Dynamics for the Mean-field Random-cluster Model*

Antonio Blanca and Alistair Sinclair

University of California, Berkeley, USA
{ablanca,sinclair}@cs.berkeley.edu

Abstract

The random-cluster model has been widely studied as a unifying framework for random graphs, spin systems and random spanning trees, but its dynamics have so far largely resisted analysis. In this paper we study a natural non-local Markov chain known as the Chayes-Machta dynamics for the mean-field case of the random-cluster model, and identify a critical regime (λ_s, λ_S) of the model parameter λ in which the dynamics undergoes an exponential slowdown. Namely, we prove that the mixing time is $\Theta(\log n)$ if $\lambda \notin [\lambda_s, \lambda_S]$, and $\exp(\Omega(\sqrt{n}))$ when $\lambda \in (\lambda_s, \lambda_S)$. These results hold for all values of the second model parameter $q > 1$. In addition, we prove that the local heat-bath dynamics undergoes a similar exponential slowdown in (λ_s, λ_S).

1998 ACM Subject Classification G.3 Probability and Statistics

Keywords and phrases random-cluster model, random graphs, Markov chains, statistical physics, dynamics

Digital Object Identifier 10.4230/LIPIcs.APPROX-RANDOM.2015.528

1 Introduction

Background and previous work. Let $H = (V, E)$ be a finite graph. The *random-cluster model* on H with parameters $p \in (0, 1)$ and $q > 0$ assigns to each subgraph $(V, A \subseteq E)$ a probability

$$\mu_{p,q}(A) \propto p^{|A|}(1-p)^{|E|-|A|}q^{c(A)},$$

where $c(A)$ is the number of connected components in (V, A). A is a configuration of the model.

The random-cluster model was introduced in the late 1960s by Fortuin and Kasteleyn [10] as a unifying framework for studying random graphs, spin systems in physics and random spanning trees; see the book [14] for extensive background. When $q = 1$ this model corresponds to the standard Erdős-Rényi model on subgraphs of H, but when $q > 1$ (resp., $q < 1$) the resulting probability measure favors subgraphs with more (resp., fewer) connected components, and is thus a strict generalization.

For the special case of integer $q \geq 2$ the random-cluster model is, in a precise sense, dual to the classical ferromagnetic *q-state Potts model*, where configurations are assignments of spin values $\{1, \ldots, q\}$ to the vertices of H; the duality is established via a coupling of the models (see, e.g., [9]). Consequently, the random-cluster model illuminates much of the physical theory of the Ising/Potts models. Indeed, recent breakthrough work by Beffara

* Antonio Blanca was supported by a U.C. Berkeley Chancellor's Fellowship, an NSF Graduate Research Fellowship and NSF grants CCF-1016896 and CCF-1420934. Alistair Sinclair was supported by NSF grants CCF-1016896 and CCF-1420934.

18th Int'l Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX'15) / 19th Int'l Workshop on Randomization and Computation (RANDOM'15).
Editors: Naveen Garg, Klaus Jansen, Anup Rao, and José D. P. Rolim; pp. 528–543

Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

and Duminil-Copin [1] uses the geometry of the random-cluster model in $\mathbb{Z}^2$ to establish the critical temperature of the q-state Potts model, settling a long-standing conjecture.

At the other extreme, when $q, p \to 0$ and p approaches zero at a slower rate (i.e., $q/p \to 0$) the random-cluster measure $\mu_{p,q}$ converges to the uniform random spanning tree measure on H. Random spanning trees are fundamental probabilistic objects, whose relevance goes back to Kirchhoff's work on electrical networks [16]. In this paper we investigate the dynamics of the random-cluster model, i.e., Markov chains on random-cluster configurations that are reversible w.r.t. $\mu_{p,q}$ and thus converge to it. The dynamics of physical models are of fundamental interest, both as evolutionary processes in their own right and as Markov chain Monte Carlo (MCMC) algorithms for sampling configurations in equilibrium. In both these contexts the central object of study is the *mixing time*, i.e., the number of steps until the dynamics is close to the equilibrium measure $\mu_{p,q}$ starting from any initial configuration. While dynamics for the Ising and Potts models have been widely studied, very little is known about random-cluster dynamics. The main reason for this appears to be the fact that connectivity is a global property which has led to the failure of existing Markov chains analysis tools.

We focus on the *mean-field* case, where H is the complete graph on n vertices. In this case the random-cluster model may be viewed as the standard random graph model $\mathcal{G}_{n,p}$, enriched by a factor that depends on the component structure. As we shall see, the mean-field case is already quite non-trivial; moreover, it has historically proven to be a useful starting point in understanding the dynamics on more general graphs. The structural properties of the mean-field model are already well understood [3, 19]; in particular, it exhibits a phase transition (analogous to that in $\mathcal{G}_{n,p}$) corresponding to the appearance of a "giant" component of linear size. It is natural here to re-parameterize by setting $p = \lambda/n$; the phase transition then occurs at the critical value $\lambda = \lambda_c(q)$ given by

$$\lambda_c(q) = \begin{cases} q & \text{for } 0 < q \leq 2; \\ 2\left(\frac{q-1}{q-2}\right)\log(q-1) & \text{for } q > 2. \end{cases}$$

For $\lambda < \lambda_c(q)$ all components are of size $O(\log n)$ w.h.p.[1], while for $\lambda > \lambda_c(q)$ there is a unique giant component of size θn (for some constant θ that depends on q and λ). The former regime is called the *disordered phase*, and the latter is the *ordered phase*. Henceforth we assume $q > 1$, since the $q < 1$ regime is structurally quite different; the dynamics are trivial for $q = 1$.

Our main object of study is a non-local dynamics known as the *Chayes-Machta (CM) dynamics* [6]. Given a random-cluster configuration (V, A), one step of this dynamics is defined as follows:

- **(i)** *activate* each connected component of (V, A) independently with probability $1/q$;
- **(ii)** remove all edges connecting active vertices;
- **(iii)** add each edge connecting active vertices independently with probability p, leaving the rest of the configuration unchanged.

It is easy to check that this dynamics is reversible w.r.t. $\mu_{p,q}$ [6]. Until now, the mixing time of the CM dynamics has not been rigorously established for any non-trivial random-cluster measure $\mu_{p,q}$ on any graph. Our goal in this paper is to analyze the CM dynamics in the mean-field case for all values of $q > 1$ and all values of $\lambda > 0$.

[1] We say that an event occurs *with high probability (w.h.p.)* if it occurs with probability approaching 1 as $n \to \infty$.

For integer q, the CM dynamics is a close cousin of the well studied and widely used *Swendsen-Wang (SW) dynamics* [20]. The SW dynamics is primarily a dynamics for the Ising/Potts model, but it may alternatively be viewed as a Markov chain for the random-cluster model using the coupling of these measures mentioned earlier. However, the SW dynamics is only well-defined for integer q, while the random-cluster model makes perfect sense for all $q > 0$. The CM dynamics was introduced precisely in order to allow for this generalization.

The SW dynamics for the mean-field case is fully understood for $q = 2$: recent results of Long, Nachmias, Ning and Peres [18], building on earlier work of Cooper, Dyer, Frieze and Rue [7], show that the mixing time is $\Theta(1)$ for $\lambda < \lambda_c$, $\Theta(\log n)$ for $\lambda > \lambda_c$, and $\Theta(n^{1/4})$ for $\lambda = \lambda_c$. Until recently, the picture for integer $q \geq 3$ was much less complete: Huber [15] gave bounds of $O(\log n)$ and $O(n)$ on the mixing time when λ is far below and far above λ_c respectively, while Gore and Jerrum [13] showed that at the critical value $\lambda = \lambda_c$ the mixing time is $\exp(\Omega(\sqrt{n}))$. All these results were developed for the Ising/Potts model, so their relevance to the random-cluster model is limited to the case of integer q. In work that appeared after the submission of this manuscript [2], Galanis, Štefankovič and Vigoda [11] provide a more comprehensive analysis of the $q \geq 3$ mean-field case. Finally, for the very different case of the d-dimensional torus, Borgs et al. [4, 5] proved exponential lower bounds for the mixing time of the SW dynamics for $\lambda = \lambda_c$ and q sufficiently large.

Our work is the first to provide tight bounds for the mixing time of any random-cluster dynamics for general (non-integer) values of q.

Results. To state our results we identify two further critical points, $\lambda_s(q)$ and $\lambda_S(q)$, with the property that $\lambda_s(q) \leq \lambda_c(q) \leq \lambda_S(q)$. (For $1 < q \leq 2$ these three points coincide; for $q > 2$ they are all distinct.) The definitions of these points are somewhat technical and can be found in Section 2.

Our first result shows that the CM dynamics reaches equilibrium very rapidly for λ outside the "critical" window $[\lambda_s, \lambda_S]$. Moreover, our bounds are tight throughout the fast mixing regime.

▶ **Theorem 1.** *For any $q > 1$, the mixing time of the mean-field CM dynamics is $\Theta(\log n)$ for $\lambda \notin [\lambda_s, \lambda_S]$.*

Our next result shows that, *inside* the critical window (λ_s, λ_S), the mixing time is dramatically larger. (We state this result only for $q > 2$ as otherwise the window is empty.)

▶ **Theorem 2.** *For any $q > 2$, the mixing time of the mean-field CM dynamics is $e^{\Omega(\sqrt{n})}$ for $\lambda \in (\lambda_s, \lambda_S)$.*

We now provide an interpretation of the above results. When $q > 2$ the mean-field random-cluster model exhibits a *first-order* phase transition, which means that at criticality ($\lambda = \lambda_c$) the ordered and disordered phases mentioned earlier *coexist* [19], i.e., each contributes about half of the probability mass. (For $q \leq 2$, there is no phase coexistence.) Phase coexistence suggests exponentially slow mixing for most natural dynamics, because of the difficulty of moving between the phases. Moreover, by continuity we should expect that, within a constant-width interval around λ_c, the effect of the non-dominant phase (ordered below λ_c, disordered above λ_c) will still be felt, as it will form a second mode (local maximum) for the random-cluster measure. This leads to so-called *metastable* states near that local maximum from which it is very hard to escape, so slow mixing should persist throughout this interval. Intuitively, the values λ_s, λ_S mark the points at which the local maxima disappear. A similar

phenomenon was captured in the case of the Potts model by Cuff *et al.* [8]. Our results make the above picture for the dynamics rigorous for the random-cluster model for all $q > 2$; notably, in contrast to the Potts model, in the random-cluster model metastability affects the mixing time on *both* sides of λ_c. Note that our results leave open the behavior of the mixing time exactly at λ_s and λ_S.

As a byproduct of our main results above, we deduce new bounds on the mixing time of *local* dynamics for the random-cluster model (i.e., dynamics that modify only a constant-size region of the configuration at each step). For definiteness we consider the canonical *heat-bath (HB) dynamics*, which in each step updates a single edge of the current configuration (V, A) as follows:

(i) pick an edge $e \in E$ u.a.r;
(ii) replace A by $A \cup \{e\}$ with probability $\frac{\mu_{p,q}(A\cup\{e\})}{\mu_{p,q}(A\cup\{e\})+\mu_{p,q}(A\setminus\{e\})}$, else by $A \setminus \{e\}$.

Local dynamics for the random-cluster model are currently very poorly understood (but see [12] for the special case of graphs with bounded tree-width). However, in a recent surprising development, Ullrich [21, 22] showed that the mixing time of the heat-bath dynamics on any graph differs from that of the SW dynamics by at most a poly(n) factor. Thus the previously known bounds for SW translate to bounds for the heat-bath dynamics for integer q. By adapting Ullrich's technology to our CM setting, we are able to obtain a similar translation of our results, thus establishing the first non-trivial bounds on the mixing time of the mean-field heat-bath dynamics for all $q > 1$.

▶ **Theorem 3.** *For any $q > 1$, the mixing time of the heat-bath dynamics for the mean-field random-cluster model is $\tilde{O}(n^4)$ for $\lambda \notin [\lambda_s, \lambda_S]$, and $e^{\Omega(\sqrt{n})}$ for $\lambda \in (\lambda_s, \lambda_S)$.*

The $\tilde{O}$ here hides polylogarithmic factors. We conjecture that the upper bound should be $\tilde{O}(n^2)$ for all $\lambda \notin [\lambda_s, \lambda_S]$; the additional n^2 factor is inherent in Ullrich's spectral approach.

We conclude this introduction with some brief remarks about our techniques. Both our upper and lower bounds on the mixing time of the CM dynamics focus on the evolution of the one-dimensional random process given by the size of the largest component (which approaches θn for $\lambda > \lambda_c$ and $\Theta(\log n)$ for $\lambda < \lambda_c$). A key ingredient in our analysis is a function that describes the expected change, or "drift", of this random process at each step; the critical points λ_s and λ_S discussed above arise naturally from consideration of the zeros of this drift function.

For our upper bounds, we construct a multiple-phase coupling of the evolution of two arbitrary configurations, showing that they converge in $O(\log n)$ steps; this coupling is similar in flavor to that used by Long *et al.* [18] for the SW dynamics for $q = 2$, but there are additional complexities in that our analysis has to identify the "slow mixing" window (λ_s, λ_S) for $q > 2$, and also has to contend with the fact that only a subset of the vertices (rather than the whole graph, as in SW) are active at each step. This latter issue is handled using precise concentration bounds for the number of active vertices, tailored estimates for the component structure of random graphs and a new coupling for pairs of binomial random variables.

For our exponential lower bounds we use the drift function to identify the metastable states mentioned ealier from which the dynamics cannot easily escape. For both upper and lower bounds, we have to handle the sub-critical and super-critical cases, $\lambda < \lambda_c$ and $\lambda > \lambda_c$, separately, even though our final results are insensitive to λ_c, because the structure of typical configurations differs in the two cases.

2 Preliminaries

In this section we gather a number of standard definitions and background results that we will refer to repeatedly in our proofs. For those results that are not available in the literature, we provide proofs in the full version of this paper [2].

Mixing time. Let P be the transition matrix of a finite, ergodic Markov chain M with state space Ω and stationary distribution π. The mixing time of M is defined by

$$\tau_{\text{mix}} = \max_{z\in\Omega} \min_t \left\{ ||P^t(z,\cdot) - \pi(\cdot)||_{\text{TV}} \leq 1/4 \right\}$$

where $||\mu-\nu||_{\text{TV}} = \max_{A\subset\Omega} |\mu(A)-\nu(A)|$ is the total variation distance between distributions μ and ν.

A *(one step) coupling* of the Markov chain M specifies for every pair of states $(X_t, Y_t) \in \Omega^2$ a probability distribution over (X_{t+1}, Y_{t+1}) such that the processes $\{X_t\}$ and $\{Y_t\}$, viewed in isolation, are faithful copies of M, and if $X_t = Y_t$ then $X_{t+1} = Y_{t+1}$. The *coupling time* is defined by

$$T_{\text{coup}} = \max_{x,y\in\Omega} \min_t \{X_t = Y_t | X_0 = x, Y_0 = y\}.$$

For any $\delta \in (0,1)$, the following standard inequality (see, e.g., [17]) provides a bound on the mixing time:

$$\tau_{\text{mix}} \leq \min_t \left\{ \Pr[T_{\text{coup}} > t] \leq 1/4 \right\} \leq O\left(\delta^{-1}\right) \cdot \min_t \left\{ \Pr[T_{\text{coup}} > t] \leq 1-\delta \right\}. \tag{1}$$

Random graphs. Let G_d be distributed as a $G(n, p = d/n)$ random graph where $d > 0$. Let $\mathcal{L}(G_d)$ denote the largest component of G_d and let $L_i(G_d)$ be the *size* of the i-th largest component of G_d. (Thus, $L_1(G_d) = |\mathcal{L}(G_d)|$.) In our proofs we will use several facts about the random variables $L_i(G_d)$, which we gather here for convenience.

▶ **Lemma 4** ([18, Lem. 5.7]). *Let $I(G_d)$ denote the number of isolated vertices in G_d. If $d = O(1)$, then there exists a constant $A > 0$ such that $\Pr[I(G_d) > An] = 1 - O\left(n^{-1}\right)$.*

▶ **Lemma 5.** *If $d = O(1)$, then $L_2(G_d) < 2n^{11/12}$ with probability $1 - O\left(n^{-1/12}\right)$ for sufficiently large n.*

▶ **Lemma 6** ([7, Lem. 7]). *If $d < 1$ is bounded away[2] from 1, then $L_1(G_d) = O(\log n)$ with probability $1 - O\left(n^{-1}\right)$.*

For $d > 1$, let $\beta = \beta(d)$ be the unique positive root of the equation

$$e^{-dx} = 1 - x. \tag{2}$$

▶ **Lemma 7.** *Let $\widetilde{G}_{d_n}$ be distributed as a $G(n+m, d_n/n)$ random graph where $\lim_{n\to\infty} d_n = d$ and $|m| = o(n)$. Assume $1 < d_n = O(1)$ and d_n is bounded away from 1 for all $n \in \mathbb{N}$. Then, $L_2(\widetilde{G}_{d_n}) = O(\log n)$ with probability $1 - O\left(n^{-1}\right)$, and for $A = o(\log n)$ and n large enough, there exists a constant $c > 0$ such that*

$$\Pr[|L_1(\widetilde{G}_{d_n}) - \beta(d)n| > |m| + A\sqrt{n}] \leq e^{-cA^2}.$$

▶ **Corollary 8.** *With the same notation as in Lemma 7, $|\operatorname{E}[L_1(\widetilde{G}_{d_n})] - \beta(d)n| < |m| + O(\sqrt{n})$.*

▶ **Lemma 9** ([13, Lem. 6]). *If $d < 1$ is bounded away from 1, then $L_1(G_d) = O(\sqrt{n})$ with probability $1 - e^{-\Omega(\sqrt{n})}$.*

[2] We say that d is bounded away from a if there exists a constant ξ such that $|d - a| \geq \xi$.

The random-cluster model. Recall from the introduction that the mean-field random-cluster model exhibits a phase transition at $\lambda = \lambda_c(q)$ (see [3]): in the sub-critical regime $\lambda < \lambda_c$ the largest component is of size $O(\log n)$, while in the super-critical regime $\lambda > \lambda_c$ there is a unique giant component of size $\sim\theta_r n$, where $\theta_r = \theta_r(\lambda, q)$ is the largest $x > 0$ satisfying the equation

$$e^{-\lambda x} = 1 - \frac{qx}{1 + (q-1)x}. \tag{3}$$

(Note that, as expected, this equation is identical to (2) when $q = 1$.)

Drift function. As indicated in the introduction, our analysis relies heavily on understanding the evolution of the size of the largest component under the CM dynamics. To this end, for fixed λ and q let $\phi(\theta)$ be the largest $x > 0$ satisfying the equation

$$e^{-\lambda x} = 1 - \frac{qx}{1 + (q-1)\theta}. \tag{4}$$

This equation corresponds to (2) for a $G(\alpha n, \lambda/n)$ random graph where $\alpha = (1+(q-1)\theta)q^{-1}$. Thus, $\phi(\theta) = \beta(\alpha\lambda)$ and consequently ϕ is well-defined when $\alpha\lambda > 1$. In particular, ϕ is well-defined in the interval $(\theta_{\min}, 1]$, where $\theta_{\min} = \max\{(q-\lambda)/\lambda(q-1), 0\}$.

We will see in Section 3 that for a configuration with a unique "large" component of size θn, the expected "drift" in the size of the largest component will be determined by the sign of the function $f(\theta) = \theta - \phi(\theta)$: $f(\theta) > 0$ corresponds to a negative drift and $f(\theta) < 0$ to a positive drift. Thus, let

$$\lambda_s = \max\{\lambda \le \lambda_c : f(\theta) > 0 \ \ \forall\theta \in (\theta_{\min}, 1]\}, \text{ and}$$

$$\lambda_S = \min\{\lambda \ge \lambda_c : f(\theta)(\theta - \theta_r) > 0 \ \ \forall\theta \in (\theta_{\min}, 1]\}.$$

In words, λ_s and λ_S are the maximum and minimum values, respectively, of λ for which the drift in the size of the largest component is always in the right direction (i.e., towards 0 in the sub-critical case and towards $\theta_r n$ in the super-critical case). The following lemma reveals basic information about these quantities.

▶ **Lemma 10.** *For $q \le 2$, $\lambda_s = \lambda_c = \lambda_S = q$; and for $q > 2$, $\lambda_s < \lambda_c < \lambda_S = q$.*

For integer $q \ge 3$, λ_s corresponds to the threshold β_s in the mean-field q-state Potts model at which the local (Glauber) dynamics undergoes an exponential slowdown [8]. In fact, a change of variables reveals that $\lambda_s = 2\beta_s$ for the specific mean-field Potts model normalization in [8].

In Figure 1 we sketch f in its only two qualitatively different regimes: $q \le 2$ and $q > 2$. The following lemma provides bounds for the drift of the size of the largest component under CM steps.

▶ **Lemma 11.** *For all $\theta \in (\theta_{\min}, 1]$,*

(i) *If $\lambda < \lambda_s$, there exists a constant $\delta > 0$ such that $f(\theta) \ge \delta$.*
(ii) *When $\lambda > \lambda_S$, if $\theta > \theta_r$, then $\theta \ge \phi(\theta) \ge \theta_r$ and if $\theta < \theta_r$, then $\theta \le \phi(\theta) \le \theta_r$.*
(iii) *If $\lambda > \lambda_S$, there exists a constant $\delta \in (0, 1)$ such that $\delta|\theta - \theta_r| \le |\phi(\theta) - \theta|$.*

Binomial coupling. In our coupling constructions we will use the following fact about the coupling of two binomial random variables.

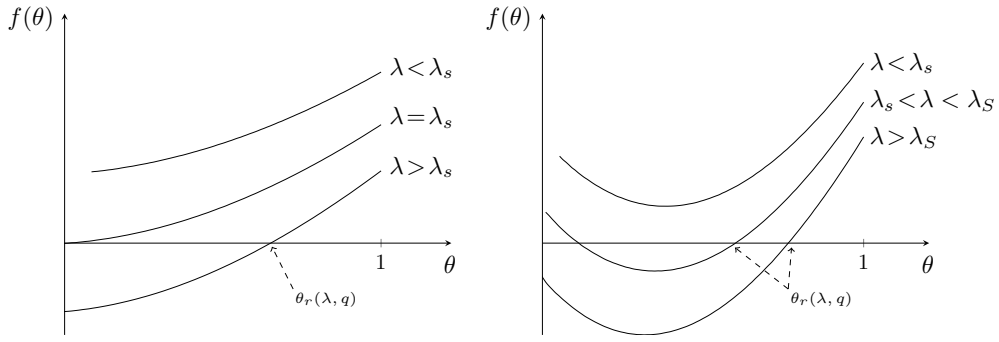

Figure 1 Sketch of the function f. (Left figure corresponds to $q \leq 2$ and right figure to $q > 2$.)

▶ **Lemma 12.** *Let X and Y be binomial random variables with parameters m and r, where $r \in (0,1)$ is a constant. Then, for any $y \in \mathbb{N}$, there exists a coupling (X, Y) such that for a suitable constant $\gamma = \gamma(r) > 0$,*

$$\Pr[X - Y = y] \geq 1 - \frac{\gamma y}{\sqrt{m}}.$$

Moreover if $y = a\sqrt{m}$ for a fixed constant a, then $\gamma a < 1$.

Hitting time for supermartingales. We will require the following easily derived hitting time estimate.

▶ **Lemma 13.** *Consider the stochastic process $\{Z_t\}$ such that $Z_t \in [-n, n]$ for all $t \geq 0$. Assume $Z_0 > a$ for some $a \in [-n, n]$ and let $T = \min\{t > 0 : Z_t \leq a\}$. Suppose $\mathrm{E}[Z_{t+1} - Z_t | \mathcal{F}_t] \leq -A$, where $A > 0$ and $\mathcal{F}_t$ is the history of the first t steps. Then, $\mathrm{E}[T] \leq 4n/A$.*

3 Mixing time upper bounds

In this section we prove the upper bound portion of Theorem 1 from the introduction.

▶ **Theorem 14.** *Consider the CM dynamics for the mean-field random-cluster model with parameters $p = \lambda/n$ and q where $\lambda > 0$ and $q > 1$ are constants independent of n. If $\lambda \notin [\lambda_s, \lambda_S]$, then $\tau_{\mathrm{mix}} = O(\log n)$.*

Proof Sketch. Consider two copies $\{X_t\}$ and $\{Y_t\}$ of the CM dynamics starting from two arbitrary configurations X_0 and Y_0. We design a coupling (X_t, Y_t) of the CM steps and show that $\Pr[X_T = Y_T] = \Omega(1)$ for some $T = O(\log n)$; the result then follows from (1). The coupling consists of four phases. In the first phase $\{X_t\}$ and $\{Y_t\}$ are run independently. In Section 3.1 we establish that after $O(\log n)$ steps $\{X_t\}$ and $\{Y_t\}$ each have at most one large component with probability $\Omega(1)$. We call a component *large* if it contains at least $2n^{11/12}$ vertices; otherwise it is *small*.

In the second phase, $\{X_t\}$ and $\{Y_t\}$ also evolve independently. In Sections 3.2 and 3.3 we show that, conditioned on the success of Phase 1, after $O(\log n)$ steps with probability $\Omega(1)$ the largest components in $\{X_t\}$ and $\{Y_t\}$ have sizes close to their expected value: $O(\log n)$ in the sub-critical case and $\sim\theta_r n$ in the super-critical case. In the third phase, $\{X_t\}$ and $\{Y_t\}$ are coupled to obtain two configurations with the same component structure. This

coupling, described in Section 3.4, makes crucial use of the binomial coupling of Section 2, and conditioned on a successful conclusion of Phase 2 succeeds with probability $\Omega(1)$ after $O(\log n)$ steps. In the last phase, a straightforward coupling is used to obtain two identical configurations from configurations with the same component structure. This coupling is described in Section 3.5 and succeeds w.h.p. after $O(\log n)$ steps, conditioned on the success of the previous phases.

Putting all this together, there exists a coupling (X_t, Y_t) such that, after $T = O(\log n)$ steps, $X_T = Y_T$ with probability $\Omega(1)$. The reminder of this section fleshes out the above proof sketch. ◀

We now introduce some notation that will be used throughout the rest of the paper. As before, we will use $\mathcal{L}(X_t)$ for the largest component in X_t and $L_i(X_t)$ for the *size* of the i-th largest component of X_t. (Thus, $L_1(X_t) = |\mathcal{L}(X_t)|$.) For convenience, we will sometimes write $\theta_t n$ for $L_1(X_t)$. Also, we will use $\mathcal{E}_t$ for the event that $\mathcal{L}(X_t)$ is activated, and A_t for the number of activated vertices at time t.

3.1 Convergence to configurations with a unique large component

▶ **Lemma 15.** *For any starting random-cluster configuration X_0, there exists $T = O(\log n)$ such that X_T has at most one large component with probability $\Omega(1)$.*

Proof. Let N_t be the number of new large components created in sub-step (iii) of the CM dynamics at time t. If $A_t < 2n^{11/12}$, then $N_t = 0$. Together with Lemma 5 this implies that $\Pr[N_t > 1 | X_t, A_t = a] \le a^{-1/12}$ for all $a \in [0, n]$. Thus,

$$\begin{aligned}
\mathrm{E}[N_t|X_t] &= \sum_{a=0}^{n} \mathrm{E}[N_t|X_t, A_t = a] \Pr[A_t = a|X_t] \\
&\le \sum_{a=0}^{n} \left(\Pr[N_t \le 1|X_t, A_t = a] + \frac{a}{2n^{11/12}} \Pr[N_t > 1|X_t, A_t = a] \right) \Pr[A_t = a|X_t] \\
&\le \sum_{a=0}^{n} \left(1 + \frac{a}{2n^{11/12}} \frac{1}{a^{1/12}} \right) \Pr[A_t = a|X_t] \le 2.
\end{aligned}$$

Let K_t be the number of large components in X_t and let C_t be the number of activated large components in sub-step (i) of the CM dynamics at time t. Then,

$$\mathrm{E}[K_{t+1}|X_t] = K_t - \mathrm{E}[C_t|X_t] + \mathrm{E}[N_t|X_t] \le K_t - \frac{K_t}{q} + 2 \le \left(1 - \frac{1}{2q}\right) K_t$$

provided $K_t \ge 4q$. Assuming that $K_t \ge 4q$ for all $t < T$, we have

$$\mathrm{E}[K_T|X_0] \le \left(1 - \frac{1}{2q}\right)^T K_0.$$

Hence, Markov's inequality implies that $K_T < 4q$ w.h.p. for some $T = O(\log n)$. If at time T the remaining K_T large components become active, then $K_{T+1} \le 1$ w.h.p. by Lemma 5. All K_T components become active simultaneously with probability at least q^{-4q} and thus $K_{T+1} \le 1$ with probability $\Omega(1)$, as desired. ◀

3.2 Convergence to typical configurations: the sub-critical case

▶ **Lemma 16.** *Let $\lambda < \lambda_s$; if X_0 has at most one large component, then there exists $T = O(\log n)$ such that $L_1(X_T) = O(\log n)$ with probability $\Omega(1)$.*

The following fact will be used in the proof. Let $\xi = \sqrt{2n^{23/12}\log n}$.

▶ **Fact 17.** *If X_t has at most one large component, then for sufficiently large n each of the following holds with probability $1 - O(n^{-1})$:*

(i) *If $\mathcal{L}(X_t)$ is inactive, then all new components in X_{t+1} have size $O(\log n)$.*

(ii) *If $\mathcal{L}(X_t)$ is active, then $A_t \in J_t := \left[L_1(X_t) + \frac{n-L_1(X_t)}{q} - \xi, L_1(X_t) + \frac{n-L_1(X_t)}{q} + \xi\right]$.*

(iii) *If there is no large component in X_t, then the largest new component in X_{t+1} have size $O(\log n)$.*

Proof of Lemma 16. If X_0 has at most one large component, then it is easy to check that X_t retains this property for $O(\log n)$ CM steps w.h.p. Thus, we condition on this event throughout this phase. We show first that one step of the CM dynamics contracts the size of the largest component in expectation.

For ease of notation set $\Theta_{\mathrm{s}} := \theta_{\min}$, with $\theta_{\min}$ defined as in Section 2. Note that $(\Theta_{\mathrm{s}} + (1-\Theta_{\mathrm{s}})q^{-1})\lambda = 1$. Hence, if $L_1(X_t) = \Theta_{\mathrm{s}} n$ and $\mathcal{L}(X_t)$ is activated, then the percolation step (sub-step (iii) of the CM dynamics) is critical with non-negligible probability. This makes the analysis in the neighborhood of $\Theta_{\mathrm{s}} n$ more delicate.

We consider first the case where $\theta_t \geq \Theta_{\mathrm{s}} + \varepsilon$ for some small constant $\varepsilon > 0$ to be chosen later. By Fact 17(i), if $\mathcal{L}(X_t)$ is inactive all the new components have size $O(\log n)$ with probability $1 - O(n^{-1})$. Thus,

$$\mathrm{E}[L_1(X_{t+1}) \,|\, X_t, \neg\mathcal{E}_t] \leq L_1(X_t) + O(1) = \theta_t n + O(1). \tag{5}$$

To bound $\mathrm{E}[L_1(X_{t+1}) \,|\, X_t, \mathcal{E}_t]$, let $h^+(\theta_t) = \theta_t n + (1-\theta_t)q^{-1}n + \xi$ and let $\ell^+(\theta_t)$ be a random variable distributed as the size of the largest component of a $G(h^+(\theta_t), p)$ random graph. Then, by Fact 17(ii) we have

$$\begin{aligned}\mathrm{E}[L_1(X_{t+1}) \,|\, X_t, \mathcal{E}_t] &\leq \sum_{a \in J_t} \mathrm{E}[L_1(X_{t+1}) \,|\, X_t, \mathcal{E}_t, A_t = a] \Pr[A_t = a \,|\, X_t, \mathcal{E}_t] + O(1) \\ &\leq \mathrm{E}[L_1(X_{t+1}) \,|\, X_t, \mathcal{E}_t, A_t = h^+(\theta_t)] + O(1) = \mathrm{E}[\ell^+(\theta_t)] + O(1).\end{aligned}$$

When $\theta_t \geq \Theta_{\mathrm{s}} + \varepsilon$, $G(h^+(\theta_t), p)$ is a super-critical random graph. Thus, Corollary 8 implies

$$\mathrm{E}[L_1(X_{t+1}) \mid X_t, \mathcal{E}_t] \leq \phi(\theta_t) n + O(\xi), \tag{6}$$

where $\phi(\theta_t)$ is defined as in (4). Since $\lambda < \lambda_s$, by Lemma 11 there exists a constant $\delta > 0$ such that $\theta_t - \phi(\theta_t) \geq \delta$. Therefore, putting (5) and (6) together, we have

$$\mathrm{E}[L_1(X_{t+1}) \mid X_t] \leq (1-q^{-1})\theta_t n + q^{-1}\phi(\theta_t) n + O(\xi) \leq \theta_t n - \delta q^{-1} n + O(\xi). \tag{7}$$

As mentioned before, in a close neighborhood of Θ_{s} the percolation step is critical with non-negligible probability, so when $\theta_t \in (\Theta_{\mathrm{s}} - \varepsilon, \Theta_{\mathrm{s}} + \varepsilon)$ we use monotonicity to simplify the analysis. Namely, we assume that $\theta_t = \Theta_{\mathrm{s}} + \varepsilon$ and use the previous steps to obtain (7). Thus, there exists a constant $\gamma > 0$ such that for all $\theta_t > \Theta_{\mathrm{s}} - \varepsilon$:

$$\mathrm{E}[L_1(X_{t+1}) - L_1(X_t) \,|\, X_t] \leq -\gamma n.$$

Let $\tau = \min\{t > 0 : L_1(X_t) \leq (\Theta_{\mathrm{s}} - \varepsilon)n\}$. Lemma 13 implies $\mathrm{E}[\tau] \leq 4/\gamma$ and thus $\Pr[\tau > 8/\gamma] \leq 1/2$ by Markov's inequality. Hence, $L_1(X_T) \leq (\Theta_{\mathrm{s}} - \varepsilon)n$ for some $T = O(1)$ with probability $\Omega(1)$.

To conclude, we show that after $O(\log n)$ additional steps the largest component has size $O(\log n)$ with probability $\Omega(1)$. If $L_1(X_T) \le (\Theta_s - \varepsilon)n$ and $\mathcal{L}(X_T)$ is activated, then the definition of Θ_s implies that the percolation step of the CM dynamics is sub-critical, and thus X_{T+1} has no large component w.h.p. Hence, X_{T+1} has no large component with probability $\Omega(1)$. Now, by Fact 17(iii) and a union bound, all the new components created during the $O(\log n)$ steps immediately after time $T+1$ have size $O(\log n)$ w.h.p. Another union bound over components shows that during these $O(\log n)$ steps, every component in X_{T+1} is activated w.h.p. Thus, after $O(\log n)$ steps the largest component in the configuration has size $O(\log n)$ with probability $\Omega(1)$, which establishes Lemma 16. ◀

The reader is referred to the full version [2] for the proof of Fact 17, as well as some of the details omitted from the proof of Lemma 16.

3.3 Convergence to typical configurations: the super-critical case

▶ **Lemma 18.** *Let $\lambda > \lambda_S = q$ and $\Delta_t := |L_1(X_t) - \theta_r n|$. If X_0 has at most one large component, then for some $T = O(\log n)$ there exists a constant $c > 0$ such that $\Pr[\,\Delta_T > A\sqrt{cn}\,] < 1/A$ for all $A > 0$.*

The following facts, whose proofs can be found in the full version [2], will be useful. Let $\xi(r) = \sqrt{nr\log n}$, $\Theta_S := 1 - q/\lambda$ and $\mu_t = L_1(X_t) + \frac{n - L_1(X_t)}{q}$.

▶ **Fact 19.** *If X_0 has at most one large component, then there exists $T = O(\log n)$ such that with probability $\Omega(1)$: $L_1(X_T) > (\Theta_S + \varepsilon)n$, $L_2(X_T) = O(\log n)$ and $\sum_{j\ge 2} L_j(X_T)^2 = O(n)$. Moreover, once these properties are obtained they are preserved for a further $T' = O(\log n)$ CM steps w.h.p.*

▶ **Fact 20.** *Assume X_t has exactly one large component and all its other components have size at most $r < 2n^{11/12}$. Then, for a small constant $\varepsilon > 0$ and sufficiently large n, each of the following holds with probability $1 - O\left(n^{-1}\right)$:*

- **(i)** *If $\mathcal{L}(X_t)$ is inactive and $L_1(X_t) > (\Theta_S + \varepsilon)n$, then all new components in X_{t+1} have size $O(\log n)$.*
- **(ii)** *If $\mathcal{L}(X_t)$ is active, then $A_t \in J_{t,r} := [\mu_t - \xi(r), \mu_t + \xi(r)]$ and $G(A_t, p)$ is a super-critical random graph.*

Proof of Lemma 18. We show that one step of the CM dynamics contracts Δ_t in expectation. Observe that by Fact 19 we may assume X_0 is such that $L_1(X_0) > (\Theta_S + \varepsilon)n$, $L_2(X_0) = O(\log n)$ and $\sum_{j\ge 2} L_j(X_0)^2 = O(n)$, and that X_t retains these properties for the $O(\log n)$ steps of this phase w.h.p. Consequently, if $\mathcal{L}(X_t)$ is inactive, then $L_1(X_{t+1}) = L_1(X_t)$ with probability $1 - O(n^{-1})$ by Fact 20(i). Hence,

$$\mathrm{E}[\Delta_{t+1} \mid X_t, \neg\mathcal{E}_t] \le \mathrm{E}\left[\,|L_1(X_{t+1}) - L_1(X_t)| \mid X_t, \neg\mathcal{E}_t\right] + |L_1(X_t) - \theta_r n| \le \Delta_t + O(1). \quad (8)$$

To bound $\mathrm{E}[\Delta_{t+1} \mid X_t, \mathcal{E}_t]$, let $M_t = A_t - \mu_t$ and let $\ell_t(m)$ denote the size of the largest component of a $G(\mu_t + m, p)$ random graph. Also, let $\Delta'_{t+1} := |L_1(X_{t+1}) - \phi(\theta_t)n|$. Note that, conditioned on $M_t = m$, $L_1(X_{t+1})$ and $\ell_t(m)$ have the same distribution. Moreover, if $A_t \in J_{t,r}$ then $M_t \in J'_{t,r} := [-\xi(r), \xi(r)]$. Hence, Fact 20(ii) with $r = O(\log n)$ implies

$$\begin{aligned} E[\Delta'_{t+1} \mid X_t, \mathcal{E}_t] &\le \sum_{m \in J'_{t,r}} \mathrm{E}[\Delta'_{t+1} \mid X_t, \mathcal{E}_t, M_t = m]\ \Pr[M_t = m \mid X_t, \mathcal{E}_t] + O(1) \\ &= \sum_{m \in J'_{t,r}} \mathrm{E}[|\ell_t(m) - \phi(\theta_t)n|]\ \Pr[M_t = m \mid X_t, \mathcal{E}_t] + O(1). \end{aligned}$$

Now, by Fact 20(ii), $G(\mu_t + m, p)$ is a super-critical random graph, and thus $\mathrm{E}[|\ell_t(m) - \phi(\theta_t)n|] \le |m| + O(\sqrt{n})$ by Corollary 8. Hence,

$$E[\Delta'_{t+1} \mid X_t, \mathcal{E}_t] \le \mathrm{E}[|M_t| \mid X_t, \mathcal{E}_t] + O(\sqrt{n}).$$

Since $\sum_{j\ge 2} L_j(X_t)^2 = O(n)$, it follows from Hoeffding's inequality that $\mathrm{E}[|M_t| \,|\, X_t, \mathcal{E}_t] = O(\sqrt{n})$ (the explicit calculation is provided in the full version [2]), and thus $E[\Delta'_{t+1} \,|\, X_t, \mathcal{E}_t] = O(\sqrt{n})$. The triangle inequality then implies

$$\mathrm{E}[\Delta_{t+1} \mid X_t, \mathcal{E}_t] \;\le\; \mathrm{E}[\Delta'_{t+1} \mid X_t, \mathcal{E}_t] + |\theta_r - \phi(\theta_t)|n \;\le\; |\theta_r - \phi(\theta_t)|n + O(\sqrt{n}). \tag{9}$$

Putting (8) and (9) together, we have

$$\mathrm{E}[\Delta_{t+1} \mid X_t] \le (1 - q^{-1})\Delta_t + q^{-1}|\theta_r - \phi(\theta_t)|n + O(\sqrt{n}).$$

By Lemma 11(iii), there exists a constant $\delta \in (0, 1)$ such that $\delta|\theta_t - \theta_r| \le |\theta_t - \phi(\theta_t)|$. Together Lemma 11(ii), this implies $|\theta_r - \phi(\theta_t)| \le (1 - \delta)|\theta_t - \theta_r|$. Thus, there exists a constant $\delta' > 0$ such that

$$\mathrm{E}[\Delta_{t+1} \mid X_t] \le (1 - \delta')\Delta_t + \xi$$

where $\xi = O(\sqrt{n})$. Inducting, $\mathrm{E}[\Delta_t] \le (1 - \delta')^t\Delta_0 + \xi/\delta'$. Hence, for some $t = O(\log n)$, $\mathrm{E}[\Delta_t] = O(\sqrt{n})$ and so Markov's inequality implies $\Pr[\Delta_t > A\sqrt{cn}] \le 1/A$ for some constant $c > 0$ and any $A > 0$. ◀

3.4 Coupling to the same component structure

In this section we design a coupling of the CM steps which, starting from two configurations with certain properties (namely, those obtained in Sections 3.2 and 3.3 for the sub-critical and super-critical case respectively), quickly converges to a pair of configurations with the same component structure. (We say that two random-cluster configurations X and Y have the same component structure if $L_j(X) = L_j(Y)$ for all $j \ge 1$.)

The only additional property we will require is that the starting configurations should have a linear number of isolated vertices. Although in Sections 3.2 and 3.3 we do not guarantee this, observe that a single CM step from a configuration with at most one large component activates a linear number of vertices w.h.p., and thus Lemma 4 implies that the new configuration has a linear number of isolated vertices w.h.p. We will focus first on the super-critical case, since a simplified version of the arguments works in the sub-critical case.

▶ **Lemma 21.** *Let $\lambda > q$ and let X_0, Y_0 be random-cluster configurations with $\Omega(n)$ isolated vertices such that: $L_2(X_0) = O(\log n)$, $|L_1(X_0) - \theta_r n| = O(\sqrt{n}\log^2 n)$, $\sum_{j\ge 2} L_j(X_0)^2 = O(n)$ and similarly for Y_0. Then, there exists a coupling of the CM steps such that X_T and Y_T have the same component structure after $T = O(\log n)$ steps with probability $\Omega(1)$.*

Proof. It is straightforward to check that X_t, Y_t retain the above structural properties of X_0, Y_0 for $O(\log n)$ CM steps w.h.p. (The details are provided in the full version [2].)

Our coupling will be a composition of three couplings. Coupling I contracts a certain notion of distance between $\{X_t\}$ and $\{Y_t\}$. This contraction will boost the probability of success of the other two couplings. Coupling II is a one-step coupling which guarantees that the largest components from $\{X_t\}$ and $\{Y_t\}$ have the same size with probability $\Omega(1)$. Coupling III uses the binomial coupling from Lemma 12 to achieve two configurations with the same component structure with probability $\Omega(1)$.

Coupling I: Excluding $\mathcal{L}(X_t)$ and $\mathcal{L}(Y_t)$, consider a maximal matching W_t between the components of X_t and Y_t with the restriction that only components of equal size are matched to each other. Let $M(X_t)$ and $M(Y_t)$ be the components in the matching from X_t and Y_t respectively. Let $D(X_t)$ and $D(Y_t)$ be the complements of $\mathcal{L}(X_t) \cup M(X_t)$ and $\mathcal{L}(Y_t) \cup M(Y_t)$ respectively, and let $d_t := |D(X_t)| + |D(Y_t)|$ where $|\cdot|$ denotes the total number of vertices in the respective components.

The activation of the components in $M(X_t)$ and $M(Y_t)$ is coupled using the matching W_t. That is, $c \in M(X_t)$ and $W_t(c) \in M(Y_t)$ are activated simultaneously with probability $1/q$. The activations of $\mathcal{L}(X_t)$ and $\mathcal{L}(Y_t)$ are also coupled, and the components in $D(X_t)$ and $D(Y_t)$ are activated independently. Let $A(X_t)$ and $A(Y_t)$ denote the set of active *vertices* in X_t and Y_t respectively, and w.l.o.g. assume $|A(X_t)| \geq |A(Y_t)|$. Let R_t be an arbitrary subset of $A(X_t)$ such that $|R_t| = |A(Y_t)|$ and let $Q_t = A(X_t) \setminus R_t$. The percolation step is coupled by establishing an arbitrary vertex bijection $b_t : R_t \to A(Y_t)$ and coupling the re-sampling of each edge $(u, v) \in R_t \times R_t$ with $(b_t(u), b_t(v)) \in A(Y_t) \times A(Y_t)$. Edges within Q_t and in the cut $C_t = R_t \times Q_t$ are re-sampled independently. The following claim establishes the desired contraction in d_t.

▶ **Claim 22.** *Let $\omega(n) = n/\log^4 n$; after $T = O(\log\log n)$ steps, $d_T \leq \omega(n)$ w.h.p.*

Proof. Let $D_a(X_t)$ and $D_a(Y_t)$ be the number of active vertices from $D(X_t)$ and $D(Y_t)$ respectively, and let $\mathcal{F}_t$ be the history of the first t steps. Observe that Coupling I guarantees that R_t and $A(Y_t)$ will have the same component structure internally. However, the vertices in Q_t will contribute to d_{t+1} unless they are part of the new large component, and each edge in C_t could increase d_{t+1} by at most (twice) the size of one component of R_t, which is $O(\log n)$. Thus,

$$\mathrm{E}[d_{t+1} \mid A(X_t), A(Y_t), C_t, \mathcal{F}_t] \leq d_t - (|D_a(X_t)| + |D_a(Y_t)|) + |Q_t| + 2|C_t| \times O(\log n). \quad (10)$$

Observe that $\mathrm{E}[|D_a(X_t)| + |D_a(Y_t)| \mid \mathcal{F}_t] = d_t/q$, and $\mathrm{E}[|C_t| \mid A(X_t), A(Y_t), \mathcal{F}_t] = |R_t||Q_t|p \leq \lambda|Q_t|$. Since $|Q_t| = O(\sqrt{n}\log^2 n)$, taking expectations in (10) we get

$$\mathrm{E}[d_{t+1} \mid \mathcal{F}_t] \leq d_t - \frac{d_t}{q} + O\left(\sqrt{n}\log^3 n\right) \leq \left(1 - \frac{1}{2q}\right) d_t$$

provided $d_t > \omega(n)$. Thus, Markov's inequality implies $d_T \leq \omega(n)$ for some $T = O(\log\log n)$ w.h.p. Note that for larger values of T, this argument immediately provides stronger bounds for d_T, but neither our analysis nor the order of the coupling time benefits from this. ◀

Coupling II: Assume now that $d_t \leq \omega(n)$ and let $I_{\mathrm{M}}(X_t)$ and $I_{\mathrm{M}}(Y_t)$ denote the isolated vertices in $M(X_t)$ and $M(Y_t)$ respectively. The activation in $X_t \setminus I_{\mathrm{M}}(X_t)$ and $Y_t \setminus I_{\mathrm{M}}(Y_t)$ is coupled as in Coupling I, except we condition on the event that $\mathcal{L}(X_t)$ and $\mathcal{L}(Y_t)$ are activated, which occurs with probability $1/q$. This first part of the activation could activate a different number of vertices from each copy of the chain; let ρ_t be this difference.

First we show that $\rho_t = O(\sqrt{n})$ with probability $\Omega(1)$. By Lemma 18 (with $A = 2$), we have $|L_1(X_t) - L_1(Y_t)| = O(\sqrt{n})$ with probability $\Omega(1)$. If this is the case, then $||D(X_t)| - |D(Y_t)|| = O(\sqrt{n})$. Also, since $\sum_{j\geq 2} L_j(X_t)^2 = O(n)$ and $\sum_{j\geq 2} L_j(Y_t)^2 = O(n)$, by Hoeffding's inequality the numbers of active vertices from $D(X_t)$ and $D(Y_t)$ differ by at most $O(\sqrt{n})$ with probability $\Omega(1)$. Thus, $\rho_t = O(\sqrt{n})$ with probability $\Omega(1)$.

Now we show how to couple the activation in $I_{\mathrm{M}}(X_t)$, $I_{\mathrm{M}}(Y_t)$ in a way such that $|A(X_t)| = |A(Y_t)|$ with probability $\Omega(1)$. The number of active isolated vertices from $I_{\mathrm{M}}(X_t)$ is binomially distributed with parameters $|I_{\mathrm{M}}(X_t)|$ and $1/q$, and similarly for $I_{\mathrm{M}}(Y_t)$. Hence, the activation

of the isolated vertices may be coupled using the binomial coupling from Section 2. Since $|I_{\text{M}}(X_t)|=|I_{\text{M}}(Y_t)|=\Omega(n)$ and $\rho_t=O(\sqrt{n})$, Lemma 12 implies that this coupling corrects the difference ρ_t with probability $\Omega(1)$. If this is the case, then by coupling the edge sampling bijectively as in Coupling I, we ensure that $L_1(X_{t+1}) = L_1(Y_{t+1})$ and $d_{t+1} \leq \omega(n)$ with probability $\Omega(1)$.

Coupling III: Assume $L_1(X_0) = L_1(Y_0)$ and $d_0 \leq \omega(n)$. The component activation is coupled as in Coupling II, but we do not require the two large components to be active; rather, we just couple their activation together.

If $L_1(X_t) = L_1(Y_t)$, then $|D(X_t)| = |D(Y_t)|$ and thus the expected number of active vertices from $D(X_t)$ and $D(Y_t)$ is the same. Consequently, since $d_t \leq \omega(n)$, Hoeffding's inequality implies $\rho_t = O\left(\sqrt{n}\log^{-1} n\right)$ w.h.p. Let $\mathcal{F}_t$ be the event that the coupling of the isolated vertices succeeds in correcting the error ρ_t. Since $|I_{\text{M}}(X_t)| = |I_{\text{M}}(Y_t)| = \Omega(n)$, $\mathcal{F}_t$ occurs with probability $1 - O(\log^{-1} n)$ by Lemma 12. If this is the case, the updated part of both configurations will have the same component structure; thus, $L_1(X_{t+1}) = L_1(Y_{t+1})$ and $d_{t+1} \leq d_t$. Hence, if $\mathcal{F}_t$ occurs for all $0 \leq t \leq T$, then X_T and Y_T fail to have the same component structure only if at least one of the initial components was never activated. For $T = O(\log n)$ this occurs with at most constant probability. Since $\mathcal{F}_t$ occurs for $T = O(\log n)$ consecutive steps with at least constant probability, then X_T and Y_T have the same component structure with probability $\Omega(1)$. ◀

In the sub-critical case we may assume also that $L_1(X_0)$ and $L_1(Y_0)$ are $O(\log n)$. Therefore, a simplified version of the same coupling works since Coupling II is not necessary.

▶ **Corollary 23.** *If $\lambda < \lambda_s$ and X_0, Y_0 are as in Lemma 21, then there exists a coupling of the CM steps such that X_T and Y_T have the same component structure with probability $\Omega(1)$, for some $T = O(\log n)$.*

3.5 Coupling to the same configuration

▶ **Lemma 24.** *Let X_0 and Y_0 be two random-cluster configurations with the same component structure. Then, there exists a coupling of the CM steps such that after $T = O(\log n)$ steps $X_T = Y_T$ w.h.p.*

Proof. Let B_t a bijection between the vertices of X_t and Y_t. We first describe how to construct B_0. Consider a maximal matching between the components of X_0 and Y_0 with the restriction that only components of equal size are matched to each other. Since the two configurations have the same component structure all components are matched. Using this matching, vertices between matched components are mapped arbitrarily to obtain B_0.

Vertices mapped to themselves we call "fixed". At time t, the component activation is coupled according to B_t. That is, if $B_t(u) = v$ for $u \in X_t$ and $v \in Y_t$, then the components containing u and v are simultaneously activated with probability $1/q$. B_{t+1} is adjusted such that if a vertex w becomes active in both configurations then $B_{t+1}(w) = w$; the rest of the activated vertices are mapped arbitrarily in B_{t+1} and the inactive vertices are mapped like in B_t. The percolation step at time t is then coupled using B_{t+1}. That is, the re-sampling of the active edge $(u, v) \in X_t$ is coupled with the re-sampling of the active edge $(B_{t+1}(u), B_{t+1}(v)) \in Y_t$.

This coupling ensures that the component structures of X_t and Y_t remain the same for all $t \geq 0$. Moreover, once a vertex is fixed it remains fixed forever. The probability that a vertex is fixed in one step is $1/q^2$. Therefore, after $O(\log n)$ steps the probability that a

vertex is not fixed is at most $1/n^2$. A union bound over all vertices implies that $X_T = Y_T$ w.h.p. after $T = O(\log n)$ steps. ◀

4 Mixing time lower bounds

In this section we prove the exponential lower bound on the mixing time of the CM dynamics for λ in the critical window (λ_s, λ_S), as stated in Theorem 2 in the introduction. We also prove a $\Omega(\log n)$ lower bound in the "fast mixing" regime, showing that our upper bounds in Section 3 are tight.

Recall from the introduction that when $q = 2$ and $\lambda < \lambda_s = \lambda_c$, the SW dynamics mixes in $\Theta(1)$ steps and thus the CM dynamics requires $\Theta(\log n)$ additional steps to mix. This is due to the fact that the CM dynamics may require as many steps to activate all the components from the initial configuration.

▶ **Theorem 25.** *For any $q > 1$, the mixing time of the CM dynamics is $\exp(\Omega(\sqrt{n}))$ for $\lambda \in (\lambda_s, \lambda_S)$, and $\Omega(\log n)$ for $\lambda \notin [\lambda_s, \lambda_S]$.*

Proof. Note that when $q \leq 2$ the interval (λ_s, λ_S) is empty and the exponential lower bound is vacuously true. It is natural to divide the proof into four cases: $\lambda < \lambda_s$, $\lambda \in (\lambda_s, \lambda_c)$, $\lambda \in [\lambda_c, \lambda_S)$ and $\lambda > \lambda_S$. First consider the case when $\lambda < \lambda_s$. Let X_0 be a configuration where all the components have size $\Theta(\log^2 n)$ and let $b = q/(q-1)$. The probability that a particular component is not activated in any of the first $T = \frac{1}{2}\log_b n$ steps is $(1-1/q)^T = n^{-1/2}$. Therefore, the probability that all initial components are activated in the first T steps is $(1-n^{-1/2})^K$ with $K = \Theta(n/\log^2 n)$. Thus, after T steps, $L_1(X_T) = \Theta(\log^2 n)$ w.h.p., and the result follows.

Consider now the case $q > 2$ and $\lambda_c \leq \lambda < \lambda_S = q$. Let S be the set of graphs G such that $L_1(G) = \Theta(\sqrt{n})$ and let $X_0 \in S$. Let $\mu := \mathrm{E}[A_0] = n/q$; then by Hoeffding's inequality $\Pr\left[|A_0 - \mu| > \varepsilon n\right] \leq 2\exp\left(-2\varepsilon^2\sqrt{n}\right)$. If $A_0 < \mu + \varepsilon n$, the percolation step is sub-critical for sufficiently small ε. Therefore, Lemma 9 implies that $\Pr[X_1 \notin S | X_0 \in S] \leq e^{-c\sqrt{n}}$ for some constant $c > 0$. Hence, $\Pr[X_1, ..., X_t \in S | X_0 \in S] \geq 1 - te^{-c\sqrt{n}} \geq 3/4$ for $t = \lfloor e^{c\sqrt{n}}/4 \rfloor$, and the result follows.

The intuition for the other two cases, which are more technically involved, comes directly from Figure 1. When $q > 2$ and $\lambda_s < \lambda < \lambda_c$, the function $f(\theta) = \theta - \phi(\theta)$ has two positive zeros θ^* and θ_r in $(\theta_{\min}, 1]$. Moreover, f is negative in the interval (θ^*, θ_r). Therefore, any configuration with a unique large component of size θn with $\theta \in (\theta^*, \theta_r)$ will drift towards a configuration with a bigger large component. However, a typical random-cluster configuration in this regime does not have a large component. This drift in the incorrect direction is sufficient to prove the exponential lower bound in this regime.

When $\lambda > \lambda_S$, we show that the derivative of f between its unique zero θ_r and 1 is bounded above by a constant. This implies that, starting from the complete graph, it takes at least $\Omega(\log n)$ steps for the size of the largest component to shrink to close to $\theta_r n$. The reader is referred to the full version [2] for the proofs of the last two cases. ◀

5 Local dynamics

In this section we sketch the proof of Theorem 3 from the introduction; the full proof is included in the full version [2]. Consider an arbitrary finite graph $H = (V, E)$ and let $\Omega_{\mathrm{E}} = \{(V, A) : A \subseteq E\}$ be the set of random-cluster configurations on H. Let P be the transition matrix of a finite, ergodic and reversible Markov chain over Ω_{E} with stationary

distribution $\mu = \mu_{p,q}$ and eigenvalues $1 = \lambda_1 \geq \lambda_2 \geq ... \geq \lambda_n$. The *spectral gap* of P is defined by $\lambda(P) := 1 - \lambda^*$, where $\lambda^* = \max\{|\lambda_2|, |\lambda_n|\}$. Let $\mu_{\min} = \min_{x \in \Omega_{\mathrm{E}}} \mu(x)$; the following bounds on the mixing time are standard (see, e.g., [17]):

$$\lambda^{-1}(P) - 1 \leq \tau_{\mathrm{mix}}(P) \leq \log\left(2e\mu_{\min}^{-1}\right) \lambda^{-1}(P). \tag{11}$$

For $r \in \mathbb{N}$, let $\Omega_{\mathrm{V}} = \{0, 1..., r-1\}^V$ be the set of "r-labelings" of V, and let $\Omega_{\mathrm{J}} = \Omega_{\mathrm{V}} \times \Omega_{\mathrm{E}}$. Assume P can be decomposed as a product of stochastic matrices of the form $P = M(\prod_{i=1}^m T_i)M^*$, where:

(i) M is a $|\Omega_{\mathrm{E}}| \times |\Omega_{\mathrm{J}}|$ matrix indexed by the elements of Ω_{E}, Ω_{J} such that $M(A, (\sigma, B)) \neq 0$ only if $A = B$.

(ii) Each T_i is a $|\Omega_{\mathrm{J}}| \times |\Omega_{\mathrm{J}}|$ matrix indexed by the elements of Ω_{J}, reversible w.r.t. the distribution $\nu = \mu M$ and such that $T_i((\sigma, A), (\tau, B)) \neq 0$ only if $\sigma = \tau$.

(iii) M^* is a $|\Omega_{\mathrm{J}}| \times |\Omega_{\mathrm{E}}|$ matrix indexed by the elements of Ω_{J}, Ω_{E} such that $M^*((\sigma, A), B) = \mathbb{1}(A = B)$.

In words, M assigns a (random) r-labeling to the vertices of H; $(\prod_{i=1}^m T_i)$ performs a sequence of m operations T_i, each of which updates some edges of H; and M^* drops the labels from the vertices.

Consider now the matrix $P_{\mathrm{L}} = M(\frac{1}{m}\sum_{i=1}^m T_i)M^*$. It is straightforward to verify that P_{L} is also reversible w.r.t. μ. The following theorem, which generalizes a recent result of Ullrich [21, 22], relates the spectral gaps of P and P_{L} up to a factor of $O(m \log m)$.

▶ **Theorem 26.** *If M, M^* and T_i are stochastic matrices satisfying (i)–(iii) above, and the T_i's are idempotent commuting matrices, then $\lambda(P_{\mathrm{L}}) \leq \lambda(P) \leq 8m \log m \cdot \lambda(P_{\mathrm{L}})$.*

We pause to note that this fact has a very attractive intuitive basis: P_{L} performs a single update T_i chosen u.a.r., while P performs all m updates T_i, so by coupon collecting one might expect that $O(m \log m)$ P_{L} steps should suffice to simulate a single P step. However, the proof has to take account of the fact that the T_i updates are interleaved with the vertex re-labeling operations M and M^* in P_{L}. The proofs in [21] and [22] are specific to the case where P corresponds to the SW dynamics. Our contribution is the realization that these proofs still go through (without essential modification) under the more general assumptions of Theorem 26, as well as the framework described above that provides a systematic way of deriving P_{L} from P.

The key observation in the proof of Theorem 3 is that we can express P_{CM}, the transition matrix of the CM dynamics, as a product of stochastic matrices as above: specifically, $P_{\mathrm{CM}} = L(\prod_{e \in E} T_e)L^*$ where L is the matrix that assigns a random active-inactive labeling to a random-cluster configuration, T_e samples e with probability p provided both its endpoints are active, and L^* drops the active-inactive labeling from a joint configuration.

Now consider the Markov chain given by the matrix $P_{\mathrm{SU}} = L(\frac{1}{|E|}\sum_{e \in E} T_e)L^*$, which we call the *Single Update (SU) dynamics* (P_{SU} plays the role of the matrix P_{L} above.) The matrices L, L^* and the T_e's satisfy the assumptions in Theorem 26, so we have $\lambda(P_{\mathrm{SU}}) \leq \lambda(P_{\mathrm{CM}}) \leq 8|E| \log |E| \cdot \lambda(P_{\mathrm{SU}})$.

The SU dynamics is very closely related to the heat-bath dynamics defined in the introduction; in fact, their spectral gaps differ by a constant. Hence, Theorem 3 now follows from (11) and Theorems 1 and 2 since in the mean-field case $\log(\mu_{\min}^{-1}) = \tilde{O}(n^2)$.

References

1 V. Beffara and H. Duminil-Copin. The self-dual point of the two-dimensional random-cluster model is critical for $q \geq 1$. *Probability Theory and Related Fields*, 153:511–542, 2012.

2 A. Blanca and A. Sinclair. Dynamics for the mean-field random-cluster model, 2014. Preprint. arXiv:1412.6180v1 [math.PR].

3 B. Bollobás, G. Grimmett, and S. Janson. The random-cluster model on the complete graph. *Probability Theory and Related Fields*, 104(3):283–317, 1996.

4 C. Borgs, J. Chayes, and P. Tetali. Swendsen-Wang algorithm at the Potts transition point. *Probability Theory and Related Fields*, 152:509–557, 2012.

5 C. Borgs, A. Frieze, J.H. Kim, P. Tetali, E. Vigoda, and V. Vu. Torpid mixing of some Monte Carlo Markov chain algorithms in statistical physics. *Proceedings of the 40th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 218–229, 1999.

6 L. Chayes and J. Machta. Graphical representations and cluster algorithms II. *Physica A*, 254:477–516, 1998.

7 C. Cooper, M.E. Dyer, A.M. Frieze, and R. Rue. Mixing properties of the Swendsen-Wang process on the complete graph and narrow grids. *Journal of Mathematical Physics*, 41:1499–1527, 2000.

8 P. Cuff, J. Ding, O. Louidor, E. Lubetzky, Y. Peres, and A. Sly. Glauber dynamics for the mean-field Potts model. *Journal of Statistical Physics*, 149(3):432–477, 2012.

9 R.G. Edwards and A.D. Sokal. Generalization of the Fortuin-Kasteleyn-Swendsen-Wang representation and Monte Carlo algorithm. *Physical Review D*, 38(6):2009–2012, 1988.

10 C.M. Fortuin and P.W. Kasteleyn. On the random-cluster model I. Introduction and relation to other models. *Physica*, 57(4):536–564, 1972.

11 A. Galanis, D. Štefankovič, and E. Vigoda. Swendsen-Wang Algorithm on the Mean-Field Potts Model, 2015. Preprint. arXiv:1502.06593v1 [cs.DM].

12 Q. Ge and D. Štefankovič. A graph polynomial for independent sets of bipartite graphs. *Combinatorics, Probability and Computing*, 21(5):695–714, 2012.

13 V.K. Gore and M.R. Jerrum. The Swendsen-Wang process does not always mix rapidly. *Journal of Statistical Physics*, 97(1-2):67–86, 1999.

14 G.R. Grimmett. *The Random-Cluster Model*, volume 333 of *Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 2006.

15 M. Huber. A bounding chain for Swendsen-Wang. *Random Structures & Algorithms*, 22(1):43–59, 2003.

16 G. Kirchhoff. Über die Auflösung der Gleichungen, auf welche man bei der Untersuchung der linearen Verteilung galvanischer Strome gefuhrt wird. *Annalen der Physik und Chemie*, pages 497–508, 1847.

17 D.A. Levin, Y. Peres, and E.L. Wilmer. *Markov Chains and Mixing Times*. American Mathematical Society, 2008.

18 Y. Long, A. Nachmias, W. Ning, and Y. Peres. A power law of order 1/4 for critical mean-field Swendsen-Wang dynamics. *Memoirs of the American Mathematical Society*, 232(1092), 2011.

19 M.J. Luczak and T. Łuczak. The phase transition in the cluster-scaled model of a random graph. *Random Structures & Algorithms*, 28(2):215–246, 2006.

20 R.H. Swendsen and J.S. Wang. Nonuniversal critical dynamics in Monte Carlo simulations. *Physical Review Letters*, 58:86–88, 1987.

21 M. Ullrich. Rapid mixing of Swendsen-Wang dynamics in two dimensions. *Dissertationes Mathematicae*, 502, 2014.

22 M. Ullrich. Swendsen-Wang is faster than single-bond dynamics. *SIAM Journal on Discrete Mathematics*, 28(1):37–48, 2014.