
Polynomial Mixing of the Edge-Flip Markov Chain for Unbiased Dyadic Tilings

S. CANNON^{1,†}, D. A. LEVIN² and A. STAUFFER^{3‡}

¹Computer Science Division, University of California Berkeley, Berkeley, CA 94720-1776, USA
(e-mail: sarah.cannon@berkeley.edu)

²Department of Mathematics, University of Oregon, Eugene, OR 97403-1222, USA
(e-mail: dlevin@uoregon.edu)

³Department of Mathematical Sciences, University of Bath, Claverton Down, Bath BA2 7AY, UK
(e-mail: a.stauffer@bath.ac.uk)

Received 1 May 2017; revised 21 August 2018; first published online 31 October 2018

We give the first polynomial upper bound on the mixing time of the edge-flip Markov chain for unbiased dyadic tilings, resolving an open problem originally posed by Janson, Randall and Spencer in 2002 [14]. A *dyadic tiling* of size n is a tiling of the unit square by n non-overlapping dyadic rectangles, each of area $1/n$, where a *dyadic rectangle* is any rectangle that can be written in the form $[a2^{-s}, (a+1)2^{-s}] \times [b2^{-t}, (b+1)2^{-t}]$ for $a, b, s, t \in \mathbb{Z}_{\geq 0}$. The edge-flip Markov chain selects a random edge of the tiling and replaces it with its perpendicular bisector if doing so yields a valid dyadic tiling. Specifically, we show that the relaxation time of the edge-flip Markov chain for dyadic tilings is at most $O(n^{4.09})$, which implies that the mixing time is at most $O(n^{5.09})$. We complement this by showing that the relaxation time is at least $\Omega(n^{1.38})$, improving upon the previously best lower bound of $\Omega(n \log n)$ coming from the diameter of the chain.

2010 *Mathematics subject classification*: Primary 60J10
Secondary 68Q87, 05B45, 52C20

1. Introduction

We study the edge-flip Markov chain for dyadic tilings. An interval is *dyadic* if it can be written in the form $[a2^{-s}, (a+1)2^{-s}]$ for non-negative integers a and s with $0 \leq a < 2^s$.

[†] Supported by a grant from the Simons Foundation (#361047 to Sarah Cannon) and the National Science Foundation Graduate Research Fellowship Program under grant 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.

[‡] Supported by a Marie Curie Career Integration Grant PCIG13-GA-2013-618588 DSRELIS, and an EPSRC Early Career Fellowship.




Figure 1. (a) A dyadic tiling of size 16 with a vertical bisector. (b) A dyadic tiling of size 16 with both a vertical and horizontal bisector. (c) A tiling that is not dyadic; the vertical component of the shaded rectangles is not a dyadic interval.

A rectangle is dyadic if it is the Cartesian product of two dyadic intervals. A *dyadic tiling of size n* is a tiling of the unit square by n non-overlapping dyadic rectangles with the same area $1/n$; see Figure 1. Lagarias, Spencer and Vinson [15] showed that dyadic tilings are precisely those tilings that can be constructed by bisecting the unit square, either horizontally or vertically; bisecting each half again, either horizontally or vertically; and repeatedly bisecting all remaining rectangular regions until there are n total dyadic rectangles, each of equal area. We necessarily assume n is a power of 2. There is a natural Markov chain which connects the state space of all dyadic tilings of size n by moves we refer to as *edge flips*.

We analyse this *edge-flip Markov chain* over the set of dyadic tilings of size n . Given any dyadic tiling, this chain evolves by selecting an edge of the tiling uniformly at random and replacing it by its perpendicular bisector, if doing so yields a valid dyadic tiling of size n ; an illustration is given in Figure 2(a). Our main result gives the first polynomial upper bound for the mixing time of this Markov chain. (The precise definitions of mixing time and relaxation time are deferred to Section 2.2.) *In this paper, all logarithms have base 2.*

Theorem 1.1. *The relaxation time of the edge-flip Markov chain for dyadic tilings of size n is at most $O(n^{\log 17})$. As a consequence, the mixing time of this chain is at most $O(n^{1+\log 17})$.*

The best previously known lower bound for the mixing time is $\Omega(n \log n)$, which is a simple consequence of the fact that the diameter of the Markov chain is of order $n \log n$ [14]. In the theorem below we improve this bound.

Theorem 1.2. *The relaxation time and mixing time of the edge-flip Markov chain for dyadic tilings of size n are both at least $\Omega(n^{2 \log \phi})$, where $\phi = (\sqrt{5} + 1)/2$ is the golden ratio.*

We note that $\log 17 \approx 4.09$ and $2 \log \phi \approx 1.38$.

1.1. Related work

The edge-flip Markov chain for dyadic tilings was first considered by Janson, Randall and Spencer in 2002 [14], who showed that it is irreducible but left it as an open

problem to derive that the mixing time is polynomial in n . Instead, they presented another Markov chain, which has additional global moves consisting of rotations at all scales, and showed that this chain mixes in polynomial time. However, applications of the comparison technique of Diaconis and Saloff-Coste [9] have failed to extend this polynomial mixing bound to the more natural edge-flip Markov chain (which, in fact, corresponds to only performing rotations at the smallest scale).

Cannon, Miracle and Randall considered the mixing time of the edge-flip Markov chain for a weighted version of dyadic tilings [2]. In this version, given a parameter $\lambda > 0$, the stationary probability of a dyadic tiling x is proportional to $\lambda^{|x|}$, where $|x|$ is the sum of the length of the edges of x . The Metropolis rule [22] is incorporated into the edge-flip Markov chain so that the chain has the desired stationary distribution. They showed that the mixing time of this chain is at least exponential in n^2 for any $\lambda > 1$, and at most $O(n^2 \log n)$ for any $\lambda < 1$. This establishes a phase transition at critical point $\lambda = 1$, which corresponds to the unweighted case considered here. However, their techniques did not extend to the critical point, and they left as an open problem bounding the mixing time when $\lambda = 1$; it is notoriously often quite difficult to bound mixing times at or near critical points. Our main result, Theorem 1.1, uses a different, non-local approach to finally answer the question of [14] and [2] by showing that the mixing time of the edge-flip Markov chain at critical point $\lambda = 1$ is at most polynomial in n , substantially less than the mixing time when $\lambda > 1$. Furthermore, our Theorem 1.2 combined with the result for the weighted case in [2] shows that the behaviour at the (unweighted) critical point $\lambda = 1$ is also substantially different than when $\lambda < 1$. While it follows from the path coupling analysis in [2] that the relaxation time is $O(n)$ for all fixed $\lambda < 1$, Theorem 1.2 establishes a super-linear lower bound on the relaxation time when $\lambda = 1$.

It is a general principle in statistical physics that in systems with some bias parameter (*temperature*) that induces different phases, the mixing time of natural heat-bath dynamics should be as fast as possible (the diameter of the state space) at high temperature, a larger polynomial at the critical temperature, and exponential at low temperature. However, there are very few instances for which this behaviour has been rigorously confirmed. Exceptions are the Ising model on complete graphs [10, 16], regular trees [11], and the two-dimensional lattice [18], and the Potts model on the complete graph [8] and the two-dimensional lattice [12], all of which required significant effort to analyse. The edge-flip Markov chain for dyadic tilings is an example of heat-bath dynamics, and the parameter λ introduced by Cannon, Miracle and Randall can be viewed as a function of inverse temperature. Their work confirms exponential mixing at low temperature ($\lambda > 1$) and polynomial mixing at high temperature ($\lambda < 1$). Our work shows that the mixing time at the critical point ($\lambda = 1$) is indeed polynomial but strictly larger than the diameter of the state space (which is $n \log(n)/2$), providing further evidence for the general statistical physics principle above.

Variants of the edge-flip Markov chain offer a natural way to sample from many structures, but establishing rigorous polynomial upper bounds on the mixing time has often proved difficult, even in simple cases. Perhaps the most studied case is that of triangulations of a given point set; efficiently generating uniformly random triangulations of general planar point sets has been a problem of great interest in computer graphics

and computational geometry. However, the mixing time of the edge-flip Markov chain for triangulations remains open in the general case, and no polynomial upper bound is known. The only known exception is for n points in *convex position*, which corresponds to triangulations of a convex polygon. In this case, the edge-flip Markov chain is known to mix in at most $O(n^5)$ steps [21], but the correct order of the mixing time is still unknown. For the case of lattice triangulations, which are triangulations of an $m \times n$ grid of points, no polynomial upper bound on the mixing time is known even when $m \geq 2$ is kept fixed as $n \rightarrow \infty$. The only known results in this case are limited to the weighted case [3, 4, 25].

Another example of a related Markov chain that uses natural edge-flip type moves is the *switch Markov chain* for sampling from graphs with a given degree sequence. In this chain, at each iteration two random non-adjacent edges are removed and their four endpoints are randomly rematched; the move is rejected if it results in a multiple edge. Again, in the general case the mixing time of this Markov chain is unknown, though polynomial upper bounds exist when certain restrictions are placed on the degree sequence [7, 13].

For the case of rectangular tilings, results for the mixing time of the edge-flip Markov chain have been quite rare. One important result was obtained for domino tilings, which are tilings of an $n \times n$ square by rectangles of dimensions 1×2 or 2×1 . In this case, the edge-flip Markov chain is known to mix in time polynomial in the number of dominoes, a result that heavily relies on the connection between domino tilings and random lattice paths [19, 23].

The case of dyadic tilings exhibits interesting asymptotic properties that have been studied by combinatorialists [14, 15]. Tilings in which all rectangles are dyadic, but may have different areas, have been used as a basis for subdivision algorithms to solve problems such as approximating singular algebraic curves [1] and classifying data using decision trees [24]. In both of these examples, the unit square is repeatedly subdivided into smaller and smaller dyadic rectangles until the desired approximation or classification is achieved, with more subdivisions in the areas of the most interest (e.g. near the algebraic curve or where data classified differently are close together).

1.2. Proof ideas

We identify a certain block structure on dyadic tilings that allows us to relate the spectral gap of the edge-flip Markov chain to that of another, simpler Markov chain. In the simpler Markov chain, which we refer to as the block dynamics, for each transition a large region of the tiling is selected and retiled uniformly at random, if possible. At the smallest scale, $n = 4$, these correspond to exactly the moves of the (lazy) edge-flip Markov chain. The structure of these block moves allows us to set up a recursion that relates the spectral gap of the edge-flip Markov chain for tilings of size n with that of sizes smaller than n and that of the block dynamics. This produces an inverse polynomial lower bound on the spectral gap of the edge-flip Markov chain.

Specifically, we adapt a bisection approach inspired by spin system analysis [5, 20]. We bound the spectral gap γ_k of the Markov chain \mathcal{M}_k for dyadic tilings of size $n = 2^k$ by the product of the spectral gap γ_{block} of the block dynamics Markov chain and the spectral gap γ_{k-1} of \mathcal{M}_{k-1} , and then use recursion to obtain $\gamma_k \geq (\gamma_{\text{block}})^k = (\gamma_{\text{block}})^{\log n}$. As γ_{block} is constant, this implies a polynomial relaxation time and thus a polynomial mixing time.

For the explicit upper bound in Theorem 1.1, we use a coupling argument to bound γ_{block} ; see e.g. [17, Chapter 13]. The distance metric we use is a carefully weighted average of two different notions of distance between tilings. We do a case analysis and show that this distance metric contracts by a factor of at least $1 - 1/17$ in each step, implying that the spectral gap γ_{block} is at least $1/17$.

We use a distinguishing statistic to show that the mixing time and relaxation time of the edge-flip Markov chain for dyadic tilings are at least $\Omega(n^{1.38})$; again; see [17, Chapter 13]. That is, we define a specific function f on the state space of all dyadic tilings of size $n = 2^k$. By considering the variance and Dirichlet form of f , and using combinatorial properties of dyadic tilings, we can give an upper bound on the spectral gap and thus a lower bound on the relaxation and mixing times.

2. Background

Here we present some necessary information on dyadic tilings, including their asymptotic behaviour, and on Markov chains, including mixing time and local variance.

2.1. Dyadic tilings

A *dyadic interval* is an interval that can be written in the form $[a2^{-s}, (a + 1)2^{-s}]$ for non-negative integers a and s with $0 \leq a < 2^s$. A *dyadic rectangle* is the product of two dyadic intervals. A *dyadic tiling of size $n = 2^k$* is a tiling of the unit square by n dyadic rectangles of equal area $1/n = 2^{-k}$ that do not overlap except on their boundaries; see Figure 1. Let Ω_k be the set of all dyadic tilings of size $n = 2^k$. We say a dyadic tiling has a *vertical bisector* if the line $x = 1/2$ does not intersect the interior of any dyadic rectangle in the tiling. We say it has a *horizontal bisector* if the same is true of the line $y = 1/2$. It is easy to prove that every dyadic tiling of size $n > 1$ has a horizontal bisector or a vertical bisector.

The asymptotics of dyadic tilings were first explored by Lagarias, Spencer and Vinson [15], and we present a summary of their results. Let $A_k = |\Omega_k|$ denote the number of dyadic tilings of size $n = 2^k$. The unit square is the unique dyadic tiling consisting of one dyadic rectangle, so $A_0 = 1$. There are two dyadic tilings of size 2, since the unit square may be divided by either a horizontal or vertical bisector, so $A_1 = 2$. One can also observe that $A_2 = 7$, $A_3 = 82$, $A_4 = 11\,047$, ... In fact, the values A_k can be shown to satisfy the recurrence $A_k = 2A_{k-1}^2 - A_{k-2}^4$; we include a proof of this fact as presented in [14], because we will use these ideas later.

Proposition 2.1 ([15]). For $k \geq 2$, the number of dyadic tilings of size 2^k is

$$A_k = 2A_{k-1}^2 - A_{k-2}^4.$$

Proof. A dyadic tiling of size 2^k has a horizontal bisector, a vertical bisector, or both. If it has a vertical bisector, the number of ways to tile the left half of the unit square is A_{k-1} ; by mapping $x \rightarrow 2x$, we can see that the left half of a dyadic tiling of size 2^k is equivalent to a dyadic tiling of the unit square of size 2^{k-1} because dyadic rectangles

scaled by factors of two remain dyadic. Similarly, mapping $x \rightarrow 2x - 1$, the right half of a dyadic tiling of size 2^k is equivalent to a dyadic tiling of size 2^{k-1} . We conclude that the number of dyadic tilings of size 2^k with a vertical bisector is A_{k-1}^2 . Similarly, by appealing to the maps $y \rightarrow 2y$ and $y \rightarrow 2y - 1$, the number of dyadic tilings of size 2^k with a horizontal bisector is A_{k-1}^2 . The number of dyadic tilings of size 2^k with both a horizontal and a vertical bisector is A_{k-2}^4 , as each quadrant of any such tiling is equivalent to a dyadic tiling of size 2^{k-2} . This follows from appealing to the map $(x, y) \rightarrow (2x, 2y)$ for the lower left quadrant, and appropriate translations of this for the other three quadrants. Altogether, we see $A_k = A_{k-1}^2 + A_{k-1}^2 - A_{k-2}^4 = 2A_{k-1}^2 - A_{k-2}^4$. \square

It is believed that this recurrence does not have a closed-form solution. As proved in [15], $A_k \approx \phi^{-1}\omega^{2^k} = \phi^{-1}\omega^n$, where $\phi = (1 + \sqrt{5})/2$ is the golden ratio and $\omega = 1.84454757\dots$; an exact value for ω is not known.

We now define a recurrence for another useful statistic. We say that a dyadic tiling has a *left half-bisector* if the straight line segment from $(0, 1/2)$ to $(1/2, 1/2)$ does not intersect the interior of any dyadic rectangles. Figure 1(a) does not have a left half-bisector, while Figure 1(b) does. We are interested in the number of ways to tile the left half of a vertically bisected dyadic tiling of size 2^k such that it has a left half-bisector. Appealing to the dilation maps defined in the proof of Proposition 2.1, this number is A_{k-2}^2 . Among all possible ways to tile the left half of a vertically bisected tiling $\sigma \in \Omega_k$, we define f_k to be the fraction with a left half-bisector. We see

$$f_k = \frac{A_{k-2}^2}{A_{k-1}}$$

We can similarly define *right half-bisectors*, *top half-bisectors* and *bottom half-bisectors* by considering the straight line segments between $(1/2, 1/2)$ and, respectively, $(1, 1/2)$, $(1/2, 1)$ and $(1/2, 0)$. Then f_k is also the fraction of tilings of the right half of vertically bisected tiling σ with a right half-bisector, or the fraction of tilings of the top or bottom halves of a horizontally bisected tiling σ with a top or bottom half-bisector, respectively. Note that $f_2 = 0.5$, $f_3 = 4/7 \approx 0.571$ and $f_4 = 49/82 \approx 0.598$. We now examine the asymptotic behaviour of f_k .

Lemma 2.2. For all $k \geq 3$, $f_k = 1/(2 - f_{k-1}^2)$.

Proof. This follows from the recurrence for A_k given in Proposition 2.1:

$$f_k = \frac{A_{k-2}^2}{A_{k-1}} = \frac{A_{k-2}^2}{2A_{k-2}^2 - A_{k-3}^4} = \frac{1}{2 - A_{k-3}^4/A_{k-2}^2} = \frac{1}{2 - f_{k-1}^2}. \quad \square$$

We can use this recurrence to study the asymptotic behaviour of the sequence $\{f_k\}_{k=2}^\infty$.

Lemma 2.3. The sequence $\{f_k\}_{k=2}^\infty$ is strictly increasing and bounded above by $(\sqrt{5} - 1)/2$. Furthermore, $\lim_{k \rightarrow \infty} f_k = (\sqrt{5} - 1)/2$.

Proof. Note that $f_2 = 0.5 < (\sqrt{5} - 1)/2$. Suppose by induction that $f_{k-1} < (\sqrt{5} - 1)/2$. Then

$$f_k = \frac{1}{2 - f_{k-1}^2} < \frac{1}{2 - ((\sqrt{5} - 1)/2)^2} = \frac{4}{8 - (6 - 2\sqrt{5})} = \frac{4}{2 + 2\sqrt{5}} = \frac{2}{1 + \sqrt{5}} = \frac{\sqrt{5} - 1}{2}.$$

To show $f_k < f_{k+1}$ for all $k \geq 3$, it suffices to show $x < 1/(2 - x^2)$ for all $x \in [0.5, (\sqrt{5} - 1)/2]$. This is equivalent to showing that the polynomial $x^3 - 2x + 1$ is positive in that range. Factoring shows that this polynomial has roots at $1, (\sqrt{5} - 1)/2$ and $-(\sqrt{5} + 1)/2$, and is positive in the range $(-(\sqrt{5} + 1)/2, (\sqrt{5} - 1)/2)$. This implies $f_k < f_{k+1}$, so the sequence is strictly increasing.

The sequence $\{f_k\}_{k=2}^\infty$ is bounded and monotone, so it must converge to some limit β . To find β , we consider the function $g(x) = 1/(2 - x^2)$, which is the recurrence for the f_k . This function is continuous away from $\sqrt{2}$ and $-\sqrt{2}$, and thus certainly is continuous on $[0.5, (\sqrt{5} - 1)/2]$, the range of possible values for the f_k and their limit β . This continuity implies

$$g(\beta) = g(\lim_{k \rightarrow \infty} f_k) = \lim_{k \rightarrow \infty} g(f_k) = \lim_{k \rightarrow \infty} f_{k+1} = \beta.$$

Thus the limit β is necessarily a fixed point of $g(x)$. The fixed points of $g(x)$ are exactly the three roots of $x^3 - 2x + 1$ found above, and the only one in $[0.5, (\sqrt{5} - 1)/2]$ is $(\sqrt{5} - 1)/2$. We conclude $\lim_{k \rightarrow \infty} f_k = (\sqrt{5} - 1)/2$, as desired. □

2.2. Markov chains

We will consider only discrete-time Markov chains in this paper, though identical results hold for the analogous continuous-time Markov chains. Any finite ergodic Markov chain converges to a unique stationary distribution π . The time a Markov chain with transition matrix P takes to converge to its stationary distribution is measured by the *total variation distance*, which captures how far the distribution after t steps is from the stationary distribution given a worst-case starting configuration:

$$\|P^t - \pi\|_{TV} = \max_{x \in \Omega} \frac{1}{2} \sum_{y \in \Omega} |P^t(x, y) - \pi(y)|.$$

The mixing time of a Markov chain \mathcal{M} is defined to be

$$t_{\text{mix}}(\varepsilon) = \min\{t : \|P^t - \pi\|_{TV} \leq \varepsilon \ \forall t' \geq t\}.$$

For convenience, as is standard we define $t_{\text{mix}} = t_{\text{mix}}(1/4)$.

We will bound the mixing time of the edge-flip Markov chain for dyadic tilings by studying its relaxation time and spectral gap. The *spectral gap* γ of a Markov chain \mathcal{M} with transition matrix P is $1 - \lambda_2$, where λ_2 is the second largest eigenvalue of P . A *lazy* Markov chain is one where $P(x, x) \geq 1/2$ for all $x \in \Omega$; for a lazy Markov chain \mathcal{M} , the relaxation time, denoted by t_{rel} , is then the inverse of this spectral gap. We will see in the next section that the edge-flip Markov chain for dyadic tilings is lazy. The following

well-known proposition relates the relaxation time and mixing time for Markov chains; for a proof, see e.g. [17, Theorems 12.3 and 12.4].

Proposition 2.4. *Let \mathcal{M} be an ergodic Markov chain on state space Ω with reversible transition matrix P and stationary distribution π . Let $\pi_{\min} = \min_{x \in \Omega} \pi(x)$. Then*

$$(t_{rel} - 1) \log\left(\frac{1}{2\varepsilon}\right) \leq t_{mix}(\varepsilon) \leq \log\left(\frac{1}{\varepsilon\pi_{\min}}\right) t_{rel}.$$

We will bound the spectral gap, and thus the relaxation and mixing times, of the edge-flip Markov chain for dyadic tilings by considering functions on the chain’s state space. For $f : \Omega \rightarrow \mathbb{R}$, the variance of f with respect to a distribution π on Ω can be expressed as

$$\text{var}_{\pi}(f) = \sum_{x \in \Omega} \pi(x)(f(x) - \mathbb{E}_{\pi}[f(x)])^2 = \frac{1}{2} \sum_{x,y \in \Omega} \pi(x)\pi(y)(f(x) - f(y))^2.$$

We will only be considering the variance with respect to the uniform distribution on Ω , so the subscript π will be omitted. For a given reversible transition matrix P on state space Ω with stationary distribution π , the Dirichlet form, also known as the local variance, associated with the pair (P, π) is, for any function $f : \Omega \rightarrow \mathbb{R}$,

$$\mathcal{E}(f) = \frac{1}{2} \sum_{x,y \in \Omega} [f(x) - f(y)]^2 \pi(x)P(x,y).$$

As we see in the following well-known proposition, the Dirichlet form and variance of a function f can be used to bound the spectral gap of a transition matrix, and therefore the relaxation time and mixing time of a Markov chain; see e.g. [17, Lemma 13.12].

Proposition 2.5. *Given a Markov chain with reversible transition matrix P and stationary distribution π , the spectral gap $\gamma = 1 - \lambda_2$ of P satisfies*

$$\gamma = \min_{\substack{f : \Omega \rightarrow \mathbb{R} \\ \text{var}_{\pi}(f) \neq 0}} \frac{\mathcal{E}(f)}{\text{var}_{\pi}(f)}.$$

3. The edge-flip Markov chain \mathcal{M}_k

Let $n = 2^k$. For $k \geq 1$, the edge-flip Markov chain \mathcal{M}_k on the state space Ω_k of all dyadic tilings of size 2^k is given by the following rules.

Beginning at any $\sigma_0 \in \Omega_k$, repeat:

- Choose a rectangle R of σ_i uniformly at random.
- Choose *left, right, top* or *bottom* uniformly at random; let e be the corresponding side of R .
- If e bisects a rectangle of area 2^{-k+1} , remove e and replace it with its perpendicular bisector to obtain σ_{i+1} if the result is a valid dyadic tiling; else, set $\sigma_{i+1} = \sigma_i$.

An example of an edge-flip move of \mathcal{M}_k is shown in Figure 2(a); two selections of R and e that do not yield valid moves are shown in (b) and (c). Let $P_{k,\text{edge}}$ denote the transition

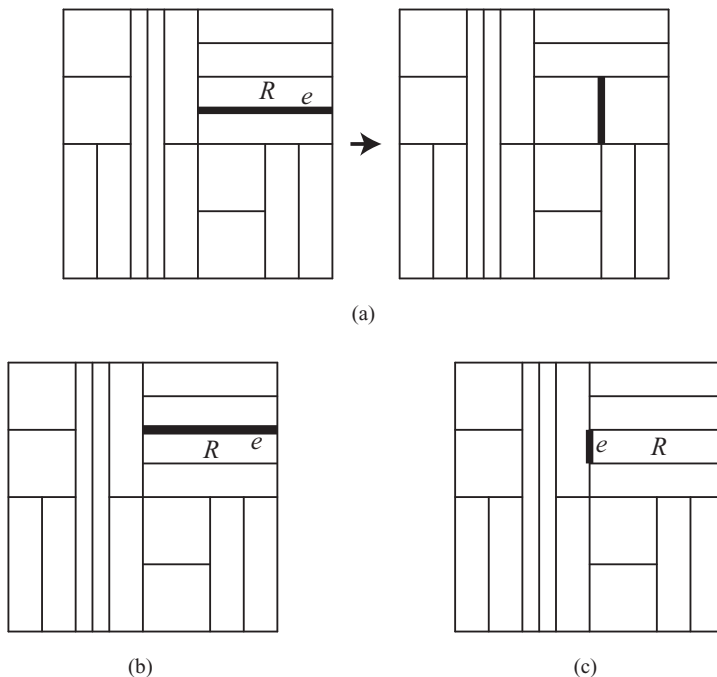


Figure 2. A random rectangle R and one of its edges e are selected in each iteration of \mathcal{M}_k . (a) Random choices of R and e as shown yield a valid edge flip. (b) Random choices of R and e as shown do not yield a valid edge flip as flipping edge e results in a tiling that is not dyadic. (c) Random choices of R and e as shown do not yield a valid edge flip as flipping edge e does not produce a tiling of the unit square by rectangles.

matrix of this edge-flip Markov chain and γ_k its spectral gap. For every valid edge flip, there are two choices of (R, e) that produce that move. This implies that every move between two tilings differing by an edge flip occurs with probability $1/(2n) = 2^{-k-1}$, so all off-diagonal entries of $P_{k, \text{edge}}$ are 2^{-k-1} or 0.

The Markov chain \mathcal{M}_k , in a slightly different form, was introduced by Janson, Randall and Spencer [14]. Note that \mathcal{M}_k is lazy, as for any rectangle R of a dyadic tiling at most one of its left and right edges can be flipped to produce another valid dyadic tiling. This is because if R 's projection onto the x -axis is dyadic interval $[a2^{-s}, (a + 1)2^{-s}]$ for $a, s \in \mathbb{Z}_{\geq 0}$, then flipping its left edge yields a rectangle with x -projection $[(a - 1)2^{-s}, (a + 1)2^{-s}]$ and flipping its right edge yields a rectangle with x -projection $[a2^{-s}, (a + 2)2^{-s}]$. If a is even, the first of these intervals is not dyadic, while if a is odd, the second is not, so at most one of these edge flips produces a valid dyadic tiling. Similarly, at most one of R 's top and bottom edges yields a valid edge flip. This implies in each iteration with probability at least $1/2$ a pair (R, e) is selected that does not yield a valid edge-flip move.

It was previously shown that this Markov chain is irreducible [14], so \mathcal{M}_k is ergodic and thus has a unique stationary distribution. The uniform distribution satisfies the detailed balance equation, implying both that \mathcal{M}_k is reversible and that its stationary distribution is uniform on Ω_k .

3.1. The block dynamics Markov chain $\mathcal{M}_k^{\text{block}}$

To analyse the mixing time of Markov chain \mathcal{M}_k , we will appeal to a similar Markov chain that uses larger block moves instead of single edge flips. We use in a crucial way the bijection between tilings in Ω_{k-1} and the left or right (resp. top or bottom) half of a tiling in Ω_k that has a vertical (resp. horizontal) bisector, as discussed in the proof of Proposition 2.1. For $k \geq 2$, the block dynamics Markov chain $\mathcal{M}_k^{\text{block}}$ on the state space Ω_k of all dyadic tilings of size 2^k is given by the following rules.

Beginning at any dyadic tiling σ_0 , repeat:

- Uniformly at random choose a tiling $\rho \in \Omega_{k-1}$.
- Uniformly at random choose *Left*, *Right*, *Top* or *Bottom*.
- To obtain σ_{i+1} :
 - If *Left* was chosen and σ_i has a vertical bisector, retiling σ_i 's left half with ρ , under the mapping $x \rightarrow x/2$.
 - If *Right* was chosen and σ_i has a vertical bisector, retiling σ_i 's right half with ρ , under the mapping $x \rightarrow (x + 1)/2$.
 - If *Bottom* was chosen and σ_i has a horizontal bisector, retiling σ_i 's bottom half with ρ , under the mapping $y \rightarrow y/2$.
 - If *Top* was chosen and σ_i has a horizontal bisector, retiling σ_i 's top half with ρ , under the mapping $y \rightarrow (y + 1)/2$.
- Else, set $\sigma_{i+1} = \sigma_i$.

Let $P_{k,\text{block}}$ be the transition matrix of this Markov chain and let $\gamma_{k,\text{block}}$ be its spectral gap. Any valid non-stationary transition of $\mathcal{M}_k^{\text{block}}$ occurs with probability $1/(4|\Omega_{k-1}|)$. This Markov chain is not lazy, but it is aperiodic, irreducible and reversible. This implies that it is ergodic and thus has a unique stationary distribution, which by detailed balance is uniform on Ω_k .

4. A polynomial upper bound on the mixing time of \mathcal{M}_k

Recall that we wish to show that the mixing time of \mathcal{M}_k is polynomial in $n = 2^k$, not polynomial in k . We show that the spectral gap γ_k of \mathcal{M}_k and the spectral gap γ_{k-1} of \mathcal{M}_{k-1} differ by a multiplicative constant (specifically, $1/17$) by appealing to the Dirichlet forms of both of these Markov chains as well as the block dynamics Markov chain $\mathcal{M}_k^{\text{block}}$. We can then use recursion to show that γ_k is bounded below by $(1/17)^k$, which, because $k = \log n$, gives a polynomial upper bound on the relaxation time and thus on the mixing time of \mathcal{M}_k .

For any function $f : \Omega_k \rightarrow \mathbb{R}$, we will denote the Dirichlet form of f with respect to transition matrix $P_{k,\text{edge}}$ and the uniform stationary distribution as $\mathcal{E}_{k,\text{edge}}(f)$. The Dirichlet form of f with respect to transition matrix $P_{k,\text{block}}$ and the uniform stationary distribution will be $\mathcal{E}_{k,\text{block}}(f)$. We will let the variance of function f on Ω_k with respect to the uniform stationary distribution be $\text{var}_k(f)$. Here the k indicates which state space Ω_k we are considering, rather than which distribution on Ω_k the variance is taken with respect to; all variances we consider will be with respect to the uniform distribution.

Because we consider two different Markov chains on the same state space Ω_k , there are two different notions of adjacencies on this state space, each corresponding to the moves of one of these Markov chains. For $x, y \in \Omega_k$, we say $x \sim_e y$ if x and y differ by a single edge-flip move of \mathcal{M}_k and $x \sim_b y$ if x and y differ by a single move of the block dynamics chain $\mathcal{M}_k^{\text{block}}$. More specifically, if x and y differ by a retiling of their left half (implying that x and y both have a vertical bisector and are the same on their right half), we say $x \sim_L y$; then $x \sim_R y$, $x \sim_T y$ and $x \sim_B y$ are defined similarly for the right, top and bottom halves.

Theorem 4.1. *For any $k \geq 2$, the spectral gap γ_k of the edge-flip Markov chain \mathcal{M}_k satisfies*

$$\gamma_k \geq \gamma_{k,\text{block}} \cdot \gamma_{k-1}.$$

Proof. We begin by relating the Dirichlet forms for block dynamics and for the edge-flip dynamics, which will allow comparison of their spectral gaps. Recall that for any function $f : \Omega_k \rightarrow \mathbb{R}$,

$$\mathcal{E}_{k,\text{block}}(f) = \frac{1}{2} \sum_{\substack{x,y \in \Omega_k \\ x \sim_b y}} \pi(x) P_{k,\text{block}}(x,y) (f(x) - f(y))^2.$$

This sum can be split into four terms, corresponding to the type of block move (left, right, top or bottom) transforming x into y . If x and y differ only in their top-left quadrants, then x could transition to y via either a left block move or a top block move; each of these moves occurs with probability $1/(4|\Omega_{k-1}|)$, and the total probability of $P_{k,\text{block}}(x,y) = 1/(2|\Omega_{k-1}|)$ will be split correspondingly between the terms for left block moves and top block moves.

We now analyse the first of these terms, containing all x, y differing by a retiling of their left halves. For $x_L, x_R \in \Omega_{k-1}$, by $x_L x_R$ below we mean the tiling in Ω_k with a vertical bisector whose left half is x_L under the map $x \rightarrow x/2$ and whose right half is x_R under the map $x \rightarrow (x + 1)/2$. Thus

$$\begin{aligned} \mathcal{E}_{k,\text{block}}^L &= \frac{1}{2} \sum_{\substack{x,y \in \Omega_k \\ x \sim_L y}} \frac{1}{|\Omega_k|} \frac{1}{4|\Omega_{k-1}|} (f(x) - f(y))^2 \\ &= \frac{1}{8} \sum_{x_R \in \Omega_{k-1}} \sum_{x_L, y_L \in \Omega_{k-1}} \frac{1}{|\Omega_k|} \frac{1}{|\Omega_{k-1}|} (f(x_L x_R) - f(y_L x_R))^2 \\ &= \frac{1}{4} \sum_{x_R \in \Omega_{k-1}} \frac{|\Omega_{k-1}|}{|\Omega_k|} \left(\frac{1}{2} \sum_{x_L, y_L \in \Omega_{k-1}} \frac{1}{|\Omega_{k-1}|^2} (f(x_L x_R) - f(y_L x_R))^2 \right). \end{aligned}$$

We note that the second sum above is over all pairs of tilings in Ω_{k-1} . While the Dirichlet form of a function sums over all pairs of states that differ by a transition of a Markov chain, the variance of a function sums over all pairs of states, regardless of the local structure imposed on the state space by the Markov chain. In fact, we have written the second sum above suggestively, and note that it is in fact a variance of a function over the state space Ω_{k-1} . For each $x_R \in \Omega_{k-1}$, the function $f|_{x_R} : \Omega_{k-1} \rightarrow \mathbb{R}$ given by

$f|_{x_R}(z) = f(zx_R)$ has variance $\text{var}_{k-1}(f|_{x_R})$ (with respect to the uniform distribution) that is exactly equal to the term in parentheses above. Because the variance of a function is the same regardless of which transitions on the state space we are considering, it is through this variance we can relate $\mathcal{E}_{k,\text{block}}$, which we have calculated above, to a Dirichlet form for edge-flip moves. That is, by Proposition 2.5, we can bound this variance with the Dirichlet form of $f|_{x_R}$ associated with $P_{k-1,\text{edge}}$ and the spectral gap γ_{k-1} of \mathcal{M}_{k-1} . Thus,

$$\mathcal{E}_{k,\text{block}}^L = \frac{1}{4} \sum_{x_R \in \Omega_{k-1}} \frac{|\Omega_{k-1}|}{|\Omega_k|} \text{var}_{k-1}(f|_{x_R}) \leq \frac{1}{4} \sum_{x_R \in \Omega_{k-1}} \frac{|\Omega_{k-1}|}{|\Omega_k|} \frac{\mathcal{E}_{k-1,\text{edge}}(f|_{x_R})}{\gamma_{k-1}}.$$

We now see that the Dirichlet form for the edge-flip Markov chain on Ω_{k-1} is

$$\begin{aligned} \mathcal{E}_{k-1,\text{edge}}(f|_{x_R}) &= \frac{1}{2} \sum_{\substack{x_L, y_L \in \Omega_{k-1} \\ x_L \sim_e y_L}} \pi(x_L) P_{k-1,\text{edge}}(x_L, y_L) (f(x_L x_R) - f(y_L x_R))^2 \\ &= \sum_{\substack{x_L, y_L \in \Omega_{k-1} \\ x_L \sim_e y_L}} \frac{1}{|\Omega_{k-1}|} \frac{1}{2n} (f(x_L x_R) - f(y_L x_R))^2. \end{aligned}$$

Using this expression, we see that

$$\begin{aligned} \mathcal{E}_{k,\text{block}}^L(f) &\leq \frac{1}{4\gamma_{k-1}} \sum_{x_R \in \Omega_{k-1}} \frac{|\Omega_{k-1}|}{|\Omega_k|} \left(\sum_{\substack{x_L, y_L \in \Omega_{k-1} \\ x_L \sim_e y_L}} \frac{1}{|\Omega_{k-1}|} \frac{1}{2n} (f(x_L x_R) - f(y_L x_R))^2 \right) \\ &= \frac{1}{4\gamma_{k-1}} \sum_{\substack{x, y \in \Omega_k \\ x \sim_e y \\ x \sim_L y}} \frac{1}{|\Omega_k|} \frac{1}{2n} (f(x) - f(y))^2. \end{aligned}$$

We now compare this to the Dirichlet form for the edge-flip Markov chain on Ω_k , which we recall is

$$\mathcal{E}_{k,\text{edge}}(f) = \frac{1}{2} \sum_{\substack{x, y \in \Omega_k \\ x \sim_e y}} \frac{1}{|\Omega_k|} \frac{1}{2n} (f(x) - f(y))^2.$$

We note that for every $x, y \in \Omega_k$ such that $x \sim_e y$, at least one of and at most two of $x \sim_L y$, $x \sim_R y$, $x \sim_T y$ and $x \sim_B y$ hold. Thus each summand of $\mathcal{E}_{k,\text{edge}}(f)$ appears at most twice as a summand of

$$\mathcal{E}_{k,\text{block}}(f) = \mathcal{E}_{k,\text{block}}^L(f) + \mathcal{E}_{k,\text{block}}^R(f) + \mathcal{E}_{k,\text{block}}^T(f) + \mathcal{E}_{k,\text{block}}^B(f).$$

It follows that

$$\mathcal{E}_{k,\text{block}}(f) \leq \frac{1}{4\gamma_{k-1}} \cdot 2 \cdot (2\mathcal{E}_k(f)) = \frac{\mathcal{E}_{k,\text{edge}}(f)}{\gamma_{k-1}}.$$

Note that this implies that, for any f ,

$$\text{var}_k(f) \leq \frac{\mathcal{E}_{k,\text{block}}(f)}{\gamma_{k,\text{block}}} \leq \frac{\mathcal{E}_{k,\text{edge}}(f)}{\gamma_{k,\text{block}} \cdot \gamma_{k-1}}.$$

Let f be chosen to be the function achieving equality in $\text{var}_k(f) \leq \mathcal{E}_{k,\text{edge}}(f)/\gamma_k$. We conclude that

$$\gamma_k = \frac{\mathcal{E}_{k,\text{edge}}(f)}{\text{var}_k(f)} \geq \gamma_{k,\text{block}} \cdot \gamma_{k-1}. \quad \square$$

In Section 5 we prove that $\gamma_{k,\text{block}}$ is at least $1/17$ for sufficiently large k . This can be used to bound the spectral gap, the relaxation time, and finally the mixing time of \mathcal{M}_k .

Theorem 4.2. *There exists a positive integer k_0 such that for all $k \geq k_0$, $\gamma_{k,\text{block}} \geq 1/17$.*

Proof. See Section 5. We introduce a distance metric on dyadic tilings, and then give a coupling where the distance between two tilings decreases in expectation after one iteration by a multiplicative factor of $1 - 1/17$ for all k sufficiently large. By a result of Chen [6] (see also [17, Theorem 13.1]), this implies the theorem. \square

We are now ready to prove our first main theorem, Theorem 1.1, which states that the relaxation time of \mathcal{M}_k for $n = 2^k$ is $O(n^{\log 17})$ and its mixing time is $O(n^{1+\log 17})$.

Proof of Theorem 1.1. By Theorems 4.1 and 4.2, the spectral gap of \mathcal{M}_k satisfies

$$\gamma_k \geq \frac{1}{17} \gamma_{k-1} \geq 17^{-(k-k_0)} \gamma_{k_0},$$

where k_0 is the value from Theorem 4.2. Since γ_{k_0} is a constant that does not depend on n ,

$$\gamma_k = \Omega(17^{-k}) = \Omega(n^{-\log 17}) = \Omega(n^{-4.09}).$$

Because \mathcal{M}_k is a lazy Markov chain, its relaxation time satisfies

$$t_{\text{rel}} = O(n^{\log 17}).$$

To use this to bound the mixing time of \mathcal{M}_k , we appeal to Proposition 2.4, though we must first calculate π_{\min} . For π the uniform distribution, $\min_{x \in \Omega_k} \pi(x) = 1/|\Omega_k|$. By Proposition 2.1, $|\Omega_k| < 2|\Omega_{k-1}|^2$, so a loose bound is $1/\pi_{\min} = |\Omega_k| < 2^{2^k} = 2^n$. This implies

$$t_{\text{mix}} = O(n^{1+\log 17}). \quad \square$$

5. The spectral gap of the block dynamics

We now present the proof of Theorem 4.2, which states that there exists a positive integer k_0 such that for all $k \geq k_0$, the spectral gap $\gamma_{k,\text{block}}$ is at least $1/17$.

Proof of Theorem 4.2. We start by defining the distance between two dyadic tilings $x, y \in \Omega_k$. In order to do this, we recall the notion of *half-bisectors*. We say that a tiling x has a *left half-bisector* if the line segment from $(0, 1/2)$ to $(1/2, 1/2)$ does not intersect the interior of any dyadic rectangle. In an analogous way we can define a *right half-bisector* using the line segment from $(1/2, 1/2)$ to $(1, 1/2)$, a *top half-bisector* using the

line segment from $(1/2, 1)$ to $(1/2, 1/2)$, and a *bottom half-bisector* using the line segment from $(1/2, 1/2)$ to $(1/2, 0)$. Note that if x has a horizontal bisector, then it has both a left half-bisector and a right half-bisector. However, x may have a left half-bisector but no horizontal bisector. For example, the dyadic tiling in Figure 1(a) has top, right and bottom half-bisectors, but no left half-bisector.

Now we define the distance between x and y as follows. For each of the four possible half-bisectors, let ℓ_1 be the number of such half-bisectors that are present in either x or y , but not in both of them. Also, for each of the four possible quadrants (top-left, top-right, bottom-left and bottom-right) of x and y , let ℓ_2 denote the number of such quadrants for which the rectangles in x intersecting that quadrant are not the same as the rectangles in y intersecting that quadrant. Then, introducing a parameter $b > 0$ that we will take to be sufficiently large later, we define the distance between x and y as

$$d(x, y) = b\ell_1 + \ell_2.$$

For instance, consider the two dyadic tilings in Figure 1(a,b). In this case we have $\ell_1 = 1$ due to the left half-bisector that is present in (b) but not in (a), and $\ell_2 = 3$ for top-left, top-right and bottom-left quadrants. The distance between these two tilings is then $b + 3$.

Our goal is to couple two instances of the block dynamics $\mathcal{M}_k^{\text{block}}$, one starting from a state $x \in \Omega_k$ and the other from a state $y \in \Omega_k$, such that the distance between x and y contracts after one step of the chains. More precisely, letting $\mathbb{E}_{x,y}$ denote the expectation with respect to the coupling, and if x' and y' are the dyadic tilings obtained after one step of each chain, respectively, we want to obtain a coupling and a value $\Delta > 0$ such that

$$\mathbb{E}_{x,y}[d(x', y')] \leq (1 - \Delta)d(x, y) \quad \text{for all } x, y \in \Omega_k. \tag{5.1}$$

Once we have the above inequality, then a result of Chen [6] (see also [17, Theorem 13.1]) implies that $\gamma_{k,\text{block}} \geq \Delta$.

We will use the following simple coupling between x' and y' :

- Uniformly at random choose a tiling $\rho \in \Omega_{k-1}$.
- Uniformly at random choose *Left*, *Right*, *Top* or *Bottom*.
- Retile the chosen half (left, right, top or bottom) of x with ρ , if possible.
- Retile the chosen half (left, right, top or bottom) of y with ρ , if possible.

For a more detailed description of the retiling step, see the definition of the transition rule of $\mathcal{M}_k^{\text{block}}$ in Section 3.1. When we update the left (resp. right) half of x and ρ contains a horizontal bisector, note that x' will contain a left (resp. right) half-bisector. Similarly, if we update the top (resp. bottom) half of x and ρ contains a vertical bisector, then x' will contain a top (resp. bottom) half-bisector. In any of these cases, we say that the retiling yields a half-bisector of x .

The remainder of the proof is devoted to showing that we can set b large enough so that (5.1) holds with $\Delta = 1/17$. In order to see this, we will split into three cases, and show that (5.1) holds with $\Delta = 1/17$ for each case.

Case 1: x and y have no common bisector. The maximum number of common half-bisectors of x and y in this case is two. Figure 3 illustrates the three possible configurations for the number of common half-bisectors of x and y . Consider first that x and y have no




Figure 3. Possible configurations for the half-bisectors of x and y in Case 1. In (c), $i \in \{0, 1\}$ denotes how many grey quadrants are tiled identically in x and y .

common half-bisector, which is illustrated in Figure 3(a) and has $d(x, y) = 4b + 4$. Then, whichever half (left, right, top or bottom) is chosen to be retiled, note that either x or y is actually retiled, but never both. With probability $|\Omega_{k-2}|^2/|\Omega_{k-1}| = f_k$ the retiling yields a half-bisector, which increases the number of common half-bisectors between x and y , and thus decreases their distance by b . Hence, using that $f_k \geq 1/2$, we have

$$\mathbb{E}_{x,y}[d(x', y')] = d(x, y) - f_k b \leq 4b + 4 - \frac{b}{2} < \left(1 - \frac{1}{17}\right)(4b + 4),$$

where the last step is true by setting b large enough (in this case, $b \geq 1$ suffices).

Now consider that x and y have one common half-bisector, and use Figure 3(b) as a reference, with x being the left tiling and y being the right tiling. We have $d(x, y) = 3b + 4$. If we retiling the left or right halves, so only x gets retiled, and the retiling yields a half-bisector, then the number of common half-bisectors of x and y decreases by 1. Similar behaviour happens if we retiling the top half. However, if we retiling the bottom half, and the retiling does not yield a half-bisector, then the number of common half-bisectors decreases by 1. Hence, using that $f_k \geq 1/2$, we obtain

$$\mathbb{E}_{x,y}[d(x', y')] \leq d(x, y) - \frac{3f_k b}{4} + \frac{(1 - f_k)b}{4} \leq 3b + 4 - \frac{b}{4} < \left(1 - \frac{1}{17}\right)(3b + 4),$$

where the last step is true by setting b large enough (in this case, $b \geq 4$ suffices).

Finally, suppose x and y have two common half-bisectors, as illustrated in Figure 3(c), where they may or may not be tiled the same in the quadrant bounded by these common half-bisectors. In this case $d(x, y) = 2b + 4 - i$, where $i = 1$ if they agree on this quadrant and $i = 0$ otherwise. Retiling the left and top halves can yield a new common half-bisector, while retiling the right and bottom halves may remove a common half-bisector. Moreover, if $i = 1$ and we retiling the right or bottom halves, the tilings of the bottom-right quadrant of x and of y may become different, increasing the distance between x and y by 1. Putting these together, we have

$$\begin{aligned} \mathbb{E}_{x,y}[d(x', y')] &\leq d(x, y) - \frac{2f_k b}{4} + \frac{2(1 - f_k)b}{4} + i\frac{2}{4} \\ &\leq 2b + 4 - \frac{i}{2} - \frac{(2f_k - 1)b}{2} = \frac{(5 - 2f_k)b}{2} + 4 - \frac{i}{2}. \end{aligned}$$

Since $f_k \rightarrow (\sqrt{5} - 1)/2$ as $k \rightarrow \infty$, the right-hand side above goes to

$$\left(\frac{6 - \sqrt{5}}{2}\right)b + 4 - \frac{i}{2}.$$




Figure 4. Possible configurations for the half-bisectors of x and y in Case 2. The value of $i \in \{0, 1, 2, 3\}$ denotes the number of grey quadrants which are tiled identically in x and y .

In particular, for $k \geq 10$, the coefficient of b above satisfies

$$\frac{5 - 2f_k}{2} < 2\left(1 - \frac{1}{17}\right),$$

and so we can set b large enough so that

$$\mathbb{E}_{x,y}[d(x', y')] \leq \left(1 - \frac{1}{17}\right)(2b + 4 - i).$$

We note that this is the tight case, as

$$\frac{6 - \sqrt{5}}{2} > 2\left(1 - \frac{1}{16}\right),$$

so this particular coupling and distance metric cannot be used to show that the spectral gap is at least $1/16$. This concludes the first case.

Case 2: x and y have a common bisector, but neither x nor y has both bisectors. Without loss of generality we assume x and y both have a vertical bisector and neither has a horizontal bisector. Each of x and y has at least two and at most three half-bisectors. Figure 4 illustrates the four possible configurations for the number of half-bisectors of x and y ; the shaded quadrants are those where x and y could have the same tiling. In all the situations of Figure 4, if we retiling the left or right halves, then we match up the configuration of x and y in that half. In particular, if x and y do not agree on the presence of left half-bisector, then they also do not have the same tiling of the top left or bottom left quadrants, so the decrease in distance due to a retiling of the left half, a move that occurs with probability $1/4$, is $(b + 2)$. If x and y agree on the presence of a left half-bisector and have the same tiling on $i' \in \{0, 1, 2\}$ of the two left quadrants, then the decrease in distance due to a retiling of the left half is $(2 - i')$. The same holds for right half-bisectors and retilings of the right half. As there are no moves of the coupling that can increase the distance between x and y , it can be shown that in all of the cases shown in Figure 4 the distance decreases by $1/4$ in expectation. Hence,

$$\mathbb{E}_{x,y}[d(x', y')] \leq d(x, y) - \frac{d(x, y)}{4} \leq \left(1 - \frac{1}{17}\right)d(x, y),$$

which concludes the second case.




Figure 5. Possible configurations for the half-bisectors of x and y in Case 3. The value of $i \in \{0, 1, 2, 3\}$ denotes the number of grey quadrants which are tiled identically in x and y .

Case 3: y has both vertical and horizontal bisectors. Here there are three situations, depending on whether x has two, three or four half-bisectors; see Figure 5. In the situation of Figure 5(a), if the left or right halves are retiled, then we match up x and y in that half, decreasing the distance by $b + 2$. But if we retiling the top or bottom halves, then we may increase the distance by b if the retiling does not yield a half-bisector. Hence,

$$\mathbb{E}_{x,y}[d(x', y')] \leq d(x, y) - \frac{2(b + 2)}{4} + \frac{2(1 - f_k)b}{4} = \frac{(4 - f_k)b}{2} + 3.$$

Since

$$\frac{4 - f_k}{2} \rightarrow \frac{9 - \sqrt{5}}{4} < \left(1 - \frac{1}{17}\right)2,$$

the right-hand side above is smaller than $(1 - 1/17)(2b + 4)$ when k and b are large enough. A similar situation occurs in Figure 5(b), but the distance increases a bit more when the top or bottom half is retiled as quadrants that were equal in x and y may become different. In this case, we have

$$\mathbb{E}_{x,y}[d(x', y')] \leq d(x, y) - \frac{(b + 4 - i)}{4} + \frac{2(1 - f_k)b}{4} + \frac{2}{4} = \frac{(5 - 2f_k)b}{4} + \frac{6 - i}{4}.$$

Since

$$\frac{5 - 2f_k}{4} \rightarrow \frac{6 - \sqrt{5}}{4} < \left(1 - \frac{1}{17}\right),$$

the right-hand side above is smaller than $(1 - 1/17)(b + 4 - i)$ when k and b are large enough; this is the second tight case, where we see contraction by a factor of $1 - 1/17$ but not by $1 - 1/16$. Finally, for the situation in Figure 5(c), regardless of which half we choose to retiling, the distance will not increase; if we choose a half containing a quadrant on which x and y differ, the distance will decrease. Each quadrant on which x and y differ is contained in two halves and thus is retiled so that x and y agree there with probability $1/2$. That is,

$$\mathbb{E}_{x,y}[d(x', y')] \leq d(x, y) - \frac{d(x, y)}{2} \leq \left(1 - \frac{1}{17}\right)d(x, y).$$

This concludes the third case. We have shown that for all possible tilings x and y , it holds that

$$\mathbb{E}_{x,y}[d(x', y')] \leq \left(1 - \frac{1}{17}\right)d(x, y).$$

This implies $\gamma_{k, \text{block}} \geq 1/17$ for all k sufficiently large, as desired. □

6. Lower bound on the mixing time of \mathcal{M}_n

In this section we give the proof of Theorem 1.2, which gives improved lower bounds on the mixing and relaxation times of \mathcal{M}_k . For this, we define the following subsets of Ω_k :

$$\begin{aligned} \Omega_k^+ &= \{x \in \Omega_k : x \text{ has both a horizontal and a vertical bisector}\}, \\ \Omega_k^| &= \{x \in \Omega_k : x \text{ has a vertical bisector}\}, \text{ and} \\ \Omega_k^- &= \{x \in \Omega_k : x \text{ has a horizontal bisector}\}. \end{aligned}$$

By definition, we have $\Omega_k^+ = \Omega_k^| \cap \Omega_k^-$. We start with the following simple lemma.

Lemma 6.1. *For all $k \geq 2$, we have*

$$\frac{|\Omega_k|}{|\Omega_k^+|} = \frac{2}{f_k^2} - 1 \geq 2\phi + 1.$$

Furthermore,

$$\lim_{k \rightarrow \infty} \frac{|\Omega_k|}{|\Omega_k^+|} = 2\phi + 1,$$

where $\phi = (\sqrt{5} + 1)/2$ is the golden ratio.

Proof. Using that $|\Omega_k^+| = |\Omega_{k-2}|^4$, and Proposition 2.1, we have

$$\frac{|\Omega_k|}{|\Omega_k^+|} = \frac{2|\Omega_{k-1}|^2 - |\Omega_{k-2}|^4}{|\Omega_{k-2}|^4} = \frac{2}{f_k^2} - 1.$$

By Lemma 2.3,

$$f_k \leq \frac{\sqrt{5} - 1}{2} = \frac{1}{\phi} = \lim_{k \rightarrow \infty} f_k.$$

This, along with the identity $\phi^2 = 1 + \phi$, implies the lemma. □

We will also require the following technical estimate.

Lemma 6.2. *For any $k \geq 2$, we have*

$$\frac{1}{|\Omega_k|} \prod_{i=0}^{k-2} |\Omega_i|^2 \leq \phi^{-2k+2}.$$

Proof. We will show how to estimate $\prod_{i=0}^{k-2} |\Omega_i|^2$ via the construction of a tiling in Ω_k . We start with a tiling with both a horizontal and a vertical bisector, as in Figure 6(a). Then we inductively do the following. Both quadrants of the left half are tiled independently with a uniformly random tiling from Ω_{k-2} . In the top-right quadrant, we add a vertical bisector and complete the two halves of this quadrant with independent, uniformly random tilings from Ω_{k-3} . Finally, in the bottom-right quadrant, we create a horizontal and a vertical bisector, reaching the tiling in Figure 6(b). Then we take this bottom-right quadrant, and




Figure 6. The construction of a tiling to count $\prod_{i=0}^{k-2} |\Omega_i|^2$. A rectangle with number a indicates that we tile it with a tiling from Ω_{k-a} , appropriately scaled.

iterate the procedure above; see Figure 6(c,d) for the configurations after one and two more iterations. This iteration continues until creating a bisector will result in rectangles of area less than 2^{-k} . In the case where an attempt is made to divide a rectangle of area 2^{-k+1} into four rectangles of equal area by adding both a horizontal and vertical bisector, we instead add just a horizontal bisector, resulting in two rectangles each of area 2^{-k} .

Let $\Upsilon_k \subset \Omega_k$ be the set of tilings obtained in this way. Note that the number of tilings in Υ_k is exactly $\prod_{i=0}^{k-2} |\Omega_i|^2$. Since $\Upsilon_k \subset \Omega_k^+$, we have that

$$\frac{|\Upsilon_k|}{|\Omega_k|} \leq \frac{|\Omega_k^+|}{|\Omega_k|},$$

where the first expression is exactly the value we wish to bound. Using the construction above until Figure 6(b), we obtain that

$$\frac{|\Upsilon_k|}{|\Omega_k|} \leq \frac{|\Omega_k^+| |\Omega_{k-2}^|}{|\Omega_k| |\Omega_{k-2}|},$$

where the second factor stands for the fact that the top-right quadrant must contain a vertical bisector. Iterating this in the bottom-right quadrant, we obtain

$$\frac{|\Upsilon_k|}{|\Omega_k|} \leq \frac{|\Omega_k^+| |\Omega_{k-2}^| |\Omega_{k-2}^+| |\Omega_{k-4}^|}{|\Omega_k| |\Omega_{k-2}| |\Omega_{k-2}| |\Omega_{k-4}|} \dots \tag{6.1}$$

Proposition 2.1 gives that

$$\frac{|\Omega_k^|}{|\Omega_k|} = \frac{|\Omega_k| + |\Omega_{k-2}|^4}{2|\Omega_k|} = \frac{1}{2} \left(1 + \frac{|\Omega_k^+|}{|\Omega_k|} \right) \leq \frac{1}{2} \left(1 + \frac{1}{2\phi + 1} \right) = \frac{\phi^2}{2\phi + 1},$$

where the inequality follows from Lemma 6.1. For even k , because $|\Omega_0^| = 0$ the last term we can obtain in (6.1) is $|\Omega_2^+|/|\Omega_2|$, so we can write

$$\begin{aligned} \frac{|\Upsilon_k|}{|\Omega_k|} &\leq \left(\prod_{i=0}^{k/2-2} \frac{|\Omega_{k-2i}^+|}{|\Omega_{k-2i}|} \cdot \frac{|\Omega_{k-2i-2}^|}{|\Omega_{k-2i-2}|} \right) \frac{|\Omega_2^+|}{|\Omega_2|} \\ &\leq \frac{1}{2\phi + 1} \left(\frac{1}{2\phi + 1} \cdot \frac{\phi^2}{2\phi + 1} \right)^{k/2-1} \end{aligned}$$

$$\begin{aligned} &= \frac{\phi^{-2k+4}}{2\phi + 1} \\ &\leq \phi^{-2k+2}, \end{aligned}$$

where the last expressions come from, respectively, identities for ϕ and the easily checked inequality $2\phi + 1 > \phi^2$. When k is odd, the last term in (6.1) is $|\Omega_1^-|/|\Omega_1|$ because $|\Omega_1^+| = 0$, so we can write

$$\frac{|Y_k|}{|\Omega_k|} \leq \left(\prod_{i=0}^{(k-3)/2} \frac{|\Omega_{k-2i}^+|}{|\Omega_{k-2i}|} \cdot \frac{|\Omega_{k-2i-2}^-|}{|\Omega_{k-2i-2}|} \right) \leq \left(\frac{1}{2\phi + 1} \cdot \frac{\phi^2}{2\phi + 1} \right)^{(k-1)/2} \leq \phi^{-2k+2},$$

where again the last expression is the result of applying identities for ϕ and simplifying. □

We are now ready to prove our second main theorem, giving a lower bound on the mixing and relaxation times of \mathcal{M}_k of $\Omega(n^{2 \log \phi})$, where $\phi = (\sqrt{5} + 1)/2$ is the golden ratio.

Proof of Theorem 1.2. We will derive an upper bound on the spectral gap γ_k . To do this, we consider the test function $f : \Omega_k \rightarrow \{0, 1\}$ such that

$$f(x) \text{ is } 1 \text{ if } x \in \Omega_k^|, \text{ and } 0 \text{ otherwise.} \tag{6.2}$$

We will apply this function to the characterization of the spectral gap in Proposition 2.5.

We start by showing that the variance of f is bounded away from 0 as $k \rightarrow \infty$. Recall that var_k denotes variance with respect to the uniform measure on Ω_k .

Claim 6.3. *With $f : \Omega_k \rightarrow \{0, 1\}$ as in (6.2), we have that*

$$\lim_{k \rightarrow \infty} \text{var}_k(f) = \sqrt{5} - 2.$$

Proof of claim. We start by writing

$$\text{var}_k(f) = \sum_{x \in \Omega_k^|} \sum_{y \in \Omega_k \setminus \Omega_k^|} \frac{1}{|\Omega_k|^2} = \frac{|\Omega_k^| \cdot |\Omega_k \setminus \Omega_k^|}{|\Omega_k|^2}. \tag{6.3}$$

Since $|\Omega_k^| = |\Omega_{k-1}|^2$, using Proposition 2.1 we obtain

$$|\Omega_k^| = \frac{|\Omega_k| + |\Omega_{k-2}|^4}{2} = \frac{|\Omega_k| + |\Omega_k^+|}{2}, \tag{6.4}$$

and

$$|\Omega_k \setminus \Omega_k^| = |\Omega_k| - |\Omega_k^| = \frac{|\Omega_k| - |\Omega_k^+|}{2}. \tag{6.5}$$

Plugging (6.4) and (6.5) into (6.3), we get

$$\text{var}_k(f) = \frac{1}{4} \left(1 + \frac{|\Omega_k^+|}{|\Omega_k|} \right) \left(1 - \frac{|\Omega_k^+|}{|\Omega_k|} \right) = \frac{1}{4} \left(1 - \left(\frac{|\Omega_k^+|}{|\Omega_k|} \right)^2 \right).$$




Figure 7. A tiling in $\partial\Omega_k^|$, with the dashed edge being the flip that brings the tiling into $\Omega_k^|$.

Then Lemma 6.1 yields

$$\lim_{k \rightarrow \infty} \text{var}_k(f) = \frac{1}{4} \left(1 - \frac{1}{(2\phi + 1)^2} \right).$$

Plugging in the value of ϕ completes the proof of the claim. □

Now it remains to obtain an upper bound for $\mathcal{E}(f)$. Let $\partial\Omega_k^|$ be the set of tilings in $\Omega_k \setminus \Omega_k^|$ that can be obtained from a tiling in $\Omega_k^|$ via one edge flip. Recall that for two tilings $x, y \in \Omega_k$, we write $x \sim_e y$ if x can be obtained from y by one edge flip. Hence,

$$\mathcal{E}(f) = \sum_{x \in \partial\Omega_k^|} \sum_{y \in \Omega_k^| : y \sim_e x} \frac{1}{|\Omega_k^|} \frac{1}{2n}.$$

Note that each tiling in $\partial\Omega_k^|$ has a horizontal bisector and is not in Ω_k^+ . This means that it has exactly one edge flip that can bring it into $\Omega_k^|$, which is the flip that creates a vertical bisector. Then, we have

$$\mathcal{E}(f) = \frac{|\partial\Omega_k^|}{2n \cdot |\Omega_k^|}.$$

Now we need to describe the set $\partial\Omega_k^|$. It is a set of tilings with no vertical bisector, but with one edge flip that creates a vertical bisector; see Figure 7. Note that the edge whose flip creates a vertical bisector must be a horizontal edge of length 1 which flips to a vertical edge of length $2/n$. From now on we will refer to this edge as the *pivotal edge*.

In order to estimate the cardinality of $\partial\Omega_k^|$, we will describe a procedure to construct a tiling $x \in \partial\Omega_k^|$, observing the position of the pivotal edge. Note that x must have a horizontal bisector, which splits $[0, 1]^2$ into its top and bottom halves. Assume that the pivotal edge is in the top half of x . This implies that the bottom half of x must itself contain a vertical bisector since the pivotal edge must be the only edge that forbids a vertical bisector to exist; see Figure 8(a). The two quadrants in the bottom half are simply any tilings of Ω_{k-2} . Note also that the top half of x must contain a horizontal bisector, otherwise $x \notin \partial\Omega_k^|$; see Figure 8(b). Then we iterate the above construction: among the

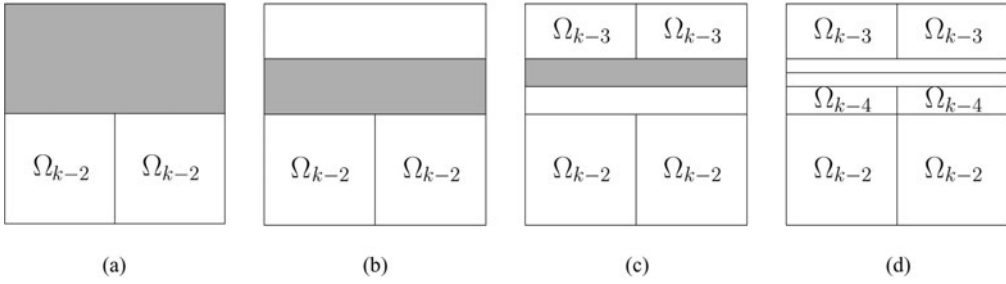


Figure 8. The construction of a tiling in $\partial\Omega_k^|$. The grey areas represent the part that contains the pivotal edge.

two halves of the top half, one must contain the pivotal edge, say the bottom one, while the other contains a vertical bisector, each side of which is completed with a tiling from Ω_{k-3} , which gives the configuration in Figure 8(c). Continuing this for $k - 2$ steps concludes the construction.

To estimate the cardinality of $\partial\Omega_k^|$, note that in each step of the construction we have two choices for where the pivotal edge is: either in the top half or the bottom half of the corresponding region. Therefore, the number of tilings in $\partial\Omega_k^|$ is

$$|\partial\Omega_k^| = \prod_{i=2}^k (2|\Omega_{k-i}|^2) = 2^{k-1} \prod_{i=0}^{k-2} |\Omega_i|^2 = \frac{n}{2} \prod_{i=0}^{k-2} |\Omega_i|^2.$$

Hence,

$$\mathcal{E}(f) = \frac{1}{4|\Omega_k|} \prod_{i=0}^{k-2} |\Omega_i|^2 \leq \frac{1}{4} \phi^{-2k+2}$$

where the last step follows from Lemma 6.2. Therefore, there exists a constant $c > 0$ such that

$$\gamma_k \leq c\phi^{-2k}.$$

This implies that the relaxation time and mixing time satisfy

$$t_{\text{rel}}, t_{\text{mix}} \geq \frac{1}{c} \phi^{2k} = \frac{1}{c} \phi^{2 \log n} = \frac{1}{c} n^{2 \log \phi} = \Omega(n^{2 \log \phi}).$$

This completes the proof of the theorem. □

Acknowledgements

This work started during the 2016 AIM workshop ‘Markov chain mixing times’. We thank the organizers for the invitation and the stimulating atmosphere.

References

[1] Burr, M., Choi, S. W., Galehouse, B. and Yap, C. K. (2012) Complete subdivision algorithms, II: Isotopic meshing of singular algebraic curves. *J. Symbol. Comput.* **47** 131–152.

- [2] Cannon, S., Miracle, S. and Randall, D. (2018) Phase transitions in random dyadic tilings and rectangular dissections. *SIAM J. Discrete Math.* **32** 1966–1992.
- [3] Caputo, P., Martinelli, F., Sinclair, A. and Stauffer, A. (2015) Random lattice triangulations: Structure and algorithms. *Ann. Appl. Probab.* **25** 1650–1685.
- [4] Caputo, P., Martinelli, F., Sinclair, A. and Stauffer, A. (2016) Dynamics of lattice triangulations on thin rectangles. *Electron. J. Probab.* **21** 29.
- [5] Cesi, F. (2001) Quasi-factorization of the entropy and logarithmic Sobolev inequalities for Gibbs random fields. *Probab. Theory Rel. Fields* **120** 569–584.
- [6] Chen, M.-F. (1998) Trilogy of couplings and general formulas for lower bound of spectral gap. In *Probability Towards 2000* (L. Accardi and C. C. Heyde, eds), Vol. 128 of Lecture Notes in Statistics, Springer, pp. 123–136.
- [7] Cooper, C., Dyer, M. and Greenhill, C. (2007) Sampling regular graphs and a peer-to-peer network. *Combin. Probab. Comput.* **16** 557–593.
- [8] Cuff, P., Ding, J., Louidor, O., Lubetzky, E., Peres, Y. and Sly, A. (2012) Glauber dynamics for the mean-field Potts model. *J. Statist. Phys.* **149** 432–477.
- [9] Diaconis, P. and Saloff-Coste, L. (1993) Comparison theorems for reversible Markov chains. *Ann. Appl. Probab.* **3** 696–730.
- [10] Ding, J., Lubetzky, E. and Peres, Y. (2009) The mixing time evolution of Glauber dynamics for the mean-field Ising model. *Comm. Math. Phys.* **289** 725–764.
- [11] Ding, J., Lubetzky, E. and Peres, Y. (2010) Mixing time of critical Ising model on trees is polynomial in the height. *Comm. Math. Phys.* **295** 161–207.
- [12] Gheissari, R. and Lubetzky, E. (2017) Mixing times of critical two-dimensional Potts models. *Comm. Pure Appl. Math.* **71** 994–1046.
- [13] Greenhill, C. and Sfragara, M. (2018) The switch Markov chain for sampling irregular graphs and digraphs. *Theoret. Comput. Sci.* **719** 1–20.
- [14] Janson, S., Randall, D. and Spencer, J. (2002) Random dyadic tilings of the unit square. *Random Struct. Alg.* **21** 225–251.
- [15] Lagarias, J. C., Spencer, J. H. and Vinson, J. P. (2002) Counting dyadic equipartitions of the unit square. *Discrete Math.* **257** 481–499.
- [16] Levin, D. A., Luczak, M. J. and Peres, Y. (2010) Glauber dynamics for the mean-field Ising model: Cut-off, critical power law, and metastability. *Probab. Theory Rel. Fields* **146** 223–265.
- [17] Levin, D. A., Peres, Y. and Wilmer, E. L. (2009) *Markov Chains and Mixing Times*, American Mathematical Society.
- [18] Lubetzky, E. and Sly, A. (2012) Critical Ising on the square lattice mixes in polynomial time. *Comm. Math. Phys.* **313** 815–836.
- [19] Luby, M., Randall, D. and Sinclair, A. (2001) Markov chain algorithms for planar lattice structures. *SIAM J. Comput.* **31** 167–192.
- [20] Martinelli, F. (1999) Lectures on Glauber dynamics for discrete spin models. In *Lectures on Probability Theory and Statistics: Ecole d'Eté de Probabilités de Saint-Flour XXVII, 1997* (P. Bernard, ed.), Springer, pp. 93–191.
- [21] McShine, L. and Tetali, P. (1998) On the mixing time of the triangulation walk and other Catalan structures. In *Randomization Methods in Algorithm Design*, Vol. 43 of DIMACS Series in Discrete Mathematics and Theoretical Computer Science, pp. 147–160.
- [22] Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H. and Teller, E. (1953) Equation of state calculations by fast computing machines. *J. Chem. Phys.* **21** 1087–1092.
- [23] Randall, D. and Tetali, P. (2000) Analyzing Glauber dynamics by comparison of Markov chains. *J. Math. Phys.* **41** 1598–1615.
- [24] Scott, C. and Nowak, R. D. (2006) Minimax-optimal classification with dyadic decision trees. *IEEE Trans. Inform. Theory* **52** 1335–1353
- [25] Stauffer, A. (2017) A Lyapunov function for Glauber dynamics on lattice triangulations. *Probab. Theory Rel. Fields* **169** 469–521.