

STRONG SPATIAL MIXING AND RAPID MIXING WITH FIVE COLOURS FOR THE KAGOME LATTICE

MARKUS JALSENIUS

Abstract

We consider proper 5-colourings of the kagome lattice. Proper q -colourings correspond to configurations in the zero-temperature q -state anti-ferromagnetic Potts model. Salas and Sokal have given a computer assisted proof of strong spatial mixing on the kagome lattice for $q \geq 6$ under any temperature, including zero temperature. It is believed that there is strong spatial mixing for $q \geq 4$. Here we give a computer assisted proof of strong spatial mixing for $q = 5$ and zero temperature. It is commonly known that strong spatial mixing implies that there is a unique infinite-volume Gibbs measure and that the Glauber dynamics is rapidly mixing. We give a proof of rapid mixing of the Glauber dynamics on any finite subset of the vertices of the kagome lattice, provided that the boundary is free (not coloured). The Glauber dynamics is not necessarily irreducible if the boundary is chosen arbitrarily for $q = 5$ colours. The Glauber dynamics can be used to uniformly sample proper 5-colourings. Thus, a consequence of rapidly mixing Glauber dynamics is that there is fully polynomial randomised approximation scheme for counting the number of proper 5-colourings.

1. *Introduction*

Proper colourings correspond to configurations in the *zero-temperature anti-ferromagnetic Potts model*. In this paper we will show that the system specified by proper 5-colourings of the *kagome lattice* has *strong spatial mixing*, and that the *Glauber dynamics* is *rapidly mixing*. The previously best known result [20] on mixing on the kagome lattice was for 6 colours. It is believed [20] that there is strong spatial mixing for 4 or more colours, and hence our result is narrowing the gap between what is believed and known. In Section 1.1 below we give an introduction to mixing, and in Section 2 we state our results and discuss related work.

1.1. *Definitions and background*

The kagome lattice, Figure 1(a), is a natural lattice of interest in statistical physics [20]. Instead of drawing graphs in the traditional way, with a vertex denoted with a solid circle and an edge denoted with a line segment, we draw graphs such that faces represent vertices. Two adjacent faces therefore represent two adjacent

Some of the work has been done at the Department of Computer Science, University of Warwick, Coventry, CV4 7AL, United Kingdom.

Received 21 September 2006, revised 30 October 2008; *published* 20 November 2009.

2000 Mathematics Subject Classification 05C15.

© 2009, Markus Jalsenius




Figure 1: (a) The kagome lattice, here drawn in a coordinate system. We illustrate graphs such that a face represents a vertex. (b) A vertex $(x, y) \in V_{\text{odd}}$ and its right neighbour. (c) A vertex $(x, y) \in V_{\text{even}}$ and its four neighbours.

vertices. Let $\mathcal{G} = (V_G, E_G)$ denote the kagome lattice with vertex set V_G and edge set E_G . We have $V_G = V_{\text{odd}} \cup V_{\text{even}}$, where

$$V_{\text{odd}} = \{(x, y) \mid x, y \in \mathbb{Z} \text{ are both odd}\},$$

$$V_{\text{even}} = \{(x, y) \mid x = 4k_1 + r, y = 4k_2 + r, k_1, k_2 \in \mathbb{Z}, r \in \{0, 2\}\}.$$

The edge set

$$E_G = \{((x, y), (x + 2, y)) \mid (x, y) \in V_{\text{odd}}\} \cup \{((x, y), (x + k_1, y + k_2)) \mid (x, y) \in V_{\text{even}}, k_1 \in \{-1, 1\}, k_2 \in \{-1, 1\}\}.$$

Note that both vertices in Figure 1(b) are in V_{odd} , and we see that two adjacent vertices in V_{odd} differ by 2 in their x -coordinate. In Figure 1(c), the centre vertex is in V_{even} , and we see that its four neighbours are in V_{odd} .

A region $R \subseteq V_G$ is a *finite non-empty* subset of the vertex set of the kagome lattice. The subset $\partial R \subseteq V_G$ denotes the *vertex boundary* of R such that ∂R is the set of vertices that are not in R but are adjacent to any vertex in R . The edge set $E(R)$ is the set of all edges $(u, v) \in E_G$ such that at least one of the vertices u and v is in R . The *edge boundary* $\mathcal{E}R$ of R is the set of all edges $(u, v) \in E(R)$ such that exactly one of the vertices u and v is in R and the other one is in ∂R .

The set $Q = \{1, \dots, q\}$ denotes the set of q colours, and the set $Q_0 = \{0\} \cup Q$. The colour 0 represents “no colour”. A q -colouring of a region R is a function from R to the set Q , and a q_0 -colouring of R is a function from R to Q_0 . A 0-colouring of R is a function from R to the set $\{0\}$, which means that all vertices in R are assigned colour 0. We often write only colouring when it is obvious from the context if it is a q -, q_0 - or 0-colouring, or if any colouring will do. Let σ be a colouring of a region R . If R' is a subset of R then $\sigma(R')$ is the colouring of R' induced by σ . Furthermore, for a vertex $v \in R$, $\sigma(v)$ is the colour of v under σ . Let Ω_R^+ denote the set of all q -colourings of the region R . For two colourings $\sigma, \sigma' \in \Omega_R^+$, the *Hamming distance* between σ and σ' is the number of vertices in R on which σ and σ' differ. A colouring σ of R is *proper* if no adjacent vertices receive the same colour. That is, $\sigma(u) \neq \sigma(v)$ for all adjacent vertices u and v in R . Let Ω_R denote the set of all proper q -colourings of the region R . Given a q_0 -colouring \mathcal{B} of ∂R , a proper

q -colouring σ of R agrees with \mathcal{B} if $\sigma(u) \neq \mathcal{B}(v)$ for all $(u, v) \in \mathcal{E}R$, where $u \in R$. We let $\Omega_R(\mathcal{B})$ denote the set of all proper q -colourings of R that agree with \mathcal{B} . The uniform distribution on $\Omega_R(\mathcal{B})$ is denoted $\pi_{\mathcal{B}}$, and for any subregion $R' \subseteq R$, let $\pi_{\mathcal{B}, R'}$ denote the distribution on proper q -colourings of R' induced by $\pi_{\mathcal{B}}$.

In this paper we will show that the system specified by proper 5-colourings of the kagome lattice has *strong spatial mixing*. Informally, strong spatial mixing means that if R is a region and \mathcal{B} is a q_0 -colouring of ∂R , then the effect the colour of a vertex $w \in \partial R$ has on a vertex $v \in R$ decays exponentially with the distance between w and v . The effect is measured with the *total variation distance*. For two distributions D_1 and D_2 on a set S , the total variation distance between D_1 and D_2 is defined as

$$d_{\text{TV}}(D_1, D_2) = \frac{1}{2} \sum_{s \in S} |D_1(s) - D_2(s)| = \max_{A \subseteq S} |D_1(A) - D_2(A)|.$$

The following definition of strong spatial mixing is taken from [12] and is adapted to the kagome lattice.

DEFINITION 1 (Strong spatial mixing). The system specified by proper q -colourings of the kagome lattice has *strong spatial mixing* if there are two constants $\alpha > 0$ and $\varepsilon \in (0, 1)$ such that, for any region R , any subregion $R' \subseteq R$, any two q_0 -colourings \mathcal{B} and \mathcal{B}' of ∂R which differ on exactly one vertex $w \in \partial R$ and such that $\mathcal{B}(w) \neq 0$ and $\mathcal{B}'(w) \neq 0$,

$$d_{\text{TV}}(\pi_{\mathcal{B}, R'}, \pi_{\mathcal{B}', R'}) \leq \alpha |R'| (1 - \varepsilon)^{d(w, R')},$$

where $d(w, R')$ is the minimal distance within R from w to some vertex of R' .

A distribution π on the set of proper q -colourings of the infinite kagome lattice is an *infinite-volume Gibbs distribution* if, for any region R and any proper q -colouring σ of the kagome lattice, the conditional distribution $\pi(\cdot | \sigma(V_G \setminus R))$ on Ω_R (conditioned on the colouring $\sigma(V_G \setminus R)$ of all vertices other than those in R) is $\pi_{\mathcal{B}}$, where $\mathcal{B} = \sigma(\partial R)$. It is known that there is always at least one infinite-volume Gibbs distribution, and the question of interest is to determine whether it is unique or not. This question is central in statistical physics because it corresponds to the number of macroscopic equilibria for a given system. The phenomenon of non-uniqueness corresponds to what is referred to as a *phase transition*. A consequence of strong spatial mixing is that the infinite-volume Gibbs distribution is unique [8, 23, 24]. For more on Gibbs distributions, see for example [9] or [10].

Another question of interest is to determine how quickly the system converges to equilibrium. The answer to this question is connected to the quantities α and ε in Definition 1 above. From a statistical physics point of view, this question is important for understanding phenomena such as how the system returns to equilibrium after a shock forces it out of it. In this paper we consider a famous dynamical process called the *Glauber dynamics* which models how the system converges. The Glauber dynamics, defined next, is a Markov chain that performs *single-vertex heat-bath* updates.

DEFINITION 2 (Glauber dynamics). For any region R and any q_0 -colouring \mathcal{B} of ∂R , the *Glauber dynamics* is a Markov chain with state space $\Omega_R(\mathcal{B})$, and a transition is made from a state σ to σ' in the following way:




Figure 2: A 2-vertex region of the kagome lattice with different colourings. A vertex is labelled with its colour. The two colourings in (a) are “frozen” and do not communicate in the Glauber dynamics for $q = 5$. However, when restricting the boundary to the 0-colouring in (b), the two colourings do communicate.

1. Choose a vertex v uniformly at random from R .
2. Let Q_v be the set of colours which are assigned to the neighbours of v (either in σ or \mathcal{B}).
3. Choose a colour c uniformly at random from $Q \setminus Q_v$ and obtain the new colouring σ' from σ by assigning colour c to vertex v .

A sufficient condition for the Glauber dynamics to be *connected* (that is, any proper colouring can be obtained from another proper colouring by a series of transitions) is to have $q \geq 6$. In general, with the Glauber dynamics defined similarly on any underlying infinite graph of maximum degree Δ , having $q \geq \Delta + 2$ is a sufficient condition for the dynamics to be connected. In this paper we focus on 5-colourings and in order to guarantee that the Glauber dynamics is connected we will have to restrict the colourings \mathcal{B} of the boundary to the 0-colouring (see Figure 2). For this reason, our 5-colour mixing result for the Glauber dynamics is restricted to the 0-colouring of the boundary. It is worth pointing out that if we add moves to the Glauber dynamics that allow swapping the colours of two neighbouring vertices (when this move is allowed with respect to the colouring of the rest of the vertices) then the new dynamics is connected for any region R and any q_0 -colouring of ∂R if $q \geq 5$. This fact is true for any graph of maximum degree Δ and $q \geq \Delta + 1$ colours. This augmented Glauber dynamics can be simulated by the *heat-bath dynamics on edges* which we define as follows: Choose an edge $e = (v_1, v_2)$ uniformly at random and simultaneously recolour v_1 and v_2 uniformly at random from the allowed colourings.

If the Glauber dynamics is connected, and hence ergodic, then $\pi_{\mathcal{B}}$ is the unique stationary distribution. This follows from the fact that the Glauber dynamics is reversible with respect to $\pi_{\mathcal{B}}$. For the same reason, $\pi_{\mathcal{B}}$ is the unique stationary distribution of the heat-bath dynamics on edges. The Glauber dynamics can be used as a sampler to sample colourings from the uniform distribution on $\Omega_R(\mathcal{B})$. This can be done efficiently if the Glauber dynamics is *rapidly mixing* (see definition below), which means that it quickly reaches its stationary distribution.

DEFINITION 3 (Mixing time). Consider the Glauber dynamics on a region R with boundary colouring \mathcal{B} . Let $P^t(\sigma, \sigma')$ be the probability of going from state σ to σ' in exactly t steps. For any $\delta > 0$, the *mixing time*

$$\tau(\delta) = \max_{\sigma \in \Omega_R(\mathcal{B})} \min\{t_0 : d_{TV}(P^t(\sigma, \cdot), \pi_{\mathcal{B}}) \leq \delta \text{ for all } t \geq t_0\}.$$

The Glauber dynamics is *rapidly mixing* if $\tau(\delta)$ is upper-bounded by a polynomial in the region-size $|R|$ and $\log(1/\delta)$.

It is a well-known fact that if the system has strong spatial mixing then the Glauber dynamics is (often) rapidly mixing [8, 17, 23]. In Section 7 we will study this fact and see how strong spatial mixing and rapid mixing are closely related. For $q \geq 6$ colours (or $q \geq \Delta + 2$ in general) it is straightforward to apply Theorem 8 in [12] in order to infer rapid mixing from strong spatial mixing. However, with $q = 5$ colours we cannot rely entirely on previous results. We will establish certain properties of 5-colourings of the kagome lattice and show that the Glauber dynamics is rapidly mixing for $q = 5$ under the 0-colouring of the boundary.

In [15] it is explained how *approximate counting* and almost uniform sampling are related. If there is a method for sampling (almost) uniformly at random in polynomial time from the set of proper colouring of a finite region R , then we can construct a *fully polynomial randomised approximation scheme*, or *FPRAS*, for counting the number of proper colourings of R . Thus, if the Glauber dynamics is rapidly mixing then we could use it to construct (in a non-trivial way) an FPRAS for estimating $|\Omega_R|$. For details on the topic of how sampling and counting are related, see Jerrum [15] and Jerrum, Valiant and Vazirani [16].

2. The results and related work

We will prove the following theorems, which improve previously known results on mixing for proper colourings of the kagome lattice.

THEOREM 4. *The system specified by proper 5-colourings of the kagome lattice has strong spatial mixing.*

THEOREM 5. *For any region R of the kagome lattice and $q = 5$ colours, the Glauber dynamics is rapidly mixing on R under the 0-colouring of ∂R . The mixing time $\tau(\delta) \in O(n^2 + n \log \frac{1}{\delta})$, where n is the number of vertices in R .*

THEOREM 6. *For any region R of the kagome lattice and $q = 5$ colours, the heat-bath dynamics on edges is rapidly mixing on R under any q_0 -colouring of ∂R . The mixing time $\tau(\delta) \in O(n^2 + n \log \frac{1}{\delta})$, where n is the number of vertices in R .*

The previously best known result on mixing on the kagome lattice is that of Salas and Sokal [20]. They provided a computer assisted proof of strong spatial mixing for $q = 6$ colours. It is believed [20] that there is strong spatial mixing for $q \geq 4$ colours.

It is worth mentioning some previous general results on mixing. Independently, Jerrum [14] and Salas and Sokal [20] proved that for proper q -colourings on a graph of maximum degree Δ the Glauber dynamics has $O(n \log n)$ -mixing for $q > 2\Delta$, where n is the number of vertices of the region. For $q = 2\Delta$, Bubley and Dyer [3] showed that it mixes in $O(n^3)$ time and Molloy [18] showed that it mixes $O(n \log n)$ time. Vigoda [22] used a Markov chain that differs from the Glauber dynamics and showed that it has $O(n \log n)$ -mixing for $q > (11/6)\Delta$. This result implied that the Glauber dynamics is rapidly mixing for $q > (11/6)\Delta$. Goldberg, Martin and Paterson [12] showed that any triangle-free graph has strong spatial mixing provided $q > \alpha\Delta - \gamma$, where α is the solution to $\alpha^\alpha = e$ ($\alpha \approx 1.76322$) and

$\gamma = \frac{4\alpha^3 - 6\alpha^2 - 3\alpha + 4}{2(\alpha^2 - 1)} \approx 0.47031$. Note that their result cannot be applied to the kagome lattice since its edge set contains triangles. However, for other 4-regular graphs, such as the square lattice \mathbb{Z}^2 , it follows that mixing occurs for $q \geq 7$ colours. The technique Goldberg, Martin and Paterson used in [12] is well suited to be extended to involve special cases that depend on the particular graph under consideration. Involving such special cases can improve the mixing bounds. In order to deal with all special cases it might be helpful to incorporate computer assistance. This has been done in [12] for the lattice \mathbb{Z}^3 . The general result gives mixing for $q \geq 11$ colours but by taking advantage of the geometry of the lattice it has been shown that mixing occurs for $q \geq 10$. This proof is computer assisted. Another computer assisted proof of mixing in [12] is given for the triangular lattice and $q = 10$ colours. This result was improved by Jalsenius [13] to $q = 9$ by exploiting the geometry of the lattice even further. Goldberg, Jalsenius, Martin and Paterson used the technique from [12] and gave in [11] a computer assisted proof of mixing for $q = 6$ on the square lattice \mathbb{Z}^2 . This is an alternative proof of the result of Achlioptas, Molloy, Moore and van Bussel [1] (who also used computer assistance). In this paper we will refine the technique Goldberg, Martin and Paterson introduced in [12] to show mixing on the kagome lattice for $q = 5$ colours. Both the square lattice and the kagome lattice are 4-regular graphs, but the kagome lattice contains triangles whereas the square lattice does not. An interesting observation is that the presence of triangles seem to have a positive effect on the technique we use to show strong spatial mixing. Attempts to prove mixing with 5 colours on the square lattice with this technique has failed so far. The absence of triangles seem to be one strong reason why (assuming the square lattice does have strong spatial mixing with 5 colours).

3. The framework

When Goldberg, Martin and Paterson [12] derived improved mixing bounds for spin systems consisting of proper colourings, they introduced the notion of a *vertex-boundary pair*. A vertex-boundary pair is a data structure holding information about a region R and colourings of ∂R . The idea is to derive certain properties of the vertex-boundary pairs which can be easily translated into properties such as whether there is strong spatial mixing or not. When Goldberg, Martin and Paterson derived these properties, it turned out to be convenient to work with *edge-boundary pairs*. An edge-boundary pair (defined in the next section) contains colourings of the edge boundary $\mathcal{E}R$ rather than the vertex boundary ∂R .

DEFINITION 7 (Vertex-boundary pair). A *vertex-boundary pair* \mathcal{X} consists of

- a region $R_{\mathcal{X}}$,
- a distinguished boundary vertex $w_{\mathcal{X}} \in \partial R_{\mathcal{X}}$, and
- a pair $(\mathcal{B}_{\mathcal{X}}, \mathcal{B}'_{\mathcal{X}})$ of q_0 -colourings of $\partial R_{\mathcal{X}}$ that are identical on all vertices except on $w_{\mathcal{X}}$, where they differ. The colour of $w_{\mathcal{X}}$ is in Q for both $\mathcal{B}_{\mathcal{X}}$ and $\mathcal{B}'_{\mathcal{X}}$.

Note that the colour of the distinguished vertex $w_{\mathcal{X}}$ has to be in the set Q . That is, $\mathcal{B}_{\mathcal{X}}(w_{\mathcal{X}}) \neq 0$ and $\mathcal{B}'_{\mathcal{X}}(w_{\mathcal{X}}) \neq 0$. Definition 1 of strong spatial mixing can be rephrased using the definition of a vertex-boundary pair. That is, in order to show

strong spatial mixing, we will show that there are two constants $\alpha > 0$ and $\varepsilon \in (0, 1)$ such that for every vertex-boundary pair \mathcal{X} and every subregion $R' \subseteq R_{\mathcal{X}}$,

$$d_{\text{TV}}(\pi_{\mathcal{B}_{\mathcal{X},R'}}, \pi_{\mathcal{B}'_{\mathcal{X},R'}}) \leq \alpha |R'| (1 - \varepsilon)^{d(w_{\mathcal{X}}, R')}$$

One approach to show exponential decay of the total variation distance in the distance between $w_{\mathcal{X}}$ and R' is to construct a suitable *coupling* (defined next) of the distributions $\pi_{\mathcal{B}_{\mathcal{X}}}$ and $\pi_{\mathcal{B}'_{\mathcal{X}}}$. For two distributions D_1 and D_2 on a set S , a coupling Ψ of D_1 and D_2 is a joint distribution on $S \times S$ with marginal distributions D_1 and D_2 . If the pair (X_1, X_2) is a random variable drawn from Ψ then

$$d_{\text{TV}}(D_1, D_2) \leq \Pr[X_1 \neq X_2].$$

Thus, in order to upper-bound the total variation distance, one can find some suitable coupling Ψ and compute the probability of having $X_1 \neq X_2$. The aim here is to construct a coupling $\Psi_{\mathcal{X}}$ of $\pi_{\mathcal{B}_{\mathcal{X}}}$ and $\pi_{\mathcal{B}'_{\mathcal{X}}}$ such that if the pair (σ, σ') of colourings is drawn from $\Psi_{\mathcal{X}}$ then the probability that σ and σ' differ on $R' \subseteq R_{\mathcal{X}}$ decreases exponentially with the distance between the discrepancy vertex $w_{\mathcal{X}} \in \partial R_{\mathcal{X}}$ and R' . For a vertex $v \in R_{\mathcal{X}}$ we define the indicator random variable $1_{\Psi_{\mathcal{X}},v}$ for the event that the colour of v differs in a pair of colourings drawn from $\Psi_{\mathcal{X}}$. Hence, the quantity $\sum_{v \in R_{\mathcal{X}}} \mathbb{E}[1_{\Psi_{\mathcal{X}},v}]$ is the expected number of vertices in $R_{\mathcal{X}}$ on which the colours differ in a pair of colourings drawn from $\Psi_{\mathcal{X}}$. If $\mathbb{E}[1_{\Psi_{\mathcal{X}},v}]$ is small enough for all vertex-boundary pairs \mathcal{X} and vertices $v \in R_{\mathcal{X}}$ then we can infer strong spatial mixing (Section 6) and rapid mixing (Section 7).

4. Edge discrepancies

Similarly to the definition of a vertex colouring we define a q -, q_0 - and 0 -colouring of a set $E \subseteq E_G$ of edges to be a function from E to Q , Q_0 and $\{0\}$, respectively. If B is an edge colouring of E , and E' is a subset of E then $B(E')$ is the colouring of E' induced by B . For an edge $e \in E$, $B(e)$ is the colour of e under B . Given a region R and a q_0 -colouring B of $\mathcal{E}R$, a proper q -colouring σ of R agrees with B if $\sigma(u) \neq B(e)$ for all edges $e \in \mathcal{E}R$, where $u \in R$ is incident to e . We let $\Omega_R(B)$ denote the set of all proper q -colourings of R that agree with B . The uniform distribution on $\Omega_R(B)$ is denoted π_B .

Let E be a set that contains the four edges that are incident to some vertex $v \in V_G$. Two edges $e, e' \in E$ are *adjacent* if there is a clockwise ordering around v of the edges in E such that e' follows immediately after e . Similarly to a vertex-boundary pair \mathcal{X} we define an edge-boundary pair X as follows. Note that this definition is equivalent to the notion of a relevant boundary-pair in [12].

DEFINITION 8 (Edge-boundary pair). An *edge-boundary pair* X consists of

- a region R_X ,
- a distinguished boundary edge $e_X = (w_X, v_X) \in \mathcal{E}R_X$ with $w_X \in \partial R_X$, $v_X \in R_X$, and
- a pair (B_X, B'_X) of q_0 -colourings of $\mathcal{E}R_X$ that are identical on all edges except on e_X , where they differ.

We require

- $B_X(e_X) \in Q$ and $B'_X(e_X) \in Q$, and

- any two adjacent boundary edges that share a vertex in ∂R_X have the same colour in at least one of the two colourings B_X and B'_X (and so in both of B_X and B'_X except when edge e_X is involved).

Suppose X is an edge-boundary pair. For a coupling Ψ_X of π_{B_X} and $\pi_{B'_X}$ we define $1_{\Psi_X, v}$ to be the indicator random variable for the event that, when a pair of colourings is drawn from Ψ_X , the colour of vertex $v \in R_X$ differs in these two colourings. For any edge-boundary pair X we define Ψ_X^{\min} to be some coupling of π_{B_X} and $\pi_{B'_X}$ minimising $\mathbb{E}[1_{\Psi_X, v_X}]$. For every pair of colours $c, c' \in Q$, let $p_X^{\min}(c, c')$ be the probability that $C(v_X) = c$ and $C'(v_X) = c'$, where (C, C') is a pair of colourings drawn from Ψ_X^{\min} . For a vertex $v \in R_X$, let $d(e_X, v)$ denote the distance within R_X from edge e_X to v . Thus, $d(e_X, v_X) = 1$ and if $v \in R_X$ adjoins v_X then $d(e_X, v) = 2$, and so on. We wish to construct a coupling Ψ_X of π_{B_X} and $\pi_{B'_X}$ such that $\mathbb{E}[1_{\Psi_X, v}]$ decreases exponentially in the distance $d(e_X, v)$. In order to do this we use a recursive coupling. To aid the analysis we define a labelled tree T_X associated with each edge-boundary pair X . The notion of T_X was introduced by Goldberg, Martin and Paterson in [12].

Suppose X is an edge-boundary pair. We will now construct the tree T_X . Start with a node r which will be the root of T_X . For every pair $c, c' \in Q$ of distinct colours, add an edge labelled $(p_X^{\min}(c, c'), v_X)$ from r to a new node $r_{c, c'}$. Let e_1, e_2, e_3 be the clockwise ordering of the edges incident to v_X (excluding edge e_X) such that e_X appears between e_3 and e_1 . The i -th neighbour of v_X denotes the vertex that is incident to e_i . If the i -th neighbour of v_X is not in R_X then we define $X_i(c, c') = \emptyset$. If the i -th neighbour of v_X is in R_X then let $X_i(c, c')$ be the edge-boundary pair consisting of

- The region $R_{X_i(c, c')} = R_X \setminus \{v_X\}$,
- the distinguished boundary edge $e_{X_i(c, c')} = e_i$, and
- the pair $(B_{X_i(c, c')}, B'_{X_i(c, c')})$ of q_0 -colourings of $\mathcal{E}R_{X_i(c, c')}$ such that both colourings are identical to B_X on all edges in $\mathcal{E}R_{X_i(c, c')} \setminus \{e_1, e_2, e_3\}$. The colours of the boundary edges in $\{e_1, e_2, e_3\}$ are assigned as follows.
 - $B_{X_i(c, c')}(e_i) = c'$ and $B'_{X_i(c, c')}(e_i) = c$.
 - For the boundary edge $e_j \in \{e_1, e_2, e_3\}$ such that $j < i$, both $B_{X_i(c, c')}(e_j)$ and $B'_{X_i(c, c')}(e_j)$ are c' .
 - For the boundary edge $e_j \in \{e_1, e_2, e_3\}$ such that $j > i$, both $B_{X_i(c, c')}(e_j)$ and $B'_{X_i(c, c')}(e_j)$ are c .

If the i -th neighbour of v_X is in R_X , recursively construct the tree $T_{X_i(c, c')}$ and join it to T_X by adding an edge with label $(1, \cdot)$ from $r_{c, c'}$ to the root of $T_{X_i(c, c')}$. Note that if v_X has no neighbours in R_X then $r_{c, c'}$ is a leaf. That completes the construction of T_X .

We say that an edge e of T_X is *degenerate* if the second component of its label is “.”. For edges e and e' of T_X , we write $e \rightarrow e'$ to denote the fact that e is an ancestor of e' . That is, either $e = e'$, or e is a proper ancestor of e' . Define the *level* of an edge e of T_X to be the number of non-degenerate edges on the path from the root down to, and including, e . Suppose that e is an edge of T_X with label (p, v) . We say that the *weight* $w(e)$ of edge e is p . Also the *name* $n(e)$ of edge e is v . The *likelihood* $l(e)$ of e is $\prod_{e': e' \rightarrow e} w(e')$. The *cost* $\gamma(v, T_X)$ of a vertex $v \in R_X$

is $\sum_{e:n(e)=v} l(e)$. If the region R_X is not connected and vertex v_X and a vertex $v \in R_X$ belong to different connected components, then there will be no edge with name v in T_X and we define $\gamma(v, T_X) = 0$. We have the following lemma, which is proved in [12] as Lemma 12.

LEMMA 9 ([12, Lemma 12]). *For every edge-boundary pair X there exists a coupling Ψ_X of π_{B_X} and $\pi_{B'_X}$ such that $\mathbb{E}[1_{\Psi_X, v}] \leq \gamma(v, T_X)$ for all $v \in R_X$.*

A key ingredient from the construction of T_X that affects $\gamma(v, T_X)$ is the quantity $\mathbb{E}[1_{\Psi_X^{\min}, v_X}]$, which we denote $\nu(X)$. Thus,

$$\nu(X) = \mathbb{E}[1_{\Psi_X^{\min}, v_X}] = \sum_{\substack{c, c' \in Q, \\ c \neq c'}} p_X^{\min}(c, c').$$

For an edge-boundary pair X and an integer $d \geq 1$, let $E_d(X)$ denote the set of level- d edges in T_X , and define $\Gamma^d(X) = \sum_{e \in E_d(X)} l(e)$. We define $\Gamma^d(\emptyset) = 0$ for $d \geq 1$. Equivalently, we can define $\Gamma^d(X)$ recursively:

$$\Gamma^1(X) = \nu(X) = \sum_{\substack{c, c' \in Q, \\ c \neq c'}} p_X^{\min}(c, c'), \tag{1}$$

and for $d > 1$ we have

$$\Gamma^d(X) = \sum_{\substack{c, c' \in Q, \\ c \neq c'}} p_X^{\min}(c, c') \sum_{i=1}^3 \Gamma^{d-1}(X_i(c, c')). \tag{2}$$

LEMMA 10. *Suppose X is an edge-boundary pair and $R' \subseteq R_X$. Then there is a coupling Ψ_X of π_{B_X} and $\pi_{B'_X}$ such that*

$$\sum_{v \in R'} \mathbb{E}[1_{\Psi_X, v}] \leq \sum_{d \geq d(e_X, R')} \Gamma^d(X).$$

Proof. By Lemma 9 there is a coupling Ψ_X of π_{B_X} and $\pi_{B'_X}$ such that $\mathbb{E}[1_{\Psi_X, v}] \leq \gamma(v, T_X)$ for $v \in R_X$. Thus,

$$\begin{aligned} \sum_{v \in R'} \mathbb{E}[1_{\Psi_X, v}] &\leq \sum_{v \in R'} \gamma(v, T_X) \leq \sum_{v \in R'} \sum_{e:n(e)=v} l(e) \\ &\leq \sum_{d \geq d(e_X, R')} \sum_{e \in E_d(X)} l(e) \leq \sum_{d \geq d(e_X, R')} \Gamma^d(X). \end{aligned}$$

□

5. Exponential decay of $\Gamma^d(X)$

Suppose X is an edge-boundary pair. Let B be the colouring of $\mathcal{E}R_X$ such that $B(e) = B_X(e)$ for $e \in \mathcal{E}R_X \setminus \{e_X\}$ and $B(e_X) = 0$. For $i \in Q$, we define $n_i(X)$ to be the number of proper q -colourings σ in $\Omega_{R_X}(B)$ such that $\sigma(v_X) = i$. For $i, i' \in Q$, we define $N_{i, i'}(X) = \sum_{j \in Q \setminus \{i, i'\}} n_j(X)$ and

$$\mu_{i, i'}(X) = \frac{n_i(X)}{n_i(X) + N_{i, i'}(X)}.$$

Suppose $c = B_X(e_X)$ and $c' = B'_X(e_X)$. Then $\mu_{c,c'}(X)$ is the probability that v_X receives colour c in $\pi_{B'_X}$, and $\mu_{c',c}(X)$ is the probability that v_X receives colour c' in π_{B_X} . We now define $\mu(X) = \max[\mu_{c,c'}(X), \mu_{c',c}(X)]$.

LEMMA 11. *For every edge-boundary pair X , $\nu(X) \leq \mu(X)$.*

Proof. Let X be an edge-boundary pair and suppose without loss of generality that $B_X(e_X) = c$ and $B'_X(e_X) = c'$. Suppose first that $\mu_{c,c'}(X) \geq \mu_{c',c}(X)$. Then $n_c(X) \geq n_{c'}(X)$. We define a coupling Ψ_X of π_{B_X} and $\pi_{B'_X}$ as follows. Let (C, C') be a pair of colourings drawn from Ψ_X such that C is drawn from π_{B_X} and C' from $\pi_{B'_X}$. We have $\Pr[C(v_X) = c] = 0$, $\Pr[C'(v_X) = c'] = 0$, $\Pr[C'(v_X) = c] \geq \Pr[C(v_X) = c']$ and $\Pr[C'(v_X) = i] \leq \Pr[C(v_X) = i]$ for $i \in Q \setminus \{c, c'\}$. We pair up colourings in (C, C') such that $C(v_X) = C'(v_X)$ when $C'(v_X) = i$ for $i \in Q \setminus \{c, c'\}$. Then $C(v_X) \neq C'(v_X)$ only when $C'(v_X) = c$. Thus, $\Pr[C(v_X) \neq C'(v_X)] = \mu_{c,c'}(X)$ and $\nu(X) \leq \mu_{c,c'}(X)$. Suppose second that $\mu_{c',c}(X) \geq \mu_{c,c'}(X)$. Similarly to above, $\nu(X) \leq \mu_{c',c}(X)$. Thus, $\nu(X) \leq \mu(X)$. □

Suppose X is an edge-boundary pair and $c = B_X(e_X)$ and $c' = B'_X(e_X)$. In order to obtain sufficiently good upper bounds on $\nu(X)$ we use the previous lemma together with Lemma 12 below, which we first describe in words. Suppose we want to upper-bound $\mu_{c,c'}(X)$. The idea is to pick a subregion $R' \subseteq R_X$ that contains vertex v_X . Then we compute the maximum value of $\mu_{c,c'}$ for that subregion, where we maximise over colourings of the boundary of R' that are identical to B_X on the overlapping boundary edges $\mathcal{E}R_X \cap \mathcal{E}R'$. This maximum value is an upper bound on $\mu_{c,c'}(X)$. Note that Goldberg, Martin and Paterson [12, Lemma 13] gave a similar lemma in terms of $\mu(X)$. However, in this paper it is crucial to be precise about the order of the colours c and c' in $\mu_{c,c'}(X)$.

LEMMA 12. *Suppose that X is an edge-boundary pair and let $c = B_X(e_X)$, $c' = B'_X(e_X)$. Let R' be any subset of R_X which includes v_X . Let S be the set of edge-boundary pairs X' such that $R_{X'} = R'$, the distinguished edge $e_{X'} = e_X$, and for the boundary colourings $B_{X'}$ and $B'_{X'}$, we have $B_{X'}(e) = B_X(e)$ and $B'_{X'}(e) = B'_X(e)$ on $e \in \mathcal{E}R_X \cap \mathcal{E}R'$. Then $\mu_{c,c'}(X) \leq \max_{X' \in S} \mu_{c,c'}(X')$.*

Proof. Let X be an edge-boundary pair and let $c = B_X(e_X)$ and $c' = B'_X(e_X)$. For a subregion $R' \subseteq R_X$ that contains v_X , let $H = R_X \setminus R'$. For $i \in Q \setminus \{c\}$ and $\theta \in \Omega_H$, let $n_{i,\theta}$ denote the number of colourings in $\Omega_{R_X}(B_X)$ which colour v_X with colour i and H with colouring θ . For $\theta \in \Omega_H$, let $n_{c,\theta}$ denote the number of colourings in $\Omega_{R_X}(B'_X)$ which colour v_X with colour c and H with colouring θ . Let $N_{c,c',\theta} = \sum_{i \in Q \setminus \{c,c'\}} n_{i,\theta}$. Then

$$\begin{aligned} \mu_{c,c'}(X) &= \frac{n_c(X)}{n_c(X) + N_{c,c'}(X)} = \frac{\sum_{\theta \in \Omega_H} n_{c,\theta}}{\sum_{\theta \in \Omega_H} (n_{c,\theta} + N_{c,c',\theta})} \\ &\leq \max_{\theta \in \Omega_H} \frac{n_{c,\theta}}{n_{c,\theta} + N_{c,c',\theta}} \leq \max_{X' \in S} \mu_{c,c'}(X'). \end{aligned}$$

To see the last inequality, take any $\theta \in \Omega_H$ and construct the edge-boundary pair X' in S with the following parameters: $R_{X'} = R'$, $B_{X'} = B_X$ on $\mathcal{E}R_X \cap \mathcal{E}R'$ and




Figure 3: (a) An extended region \mathcal{R} . A non-shaded vertex is labelled “in”, and a shaded vertex is labelled “out”. (b) A region R with a marked edge $e \in E(R)$. (c) We see that \mathcal{R} matches R with respect to edge e in R .




Figure 4: (a) The extended region $\mathcal{R}_{M(1,2)}$. (b) The extended region $\mathcal{R}_{M(3,4)}$. (c) Labelling of vertices and edges.

$B'_{X'} = B'_X$ on $\mathcal{E}R_X \cap \mathcal{E}R'$. For each boundary edge $e \in \mathcal{E}R'$ such that $e \notin \mathcal{E}R_X$, let $B_{X'}(e) = B'_{X'}(e) = \theta(v)$, where vertex $v \in H$ is the endpoint of e in $\partial R'$. Now,

$$\frac{n_{c,\theta}}{n_{c,\theta} + N_{c,c',\theta}} = \frac{n_c(X')}{n_c(X') + N_{c,c'}(X')} = \mu_{c,c'}(X').$$

□

5.1. Extended regions

It will be convenient to introduce the notion of an *extended region* \mathcal{R} , which is a region with the following additional information: (i) Every vertex in \mathcal{R} is labelled either “in” or “out”, and (ii) one of the boundary edges of \mathcal{R} is referred to as the designated edge.

An extended region \mathcal{R} and a region R are *matching* with respect to an edge $e \in E(R)$ if there is a way of overlapping R with \mathcal{R} such that the designated edge of \mathcal{R} coincides with the edge e , and every vertex that is labelled “in” in \mathcal{R} coincides with a vertex that is in R , and every vertex that is labelled “out” in \mathcal{R} coincides with a vertex that is not in R . When illustrating extended regions in the figures, we let non-shaded faces represent vertices that are labelled “in”, and we let shaded faces represent vertices that are labelled “out”. We mark the designated boundary edge with a short and thick line segment. Figure 3 illustrates how an extended region \mathcal{R} matches a region R with respect to an edge e . Note that the overlapping takes place under any rotation or reflection of the regions.

Suppose \mathcal{R} is an extended region. An extended region \mathcal{R}' is an *extended subregion* of \mathcal{R} if \mathcal{R}' is obtained from \mathcal{R} by removing vertices, except for the vertex that is incident to the designated edge. The labelling of the vertices in \mathcal{R}' is identical to the labelling of the same vertices in \mathcal{R} .

5.2. A collection \mathcal{F} of edge-boundary pairs

Let $\mathcal{R}_{M(1,2)}$ be the extended region in Figure 4(a) and let $\mathcal{R}_{M(3,4)}$ be the extended




Figure 5: (a) The extended region \mathcal{R}_F . (a)–(c) Three of the extended regions $\mathcal{R}_{F_1}, \dots, \mathcal{R}_{F_{4720}}$. All vertices are labelled “in”.

region in Figure 4(b). Let $M_{(1,2)}$ be the set of edge-boundary pairs X such that R_X and $\mathcal{R}_{M_{(1,2)}}$ are matching with respect to e_X . Let M_4 be the set of edge-boundary pairs X such that R_X and $\mathcal{R}_{M_{(3,4)}}$ are matching with respect to e_X . Let X be an edge-boundary pair and suppose $c = B_X(e_X)$ and $c' = B'_X(e_X)$. Let v be the vertex that is a neighbour to both v_X and w_X , let e be the edge between w_X and v , and let e' be the edge between v and v_X (see Figure 4(c)). The three sets $M_1 \subseteq M_{(1,2)}$, $M_2 \subseteq M_{(1,2)}$ and $M_3 \subseteq M_4$ of edge-boundary pairs are defined as follows.

- $X \in M_1$ if $v \in R_X$ and either $\mu_{c,c'} \geq \mu_{c',c}$ and $B_X(e) = c$, or $\mu_{c',c} \geq \mu_{c,c'}$ and $B_X(e) = c'$.
- $M_2 = M_{(1,2)} \setminus M_1$.
- $X \in M_3$ if $v \notin R_X$ and either $B_X(e') = c$ or $B_X(e') = c'$.

Let \mathcal{R}_F be the extended region in Figure 5(a). For $f \in \{1, \dots, 4720\}$ we define the extended region \mathcal{R}_{F_f} such that it is an extended subregion of \mathcal{R}_F . Note that all vertices in \mathcal{R}_{F_f} are labelled “in”. The remark on page 209 explains why we define exactly these 4720 extended regions. Due to the large number of extended regions we only illustrate three of them here (Figure 5). For $f \in \{1, \dots, 4720\}$, let F_f be the set of edge-boundary pairs X such that R_X and \mathcal{R}_{F_f} are matching with respect to edge e_X . For $m \in \{1, \dots, 4\}$, let $F_{f,m} = F_f \cap M_m$. Let \mathcal{F} be the collection of all sets $F_{f,m}$. One of the extended regions $\mathcal{R}_{F_1}, \dots, \mathcal{R}_{F_{4720}}$ is defined to contain only the single vertex that is incident to the designated edge. Hence any edge-boundary pair X is guaranteed to belong to at least one of the sets in \mathcal{F} . Note that many of the sets $F_{f,m}$ are empty. For instance, if \mathcal{R}_{F_i} is the extended region in Figure 5(b) for some $i \in \{1, \dots, 4720\}$ then obviously no edge-boundary pair X can belong to both F_i and M_4 . Hence $F_{i,4} = \emptyset$.

5.3. The constants $\mu_{F_{f,m}}$

For $f \in \{1, \dots, 4720\}$ we define F'_f to be the set of edge-boundary pairs X such that the vertices of R_X are exactly those of \mathcal{R}_{F_f} , e_X is the designated edge of \mathcal{R}_{F_f} , $B_X(e_X) = 1$, $B'_X(e_X) = 2$, and the number of colours $q = 5$. For an edge-boundary pair $X \in F'_f$, let v be the vertex that is a neighbour to both v_X and w_X , let e be the edge between w_X and v , and let e' be the edge between v and v_X (see Figure 4(c)). Suppose first that $\mathcal{R}_{M_{(1,2)}}$ is an extended subregion of \mathcal{R}_{F_f} . Then we define $F'_{f,1} \subseteq F'_f$ to be the set of edge-boundary pairs $X \in F'_f$ such that $B_X(e) = 1$, we define $F'_{f,2} \subseteq F'_f$ to be the set of edge-boundary pairs $X \in F'_f$ such that $B_X(e) = 2$, and we define $F'_{f,3} = F'_{f,4} = \emptyset$. Suppose second that

$\mathcal{R}_{M_{(1,2)}}$ is not an extended subregion of \mathcal{R}_{F_f} . Then we define $F'_{f,1} = F'_{f,2} = \emptyset$, we define $F'_{f,3} \subseteq F'_f$ to be the set of edge-boundary pairs $X \in F'_f$ such that either $B_X(e') = 1$ or $B_X(e') = 2$, and we define $F'_{f,4} = F'_f$. Now, for $f \in \{1, \dots, 4720\}$ and $m \in \{1, \dots, 4\}$, we define

$$\mu_{F_f,m} = \max_{X \in F'_{f,m}} \mu_{1,2}(X)$$

if $F'_{f,m} \neq \emptyset$, and $\mu_{F_f,m} = 0$ if $F'_{f,m} = \emptyset$.

LEMMA 13. Suppose $q = 5$, $f \in \{1, \dots, 4720\}$, $m \in \{1, \dots, 4\}$ and $F_{f,m} \neq \emptyset$. Then $\nu(X) \leq \mu_{F_f,m}$ for every edge-boundary pair $X \in F_{f,m}$.

Proof. Suppose $f \in \{1, \dots, 4720\}$ and $m \in \{1, \dots, 4\}$ such that $F_{f,m} \neq \emptyset$. Let X be an edge-boundary pair in $F_{f,m}$. Let v be the vertex that is a neighbour to both v_X and w_X , let e be the edge between w_X and v , and let e' be the edge between v and v_X (see Figure 4(c)). From Lemma 11 we have that $\nu(X) \leq \mu(X)$. In order to upper-bound $\mu(X)$ we may assume without loss of generality that $B_X(e_X) = 1$ and $B'_X(e_X) = 2$.

Suppose first that $m = 1$. Without loss of generality we may assume that $B_X(e) = 1$ and hence $\mu_{1,2}(X) \geq \mu_{2,1}(X)$. Then $\mu(X) = \mu_{1,2}(X)$. Let R' be the subset of R_X such that the vertices of R' are exactly those of \mathcal{R}_{F_f} . Let S be the set of edge-boundary pairs X' such that $R_{X'} = R'$, the distinguished edge $e_{X'} = e_X$, and for the boundary colourings $B_{X'}$ and $B'_{X'}$ we have $B_{X'}(e'') = B_X(e'')$ and $B'_{X'}(e'') = B'_X(e'')$ on $e'' \in \mathcal{E}R_X \cap \mathcal{E}R'$. Note that $S \subseteq F'_{f,m}$. We have

$$\mu_{1,2}(X) \leq \max_{X' \in S} \mu_{1,2}(X') \leq \max_{X' \in F'_{f,m}} \mu_{1,2}(X') = \mu_{F_f,m},$$

where the first inequality is from Lemma 12.

Suppose second that $m = 2$. Without loss of generality we may assume that $B_X(e) = 2$ and hence $\mu_{1,2}(X) \geq \mu_{2,1}(X)$. Proceeding as above we see that $\mu(X) \leq \mu_{F_f,m}$. Now suppose $m = 3$. Without loss of generality we may assume that $B_X(e') = 1$ or $B_X(e') = 2$ and $\mu_{1,2}(X) \geq \mu_{2,1}(X)$. Proceeding as above we see that $\mu(X) \leq \mu_{F_f,m}$. Lastly, for $m = 4$ we make no assumption on the colour of edge e' and again we see that $\mu(X) \leq \mu_{F_f,m}$. □

5.4. A collection \mathcal{A} of edge-boundary pairs

Let \mathcal{R}_A be the extended region in Figure 6(a). For $a \in \{1, \dots, 342\}$ we define the extended region \mathcal{R}_{A_a} to be a subregion of \mathcal{R}_A . The extended regions are defined such that for any edge-boundary pair X , the region R_X matches exactly one of $\mathcal{R}_{A_1}, \dots, \mathcal{R}_{A_{342}}$ with respect to edge e_X . The remark on page 209 explains why we define exactly these 342 extended regions. In Figure 6 we illustrate three of the 342 extended regions. For $a \in \{1, \dots, 342\}$, let A_a be the set of edge-boundary pairs X such that R_X matches \mathcal{R}_{A_a} with respect to edge e_X . Furthermore, for $m \in \{1, \dots, 4\}$ we define $A_{a,m} = A_a \cap M_m$, and define \mathcal{A} to be the collection of all sets $A_{a,m}$. Note that many of the sets $A_{a,m}$ are empty.

5.5. Exponential decay

A set $\mathcal{S} \subseteq \mathcal{A} \times \mathcal{F} \times \mathcal{A} \times \mathcal{A} \times \mathcal{A}$ is called an $(\mathcal{A}, \mathcal{F})$ -set if the following is true about \mathcal{S} : For every set $A_{a,m} \in \mathcal{A}$, every edge-boundary pair $X \in A_{a,m}$, and every two




Figure 6: (a) The extended region \mathcal{R}_A . (a)–(c) Three of the extended regions $\mathcal{R}_{A_1}, \dots, \mathcal{R}_{A_{342}}$.

distinct colours $c, c' \in Q$ such that $p_X^{\min}(c, c') > 0$, there is a 5-tuple $(A_{a,m}, F_{f,m}, A_{a_1,m_1}, A_{a_2,m_2}, A_{a_3,m_3})$ in \mathcal{S} , such that $X \in F_{f,m}$, and for $i \in \{1, 2, 3\}$ the edge-boundary pair $X_i(c, c')$ constructed recursively in the tree T_X belongs to A_{a_i,m_i} . For values of i such that $X_i(c, c') = \emptyset$, $A_{a_i,m_i} = \emptyset$.

Suppose $\varepsilon \in (0, 1)$ is a constant. An $(\mathcal{A}, \mathcal{F})$ -set \mathcal{S} is good with respect to ε if the following is true: For $i \in \{1, \dots, 342\}$ and $j \in \{1, \dots, 4\}$ there is a constant $\alpha_{A_{i,j}}$ such that $\alpha_{A_{i,j}} \geq 1/(1 - \varepsilon)$ if $A_{i,j} \neq \emptyset$ and $\alpha_{A_{i,j}} \geq 0$ if $A_{i,j} = \emptyset$, and for every 5-tuple $(A_{a,m}, F_{f,m}, A_{a_1,m_1}, A_{a_2,m_2}, A_{a_3,m_3})$ in \mathcal{S} ,

$$\mu_{F_{f,m}}(\alpha_{A_{a_1,m_1}} + \alpha_{A_{a_2,m_2}} + \alpha_{A_{a_3,m_3}}) \leq \alpha_{A_{a,m}}(1 - \varepsilon). \tag{3}$$

LEMMA 14. Suppose $q = 5$, $\varepsilon \in (0, 1)$ is a constant, and \mathcal{S} is an $(\mathcal{A}, \mathcal{F})$ -set that is good with respect to ε . Then there is a constant $\alpha \geq 0$ such that $\Gamma^d(X) \leq \alpha(1 - \varepsilon)^d$ for all edge-boundary pairs X .

Proof. Since \mathcal{S} is good with respect to ε , there are constants $\alpha_{A_{a,m}}$, $a \in \{1, \dots, 342\}$ and $m \in \{1, \dots, 4\}$, such that Equation (3) is satisfied for every 5-tuple in \mathcal{S} . For $A_{a,m} \in \mathcal{A}$, let $\Gamma^d(A_{a,m})$ denote the maximum of $\Gamma^d(X)$ over all $X \in A_{a,m}$. Remember $\Gamma^d(\emptyset) = 0$ for $d \geq 1$. In order to show that there is a constant α such that $\Gamma^d(X) \leq \alpha(1 - \varepsilon)^d$ for every edge-boundary pair X , we will show that $\Gamma^d(A_{a,m}) \leq \alpha_{A_{a,m}}(1 - \varepsilon)^d$ for every non-empty set $A_{a,m} \in \mathcal{A}$. Then we let α be the maximum of $\alpha_{A_{a,m}}$ over all $a \in \{1, \dots, 342\}$ and $m \in \{1, \dots, 4\}$. Note that any edge-boundary pair X belongs to at least one of the sets in \mathcal{A} .

Consider any non-empty set $A_{a,m} \in \mathcal{A}$ and any edge-boundary pair $X \in A_{a,m}$. We are going to show that $\Gamma^d(X) \leq \alpha_{A_{a,m}}(1 - \varepsilon)$ by induction on d . We start with the base case $d = 1$. Since $\alpha_{A_{a,m}} \geq 1/(1 - \varepsilon)$, we have

$$\Gamma^1(X) = \nu(X) \leq \mu(X) \leq 1 \leq \alpha_{A_{a,m}}(1 - \varepsilon),$$

where the first inequality is from Lemma 11. Now consider the inductive step. We repeat Equation (2):

$$\Gamma^d(X) = \sum_{\substack{c, c' \in Q, \\ c \neq c'}} p_X^{\min}(c, c') \sum_{i=1}^3 \Gamma^{d-1}(X_i(c, c')), \tag{4}$$

where $X_i(c, c')$ is the edge-boundary pair constructed recursively in the tree T_X . Here $Q = \{1, \dots, 5\}$. For every two distinct colours $c, c' \in Q$ such that $p_X^{\min}(c, c') > 0$, we know that there is a 5-tuple $(A_{a,m}, F_{f,m}, A_{a_1,m_1}, A_{a_2,m_2}, A_{a_3,m_3})$ in \mathcal{S} such

that $X \in F_{f,m}$ and $X_i(c, c') \in A_{a_i, m_i}$, where $i \in \{1, 2, 3\}$. If the i -th neighbour of v_X is not in R_X then we have $A_{a_i, m_i} = \emptyset$. By the induction hypothesis we have

$$\sum_{i=1}^3 \Gamma^{d-1}(X_i(c, c')) \leq \sum_{i=1}^3 \Gamma^{d-1}(A_{a_i, m_i}) \leq \sum_{i=1}^3 \alpha_{A_{a_i, m_i}} (1 - \varepsilon)^{d-1}. \tag{5}$$

Using Equation (4) with Equation (5) gives

$$\begin{aligned} \Gamma^d(X) &\leq \sum_{\substack{c, c' \in Q, \\ c \neq c'}} p_X^{\min}(c, c') \sum_{i=1}^3 \alpha_{A_{a_i, m_i}} (1 - \varepsilon)^{d-1} = \nu(X) \sum_{i=1}^3 \alpha_{A_{a_i, m_i}} (1 - \varepsilon)^{d-1} \\ &\leq \mu_{F_{f,m}} \sum_{i=1}^3 \alpha_{A_{a_i, m_i}} (1 - \varepsilon)^{d-1} \leq \alpha_{A_{a,m}} (1 - \varepsilon)^d, \end{aligned}$$

where $\nu(X) \leq \mu_{F_{f,m}}$ is from Lemma 13, and the last inequality follows from Equation (3). □

The next lemma is proved by computer assistance and we will explain the details in Section 8.

LEMMA 15. *Suppose $q =$ and $\varepsilon = 1/1000$. Then $\Gamma^d(X) \leq 5(1 - \varepsilon)^d$ for every edge-boundary pair X .*

Proof. In order to prove this lemma we use computer assistance. The computerised steps are to first calculate all the constants $\mu_{F_{f,m}}$, then generate an $(\mathcal{A}, \mathcal{F})$ -set \mathcal{S} that is good with respect to $\varepsilon = 1/1000$. This last step is broken into the following steps. First we generate an $(\mathcal{A}, \mathcal{F})$ -set \mathcal{S} . Then, for every 5-tuple $(A_{a,m}, F_{f,m}, A_{a_1, m_1}, A_{a_2, m_2}, A_{a_3, m_3})$ in \mathcal{S} , we add the inequality

$$\mu_{F_{f,m}} (\alpha_{A_{a_1, m_1}} + \alpha_{A_{a_2, m_2}} + \alpha_{A_{a_3, m_3}}) \leq \alpha_{A_{a,m}} (1 - \varepsilon)$$

to a linear program. The unknowns in this linear program are the variables $\alpha_{A_{a,m}}$. A solution to the linear program is found with $\alpha_{A_{a,m}} \in [2, 5]$ for $A_{a,m} \neq \emptyset$ and $\alpha_{A_{a,m}} \in [0, 5]$ for $A_{a,m} = \emptyset$. Hence \mathcal{S} is good with respect to ε . By Lemma 14 it follows that $\Gamma^d(X) \leq \alpha(1 - \varepsilon)^d$ for every edge-boundary pair X , where $\alpha \geq 0$ is a constant. From the proof of Lemma 14 we see that we can choose α to be the maximum of all $\alpha_{A_{a,m}}$, which is 5. □

REMARK. One probably asks why the sets in \mathcal{A} and \mathcal{F} are the sets we use to prove mixing. The sets in \mathcal{A} and \mathcal{F} , or the extended regions \mathcal{R}_{A_i} and \mathcal{R}_{F_i} to be more precise, have arisen from a lengthy process of trial and error and experiments. One part of the proof of Lemma 15 above is to find a solution to a linear program. If the values $\mu_{F_{f,m}}$ are too large then there will be no solution to this linear program. In order to obtain smaller values $\mu_{F_{f,m}}$ we must increase the size of the regions \mathcal{R}_{F_i} . Small extended regions \mathcal{R}_{A_i} contain only little information about which vertices are in and not in the region R_X for an edge-boundary pair $X \in A_i$. In particular, with small regions \mathcal{R}_{A_i} we quickly lose information about which vertices are in and not in the regions $R_{X_i(c, c')}$ for the recursively constructed edge-boundary pairs $X_i(c, c')$. Thus, too small extended regions \mathcal{R}_{A_i} will result in a linear program that is too small and has no solution. We started with a few small extended regions \mathcal{R}_{A_i}

and \mathcal{R}_{F_i} and slowly increased the sizes of them until we obtained a linear program that could be successfully solved. We let the regions grow in a way that seemed reasonable based on experiments and intuition.

6. Strong spatial mixing

LEMMA 16. *Suppose $\varepsilon = 1/1000$ and $q = 5$. Suppose \mathcal{X} is a vertex-boundary pair and $R' \subseteq R_{\mathcal{X}}$. Then there is a coupling $\Psi_{\mathcal{X}}$ of $\pi_{\mathcal{B}_{\mathcal{X}}}$ and $\pi_{\mathcal{B}'_{\mathcal{X}}}$ such that*

$$\sum_{v \in R'} \mathbb{E}[1_{\Psi_{\mathcal{X},v}}] \leq \frac{30}{\varepsilon(1-\varepsilon)}(1-\varepsilon)^{d(w_{\mathcal{X}},R')}.$$

Proof. First suppose that $w_{\mathcal{X}}$ has a neighbour $y \notin R_{\mathcal{X}}$. Let $k = |E| \leq 3$, where $E = \{e_1, \dots, e_k\} \subseteq \mathcal{E}R_{\mathcal{X}}$ is the set of boundary edges incident to $w_{\mathcal{X}}$. Label the edges in E clockwise around $w_{\mathcal{X}}$ so that edge $(w_{\mathcal{X}}, y)$ appears between edge e_k and e_1 when traversing edges around $w_{\mathcal{X}}$ in clockwise direction. This guarantees that e_i and e_j are adjacent only if i and j differ by 1.

For $i = 1, \dots, k$, let X_i be the edge-boundary pair consisting of region $R_{X_i} = R_{\mathcal{X}}$, the distinguished edge $e_{X_i} = e_i$, and boundary colourings B_{X_i} and B'_{X_i} . For every boundary edge $e = (w, v) \in \mathcal{E}R_{X_i} \setminus E$, where $w \in \partial R_{\mathcal{X}}$, we have $B_{X_i}(e) = B'_{X_i}(e) = \mathcal{B}_{\mathcal{X}}(w)$. The colours of the edges in E are assigned as follows.

- $B_{X_i}(e_j) = B'_{X_i}(e_j) = \mathcal{B}'_{\mathcal{X}}(w_{\mathcal{X}})$ for $j = 1, \dots, i - 1$,
- $B_{X_i}(e_j) = \mathcal{B}_{\mathcal{X}}(w_{\mathcal{X}})$ and $B'_{X_i}(e_j) = \mathcal{B}'_{\mathcal{X}}(w_{\mathcal{X}})$ for $j = i$, and
- $B_{X_i}(e_j) = B'_{X_i}(e_j) = \mathcal{B}_{\mathcal{X}}(w_{\mathcal{X}})$ for $j = i + 1, \dots, k$.

By Lemma 10 there is a coupling Ψ_i of $\pi_{B_{X_i}}$ and $\pi_{B'_{X_i}}$ such that

$$\sum_{v \in R'} \mathbb{E}[1_{\Psi_i,v}] \leq \sum_{d \geq d(e_{X_i}, R')} \Gamma^d(X_i). \tag{6}$$

Let $\Psi_{\mathcal{X}}$ be the coupling of $\pi_{\mathcal{B}_{\mathcal{X}}}$ and $\pi_{\mathcal{B}'_{\mathcal{X}}}$ defined by composing the couplings Ψ_1, \dots, Ψ_k . More precisely, in order to choose a pair (σ_0, σ_k) of colourings from $\Psi_{\mathcal{X}}$, first draw the pair (σ_0, σ_1) from Ψ_1 . Say $\sigma_0 = x_0$ and $\sigma_1 = x_1$. Then choose the pair (σ_1, σ_2) from the conditional distribution Ψ_2 , conditioned on $\sigma_1 = x_1$. Say $\sigma_2 = x_2$. Then choose the pair (σ_2, σ_3) from the conditional distribution Ψ_3 , conditioned on $\sigma_2 = x_2$, and so on. Hence, σ_0 is drawn from $\pi_{B_{X_1}} = \pi_{\mathcal{B}_{\mathcal{X}}}$ and σ_k is drawn from $\pi_{B'_{X_k}} = \pi_{\mathcal{B}'_{\mathcal{X}}}$. By the construction of the coupling $\Psi_{\mathcal{X}}$ it follows that if the colour of a vertex $v \in R_{\mathcal{X}}$ differs in a pair (σ_0, σ_k) drawn from $\Psi_{\mathcal{X}}$ then it must differ in at least one of the pairs (σ_{i-1}, σ_i) drawn from Ψ_i , where $i = 1, \dots, k$. Using Equation (6) and Lemma 15 we have

$$\begin{aligned} \sum_{v \in R'} \mathbb{E}[1_{\Psi_{\mathcal{X},v}}] &\leq \sum_{v \in R'} \sum_{i=1}^k \mathbb{E}[1_{\Psi_i,v}] = \sum_{i=1}^k \sum_{v \in R'} \mathbb{E}[1_{\Psi_i,v}] \\ &\leq \sum_{i=1}^k \sum_{d \geq d(e_{X_i}, R')} \Gamma^d(X_i) \leq \sum_{i=1}^k \sum_{d \geq d(w_{\mathcal{X}}, R')} 5(1-\varepsilon)^d \\ &= \sum_{i=1}^k \frac{5}{\varepsilon} (1-\varepsilon)^{d(w_{\mathcal{X}}, R')} \leq \frac{15}{\varepsilon} (1-\varepsilon)^{d(w_{\mathcal{X}}, R')}. \end{aligned}$$

Now suppose all neighbours of $w_{\mathcal{X}}$ are in $R_{\mathcal{X}}$. Breaking the discrepancy at vertex $w_{\mathcal{X}}$ into edge-boundary pairs X_i as above is not possible because the induced edge-boundary pairs are not valid with respect to the colouring of adjacent boundary edges.

Let $u \in R_{\mathcal{X}}$ be a neighbour of $w_{\mathcal{X}}$. Suppose $u \notin R'$. Let $R_{\mathcal{X},u} = R_{\mathcal{X}} \setminus \{u\}$ be the region $R_{\mathcal{X}}$ after removing vertex u . For $c \in Q$, let $\mathcal{B}_{\mathcal{X},c}$ be the colouring of the vertex-boundary $\partial R_{\mathcal{X},u}$ such that for all $v \in \partial R_{\mathcal{X}} \cap \partial R_{\mathcal{X},u}$, $\mathcal{B}_{\mathcal{X},c}(v) = \mathcal{B}_{\mathcal{X}}(v)$, and $\mathcal{B}_{\mathcal{X},c}(u) = c$ (if $u \in \partial R_{\mathcal{X},u}$). Similarly, for $c' \in Q$, let $\mathcal{B}'_{\mathcal{X},c'}$ be the colouring of the vertex-boundary $\partial R_{\mathcal{X},u}$ such that for all $v \in \partial R_{\mathcal{X}} \cap \partial R_{\mathcal{X},u}$, $\mathcal{B}'_{\mathcal{X},c'}(v) = \mathcal{B}'_{\mathcal{X}}(v)$, and $\mathcal{B}'_{\mathcal{X},c'}(u) = c'$. Note that the colourings $\mathcal{B}_{\mathcal{X},c}$ and $\mathcal{B}'_{\mathcal{X},c'}$ can differ on up to two vertices, namely on vertex $w_{\mathcal{X}}$ and u . We break the difference in the (up to) two vertices $w_{\mathcal{X}}$ and u on the boundary $\partial R_{\mathcal{X},u}$ into differences in the edges that bound them.

Let $k = |E| \leq 6$, where $E = \{e_1, \dots, e_k\} \subseteq \mathcal{E}R_{\mathcal{X},u}$ is the set of boundary edges incident to $w_{\mathcal{X}}$ or u . Label the edges in E clockwise around $w_{\mathcal{X}}$ and u so that e_k and e_1 are not adjacent. Such a labelling is always possible since $w_{\mathcal{X}}$ and u are neighbours. This guarantees that e_i and e_j are only adjacent if i and j differ by 1.

Let $c \in Q$ and $c' \in Q$ be two (not necessarily different) colours. Similarly to above, for $i = 1, \dots, k$, let X_i be the edge-boundary pair consisting of region $R_{X_i} = R_{\mathcal{X},u}$, the distinguished edge $e_{X_i} = e_i$, and boundary colourings B_{X_i} and B'_{X_i} . The colourings B_{X_i} and B'_{X_i} are defined similarly to above, as a sequence of colourings differing only on the distinguished edge e_i . That is, for a boundary edge $e = (w, v) \in \mathcal{E}R_{\mathcal{X},u}$, where $w \in \partial R_{\mathcal{X},u}$, we have $B_{X_1}(e) = \mathcal{B}_{\mathcal{X},c}(w)$ and $B'_{X_k}(e) = \mathcal{B}'_{\mathcal{X},c'}(w)$. Let Ψ_i be a coupling of $\pi_{B_{X_i}}$ and $\pi_{B'_{X_i}}$ such that Equation (6) is satisfied, which is possible due to Lemma 10. We now construct a coupling $\Psi_{\mathcal{X}}$ of $\pi_{\mathcal{B}_{\mathcal{X}}}$ and $\pi_{\mathcal{B}'_{\mathcal{X}}}$ in the following way.

Let $\Psi'_{\mathcal{X}}$ be any coupling of $\pi_{\mathcal{B}_{\mathcal{X}}}$ and $\pi_{\mathcal{B}'_{\mathcal{X}}}$. Let (C, C') be the random variable corresponding to the pair of colourings drawn from $\Psi_{\mathcal{X}}$ (yet to be constructed). We will choose the colour of u in C and C' according to $\Psi'_{\mathcal{X}}$. Let c and c' be the colour of u drawn from $\Psi'_{\mathcal{X}}$. Let $\Psi_{\mathcal{X},c,c'}$ be a coupling of $\pi_{\mathcal{B}_{\mathcal{X},c}}$ and $\pi_{\mathcal{B}'_{\mathcal{X},c'}}$. To complete the construction of $\Psi_{\mathcal{X}}$ we colour the remaining vertices in $R_{\mathcal{X}}$ by choosing two colourings from $\Psi_{\mathcal{X},c,c'}$. The coupling $\Psi_{\mathcal{X},c,c'}$ is constructed by composing the k couplings Ψ_{X_i} as above. We have

$$\begin{aligned} \sum_{v \in R'} \mathbb{E}[1_{\Psi_{\mathcal{X}},v}] &\leq \sum_{v \in R'} \sum_{i=1}^k \mathbb{E}[1_{\Psi_i,v}] = \sum_{i=1}^k \sum_{v \in R'} \mathbb{E}[1_{\Psi_i,v}] \\ &\leq \sum_{i=1}^k \sum_{d \geq d(e_{X_i}, R')} \Gamma^d(X_i) \leq \sum_{i=1}^k \sum_{d \geq d(w_{\mathcal{X}}, R') - 1} 5(1 - \varepsilon)^d \\ &= \sum_{i=1}^k \frac{5}{\varepsilon} (1 - \varepsilon)^{d(w_{\mathcal{X}}, R') - 1} \leq \frac{30}{\varepsilon(1 - \varepsilon)} (1 - \varepsilon)^{d(w_{\mathcal{X}}, R')}, \end{aligned}$$

where the -1 in “ $d(w_{\mathcal{X}}, R') - 1$ ” comes from the fact that the distance from the discrepancy edge e_{X_i} to R' may be one less than $d(w_{\mathcal{X}}, R')$. Since we sum over all distances greater than or equal to $d(w_{\mathcal{X}}, R') - 1$, and $(1 - \varepsilon)^0 = 1$, we note that the bound also holds when $u \in R'$. □

We now prove Theorem 4 of strong spatial mixing for $q = 5$ colours.

THEOREM (4, repeated). *The system specified by proper 5-colourings of the kagome lattice has strong spatial mixing.*

Proof. Consider the vertex-boundary pair \mathcal{X} such that, from Definition 1 of strong spatial mixing, we have $R_{\mathcal{X}} = R$, $\mathcal{B}_{\mathcal{X}} = \mathcal{B}$, $\mathcal{B}'_{\mathcal{X}} = \mathcal{B}'$ and $w_{\mathcal{X}} = w$. Let R' be any subregion of R . The total variation distance between $\pi_{\mathcal{B},R'}$ and $\pi_{\mathcal{B}',R'}$ is upper-bounded by the probability that R' differ under any coupling Ψ of $\pi_{\mathcal{B},R'}$ and $\pi_{\mathcal{B}',R'}$. This probability is upper-bounded by $\sum_{v \in R'} \mathbb{E}[1_{\Psi}, v]$. Using the coupling $\Psi_{\mathcal{X}}$ in Lemma 16, we have

$$d_{\text{TV}}(\pi_{\mathcal{B},R'}, \pi_{\mathcal{B}',R'}) = d_{\text{TV}}(\pi_{\mathcal{B}_{\mathcal{X}},R'}, \pi_{\mathcal{B}'_{\mathcal{X}},R'}) \leq \sum_{v \in R'} \mathbb{E}[1_{\Psi_{\mathcal{X}}}, v] \leq \alpha |R'| (1 - \varepsilon)^{d(w_{\mathcal{X}}, R')},$$

where $\varepsilon = 1/1000$ and $\alpha = 30/(\varepsilon(1 - \varepsilon))$. □

7. Rapid mixing

The implication from strong spatial mixing to rapidly mixing Glauber dynamics is only known to hold for graphs of *sub-exponential growth* [25], meaning that the number of vertices at distance d from any vertex v is sub-exponential in d . This is an important property we make use of in the proof of rapid mixing in this section. For further discussion on this topic in general, see [12], in particular [12, Section 7.5].

LEMMA 17. *Let $v \in V_{\mathcal{G}}$ be any vertex in the kagome lattice and let $n_d(v)$ denote the number of vertices at distance d from v . Then $n_d(v) \in \Theta(d)$.*

Proof. Recall the definition of the kagome lattice in Section 1.1, in particular Figure 1. First assume that $v \in V_{\text{odd}}$. In order to derive lower and upper bounds on $n_d(v)$, we assume without loss of generality that $v = (1, 1)$ is the vertex at x -coordinate 1 and y -coordinate 1. Fix any positive integer d .

We first derive a lower bound on $n_d((1, 1))$. For each odd value of $y \in \{1, \dots, d\}$, let (x, y) be the vertex at distance d from $(1, 1)$ that is reached with the following path: $(1, 1), (2, 2), (3, 3), \dots, (y, y), (y + 2, y), (y + 4, y) \dots, (x, y)$. Note that vertex $(y, y) \in V_{\text{odd}}$, and from (y, y) we go as far as possible to the right. Also note that there is no path from $(1, 1)$ to (x, y) that is shorter than length d . Thus, there are at least $\lfloor d/2 \rfloor$ vertices at distance d from $(1, 1)$, and we have $n_d((1, 1)) \geq \lfloor d/2 \rfloor$.

When deriving an upper bound on $n_d((1, 1))$ we will use two claims:

Claim 1. For any two vertices (x, y_{low}) and (x, y_{high}) , where $1 \leq y_{\text{low}} < y_{\text{high}}$, the distance between $(1, 1)$ and (x, y_{low}) is strictly smaller than the distance between $(1, 1)$ and (x, y_{high}) . We prove the claim by considering two cases:

Case (i). Assume that x is odd, and hence both (x, y_{low}) and (x, y_{high}) are in V_{odd} . Consider a shortest path from $(1, 1)$ to (x, y_{high}) . The path must use a vertex $(x_{\text{pass}}, y_{\text{low}}) \in V_{\text{odd}}$ at y -coordinate y_{low} . From $(x_{\text{pass}}, y_{\text{low}})$ we can reach (x, y_{low}) in exactly $|x - x_{\text{pass}}|/2$ steps. The number of steps required to reach (x, y_{high}) from $(x_{\text{pass}}, y_{\text{low}})$ is strictly greater than $|x - x_{\text{pass}}|/2$ since some steps must be used to increase the y -coordinate so it will eventually reach y_{high} , and for each such up-move the x -coordinate is increased/decreased only by 1. Thus, if x is odd then the distance between $(1, 1)$ and (x, y_{low}) is strictly smaller than the distance between (x, y_{high}) .

Case (ii). Assume that x is even, and hence both (x, y_{low}) and (x, y_{high}) are in V_{even} . We will use the same argument as for odd values of x , only with the difference that we consider a vertex $(x_{\text{pass}}, y_{\text{low}} - 1) \in V_{\text{odd}}$ on a shortest path from $(1, 1)$ to (x, y_{high}) . From $(x_{\text{pass}}, y_{\text{low}} - 1)$ we can reach (x, y_{low}) in at most $\lfloor |x - x_{\text{pass}}|/2 \rfloor + 1$ steps, where the $+1$ comes from the fact that we need to go up one y -coordinate. The number of steps required to reach (x, y_{high}) from $(x_{\text{pass}}, y_{\text{low}} - 1)$ is strictly greater than $\lfloor |x - x_{\text{pass}}|/2 \rfloor + 1$ since some steps must be used to increase the y -coordinate so it will eventually reach y_{high} , and for each such up-move the x -coordinate is increased/decreased only by 1. Thus, also for even values of x we have that the distance between $(1, 1)$ and (x, y_{low}) is strictly smaller than the distance between $(1, 1)$ and (x, y_{high}) .

Claim 2. For any two vertices (x, y_{low}) and (x, y_{high}) , where $y_{\text{low}} < y_{\text{high}} \leq 1$, the distance between $(1, 1)$ and (x, y_{high}) is strictly smaller than the distance between $(1, 1)$ and (x, y_{low}) . We prove the claim by using exactly the same reasoning as for Claim 1.

Using Claim 1 and 2 we conclude that there are at most two vertices (x, y) and (x, y') , with the same x -coordinate, at distance d from $(1, 1)$. The leftmost vertex that is at distance d from $(1, 1)$ is $(1 - 2d, 1)$. It is reached by making d consecutive left-moves. Similarly, the rightmost vertex at distance d from $(1, 1)$ is $(1 + 2d, 1)$. Thus, the x -coordinate of any vertex at distance d from $(1, 1)$ is in the set $\{1 - 2d, \dots, 1 + 2d\}$, and hence there are at most $2 \times (4d + 1) = 8d + 2$ vertices at distance d from $(1, 1)$. That is, $n_d((1, 1)) \leq 8d + 2$. We have now showed that for any vertex $v \in V_{\text{odd}}$, $\lfloor d/2 \rfloor \leq n_d(v) \leq 8d + 2$.

It remains to derive upper and lower bounds on $n_d(v)$ for $v \in V_{\text{even}}$. Without loss of generality we assume that $v = (0, 0)$ is the vertex at x -coordinate 0 and y -coordinate 0. Fix any positive integer d .

We derive a lower bound on $n_d((0, 0))$ in the same way as when $v = (1, 1)$. For each odd value of $y \in \{1, \dots, d\}$, let (x, y) be the vertex at distance d from $(0, 0)$ that is reached with the following path: $(0, 0), (1, 1), (2, 2), \dots, (y, y), (y + 2, y), (y + 4, y) \dots, (x, y)$. Thus, there are at least $\lfloor d/2 \rfloor$ vertices at distance d from $(0, 0)$, and we have $n_d((0, 0)) \geq \lfloor d/2 \rfloor$.

We now derive an upper bound on $n_d((0, 0))$. Vertex $(0, 0)$ has exactly four neighbours: $(1, 1), (1, -1), (-1, -1)$ and $(-1, 1)$, which are all in V_{odd} . The shortest path from $(0, 0)$ to any vertex at distance d from $(0, 0)$ must use one of these four vertices. Thus, an upper bound on the number of vertices at distance d from $(0, 0)$ is $n_d((0, 0)) \leq n_{d-1}((1, 1)) + n_{d-1}((1, -1)) + n_{d-1}((-1, -1)) + n_{d-1}((-1, 1))$. From the upper bound above we have that there are at most $8(d - 1) + 2$ vertices at distance $d - 1$ from a vertex in V_{odd} . Hence there are at most than $4 \times (8(d - 1) + 2) = 32d - 24$ vertices at distance d from $(0, 0)$, and we have $n_d((0, 0)) \leq 32d - 24$.

Finally, for any vertex $v \in V_G$ and any positive integer d we have shown that $\lfloor d/2 \rfloor \leq n_d(v) \leq 32d - 24$. □

For a vertex $v \in V_G$ and an integer $d \geq 0$, let $\text{Ball}_d(v)$ denote the set of vertices that are at most distance d from v . Thus we have $\text{Ball}_0(v) = \{v\}$.

LEMMA 18. *For any real number $a > 0$ there is an integer $d \geq 0$ such that*

$$\frac{|\partial \text{Ball}_d(v)|}{|\text{Ball}_d(v)|} \leq a,$$

uniformly in $v \in V_G$.

Proof. Let v be a vertex in V_G and let $a > 0$ be a real number. For an integer $d \geq 0$, let $n_d(v)$ denote the number of vertices at distance d from v . By Lemma 17, $n_d(v) \in \Theta(d)$. We have $|\partial \text{Ball}_d(v)| = n_{d+1}(v) \in \Theta(d)$ and $|\text{Ball}_d(v)| = \sum_{i=0}^d n_i(v) \in \Theta(d^2)$. Hence there is an integer $d_0 \geq 0$ such that $|\partial \text{Ball}_d(v)|/|\text{Ball}_d(v)| \leq a$ for $d \geq d_0$. \square

7.1. *The Markov chain \mathcal{M}_d*

In order to analyse the mixing time of the Glauber dynamics we first define a similar Markov chain that corresponds to heat-bath dynamics on small subregions instead of single vertices. For a region R , vertex $v \in V_G$ and integer $d \geq 0$, let $R_v^d = R \cap \text{Ball}_d(v)$. Let $R^d = \{v \in V_G \mid R_v^d \neq \emptyset\}$. For a region R , q_0 -colouring \mathcal{B} of ∂R and integer $d \geq 0$, we define the heat-bath Markov chain \mathcal{M}_d as follows. The state space is $\Omega_R(\mathcal{B})$ and a transition from a state σ is made in the following way: First choose a vertex v uniformly at random from R^d . Let \mathcal{B}_v^d be the colouring of ∂R_v^d induced by σ and \mathcal{B} . To make the transition from σ , recolour the vertices in R_v^d by sampling a colouring from $\pi_{\mathcal{B}_v^d}$, the uniform distribution on proper colourings of the region R_v^d that agree with ∂R_v^d . As for the Glauber dynamics, the stationary distribution of \mathcal{M}_d is $\pi_{\mathcal{B}}$. Since $\text{Ball}_0(v) = \{v\}$, Glauber dynamics is \mathcal{M}_0 . In order to prove rapid mixing of the Glauber dynamics, we will use the mixing time of \mathcal{M}_d for some constant d and use a Markov chain comparison method to infer rapid mixing of \mathcal{M}_0 .

To establish the mixing time of \mathcal{M}_d we use path coupling, due to Bubley and Dyer [3]. Let σ_1 and σ_2 be two states of \mathcal{M}_d , where d is to be specified. Using the path-coupling method, we only need to consider two colourings σ_1 and σ_2 that differ on exactly one vertex, which we refer to as w . That is, the Hamming distance between σ_1 and σ_2 is 1. Let \mathcal{M}_d make a transition from σ_1 to σ'_1 , and from σ_2 to σ'_2 . We want to correlate (or couple) these two transitions such that the expected Hamming distance between σ'_1 and σ'_2 is less than 1. If we can do this then we use the path-coupling theorem (see for instance [3, 7]) to infer the mixing time of \mathcal{M}_d . It is possible to construct such a coupling of the transitions provided d is sufficiently large. The idea is that we update the same vertices R_v^d in both the transition from σ_1 to σ'_1 and σ_2 to σ'_2 . If the vertices we update do not include w , and w is not in ∂R_v^d , then we choose the same colouring of R_v^d in both transitions, and hence the Hamming distance between σ'_1 and σ'_2 remains 1. If the vertices R_v^d we update contain w then again we choose the same colouring of R_v^d in both transitions, and the Hamming distance drops to 0. The only situation when the Hamming distance can increase is when w is on the boundary ∂R_v^d of the vertices R_v^d we update. In this case we use the coupling in Lemma 16 to colour the vertices in R_v^d . This guarantees that the expected Hamming distance between σ'_1 and σ'_2 will only increase by at most a constant $K = 30/(\varepsilon(1 - \varepsilon))$. Due to Lemma 18 we can choose a radius d such that the ratio of the probability of having $w \in \partial R_v^d$ and the probability of having $w \in R_v^d$ is arbitrarily small. Thus, we choose d such that the probability of decreasing the Hamming distance by 1 is so much bigger than the probability of increasing it by K that the expected Hamming distance between σ'_1 and σ'_2 is less than 1. The exact details of how to achieve this is explained in Sections 7.1 and 7.2 in [12]. In Section 7.2 in [12] a proof of the following lemma is found. Note that the notation in [12] differ slightly and of course we make use of Lemmas 16 and 18

as explained above rather than using equivalent lemmas in [12].

LEMMA 19. *Suppose $q = 5$. There is an integer $d \geq 0$ such that the Markov chain \mathcal{M}_d is rapidly mixing on any region R under any q_0 -colouring \mathcal{B} of ∂R . The mixing time $\tau_{\mathcal{M}_d}(\delta) \in O(n \log \frac{n}{\delta})$, where n is the number of vertices in R .*

7.2. Rapidly mixing Glauber dynamics

We will compare the mixing time of the Markov chain \mathcal{M}_d and the Glauber dynamics \mathcal{M}_0 by using a method of Diaconis and Saloff-Coste [4]. Their method has been used before by Goldberg, Martin and Paterson in [12] to compare the mixing time of \mathcal{M}_d and \mathcal{M}_0 under the assumption that $q \geq \Delta + 2$, where Δ is the maximum degree of the lattice. Here we consider $q = 5$ on the kagome lattice ($\Delta = 4$) and therefore we cannot make direct use of the comparison in [12]. Next we review the comparison described in [12] and provide a proof of rapidly mixing Glauber dynamics with $q = 5$ colours. For a survey on Markov chain comparison in general, see [6].

Let P_d and P_0 denote the transition matrix for the chain \mathcal{M}_d and \mathcal{M}_0 , respectively. For $i \in \{0, d\}$, let E_i be the set of pairs of distinct colourings (σ_1, σ_2) with $P_i(\sigma_1, \sigma_2) > 0$. The set E_i can be thought of as containing the edges of the transition graph of \mathcal{M}_i , and hence we sometimes refer to a pair in E_i as an edge. For every edge $(\sigma_1, \sigma_2) \in E_d$, let $\mathcal{P}_{\sigma_1, \sigma_2}$ be the set of paths from σ_1 to σ_2 using transitions of \mathcal{M}_0 . More formally, let $\mathcal{P}_{\sigma_1, \sigma_2}$ be the set of paths $\gamma = (\sigma_1 = \theta_0, \theta_1, \dots, \theta_k = \sigma_2)$ such that

- (1) each (θ_i, θ_{i+1}) is in E_0 , and
- (2) each edge in E_0 appears at most once on γ .

We write $|\gamma|$ to denote the length of path γ . So, for example, if $\gamma = (\theta_0, \dots, \theta_k)$ we have $|\gamma| = k$. Let $\mathcal{P} = \cup_{(\sigma_1, \sigma_2) \in E_d} \mathcal{P}_{\sigma_1, \sigma_2}$ be the set of all paths for all edges in E_d .

A *flow* is a function ϕ from \mathcal{P} to the interval $[0, 1]$ such that for every $(\sigma_1, \sigma_2) \in E_d$,

$$\sum_{\gamma \in \mathcal{P}_{\sigma_1, \sigma_2}} \phi(\gamma) = P_d(\sigma_1, \sigma_2) \pi_{\mathcal{B}}(\sigma_1).$$

For every $(\theta_1, \theta_2) \in E_0$, the *congestion of edge* (θ_1, θ_2) in the flow ϕ is the quantity

$$A_{\theta_1, \theta_2}(\phi) = \frac{1}{\pi_{\mathcal{B}}(\theta_1) P_0(\theta_1, \theta_2)} \sum_{\gamma \in \mathcal{P}: (\theta_1, \theta_2) \in \gamma} |\gamma| \phi(\gamma).$$

The *congestion of the flow* is the quantity

$$A(\phi) = \max_{(\theta_1, \theta_2) \in E_0} A_{\theta_1, \theta_2}(\phi).$$

Theorem 20 below describes how the mixing times of \mathcal{M}_d and \mathcal{M}_0 are related. A proof of this theorem can be found in [6, Observation 13]. As pointed out in [12], this theorem is similar to Proposition 4 of Randall and Tetali [19] except that [19, Proposition 4] requires the eigenvalues of transition matrices to be non-negative. Both results are based closely on the ideas of Aldous [2], Diaconis and Stroock [5], and Sinclair [21]. Let $\tau_{\mathcal{M}_d}(\delta)$ be the mixing time of \mathcal{M}_d and let $\tau_{\mathcal{M}_0}(\delta)$ be the mixing time of the Glauber dynamics \mathcal{M}_0 .

THEOREM 20. *Suppose that ϕ is a flow. Let $p = \min_{\theta \in \Omega_R(\mathcal{B})} P_0(\theta, \theta)$ and assume that $p > 0$. Then for any $0 < \delta' < \frac{1}{2}$*

$$\tau_{\mathcal{M}_0}(\delta) \leq \ln \frac{1}{\delta \cdot \pi_{\min}} \cdot \max \left[A(\phi) \left(\frac{\tau_{\mathcal{M}_d}(\delta')}{\ln \frac{1}{2\delta'}} + 1 \right), \frac{1}{2p} \right]$$

where $\pi_{\min} = \min_{\sigma \in \Omega_R(\mathcal{B})} \pi_{\mathcal{B}}(\sigma)$.

LEMMA 21. *Suppose that there is a flow ϕ such that the congestion $A(\phi) \in O(1)$. Then the mixing time of the Glauber dynamics \mathcal{M}_0 on a region R is $\tau_{\mathcal{M}_0}(\delta) \in O(n(n + \log \frac{1}{\delta}))$, where n is the number of vertices in R .*

Proof. From Definition 2 of Glauber dynamics, $p = \min_{\theta \in \Omega_R(\mathcal{B})} P_0(\theta, \theta) \geq 1/q$. Suppose $\delta' = 1/n$. Then by Lemma 19 we have $\tau_{\mathcal{M}_d}(\delta') \in O(n \log n)$. With $A(\phi) \in O(1)$, Theorem 20 gives

$$\tau_{\mathcal{M}_0}(\delta) \leq \ln \frac{1}{\delta \cdot \pi_{\min}} \cdot O(1) \cdot O(n) = O(n(n + \log \frac{1}{\delta}))$$

since $\pi_{\min} \geq 1/q^n$ and hence $\ln(1/\pi_{\min}) \in O(n)$. □

In order to establish the mixing time of the Glauber dynamics \mathcal{M}_0 by applying Lemma 21 we have to construct a flow ϕ such that the congestion $A(\phi) \in O(1)$. Given a q -colouring σ of a region R and a q_0 -colouring \mathcal{B} of ∂R , a *single-vertex update* of a vertex $v \in R$ is a recolouring of v to a colour $c \in Q$ such that no neighbour of v has colour c in either σ or \mathcal{B} . Suppose R is a region and σ_1 and σ_2 are two proper 5-colourings of R that differ on m vertices. The next two lemmas tell us how a series of $O(m)$ single-vertex updates applied to σ_1 can transform σ_1 to σ_2 . This sequence of single-vertex updates will be used when constructing the flow ϕ .

LEMMA 22. *Consider the region in Figure 7(a). In every proper 5-colouring of this region there is a vertex that has two neighbours with the same colour.*

Proof. Suppose σ is a proper 5-colouring of the region in Figure 7(a) such that no two neighbours of a vertex in the region have the same colour. We will show that this leads to contradiction. Without loss of generality we may assume that five of the vertices have the colours specified in Figure 7(b). A vertex is labelled with its colour. It follows that the two vertices adjacent to the vertex coloured 5 must have colour 3 and 4, otherwise there would be a vertex that has two neighbours with the same colour. Similarly, the vertices adjacent to the vertex coloured 3 must have colour 1 and 5, and therefore the two bottom left vertices must have colour 2 and 4 in σ . Figure 7(c) illustrates this fact, where a square contains the two colours of the two vertices it is overlapping. From the two left squares we see that the colour 4 must be on the vertices that are as far apart as possible. Thus, σ must agree with the colouring in Figure 7(d). Figure 7(e) illustrates how other vertices of the region must be coloured in σ , and Figure 7(f) shows the necessary colouring of the four rightmost vertices at the top. To finish the proof we note that it is impossible to assign colours to the two leftmost vertices at the top without introducing a vertex such that two of its neighbours receive the same colour. □

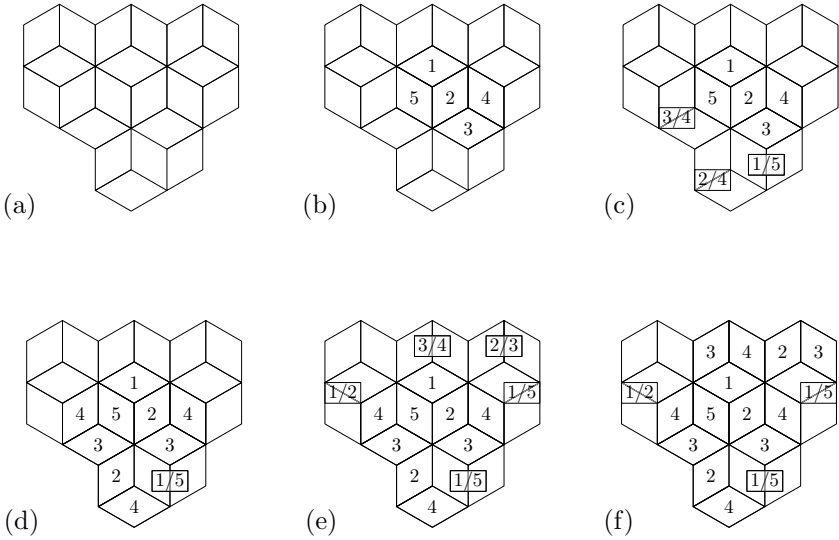


Figure 7: In every proper 5-colouring of the region in (a) there is a vertex that has two neighbours with the same colour.

LEMMA 23. Let R be a region of the kagome lattice and let \mathcal{B} be the 0-colouring of the boundary ∂R . Suppose that $q = 5$ and let σ_1 and σ_2 be any two proper q -colourings of R that differ on m vertices. We can go from σ_1 to σ_2 by applying a series of $O(m)$ single-vertex updates.

Proof. Let $v \in R$ be a vertex on which σ_1 and σ_2 differ. We will show how to recolour v to the colour it has in σ_2 by doing at most a constant number of single-vertex updates. A vertex in R that has the same colour in both σ_1 and σ_2 will not change colour after v has been updated. First we analyse situations where no boundary vertices in ∂R are involved. We note at the end of the proof that if boundary vertices are present, then it only makes it easier to recolour v . That is, assume for now that all vertices we consider belong to the region R . The proof goes through a series of cases.

If possible, simply recolour v to the colour it has in σ_2 . If this is not possible then there must be one or two neighbours of v that have colour $\sigma_2(v)$ in σ_1 . It cannot be more than two such neighbours since σ_1 is a proper colouring.

Without loss of generality, assume that $\sigma_1(v) = 1$ and $\sigma_2(v) = 2$. If v has two neighbours with colour 2 in σ_1 then we will first recolour one of these two neighbours to some other colour than 2. Let w be the neighbour of v with colour 2 that we are going to recolour. Note that $\sigma_2(w) \neq 2$ since σ_2 is a proper colouring. If possible, recolour w to some other colour than 2. If this is not possible then w is “locked” and must have three neighbours coloured 3, 4 and 5, respectively. In this case, first recolour v (which is possible since v has two neighbours with colour 2) and then recolour w to colour 1. Now only one neighbour of v has colour 2. We deal with this case next.

Without loss of generality, assume that $\sigma_1(v) = 1$ and $\sigma_2(v) = 2$, and exactly




Figure 8: The colours 1 and 2 are going to swap place. Lemma 23 guarantees that this can be done with a constant number of single-vertex updates.

one neighbour w of v has colour 2 in σ_1 . Note that $\sigma_2(w) \neq 2$ since σ_2 is a proper colouring. If possible, recolour w to something else than 2 and then recolour v to 2. If this is not possible then w is “locked” and must have four neighbours (including v) with colours 1, 3, 4 and 5, respectively, in σ_1 . Without loss of generality, consider the region in Figure 8(a), which is a subregion of R . Call this region R' . The vertices of R' are labelled with their colours in σ_1 . The vertex with colour 1 is v and the vertex with colour 2 is w . We assume without loss of generality that the two neighbours of v that are below v are the two neighbours with colour 3 and 4 in σ_1 .

Three of the vertices in R' are given the colours a, b and c , which are to be determined. Since w is “locked”, the colours a, b and c is any permutation of the colours 3, 4 and 5. If a is 3 or 4 then we recolour v to 5 and then recolour w to 1, and then recolour v to 2. If this is not the case then a must be 5, and hence the colours b and c are 3 and 4 in any order. Figure 8(b) illustrates this. We now analyse this case.

We will use Lemma 22 to show that we can recolour v to 2 without changing the colour of any other vertex except w (which will be recoloured to 1). Consider Figure 8(c) which illustrates the region R' extended with vertices in R . The vertices we extend R' with correspond to the region that we used in Lemma 22. From Lemma 22 we know that there must be at least one vertex u among the vertices we extend R' with such that u has two neighbours with the same colour. Let P be a shortest path from v to u such that the path goes from v to the neighbour above that has colour 5 and then is entirely inside the region we added to R' . Figure 8(d) illustrates an example of such a path. The path is shaded in the figure. Suppose that the vertex u is chosen such that all vertices on the path P (except from u itself) are “locked” (have four neighbours of different colours). Note that if the vertex coloured 5 above v does not have four neighbours of different colours then we let u be this vertex and hence the path P consists only of the two vertices u and v .

Suppose that the path P contains k vertices. Let c_1, \dots, c_k be the colours in σ_1 of the vertices from u to v along the path. That is, $\sigma_1(u) = c_1, c_{k-1} = 5$ and $c_k = 1$. Since u has two neighbours with the same colour, we recolour u from c_1 to another colour c'_1 . Now the vertex after u on P has two neighbours with the same colour (namely c'_1), since all its neighbours had different colours before recolouring u . We

recolour this vertex from c_2 to c'_2 . We continue this recolouring procedure along the path P all the way to vertex v , which will be recoloured to 3. Note that the vertex above v which had previously colour 5 now must have colour 3 or 4. We can now recolour w to 1 and then recolour v to 2. It remains to recolour the vertices on the path back to their original colours in σ_1 . We do this by reversing the recolouring procedure, starting with the vertex above v , which is recoloured back to 5. When u is recoloured back to c_1 we are done.

We have now shown how a constant number of single-vertex updates are applied in order to recolour a vertex v to the colour it has in σ_2 without changing the colour on vertices that have the same colour in σ_1 and σ_2 .

We note that if any vertices involved in the recolouring procedure of v are boundary vertices then this will only make it easier. Note from the statement of the lemma that we assume that a boundary vertex has colour 0. As we have seen, the tricky situations arise when a vertex is “locked” with four neighbours of different colours (excluding colour 0). Such a vertex is tricky because we cannot just change its colour to another colour in $Q = \{1, \dots, 5\}$. A vertex that is adjacent to a boundary vertex can never be “locked” since there is always at least one colour in Q that it can be recoloured to. Thus, although the part of the proof above assumes that all vertices are in R , we note that the presence of boundary vertices only makes the recolouring procedure easier. Of course, depending on which vertex v we are going to recolour, and which neighbour w is “locked”, the path P might go in a direction that is different from the one in Figure 8(d). However, the same technique is applied in order to successfully recolour v .

Finally, in order to transform σ_1 to σ_2 , we recolour each vertex v at which σ_1 and σ_2 differ. For each such vertex it takes only a constant number of single-vertex updates to do so. Since σ_1 and σ_2 differ only at m vertices, the total number of updates is $O(m)$. Notice that in recolouring a vertex v we might have changed the colours of neighbours of v as well. However, we never change the colour of a vertex whose colour agrees with the destination colouring σ_2 , a fact that ensures that the process described above indeed terminates with the colouring σ_2 . \square

We are now able to show how to construct a flow ϕ such that $A(\phi) \in O(1)$ for $q = 5$ colours. This only holds when the boundary colouring \mathcal{B} of ∂R is the 0-colouring.

LEMMA 24. *Suppose $q = 5$. Consider any region R and let \mathcal{B} be the 0-colouring of ∂R . There is a flow ϕ such that the congestion $A(\phi) \in O(1)$.*

Proof. For every pair $(\sigma_1, \sigma_2) \in E_d$ we know that σ_1 and σ_2 differ only on vertices that are contained in the ball $\text{Ball}_d(v)$ for some vertex $v \in R^d$. Let \prec be a fixed canonical ordering of the vertices in R . Let $\gamma_{\sigma_1, \sigma_2} \in \mathcal{P}_{\sigma_1, \sigma_2}$ be the path from σ_1 to σ_2 constructed according to the proof of Lemma 23. We consider vertices in order specified by \prec to make sure that $\gamma_{\sigma_1, \sigma_2}$ is well defined.

Assign all of the flow from σ_1 to σ_2 to path $\gamma_{\sigma_1, \sigma_2} \in \mathcal{P}_{\sigma_1, \sigma_2}$. That is, $\phi(\gamma_{\sigma_1, \sigma_2}) = P_d(\sigma_1, \sigma_2)\pi_{\mathcal{B}}(\sigma_1)$ and $\phi(\gamma) = 0$ for all paths $\gamma \in \mathcal{P}_{\sigma_1, \sigma_2} \setminus \{\gamma_{\sigma_1, \sigma_2}\}$. Let θ_1 and θ_2 , where $(\theta_1, \theta_2) \in E_0$, be two colourings that disagree on a vertex w . Then the

congestion of edge (θ_1, θ_2) is

$$\begin{aligned} A_{\theta_1, \theta_2}(\phi) &= \frac{1}{\pi_{\mathcal{B}}(\theta_1)P_0(\theta_1, \theta_2)} \sum_{\substack{(\sigma_1, \sigma_2) \in E_d: \\ (\theta_1, \theta_2) \in \gamma_{\sigma_1, \sigma_2}}} |\gamma_{\sigma_1, \sigma_2}| P_d(\sigma_1, \sigma_2) \pi_{\mathcal{B}}(\sigma_1) \\ &= \sum_{\substack{(\sigma_1, \sigma_2) \in E_d: \\ (\theta_1, \theta_2) \in \gamma_{\sigma_1, \sigma_2}}} |\gamma_{\sigma_1, \sigma_2}| \cdot \frac{P_d(\sigma_1, \sigma_2)}{P_0(\theta_1, \theta_2)} \cdot \frac{\pi_{\mathcal{B}}(\sigma_1)}{\pi_{\mathcal{B}}(\theta_1)} \\ &\leq \sum_{\substack{(\sigma_1, \sigma_2) \in E_d: \\ (\theta_1, \theta_2) \in \gamma_{\sigma_1, \sigma_2}}} k_1 \cdot \frac{P_d(\sigma_1, \sigma_2)}{P_0(\theta_1, \theta_2)} \leq k_1 \cdot k_2 \leq O(1), \end{aligned}$$

where k_1 and k_2 are constants, specified next. Note that $\pi_{\mathcal{B}}(\sigma_1)/\pi_{\mathcal{B}}(\theta_1) = 1$.

The path length $|\gamma_{\sigma_1, \sigma_2}|$ is upper-bounded by a constant k_1 since σ_1 and σ_2 differ only on vertices inside a ball of fixed radius d . The path $\gamma_{\sigma_1, \sigma_2}$ is constructed such that for each vertex v that is updated, we do at most a constant number of recolourings of vertices that are within constant distance from v .

To see that the last sum is bounded by a constant k_2 , note that there are only a constant number of pairs (σ_1, σ_2) in the summation. This is true since σ_1 and σ_2 agree with θ_1 on all vertices in R except in a constant-sized ball around a vertex w on which θ_1 and θ_2 differ. Let m be the number of vertices u such that R_u^d contains all vertices on which σ_1 and σ_2 differ. Note that m is bounded by a constant since σ_1 and σ_2 differ only on vertices inside a ball of fixed radius d . We have

$$P_d(\sigma_1, \sigma_2) \leq \frac{m}{|R^d|} \in O\left(\frac{1}{|R^d|}\right).$$

Furthermore,

$$P_0(\theta_1, \theta_2) \geq \frac{1}{|R^d|} \cdot \frac{1}{q} \in \Omega\left(\frac{1}{|R^d|}\right)$$

since $1/q$ is the smallest probability of making a transition in \mathcal{M}_0 from colouring θ_1 to θ_2 once vertex w on which θ_1 and θ_2 differ has been chosen for an update. Thus,

$$\frac{P_d(\sigma_1, \sigma_2)}{P_0(\theta_1, \theta_2)} \in O(1)$$

and we have that the sum is bounded by a constant k_2 .

Now, $A_{\theta_1, \theta_2}(\phi) \in O(1)$ for all $(\theta_1, \theta_2) \in E_0$ and it follows that the congestion $A(\phi) \in O(1)$. □

Finally we have the machinery for proving Theorem 5.

THEOREM (5, repeated). *For any region R of the kagome lattice and $q = 5$ colours, the Glauber dynamics is rapidly mixing on R under the 0-colouring of ∂R . The mixing time $\tau(\delta) \in O(n^2 + n \log \frac{1}{\delta})$, where n is the number of vertices in R .*

Proof. The theorem is proved by using Lemmas 21 and 24. □

The proof of Theorem 6 is similar to the proof of Theorem 5. The implications from rapid mixing of \mathcal{M}_d to rapid mixing of the heat-bath dynamics on edges hold.




Figure 9: The region R_X of the edge-boundary pair X in (a) is split into two halves (b) and (c). The split is along the vertices labelled v and v_X .

Lemma 21 has to be stated with \mathcal{M}_0 replaced by the heat-bath dynamics on edges (which slightly changes the proof) and Lemma 24 has to be adjusted to deal with an arbitrary q_0 -colouring of the boundary of the region, where $q = 5$. Showing that the congestion is constant under any q_0 -colouring of the boundary is not difficult since we are allowed to update two vertices at the same time.

8. The computational part of Lemma 15

The computational part of the proof of Lemma 15 consists of two tasks: calculating the values $\mu_{F_{f,m}}$ and constructing an $(\mathcal{A}, \mathcal{F})$ -set \mathcal{S} that is good with respect to $\varepsilon = 1/1000$. These two tasks are explained in the next sections. Both tasks are carried out using computer assistance. We have written programs in C, and the source code can be found on the webpage <http://www.csc.liv.ac.uk/~markus/kagome5colours/>

8.1. Computing $\mu_{f,m}$

Calculating the values $\mu_{F_{f,m}}$ is a computationally challenging task. We are going to calculate $\mu_{F_{f,m}}$ for $f \in \{1, \dots, 4720\}$ and $m \in \{1, \dots, 4\}$. From the definition of $\mu_{F_{f,m}}$ in Section 5.3, $\mu_{F_{f,m}} = 0$ if $F'_{f,m} = \emptyset$. For every fixed $f \in \{1, \dots, 4720\}$, $F'_{f,m} = \emptyset$ for exactly two values of $m \in \{1, \dots, 4\}$. Thus, we will have to calculate the value of $2 \times 4720 = 9440$ constants $\mu_{F_{f,m}}$. We must be able to compute a single value rather quickly, otherwise the total running time for all values will be too long. A brute-force approach would result in a running time of several months, maybe even years. We use a technique that is illustrated with the following example.

Suppose \mathcal{R}_{F_f} is the extended region in Figure 9(a) and suppose $m = 1$. Hence the set $F'_{f,m} \neq \emptyset$. The value $\mu_{F_{f,m}}$ is obtained by maximising $\mu_{1,2}(X)$ over all edge-boundary pairs $X \in F'_{f,m}$. Let v' be the vertex that is a neighbour to both v_X and w_X . Since $m = 1$, the boundary edge between w_X and v' has colour 1 in every edge-boundary pair $X \in F'_{f,m}$. An edge-boundary pair $X \in F'_{f,m}$ is therefore uniquely specified by the colour of its remaining boundary edges. Figure 9(a) illustrates an arbitrary edge-boundary pair $X \in F'_{f,m}$, where boundary edges are labelled with their colour (z_1, \dots, z_{13}) . Thus, in order to maximise $\mu_{1,2}(X)$ over all $X \in F'_{f,m}$, we could loop through all combinations of the colours z_1, \dots, z_{13} and compute $\mu_{1,2}(X)$ for each such combination. This process will take very long. Next we explain how to speed up the process.

By computing $\mu_{1,2}(X)$ for many colourings of the boundary, one quickly makes the observation that only some particular colourings of the boundary result in a large value of $\mu_{1,2}(X)$. For other colourings, $\mu_{1,2}(X)$ tends to be rather small. For example, it turns out that if z_1, \dots, z_6 are all colour 1, then $\mu_{1,2}(X)$ will be rather small regardless of the remaining colours z_7, \dots, z_{13} . Thus, setting the colours z_1, \dots, z_6 to 1 is a “bad” choice if we want to maximise $\mu_{1,2}(X)$. From this observation we conclude that if we can filter out certain “bad” colourings of the boundary then we could speed up the process of finding the maximum value $\mu_{1,2}(X)$.

We “split” the extended region \mathcal{R}_{F_f} into two extended regions $\mathcal{R}^{\text{left}}$ and $\mathcal{R}^{\text{right}}$. Figure 9(b) and (c) illustrate $\mathcal{R}^{\text{left}}$ and $\mathcal{R}^{\text{right}}$, respectively. The two regions share the vertices in the split. In this case it is vertex v_X and v , both labelled in the figure. Let X^{left} be the edge-boundary pair such that $R_{X^{\text{left}}} = \mathcal{R}^{\text{left}}$, $e_{X^{\text{left}}} = e_X$ and the boundary edges receive the same colours as in X . Boundary edges that are introduced from the split are given colour 0. Let X^{right} be the edge-boundary pair defined similarly to X^{left} but with $R_{X^{\text{right}}} = \mathcal{R}^{\text{right}}$. Figure 9(b) and (c) illustrate X^{left} and X^{right} .

Let B be the colouring of $\mathcal{E}R_X$ such that $B(e) = B_X(e)$ for $e \in \mathcal{E}R_X \setminus \{e_X\}$ and $B(e_X) = 0$. Recall from Section 5 that for $i \in Q$, $n_i(X)$ denotes the number of proper q -colourings σ in $\Omega_{R_X}(B)$ such that $\sigma(v_X) = i$. For two colours $i, i' \in Q$ we now define $n_{i,i'}^{\text{both}}$ to be the number of proper q -colourings σ in $\Omega_{R_X}(B)$ such that $\sigma(v_X) = i$ and $\sigma(v) = i'$, where v is the second vertex in the split. Thus,

$$n_i(X) = \sum_{i'=1}^q n_{i,i'}^{\text{both}}.$$

Let B^{left} be the colouring of $\mathcal{E}R_{X^{\text{left}}}$ such that $B^{\text{left}}(e_{X^{\text{left}}}) = 0$ and $B^{\text{left}}(e) = B_{X^{\text{left}}}(e)$ for $e \in \mathcal{E}R_{X^{\text{left}}} \setminus \{e_{X^{\text{left}}}\}$. For two colours $i, i' \in Q$ we define $n_{i,i'}^{\text{left}}$ to be the number of proper q -colourings σ in $\Omega_{R_{X^{\text{left}}}}(B^{\text{left}})$ such that $\sigma(v_X) = i$ and $\sigma(v) = i'$. We define $n_{i,i'}^{\text{right}}$ similarly for the edge-boundary pair X^{right} . It follows that

$$n_{i,i'}^{\text{both}} = n_{i,i'}^{\text{left}} n_{i,i'}^{\text{right}},$$

and hence

$$n_i(X) = \sum_{i'=1}^q n_{i,i'}^{\text{left}} n_{i,i'}^{\text{right}}.$$

With $q = 5$ colours, we have

$$\begin{aligned} \mu_{1,2}(X) &= \frac{n_1(X)}{\sum_{i \in \{1,3,4,5\}} n_i(X)} = \frac{\sum_{j=1}^5 n_{1,j}^{\text{left}} n_{1,j}^{\text{right}}}{\sum_{i \in \{1,3,4,5\}} \sum_{k=1}^5 n_{i,k}^{\text{left}} n_{i,k}^{\text{right}}} \\ &= \sum_{j=1}^5 \frac{n_{1,j}^{\text{left}} n_{1,j}^{\text{right}}}{\sum_{i \in \{1,3,4,5\}} \sum_{k=1}^5 n_{i,k}^{\text{left}} n_{i,k}^{\text{right}}} \\ &= \sum_{j=1}^5 \frac{1}{\sum_{i \in \{1,3,4,5\}} \sum_{k=1}^5 \left(\frac{n_{i,k}^{\text{left}}}{n_{1,j}^{\text{left}}} \times \frac{n_{i,k}^{\text{right}}}{n_{1,j}^{\text{right}}} \right)}. \end{aligned} \tag{7}$$

Note that the colours z_1, \dots, z_6 specify the quantity $n_{i,i'}^{\text{left}}$, and z_7, \dots, z_{13} specify

the quantity $n_{i,i'}^{\text{right}}$. In order to maximise $\mu_{1,2}(X)$ over edge-boundary pairs X , we could consider all combinations of the colours z_1, \dots, z_{13} and use Equation (7). There are $5^{13} \approx 1.2 \times 10^9$ such combinations, so considering them all will take very long. Now, consider two different sets of the six colours z_1, \dots, z_6 . For $i, i' \in Q$, let $n_{i,i'}^{\text{left-1}}$ be the value of $n_{i,i'}$ for the first set of colours, and let $n_{i,i'}^{\text{left-2}}$ be the value of $n_{i,i'}$ for the second set of colours. Suppose

$$\frac{n_{i,k}^{\text{left-1}}}{n_{1,j}^{\text{left-1}}} \leq \frac{n_{i,k}^{\text{left-2}}}{n_{1,j}^{\text{left-2}}} \tag{8}$$

for all $i \in \{1, 3, 4, 5\}$, $j \in \{1, \dots, 5\}$ and $k \in \{1, \dots, 5\}$. Then we have from Equation (7) that $\mu_{1,2}(X)$ can only get smaller if we take $n_{i,i'}^{\text{left}} = n_{i,i'}^{\text{left-2}}$ instead of $n_{i,i'}^{\text{left}} = n_{i,i'}^{\text{left-1}}$. In other words, there is no point considering the colours specified by the second set of colours z_1, \dots, z_6 when maximising $\mu_{1,2}(X)$. This observation suggests that we loop through all combinations of colours z_1, \dots, z_6 and compare each pair of combinations like in Equation (8). We only keep the sets of colours that cannot be ruled out in some pairwise comparison like the second set above. This gives us a collection C^{left} of colours z_1, \dots, z_6 that turns out to be much smaller than the collection of all $5^6 = 15,625$ sets of colours. Similarly we obtain a collection C^{right} of colours z_7, \dots, z_{13} for the right part of the region. In order to find which colours z_1, \dots, z_{13} that maximise $\mu_{1,2}(X)$ we combine C^{left} with C^{right} . That is, we use Equation (7) to compute $\mu_{1,2}(X)$ for each set z_1, \dots, z_6 of colours in C^{left} with each set z_7, \dots, z_{13} of colours in C^{right} .

The technique of splitting regions and filtering out boundary colourings that are guaranteed not to maximise $\mu_{1,2}(X)$ has a huge impact on the running time of the program. On a fairly powerful home-PC as of year 2006, it takes about two days to to obtain all 9440 values $\mu_{F_f,m}$.

8.2. Constructing an $(\mathcal{A}, \mathcal{F})$ -set

We describe how to construct an $(\mathcal{A}, \mathcal{F})$ -set. Let \mathcal{R}_{big} be the extended region in Figure 10(a) with some combination of labels “in” and “out” on the vertices. From \mathcal{R}_{big} we will derive 5-tuples that are added to a set \mathcal{S} . By considering all possible combinations of labels “in” and “out” on the vertices of \mathcal{R}_{big} , we construct the $(\mathcal{A}, \mathcal{F})$ -set \mathcal{S} . We describe the process by first giving a concrete example.

Fix an “in/out”-labelling of the vertices of the extended region \mathcal{R}_{big} . Let $a \in \{1, \dots, 342\}$ be the value such that \mathcal{R}_{A_a} is an extended subregion of \mathcal{R}_{big} . Note that the extended regions \mathcal{R}_{A_i} , $i \in \{1, \dots, 342\}$, are defined such that there is exactly one value $a \in \{1, \dots, 342\}$ for which this is true. Figure 10(b) shows the largest possible \mathcal{R}_{A_a} and Figure 10(c) shows the overlapping of \mathcal{R}_{big} and \mathcal{R}_{A_a} . We see from this figure that only some of the vertices of \mathcal{R}_{big} define \mathcal{R}_{A_a} . Similarly to how the extended region \mathcal{R}_{A_a} is obtained from \mathcal{R}_{big} , let $a_1, a_2, a_3 \in \{1, \dots, 342\}$ be the three unique values such that $\mathcal{R}_{A_{a_1}}$ is obtained from \mathcal{R}_{big} by the overlapping in Figure 10(d), $\mathcal{R}_{A_{a_2}}$ is obtained from \mathcal{R}_{big} by the overlapping in Figure 10(e), and $\mathcal{R}_{A_{a_3}}$ is obtained from \mathcal{R}_{big} by the overlapping in Figure 10(f). It is possible that neighbours of vertex v in Figure 10(d)–(f) are labelled “out”, meaning that some of the extended regions \mathcal{R}_{A_i} might not exist. If this is the case we define $a_i = 0$ and $A_0 = \emptyset$. For example, if the vertex to the left of vertex v in Figure 10(d) is “out” then \mathcal{R}_{A_1} cannot exist and hence $a_1 = 0$.




Figure 10: (a) The extended region \mathcal{R}_{big} (here all vertices are labelled “in”). Note that the shaded vertex is not a part of the region. (b) The extended region \mathcal{R}_A repeated. (c)–(f) Intersections of \mathcal{R}_A and \mathcal{R}_{big} .

Suppose X is an edge-boundary pair such that R_X and \mathcal{R}_{big} are matching with respect to edge e_X . Then $X \in A_a$. For $i \in \{1, 2, 3\}$ and any two distinct colours $j, j' \in Q$ such that $p_X^{\text{min}}(j, j') > 0$, suppose $X_i(j, j')$ is the extended edge-boundary pair that is constructed recursively in the tree T_X . If $X_i(j, j') = \emptyset$ then $X_i(j, j') \in A_{a_i}$. We will now be more precise about the sets of edge-boundary pairs and incorporate the sets M_1, \dots, M_4 .

Suppose without loss of generality that $B_X(e_X) = c$ and $B'_X(e_X) = c'$, and $\mu_{c,c'}(X) \geq \mu_{c',c}(X)$ for some $c, c' \in Q$. Suppose the extended region in Figure 11(a) is an extended subregion of \mathcal{R}_{big} . Suppose that the colour of the edge between w_X and v in Figure 11(a) has colour c in B_X and B'_X . Then $X \in M_1$ and hence $X \in A_{a,1}$. From Figure 11(a) we see that the extended region $\mathcal{R}_{M_{(3,4)}}$ in Figure 4(b) is an extended subregion of $\mathcal{R}_{A_{a,1}}$. Hence $X_1(j, j')$ belongs to M_3 or M_4 (or both). The crucial observation here is that $p_X^{\text{min}}(j, j') > 0$ if and only if $j' = c$. This follows from the fact that $\mu_{c,c'}(X) \geq \mu_{c',c}(X)$ and hence there is a discrepancy at v_X only when the colour c is drawn from $\pi_{B'_X}$ in the coupling Ψ_X^{min} . We therefore conclude that $X_1(j, j') \in M_3$. Thus, $X_1(j, j') \in A_{a,1,3}$. For $X_2(j, j')$ and $X_3(j, j')$ we see in Figure 11(a) that these edge-boundary pairs belong to either M_1 or M_2 . However, we are unable to tell exactly to which of the two sets these edge-boundary pairs belong. We therefore assume that any combination of the two sets is possible. The 3-tuples listed in Figure 11(a) indicate to which possible sets M_1, \dots, M_4 the edge boundary pairs $X_1(j, j')$, $X_2(j, j')$ and $X_3(j, j')$ belong. That is, a 3-tuple (m_1, m_2, m_3) means that $X_1(j, j') \in A_{a,1,m_1}$, $X_2(j, j') \in A_{a,2,m_2}$ and $X_3(j, j') \in A_{a,3,m_3}$.

Let $\mathcal{F}' \subseteq \mathcal{F}$ be the set of edge-boundary pairs $F_{i,1}$ such that \mathcal{R}_{F_i} is an extended subregion of \mathcal{R}_{big} and $i \in \{1, \dots, 4720\}$. Then $X \in F_{i,1}$ for every $F_{i,1} \in \mathcal{F}'$. Remember that we have assumed above that $X \in M_1$. Let f be the value such that $F_{f,1} \in \mathcal{F}'$ is the set that minimises $\mu_{F_{f,1}}$ over all $F_{i,1} \in \mathcal{F}'$. If the minimiser

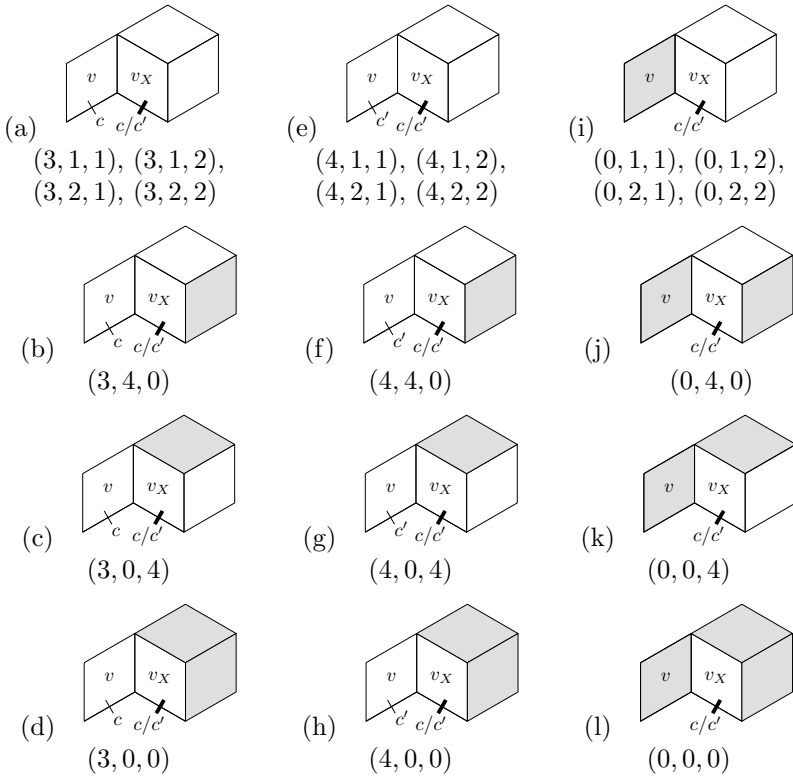


Figure 11: Twelve cases which cover all possible combinations of the sets M_1, \dots, M_4 to which the recursively constructed edge-boundary pairs $X_1(c_1, c_2)$, $X_2(c_1, c_2)$ and $X_3(c_1, c_2)$ belong. If $X \in M_1$ then (a)–(d) apply. If $X \in M_2$ then (e)–(h) apply. If $X \in M_3 \cup M_4$ then (i)–(l) apply.

is not unique, let f be the smallest i among the minimisers. Now, for each 3-tuple (m_1, m_2, m_3) in Figure 11(a) we add the following 5-tuple to the set \mathcal{S} : $(A_{a,1}, F_{f,1}, A_{a_1,m_1}, A_{a_2,m_2}, A_{a_2,m_2})$.

Summing it all up, we construct the set \mathcal{S} as follows. First take an extended region \mathcal{R}_{big} . From \mathcal{R}_{big} we uniquely derive the sets A_a, A_{a_1}, A_{a_2} and A_{a_3} . If $\mathcal{R}_{M_{(1,2)}}$ is an extended subregion of \mathcal{R}_{big} then we consider two values of m : $m = 1$ and $m = 2$. If $\mathcal{R}_{M_{(3,4)}}$ is an extended subregion of \mathcal{R}_{big} then we also consider two values of m : $m = 3$ and $m = 4$. Now suppose $X \in A_{a,m}$. The twelve cases in Figure 11 cover all possible combinations of the sets M_1, \dots, M_4 to which the recursively constructed edge-boundary pairs $X_1(j, j')$, $X_2(j, j')$ and $X_3(j, j')$ belong. More precisely, if $m = 1$ then Figure 11(a)–(d) apply. if $m = 2$ then Figure 11(e)–(h) apply. if $m = 3$ or $m = 4$ then Figure 11(i)–(l) apply. From \mathcal{R}_{big} and the value of m , we uniquely derive the set $F_{f,m}$ to which X belongs. For each 3-tuple (m_1, m_2, m_3) in the relevant case in Figure 11, we add the following 5-tuple to the set \mathcal{S} : $(A_{a,m}, F_{f,m}, A_{a_1,m_1}, A_{a_2,m_2}, A_{a_2,m_2})$. If the value of m_i in a 3-tuple is 0 then $A_{a_i,m_i} = \emptyset$. By considering every possible extended region \mathcal{R}_{big} and every possible value of m (two values per region \mathcal{R}_{big}), we construct a set \mathcal{S} that is an $(\mathcal{A}, \mathcal{F})$ -set.

References

1. D. ACHLIOPTAS, M. MOLLOY, C. MOORE and F. VAN BUSSEL, ‘Sampling grid colorings with fewer colours.’ ‘LATIN 2004: Theoretical Informatics,’ (Springer, 2004), vol. 2976 of *Lecture Notes in Computer Science* pp. 80–89, pp. 80–89. [200](#)
2. D. ALDOUS, ‘Random walks on finite groups and rapidly mixing Markov chains.’ ‘Seminar on Probability, XVII,’ (Springer, 1983), vol. 986 of *Lecture Notes in Mathematics* pp. 243–297, pp. 243–297. [215](#)
3. R. BUBLEY and M. DYER, ‘Path coupling: a technique for proving rapid mixing in Markov chains.’ ‘FOCS ’97: Proceedings of the 38th Symposium on Foundations of Computer Science,’ (IEEE Computer Society Press, 1997) pp. 223–231, pp. 223–231. [199](#), [214](#)
4. P. DIACONIS and L. SALOFF-COSTE, ‘Comparison theorems for reversible Markov chains.’ *Annals of Applied Probability* 3 (1993) 696–730. [215](#)
5. P. DIACONIS and D. STROOCK, ‘Geometric bounds for eigenvalues of Markov chains.’ *Annals of Applied Probability* 1 (1991) 36–61. [215](#)
6. M. DYER, L. A. GOLDBERG, M. JERRUM and R. MARTIN, ‘Markov chain comparison.’ *Probability Surveys* 3 (2006) 89–111. [215](#)
7. M. DYER and C. GREENHILL, ‘Random walks on combinatorial objects.’ ‘Surveys in Combinatorics,’ (Cambridge University Press, 1999), vol. 267 of *London Mathematical Society Lecture Note Series* pp. 101–136, pp. 101–136. [214](#)
8. M. DYER, A. SINCLAIR, E. VIGODA and D. WEITZ, ‘Mixing in time and space for lattice spin systems: a combinatorial view.’ *Random Structures and Algorithms* 24 (2004) 461–479. [197](#), [199](#)
9. H.-O. GEORGII, *Gibbs measures and phase transitions*. de Gruyter Studies in Mathematics 9 (Walter de Gruyter & Co., Berlin, Germany, 1988). [197](#)
10. H.-O. GEORGII, O. HÄGGSTRÖM and C. MAES, ‘The random geometry of equilibrium phases.’ *Phase Transitions and Critical Phenomena* 18 (2001) 1–142. [197](#)
11. L. A. GOLDBERG, M. JALSENIUS, R. MARTIN and M. PATERSON, ‘Improved mixing bounds for the anti-ferromagnetic potts model on \mathbb{Z}^2 .’ *LMS Journal of Computation and Mathematics* 9 (2006) 1–20. [200](#)
12. L. A. GOLDBERG, R. MARTIN and M. PATERSON, ‘Strong spatial mixing with fewer colours for lattice graphs.’ *SIAM Journal on Computing* 35 (2005) 486–517. [197](#), [199](#), [200](#), [201](#), [202](#), [203](#), [204](#), [212](#), [214](#), [215](#)
13. M. JALSENIUS, ‘Strong spatial mixing and rapid mixing with 9 colours for the triangular lattice.’ arXiv:0706.0489v1 [math-ph], (2007). [200](#)
14. M. JERRUM, ‘A very simple algorithm for estimating the number of k -colorings of a low-degree graph.’ *Random Structures and Algorithms* 7 (1995) 157–165. [199](#)
15. M. JERRUM, *Counting, Sampling and Integrating: Algorithms and Complexity* (Birkhäuser, Basel, Switzerland, 2003). [199](#)

16. M. JERRUM, L. VALIANT and V. VAZIRANI, ‘Random generation of combinatorial structures from a uniform distribution.’ *Theoretical Computer Science* 43 (1986) 169–188. [199](#)
17. F. MARTINELLI, ‘Lectures on Glauber dynamics for discrete spin models.’ ‘Lectures on Probability Theory and Statistics (Saint-Flour, 1997),’ (Springer, 1999), vol. 1717 of *Lecture Notes in Mathematics* pp. 93–191, pp. 93–191. [199](#)
18. M. MOLLOY, ‘Very rapidly mixing Markov chains for 2Δ -colourings and for independent sets in a 4-regular graph.’ *Random Structures and Algorithms* 18 (2001) 101–115. [199](#)
19. D. RANDALL and P. TETALI, ‘Analyzing Glauber dynamics by comparison of Markov chains.’ *Journal of Mathematical Physics* 41 (2000) 1598–1615. [215](#)
20. J. SALAS and A. D. SOKAL, ‘Absence of phase transition for antiferromagnetic Potts models via the Dobrushin uniqueness theorem.’ *Journal of Statistical Physics* 86 (1997) 551. [195](#), [199](#)
21. A. SINCLAIR, ‘Improved bounds for mixing rates of Markov chains and multicommodity flow.’ *Combinatorics, Probability and Computing* 1 (1992) 351–370. [215](#)
22. E. VIGODA, ‘Improved bounds for sampling colourings.’ *Journal of Mathematical Physics* 41 (2000) 1555–1569. [199](#)
23. D. WEITZ, ‘Mixing in time and space for discrete spin systems.’ Ph.D. thesis, University of California, Berkeley, (2004). [197](#), [199](#)
24. D. WEITZ, ‘Combinatorial criteria for uniqueness of Gibbs measures.’ *Random Structures and Algorithms* 27 (2005) 445–475. [197](#)
25. D. WEITZ, ‘Counting independent sets up to the tree threshold.’ ‘STOC ’06: Proceedings of the 38th Annual ACM Symposium on Theory of Computing,’ (ACM Press, 2006) pp. 140–149, pp. 140–149. [212](#)

Markus Jalsenius

www.cs.bris.ac.uk/home/markus/

Department of Computer Science
University of Liverpool
Liverpool, L69 3BX
United Kingdom

Current address:

Department of Computer Science
University of Bristol
Bristol BS8 1UB
United Kingdom