

UNIVERSALITY OF CUTOFF FOR THE ISING MODEL

EYAL LUBETZKY AND ALLAN SLY

ABSTRACT. On any locally-finite geometry, the stochastic Ising model is known to be contractive when the inverse-temperature β is small enough, via classical results of Dobrushin and of Holley in the 1970's. By a general principle proposed by Peres, the dynamics is then expected to exhibit cutoff. However, so far cutoff for the Ising model has been confirmed mainly for lattices, heavily relying on amenability and log Sobolev inequalities. Without these, cutoff was unknown at any fixed $\beta > 0$, no matter how small, even in basic examples such as the Ising model on a binary tree or a random regular graph.

We use the new framework of information percolation to show that, in any geometry, there is cutoff for the Ising model at high enough temperatures. Precisely, on any sequence of graphs with maximum degree d , the Ising model has cutoff provided that $\beta < \kappa/d$ for some absolute constant κ (a result which, up to the value of κ , is best possible). Moreover, the cutoff location is established as the time at which the sum of squared magnetizations drops to 1, and the cutoff window is $O(1)$, just as when $\beta = 0$.

Finally, the mixing time from almost every initial state is not more than a factor of $1 + \varepsilon_\beta$ faster than the worst one (with $\varepsilon_\beta \rightarrow 0$ as $\beta \rightarrow 0$), whereas the uniform starting state is at least $2 - \varepsilon_\beta$ times faster.

1. INTRODUCTION

Classical results going back to Dobrushin [12] and to Holley [14] in the early 1970's and continuing with the works of Dobrushin and Shlosman [13] and of Aizenman and Holley [1] show that, if G is any graph on n vertices with maximum degree d , the Glauber dynamics for the Ising model on G exhibits a rapid convergence to equilibrium in total-variation distance at high enough temperatures. Namely, if the inverse-temperature β is at most c_0/d for some absolute $c_0 > 0$ then the continuous-time dynamics is contractive, whence coupling techniques show that the total-variation mixing time is $O(\log n)$.

A known consequence of contraction is that the spectral gap of the dynamics is bounded away from 0, and so, by a general principle proposed by Peres in 2004 (addressing whether or not the product of the spectral gap and mixing time diverges with n), one expects the *cutoff phenomenon*¹ to occur. (For more on the cutoff phenomenon, discovered in the early 80's by Aldous and Diaconis, see [3, 4].) Concretely, Peres conjectured ([16, Conjecture 1],[17, §23.2]) cutoff for the Ising model on any sequence of transitive graphs when the mixing time is $O(\log n)$, and in particular in the range $\beta < c_0/d$ as above.

This universality principle, whereby cutoff should accompany high enough temperatures in any underlying geometry, is supported by the heuristic that at small enough β the model should qualitatively behave as if $\beta = 0$. The latter, equivalent to random walk on the hypercube, was one of the first examples of cutoff, established with an $O(1)$ -cutoff window by Aldous [2], and refined in [5, 10]. Thus, one may further expect cutoff for the Ising model with an $O(1)$ -window provided that β is small enough.

In contrast, cutoff for the Ising model has so far mainly been confirmed on \mathbb{Z}^d [21, 22], via proofs that hinged on log-Sobolev inequalities (see [6–8, 30]) that are known to hold for the Ising model on the lattice [15, 24–27, 32, 33] as well as on the sub-exponential growth rate of balls in the lattice.

¹sharp transition in the L^1 -distance of a finite Markov chain from equilibrium, dropping quickly from near 1 to near 0.

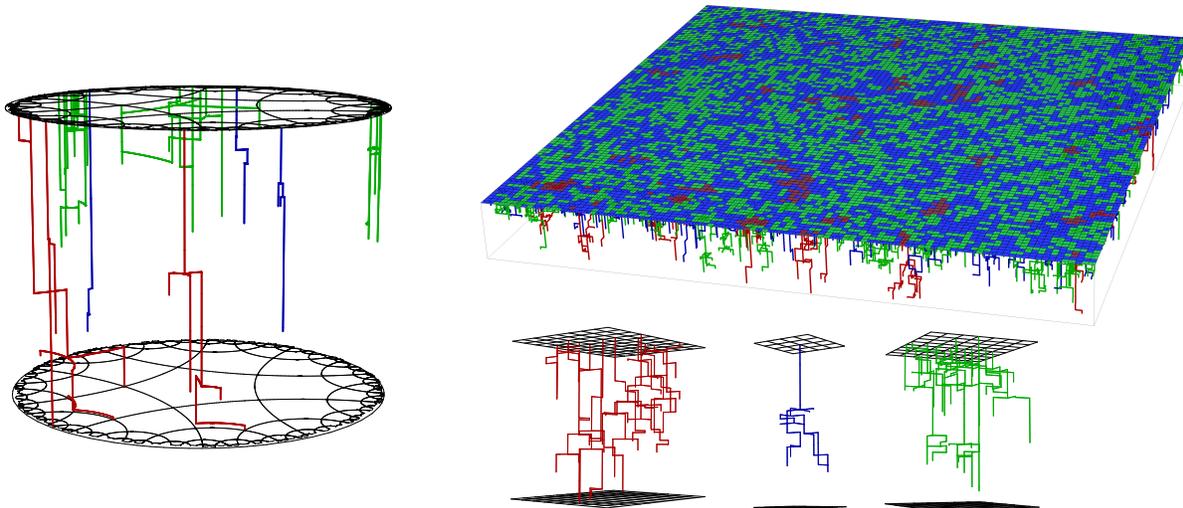


FIGURE 1. Information percolation clusters for the stochastic Ising model on two geometries: hyperbolic graph (left, showing largest 3 clusters of each type) and the lattice \mathbb{Z}^2_{100} (right). A cluster is red if it survives to time 0, blue if it dies out and is the history of a single vertex, and green o/w.

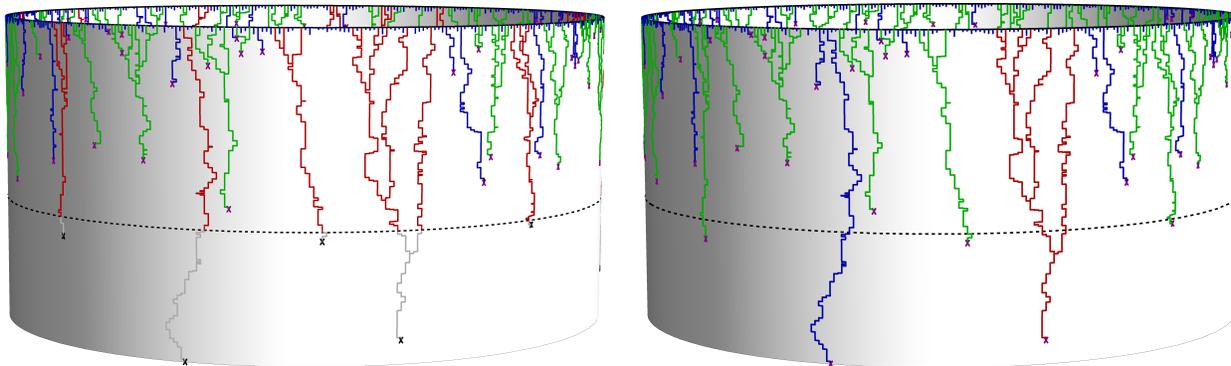


FIGURE 2. Flavor of information percolation for analyzing random initial states in 1D Ising model: On the left, the standard framework (red clusters are those reaching $t = 0$) for worst-case analysis. On the right, red clusters are redefined as those coalescing *below* $t = 0$ for the annealed analysis.

Even before requiring these powerful log-Sobolev inequalities, the restriction to sub-exponential growth rate automatically precluded the analysis of examples as basic as the Ising model on a binary tree at *any* small $\beta > 0$, or on an expander graph (e.g., a random regular graph), the hypercube, etc.

Here, using the framework of information percolation that we introduced in the companion paper [23], we confirm that on any sequence of graphs with maximum degree d , cutoff indeed occurs whenever βd is small enough, and with an $O(1)$ -window (just as when $\beta = 0$). Furthermore, we analyze the effect of the initial state on the mixing time (e.g., a warm start of i.i.d. spins vs. the all-plus starting state).

1.1. Results. Our first result establishes that, on any geometry, at high enough temperature there is cutoff within an $O(1)$ -window around the point

$$t_m = \inf \{ t > 0 : \sum_v \mathbf{m}_t(v)^2 \leq 1 \}, \quad (1.1)$$

where $\mathbf{m}_t(v)$ is the magnetization at a vertex $v \in V$ at time $t > 0$, i.e.,

$$\mathbf{m}_t(v) = \mathbb{E}X_t^+(v), \quad (1.2)$$

with X_t^+ denoting the dynamics started from all-plus. Note that on a transitive graph (such as \mathbb{Z}_n^d), the point t_m coincides with the time at which $\sum_v \mathbf{m}_t(v)$ drops to a square-root of the volume, which has the intuitive interpretation that mixing occurs once the expected sum of spins in X_t^+ drops within the normal deviations in the Ising measure. However, it turns out that for general (non-transitive) geometries (such as trees) it is the sum of *squared* magnetizations $\sum_v \mathbf{m}_t(v)^2$ that governs the mixing.

Theorem 1. *There exist absolute constants $\kappa, C > 0$ such that the following holds. Let G be a graph on n vertices with maximum degree d . For any fixed $0 < \varepsilon < 1$ and large enough n , the continuous-time Glauber dynamics for the Ising model on G with inverse-temperature $0 \leq \beta < \kappa/d$ satisfies*

$$\begin{aligned} t_{\text{MIX}}(1 - \varepsilon) &\geq t_m - C \log(1/\varepsilon), \\ t_{\text{MIX}}(\varepsilon) &\leq t_m + C \log(1/\varepsilon). \end{aligned}$$

In particular, on any sequence of such graphs the dynamics has cutoff with an $O(1)$ -window around t_m .

Apart from giving a first proof of cutoff for the Ising model on any tree / expander graph at $\beta > 0$, note that the above theorem allows the maximum degree d to depend on n in any way, and so it applies, e.g., to the Ising model on the hypercube (with $d = \log_2 n$), a dense Erdős-Rényi graph $\mathcal{G}(n, \frac{1}{2})$, etc.

As mentioned above, the proof uses the new information percolation framework, which analyzes interactions between spins viewed as a percolation process in the space-time slab. As opposed to the application of this method in the companion paper [23] for the torus, various obstacles arise in the present setting due to the asymmetry between vertices and lack of amenability. Moreover, a naïve application of the method would require β to be as small as about d^{-d} , and carrying it up to κ/d (the correct dependence in d up to the value of κ) required several novel ingredients, notably using a discrete Fourier expansion (see §4.2) to prescribe update rules for the dynamics that would endow the resulting percolation clusters with a subcritical behavior.

Roughly put, the framework considers the dynamics at a designated time around t_m , and for each site develops the history of updates that led to its final spin (tracing back branching to its neighbors). The resulting “information percolation” clusters in the space-time slab are then categorized into three types — RED (those surviving to time zero and nontrivially depending on the initial state), BLUE (those remaining which involve a unique “ancestor”) and GREEN (all remaining clusters), as illustrated in Figure 1. The green clusters (which may exhibit complicated dependencies but are independent of the initial state) are taken out of the equation via conditioning, leaving behind a competition between blue clusters (whose ancestor vertices are i.i.d. uniform spins by symmetry) and red clusters. Controlling the latter, namely an exponential moment of their cumulative size, then establishes mixing.

Overall, the information percolation framework allows one to reduce challenging problems involving mixing and cutoff for the Ising model into simpler and tractable problems on subcritical percolation.

Furthermore, by analyzing not only on the size of the red clusters, but rather *where* these hit the initial state at time zero, this framework opens the door to understanding the effect of the starting configuration on the mixing time (where sharp results on total-variation mixing for the Ising model were only applicable to worst-case starting states, usually via coupling techniques).

Our next result demonstrates this by comparing the worst-case mixing time (which is matched by the all-plus starting state up to an additive $O(1)$ -term) with a typical starting configuration, and finally with the uniform starting configuration, i.e., each site is initialized by an independent uniform ± 1 spin. Informally, we show that the uniform starting state is roughly at least twice faster compared to all-plus, but perhaps surprisingly, almost every deterministic starting state is about as slow as the worst one.

Formally, if $\mu_t^{(x_0)}$ is the distribution of the dynamics at time t started from x_0 then $t_{\text{MIX}}^{(x_0)}(\varepsilon)$ is the minimal t for which $\mu_t^{(x_0)}$ is within distance ε from equilibrium, and $t_{\text{MIX}}^{(u)}(\varepsilon)$ is the analogue for the average $2^{-n} \sum_{x_0} \mu_t^{(x_0)}$ (i.e., the *annealed* version, as opposed to the *quenched* $t_{\text{MIX}}^{(X_0)}$ for a uniform X_0).

Theorem 2. *Consider continuous-time Glauber dynamics for the Ising model on an n -vertex graph G with maximum degree at most some fixed $d > 0$, and define t_m as in (1.1). For every $\varepsilon > 0$ there exists $\beta_0 > 0$ such that the following hold for any $0 < \beta < \beta_0$ and any fixed $0 < \alpha < 1$ at large enough n .*

1. (Annealed) Uniform initial state: $t_{\text{MIX}}^{(u)}(\alpha) \leq (\frac{1}{2} + \varepsilon)t_m$.
2. (Quenched) Deterministic initial state: $t_{\text{MIX}}^{(x_0)}(\alpha) \geq (1 - \varepsilon)t_m$ for almost every x_0 , while $t_{\text{MIX}}^{(+)}(\alpha) \sim t_m$.

The delicate part in the proof of the above theorem is comparing the distribution at time t directly to the Ising measure. One often bypasses this point by coupling the distributions started at worst-case states; here, however, that would fail as we are analyzing the dynamics well before these distributions can couple with high probability. Instead (and as demonstrated in the companion paper for analyzing the effect of initial states in the 1D Ising model), we appeal to the *Coupling From The Past* method [29].

Rather than developing the information percolation clusters until reaching time zero, we continue until time $-\infty$, letting all clusters eventually die. The beautiful Coupling From The Past argument implies that, if we ignore the initial state altogether, the final configuration would be a perfect simulation of the Ising measure. Thus, the natural coupling the information percolation clusters allows one to compare the dynamics with the Ising measure, simply by considering the effect of replacing the spins generated along the interval $(-\infty, 0]$ by those of the initial state.

Specifically for the annealed analysis, even if a cluster survives to time zero (and beyond) it might still be perfectly coupled to the stationary measure, e.g., a singleton strand (and more generally, a blue cluster) would receive a uniform spin both from the Ising measure and from the random initial state. Hence, we modify the framework by redefining red clusters as those in which *at least two branches* of the cluster reach time zero, then proceed to merge in the interval $(-\infty, 0)$, as illustrated in Figure 2. It is this factor of 2 that eventually transforms into the factor of $2 - \varepsilon$ improvement in the mixing time.

Organization. The rest of this paper is organized as follows. In §2 we give the formal definitions of the above described framework, including several modification needed here (e.g., custom update rules to be derived from a Fourier expansion) and two lemmas analyzing the information percolation clusters. In §3 we prove the cutoff result in Theorem 1 modulo these technical lemmas, which are proved in §4. The final section, §5, is devoted to the effect of the initial states on mixing and the proof of Theorem 2.

2. INFORMATION PERCOLATION FOR THE ISING MODEL

2.1. Preliminaries. In what follows we set up standard notation for analyzing the mixing of Glauber dynamics for the Ising model; see [21, 23] and the references therein for additional information.

Mixing time and cutoff. Let (X_t) be an ergodic finite Markov chain with stationary measure π . An important gauge in MCMC theory for measuring the convergence of a Markov chain to stationarity is its total-variation mixing time. Denoted $t_{\text{MIX}}(\varepsilon)$ for a precision parameter $0 < \varepsilon < 1$, it is defined as

$$t_{\text{MIX}}(\varepsilon) \triangleq \inf \left\{ t : \max_{x_0 \in \Omega} \|\mathbb{P}_{x_0}(X_t \in \cdot) - \pi\|_{\text{TV}} \leq \varepsilon \right\},$$

where here and in what follows \mathbb{P}_{x_0} denotes the probability given $X_0 = x_0$, and the total-variation distance $\|\cdot\|_{\text{TV}}$ between two probability measures ν_1, ν_2 on a finite space Ω is given by

$$\|\nu_1 - \nu_2\|_{\text{TV}} = \max_{A \subset \Omega} |\nu_1(A) - \nu_2(A)| = \frac{1}{2} \sum_{\sigma \in \Omega} |\nu_1(\sigma) - \nu_2(\sigma)|,$$

i.e., half the L^1 -distance between the two measures.

Addressing the role of the parameter ε , the cutoff phenomenon is essentially the case where the choice of any fixed ε does not affect the asymptotics of $t_{\text{MIX}}(\varepsilon)$ as the system size tends to infinity. Formally, a family of ergodic finite Markov chains (X_t) , indexed by an implicit parameter n , is said to exhibit *cutoff* (a concept going back to the pioneering works [2, 9]) iff the following sharp transition in its convergence to stationarity occurs:

$$\lim_{n \rightarrow \infty} \frac{t_{\text{MIX}}(\varepsilon)}{t_{\text{MIX}}(1 - \varepsilon)} = 1 \quad \text{for any } 0 < \varepsilon < 1. \quad (2.1)$$

That is, $t_{\text{MIX}}(\alpha) = (1 + o(1))t_{\text{MIX}}(\beta)$ for any fixed $0 < \alpha < \beta < 1$. The *cutoff window* addresses the rate of convergence in (2.1): a sequence $w_n = o(t_{\text{MIX}}(e^{-1}))$ is a cutoff window if $t_{\text{MIX}}(\varepsilon) = t_{\text{MIX}}(1 - \varepsilon) + O(w_n)$ holds for any $0 < \varepsilon < 1$ with an implicit constant that may depend on ε . Equivalently, if t_n and w_n are sequences with $w_n = o(t_n)$, we say that a sequence of chains exhibits cutoff at t_n with window w_n if

$$\begin{cases} \lim_{\gamma \rightarrow \infty} \liminf_{n \rightarrow \infty} \max_{x_0 \in \Omega} \|\mathbb{P}_{x_0}(X_{t_n - \gamma w_n} \in \cdot) - \pi\|_{\text{TV}} = 1, \\ \lim_{\gamma \rightarrow \infty} \limsup_{n \rightarrow \infty} \max_{x_0 \in \Omega} \|\mathbb{P}_{x_0}(X_{t_n + \gamma w_n} \in \cdot) - \pi\|_{\text{TV}} = 0. \end{cases}$$

Verifying cutoff is often quite challenging, e.g., even for simple random walk on an expander graph, no examples were known prior to [19, 20] (while this had been conjectured for almost all such graphs), and to date there is no known transitive example (while conjectured to hold for all transitive expanders).

Glauber dynamics for the Ising model. Let G be a finite graph G with vertex-set V and edge-set E . The Ising model on G is a distribution over the set $\Omega = \{\pm 1\}^V$ of possible configurations, each corresponding to an assignment of plus/minus spins to the sites in V . The probability of $\sigma \in \Omega$ is given by

$$\pi(\sigma) = Z^{-1} e^{\beta \sum_{uv \in E} \sigma(u)\sigma(v)}, \quad (2.2)$$

where the normalizer $Z = Z(\beta, h)$ is the partition function. The parameter β is the inverse-temperature, which we always take to be non-negative (ferromagnetic). These definitions extend to infinite locally finite graphs (see, e.g., [18, 24]).

The Glauber dynamics for the Ising model (the *Stochastic Ising* model) is a family of continuous-time Markov chains on the state space Ω , reversible w.r.t. the Ising measure π , given by the generator

$$(\mathcal{L}f)(\sigma) = \sum_u c(u, \sigma) (f(\sigma^u) - f(\sigma)) \quad (2.3)$$

where σ^u for $u \in V$ is the configuration σ with the spin at the vertex u flipped. We will focus on the two most notable examples of Glauber dynamics, each having an intuitive and useful graphical interpretation where each site experiences updates via an associated i.i.d. rate-one Poisson clock:

- (i) *Metropolis*: flip $\sigma(u)$ if the new state σ^u has a lower energy (i.e., $\pi(\sigma^u) \geq \pi(\sigma)$), otherwise perform the flip with probability $\pi(\sigma^u)/\pi(\sigma)$. This corresponds to $c(u, \sigma) = \exp(2\beta\sigma(u) \sum_{v \sim u} \sigma(v)) \wedge 1$.
- (ii) *Heat-bath*: erase $\sigma(u)$ and replace it with a sample from the conditional distribution given the spins at its neighboring sites. This corresponds to $c(u, \sigma) = 1/[1 + \exp(-2\beta\sigma(u) \sum_{v \sim u} \sigma(v))]$.

It is easy to verify that these chains are indeed ergodic and reversible w.r.t. the Ising distribution π . Until recently, sharp mixing results for this dynamics were obtained in relatively few cases, with cutoff only known for the complete graph [11, 16] prior to the works [21, 22].

2.2. Red, green and blue information percolation clusters. In what follows, we describe the basic setting of the framework, which will be enhanced in §2.3 to support the setting of Theorem 1 (where the underlying geometry may feature exponential growth rate and we are in the range $\beta < \kappa/d$).

The *update sequence* of the Glauber dynamics along an interval $(t_0, t_1]$ is the set of tuples of the form (J, U, τ) , where $t_0 < \tau \leq t_1$ is the update time, $J \in V$ is the site to be updated and U is a uniform unit variable. Given this update sequence, X_{t_1} is a deterministic function of X_{t_0} , right-continuous w.r.t. t_1 .

We call a given update (J, U, τ) an *oblivious update* iff $U \leq \theta$ for

$$\theta = \theta_{\beta, d} := 1 - \tanh(\beta d), \quad (2.4)$$

since in that situation one can update the spin at J to plus/minus with equal probability (that is, with probability $\theta/2$ each) independently of the spins at the neighbors of the vertex J , and a properly chosen rule for the case $U > \theta$ legally extends this protocol to the Glauber dynamics.

Consider some designated target time t_* for analyzing the spin distribution of the dynamics on G . The *update history* of $X_{t_*}(v)$ going back to time t , denoted $\mathcal{H}_v(t)$, is a subset $A \times \{t\}$ of the space-time slab $V \times \{t\}$, such that one we can determine $X_{t_*}(v)$ from the update sequence and spin-set $X_t(A)$. The most basic way of defining $\{\mathcal{H}_v(t) : 0 \leq t \leq t_*\}$ is as follows:

- List the updates in reverse chronological order as $\{(J_i, U_i, t_i)\}_{i \geq 1}$ (i.e., $t_i > t_{i+1}$ for all i), and initialize the update history by $\mathcal{H}_v(t) = \{v\}$ for all $t \in [t_1, t_*]$.
- In step $i \geq 1$, process the update (J_i, U_i, t_i) to determine $\mathcal{H}_v(t)$ for $t \in [t_{i+1}, t_i]$:
 - If $J_i \notin \mathcal{H}_v(t_i)$ then the history is unchanged, i.e., $\mathcal{H}_v(t) = \mathcal{H}_v(t_i)$ for all $t \in [t_{i+1}, t_i]$.
 - If $J_i \in \mathcal{H}_v(t_i)$ but $U_i \leq \theta$ then J_i is removed, i.e., $\mathcal{H}_v(t) = \mathcal{H}_v(t_i) \setminus \{J_i\}$ for all $t \in [t_{i+1}, t_i]$.
 - Otherwise, replace J_i by its neighbors $N(J_i)$, i.e., $\mathcal{H}_v(t) = \mathcal{H}_v(t_i) \cup N(J_i) \setminus \{J_i\}$ for all $t \in [t_{i+1}, t_i]$.

The *information percolation clusters* are the connected components of the graph on the vertex set V where (u, v) is an edge if $\mathcal{H}_u(t) \cap \mathcal{H}_v(t) \neq \emptyset$ for some $t \geq 0$. Denote by \mathcal{C}_v the cluster containing $v \in V$.

We will also consider clusters in the context of the full space-time slab. The cluster of a point $(w, r) \in V \times [0, t_*]$, denoted $\mathcal{X}_{w, r}$, is the connected component of $\bigcup\{\mathcal{H}_v(t) : v \in V, 0 \leq t \leq t_*\}$ that contains (w, r) . (Thus, the cluster \mathcal{C}_v is identified with the intersection of \mathcal{X}_{v, t_*} with the slab $V \times \{t_*\}$.)

For any $A \subset V$ we use the notation $\mathcal{H}_A(t) = \bigcup_{v \in A} \mathcal{H}_v(t)$, as well as $\mathcal{H}_A(t_1, t_2) = \bigcup_{t_1 \leq t \leq t_2} \mathcal{H}_A(t)$ (both cases describing subsets of V). Omitting the time subscript altogether would refer to the full time interval: $\mathcal{H}_A := \mathcal{H}_A(0, t_\star)$, so that, for instance, if $\mathcal{C} \subset V$ is a cluster then $\mathcal{H}_{\mathcal{C}}$ is the set of all vertices ever visited by this cluster. (By a slight abuse of notation, we may write $\mathcal{H}_{\mathcal{X}} := \mathcal{H}_{\mathcal{C}}$ with \mathcal{C} the cluster that \mathcal{X} identifies with (i.e., the cluster \mathcal{C} such that $\mathcal{C} \times \{t_\star\} = (V \times \{t_\star\}) \cap \mathcal{X}$.) A final useful notation in this context is the collective history of $V \setminus A$, defined as

$$\mathcal{H}_A^- = \{\mathcal{H}_v(t) : v \notin A, t \leq t_\star\}.$$

The clusters are classified into three classes (identifying for this purpose \mathcal{C}_v and \mathcal{X}_{v, t_\star}) as follows:

- A cluster \mathcal{C} is RED if, given the update sequence, its final state $X_{t_\star}(\mathcal{C})$ is a nontrivial function of the initial configuration X_0 ; in particular, its history must survive to time zero ($\mathcal{H}_{\mathcal{C}}(0) \neq \emptyset$).
- A cluster \mathcal{C} is BLUE if it is a singleton — i.e., $\mathcal{C} = \{v\}$ for some $v \in V$ — whose history does not survive to time zero ($\mathcal{H}_v(0) = \emptyset$).
- Every other cluster \mathcal{C} is GREEN.

Note that if a cluster is blue then its single spin at time t_\star does not depend on the initial state X_0 , and so, by symmetry, it is a uniform ± 1 spin. (While a green cluster is similarly independent of X_0 , as multiple update histories intersect, the distribution of its spin set $X_{t_\star}(\mathcal{C})$ may become quite nontrivial.)

Let V_{RED} denote the union of the red clusters, and let \mathcal{H}_{RED} be the its collective history — the union of $\mathcal{H}_v(t)$ for all $v \in V_{\text{RED}}$ and $0 \leq t \leq t_\star$ (with analogous definitions for blue/green).

A beautiful short lemma of Miller and Peres [28] shows that, if a measure μ on $\{\pm 1\}^V$ is given by sampling a variable $R \subset V$ and using an arbitrary law for its spins and a product of Bernoulli($\frac{1}{2}$) for $V \setminus R$, then the L^2 -distance of μ from the uniform measure is at most $\mathbb{E}2^{|R \cap R'|} - 1$ for i.i.d. copies R, R' . (See Lemma 3.1 below; also see [23, Lemma 4.3] for a generalization of this to a product of general measures, which becomes imperative for the information percolation framework at β near criticality.) Applied to our setting, if we condition on $\mathcal{H}_{\text{GREEN}}$ and look at the spins of $V \setminus V_{\text{GREEN}}$ then V_{RED} can assume the role of the variable R , as the remaining blue clusters are a product of Bernoulli($\frac{1}{2}$) variables.

In this conditional space, since the law of the spins of V_{GREEN} , albeit potentially complicated, is independent of the initial state, we can safely project the configurations on $V \setminus V_{\text{GREEN}}$ without it increasing the total-variation distance between the distributions started at the two extreme states. Hence, a sharp upper bound on worst-case mixing will follow by showing for this exponential moment

$$\mathbb{E} \left[2^{|V_{\text{RED}} \cap V'_{\text{RED}}|} \mid \mathcal{H}_{\text{GREEN}} \right] \rightarrow 1 \quad \text{in probability as } n \rightarrow \infty, \quad (2.5)$$

by coupling the distribution of the dynamics at time t_\star from any initial state to the uniform measure. Finally, with the green clusters out of the picture by the conditioning (which has its own toll, forcing various updates along history so that no other cluster would intersect with those nor become green), we can bound the probability that a subset of sites would become a red cluster by its ratio with the probability of all sites being blue clusters. Being red entails connecting the subset in the space-time slab, hence the exponential decay needed for (2.5).

2.3. Enhancements of the framework: custom update rules and modified last unit interval.

We will consider the information percolation clusters developed as above from the designated time

$$t_\star = t_m + s_\star \quad \text{for} \quad s_\star = C \log(1/\varepsilon)$$

where $C > 0$ will be specified later, and $\varepsilon > 0$ is the parameter for the mixing time. However, instead of the standard procedure of developing the history, where an update at v either deletes it from the history (via an oblivious update) or replaces it by its set of neighbors $N(v)$, we will allow v to be replaced (with varying probabilities) by any subset of its neighbors, in the following way.

Recall that an update of the form $(J, U, t) \in V \times [0, 1] \times [0, t_\star]$ results in replacing the spin at J at time t by some deterministic function $\Upsilon(x, U)$, where $x = \sum_{u \in N(J)} X_t(u)$. A *generalized* update rule observes updates of the form (J, A, U, t) where (J, U, t) is as before and the additional variable $A \subset [d]$ corresponds to a subset of the neighbors of vertex J . The new update rule exposes the spins $\{\sigma_1, \dots, \sigma_{|A|}\}$ of these neighbors at time t , then generates the new spin at J via $\Phi_A(\sigma_1, \dots, \sigma_{|A|}, U)$.

With this generalized update rule, one unfolds the update history of a vertex $\{\mathcal{H}_v(t) : 0 \leq t \leq t_\star\}$ as before, with the one difference that an update (J_i, A_i, U_i, t_i) for which $J_i \in \mathcal{H}_v(t_i)$ now results in $\mathcal{H}_v(t) = \mathcal{H}_v(t_i) \cup A_i \setminus \{J_i\}$ for all $t \in [t_{i+1}, t_i)$. The functions $\{\Phi_A : A \subset [d]\}$, as well as the probability distribution over the subsets $A \subset [d]$ to be exposed, will be derived from a discrete Fourier expansion of the original rule Υ (see Lemma 4.1), so that the new update procedure would, one on hand, couple with the Glauber dynamics, and on the other, endow our percolation clusters with a subcritical behavior.

A final ingredient needed for coping with the arbitrary underlying geometry is a modification of the update history, denoted by $\hat{\mathcal{H}}$: in the modified version, every vertex $v \in V$ receives an (extra) update at time t_\star , and no vertex is removed from the history along the unit interval $(t_\star - 1, t_\star]$. (For a given update sequence, this operation can only increase any information percolation cluster, and forbidding vertices to die in the first unit interval will be useful in the context of conditioning on other clusters.) We will write $\hat{\mathcal{C}}, \hat{\mathcal{X}}$, as well as $\hat{\mathcal{H}}_A(t)$ etc. for the corresponding notation w.r.t. the modified history $\hat{\mathcal{H}}$.

We end this section with two results on the information percolation clusters — Lemmas 2.1 and 2.2 — which will be central in the proof of Theorem 1. The proofs of these lemmas are postponed to §4.

As explained following the definition of the three cluster types, at the heart of the matter is estimating an exponential moment of the size of the red clusters given $\mathcal{H}_{\text{GREEN}}$, the joint history of all green clusters. To this end, we wish to bound the probability that a subset A is a red cluster given $\mathcal{H}_{\text{GREEN}}$. Define

$$\Psi_A = \sup_{\mathcal{H}_A^-} \mathbb{P}(A \in \text{RED} \mid \mathcal{H}_A^-, \{A \in \text{RED}\} \cup \{A \subset V_{\text{BLUE}}\}), \quad (2.6)$$

noting that, towards estimating the probability of $A \in \text{RED}$, the effect of conditioning on \mathcal{H}_A^- amounts to requiring that \mathcal{H}_A must not intersect \mathcal{H}_A^- .

Lemma 2.1. *If $\beta < 1/(5d)$ then for any $A \subset V$ and $v \in A$,*

$$\Psi_A \leq 2^{|A|} \mathbb{E} \left[\mathbb{1}_{\{A \subset \hat{\mathcal{C}}_v\}} e^{\hat{\tau}_v} \sum_w \mathbb{1}_{\{w \in \hat{\mathcal{H}}_A(t_\star - \hat{\tau}_v, t_\star)\}} \mathbf{m}_{t_\star}(w) \right]$$

where $\hat{\tau}_v$ is the time it takes the history of $\hat{\mathcal{C}}_v$ to first coalesce into a single point (if at all), i.e.,

$$\hat{\tau}_v = \min \left\{ t \geq 1 : |\hat{\mathcal{H}}_{\hat{\mathcal{C}}_v}(t_\star - t)| = 1 \right\} \wedge t_\star. \quad (2.7)$$

It is worthwhile noting in the context of the parameter $\hat{\tau}_v$ that, when developing the update history backward in time, $\hat{\tau}_v$ is not a stopping time, since $\hat{\mathcal{C}}_v$ is affected by any potential coalescence points for $t < t_\star - \hat{\tau}_v$; instead, one can determine $\hat{\tau}_v$ as soon as $\hat{\mathcal{H}}_{\hat{\mathcal{C}}_v}(t) = \emptyset$. Also observe that $\hat{\tau}_v = 1$ iff $|\mathcal{C}_v| = 1$. Finally, the coalescence point w at time $t = t_\star - \hat{\tau}_v$ (when $t > 0$) need not belong to $\hat{\mathcal{H}}_v$ — e.g., we may have $\hat{\mathcal{H}}_v(t) = \emptyset$ while $w \in \hat{\mathcal{H}}_u$ for some $u \neq v$ whose history intersected that of v at time $t' > t$.

The subcritical nature of the information percolation clusters (prompted by our modified update functions Φ_A) allows one to control exponential moments of the cluster sizes, as in the following lemma.

Lemma 2.2. *Fix $0 < \eta < 1$ and $\lambda > 0$. There exist constants $\kappa, \gamma > 0$ such that the following holds. For any point (w_0, t_0) in the space-time slab $V \times (0, t_\star]$, if $\beta < \kappa/d$ then*

$$\mathbb{E} \left[\exp \left(\eta \mathfrak{L}(\hat{\mathcal{X}}_{w_0, t_0}) + \lambda |\hat{\mathcal{H}}_{\hat{\mathcal{X}}_{w_0, t_0}}| \right) \right] < \gamma,$$

where

$$\mathfrak{L}(\hat{\mathcal{X}}) = \sum_{u \in V} \int_0^{t_\star} \mathbb{1}_{\{(u, t) \in \hat{\mathcal{X}}\}} dt.$$

The above lemma, whose proof follows standard arguments from percolation theory, will be applied for absolute constants η and λ in the proof of Theorem 1 (any $1/2 < \eta < 1$ and $\lambda > \log 8$ would do), leading to the absolute constant κ in the statement of that theorem. The above formulation will be important in the context of Theorem 2, where one requires η that may be very close to 1 (as a function of ε from the statement of that theorem) and λ that depends on the maximum degree.

3. CUTOFF WITH CONSTANT WINDOW FROM A WORST STARTING STATE

In this section we prove Theorem 1 via the framework defined in §2. As is often the case in proofs of cutoff, the upper bound will require the lion's share of the efforts.

3.1. Upper bound modulo Lemmas 2.1 and 2.2. Define the coupling distance $\bar{d}_{\text{TV}}(t)$ to be

$$\bar{d}_{\text{TV}}(t) = \max_{x_0, y_0} \|\mathbb{P}_{x_0}(X_t \in \cdot) - \mathbb{P}_{y_0}(X_t \in \cdot)\|_{\text{TV}}$$

(so that $\frac{1}{2}\bar{d}_{\text{TV}}(t) \leq d_{\text{TV}}(t) \leq \bar{d}_{\text{TV}}(t)$), and observe that

$$\begin{aligned} \bar{d}_{\text{TV}}(t) &\leq \mathbb{E} \left[\max_{x_0, y_0} \|\mathbb{P}_{x_0}(X_t \in \cdot \mid \mathcal{H}_{\text{GREEN}}) - \mathbb{P}_{y_0}(X_t \in \cdot \mid \mathcal{H}_{\text{GREEN}})\|_{\text{TV}} \right] \\ &\leq \sup_{\mathcal{H}_{\text{GREEN}}} \max_{x_0, y_0} \|\mathbb{P}_{x_0}(X_t(V \setminus V_{\text{GREEN}}) \in \cdot \mid \mathcal{H}_{\text{GREEN}}) - \mathbb{P}_{y_0}(X_t(V \setminus V_{\text{GREEN}}) \in \cdot \mid \mathcal{H}_{\text{GREEN}})\|_{\text{TV}}, \end{aligned}$$

where the first inequality follows by Jensen's Inequality and the second follows since $X_t(V_{\text{GREEN}})$ is independent of the initial condition and so taking a projection onto $V \setminus V_{\text{GREEN}}$ does not change the total-variation distance between the distributions started at x_0 and y_0 . Thus,

$$\bar{d}_{\text{TV}}(t) \leq 2 \sup_{\mathcal{H}_{\text{GREEN}}} \max_{x_0} \|\mathbb{P}_{x_0}(X_t(V \setminus V_{\text{GREEN}}) \in \cdot \mid \mathcal{H}_{\text{GREEN}}) - \nu_{V \setminus V_{\text{GREEN}}}\|_{\text{TV}}, \quad (3.1)$$

where ν_A is the uniform measure on configurations on the sites in A . At this point we appeal to the exponential-moment bound of [28], whose short proof is included here for completeness.

Lemma 3.1 ([28]). *Let $\Omega = \{\pm 1\}^V$ for a finite set V . For each $S \subset V$, let φ_S be a measure on $\{\pm 1\}^S$. Let ν be the uniform measure on Ω , and let μ be the measure on Ω obtained by sampling a subset $S \subset V$ via some measure $\tilde{\mu}$, generating the spins of S via φ_S , and finally sampling $V \setminus S$ uniformly. Then*

$$\|\mu - \nu\|_{L^2(\nu)}^2 \leq \mathbb{E} \left[2^{|S \cap S'|} \right] - 1,$$

where the variables S and S' are i.i.d. with law $\tilde{\mu}$.

Proof. Write $n = |V|$, and let x_S ($S \subset V$) denote the projection of x onto S . With this notation, by definition of the $L^2(\nu)$ metric (see, e.g., [31]) one has that $\|\mu - \nu\|_{L^2(\nu)}^2 + 1 = \int |\mu/\nu - 1|^2 d\nu + 1$ equals

$$\sum_{x \in \Omega} \frac{\mu^2(x)}{\nu(x)} = 2^n \sum_{x \in \Omega} \sum_S \tilde{\mu}(S) \frac{\varphi_S(x_S)}{2^{n-|S|}} \sum_{S'} \tilde{\mu}(S') \frac{\varphi_{S'}(x_{S'})}{2^{n-|S'|}}$$

by the definition of μ . Since $\sum_x \varphi_S(x_S) \varphi_{S'}(x_{S'}) \leq 2^{n-|S \cup S'|}$ it then follows that

$$\sum_{x \in \Omega} \frac{\mu^2(x)}{\nu(x)} \leq \sum_{S, S'} 2^{|S|+|S'|-|S \cup S'|} \tilde{\mu}(S) \tilde{\mu}(S') = \sum_{S, S'} 2^{|S \cap S'|} \tilde{\mu}(S) \tilde{\mu}(S'). \quad \blacksquare$$

Remark 3.2. In the special case where the distribution φ_S is a point-mass on all-plus for every S , the single inequality in the above proof is an equality (since then $\sum_x \varphi_S(x_S) \varphi_{S'}(x_{S'}) = \#\{x : x_{S \cup S'} \equiv 1\}$) and so in that situation the L^2 -distance $\|\mu - \nu\|_{L^2(\nu)}^2$ is precisely equal to $\mathbb{E}[2^{|S \cap S'|}] - 1$.

For example, consider Glauber dynamics for an n -vertex graph at $\beta = 0$ (i.e., continuous-time lazy random walk on the hypercube $\{\pm 1\}^n$) starting (say) from all-plus, and let S be the set of coordinates which were not updated: here $\mathbb{P}(v \in S) = e^{-t}$ at time t , and $\|\mathbb{P}(X_t^+ \in \cdot) - \nu\|_{L^2(\nu)}^2 = (1 + e^{-2t})^n - 1$.

Applying the above lemma to the right-hand side of (3.1), while recalling that any two measures μ and ν on a finite probability space satisfy $\|\mu - \nu\|_{\text{TV}} = \frac{1}{2} \|\mu - \nu\|_{L^1(\nu)} \leq \frac{1}{2} \|\mu - \nu\|_{L^2(\nu)}$, we find that

$$\bar{d}_{\text{TV}}(t_*) \leq \left(\sup_{\mathcal{H}_{\text{GREEN}}} \mathbb{E} \left[2^{|V_{\text{RED}} \cap V_{\text{RED}}'|} \mid \mathcal{H}_{\text{GREEN}} \right] - 1 \right)^{1/2}, \quad (3.2)$$

where V_{RED} and V_{RED}' are i.i.d. copies of the variable $\bigcup\{v \in V : \mathcal{C}_v \in \text{RED}\}$.

Let $\{Y_{A, A'} : A, A' \subset V\}$ be a family of independent indicators satisfying

$$\mathbb{P}(Y_{A, A'} = 1) = \Psi_A \Psi_{A'} \quad \text{for any } A, A' \subset V. \quad (3.3)$$

We claim that it is possible to couple the conditional distribution of $(V_{\text{RED}}, V_{\text{RED}}')$ given $\mathcal{H}_{\text{GREEN}}$ to the variables $Y_{A, A'}$ in such a way that

$$|V_{\text{RED}} \cap V_{\text{RED}}'| \leq \sum_{A \cap A' \neq \emptyset} |A \cup A'| Y_{A, A'}.$$

To do so, let $\{(A_l, A'_l)\}_{l \geq 1}$ denote all pairs of intersecting subsets $(A, A' \subset V \setminus V_{\text{GREEN}}$ with $A \cap A' \neq \emptyset$) arbitrarily ordered, associate each pair with a variable R_l initially set to 0, then process these in order:

- If (A_l, A'_l) is such that, for some $j < l$, one has $R_j = 1$ and either $A_j \cap A_l \neq \emptyset$ or $A'_j \cap A'_l \neq \emptyset$, then skip this pair (keeping $R_l = 0$).
- Otherwise, set R_l to the indicator of $\{A_l \in \text{RED}, A'_l \in \text{RED}'\}$.

The claim is that $\mathbb{P}(R_l = 1 \mid \mathcal{F}_{l-1}) \leq \mathbb{P}(Y_{A_l, A'_l} = 1)$ for all l , where \mathcal{F}_l denotes the natural filtration associated to the above process. Indeed, consider some (A_l, A'_l) for which we are about to set R_l to the value of $\mathbb{1}_{\{A_l \in \text{RED}, A'_l \in \text{RED}'\}}$, and take any A_j ($j < l$) such that $A_j \cap A_l \neq \emptyset$ and $\mathbb{1}_{\{C_j \in \text{RED}, C'_j \in \text{RED}'\}}$ was revealed (and necessarily found to be zero, by definition of the above process). The supremum over $\mathcal{H}_{A_l}^-$ in the definition of Ψ_{A_l} implies that we need only consider the information \mathcal{F}_{l-1} offers on \mathcal{H}_{A_l} :

- If $A_j \cap A_l \neq A_l$ then the event $\{A_j \in \text{RED}\}$ does not intersect the event $\{A_l \in \text{RED}\} \cup \{A_l \subset V_{\text{BLUE}}\}$ (on which we condition in Ψ_{A_l}) as it requires A_j to be a full red cluster (so a strict subset of A_j cannot belong to a separate red cluster, nor can it contain any blue singleton).
- If $A_j = A_l$, conditioning on $\{A_j \in \text{RED}, A'_j \in \text{RED}'\}^c$ will not increase the probability of $\{A_l \in \text{RED}\}$. Either way, $\mathbb{P}(A_l \in \text{RED} \mid \mathcal{F}_{l-1}) \leq \Psi_{A_l}$. Similarly, $\mathbb{P}(A'_l \in \text{RED}' \mid \mathcal{F}_{l-1}, \mathbb{1}_{\{A_l \in \text{RED}\}}) \leq \Psi_{A'_l}$, and together these inequalities support the desired coupling, since if $v \in V_{\text{RED}} \cap V_{\text{RED}'}$ then there is some l for which $v \in A_l \cup A'_l$ and $A_l \in \text{RED}, A'_l \in \text{RED}'$, in which case every A_j intersecting A_l nontrivially will receive $R_j = 0$ (it cannot be red) and the first j with $A_j = A_l$ to receive $R_j = 1$ will account for v in $A_j \cup A'_j$.

Relaxing $|A \cup A'|$ into $|A| + |A'|$ (which will be convenient for factorization), we get

$$\sup_{\mathcal{H}_{\text{GREEN}}} \mathbb{E} \left[2^{|V_{\text{RED}} \cap V_{\text{RED}'|} \mid \mathcal{H}_{\text{GREEN}}} \right] \leq \mathbb{E} \left[2^{\sum_{A \cap A' \neq \emptyset} (|A| + |A'|) Y_{A, A'}} \right] = \prod_{A \cap A' \neq \emptyset} \mathbb{E} \left[2^{(|A| + |A'|) Y_{A, A'}} \right],$$

with the equality due to the independence of the $Y_{A, A'}$'s. By the definition of these indicators in (3.3), this last expression is at most

$$\prod_v \prod_{\substack{A, A' \\ v \in A \cap A'}} \left((2^{|A| + |A'|} - 1) \Psi_A \Psi_{A'} + 1 \right) \leq \exp \left[\sum_v \left(\sum_{A \ni v} 2^{|A|} \Psi_A \right)^2 \right],$$

and so, revisiting (3.2), we conclude that

$$\bar{d}_{\text{TV}}(t_\star)^2 \leq \left(\exp \left[\sum_v \left(\sum_{A \ni v} 2^{|A|} \Psi_A \right)^2 \right] - 1 \right) \wedge 1 \leq 2 \sum_v \left(\sum_{A \ni v} 2^{|A|} \Psi_A \right)^2, \quad (3.4)$$

where we used that $e^x - 1 \leq 2x$ for $x \in [0, 1]$. We have thus reduced the upper bound in Theorem 1 into showing that the right-hand of (3.4) is at most ε if $s_\star = C \log(1/\varepsilon)$ for some large enough $C = C(\beta)$.

Plugging the bound on Ψ_A from Lemma 2.1 shows that the sum in the right-hand of (3.4) is at most

$$\sum_v \left(\sum_{A \ni v} 4^{|A|} \mathbb{E} \left[\mathbb{1}_{\{A \subset \hat{C}_v\}} e^{\hat{\tau}_v} \sum_w \mathbb{1}_{\{w \in \hat{\mathcal{H}}_A(t_\star - \hat{\tau}_v, t_\star)\}} \mathbf{m}_{t_\star}(w) \right] \right)^2.$$

In each of the two sums over $A \ni v$ we can specify the size of \hat{C}_v , and then relax $\{w \in \hat{\mathcal{H}}_A(t_\star - \hat{\tau}_v, t_\star)\}$ into $\{w \in \hat{\mathcal{H}}_{\hat{C}_v}\}$ (thus permitting all $2^{|\hat{C}_v|}$ subsets to play the role of A); thus, the last display is at most

$$\sum_v \sum_{k, k'} \sum_{w, w'} 8^k \mathbb{E} \left[\mathbb{1} \left\{ |\hat{C}_v| = k, w \in \hat{\mathcal{H}}_{\hat{C}_v} \right\} e^{\hat{\tau}_v} \mathbf{m}_{t_\star}(w) \right] 8^{k'} \mathbb{E} \left[\mathbb{1} \left\{ |\hat{C}_v| = k', w' \in \hat{\mathcal{H}}_{\hat{C}_v} \right\} e^{\hat{\tau}_v} \mathbf{m}_{t_\star}(w') \right]. \quad (3.5)$$

Denoting the indicators above by $\Xi(v, w, k)$ and $\Xi(v, w', k')$ respectively, and using the fact that

$$\sum_{w, w'} \mathbf{m}_t(w) \mathbf{m}_t(w') \leq \frac{1}{2} \sum_{w, w'} (\mathbf{m}_t(w)^2 + \mathbf{m}_t(w')^2) = \sum_{w, w'} \mathbf{m}_t(w)^2$$

in (3.5) culminates in the following bound on sum in the right-hand of (3.4):

$$\sum_v \left(\sum_{A \ni v} 2^{|A|} \Psi_A \right)^2 \leq \sum_w \mathbf{m}_{t_\star}(w)^2 \sum_k \sum_v \mathbb{E} \left[8^k \Xi(v, w, k) e^{\hat{\tau}_v} \right] \sum_{k'} \mathbb{E} \left[e^{\hat{\tau}_v} \sum_{w'} 8^{k'} \Xi(v, w', k') \right]. \quad (3.6)$$

For the summation over k' in (3.6), we combine the facts that $\hat{\tau}_v \leq \frac{1}{2} \mathfrak{L}(\hat{\mathcal{H}}_{\hat{\mathcal{C}}_v}(t_\star - \hat{\tau}_v, t_\star)) + 1 \leq \frac{1}{2} \mathfrak{L}(\hat{\mathcal{H}}_{\hat{\mathcal{C}}_v}) + 1$ (either $|\hat{\mathcal{C}}_v| = 1$ and then $\hat{\tau}_v = 1$, or $|\hat{\mathcal{C}}_v| \geq 2$ whence at least two strands survive for a period of $\hat{\tau}_v$), that at most $|\hat{\mathcal{H}}_{\hat{\mathcal{C}}_v}|$ choices for w' support $\Xi(v, w', k) = 1$ and that $\sum_{k'} \Xi(v, w', k) \leq 1$, to get

$$\sum_{k'} \mathbb{E} \left[e^{\hat{\tau}_v} \sum_{w'} 8^{k'} \Xi(v, w', k') \right] \leq \mathbb{E} \left[|\hat{\mathcal{H}}_{\hat{\mathcal{C}}_v}| 8^{|\hat{\mathcal{H}}_{\hat{\mathcal{C}}_v}|} e^{\frac{1}{2} \mathfrak{L}(\hat{\mathcal{H}}_{\hat{\mathcal{C}}_v}) + 1} \right] \leq \gamma_1 \quad (3.7)$$

for some absolute constant $\gamma_1 > 0$, where the last inequality applied Lemma 2.2.

Next, to treat the summation over k in (3.6), recall that $\hat{\mathcal{X}}_{w,r}$ for $(w, r) \in V \times [0, t_\star]$ is the information percolation cluster containing the point (w, r) in the space-time slab (i.e., the cluster is exposed from time r instead of time t_\star and the process of developing it moves both forward and backward in time). Further write $\hat{\mathcal{X}}_{w,r}^+ = \lim_{t \rightarrow r^+} \hat{\mathcal{X}}_{w,t}$ and $\hat{\mathcal{X}}_{w,r}^- = \lim_{t \rightarrow r^-} \hat{\mathcal{X}}_{w,t}$.

We claim that whenever $\Xi(v, w, k) = 1$, necessarily $v \in \hat{\mathcal{X}}_{w,r}^-$ for some $r \in \Pi_w$, where Π_w records the update times for the vertex w (always including t_\star , by definition of $\hat{\mathcal{H}}$). Indeed, if $w \in \hat{\mathcal{H}}(t_\star - \hat{\tau}_v, t_\star)$ then by definition we can find some $q \in (t_\star - \hat{\tau}_v, t_\star)$ such that (w, q) shares the same information percolation cluster as (v, t_\star) . Furthermore, if r is the earliest update of w after time q then the cluster of (w, t) for any $t \in (q, r)$ will contain (w, q) , and thus (v, t_\star) as-well. (It is for this reason that we addressed $\hat{\mathcal{X}}_{w,r}^-$, in case the update at (w, r) should cut its information percolation cluster from (w, q) .) For that r , we further have $t_\star - r \leq \hat{\tau}_v \leq \frac{1}{2} \mathfrak{L}(\hat{\mathcal{X}}_{w,r}^-) + 1$, and so

$$\begin{aligned} \sum_k \sum_v \mathbb{E} \left[8^k \Xi(v, w, k) e^{\hat{\tau}_v} \right] &\leq \mathbb{E} \left[\sum_{r \in \Pi_w} |\hat{\mathcal{H}}_{\hat{\mathcal{X}}_{w,r}^-}| 8^{|\hat{\mathcal{H}}_{\hat{\mathcal{X}}_{w,r}^-}|} e^{\frac{1}{2} \mathfrak{L}(\hat{\mathcal{X}}_{w,r}^-)} \mathbb{1}_{\left\{ \frac{1}{2} \mathfrak{L}(\hat{\mathcal{X}}_{w,r}^-) \geq t_\star - r - 1 \right\}} \right] \\ &\leq \mathbb{E} \left[\sum_{r \in \Pi_w} |\hat{\mathcal{H}}_{\hat{\mathcal{X}}_{w,r}^-}| 8^{|\hat{\mathcal{H}}_{\hat{\mathcal{X}}_{w,r}^-}|} e^{\frac{3}{4} \mathfrak{L}(\hat{\mathcal{X}}_{w,r}^-) - \frac{1}{2}(t_\star - r - 1)} \right], \end{aligned}$$

which, recalling that Π_w is the union of $\{t_\star\}$ and a rate-1 Poisson process, is at most

$$\begin{aligned} &\mathbb{E} \left[|\hat{\mathcal{H}}_{\hat{\mathcal{C}}_w}| 8^{|\hat{\mathcal{H}}_{\hat{\mathcal{C}}_w}|} e^{\frac{3}{4} \mathfrak{L}(\hat{\mathcal{H}}_{\hat{\mathcal{C}}_w}) + \frac{1}{2}} \right] + \int_0^{t_\star} \mathbb{E} \left[\sum_{r \in \Pi_w} |\hat{\mathcal{H}}_{\hat{\mathcal{X}}_{w,r}^-}| 8^{|\hat{\mathcal{H}}_{\hat{\mathcal{X}}_{w,r}^-}|} e^{\frac{3}{4} \mathfrak{L}(\hat{\mathcal{X}}_{w,r}^-) - \frac{1}{2}(t_\star - r - 1)} \mid r \in \Pi_w \right] dr \\ &\leq \sqrt{e} \gamma_2 \left[1 + \int_0^{t_\star} e^{-\frac{1}{2}(t_\star - r)} dr \right] \leq 5 \gamma_2 \end{aligned}$$

for some absolute constant $\gamma_2 > 0$, using Lemma 2.2 (with γ_2 from that lemma) for the first inequality.

Substituting the last two displays together with (3.7) in (3.6), while recalling (3.4), finally gives

$$\bar{d}_{\text{TV}}(t_\star)^2 \leq 10\gamma_1\gamma_2 \sum_w \mathbf{m}_{t_\star}(w)^2. \quad (3.8)$$

The proof will be concluded with the help of the next simple claim that establishes a submultiplicative bound for the second moment of the magnetization.

Claim 3.3. *For any $t, s > 0$ we have*

$$e^{-2s} \leq \frac{\sum_w \mathbf{m}_{t+s}(w)^2}{\sum_w \mathbf{m}_t(w)^2} \leq e^{-2(1-\beta d)s}.$$

Proof. The lower bound follows from the straightforward fact that $\mathbf{m}_{t+s}(w) \geq e^{-s}\mathbf{m}_t(w)$ for any $s, t > 0$ and w , since the probability of observing no updates to w along the interval $(t, t+s)$ (thus maintaining the magnetization without a change) is e^{-s} . It therefore remains to prove the upper bound.

By expanding the probability of 1 in an update, which is $\frac{1}{2} + \frac{1}{2} \tanh(\beta\sigma)$ given a sum of neighbors of σ , and using the fact that $\frac{d}{dx} \tanh(x) \leq 1$ for any $x \in \mathbb{R}$, we have that upon updating v

$$\begin{aligned} \mathbb{P}(X_t^+(v) = 1) - \mathbb{P}(X_t^-(v) = 1) &= \frac{1}{2} \mathbb{E} \left[\tanh\left(\beta \sum_{w \sim v} X_t^+(w)\right) - \tanh\left(\beta \sum_{w \sim v} X_t^-(w)\right) \right] \\ &\leq \frac{\beta}{2} \mathbb{E} \left[\sum_{w \sim v} X_t^+(w) - X_t^-(w) \right] = \beta \sum_{w \sim v} \mathbf{m}_t(w), \end{aligned}$$

and so $\frac{d}{dt} \mathbf{m}_t(v) \leq \beta \sum_{w \sim v} \mathbf{m}_t(w) - \mathbf{m}_t(v)$. Hence,

$$\frac{d}{dt} \sum_v \mathbf{m}_t(v)^2 = 2 \sum_v \mathbf{m}_t(v) \frac{d}{dt} \mathbf{m}_t(v) \leq -2 \sum_v \mathbf{m}_t(v)^2 + 2\beta \sum_v \mathbf{m}_t(v) \sum_{w \sim v} \mathbf{m}_t(w),$$

and using $\mathbf{m}_t(v)\mathbf{m}_t(w) \leq \frac{1}{2}(\mathbf{m}_t(v)^2 + \mathbf{m}_t(w)^2)$ it follows that

$$\frac{d}{dt} \sum_v \mathbf{m}_t(v)^2 \leq -2(1 - \beta d) \sum_v \mathbf{m}_t(v)^2,$$

which implies the desired upper bound. ■

Recalling that $t_\star = t_m + s_\star$, we apply the above claim for $t = t_m$ (at which point $\sum_w \mathbf{m}_{t_m}(w)^2 = 1$ by definition) and $s = s_\star$ to find that $\sum_w \mathbf{m}_{t_\star}(w)^2 \leq \exp(-2(1 - \beta d)s_\star) \leq \exp(-s_\star)$, with the last inequality via $\beta d \leq \frac{1}{2}$. By (3.8) (keeping in mind that γ_1 and γ_2 are absolute constants) this implies that $\bar{d}_{\text{TV}}(t_\star) \leq \varepsilon$ if we take $s_\star \geq C \log(1/\varepsilon)$ for some absolute constant $C > 0$, as required. ■

3.2. Lower Bound. We now estimate the correlation of two vertices at an arbitrary time.

Claim 3.4. *There exist absolute constants $\kappa, \gamma > 0$ such that, for any initial state, if $\beta < \kappa/d$ then*

$$\sum_u \text{Cov}(X_t(u), X_t(v)) \leq \gamma \quad \text{for any } t > 0 \text{ and } v \in V.$$

Proof. Let X'_t and X''_t be two independent copies of the dynamics. By exploring the histories of the support we may couple X_t with X'_t and X''_t so that, on the event $\{u \notin \mathcal{C}_v\}$, the history of u in X_t is equal to the history of u in X'_t and the history of v in X_t is equal to the history of v in X''_t . Hence,

$$\begin{aligned} \mathbb{E}[X_t(u)X_t(v)] &= \mathbb{E}[X'_t(u)X''_t(v) + (X_t(u)X_t(v) - X'_t(u)X''_t(v))\mathbb{1}_{\{u \in \mathcal{C}_v\}}] \\ &\leq \mathbb{E}[X'_t(u)] \mathbb{E}[X''_t(v)] + 2\mathbb{P}(u \in \mathcal{C}_v). \end{aligned}$$

It follows that $\text{Cov}(X_t(u), X_t(v)) \leq 2\mathbb{P}(u \in \mathcal{C}_v) \leq 2\mathbb{P}(u \in \hat{\mathcal{C}}_v)$, and so

$$\sum_u \text{Cov}(X_t(u), X_t(v)) \leq 2\mathbb{E}|\hat{\mathcal{C}}_v| \leq \gamma,$$

with the final equality thanks to Lemma 2.2. \blacksquare

We are now ready to prove the lower bound on the mixing time in Theorem 1. To this end, we use the magnetization to generate a distinguishing statistic at time $t_\star^- = t_m - s_\star$, given by

$$f(\sigma) = \sum_{v \in V} \mathbf{m}_{t_\star^-}(v)\sigma(v).$$

Putting $Y = f(X_{t_\star^-}^+)$ for the dynamics started from all-plus and $Y' = f(\sigma)$ with σ drawn from the Ising distribution π , we combine Claim 3.3 with the fact that $\sum_v \mathbf{m}_{t_m}(v)^2 = 1$ (by definition) to get

$$\mathbb{E}Y = \sum_v \mathbf{m}_{t_\star^-}(v)^2 \geq e^{2(1-\beta d)s_\star} \sum_v \mathbf{m}_{t_m}(v)^2 = e^{2(1-\beta d)s_\star} \geq e^{s_\star} \quad (3.9)$$

(the last inequality using $\beta d \leq \frac{1}{2}$), whereas $\mathbb{E}Y' = 0$ (as $\mathbb{E}[\sigma(v)] = 0$ for any v).

For the variance estimate, observe that

$$\begin{aligned} \text{Var}(Y) &= \sum_{u,v} \mathbf{m}_{t_\star^-}(u)\mathbf{m}_{t_\star^-}(v) \text{Cov}\left(X_{t_\star^-}^+(u), X_{t_\star^-}^+(v)\right) \leq \frac{1}{2} \sum_{u,v} \left(\mathbf{m}_{t_\star^-}(u)^2 + \mathbf{m}_{t_\star^-}(v)^2\right) \text{Cov}\left(X_{t_\star^-}^+(u), X_{t_\star^-}^+(v)\right) \\ &\leq \gamma \sum_v \mathbf{m}_{t_\star^-}(v)^2 = \gamma \mathbb{E}Y, \end{aligned}$$

using Claim 3.4 for the inequality in the last line. Furthermore, since the law of X_t converges as $t \rightarrow \infty$ to that of σ , for any $v \in V$ we have

$$\sum_u \text{Cov}(\sigma(u), \sigma(v)) = \lim_{t \rightarrow \infty} \sum_u \text{Cov}(X_t(u), X_t(v)) \leq \gamma,$$

and so the same calculation in the above estimate for $\text{Var}(Y)$ shows that

$$\text{Var}(Y') \leq \gamma \mathbb{E}Y.$$

Altogether, by Chebyshev's inequality,

$$\mathbb{P}(Y \geq \frac{2}{3}\mathbb{E}Y) \geq 1 - 9\gamma/\mathbb{E}Y,$$

whereas

$$\mathbb{P}(Y' \leq \frac{1}{3}\mathbb{E}Y) \geq 1 - 9\gamma/\mathbb{E}Y.$$

Recalling (3.9), the expression $9\gamma/\mathbb{E}Y$ can be made less than $\varepsilon/2$ by choosing $s_\star \geq C \log(1/\varepsilon)$ for some absolute constant $C > 0$, thus concluding the proof of the lower bound. \blacksquare

4. ANALYSIS OF PERCOLATION CLUSTERS

4.1. **Red clusters: Proof of Lemma 2.1.** As we condition on the fact that either $A \in \text{RED}$ or $A \subset V_{\text{BLUE}}$, as well as on the collective history of every $v \notin A$, the history of the vertices of A must avoid \mathcal{H}_A^- — an event that we mark as \mathcal{M} — and then give rise to blue clusters or a single red one (we are interested in bounding the probability of the latter). To analyze the probability of \mathcal{M} , for each $u \in A$ we look at the latest time at which \mathcal{H}_A^- contains it (u is “undercut” by \mathcal{H}_A^-), that is,

$$s_u = s_u(\mathcal{H}_A^-) = \max \{s : u \in \mathcal{H}_{V \setminus A}(s)\},$$

and focus our attention on the vertices that are undercut in the unit interval $(t_\star - 1, t_\star]$ (which is the first unit interval to be exposed when developing \mathcal{H}_A), writing

$$A' = \{u \in A : s_u > t_\star - 1\},$$

and we denote by \mathcal{U} the event that every $u \in A'$ received an update in the interval $(s_u, t_\star]$, which is of course a necessary condition for \mathcal{M} (so as to avoid the scenario where $u \in \mathcal{H}_u(s_u)$ and intersects \mathcal{H}_A^- at that point). With this in mind, for any $A \subset V$ and \mathcal{H}_A^- we have

$$\mathbb{P}(A \in \text{RED} \mid \mathcal{H}_A^-, \{A \in \text{RED}\} \cup \{A \subset V_{\text{BLUE}}\}) = \frac{\mathbb{P}(A \in \text{RED}, \mathcal{M} \mid \mathcal{U})}{\mathbb{P}(\{A \in \text{RED}\} \cup \{A \subset V_{\text{BLUE}}\}, \mathcal{M} \mid \mathcal{U})}.$$

The numerator is at most $\mathbb{P}(A \in \text{RED} \mid \mathcal{U})$, while the denominator can be bounded from below by the probability that, in the space conditioned on \mathcal{U} , the last update to each $u \in A$ occurs in the interval $(s_u \vee t_\star - 1, t_\star]$ and it is oblivious (implying that its history amounts to the singleton $\{u\}$ dying out prior to being possibly undercut by \mathcal{H}_A^- , and so $u \in V_{\text{BLUE}}$). Hence,

$$\mathbb{P}(\{A \in \text{RED}\} \cup \{A \subset V_{\text{BLUE}}\}, \mathcal{M} \mid \mathcal{U}) \geq \theta^{|A|} (1 - 1/e)^{|A \setminus A'|} > 2^{-|A|},$$

where the term $\theta^{|A|}$ accounts for the probability that the latest most update is oblivious, the factor $(1 - 1/e)$ requires an update for vertices of $A \setminus A'$ (whose update in the last unit interval was not guaranteed by \mathcal{U}), and the last inequality used that $\theta(1 - 1/e) \geq (1 - \tanh(\frac{1}{5}))(1 - 1/e) > \frac{1}{2}$ by our assumption on β and the definition of θ in (2.4). Overall, we find that

$$\mathbb{P}(A \in \text{RED} \mid \mathcal{H}_A^-, \{A \in \text{RED}\} \cup \{A \subset V_{\text{BLUE}}\}) \leq 2^{|A|} \mathbb{P}(A \in \text{RED} \mid \mathcal{U}). \quad (4.1)$$

Recall that in order for A to form a complete red cluster, the update histories $\{\mathcal{H}_u : u \in A\}$ must belong to the same connected component of the space-time slab, and moreover, the configuration of A at time t_\star must be a nontrivial function of the initial configuration. Thus, either the histories $\{\mathcal{H}_u : u \in A\}$ coalesce to a single point w at some time $1 \leq T < t_\star$ — and then the spin there must depend nontrivially on the initial state, i.e., $X_T^+(w) \neq X_T^-(w)$ — or the histories for all $u \in A$ all join into one cluster along $(0, t_\star]$ and at least one of these survives to time 0. (The same would be true if we did not restrict the coalescence time to be at least 1, yet in this way the conditioning on \mathcal{U} , which only pertains to updates along the interval $(t_\star - 1, t_\star]$, does not cause any complications.) For the latter, we denote by $\mathcal{J}(a, b)$ the event that the histories join in the interval (a, b) , and for the former we let

$$\tau' = \min \{t \geq 1 : |\mathcal{H}_A(t_\star - t)| = 1\} \wedge t_\star, \quad T = t_\star - \tau',$$

and note that the variable τ' is a stopping time w.r.t. the natural filtration associated with exposing the update histories backward from time t_* ; indeed, in contrast to a definition of τ_v analogous to (2.7) — asking for $\{\mathcal{H}_u : u \in \mathcal{C}_v\}$ to coalesce to a single point — here one only requires this for $\{\mathcal{H}_u : u \in A\}$ (whereas \mathcal{C}_v may be affected by the histories along $(0, T]$ as these may admit additional vertices to it). With this notation, we deduce from the above discussion that

$$\mathbb{P}(A \in \text{RED} \mid \mathcal{U}) \leq \mathbb{P}\left(\bigcup_w \{\mathcal{J}(T, t_*), w \in \mathcal{H}_A(T), X_T^+(w) \neq X_T^-(w)\} \mid \mathcal{U}\right).$$

(If $T = 0$ and $A \in \text{RED}$ then $\mathcal{H}_A(0) \neq \emptyset$, whence $X_0^+(w) \neq X_0^-(w)$ trivially holds for any $w \in \mathcal{H}_A(0)$.) By conditioning on T as well as on $\mathcal{H}_A(T, t_*)$, the first two events on the right-hand side become measurable, while the event $X_T^+(w) \neq X_T^-(w)$ only depends on the histories along $(0, T]$ and satisfies

$$\mathbb{P}(X_T^+(w) \neq X_T^-(w) \mid T, \mathcal{H}_A(T, t_*)) = \mathbf{m}_T(w) \leq e^{t_*-T} \mathbf{m}_{t_*}(w),$$

where the final inequality used the fact, mentioned in the proof of Claim 3.3, that $\mathbf{m}_{t+s}(w) \geq e^{-s} \mathbf{m}_t(w)$ for any $s, t > 0$ and w , as the probability of no updates to w along the interval $(t, t+s)$ (maintaining the magnetization without a change) is e^{-s} . Now, averaging over this conditional space yields

$$\begin{aligned} \mathbb{P}(A \in \text{RED} \mid \mathcal{U}) &\leq \mathbb{E}\left[\sum_w \mathbb{1}_{\{\mathcal{J}(T, t_*)\}} \mathbb{1}_{\{w \in \mathcal{H}_A(T)\}} e^{t_*-T} \mathbf{m}_{t_*}(w) \mid \mathcal{U}\right] \\ &\leq \mathbb{E}\left[\sum_w \mathbb{1}_{\{A \subset \mathcal{C}_v\}} \mathbb{1}_{\{w \in \mathcal{H}_A(t_*-\tau', t_*)\}} e^{\tau'} \mathbf{m}_{t_*}(w) \mid \mathcal{U}\right], \end{aligned}$$

where we increased the event $\mathcal{J}(T, t_*)$ (the joining of \mathcal{H}_A along $(T, t_*]$) into $A \subset \mathcal{C}_v$ (valid for any $v \in A$) as well as the event $\{w \in \mathcal{H}_A(T)\}$ into $\{w \in \mathcal{H}_A(T, t_*)\}$, and finally plugged in that $T = t_* - \tau'$. Since by definition $\tau' \leq \tau_v = \min\{t \geq 1 : |\mathcal{H}_{\mathcal{C}_v}(t_* - t)| = 1\} \wedge t_*$ on the event $A \subset \mathcal{C}_v$, we conclude that

$$\mathbb{P}(A \in \text{RED} \mid \mathcal{U}) \leq \mathbb{E}\left[\sum_w \mathbb{1}_{\{A \subset \mathcal{C}_v\}} \mathbb{1}_{\{w \in \mathcal{H}_A(t_*-\tau_v, t_*)\}} e^{\tau_v} \mathbf{m}_{t_*}(w) \mid \mathcal{U}\right]. \quad (4.2)$$

The final step is to eliminate the conditioning on \mathcal{U} using the modified update history $\hat{\mathcal{H}}$, which we recall does not remove vertices from the history along the unit interval $(t_* - 1, t_*]$ and grants each vertex an automatic update at time t_* . As such, $\mathcal{H}_u(t) \subset \hat{\mathcal{H}}_u(t)$ for any vertex u and time t .

We claim that each of the terms in the right-hand of (4.2) is increasing in the percolation space-time slab (i.e., they can only increase when adding connections to the update histories). Indeed, this trivially holds for $\{A \subset \mathcal{C}_v\}$; the variable τ_v is increasing as it may take only longer for \mathcal{C}_v to coalesce to a single point; finally, as the interval $(t_* - \tau_v, t_*]$ does not decrease and neither does \mathcal{H}_A along it, the event $\{w \in \mathcal{H}_A(t_* - \tau_v, t_*)\}$ is also increasing.

Therefore, if we do not remove vertices from the update history along $(t_* - 1, t_*]$ then the right-hand of (4.2) could only increase. Further observe that, as long as no vertices are removed from the history along that unit interval, the connected components of the update history at time $t_* - 1$ remain exactly the same were we to modify the update times of any vertex there, while keeping them within that unit interval. In particular, should a vertex at all be updated in that period, we can move its latest update time to t_* .

In this version of the update history (retaining all vertices in the given unit interval, and letting the latest most update, if it is in that interval, be performed at time t_*), the effect of conditioning on \mathcal{U} in that every $u \in A'$ receives an update at time t_* . The fact that $\text{Po}(\lambda \mid \cdot \geq 1) \preceq \text{Po}(\lambda) + 1$ for any $\lambda > 0$ (as the ratio $\mathbb{P}(\text{Po}(\lambda) = k)/\mathbb{P}(\text{Po}(\lambda) > k)$ is monotone increasing in k) now implies (taking $\lambda \in (0, 1)$) that the number of updates that any $u \in A'$ receives along $(t_* - 1, t_*]$ conditioned on \mathcal{U} as part of \mathcal{H} is stochastically dominated by the corresponding number of updates as part of $\hat{\mathcal{H}}$.

Altogether we conclude that the right-hand of (4.2) can be increased to yield

$$\mathbb{P}(A \in \text{RED} \mid \mathcal{U}) \leq \mathbb{E} \left[\sum_w \mathbb{1}_{\{A \subset \hat{C}_v\}} \mathbb{1}_{\{w \in \hat{\mathcal{H}}_A(t_* - \hat{\tau}_v, t_*)\}} e^{\hat{\tau}_v} \mathbf{m}_{t_*}(w) \right],$$

and combining this with (4.1) completes the proof. \blacksquare

4.2. Discrete Fourier expansion for the update rules. The following lemma, which constructs the modified update rules Φ_A (as described in §2), will play a key role in the proof of Lemma 2.2.

Lemma 4.1. *For every $\varepsilon > 0$ there exists some $\kappa > 0$ such that the following holds provided $\beta d < \kappa$. For any $r \leq d$ there are nonnegative reals $\{p_{k,r} : k = 0, \dots, r\}$ satisfying*

$$p_{0,r} \geq 1 - \varepsilon, \quad \sum_k \binom{r}{k} p_{k,r} = 1, \quad \text{and} \quad \binom{r}{k} p_{k,r} \leq D_0 (2\beta r)^k \quad \text{for all } k, \quad (4.3)$$

where D_0 is an absolute constant, such that the Glauber dynamics can be coupled to an update function Φ that selects a subset $A \subset [r]$ of the neighbors of a degree- r vertex with probability $p_{|A|,r}$ and applies to it a symmetric monotone boolean function Φ_A (i.e., $\Phi_A(-x) = -\Phi_A(x)$ and $\Phi_A(x)$ is increasing in x).

Proof. Setting

$$f(x) = \frac{1}{2} (\tanh(x) + 1) = \frac{e^x}{e^x + e^{-x}}$$

we have that the Glauber dynamics update function at a given site with neighbors $\sigma_1, \dots, \sigma_r$ assigns it a new spin of 1 with probability $f(\beta \sum_{i=1}^r \sigma_i)$. Writing $f(x) = \sum_{\ell=0}^{\infty} B_\ell x^\ell$, i.e.,

$$B_\ell = [x^\ell] f(x),$$

and so, bearing in mind that $\tanh(z)$ has no singularities in the open disc of radius $\pi/2$ around 0 in \mathbb{C} and thus $\sum B_\ell$ converges absolutely,

$$B_0 = B_1 = 1/2 \quad \text{and} \quad \sum |B_\ell| = B \quad \text{for some absolute constant } B > 0.$$

Next, since $\sigma_i \in \{\pm 1\}$ the power series is multi-linear in σ_i , whence we can write

$$\left(\beta \sum \sigma_i \right)^\ell = \sum_{\substack{A \subset [r] \\ |A| \leq \ell}} C_{\ell,A} \prod_{i \in A} \sigma_i = \sum_{k=1}^{\ell \wedge r} C_{\ell,k} \sum_{|A|=k} \prod_{i \in A} \sigma_i,$$

where we used that the nonnegative coefficient $C_{\ell,A}$ depends by symmetry on $|A|$ rather than A itself, thus we can write $C_{\ell,k}$ for $|A| = k$. (Note that for $\ell = 1$ we have $C_{1,1} = \beta$.)

Now, for any particular $k \leq \ell \wedge r$, we can put $\sigma_1 = \dots = \sigma_r = 1$ to find that

$$(\beta r)^\ell = \sum_{i=0}^{\ell \wedge r} \sum_{|A|=i} C_{\ell,i} \geq \sum_{|A|=k} C_{\ell,k} = \binom{r}{k} C_{\ell,k},$$

and so

$$0 \leq C_{\ell,k} \leq \frac{(\beta r)^\ell}{\binom{r}{k}}. \quad (4.4)$$

Therefore, letting

$$C_k = \sum_{\ell=k}^{\infty} C_{\ell,k} B_\ell \quad \text{for } k \geq 1$$

and recalling that $\sum |B_\ell| = B$, we see that

$$|C_k| \leq B \sum_{\ell \geq k} \frac{(\beta r)^\ell}{\binom{r}{k}} \leq 2B \frac{(\beta r)^k}{\binom{r}{k}}, \quad (4.5)$$

with the last inequality valid as long as $\beta r \leq 1/2$.

We now define $p_{k,r}$ as follows:

$$p_{k,r} = \begin{cases} 2|C_k|(k+1) & k \geq 2, \\ 2\left(C_1 - \sum_{\substack{A' \ni 1 \\ |A'| \geq 2}} |C_{|A'|}\right) & k = 1, \\ 1 - \sum_{k \geq 1} \binom{r}{k} p_{k,r} & k = 0. \end{cases} \quad (4.6)$$

Our first step in verifying that this definition satisfies (4.6) is to show that $0 < p_{1,r} < 1$. For the upper bound, using (4.5) we have $p_{1,r} \leq 2|C_1| \leq 4B\beta < 1$ for β small enough. For the lower bound, observe that since $B_1 = 1/2$, $C_{1,1} = \beta$ and $C_{\ell,1} \leq (\beta r)^\ell / r$ using (4.4),

$$C_1 \geq \frac{\beta}{2} - \sum_{\ell=2}^{\infty} C_{\ell,1} |B_\ell| \geq \frac{\beta}{2} - \frac{B}{r} \sum_{\ell \geq 2} (\beta r)^\ell \geq \beta \left(\frac{1}{2} - 2\beta r B \right) > \beta/4. \quad (4.7)$$

as long as $\beta < 1/(4rB)$. On the other hand, again appealing to (4.5),

$$\sum_{\substack{A' \ni 1 \\ |A'| \geq 2}} |C_{|A'|}| = \sum_{k=2}^r \binom{r-1}{k-1} |C_k| \leq 2B \sum_{k=2}^r \frac{k}{r} (\beta r)^k = 2B\beta \sum_{k=2}^r k (\beta r)^{k-1} \leq \beta/8 \quad (4.8)$$

provided βr is sufficiently small. Combining the last two displays yields $p_{1,r} \geq \beta/8$.

Next, we wish to verify that $\binom{r}{k} p_{k,r} \leq D_0 (2\beta r)^k$ for some absolute constant D_0 and all k . Let $D_0 = 4B$ and note that for $k = 0$ the sought inequality is trivial since $D_0 > 1$ (recall $B \geq B_0 + B_1 = 1$) whereas $p_{0,r} < 1$ (we have shown that $p_{1,r} > 0$ and clearly $p_{k,r} \geq 0$ for all $k \geq 2$). For $k = 1$ we again recall from (4.5) that $rp_{1,r} \leq 2r|C_1| \leq 4\beta r B < D_0 (2\beta r)$, and similarly, for $k \geq 2$ we have

$$\binom{r}{k} p_{k,r} = 2 \binom{r}{k} |C_k| (k+1) \leq 4B(k+1)(\beta r)^k \leq 4B(2\beta r)^k = D_0 (2\beta r)^k.$$

For any sufficiently small βr this of course also shows that $p_{k,r} \leq 1$ for all k , as well as the final fact that $p_{0,d} \geq 1 - \varepsilon$ since

$$\sum_{k \geq 1} \binom{r}{k} p_{k,r} \leq D_0 \sum_{k \geq 1} (2\beta r)^k < 4\beta r D_0 < \varepsilon \quad (4.9)$$

for a small enough βr .

Having established that desired properties for $\{p_{k,r} : 0 \leq k \leq r\}$, define the new update function Φ which will examine a random subset A of the r neighbors of a vertex, selected with probability $p_{|A|,r}$ (giving a proper distribution over the subsets of $[r]$ since $\sum_k \binom{r}{k} p_{k,r} = 1$ as shown above), then apply the following function Φ_A to determine the probability of a plus update given $\sigma_A = \{\sigma_i : i \in A\}$.

$$\Phi_A(\sigma_A) = \begin{cases} \frac{1}{2} & A = \emptyset, \\ \frac{1}{2} + \frac{1}{2}\sigma_i & A = \{i\}, \\ \frac{1}{2} + \frac{1}{2(|A|+1)} [\sum_{i \in A} \sigma_i + \text{sign}(C_{|A|}) \prod_{i \in A} \sigma_i] & |A| \geq 2. \end{cases} \quad (4.10)$$

In order to establish that Φ can be coupled to the Glauber dynamics, we need to show that $f(\beta \sum_{i=1}^r \sigma_i)$ identifies with $\mathbb{E}[\Phi(\sigma_1, \dots, \sigma_r)]$ over all inputs $\{\sigma_i\}$. Since $B_0 = 1/2$, we must show that $\mathbb{E}[\Phi] - 1/2$ is equal to $\sum_{\ell=1}^{\infty} B_{\ell}(\beta \sum \sigma_i)^{\ell}$. Indeed,

$$\begin{aligned} \mathbb{E}[\Phi] - \frac{1}{2} &= \sum_i \sigma_i \left(C_1 - \sum_{\substack{A' \ni i \\ |A'| \geq 2}} C_{|A'|} \right) + \sum_{|A| \geq 2} |C_{|A|}| \left(\sum_{i \in A} \sigma_i + \text{sign}(C_{|A|}) \prod_{i \in A} \sigma_i \right) \\ &= \sum_{|A| \geq 1} C_{|A|} \prod_{i \in A} \sigma_i = \sum_{\ell=1}^{\infty} \sum_{k=1}^{\ell \wedge r} \sum_{|A|=k} C_{\ell,k} B_{\ell} \prod_{i \in A} \sigma_i = \sum_{\ell=1}^{\infty} B_{\ell} (\beta \sum \sigma_i)^{\ell}, \end{aligned}$$

with the last two equalities following from the definition of C_k and $C_{\ell,k}$. This completes the proof. \blacksquare

4.3. Exponential decay of cluster sizes: Proof of Lemma 2.2. Using the update rule from Lemma 4.1, the probability that an update of a vertex v of degree $r \leq d$ will examine precisely k of its neighbors is

$$\binom{r}{k} p_{k,r} \leq D_0 (2\beta r)^k \leq D_0 (2\beta d)^k,$$

with the inequality thanks to (4.3). The probability that a given neighbor of v , with degree some $r' \leq d$, receives an update in which it examines both v and $k-1$ additional neighbors is at most

$$\max_{r' \leq d} p_{k,r'} \binom{r'-1}{k-1} \leq \max_{r' \leq d} \frac{k}{r'} D_0 (2\beta r')^k = \frac{1}{d} D_0 (3\beta d)^k,$$

using that $x^{1/x} \leq e^{1/e} < 3/2$ for all $x \geq 2$. Hence, the rate at which the history of the vertex v expands to k additional vertices along the time interval $(0, t_*)$ is at most

$$D_0 (1 + r/d) (3\beta d)^k \leq 2D_0 (3\beta d)^k.$$

By the same reasoning, the extra update at time t_* that is applied to v in $\hat{\mathcal{H}}$ connects it to k of its neighbors ($k = 0, \dots, r$) with probability at most $D_0(2\beta d)^k$, while each of its r neighbors contributes at most k new points with probability at most $D_0(3\beta d)^k/d$.

We now develop the cluster of the vertex (w_0, t_0) in the space-time slab by exploring the branch at w_0 , both forwards and backward in time, examining which connections it has to new vertices — either through its own updates or through those which examine it — until it terminates via oblivious updates in both directions. We then repeat this process with one of the points discovered in the exploration process (arbitrary chosen), until all such points are exhausted and the cluster is completely revealed.

Let Y_m denote the number of vertices explored in this way after iteration m (i.e., Y_1 is the number of vertices discovered via the branch incident to (w_0, t_0) , etc.), and let Z_m be the total length of edges in the time dimension (i.e., $(z, a), (z, b)$ for $z \in V$ and $0 \leq a < b \leq t_*$) explored by then. We can stochastically dominate these by a process $(\bar{Y}_m, \bar{Z}_m) \succeq (Y_m, Z_m)$ given as follows.

First, for the length variable, we apply Lemma 4.1 with $\varepsilon = (1 - \eta)/4$, and put

$$\begin{aligned} \bar{Z}_0 &= 0, \\ \bar{Z}_m &= \bar{Z}_{m-1} + W_m \quad \text{where} \quad W_m \sim 1 + \Gamma(2, 1 - \varepsilon), \end{aligned}$$

with the gamma variable $\Gamma(2, 1 - \varepsilon)$ measuring the time until the explored branch terminates (in both ends) using the key estimate $p_{0,r} \geq 1 - \varepsilon$ from Lemma 4.1, translated by 1 to account for the unit interval $(t_* - 1, t_*]$ in which vertices are not removed from $\hat{\mathcal{H}}$.

For the vertex count variable, with the above discussion above in mind, observe that conditioned on W_m the number of new vertices exposed along the new branch is dominated by $\sum_{k=1}^d V_m^{(k)}$, in which

$$V_m^{(k)} \sim k \text{ Po} \left(2D_0(3\beta d)^k W_m \right) \quad (k = 1, \dots, d)$$

are mutually independent, while the extra update at time t_* (should the branch extend to that time) introduces at most $\sum_{k=0}^d \hat{V}_m^{(k)}$ additional vertices, where all $V_m^{(k)}$ and $\hat{V}_m^{(k)}$ are independent, given by

$$\mathbb{P} \left(\hat{V}_m^{(0)} = j \right) \leq D_0(2\beta d)^j, \quad \mathbb{P} \left(\hat{V}_m^{(k)} = j \right) \leq D_0(3\beta d)^j/d \quad (k = 1, \dots, d).$$

Therefore, with this notation, we write

$$\begin{aligned} \bar{Y}_0 &= 1, \\ \bar{Y}_m &= \bar{Y}_{m-1} + U_m \quad \text{where} \quad U_m = \sum_{k=1}^d V_m^{(k)} + \sum_{k=0}^d \hat{V}_m^{(k)}. \end{aligned}$$

Letting $\tau \geq 1$ be the iteration after which the exploration process exhausts all new vertices (so $\tau = 1$ iff both ends of the branch of (v_0, t_0) terminated before introducing any new vertices to the cluster), we wish to show that

$$\mathbb{E} \left[\exp(\eta \bar{Z}_\tau + \lambda \bar{Y}_\tau) \right] \leq \gamma \tag{4.11}$$

for $\gamma(\lambda, \eta) < \infty$. We may assume without loss of generality — recalling that $\varepsilon = (1 - \eta)/4 \leq \frac{1}{4}$ — that

$$\lambda \geq 4 \log(1/\varepsilon), \tag{4.12}$$

as the left-hand of (4.11) is monotone increasing in λ . Observe that as long as $3\beta de^\lambda < 1/2$ we have

$$\mathbb{E} \left[\exp \left(\lambda \hat{V}_m^{(0)} \right) \right] \leq 1 + D_0 \sum_{k \geq 1} (2\beta de^\lambda)^k \leq 1 + 4D_0 \beta de^\lambda$$

as well as

$$\prod_{k=1}^d \mathbb{E} \left[\exp \left(\lambda \hat{V}_m^{(k)} \right) \right] \leq 1 + 2D_0 \sum_{k \geq 1} (3\beta de^\lambda)^k \leq 1 + 12D_0 \beta de^\lambda,$$

and similarly,

$$\prod_{k=1}^d \mathbb{E} \left[\exp \left(\lambda V_m^{(k)} \right) \mid W_m \right] = \exp \left[2W_m D_0 \sum_{k=1}^d (e^{\lambda k} - 1)(3\beta d)^k \right] \leq \exp \left[12W_m D_0 \beta de^\lambda \right].$$

We can further assume that

$$h(\lambda) := D_0 \beta de^\lambda \quad \text{satisfies} \quad 12h(\lambda) < (1 - \eta)/2 = 1 - 2\varepsilon - \eta,$$

achievable by letting βd be sufficiently small. With this notation,

$$\mathbb{E} \left[e^{\lambda U_m + \eta W_m} \mid W_m \right] \leq e^{(12h(\lambda) + \eta)W_m + 16h(\lambda)},$$

and upon taking expectation over W_m , having $12h(\lambda) + \eta < 1 - 2\varepsilon$ implies that the moment-generating function of the gamma distribution will only contribute a polynomial factor, giving that

$$\mathbb{E} \left[e^{\lambda U_m + \eta W_m} \right] \leq e^{28h(\lambda) + \eta} \left(1 - \frac{12h(\lambda) + \eta}{1 - \varepsilon} \right)^{-2} \leq \varepsilon^{-2} e^{28h(\lambda)}. \quad (4.13)$$

Combining this with our definition of $\bar{Y}_m = 1 + \sum_{i=1}^m U_i$ and $\bar{Z}_m = \sum_{i=1}^m W_i$, we find that

$$\begin{aligned} \mathbb{E} \left[e^{\lambda \bar{Y}_\tau + \eta \bar{Z}_\tau} \right] &= \mathbb{E} \left[\sum_{m=1}^{\infty} e^{-\lambda m + 2\lambda \bar{Y}_m + \eta \bar{Z}_m} \mathbb{1}_{\{\tau=m\}} \right] \leq \sum_{m=1}^{\infty} e^{-\lambda m} \mathbb{E} \left[e^{2\lambda \bar{Y}_m + \eta \bar{Z}_m} \right] \\ &= \sum_{m=1}^{\infty} e^{\lambda(2-m)} \left(\mathbb{E} \left[e^{2\lambda U_1 + \eta W_1} \right] \right)^m, \end{aligned}$$

which, recalling (4.13) and plugging in the expression for $h(2\lambda)$, is at most

$$e^{2\lambda} \sum_{m=1}^{\infty} \exp \left[m \left(-\lambda + 2 \log \left(\frac{4}{1-\varepsilon} \right) + 28h(2\lambda) \right) \right] \leq e^{2\lambda} \sum_{m=1}^{\infty} \exp [m(-\lambda/2 + 28h(2\lambda))] = \gamma < \infty$$

using (4.12) for the first inequality and, say, that $28h(2\lambda) \leq \lambda/3$ (achieved by taking βd small enough) for the second one. (Note that $\gamma = \gamma(\lambda, \eta)$, as the assumption (4.12) introduces a dependence on η). This establishes (4.11) and thereby concludes the proof. \blacksquare

5. THE EFFECT OF INITIAL CONDITIONS ON MIXING

In this section we consider random initial conditions (both quenched and annealed), and prove Theorem 2. The first observation is that, thanks to Theorem 1, the worst-case mixing time satisfies

$$t_{\text{MIX}}(\alpha) = t_{\text{m}} + O(1) \quad \text{for any fixed } 0 < \alpha < 1,$$

with t_{m} as defined in (1.1), and moreover, the same holds for $t_{\text{MIX}}^{(+)}(\alpha)$, the mixing time started from all-plus. By Claim 3.3 we have $\frac{1}{2} \log n \leq t_{\text{m}} \leq (\frac{1}{2} + \varepsilon_{\beta}) \log n$ with $\varepsilon_{\beta} = \beta d / (2 - 2\beta d)$ vanishing as $\beta \downarrow 0$. Thus, we may prove the bounds on the annealed / quenched mixing times when replacing t_{m} by $\frac{1}{2} \log n$.

5.1. Annealed analysis. As mentioned in the introduction, rather than comparing two worst case boundary conditions we will compare a random one directly with the stationary distribution: By considering updates in the range $t \in (-\infty, t_{\text{m}}]$ we can use the coupling from the past construction to generate a coupling with the stationary distribution. Let X_t denote the process started from uniform initial conditions at time 0 and let Y_t be the process generated by coupling from the past.

The information percolation clusters of V will now be defined as the connected components of the graph on the vertex set V where (u, v) is an edge iff $\mathcal{H}_u(t) \cap \mathcal{H}_v(t) \neq \emptyset$ for some $-\infty < t \leq t_{\text{m}}$ (in contrast to the previous definition where we had $0 < t \leq t_{\text{m}}$). The notion of being a red cluster is redefined to be any \mathcal{C}_v such that $|\bigcup_{u \in \mathcal{C}_v} \mathcal{H}_u(t')| \geq 2$ for all $0 \leq t' \leq t_{\text{m}}$. Blue clusters will be defined as before and green clusters will again be the remaining clusters. We claim that we can couple the spins at time t_{m} of all non-red clusters. Indeed if a cluster \mathcal{C}_v is not red then there is some time $t' > 0$ such that $|\bigcup_{u \in \mathcal{C}_v} \mathcal{H}_u(t')| = 1$. Call this vertex w . By symmetry both $X_{t'}(w)$ and $Y_{t'}(w)$ are equally likely to be plus or minus and so we may couple them to be equal independently of the spins of the other clusters. We may then also couple the spins in that cluster to be the same in both X_t and Y_t to be equal for all $t > t'$. Thus the configurations will agree outside of the red clusters.

Let $\mathfrak{W}(A)$ denote the size of the smallest connected set of vertices (animal) containing A . In a graph of maximum degree d , the number of trees of size k containing the vertex v is bounded above by $(ed)^k$ and hence the number of animals A containing a specified vertex with $\mathfrak{W}(A) = k$ is at most $(ed)^k$.

Lemma 5.1. *For any $d, C, \varepsilon > 0$ there exists $\beta_0 > 0$ such that the following holds for large enough n . If $0 < \beta < \beta_0$ and $t_{\star} = (\frac{1}{4} + \varepsilon) \log n$ then for any A ,*

$$\sup_{\mathcal{H}_A^-} \mathbb{P}(A \in \text{RED} \mid \mathcal{H}_A^-, \{A \in \text{RED}\} \cup \{A \subset V_{\text{BLUE}}\}) \leq \frac{1}{\sqrt{n} \log n} e^{-C\mathfrak{W}(A)}.$$

Proof. Similarly to the proof of Lemma 2.1 we have the analogue of equation (4.1)

$$\mathbb{P}(A \in \text{RED} \mid \mathcal{H}_A^-, \{A \in \text{RED}\} \cup \{A \subset V_{\text{BLUE}}\}) \leq 2^{|A|} \mathbb{P}(A \in \text{RED} \mid \mathcal{U}),$$

where \mathcal{U} is defined as in the proof of Lemma 2.1. Then for any $v \in A$,

$$\mathbb{P}(A \in \text{RED} \mid \mathcal{U}) \leq \mathbb{P}\left(|\mathcal{H}_{\mathcal{C}_v}| \geq \mathfrak{W}(A), \mathfrak{L}(\mathcal{X}_{v,0}) \geq 2t_{\star} \mid \mathcal{U}\right),$$

since the total length of a red cluster must be at least $2t_{\star}$ and it must contain at least $\mathfrak{W}(A)$ vertices. Both $|\mathcal{H}_{\mathcal{C}_v}|$ and $\mathfrak{L}(\mathcal{X}_{v,0})$ are increasing in the component sizes, and so, by the same monotonicity

argument as Lemma 2.1, we have that

$$\mathbb{P}(A \in \text{RED} \mid \mathcal{H}_A^-, \{A \in \text{RED}\} \cup \{A \subset V_{\text{BLUE}}\}) \leq 2^{|A|} \mathbb{P}\left(|\hat{\mathcal{H}}_{\hat{c}_v}| \geq \mathfrak{W}(A), \mathfrak{L}(\hat{\mathcal{X}}_{v,0}) \geq 2t_\star\right).$$

Taking $\lambda = \log 2 + C$ and $\frac{1}{4}(\frac{1}{4} + \varepsilon)^{-1} < \eta < 1$ in Lemma 2.2 then shows that, for β_0 small enough,

$$\begin{aligned} \mathbb{P}\left(|\hat{\mathcal{H}}_{\hat{c}_v}| \geq \mathfrak{W}(A), \mathfrak{L}(\hat{\mathcal{X}}_{v,0}) \geq 2t_\star\right) &\leq \frac{\mathbb{E}\left[\exp\left(\eta \mathfrak{L}(\hat{\mathcal{X}}_{v,0}) + \lambda |\hat{\mathcal{H}}_{\hat{c}_v}|\right)\right]}{\exp(2\eta t_\star + \lambda \mathfrak{W}(A))} \\ &\leq \gamma \exp\left(-2\eta\left(\frac{1}{4} + \varepsilon\right) \log n - (\log 2 + C)\mathfrak{W}(A)\right) \\ &\leq \frac{1}{\sqrt{n} \log n} 2^{-\mathfrak{W}(A)} e^{-C\mathfrak{W}(A)} \end{aligned}$$

(with room, as we could have replaced the $\sqrt{n} \log n$ by some $n^{1/2+\varepsilon'}$), which completes the proof. \blacksquare

We now establish an upper bound on $t_{\text{MIX}}^{(v)}$, the mixing time starting from the uniform distribution.

Proposition 5.2. *For any $d, \varepsilon > 0$ there exists $\beta_0 > 0$ such that the following holds. If $0 < \beta < \beta_0$ and $t_\star = (\frac{1}{4} + \varepsilon) \log n$ then $\|\mathbb{P}(X_{t_\star} \in \cdot) - \pi\|_{\text{TV}} \rightarrow 0$ as $n \rightarrow \infty$.*

Proof. Having coupled X_t and Y_t as described above we have that

$$\begin{aligned} \|\mathbb{P}(X_{t_\star} \in \cdot) - \mathbb{P}(Y_{t_\star} \in \cdot)\|_{\text{TV}} &\leq \mathbb{E}\left[\left\|\mathbb{P}(X_{t_\star}(V \setminus V_{\text{GREEN}}) \in \cdot \mid \mathcal{H}_{\text{GREEN}}) - \nu_{V \setminus V_{\text{GREEN}}}\right\|_{\text{TV}}\right] \\ &\quad + \mathbb{E}\left[\left\|\mathbb{P}(Y_{t_\star}(V \setminus V_{\text{GREEN}}) \in \cdot \mid \mathcal{H}_{\text{GREEN}}) - \nu_{V \setminus V_{\text{GREEN}}}\right\|_{\text{TV}}\right] \end{aligned}$$

where ν_A is the uniform measure on the configurations on A . Similarly to the argument used to derive equation (3.2), we find that

$$\|\mathbb{P}(X_{t_\star} \in \cdot) - \mathbb{P}(Y_{t_\star} \in \cdot)\|_{\text{TV}} \leq \left(\sup_{\mathcal{H}_{\text{GREEN}}} \mathbb{E}\left[2^{|V_{\text{RED}} \cap V_{\text{RED}}'|} \mid \mathcal{H}_{\text{GREEN}}\right] - 1\right)^{1/2}.$$

With the same coupling as in the proof of Theorem 1, analogously to equation (3.4), we have

$$\|\mathbb{P}(X_{t_\star} \in \cdot) - \mathbb{P}(Y_{t_\star} \in \cdot)\|_{\text{TV}}^2 \leq 2 \sum_v \left(\sum_{A \ni v} 2^{|A|} \Psi_A\right)^2. \quad (5.1)$$

Applying Lemma 5.1 with $C = \lceil \log(4ed) \rceil$ (while recalling that $\#\{A \ni v : \mathfrak{W}(A) = k\} \leq (ed)^k$), we get

$$\sum_{A \ni v} 2^{|A|} \Psi_A \leq \sum_k \sum_{\substack{A: \mathfrak{W}(A)=k \\ v \in A}} \frac{2^k e^{-Ck}}{\sqrt{n} \log n} \leq \sum_k \frac{(2ed)^k e^{-Ck}}{\sqrt{n} \log n} \leq \frac{1}{\sqrt{n} \log n},$$

provided that $\beta > \beta_0$ with β_0 from that lemma. It follows that

$$\|\mathbb{P}(X_{t_\star} \in \cdot) - \mathbb{P}(Y_{t_\star} \in \cdot)\|_{\text{TV}} \leq O(\log^{-2} n),$$

and in particular $\|\mathbb{P}(X_{t_\star} \in \cdot) - \pi\|_{\text{TV}} = o(1)$, as required. \blacksquare

Remark 5.3. In the above proof one could instead carry the analysis as in the proof of Theorem 1 (partitioning the event in Lemma 2.1 according to the events $\{|\hat{\mathcal{C}}_v| = k\}$ when estimating the sum over $v \ni A$), that way replacing the factor of $(ed)^k$ lattice animals by 2^k subsets of $\hat{\mathcal{C}}_v$. Consequently, the statement of Proposition 5.2 remains valid for any $\beta < c_0/d$ where c_0 depends on ε but not on d .

5.2. Quenched analysis. Here we show that the mixing time from a typical random initial state is at most a factor of $1 + \varepsilon_\beta$ faster than that from the worst starting state. As before, let X_t be started from a uniformly chosen initial state X_0 and let Y_t be started from the stationary distribution π .

Proposition 5.4. *Let $t_\star^- = \frac{1}{2} \log n - w_n$ for some $w_n \uparrow \infty$. Then $\|\mathbb{P}_{X_0}(X_{t_\star^-} \in \cdot) - \pi\|_{\text{TV}} \xrightarrow{p} 1$ as $n \rightarrow \infty$.*

Proof. Note that by the monotonicity of the update rules, for any update history of u , the spin at u is a monotone function of X_0 . With probability e^{-t} the vertex u is never updated in which case $X_t(u) = X_0(u)$. Since by symmetry $\mathbb{E}[X_t(u)] = 0$, it follows that

$$\mathbb{E}[X_t(u) \mid X_0(u) = +1] \geq e^{-t}, \quad \mathbb{E}[X_t(u) \mid X_0(u) = -1] \leq -e^{-t}.$$

Thus we have that $\mathbb{E}[X_0(u)X_t(u)] \geq e^{-t}$ and so

$$\mathbb{E}\left[\sum_u X_0(u)X_{t_\star^-}(u)\right] \geq ne^{-t_\star^-} = \sqrt{ne^{w_n}}.$$

Let $\mathcal{E}_{u,v}$ be the event that $u \in \mathcal{C}_v$ or $v \in \mathcal{H}_u(0)$ or $u \in \mathcal{H}_v(0)$ for the history developed from time t_\star^- . Similarly to Claim 3.4, let X'_t and X''_t be two independent copies of the dynamics. By exploring the histories we may couple X_t with X'_t and X''_t so that, on the event $\mathcal{E}_{u,v}^c$, the history of v in X_t is equal to the history of v in X'_t and the history of u in X_t is equal to the history of u in X''_t . Hence,

$$\mathbb{E}\left[X_0(u)X_{t_\star^-}(u)X_0(v)X_{t_\star^-}(v)\right] \leq \mathbb{E}\left[X'_0(u)X'_{t_\star^-}(u)\right]\mathbb{E}\left[X''_0(v)X''_{t_\star^-}(v)\right] + 2\mathbb{P}(\mathcal{E}_{u,v}),$$

yielding $\text{Cov}(X_0(u)X_{t_\star^-}(u), X_0(v)X_{t_\star^-}(v)) \leq 2\mathbb{P}(\mathcal{E}_{u,v})$. By Lemma 2.2,

$$\sum_u \text{Cov}\left(X_0(u)X_{t_\star^-}(u), X_0(v)X_{t_\star^-}(v)\right) \leq c_1.$$

and so

$$\text{Var}\left(\sum_u X_0(u)X_{t_\star^-}(u)\right) \leq c_1 n.$$

Thus, by Chebyshev's inequality we infer that

$$\mathbb{P}\left(\sum_u X_0(u)X_{t_\star^-}(u) > \frac{1}{2}\sqrt{ne^{w_n}}\right) \geq 1 - O(e^{-2w_n}),$$

and so by Markov's inequality,

$$\mathbb{P}\left(\mathbb{P}\left(\sum_u X_0(u)X_{t_\star^-}(u) > \frac{1}{2}\sqrt{ne^{w_n}} \mid X_0\right) \geq 1 - e^{-w_n}\right) \geq 1 - O(e^{-w_n}) \rightarrow 1. \quad (5.2)$$

By the exponential decay of correlations of Y and the fact that it is independent of X we have that

$$\text{Var}\left(\sum_u X_0(u)Y_{t_\star^-}(u) \mid X_0\right) \leq c_2 n$$

for some $c_2 > 0$. Thus, since $\mathbb{E} \left[\sum_u X_0(u) Y_{t_*^-}(u) \mid X_0 \right] = 0$, it follows that

$$\mathbb{P} \left(\sum_u X_0(u) Y_{t_*^-}(u) > \frac{1}{2} \sqrt{n} e^{w_n} \mid X_0 \right) = O(e^{-2w_n}) \rightarrow 0 \quad (5.3)$$

uniformly in X_0 . Comparing equations (5.2) and (5.3) completes the result. \blacksquare

REFERENCES

- [1] M. Aizenman and R. Holley, *Rapid convergence to equilibrium of stochastic Ising models in the Dobrushin Shlosman regime*, (Minneapolis, Minn., 1984), IMA Vol. Math. Appl., vol. 8, Springer, New York, 1987, pp. 1–11.
- [2] D. Aldous, *Random walks on finite groups and rapidly mixing Markov chains*, Seminar on probability, XVII, 1983, pp. 243–297.
- [3] D. Aldous and P. Diaconis, *Shuffling cards and stopping times*, Amer. Math. Monthly **93** (1986), 333–348.
- [4] P. Diaconis, *The cutoff phenomenon in finite Markov chains*, Proc. Nat. Acad. Sci. U.S.A. **93** (1996), no. 4, 1659–1664.
- [5] P. Diaconis, R. L. Graham, and J. A. Morrison, *Asymptotic analysis of a random walk on a hypercube with many dimensions*, Random Structures Algorithms **1** (1990), no. 1, 51–72.
- [6] P. Diaconis and L. Saloff-Coste, *Comparison techniques for random walk on finite groups*, Ann. Probab. **21** (1993), no. 4, 2131–2156.
- [7] P. Diaconis and L. Saloff-Coste, *Comparison theorems for reversible Markov chains*, Ann. Appl. Probab. **3** (1993), no. 3, 696–730.
- [8] P. Diaconis and L. Saloff-Coste, *Logarithmic Sobolev inequalities for finite Markov chains*, Ann. Appl. Probab. **6** (1996), no. 3, 695–750.
- [9] P. Diaconis and M. Shahshahani, *Generating a random permutation with random transpositions*, Z. Wahrsch. Verw. Gebiete **57** (1981), no. 2, 159–179.
- [10] P. Diaconis and M. Shahshahani, *Time to reach stationarity in the Bernoulli-Laplace diffusion model*, SIAM J. Math. Anal. **18** (1987), no. 1, 208–218.
- [11] J. Ding, E. Lubetzky, and Y. Peres, *The mixing time evolution of Glauber dynamics for the mean-field Ising model*, Comm. Math. Phys. **289** (2009), no. 2, 725–764.
- [12] R. L. Dobrushin, *Markov processes with a large number of locally interacting components: Existence of a limit process and its ergodicity*, Problemy Peredači Informacii **7** (1971), no. 2, 70–87 (Russian).
- [13] R. L. Dobrushin and S. B. Shlosman, *Constructive criterion for the uniqueness of Gibbs field*, Statistical physics and dynamical systems (Köszeg, 1984), Progr. Phys., vol. 10, Birkhäuser Boston, Boston, MA, 1985, pp. 347–370.
- [14] R. A. Holley, *An ergodic theorem for interacting systems with attractive interactions*, Z. Wahrscheinlichkeitstheorie und Verw. Gebiete **24** (1972), 325–334.
- [15] R. A. Holley and D. W. Stroock, *Logarithmic Sobolev inequalities and stochastic Ising models*, J. Statist. Phys. **46** (1987), no. 5-6, 1159–1194.
- [16] D. A. Levin, M. Luczak, and Y. Peres, *Glauber dynamics for the Mean-field Ising Model: cut-off, critical power law, and metastability*, Probab. Theory Related Fields **146** (2010), no. 1–2, 223–265.
- [17] D. A. Levin, Y. Peres, and E. L. Wilmer, *Markov chains and mixing times*, 2008.
- [18] T. M. Liggett, *Interacting particle systems*, Classics in Mathematics, Springer-Verlag, Berlin, 2005.
- [19] E. Lubetzky and A. Sly, *Cutoff phenomena for random walks on random regular graphs*, Duke Math. J. **153** (2010), no. 3, 475–510.
- [20] E. Lubetzky and A. Sly, *Explicit expanders with cutoff phenomena*, Electron. J. Probab. **16** (2011), no. 15, 419–435.
- [21] E. Lubetzky and A. Sly, *Cutoff for the Ising model on the lattice*, Invent. Math. **191** (2013), no. 3, 719–755.
- [22] E. Lubetzky and A. Sly, *Cutoff for general spin systems with arbitrary boundary conditions*, Comm. Pure. Appl. Math. **67** (2014), no. 6, 982–1027.
- [23] E. Lubetzky and A. Sly, *Information percolation for the Ising model: cutoff in three dimensions up to criticality*, preprint. Available at [arXiv:1401.6065](https://arxiv.org/abs/1401.6065) (2014).

- [24] F. Martinelli, *Lectures on Glauber dynamics for discrete spin models*, Lectures on probability theory and statistics (Saint-Flour, 1997), Lecture Notes in Math., vol. 1717, Springer, Berlin, 1999, pp. 93–191.
- [25] F. Martinelli and E. Olivieri, *Approach to equilibrium of Glauber dynamics in the one phase region. I. The attractive case*, Comm. Math. Phys. **161** (1994), no. 3, 447–486.
- [26] F. Martinelli and E. Olivieri, *Approach to equilibrium of Glauber dynamics in the one phase region. II. The general case*, Comm. Math. Phys. **161** (1994), no. 3, 487–514.
- [27] F. Martinelli, E. Olivieri, and R. H. Schonmann, *For 2-D lattice spin systems weak mixing implies strong mixing*, Comm. Math. Phys. **165** (1994), no. 1, 33–47.
- [28] J. Miller and Y. Peres, *Uniformity of the uncovered set of random walk and cutoff for lamplighter chains*, Ann. Probab. **40** (2012), no. 2, 535–577.
- [29] J. G. Propp and D. B. Wilson, *Exact sampling with coupled Markov chains and applications to statistical mechanics*, Proceedings of the Seventh International Conference on Random Structures and Algorithms (Atlanta, GA, 1995), 1996, pp. 223–252.
- [30] L. Saloff-Coste, *Lectures on finite Markov chains*, Lectures on probability theory and statistics (Saint-Flour, 1996), Lecture Notes in Math., vol. 1665, Springer, Berlin, 1997, pp. 301–413.
- [31] L. Saloff-Coste, *Random walks on finite groups*, Probability on discrete structures, 2004, pp. 263–346.
- [32] D. W. Stroock and B. Zegarliński, *The equivalence of the logarithmic Sobolev inequality and the Dobrushin-Shlosman mixing condition*, Comm. Math. Phys. **144** (1992), no. 2, 303–323.
- [33] D. W. Stroock and B. Zegarliński, *The logarithmic Sobolev inequality for discrete spin systems on a lattice*, Comm. Math. Phys. **149** (1992), no. 1, 175–193.

EYAL LUBETZKY

MICROSOFT RESEARCH, ONE MICROSOFT WAY, REDMOND, WA 98052-6399, USA.

E-mail address: eyal@microsoft.com

ALLAN SLY

DEPARTMENT OF STATISTICS, UC BERKELEY, BERKELEY, CA 94720, USA.

E-mail address: sly@stat.berkeley.edu