be paired horizontally and edges can be flipped such that the region occupied by G and the height h rectangles adjacent to its left and right boundary is tiled with rectangles of height $h/2$. There are now fewer rectangles of height h because, at the least, the rectangles adjacent to the boundary of G have been eliminated.

After a finite number of steps, all connected components of E and O are boundary $h/2$ -regions. By (1), flip edges such that each $h/2$ -region is tiled by rectangles of height $h/2$. As in fact all vertical boundary edges of these regions are of height h , it is possible to pair all of these height $h/2$ rectangles vertically, and flip edges such that each connected component of E and O is tiled by rectangles of height h , as demonstrated in Figure 9. This yields a tiling of P exclusively by rectangles of height h . \square

THEOREM 7. The Markov chain $\widehat{\mathcal{M}}_n$ connects the state space $\widehat{\Omega}_n$ consisting of all rectangular dissections of Λ_n into n rectangles with area n .

Proof. Note that any tilings σ_1 and σ_2 of the $n \times n$ square are n -regions. By Lemma 6, there exists a sequence s_1 of edge-flips for σ_1 leading to the tiling of the $n \times n$ square by rectangles of height n and width 1,

and similarly there exists s_2 for σ_2 . Applying the edge-flips of s_1 and then the reverse of s_2 yields a sequence of edge-flips connecting σ_1 and σ_2 . \square

3. Polynomial Time Mixing for Dyadic Tilings when $\lambda < 1$. We begin by proving that \mathcal{M}_n is rapidly mixing for all $\lambda < 3^{-1/\sqrt{n}}$ (Theorem 12). This bound approaches 1 as n grows, so for any $\lambda < 1$ there is sufficiently large n for which the Markov chain \mathcal{M}_n is rapidly mixing (Theorem 13). To give some perspective, we note that for all $n \geq 4$, we have fast mixing for all $\lambda < 0.577$, a much better constant than obtained in [5]. Already for $n \geq 1024$ we have fast mixing for all $\lambda < 0.966$.

We use a path coupling argument and an exponential metric, as in [8]. A *coupling* of a Markov chain \mathcal{M} is a joint Markov process on $\Omega \times \Omega$ such that the marginals each agree with \mathcal{M} and, once the two coordinates coalesce, they move in unison. *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. By considering an appropriate metric ϕ on Ω , proving that the two marginal chains, if in a joint configuration in subset U , get no farther away in expectation after one iteration is sufficient to show that \mathcal{M} is rapidly mixing. The following theorem bounds the mixing time of a Markov chain by considering a path coupling; a version of this theorem was first proved in [3], but the statement presented in [8] will be easier for us to work with.

THEOREM 8 ([8]). *Let $\phi : \Omega \times \Omega \rightarrow \mathbb{R}^+ \cup \{0\}$ be a metric that takes on finitely many values in $\{0\} \cup [1, B]$. Let $U \subseteq \Omega \times \Omega$ be such that for all $(X_t, Y_t) \in \Omega \times \Omega$, there exists a path $X_t = Z_0, Z_1, \dots, Z_r = Y_t$ such that $(Z_i, Z_{i+1}) \in U$ for $0 \leq i < r$ and $\sum_{i=0}^{r-1} \phi(Z_i, Z_{i+1}) = \phi(X_t, Y_t)$.*

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

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

Then the mixing time satisfies

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

This theorem is particularly useful because the values taken by ϕ can be exponential in n . As long as the distance between two chains in a coupling decreases by some constant multiplicative factor with each move of the joint Markov process, the Markov chain is provably rapidly mixing.

We now apply this exponential metric theorem. Intuitively, we consider the subset U of the joint state space $\Omega_n \times \Omega_n$ of tilings that differ by one edge-flip. The main result we need to show is that for any coupling whose joint state is two configurations in U , after one iteration of the Markov chain, the expected distance between the two coupled chains decreases by a constant factor of their original distance. It is crucial to define the appropriate notion of “distance” between two tilings; we do so by carefully defining the distance between tilings that are adjacent in Ω_n , and then extending this definition to non-adjacent tilings.

DEFINITION 9. *Consider any dyadic tilings σ_1 and σ_2 that differ by one flip between edge e and edge f , both bisecting a common area $2n$ rectangle S . Without loss of generality, suppose that $|e| \geq |f|$. We define the distance between σ_1 and σ_2 to be*

$$\phi(\sigma_1, \sigma_2) = \phi(\sigma_2, \sigma_1) := \lambda^{|f| - |e|},$$

For any dyadic tilings σ_1 and σ_2 that are not adjacent in Ω_n , the distance between σ_1 and σ_2 is the minimum over all paths in Ω_n from σ_1 to σ_2 of the sum of the distances between adjacent tilings along the path.

If $\sigma_1 = \sigma_2$, then $\phi(\sigma_1, \sigma_2) = 0$.

We now examine the range of this distance metric.

LEMMA 10. *For $\lambda < 1$, the distance metric ϕ of Definition 9 takes on values in the range*

$$0 \cup [1, n \log(n) \lambda^{-n}].$$

Proof. If σ_1 and σ_2 differ by an edge-flip, then $\phi(\sigma_1, \sigma_2)$ is at least 1 as $\lambda < 1$ is being raised to a nonpositive power. If σ_1 and σ_2 differ by more than an edge-flip, then $\phi(\sigma_1, \sigma_2)$ is a sum of distances that are at least 1 so is at least 1. It only remains to show that $\phi(\sigma_1, \sigma_2)$ has the stated upper bound.




FIG. 10. A sequence of edge-flips from σ_h to σ^* .



FIG. 11. Rectangle S of area $2n$ in marginal tilings A_t and B_t .

Let σ^* denote the ground state tiling, tiling the $n \times n$ square with n smaller squares of size $\sqrt{n} \times \sqrt{n}$, careful consideration shows that the two dyadic tilings at farthest distance ϕ from σ^* are the tiling consisting of all $n \times 1$ horizontal rectangles σ_h and the tiling consisting of all $1 \times n$ vertical rectangles σ_v . We note that one path in Ω_n from σ_h to σ^* consists of $(\log n)/2 = k/2$ stages, where in each stage $n/2$ edge-flips are performed, reducing the length of each of the n rectangles by half; see Figure 10.

The contribution to $\phi(\sigma_h, \sigma^*)$ from each of these edge-flips is at most λ^{-n} , and there are $nk/4$ such moves in this particular path in Ω_n from σ_h to σ^* , giving $\phi(\sigma_h, \sigma^*) \leq (nk/4)\lambda^{-n}$. The same holds for σ_v . There is thus a path between any two tilings, through the ground state σ^* , yielding the bound

$$\phi(\sigma_1, \sigma_2) \leq (nk/2)\lambda^{-n} \leq n \log(n)\lambda^{-n}.$$

Thus ϕ takes on values in the range $\{0\} \cup [1, n \log(n)\lambda^{-n}]$, as claimed. \square

We now define a coupling and will show that at each iteration, the two coupled process are getting closer together in expectation according to distance metric ϕ . Fix $\lambda < 3^{-1/\sqrt{n}}$. Let $(\mathcal{A}, \mathcal{B})$ denote a coupling of Markov chain \mathcal{M}_n on dyadic tilings, where A_t and B_t are the states of the two chains, respectively, after t iterations. At each iteration, A_t and B_t are simultaneously updated by choosing the same random values (x, y, d, o, p) for each. Let $\phi_t = \phi_t(A_t, B_t)$ denote the distance between the two chains in the coupling $(\mathcal{A}, \mathcal{B})$ after t iterations.

LEMMA 11. *Suppose A_t and B_t differ by a single edge-flip. Then for all $\lambda < 3^{-1/\sqrt{n}}$, there is a constant c such that for $q = 1/(4n)$,*

$$\mathbb{E}[\phi_{t+1}] \leq (1 - qc)\phi_t$$

Proof. Let A_t and B_t differ by a single between edge e and edge f , where without loss of generality $|e| \geq |f|$, e is horizontal in A_t of length $2a$, f is vertical in B_t of length $2b$, and both bisect a rectangle S of area $2n$; see Figure 11.

We wish to bound $\mathbb{E}[\phi_{t+1} - \phi_t]$ in terms of ϕ_t . Any potential moves (x, y, d, o, p) that select an edge not in S and not on the boundary of S have the same effect on both A_t and B_t and thus, in these cases, $\phi_{t+1} = \phi_t$, as A_{t+1} and B_{t+1} still differ by the same single edge-flip. We next note there is a rectangle in valid dyadic tiling A_t of dimension $2a \times b$, implying that $2ab = n = 2^k$. As a and b are powers of 2, $a \geq b$ by assumption, and k is even, then $a = 2^i b$ where i is odd. We now consider two cases, $a \geq 8b$ and $a = 2b$.

Case $a \geq 8b$. We first examine the moves that decrease the distance between the two coupled chains. There are exactly two edge-flips that do this, namely flipping e to f in A_t or flipping f to e in B_t . There are $2n$ values of (x, y, d, o) that select edge e in A_t . Precisely, these are each of the $2n$ points (x, y) in S together with the appropriate direction from among t, b that selects e and the appropriate parity o such that $\log |e| = o(\text{mod } 2)$. Examining the parity o shows these same choices do not yield a flippable edge in B_t ; this is where the value of o plays a critical role, as no edges within or on the boundary of S in $B_t = \sigma_2$ are of the same length as e . As each such selection of (x, y, d, o) occurs with probability $1/(8n^2)$, potential edge-flip e




FIG. 12. An area $2n$ rectangle S bisected by horizontal edge e in A_t and vertical edge f in B_t . Four “bad” edge-flips g, h, i, j exist only if A_t and B_t are tiled in the neighborhood of S as shown.

is selected with probability $q = 1/(4n)$. In this case the condition for flipping edge e is $p < \lambda^{2b-2a}$, which always occurs as $2b - 2a \leq 0$. After such a flip, $A_{t+1} = B_t$ while $B_{t+1} = B_t$. Thus $\phi_{t+1} = 0$ and the change in distance between the two chains is $-\phi_t = -\lambda^{2b-2a}$. The total contribution to the expected change in $\phi(\mathcal{A}, \mathcal{B})$ from this move is $-q \cdot \lambda^{2b-2a}$.

Similarly, the probability (x, y, d, o) selects edge f in B_t is also $q = 1/(4n)$, and these values do not yield a flippable edge in A_t . Edge f flips only if $p < \lambda^{2a-2b}$, which occurs with probability $\lambda^{2a-2b} < 1$. If this move occurs, then $B_{t+1} = A_t = A_{t+1}$, and the change in distance between A and B is again $-\lambda^{2b-2a}$. The total contribution to the expected change in $\phi(\mathcal{A}, \mathcal{B})$ from this move is

$$-q \cdot \lambda^{2a-2b} \cdot \lambda^{2b-2a} = -q.$$

While the two potential moves above decrease the distance between the chains according to metric ϕ , there are also moves that increase it. However, dyadic constraints limit the number of such moves, because two dyadic intervals of the same length cannot overlap nontrivially (Lemma 1). For A_t , this means the top and bottom edges of S are not flippable, because the vertical dimension of any rectangle resulting from such a flip could not be dyadic: it would be an interval of length $2b$ overlapping S 's projection onto the y -axis, which is a dyadic interval of length $2b$. At first glance there are four other potential edge-flips for A_t involving S , specifically flips of the top and bottom halves of S 's left and right boundaries. However, again by Lemma 1, at most one of the left boundary and the right boundary of S contains flippable edges. Without loss of generality, assume it is the right boundary of S , and label the two potentially flippable edges as g and h . Similarly, for B_t , at first glance there exist four other potential edge-flips involving S , specifically the left and right halves of S 's top and bottom boundaries. By Lemma 1, we assume without loss of generality that only portions of S 's bottom boundary are potentially flippable, and label the two potentially flippable edges as i and j .

Such edge-flips can only occur if A_t and B_t are tiled in the neighborhood of S as in Figure 12. We suppose this worse case neighborhood tiling exists. Edges g and h are each selected by values (x, y, d, o) in A_t with probability q ; both are then flipped with probability λ^{4a-b} . The tiling A_{t+1} resulting from this flip is at distance λ^{b-4a} from configuration A_t . The same selection (x, y, d, o) does not result in any flip in B_t , so $B_{t+1} = B_t$. The change in distance between \mathcal{A} and \mathcal{B} for each of these two moves is at most λ^{b-4a} . In all, the contribution by these moves to the expected change in distance between the coupled chains is at most

$$2 \cdot q \lambda^{4a-b} \cdot \lambda^{b-4a} = 2q.$$

Similarly, edges i and j are selected to be flipped in B_t by values (x, y, d, o) with probability q , and once selected, these edge-flips occur if $p < \lambda^{4b-a}$, a bound which is at least 1 for $a \geq 8b$. The tiling B_{t+1} resulting from either flip is at distance at most λ^{4b-a} from configuration B_t . The same values of (x, y, d, o) producing these moves also yield $A_{t+1} = A_t$. Thus the change in distance between \mathcal{A} and \mathcal{B} for these two moves is at most λ^{4b-a} . In all, the contribution by these moves to the expected change in distance between the two chains in the coupling is at most $2 \cdot q \cdot \lambda^{4b-a}$.

In total, we have shown

$$\begin{aligned} \mathbb{E}[\phi_{t+1} - \phi_t] &\leq -q - q\lambda^{2b-2a} + 2q + 2q\lambda^{4b-a} \\ &= -q\lambda^{2b-2a}(\lambda^{2a-2b} + 1 - 2\lambda^{2a-2b} - 2\lambda^{2b+a}) \\ &= -q\phi_t(1 - \lambda^{2a-2b} - 2\lambda^{2b+a}). \end{aligned}$$

We first note that as $a \geq 8b$, and in particular, as $a \geq \sqrt{n}$,

$$2a - 2b \geq 2\left(a - \frac{1}{8}a\right) \geq a \geq \sqrt{n}.$$

Additionally, $2b + a \geq a \geq \sqrt{n}$. Thus,

$$\lambda^{2a-2b} + 2\lambda^{2b+a} \leq \lambda^{\sqrt{n}} + 2\lambda^{\sqrt{n}} = 3\lambda^{\sqrt{n}}.$$

Provided $\lambda < 3^{-1/\sqrt{n}}$, as we assumed at the start of this section, we have that

$$\lambda^{2a-2b} + 2\lambda^{2b+a} \leq 3\lambda^{\sqrt{n}} < 1.$$

Then,

$$\mathbb{E}[\phi_{t+1}] \leq (1 - qc)\phi_t,$$

where c is some positive constant, depending on how close λ is to the bound $3^{-1/\sqrt{n}}$ given above. This concludes the first case.

Case $a = 2b$. The analysis of potential good moves and bad moves remains the same as the first case above, though certain probabilities and distances change. Initially, $\phi(A_t, B_t) = \lambda^{2b-2a} = \lambda^{-2b}$, as in the previous case. We note that the contribution to the expected change in distance from good moves flipping edges e and f is still $-q(1 + \lambda^{2b-2a}) = -q(1 + \lambda^{-2b})$. The contributions to the expected change in distance from flipping edges g and h is still $2q$. We note now, however, that for the edges i and j , once selected by (x, y, d, o) , flips now occur with probability $q\lambda^{4b-a} = q\lambda^{2b}$ rather than probability q . Such a move results in a change in distance between the chains in the coupling of $\lambda^{a-4b} = \lambda^{-2b}$. The expected contribution to the change in distance from these moves is now $2q\lambda^{2b}\lambda^{-2b} = 2q$.

In total, we see that in this case,

$$\begin{aligned} \mathbb{E}[\phi_{t+1} - \phi_t] &\leq -q(1 + \lambda^{-2b}) + 4q \\ &= -q\lambda^{-2b}(\lambda^{2b} + 1 - 4\lambda^{2b}) \\ &= -q\phi_t(1 - 3\lambda^{-2b}). \end{aligned}$$

We note that in this case, $2ab = n$ so $a = 2b = \sqrt{n}$. Provided $\lambda < 3^{-1/\sqrt{n}}$, it follows that $3\lambda^{\sqrt{n}} < 1$, the required condition holds and we get the same bound on $\mathbb{E}[\phi_{t+1}]$ as in the previous case, though with a different constant c , also depending on λ . \square

THEOREM 12. *The mixing time of Markov chain \mathcal{M}_n is $O(n^2)$ for all $\lambda < 3^{-1/\sqrt{n}}$.*

Proof. We apply the exponential metric theorem from [8] (Theorem 8), using the coupling $(\mathcal{A}, \mathcal{B})$ and metric ϕ defined above. Metric ϕ satisfies the path requirement of Theorem 8 with U being the set of all pairs of tilings that are adjacent in Ω_n , and by Lemma 10 ϕ takes on values in $\{0\} \cup [1, B]$ for $B = n \log(n)\lambda^{-n}$. Additionally \mathcal{M}_n is lazy, as discussed in Section 2. We have also demonstrated (Lemma 11) that there is a constant c such that $\mathbb{E}[\phi_{t+1}] \leq (1 - qc)\phi_t$ whenever $\lambda < 3^{-1/\sqrt{n}}$. By Theorem 8, we conclude that

$$\tau(\varepsilon) \leq \frac{\ln(n \log(n) \lambda^{-n} \varepsilon^{-1})}{qc} = \frac{4n}{c} (\ln n + \ln(\log n) + n \ln \lambda^{-1} + \ln \varepsilon^{-1}) = O(n^2 + n \log(1/\varepsilon)).$$

When we assume $\varepsilon = 1/4$, as is standard practice, we see $\tau = \tau(1/4) = O(n^2)$. \square

Now that we know the mixing time whenever $\lambda < 3^{-1/\sqrt{n}}$, it only remains to consider sufficiently large n to extend this result to all $\lambda < 1$.

THEOREM 13. *For any constant $\lambda < 1$, the edge-flip chain \mathcal{M}_n on dyadic tilings converges in time $O(n^2)$.*

Proof. For any constant $0 < \lambda < 1$, there is an n_0 such that for all $n > n_0$, $3^{-1/\sqrt{n}} > \lambda$. By Theorem 12, this implies that for all $n > n_0$, the mixing time of \mathcal{M}_n is $O(n^2)$, which suffices to prove the claim. \square

We remark that this mixing time upper bound is within a factor of $O(n/\log n)$ of the true mixing time. We show this by presenting a lower bound on the mixing time of $(n \log n)/4$. This proof holds for Markov chain \mathcal{M}_n and all values of λ , though a better lower bound when $\lambda > 1$ will be presented in the next section. We will use the following well-known result.

THEOREM 14 (see, e.g., [13]). *Let D be the diameter of the state space of an irreducible aperiodic Markov chain. Then for any ε , the mixing time satisfies*

$$\tau(\varepsilon) \geq \frac{D}{2}.$$

THEOREM 15. *The mixing time of Markov chain \mathcal{M}_n for any value of λ is at least $(n \log n)/4$.*

Proof. Let σ_v denote the tiling consisting entirely of $1 \times n$ vertical rectangles and let σ_h denote the tiling consisting entirely of $n \times 1$ horizontal rectangles. Starting at σ_v , the number of edge-flips required to move to σ_h is at least $(n \log n)/2$: one edge-flip at most doubles the width of two rectangles, and n rectangles must have their width doubled at least $\log n$ times each to reach σ_h . This implies the diameter of Ω_n under edge-flip moves is at least $(n \log n)/2$, which by Theorem 14 gives the claimed mixing time lower bound. \square

4. Exponential Time Mixing for General and Dyadic Tilings. In this section, we prove that for certain values of λ both chains can require exponential time to converge. We begin by proving that in both the dyadic and general settings, the Markov chains \mathcal{M}_n and $\widehat{\mathcal{M}}_n$ require exponential time when $\lambda > 1$. Next, we show that for general tilings, when $\lambda < 1$ the Markov chain $\widehat{\mathcal{M}}_n$ also requires exponential time, unlike in the restricted setting of dyadic tilings. In each of these cases, we prove that the Markov chain converges slowly to equilibrium by demonstrating that the state space contains a bottleneck requiring exponential expected time to cross. We use the bottleneck to bound the conductance of the Markov chain. The *conductance* of an ergodic Markov chain \mathcal{M} with stationary distribution π is

$$\Phi_{\mathcal{M}} = \min_{\substack{S \subseteq \Omega \\ \pi(S) \leq 1/2}} \frac{1}{\pi(S)} \sum_{s_1 \in S, s_2 \in \bar{S}} \pi(s_1) \mathcal{P}(s_1, s_2).$$

We can infer a lower bound on the mixing time using the following theorem that relates conductance and mixing time (see, e.g., [21]).

THEOREM 16. *For any Markov chain with conductance $\Phi_{\mathcal{M}}$, $\forall \epsilon > 0$ we have*

$$\tau(\epsilon) \geq \left(\frac{1}{4\Phi_{\mathcal{M}}} - \frac{1}{2} \right) \log \left(\frac{1}{2\epsilon} \right).$$

A change in terminology will be convenient for the remainder of this section whereby we let $|\sigma|$ be the sum of the perimeters of the rectangles in the dissection (or tiling) σ , rather than the total edge length. This will simplify the analysis. Using detailed balance, we reformulate stationary distributions π and $\widehat{\pi}$ for \mathcal{M}_n and $\widehat{\mathcal{M}}_n$ as follows. Let $w(R)$ be the width of rectangle R and $l(R)$ be the length (height) of R . For convenience, we now let $|\sigma|$ denote the total *perimeter* of σ , that is, $|\sigma| = \sum_{R \in \sigma} 2w(R) + 2l(R)$. We note this total perimeter divided by 2 differs from the total edge length of σ by exactly $2n$. By detailed balance, we rewrite

$$\pi(\sigma) = \frac{\lambda^{|\sigma|/2}}{Z} = \frac{(\prod_{R \in \sigma} \lambda^{w(R)+l(R)})}{Z} \quad \text{and} \quad \widehat{\pi}(\sigma) = \frac{(\prod_{R \in \sigma} \lambda^{w(R)+l(R)})}{\widehat{Z}}.$$

Here \widehat{Z} and Z are new normalizing constants, differing from those in Section 2.2 by a multiplicative factor of λ^{2n} .

First, we prove the following lemma bounding the number of n -tilings in the general setting which we use in both slow mixing proofs.

LEMMA 17. *The number of general tilings of Λ_n satisfies $|\widehat{\Omega}_n| \leq (\log n)^n$.*

Proof. Consider any rectangle R in an n -tiling. By assumption R has dimensions $2^w \times 2^h$ for integers $w, h \in \{0, 1, \dots, k = \log n\}$ and thus has $\log n$ possible heights. Given the height of R , the width is uniquely determined since R has area n . To bound the total number of tilings, there are $\log n$ choices for the height of the rectangle that covers the lowest leftmost unit square of Λ_n . Next, consider the rectangle that covers

the lowest leftmost unit square not yet tiled. Given the height of all rectangles ordered in this way the rectangle tiling is uniquely determined. There are n different rectangles with $\log n$ possible heights therefore $|\widehat{\Omega}_n| \leq (\log n)^n$. \square

4.1. Slow Mixing when $\lambda > 1$. We start by showing that for both dyadic and general rectangle tilings when $\lambda > 1$, the Markov chains \mathcal{M}_n and $\widehat{\mathcal{M}}_n$ both take exponential time to converge. Informally, consider the tilings with at least one $n \times 1$ rectangle and those with at least one $1 \times n$ rectangle. In order to go between these sets we must go through a tiling where all rectangles have width and length at least 2 and thus each has perimeter at most $n + 4$. We show these tilings are exponentially unlikely and thus our state space forms a bottleneck.

THEOREM 18. *For any constant $\lambda > 1$, the edge-flip chain \mathcal{M}_n on the set of dyadic tilings Ω_n and the edge-flip chain $\widehat{\mathcal{M}}_n$ on the set of rectangular dissections $\widehat{\Omega}_n$ both require time $\exp(\Omega(n^2))$ to mix.*

Proof. We note that the proofs are identical for \mathcal{M}_n and $\widehat{\mathcal{M}}_n$ and here we show it for $\widehat{\mathcal{M}}_n$. We first partition the state space into two sets, B , the set of tilings with no rectangles of dimension $1 \times n$, and \overline{B} , the remainder. Notice that B contains the tiling σ_h where all rectangles are $n \times 1$ and \overline{B} contains the tiling σ_v where all rectangles are $1 \times n$. Both of these tilings have weight $\pi(\sigma_h) = \pi(\sigma_v) = \widehat{Z}^{-1} \lambda^{n(1+n)}$. Therefore,

$$\begin{aligned} \pi(B) &\geq \pi(\sigma_h) \geq \widehat{Z}^{-1} \lambda^{n(1+n)}, \\ \pi(\overline{B}) &\geq \pi(\sigma_v) \geq \widehat{Z}^{-1} \lambda^{n(1+n)}. \end{aligned}$$

Let $B_c \subset B$ be the set of tilings containing no $(1 \times n)$ or $(n \times 1)$ rectangles. Every rectangle in every tiling in B_c has perimeter at most $n + 4$ and thus has weight at most $\widehat{Z}^{-1} \lambda^{n(n+4)/2}$. By Lemma 17, $|B_c| \leq |\widehat{\Omega}_n| \leq (\log n)^n$; we briefly note this is true in the dyadic case as well although tighter bounds exist. Combining these, we see

$$\pi(B_c) \leq (\log n)^n \widehat{Z}^{-1} \lambda^{n(n+4)/2},$$

which is exponentially smaller than the weight of B and \overline{B} .

Using these bounds, we next bound the conductance of the Markov chain and then the mixing time using Theorem 16. If $\pi(B) \leq 1/2$, then combining the definition of conductance with the bounds on $\pi(B)$ and $\pi(B_c)$ yields

$$\begin{aligned} \Phi_{\widehat{\mathcal{M}}_n} &\leq \frac{1}{\pi(B)} \sum_{b_1 \in B, b_2 \in \overline{B}} \pi(b_1) \mathcal{P}(b_1, b_2) \\ &= \frac{1}{\pi(B)} \sum_{b_1 \in B_c, b_2 \in \overline{B}} \pi(b_1) \mathcal{P}(b_1, b_2) \\ &\leq \frac{1}{\pi(B)} \sum_{b_1 \in B_c} \pi(b_1) \\ &= \frac{\pi(B_c)}{\pi(B)} \leq \frac{(\log n)^n \widehat{Z}^{-1} \lambda^{n(n+4)/2}}{\widehat{Z}^{-1} \lambda^{n(1+n)}} = \frac{(\log n)^n}{\lambda^{n^2/2-n}} = \lambda^{-c_1 n^2}, \end{aligned}$$

for constant c_1 and n sufficiently large when λ is a constant greater than 1. Alternately, if $\pi(B) > 1/2$ then $\pi(\overline{B}) \leq 1/2$ and so by detailed balance and the bounds on $\pi(\overline{B})$ and $\pi(B_c)$,

$$\begin{aligned} \Phi_{\widehat{\mathcal{M}}_n} &\leq \frac{1}{\pi(\overline{B})} \sum_{b_1 \in B, b_2 \in \overline{B}} \pi(b_2) \mathcal{P}(b_2, b_1) \\ &= \frac{1}{\pi(\overline{B})} \sum_{b_1 \in B, b_2 \in \overline{B}} \pi(b_1) \mathcal{P}(b_1, b_2) \\ &= \frac{1}{\pi(\overline{B})} \sum_{b_1 \in B_c, b_2 \in \overline{B}} \pi(b_1) \mathcal{P}(b_1, b_2) \\ &\leq \frac{\pi(B_c)}{\pi(\overline{B})} \leq \frac{(\log n)^n \widehat{Z}^{-1} \lambda^{n(n+4)/2}}{\widehat{Z}^{-1} \lambda^{n(1+n)}} = \frac{(\log n)^n}{\lambda^{n^2/2-n}} = \lambda^{-c_1 n^2}, \end{aligned}$$




FIG. 13. The tiling s_{64} .

for constant c_1 defined above and n sufficiently large, whenever λ is a constant greater than 1.

In both cases, $\Phi_{\widehat{\mathcal{M}}_n} \leq \lambda^{-c_1 n^2}$. Applying Theorem 16 proves that for all $\epsilon > 0$, the mixing time of $\widehat{\mathcal{M}}_n$ satisfies

$$\tau(\epsilon) \geq \left(\lambda^{c_1 n^2} / 4 - \frac{1}{2} \right) \log \left(\frac{1}{2\epsilon} \right) = \Omega(\lambda^{c_1 n^2} \ln \epsilon^{-1}).$$

Letting $\epsilon = 1/4$ we have that $\tau = \Omega(\lambda^{c_1 n^2})$, as desired. \square

4.2. Slow Mixing for General Tilings, $\lambda < 1$. Finally, we consider general tilings when $\lambda < 1$ and show that in this setting $\widehat{\mathcal{M}}_n$ takes exponential time to converge by again demonstrating a bottleneck in the state space. In this case however the bottleneck is much more complex. A key fact is that if the total width of the region being filled with rectangles is $n = 2^k$, and there are at least two rectangles with width 1, then there must be many other thin rectangles in the rectangular dissection. We define the cut set to consist of rectangular dissections that are forced to have significantly more thin rectangles in order to show that there is a bad cut in the state space.

Define a *bar* to be a rectangle of width 1 and length (height) n . The bottleneck in $\widehat{\Omega}_n$ is based on the *separation* of a tiling which measures the distance between the bars in the tiling. More formally, define the *distance* between two bars to be the difference in their x -coordinates plus one. For example, two adjacent bars are at distance 2 and two bars separated by a rectangle of size $2 \times n/2$ are at distance 4. Given an n -tiling, pair the bars in order from left to right (there must be an even number of bars since $n = 2^k$). The *separation* of a tiling is the sum of the distances between each pair of bars. Let S be the set of tilings with separation greater than or equal to $n/2 + 2$ and \bar{S} be the remaining tilings, namely those with separation less than $n/2 + 2$. We show all moves from S to \bar{S} involve a tiling with at least 4 bars and separation $n/2 + 2$, and the total weight of this set of tilings is exponentially smaller than the weight of both S and \bar{S} .

THEOREM 19. *For any constant $\lambda < 1$, the edge-flip chain $\widehat{\mathcal{M}}_n$ on $\widehat{\Omega}_n$ requires time $\exp(\Omega(n \log n))$.*

Proof. We begin by proving a lower bound on $\pi(S)$ and $\pi(\bar{S})$. Let g_n be the “ground state” tiling consisting entirely of rectangles of size $\sqrt{n} \times \sqrt{n}$. This tiling has perimeter $|g_n| = 4n\sqrt{n}$. Since $g_n \in \bar{S}$ because g_n has no bars and thus separation 0, this implies that $\pi(\bar{S}) > \pi(g_n) = \widehat{Z}^{-1} \lambda^{2n\sqrt{n}}$. Next we will define a special tiling $s_n \in S$. Let s_n have one bar on the far left side of Λ_n and one bar on the far right side of Λ_n . Next to the leftmost bar there is a column with two rectangles of width 2 followed by a column with four rectangles of width 4 and so forth until there is a column with only rectangles of width $2^{k/2-1}$. The remainder of the tiling is filled with rectangles of size $\sqrt{n} \times \sqrt{n}$. Note that the combined width of these columns is $2 + \sum_{i=1}^{k/2-1} 2^i = \sqrt{n}$ so the remainder of the tiling has width $n - \sqrt{n}$ and can be tiled with $n - \sqrt{n}$ rectangles of size $\sqrt{n} \times \sqrt{n}$. Figure 13 shows s_{64} . Configuration s_n has perimeter

$$|s_n| = 4(1 + 2^k) + \sum_{i=1}^{k/2-1} (2^i 2(2^i + 2^{k-i})) + (n - \sqrt{n})4\sqrt{n} = 4n^{3/2} + n \log n - (4/3)n + (4/3).$$

As $s_n \in S$ because it has separation n , this implies $\pi(S) > \pi(s_n) = \widehat{Z}^{-1} \lambda^{|s_n|/2}$.

Let S_C be the set of tilings in S from which it is possible to transition to \bar{S} . We will prove that every tiling in S_C has at least four bars and separation exactly $n/2 + 2$. We use the following lemma.

LEMMA 20. *One move of the chain $\widehat{\mathcal{M}}_n$ changes the separation of a tiling by 0, +2 or -2.*

Proof. The only moves of the Markov chain that change the separation are when two bars are added or removed. Let's consider adding two bars first. Let P be the pairing of the bars in the configuration before the two bars are added. There are two cases; either the two bars are added between two bars that were paired in P , or between two pairs of bars. If they are added between two pairs, then they will be paired up in the new pairing and add 2 to the separation. If they are added between two bars b_l and b_r paired in P with distance d , then the new bars will be paired with b_l and b_r . The sum of the distances will remain unchanged. Next, consider the case where two bars are removed. Again, there are two cases. If the two bars are paired, then the separation decreases by two however if the two bars are paired with two other bars the distance remains unchanged. \square

The ground state configuration g_n has separation 0. Since all tilings are connected by the Markov chain $\widehat{\mathcal{M}}_n$ which by Lemma 20 changes the separation by an even number at each step, this implies that the separation of all tilings is even. Additionally, to go from S to \bar{S} we must go through a tiling with separation exactly $n/2 + 2$. Given a tiling with two bars and separation $n/2 + 2$ there is no way to decrease the separation and thus no way to transition to \bar{S} . Thus, every tiling in S_C has separation $n/2 + 2$ and at least four bars. Next, we will upper bound the weight of each tiling σ in S_C . To do this, we lower bound the perimeter of any tiling of a lattice region of size $(n/2 - 2) \times n$ and then show that every tiling in S_C has two such regions.

LEMMA 21. *Any tiling σ' of an $(n/2 - 2) \times n$ region has perimeter $|\sigma'| \geq 2n^{3/2} + n \log n - (16/3)n - (8/3)$.*

Proof. We will assign each unit square in the lattice region a weight based on the perimeter of the rectangle the square is contained in so that the combined weight of all n squares within a rectangle is equal to the perimeter of the rectangle. Assume the unit square at location (i, j) is contained in a rectangle of size $2^a \times 2^{k-a}$ then the weight $w_{i,j} = 2(2^a + 2^{k-a})/2^k$. Since each rectangle has area 2^k , the sum of all weights $\sum_{j=1}^{n/2-2} \sum_{i=1}^n w_{i,j} = |\sigma'|$. The width of the region is $n/2 - 2$, and this number has binary representation $011\dots110$. Since each rectangle has width 2^a for some integer a this implies that in each row, for each integer $\ell = 1$ to $k/2 - 1$ there must be either a rectangle of width 2^ℓ or multiple rectangles of width smaller than 2^ℓ whose widths add up to 2^ℓ . If there is a single rectangle of width 2^ℓ then the 2^ℓ unit squares in this row contained in this rectangle each have weight $2(2^\ell + 2^{k-\ell})/2^k$. If there is instead multiple smaller rectangles then they will have larger perimeter and thus larger weight. Thus, the combined weight of these unit squares in each row is at least

$$\sum_{\ell=1}^{k/2-1} 2^\ell (2(2^\ell + 2^{k-\ell})/2^k) = \log n - (4/3) - 8/(3n).$$

Since the minimum perimeter rectangle is the $2^{k/2} \times 2^{k/2}$ square, $w_{i,j} \geq 4/2^{k/2}$. Thus the remaining $2^{k-1} - 2^{k/2}$ unit squares in each row have combined weight at least

$$4(2^{k-1} - 2^{k/2})/2^{k/2} = 2\sqrt{n} - 4.$$

This implies that the total perimeter satisfies

$$|\sigma'| = \sum_{i=1}^n \sum_{j=1}^{n/2-2} w_{i,j} \geq \sum_{i=1}^n (\log n - (4/3) - 8/(3n) + 2\sqrt{n} - 4) = 2n^{3/2} + n \log n - (16/3)n - 8/3.$$

This is the desired result. \square

Consider any tiling σ with separation $n/2 + 2$ and at least four bars. Label the bars b_1, b_2, \dots, b_B from left to right. Next, label the regions between the pairs of bars $p_1, p_2, \dots, p_{B/2}$ and the gaps between the pairs $g_0, g_1, g_2, \dots, g_{B/2}$, as shown in Figure 14. Let $w(p_i), w(g_i)$ denote the widths of the regions between the bars. Now, since σ has separation $n/2 + 2$, this implies that $\sum_{i=1}^{B/2} (w(p_i) + 2) = n/2 + 2$. Reorder the tiling so that the bars and regions are in the following order: $g_0, g_1, \dots, g_{B/2}, b_1, b_2, \dots, b_{B-1}, p_1, p_2, \dots, p_{B/2}, b_B$. Notice that the region $b_4 \cup \dots \cup b_{B-1} \cup p_1 \cup \dots \cup p_{B/2}$ has width $n/2 - 2$ as does the region $g_0 \cup g_1 \cup \dots \cup g_{B/2}$.




FIG. 14. An example labeling of the bars and regions surrounding bars.

Thus we can apply Lemma 21 to show that the total perimeter of tiling σ must be at least

$$|\sigma| \geq 4(2 + 2^{k+1}) + 2 \left(2n^{3/2} + n \log n - \frac{16}{3}n - \frac{8}{3} \right) = 4n^{3/2} + 2n \log n - (8/3)n + 8/3.$$

Combining this bound with the bound on the number of tilings from Lemma 17 gives

$$\pi(S_C) \leq \widehat{Z}^{-1} (\log n)^n \lambda^{2n^{3/2} + n \log n - (4/3)n + (4/3)},$$

which is exponentially smaller than the above bound on $\pi(S)$, as desired. Using these bounds we bound the conductance of the Markov chain and then the mixing time using Theorem 16. If $\pi(S) \leq 1/2$, then combining the definition of conductance with the bounds on $\pi(S)$ and $\pi(S_C)$ yields

$$\begin{aligned} \Phi_{\widehat{\mathcal{M}}_n} &\leq \frac{1}{\pi(S)} \sum_{s_1 \in S, s_2 \in \bar{S}} \pi(s_1) \mathcal{P}(s_1, s_2) \\ &= \frac{1}{\pi(S)} \sum_{s_1 \in S_C, s_2 \in \bar{S}} \pi(s_1) \mathcal{P}(s_1, s_2) \\ &\leq \frac{1}{\pi(S)} \sum_{s_1 \in S_C} \pi(s_1) \\ &= \frac{\pi(S_C)}{\pi(S)} \\ &\leq \frac{(\log n)^n \lambda^{2n^{3/2} + n \log n - (4/3)n + (4/3)}}{\lambda^{2n^{3/2} + (n \log n)/2 - (2/3)n + (2/3)}} \\ &= (\log n)^n \lambda^{(n \log n - (4/3)n + (4/3))/2} \\ &= 2^n \log n \lambda^{(n \log n - (4/3)n + (4/3))/2} = \lambda^{c_2 n \log n}, \end{aligned}$$

for constant c_2 and n sufficiently large when $\lambda < 1$ is a constant.

If $\pi(S) > 1/2$, then $\pi(\bar{S}) \leq 1/2$, and by detailed balance and bounds on $\pi(S)$, $\pi(\bar{S})$ and $\pi(S_C)$,

$$\begin{aligned} \Phi_{\widehat{\mathcal{M}}_n} &\leq \frac{1}{\pi(\bar{S})} \sum_{s_1 \in \bar{S}, s_2 \in S} \pi(s_1) \mathcal{P}(s_1, s_2) \\ &= \frac{1}{\pi(\bar{S})} \sum_{s_1 \in \bar{S}, s_2 \in S} \pi(s_2) \mathcal{P}(s_2, s_1) \\ &= \frac{1}{\pi(\bar{S})} \sum_{s_1 \in S_C, s_2 \in \bar{S}} \pi(s_1) \mathcal{P}(s_1, s_2) \\ &\leq \frac{\pi(S_C)}{\pi(\bar{S})} \leq \frac{(\log n)^n \lambda^{4n^{3/2} + 2n \log n - 2n}}{\lambda^{4n^{3/2} + n \log n - n}} = \lambda^{c_2 n \log n}, \end{aligned}$$

for constant c_2 defined above and n sufficiently large when $\lambda < 1$ is a constant. In both cases,

$$\Phi_{\widehat{\mathcal{M}}_n} \leq \lambda^{c_2 n \log n}.$$

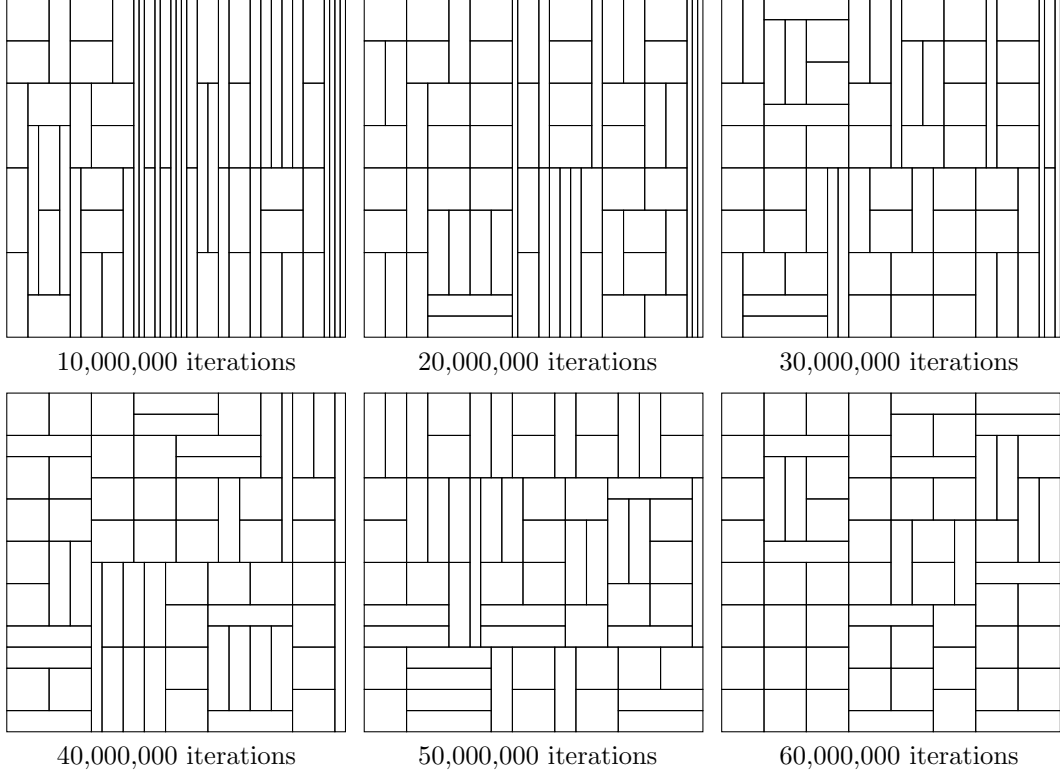


FIG. 15. A simulated execution of a Markov chain on $\widehat{\Omega}_{64}$ with $\lambda = 0.8$ that includes both edge-flip moves and additional moves that select a random rectangular region and rotate the tiling within it by 180 degrees, starting with all vertical rectangles of width 1 and height 64.

Applying Theorem 16 proves that the mixing time of $\widehat{\mathcal{M}}_n$ satisfies

$$\tau(\epsilon) \geq \left(\frac{\lambda^{-c_2 n \log n}}{4} - \frac{1}{2} \right) \log \left(\frac{1}{2\epsilon} \right) = \Omega(\lambda^{-c_2 n \log n} \ln \epsilon^{-1}).$$

Letting $\epsilon = 1/4$, we have that $\tau = \Omega(\lambda^{-c_2 n \log n})$, and we have proved Theorem 19.

5. Conclusions. In this paper, we have bounded the mixing times of both dyadic and general rectangular dissections in the entire $\lambda > 1$ and $\lambda < 1$ regimes. For dyadic tilings we have shown there is a phase transition at $\lambda = 1$ by giving polynomial upper bounds on the mixing time when $\lambda < 1$ and exponential lower bounds when $\lambda > 1$. For general rectangular dissections we have shown the mixing time is at least exponential whenever $\lambda \neq 1$, but the reasons for this are different when $\lambda < 1$ and $\lambda > 1$.

Bounding the mixing time when $\lambda = 1$ is a natural next question. For dyadic tilings, Cannon, Levin, and Stauffer [4] have shown the mixing time at $\lambda = 1$ is polynomial, but for rectangular dissections this question remains open. Simulations seem to suggest fast mixing also occurs in this case, but new approaches and techniques are needed to analyze this.

There are many other natural directions for extending this work. It will be interesting to explore whether adding additional moves can speed up the mixing time of the edge-flip Markov chain. For instance, we proved the edge-flip Markov chain for rectangular dissections is slowly mixing when $\lambda < 1$; we conjecture that if we add additional moves that rotate the tiling within a rectangular region by 180 degrees, the mixing time will become polynomial. Our simulations support this hypothesis (see Figure 15 for an example). While we proved that vertical bars are an obstruction to fast mixing for the edge-flip chain on rectangular dissections when $\lambda < 1$ (Section 4.2), when these additional rotation moves are added this bottleneck no longer exists and simulations show that bars at all scales are eliminated. Note that while the number of iterations is much higher for this simulation than the one presented at the beginning of this paper (Figure 2), we expect a slowdown by a factor of about 32 when adding rotations: each potential move is now proposed with

probability $1/64^2$, by randomly selecting grid points to be the corners of a rectangle that is rotated or flipped, instead of with probability $1/(2 \cdot 64)$ as they are in the edge-flip chain.

These simulations also provide evidence that the behavior of rectangular dissections is different when $\lambda > 1$ and when $\lambda < 1$, because the bottleneck that produces slow mixing when $\lambda > 1$ still exists when rotations of 180 degrees are added. Conversely, if 90 degree rotations of the tilings within rectangular subregions (with necessary rescalings) are added to the edge-flip chain, we expect polynomial mixing time when $\lambda > 1$ while it remains unclear if there would also be polynomial mixing when $\lambda < 1$. To distinguish further between the $\lambda < 1$ and $\lambda > 1$ regimes for rectangular dissections, we expect that from a warm start the edge-flip chain when $\lambda < 1$ will mix in polynomial time, while we do not expect this is true when $\lambda > 1$. Next steps include further investigations to rigorously differentiate between the $\lambda < 1$ and $\lambda > 1$ regimes for rectangular dissections, as well as to understand the behavior when $\lambda = 1$, where simulations seem to suggest that fast mixing occurs.

REFERENCES

- [1] O. ANGEL, A. HOLROYD, G. KOZMA, J. WÄSTLUND, AND P. WINKLER, *The phase transition for dyadic tilings*, Transactions of the AMS, 366 (2014), pp. 1029–1046.
- [2] E. BOROS AND Z. FÜREDI, *Rectangular dissections of a square*, European Journal of Combinatorics, 9 (1988), pp. 271–280.
- [3] R. BUBLEY AND M. DYER, *Path coupling: A technique for proving rapid mixing in Markov chains*, in Proceedings of the 38th Annual Symposium on Foundations of Computer Science, FOCS, IEEE Computer Society, 1997, pp. 223–231.
- [4] S. CANNON, D. LEVIN, AND A. STAUFFER, *Polynomial mixing of the edge-flip Markov chain for unbiased dyadic tilings*, in Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2017), 2017, pp. 34:1–34:21.
- [5] P. CAPUTO, F. MARTINELLI, A. SINCLAIR, AND A. STAUFFER, *Random lattice triangulations: Structure and algorithms*, The Annals of Applied Probability, 25 (2015), pp. 1650–1685.
- [6] G. DEMICHELI, P. ANTOGNETTI, AND A. SANGIOVANNI-VINCENTELLI, eds., *Design systems for VLSI circuits: logic synthesis and silicon compilation*, Springer, 1987.
- [7] N. ELKIES, G. KUPERBERG, M. LARSEN, AND J. PROPP, *Alternating-sign matrices and domino tilings*, Journal of Algebraic Combinatorics, 1 (1992), pp. 111–132.
- [8] S. GREENBERG, A. PASCOE, AND D. RANDALL, *Sampling biased lattice configurations using exponential metrics*, in 20th Symposium on Discrete Algorithms, vol. 21, 2009, pp. 225–251.
- [9] R. HAGGKVIST, P. LINDBERG, AND B. LINDSTROM, *Dissecting a square into rectangles of equal area*, Discrete Mathematics, 47 (1983), pp. 321–323.
- [10] S. JANSON, D. RANDALL, AND J. SPENCER, *Random dyadic tilings of the unit square*, in Random Structures and Algorithms, vol. 21, 2002, pp. 225–251.
- [11] W. JOCKUSCH, J. PROPP, AND P. SHOR, *Random domino tilings and the arctic circle theorem*, arXiv:math/9801068v1, (1998).
- [12] R. KENYON, *Conformal invariance of domino tilings*, The Annals of Probability, 28 (2000), pp. 759–795.
- [13] D. LEVIN, Y. PERES, AND E. WILMER, *Markov chains and mixing times*, Providence, R.I. American Mathematical Society, 2009.
- [14] J. D. LOERA, J. RAMBAU, AND F. SANTOS, *Triangulations: Structures for Algorithms and Applications*, Springer, 2010.
- [15] M. LUBY, D. RANDALL, AND A. SINCLAIR, *Markov chain algorithms for planar lattice structures*, SIAM Journal on Computing, 31 (2001), pp. 167–192.
- [16] L. MCSHINE AND P. TETALI, *On the mixing time of the triangulation walk and other Catalan structures*, DIMACS-AMS Volume on Randomization Methods in Algorithm Design, 43 (1998), pp. 147–160.
- [17] M. MOLLOY, B. REED, AND W. STEIGER, *On the mixing rate of the triangulation walk*, DIMACS-AMS Volume on Randomization Methods in Algorithm Design, 43 (1998), pp. 179–190.
- [18] H. MURATA, K. FUJIYOSHI, T. WATANABE, AND Y. KAJITANI, *A mapping from sequence-pair to rectangular dissection*, in Design Automation Conference, 1997, pp. 625–633.
- [19] J. PROPP AND D. WILSON, *Exact sampling with coupled Markov chains and applications to statistical mechanics*, Random Structures and Algorithms, 9 (1996), pp. 223–252.
- [20] D. RANDALL AND P. TETALI, *Analyzing Glauber dynamics by comparison of Markov chains*, J. Mathematical Physics, 41 (2000), pp. 1598–1615.
- [21] A. SINCLAIR, *Algorithms for Random Generation & Counting: a Markov Chain Approach*, Birkhäuser, Boston, 1993.
- [22] W. TUTTE, *A census of planar maps*, Canadian Journal of Mathematics, 15 (1963), pp. 249–271.
- [23] J. WANG AND H. ZHOU, *Linear constraint graph for floorplan optimization with soft blocks*, in International Conference on Computer-Aided Design (ICCAD’08), 2008.
- [24] E. WELZL, *The number of (and random) triangulations of planar point sets*, Presentation at Oberwolfach Workshop on Combinatorics, Probability and Computing, (2006).
- [25] D. WILSON, *Mixing times of lozenge tiling and card shuffling Markov chains*, Ann. Appl. Probab., 14 (2004), pp. 274–325.
- [26] G. ZOBRIST, *Routing, Placement, and Partitioning*, Intellect Ltd., 1994.