Finitary codings for the random-cluster model and other infinite-range monotone models[*]

Matan Harel[†]   Yinon Spinka[‡]

Abstract

A random field $X = (X_v)_{v \in G}$ on a quasi-transitive graph G is a factor of an i.i.d. process if it can be written as $X = \varphi(Y)$ for some i.i.d. process $Y = (Y_v)_{v \in G}$ and equivariant map φ. Such a map, also called a coding, is finitary if, for every vertex $v \in G$, there exists a finite (but random) set $U \subset G$ such that X_v is determined by $\{Y_u\}_{u \in U}$. We construct a coding for the random-cluster model on G, and show that the coding is finitary whenever the free and wired measures coincide. This strengthens a result of Häggström–Jonasson–Lyons [18]. We also prove that the coding radius has exponential tails in the subcritical regime. As a corollary, we obtain a similar coding for the subcritical Potts model.

Our methods are probabilistic in nature, and at their heart lies the use of coupling-from-the-past for the Glauber dynamics. These methods apply to any monotone model satisfying mild technical (but natural) requirements. Beyond the random-cluster and Potts models, we describe two further applications – the loop $O(n)$ model and long-range Ising models. In the case of $G = \mathbb{Z}^d$, we also construct finitary, translation-equivariant codings using a finite-valued i.i.d. process Y. To do this, we extend a mixing-time result of Martinelli–Olivieri [22] to infinite-range monotone models on quasi-transitive graphs of sub-exponential growth.

[†]Northeastern University, Boston, MA, United States of America. E-mail: m.harel@northeastern.edu
[‡]University of British Columbia, Vancouver, BC, Canada. E-mail: yinon@math.ubc.ca

Finitary codings for the random-cluster and other infinite-range monotone models

1 Introduction and main results

Consider an infinite graph $G = (\mathbb{V}, E)$ and a random field $X = (X_v)_{v \in \mathbb{V}}$ whose distribution is invariant under all automorphisms of G. This paper is concerned with the question of existence of codings (factor maps): is it possible to express X as an automorphism-equivariant function (which we call a coding) an i.i.d. process — i.e. a random field $Y = (Y_v)_{v \in \mathbb{V}}$ where the Y_v's are independent and identically distributed? The answer to this question depends on the graph G and the random field X. The theory of such codings traces back to the seminal work of Ornstein [24] and later Keane and Smorodisnky [20], who studied the case in which $G = \mathbb{Z}$ and X itself is an i.i.d. process. In this case, Keane and Smorodisnky showed that X and Y are finitarily isomorphic – a stronger condition than the one we require. The study of the one-dimensional problem when X is a more general process remains an active research topic.

In the setting of the d-dimensional lattice $\mathbb{Z}^d$, it is very natural to ask whether the Ising model is a factor of an i.i.d. process. This model, perhaps the most well-known of the statistical physics models, is infamously trivial on $\mathbb{Z}$, but exhibits a phase transition on $\mathbb{Z}^d$ when $d \geq 2$ – and hence, it is appropriate to study it on non one-dimensional lattices. In an unpublished work, Ornstein and Weiss [25] (see also [1] for a published version) showed that the (infinite-volume) plus state of the Ising model at any positive temperature is a factor of an i.i.d. process. Steif [31] showed a similar but stronger result for monotone spin systems. In both of these cases, the factor maps may be infinitely dependent, in the sense that determining the value of X at the origin may require knowing the value of infinitely many elements of Y. Van den Berg and Steif [2] showed that the subcritical Ising model has a *finitary* coding. Explicitly, they construct a factor map φ from an i.i.d. process Y to any subcritical Ising model such that $\varphi(Y)$ at the origin depends on a finite (but random) number of Y_v's. In fact, their work quantifies the 'amount of information' required to determine $\varphi(Y)$ at the origin in two ways. On the one hand, they show that there exists a coding whose coding radius, which controls how far one must look in the Y process, has exponential tails. On the other hand, they show that there exists a (different) coding which only requires a finite-valued input Y_v at every vertex. The same work shows that no finitary coding can exist for the supercritical Ising model. Recent works constructed finitary codings for Markov random fields with spatial mixing properties [30], or long-range interacting particle systems that satisfy a 'high noise' condition [12].

The initial goal of this project was to show that the random-cluster model on $\mathbb{Z}^d$ is a finitary factor of an i.i.d. process. Unlike the Ising model, the random-cluster model has infinite-range interactions – i.e. the state of an edge in the random-cluster model may have a nonvanishing effect on the state of an edge that is arbitrarily far away from it. Although the methods we use yield more general results, the main result of this paper is the construction of a finitary coding for the random-cluster model on an arbitrary quasi-transitive graph when the free and wired measures coincide. In the case of the random-cluster model, the factor constructed in this paper is very similar to the one discussed by Häggström–Jonasson–Lyons [18]. That paper constructs a factor map for the random-cluster model on a general quasi-transitive graph, but does not study whether it is finitary. The analysis presented herein also provides quantitative control of the coding radius of the factor. In particular, for the subcritical random-cluster model, the coding radius will have exponential tails. We further obtain results for the Potts model on such graphs using the Edwards–Sokal coupling. In the case of the subcritical random-cluster and Potts models on $\mathbb{Z}^d$, we also prove the existence of a finitary coding from a finite-valued i.i.d. process.

The general framework discussed in this paper is that of monotone specifications.

Specifications, a formalization of concepts first introduced in the work of Dobrushin [7] and Landford–Ruelle [21], are families of finite-volume measures, indexed by finite sub-sets and arbitrary configurations (boundary conditions), that satisfy certain consistency relations. They are called monotonic (or attractive) if the measures respect a partial ordering on configurations, in the sense of stochastic domination; this property is a generalization of the FKG property of the random-cluster model, or Griffiths' inequalities for the Ising model. We emphasize that we will not demand that specifications are quasi-local (which is a well-known continuity property), which is frequently assumed elsewhere and violated by the random-cluster and other infinite-range models. The generality of the framework has many possible applications. We discuss two additional applications: to the critical loop $O(n)$ model on the hexagonal lattice, and to subcritical long-range Ising models. As far as we are aware, this is the first finitary coding result for the loop $O(n)$ model and the only non-perturbative finitary coding result for long-range Ising models (see [11] for a result at sufficiently high temperatures).

We end by briefly discussing the algorithmic aspects of our results. There is an extensive literature focused on perfect simulations of infinite-range models [12, 13, 5]; for an example involving the 'high noise' regimes of the random-cluster model on $\mathbb{Z}^d$, see [6]. The proofs given in this paper rely on the method of coupling-from-the-past of Propp and Wilson [28]. This technique uses dynamics in order to get a perfect simulation of the stationary distribution of a finite-state Markov chain. In our setting, we apply coupling-from-the-past to the single-site Glauber dynamics (in the same spirit as previous works, e.g., [2, 19]). As such, there is an interest in controlling not only the spatial dependence of the factor map (i.e. the coding radius), but also the mixing time, which measures the number of steps of the dynamics required to perfectly sample X at the origin. The celebrated work of Martinelli and Olivieri [22] relates spatial and temporal mixing in the context of the finite-range, finite-energy, monotone models on $\mathbb{Z}^d$; as part of this work, we prove a generalization of this result to infinite-range, monotone models on quasi-transitive graphs of sub-exponential growth. With this perspective, it is clear that the existence of space-time finitary factors has algorithmic implications: one may create a perfect sample of X on a finite subset of $\mathbb{V}$ by applying the space-time finitary factor map to the i.i.d. process Y on some random (and possibly much larger) finite subset of $\mathbb{V}$. Controlling the size of the latter set is tantamount to quantitative control on the coding radius and mixing time.

1.1 Definitions

Let $\mathbb{V}$ be countably infinite and let Γ be a group acting on $\mathbb{V}$. The action is *quasi-transitive* if it partitions $\mathbb{V}$ into finitely many orbits. Let $(S, \mathcal{S})$ and $(T, \mathcal{T})$ be two measurable spaces, and let $X = (X_v)_{v \in \mathbb{V}}$ and $Y = (Y_v)_{v \in \mathbb{V}}$ be $(S, \mathcal{S})$-valued and $(T, \mathcal{T})$-valued Γ-invariant random fields. For the rest of the paper, we will assume that all probability spaces are standard.

A *coding* from Y to X is a measurable function $\varphi : T^{\mathbb{V}} \to S^{\mathbb{V}}$, which satisfies $\varphi(Y) \overset{d}{=} X$ and is Γ-*equivariant*, i.e., commutes with the action of every element in Γ on a Γ-invariant subset of $T^{\mathbb{V}}$ of full measure. Such a coding is also called a *factor map* or *homomorphism* from Y to X; when such a coding exists, we say that X is a Γ-*factor* of Y.

Suppose that G is a locally finite graph on vertex set $\mathbb{V}$ and that Γ acts quasi-transitively on $\mathbb{V}$ by automorphisms of G. Thus, G is a quasi-transitive graph; heuristically, such a graph has finitely many 'different types' of vertices (whereas a transitive graph has exactly one). We say that a pair of configurations $y, y' \in T^{\mathbb{V}}$ agree up to distance r around a vertex v if $y_w = y'_w$ for all v with $\mathrm{dist}(v, w) \leq r$, where $\mathrm{dist}(\cdot, \cdot)$ denotes the graph distance. We say that φ is determined at distance r around v at a configuration y if $\varphi(y)_v = \varphi(y')_v$ for any y' which agrees with y up to distance r around v. The *coding*

radius of φ at the vertex v and the configuration y, which we denote by $R_v(y)$, is the minimal distance that determines φ at v and y. It may happen that no such r exists, in which case, $R_v(y) = \infty$. Thus, associated to a coding is a random variable $R_v = R_v(Y)$ which describes the coding radius at v; a coding is called *finitary* if $R_v(Y)$ is almost surely finite for every $v \in \mathbb{V}$.[1]

We categorize factor maps as follows: When X is a Γ-factor of an i.i.d. (independent and identically distributed) process, we say it is Γ-*fiid*, and when it is a finitary Γ-factor of an i.i.d. process, we say it is Γ-*ffiid*. A still stronger notion is Γ-fv-ffiid which requires X to be a finitary Γ-factor of a *finite-valued* i.i.d. process (i.e., a finite set T). In addition, we can add a quantitative element which indicates how far a coding must look to determine the output at the origin. Explicitly, we say that a coding has *exponential tails* if $\mathbb{P}(R_v \geq r) \leq Ce^{-cr}$ for some $C, c > 0$ and all $r \geq 0$ and $v \in \mathbb{V}$, and that it has *stretched-exponential tails* if $\mathbb{P}(R_v \geq r) \leq Ce^{-r^\nu}$ for some $C > 0$, $0 < \nu < 1$, and all $r \geq 0$ and $v \in \mathbb{V}$. For the remains of the paper, when we use the notion fiid (or any variant thereof) without an explicit mention of Γ, the group may be taken to be any group acting quasi-transitively on $\mathbb{V}$ by automorphisms of the graph.

1.2 The random-cluster model

We begin with a definition of the random-cluster model; for background on the model and its fundamental properties mentioned below, we direct the reader to the monographs [16, 8].

Let $G' = (V', E')$ be a finite subgraph of G, and let $\partial V'$ denote the set of vertices in V' that have a neighbor in $\mathbb{V} \setminus V'$. The random-cluster measure in G' with parameters $p \in [0, 1]$ and $q > 0$ and boundary conditions $i \in \{0, 1\}$ is given by

$$\phi_{G',p,q}^i(\omega) = \frac{p^{o(\omega)}(1-p)^{c(\omega)}q^{k^i(\omega)}}{Z^i(G',p,q)}, \qquad \omega \in \{0,1\}^{E'},$$

where $o(\omega)$ and $c(\omega)$ are the numbers of open and closed edges, i.e., edges e such that $\omega_e = 1$ and $\omega_e = 0$, respectively, $k^0(\omega)$ is the number of open clusters in ω, $k^1(\omega)$ is the number of open clusters that do not intersect the boundary $\partial V'$, and $Z^i(G',p,q)$ is a normalizing constant, called the partition function, which makes $\phi_{G',p,q}^i$ a probability measure. We call the measures *free* and *wired* when $i = 0$ and $i = 1$, respectively. It is well-known that, when $q \geq 1$, the random-cluster model has the FKG property (a monotonicity property), which implies that $\phi_{G',p,q}^i$ converges weakly to a limiting measure $\phi_{p,q}^i$ as G' increases to G. The two limiting measures are probability measures which are supported on $\{0,1\}^E$ and are invariant under all automorphisms of G. We call $\phi_{p,q}^0$ and $\phi_{p,q}^1$ the free and wired infinite-volume random-cluster measures. Our results concern the coding properties of these two measures.

When $\phi_{p,q}^0 = \phi_{p,q}^1$, we may omit the superscript for notational clarity and write $\phi_{p,q}$ for the common measure. A standard coupling argument shows that, for each $i \in \{0, 1\}$, there exists a critical parameter $p_c^i(q) \in [0, 1]$ such that

$$\phi_{p,q}^i[\exists \text{ an infinite cluster}] = \begin{cases} 0 & p < p_c^i(q) \\ 1 & p > p_c^i(q) \end{cases}.$$

It is also a straightforward consequence of the FKG property that $p_c^0(q) \geq p_c^1(q)$, and that $\phi_{p,q}^0 = \phi_{p,q}^1$ whenever $p < p_c^1(q)$. Furthermore, the well-known Burton–Keane argument [4] implies that $p_c^0(q) = p_c^1(q)$ on any amenable graph.

---

[1]Technically, the map $R_v(y)$ may not be measurable as defined. One may deal with this by either modifying φ on a null set (with respect to Y) or by demanding instead that, for almost every y, $\varphi(y)_v = \varphi(y')_v$ for almost every y' which agrees with y up to distance r around v.

We now state our first result about codings for the random-cluster model. Recall that our definition of coding is stated for a random field on the vertex set of a graph. Thus, strictly speaking, when considering codings for the random-cluster model on G, one should think of the model as being defined on the vertices of the line graph of G. In this case, following our convention, the implicit Γ may be taken to be any group acting quasi-transitively on $\mathbb{V}$ by automorphisms of the line graph.

Theorem 1.1. *Let G be an infinite quasi-transitive graph and let $p \in [0,1]$ and $q \geq 1$.*

- *Both $\phi_{p,q}^0$ and $\phi_{p,q}^1$ are fiid.*
- *If $\phi_{p,q}^0 = \phi_{p,q}^1$, then $\phi_{p,q}$ is ffiid.*
- *If $p < p_c^1(q)$, then $\phi_{p,q}$ is ffiid with exponential tails.*

The second and third items are new, while the first item was already known; see the discussion in Section 1.6 for further details.

In the amenable case, we can also prove a partial converse for the second item in Theorem 1.1. This converse works on any amenable graph for the wired measure; we require a slightly stronger version of amenability for the free measure. We say that G is *amenable* if there exists a sequence $(F_n)_n$ of non-empty finite subsets of $\mathbb{V}$ such that $|\partial F_n|/|F_n| \to 0$ as $n \to \infty$. We say that G is *c-amenable* if there exists a sequence $(F_n, H_n)_n$ of non-empty finite subsets of $\mathbb{V}$ such that H_n is connected, $\partial F_n \subset H_n$ and $|H_n|/|F_n| \to 0$ as $n \to \infty$. The 'c' in c-amenable stands for connected. It is clear that c-amenability implies amenability.

Theorem 1.2. *Let G be an infinite quasi-transitive graph and let $p \in [0,1]$ and $q \geq 1$.*

- *If G is amenable and $\phi_{p,q}^0 \neq \phi_{p,q}^1$, then $\phi_{p,q}^1$ is not ffiid.*
- *If G is c-amenable and $\phi_{p,q}^0 \neq \phi_{p,q}^1$, then $\phi_{p,q}^0$ is not ffiid.*

Thus, for c-amenable graphs, Theorem 1.1 and Theorem 1.2 imply that the wired random-cluster measure is ffiid if and only if the free random-cluster measure is ffiid, both of which hold if and only if the two measures coincide. In particular, this is the case for $\mathbb{Z}^d$, as it is c-amenable for any $d \geq 2$, and the free and wired are always equal on $\mathbb{Z}$. We do not know whether the first item of the theorem holds for non-amenable graphs, nor do we know whether the second item holds under the standard amenability assumption. However, the second item does not necessary hold in the non-amenable case. Indeed, on the d-regular tree, Häggström [17] shows that $p_c^1(q) < p_c^0(q)$ for any $q > 2$ and $d \geq 3$, and thus, $\phi_{p,q}^0 \neq \phi_{p,q}^1$ for any $p \in (p_c^1(q), p_c^0(q))$; meanwhile, on any tree and for any p and q, $\phi_{p,q}^0$ is exactly Bernoulli percolation of parameter $p/[p + q(1-p)]$, and thus is trivially ffiid.

The next result is concerned with the existence of codings from a finite-valued i.i.d. process for the subcritical random-cluster model on $\mathbb{Z}^d$. In this case, $p_c^0(q) = p_c^1(q)$, so that we may drop the superscript.

Theorem 1.3. *Let $d \geq 2$, $G = \mathbb{Z}^d$ and Γ be the translation group of $\mathbb{Z}^d$. Let $q \geq 1$ and $p < p_c(q)$. Then $\phi_{p,q}$ is Γ-fv-ffiid with stretched-exponential tails.*

We note that the coding we produce above is *translation*-equivariant, not *automorphism*-equivariant. The construction we use does not produce a reflection/rotation-equivariant coding, though we believe that such a construction should be possible, and that a similar statement should hold for more general quasi-transitive graphs of sub-exponential growth and their full automorphism group. The proof of Theorem 1.3 relies on a result from [29] and on a new mixing-time result for a natural single-site dynamics of the subcritical random-cluster model on an arbitrary infinite quasi-transitive graph of sub-exponential growth (see Section 5).

1.3 The Potts model

The random-cluster model is closely related to the Potts model. This model, introduced by Potts [27] following a suggestion of his adviser Domb, has been the subject of intensive study by mathematicians and physicists over the last three decades. For a review of the physics results, see [32]; for proofs of the classical rigorous results quoted below, see [14].

Set $q \geq 2$ to be an integer. The Potts measure on a finite subgraph $G' = (V', E')$ of G, at inverse temperature β and boundary conditions $i \in \{0, 1, \ldots, q\}$, is defined by

$$\mu^i_{G',\beta,q}[\sigma] := \frac{e^{\beta \mathbf{H}^i_{G'}(\sigma)}}{Z^i_{G',\beta,q}}, \qquad \sigma \in \{1, \ldots, q\}^{V'}, \tag{1.1}$$

where

$$\mathbf{H}^i_{G'}(\sigma) := \sum_{\{x,y\} \in E'} \mathbf{1}[\sigma_x = \sigma_y] + \sum_{\{x,y\} \in \partial E'} \mathbf{1}[\sigma_x = i],$$

and $Z^i_{G',\beta,q}$ is a normalizing constant which makes $\mu^i_{G',\beta,q}$ a probability measure. Above, $\mathbf{1}[\cdot]$ denotes the indicator function. Note that when $i = 0$, the second sum is zero for all σ. One obtains infinite-volume measures $\mu^i_{\beta,q}$ via weak limits, which are known to exist. The case $q = 2$ is known as the Ising model.

Like the random-cluster model, the different infinite-volume Potts measures may be highly affected by their boundary conditions. However, if $\beta < \beta^{\mathrm{w}}_c(q)$, where $\beta^{\mathrm{w}}_c(q) := -\log[1 - p^1_c(q)]$, it is well-known that $\mu^0_{\beta,q}, \mu^1_{\beta,q}, \ldots, \mu^q_{\beta,q}$ coincide. In this case, we denote the common measure by $\mu_{\beta,q}$. Using the relation with the random-cluster, we obtain the following results about the coding properties of the Potts model:

Theorem 1.4. *Let G be an infinite quasi-transitive graph, $q \geq 2$ be an integer and $\beta \geq 0$.*

- *If $\mu^0_{\beta,q}, \mu^1_{\beta,q}, \ldots, \mu^q_{\beta,q}$ coincide, then $\mu_{\beta,q}$ is ffiid.*
- *If $\beta < \beta^{\mathrm{w}}_c(q)$, then $\mu_{\beta,q}$ is ffiid with exponential tails.*
- *If $\beta < \beta^{\mathrm{w}}_c(q)$, $G = \mathbb{Z}^d$ and Γ is the translation group, then $\mu_{\beta,q}$ is Γ-fv-ffiid with stretched-exponential tails.*

In the case of the Ising model ($q = 2$) on $\mathbb{Z}^d$, the results of Theorem 1.4 were shown in [2] (items one and two) and [29] (item three). For $q \neq 2$ on $\mathbb{Z}^d$, partial results were established in [19] and [30]. To the best of our knowledge, the above result is novel in all other cases.

1.4 The loop $O(n)$ model

The loop $O(n)$ model is a model for a random collection of non-intersecting loops on the hexagonal lattice, which is believed to be in the same universality class as the spin $O(n)$ model (see [26] for background on these models). Let Ω be a connected finite subset of the hexagonal lattice $\mathbb{H}$ whose complement is connected, and let $\mathrm{LoopConf}(\Omega)$ be the set of subgraphs where every vertex is of degree 0 or 2, so that the non-trivial connected components form loops. The loop $O(n)$ measure with edge-weight $x > 0$ and loop-weight $n > 0$ is the probability measure $\nu_{\Omega,n,x}$ given by

$$\nu_{\Omega,n,x}(\omega) := \frac{x^{o(\omega)} n^{\ell(\omega)}}{Z_{\Omega,n,x}} \cdot \mathbf{1}_{\mathrm{LoopConf}(\Omega)}(\omega), \qquad \omega \in \{0, 1\}^{E(\Omega)},$$

where $o(\omega)$ is the number of edges in ω, $\ell(\omega)$ is the number of loops in ω, and $Z_{\Omega,n,x}$ is a normalizing constant. The loop $O(n)$ model is conjectured to undergo a phase transition for any $0 \leq n \leq 2$ when the value of x equals

$$x_c(n) := \frac{1}{\sqrt{2 + \sqrt{2 - n}}}.$$

It is shown in [9, Theorems 1 and 2] that the loop $O(n)$ model has a unique periodic (i.e. invariant under a finite-index subgroup of the automorphism group of $\mathbb{H}$) Gibbs measure $\nu_{n,x}$ whenever $n \geq 1$ and $nx^2 \leq 1$, and moreover, that this is the unique Gibbs measure whenever $n \in [1,2]$ and $x = x_c(n)$ (we note that a unique periodic measure must be invariant under all automorphisms of $\mathbb{H}$). In the latter case, we show that this measure has a finitary coding.

Theorem 1.5. *Let $n \geq 1$ and $nx^2 \leq 1$. Then $\nu_{n,x}$ is fiid, and it is ffiid when $n \in [1,2]$ and $x = x_c(n)$.*

It would be interesting to determine whether $\nu_{n,x}$ is always ffiid in the regime $n \geq 1$ and $nx^2 \leq 1$.

1.5 Long-range (ferromagnetic) Ising models

Let $J = (J_A)_{A \subset \mathbb{V}, |A| \leq 2}$ be a collection of non-negative numbers called the *coupling constants* satisfying that

$$\sum_{\substack{A \subset \mathbb{V}, |A| \leq 2 \\ v \in A}} J_A < \infty \qquad \text{for all } v \in \mathbb{V}.$$

The Ising measure with coupling constants J in a finite volume $V \subset \mathbb{V}$ with boundary conditions τ is given by

$$\mu_V^\tau(\sigma) := \frac{\mathbf{1}_{\{\sigma_{\mathbb{V} \setminus V} = \tau_{\mathbb{V} \setminus V}\}}}{Z_V^\tau} \cdot \exp\left[\sum_{\substack{A \subset \mathbb{V}, |A| \leq 2 \\ A \cap V \neq \emptyset}} J_A \sigma_A \right], \qquad \sigma \in \{-1, +1\}^{\mathbb{V}},$$

where $\sigma_A := \prod_{v \in A} \sigma_v$. We assume that the coupling constants are automorphism-invariant in the sense that $J_{\gamma A} = J_A$ for all A and γ in the automorphism group of G.

The long-range Ising model has an important monotonicity property, known as Griffiths's inequality [15], which states that $\mathbb{E}[\sigma_A \sigma_B] \geq \mathbb{E}[\sigma_A]\mathbb{E}[\sigma_B]$ for any finite subsets $A, B \subset \mathbb{V}$. It is classical that this property implies monotonicity of the associated specifications, allowing us to define the largest and smallest Gibbs measures μ^+ and μ^-, respectively.

Theorem 1.6. *Let G be an infinite quasi-transitive graph and let J be non-negative coupling constants as above. Then μ^+ and μ^- are fiid. In addition, if $\mu^+ = \mu^-$ then μ^+ is ffiid.*

1.6 Discussion

In this section, we discuss the previously known results and their relation to ours.

We begin with the random-cluster and Potts models, which are our primary interest here. Let us start by explaining the relation between these two models. When $q \geq 2$ is an integer, the two models are closely related via the Edwards–Sokal coupling, which allows to obtain samples of one from the other via a simple procedure (which also introduces additional randomness). In this coupling, to obtain a sample from the Potts measure $\mu_{\beta,q}^0$, one first samples a configuration ω from the free random-cluster measure $\phi_{p,q}^0$ with $p = 1 - e^{-\beta}$, and then assigns a single color in $\{1, \ldots, q\}$ to all the vertices in each cluster of ω, with the colors of different clusters chosen uniformly and independently. To obtain a sample from $\mu_{\beta,q}^i$ with $i \in \{1, \ldots, q\}$, one follows the same procedure with ω sampled from the wired random-cluster measure $\phi_{p,q}^1$, except that any infinite cluster of ω (if any such clusters exist) is assigned color i. In the other direction, to obtain a sample from

$\phi_{p,q}^i$ with $i \in \{0,1\}$, one first samples σ from $\mu_{\beta,q}^i$, and then independently opens each edge with probability p if its endpoints have the same color in σ, and closes it otherwise. This relationship allows to transfer many coding properties from one model to the other. We refer the reader to [16, Theorem 4.91] for details.

Häggström–Jonasson–Lyons [18] studied (non-finitary) coding properties of the random-cluster and Potts models on general graphs. Lemma 4.5 of that paper shows that $\mu_{\beta,q}^i$ with $i \in \{1, \ldots, q\}$ is fiid for any integer $q \geq 2$ and any $\beta \geq 0$. That paper was particularly interested in the closely related notion of Bernoullicity (see Theorem 4.1 there), a classical mixing property from ergodic theory. In fact, on certain amenable graphs with mild geometric conditions – namely quasi-transitive amenable graphs that satisfy

$$\mathbb{V}_{v,r} \setminus \mathbb{V}_{v,r-1} \not\subset \mathbb{V}_{u,r} \qquad \text{for any distinct } u, v \in \mathbb{V} \text{ and infinitely many } r \in \mathbb{N},$$

the notions of Bernoullicity and fiid are equivalent ($\mathbb{V}_{v,r}$ is the graph-ball of radius r around v). Although it is not explicitly stated there, the first item of Theorem 1.1, which states that the free and wired random-cluster measures are fiid, is essentially contained in [18]. We remark that a slightly stronger version of the above condition appears in Theorem 2.2.

The second and third items of Theorem 1.1 are both new. As far as we know, previous results have been restricted to integer q, where they were deduced from results on the Potts model (see [29, Remark 5] and [30, Remark 6]). In particular, we are unaware of any finitary coding results which work directly on the random-cluster model. Let us also mention that, in the second item of the theorem, the coding radius is controlled by the rate of convergence of the finite-volume free and wired measures (see Theorem 2.1). Thus, the third item of the theorem will follow with the additional knowledge that the phase transition of the random-cluster model is sharp [10].

Theorem 1.3 is concerned with finitary codings in which the i.i.d. process is finite-valued. This is a rather natural restriction, as it makes the coding somewhat more useful for simulations. For the Ising model on $\mathbb{Z}^d$, van den Berg and Steif [2] constructed such codings, and the second author [29] showed that such codings exist with stretched-exponential tails for their coding radius. Pushing the methods used to prove Theorem 1.1, we obtain so-called space-time finitary codings, which, by combining with a general tool from [29], allow us to deduce that such codings also exist for the random-cluster model.

Theorem 1.2 is concerned with showing that the random-cluster measures are not ffiid in certain situations. While this is not the main focus of this paper, the result was given in order to provide a more complete picture. The theorem is an extension of similar results for other models (see [2, Theorem 2.1] and [30, Theorem 1.3]), which are based on exponential-rate estimates which appear in [23, 3]. Though previous results were for nearest-neighbor models on $\mathbb{Z}^d$, the key ideas apply in greater generality. We do point out however that certain technicalities may arise on other graphs, as is reflected in our need for the c-amenability property in the second item of Theorem 1.2. While we were able to verify that certain amenable graphs have this stronger property (e.g., the lamplighter group over $\mathbb{Z}^d$ for any $d \geq 1$), we do not know whether it holds for all infinite, one-ended, amenable, edge-transitive graphs; we note that two-ended graphs such as $\mathbb{Z}$ are not c-amenable.

Let us now turn to the Potts model. Theorem 1.4 was partially known in the case of $\mathbb{Z}^d$ (all results mentioned here are for Γ the group of translations). In the Ising case $q = 2$, van den Berg and Steif [2] showed the first two items (and a weakening of the third item) and the second author [29] showed the third item. In the general case $q \geq 2$, Häggström and Steif [19] showed that $\mu_{\beta,q}$ is fv-ffiid (and it also follows from the proof that it is ffiid with exponential tails) for β small enough, and the second author [30, Corollary 1.5] showed that $\mu_{\beta,q}$ is ffiid with power-law tails for all $\beta < \beta_c(q)$ and that is

ffiid when $d = 2$, $q \in \{2, 3, 4\}$ and $\beta = \beta_c(q)$.

There are very few known results regarding the coding properties of the loop $O(n)$ or long-range Ising models for general values of the parameters. Our results for these two models, as well as the random-cluster model, follow from coding results that apply to a general framework of monotone specifications, which will be discussed in Section 2. While we developed this methodology with the random-cluster model in mind, it is sufficiently broad to also include the two other models.

Van den Berg and Steif's work [2], which we mentioned above in the context of the Ising model, also proves results for monotone Markov random fields. While their methods readily extend to finite-range models, they do not apply to infinite-range models — i.e., models where the dependence of the conditional distribution at the origin on the boundary conditions is not confined to a bounded box around the origin. The main technical innovation of our technique is that it allows one to consider such models. For example, the random-cluster model is an infinite-range model. To see this, observe that the conditional distribution of an edge depends not only on its neighboring edges, but rather on the connectivity of its two endpoints, which may require looking arbitrarily far away from the given edge. In fact, this shows that the conditional distribution does not even depend continuously on the boundary condition. Our framework requires neither this continuity assumption nor a 'high-noise' assumption, both of which are used in works such as [13]. Instead, we rely only on monotonicity and uniqueness, thereby allowing us to obtain non-perturbative results.

1.7 Organization of the paper

The next Section 2 introduces the general framework of monotone specifications. In particular, it states Theorem 2.1, Theorem 2.2 and Corollary 2.3, which are the main technical results of this paper. Section 3 proves the theorems introduced above, assuming the theorems of Section 2. Section 4 defines the coupled dynamics that is the basis for constructions of all codings in this paper, and then proves Theorem 2.1. Section 5 proves a mixing time result, which is then used in the subsequent Section 6 to prove Theorem 2.2. Finally, Section 7 proves Corollary 2.3.

1.8 Notation

We now set up some notation which will be used for the rest of the paper. Let G be an infinite locally finite quasi-transitive graph on a countable set $\mathbb{V}$ (all graphs in this paper satisfy these conditions). Denote the graph distance in G by $\mathrm{dist}(\cdot, \cdot)$. For sets $U, V \subset \mathbb{V}$, we write $\mathrm{dist}(U, V) := \min_{u \in U, v \in V} \mathrm{dist}(u, v)$ and $\mathrm{dist}(u, V) := \mathrm{dist}(\{u\}, V)$. We denote $\partial V := \{u \in \mathbb{V} : \mathrm{dist}(u, \mathbb{V} \setminus V) = 1\}$ and $\partial v := \partial\{v\}$. Let

$$\mathbb{V}_{v,r} := \{u \in \mathbb{V} : \mathrm{dist}(u, v) \leq r\}$$

be the ball of radius r around v. We also denote

$$B(r) := \max_{v \in \mathbb{V}} |\mathbb{V}_{v,r}|.$$

Recall that Γ is a group acting on $\mathbb{V}$. We extend the action of Γ to $A^{\mathbb{V}}$ (for any set A) by

$$\gamma\omega := (\omega_{\gamma^{-1}v})_{v \in \mathbb{V}}.$$

Let μ and ν be probability measures on a common discrete space $\mathcal{A}$. We denote by $\|\mu - \nu\|_{TV}$ the total variation distance between μ and ν, i.e.,

$$\|\mu - \nu\|_{TV} := \tfrac{1}{2} \sum_{a \in \mathcal{A}} |\mu(a) - \nu(a)| = \max_{A \subset \mathcal{A}} |\mu(A) - \nu(A)|.$$

When $\mathcal{A}$ is partially ordered, we say that μ is *stochastically dominated* by ν, and write $\mu \leq_{\mathrm{st}} \nu$, if $\mu(A) \leq \nu(A)$ for any increasing event A.

2 Finitary codings for monotone specifications

2.1 The general framework

Let $\mathbb{V}$ be countably infinite, Γ be a group acting quasi-transitively on $\mathbb{V}$ and $(S, \leq)$ be a totally ordered discrete spin space with a maximal element $+$. We extend the order on S to the product partial order on

$$\Omega := S^{\mathbb{V}},$$

whose maximal element we denote by $+$. Thus, given two elements $\omega, \omega' \in \Omega$,

$$\omega \leq \omega' \quad \Longleftrightarrow \quad \omega_v \leq \omega'_v \quad \text{for all } v \in \mathbb{V}.$$

For a finite $V \subset \mathbb{V}$ and $\tau \in \Omega$, define

$$\Omega^\tau_V := \{\omega \in \Omega : \omega_{\mathbb{V} \setminus V} = \tau_{\mathbb{V} \setminus V}\}.$$

Denote

$$\Omega^+ := \bigcup_{V \subset \mathbb{V} \text{ finite}} \Omega^+_V = \{\omega \in \Omega : \omega \text{ agrees with } + \text{ outside a finite set}\}.$$

An *upwards specification* is a family of measures

$$\rho = \{\rho^\tau_V\}_{V \subset \mathbb{V} \text{ finite}, \ \tau \in \Omega^+},$$

where ρ^τ_V is a probability measure supported on Ω^τ_V, that satisfies the consistency relations that, for any finite $U \subset V \subset \mathbb{V}$ and any $\tau, \tau' \in \Omega^+$,

$$\rho^\tau_V = \rho^{\tau'}_V \qquad \text{whenever } \tau_{\mathbb{V} \setminus V} = \tau'_{\mathbb{V} \setminus V}$$

and

$$\rho^\tau_V(\,\cdot\,|\,\Omega^\tau_U) = \rho^\tau_U \qquad \text{whenever } \rho^\tau_V(\Omega^\tau_U) > 0.$$

If we expand this family by defining measures for any $\tau \in \Omega$ and requiring the same consistency relations, we obtain a specification. Upwards specifications are simpler objects than specifications – for one thing, there are only countably many measures in an upwards specification, whereas a specification requires uncountably many measures. For any $v \in \mathbb{V}$, we write ρ^τ_v as a shorthand for $\rho^\tau_{\{v\}}$.

An upwards specification is Γ-*invariant* if

$$\rho^{\gamma\tau}_{\gamma V}(\gamma^{-1}\omega \in \cdot) = \rho^\tau_V \qquad \text{for any } \gamma \in \Gamma, \ V \subset \mathbb{V} \text{ finite and } \tau \in \Omega^+.$$

An upwards specification ρ is *irreducible* if, for any finite V, the set $\{\omega \in \Omega^+ : \rho^+_V(\omega) > 0\}$ contains $+$ and is connected in the Hamming graph on Ω^+. An upwards specification is called *finite-energy* if there $\inf_{\tau \in \Omega^+, v \in \mathbb{V}} \rho^\tau_{\{v\}}(\tau) > 0$. Intuitively, this condition imposes a uniform lower bound on the cost of changing the configuration at a single vertex. It is straightforward to see that finite-energy is stronger than irreducibility. An upwards specification is *monotonic* if

$$\rho^\tau_V \leq_{\mathrm{st}} \rho^{\tau'}_V \qquad \text{for any } V \subset \mathbb{V} \text{ finite and } \tau, \tau' \in \Omega^+ \text{ such that } \tau \leq \tau'.$$

When S has a minimal element $-$, we similarly define a notion of a *downwards specification* by replacing $+$ with $-$, the minimal element in Ω, and replacing Ω^+ with

Ω^-, the set of configurations which equal $-$ outside a finite set. When S has both a minimal and maximal element, we may also define a notion of an *upwards-downwards specification*, where Ω^+ is replaced with $\Omega^+ \cup \Omega^-$ above. Such an upwards-downwards specification ρ may be equivalently seen as a pair (ρ^+, ρ^-), where ρ^+ is an upwards specification and ρ^- is a downwards specification. In this case, Γ-invariance of ρ is equivalent to Γ-invariance of both ρ^+ and ρ^-, while monotonicity of ρ is equivalent to monotonicity of both ρ^+ and ρ^- along with an ordering between ρ^+ and ρ^- in the sense that

$$\rho_V^\tau \leq_{\mathrm{st}} \rho_V^{\tau'} \quad \text{for any } V \subset \mathbb{V} \text{ finite and } \tau \in \Omega^-, \tau' \in \Omega^+ \text{ such that } \tau \leq \tau'.$$

On the other hand, by irreducibility of ρ, we mean that both ρ^+ and ρ^- are irreducible, without requiring a joint condition.

Let ρ be a monotone upwards specification. By monotonicity, ρ_U^+ stochastically dominates ρ_V^+ whenever $U \subset V$. Thus, there exists a weak limit

$$\mu^+ := \lim_{V \uparrow \mathbb{V}} \rho_V^+.$$

The limit μ^+ is in general a sub-probability measure on Ω (not necessarily supported on Ω^+), and is Γ-invariant when ρ is. If S is finite, then μ^+ is a probability measure. When ρ is a monotone downwards specification, we similarly define μ^-. In particular, when ρ is an upwards-downwards specification, both μ^+ and μ^- are well defined.

2.2 The general results

We now state the three general results that will be used to prove the main theorems of Section 1.

Theorem 2.1. *Let G be an infinite graph on vertex set $\mathbb{V}$ and let Γ be a group acting quasi-transitively on $\mathbb{V}$ by automorphisms of G. Let S be a totally ordered discrete spin space.*

1. *Suppose that S has a maximal element and let ρ be a monotone Γ-invariant irreducible upwards specification. If μ^+ is a probability measure, then it is Γ-fiid.*

2. *Suppose that S is finite and let ρ be a monotone Γ-invariant irreducible upwards-downwards specification. Then μ^+ is Γ-fiid with a coding radius that satisfies*

$$\mathbb{P}(R_v > r) \leq (|S| - 1) \cdot \left\| \rho_{\mathbb{V}_{v,r}}^+(\sigma_v \in \cdot) - \rho_{\mathbb{V}_{v,r}}^-(\sigma_v \in \cdot) \right\|_{TV} \qquad \text{for all } v \in \mathbb{V} \text{ and } r \geq 0. \tag{2.1}$$

   In particular, if $\mu^+ = \mu^-$ then μ^+ is Γ-ffiid.

The state spaces for the i.i.d. process Y in the above theorem are unrestricted (one may think of $(T, \mathcal{T})$ as Lebesgue space on $[0, 1]$). In the next section, we wish to control the 'amount of temporal information' used by the coding – heuristically, how many times must the factor map query a (finite-valued) input at any vertex. To this end, we equip the space $(T, \mathcal{T})$ with a more explicit structure, namely, we assume that $T = \tilde{T}^{\mathbb{N}}$, where $\tilde{T}$ is *finite*. Recall that the coding radius of a coding $\varphi \colon T^{\mathbb{V}} \to S^{\mathbb{V}}$ at a vertex $v \in \mathbb{V}$ and a configuration $y \in T^{\mathbb{V}}$ is the minimal $r \geq 0$ such that $\varphi(y)_v$ is determined by $(y_w)_{w \in \mathbb{V}_{v,r}}$. We analogously define $R_v^*(y)$, the *space-time coding radius* of φ at v and y, to be the minimal $r \geq 0$ such that $\varphi(y)_v$ is determined by $(y_w(i))_{w \in \mathbb{V}_{v,r}, 0 \leq i \leq r}$. We say that such a coding is *space-time finitary* if $R_v^* = R_v^*(Y)$ is almost surely finite for every v. In this setting, when Y is said to be an i.i.d. process, we mean that $\{Y_v(n)\}_{v \in \mathbb{V}, n \in \mathbb{N}}$ is a collection of i.i.d. random variables supported on the finite set $\tilde{T}$.

We add one final piece of notation before we state the theorem: an upwards-downwards specification ρ is *marginally finite* if $\{\rho_v^\tau(\sigma_v \in \cdot)\}_{v \in \mathbb{V}, \tau \in \Omega^+ \cup \Omega^-}$ is a finite collection of distinct measures.

Theorem 2.2. *Let G be an infinite graph on vertex set $\mathbb{V}$ and let Γ be a group acting quasi-transitively on $\mathbb{V}$ by automorphisms of G. Suppose that*

$$\mathbb{V}_{v,r} \setminus \mathbb{V}_{v,r-1} \not\subset \mathbb{V}_{u,r} \qquad \text{for any distinct } u, v \in \mathbb{V} \text{ and } r \geq 0. \tag{2.2}$$

Let S be a totally ordered finite spin space and ρ be a monotone Γ-invariant irreducible marginally finite upwards-downwards specification.

- *If $\mu^+ = \mu^-$, then there exists a space-time finitary coding from an i.i.d. process Y to μ^+.*
- *Suppose that there exist $C, c > 0$ such that*

$$\left\| \rho^+_{\mathbb{V}_{v,r}}(\sigma_v \in \cdot) - \rho^-_{\mathbb{V}_{v,r}}(\sigma_v \in \cdot) \right\|_{TV} \leq C e^{-cr} \qquad \text{for all } v \in \mathbb{V} \text{ and } r \geq 0. \tag{2.3}$$

If G has sub-exponential growth, i.e., $B(r) = \exp(o(r))$, then the tails of the space-time coding radius beats any stretched-exponential, i.e., $\mathbb{P}(R_v^* \geq r) \leq \exp(-r^{1-o(1)})$. Moreover, if G has growth $B(r) = \exp(o(\frac{r}{\log r}))$, then the space-time coding radius has exponential tails.

Condition (2.3) is commonly referred to as *weak spatial mixing*. The proof of the second item in Theorem 2.2 relies on controlling the mixing-time of a natural single-site dynamics for specifications satisfying weak spatial mixing (see Section 5).

In the case in which $\mathbb{Z}^d$ and Γ is restricted to translations of the lattice, we can use the setup of [29] to deduce the existence of fv-ffiid codings:

Corollary 2.3. *Let G be $\mathbb{Z}^d$ or its line graph, Γ be the group of translations, S be a totally ordered finite spin space, and ρ be a monotone Γ-invariant irreducible marginally finite upwards-downwards specification that satisfies (2.3). Then μ^+ is Γ-fv-ffiid with stretched-exponential tails.*

We believe that codings from a finite-valued i.i.d. process should exist for a much larger class of graphs (perhaps graphs satisfying (2.2) and having sub-exponential growth).

3 Proofs of main results

In this section, we prove the theorems stated in Section 1. All the theorems, with the exception of Theorem 1.2, will follow from the general results given in Section 2. The proofs of these general results are postponed to the subsequent sections.

3.1 The random-cluster model

In this section, we prove the three main theorems about the random-cluster model, namely, Theorem 1.1, Theorem 1.2 and Theorem 1.3. We first place the model in the general framework of Section 2.

Consider the random-cluster model with parameters $p \in [0, 1]$ and $q \geq 1$ on an infinite quasi-transitive graph $G = (\mathbb{V}, E)$. While the random-cluster model is most naturally defined on the edges of G, our abstract definitions are stated for models defined on the vertex set of a graph, and therefore we view the random-cluster model as a 'living' on the line graph $\mathcal{G}$ of G, whose vertex set is E. Let $S := \{0, 1\}$ and let Γ be a group acting quasi-transitively on $\mathbb{V}$ by automorphisms of G. The random-cluster model has two natural specifications associated to it – the free-DLR and wired-DLR specifications – corresponding to the choice of $i \in \{0, 1\}$ in the definition of the model. These are denoted by ρ^{free} and ρ^{wired} and defined by

$$\rho^{\#,\tau}_F(\omega) \propto p^{o_F(\omega)}(1-p)^{c_F(\omega)} q^{k_F^{\#}(\omega)} \mathbf{1}_{\Omega_F^\tau}(\omega), \qquad \omega \in \{0, 1\}^E, \tag{3.1}$$

where F is a finite subset of E, $o_F(\omega)$ and $c_F(\omega)$ are the numbers of open and closed edges in F, and $k_F^\#(\omega)$ is the number of open clusters that intersect an endpoint of some edge in F, with $k_F^{\mathrm{wired}}(\omega)$ counting only finite clusters and $k_F^{\mathrm{free}}(\omega)$ counting both finite and infinite clusters (we note that the usual notion of 'free' and 'wired' boundary conditions correspond to $\rho_F^{\mathrm{free},\mathbf{0}}$ and $\rho_F^{\mathrm{wired},\mathbf{1}}$, where $\mathbf{0}$ and $\mathbf{1}$ correspond to the all-closed and all-open configurations, respectively).

Remark 3.1. A specification ρ gives rise to the notion of a Gibbs measure (also called a DLR state), which is a probability measure μ such that, for every finite $F \subset E$ and μ-a.e. $\tau \in \{0,1\}^E$, the conditional law of ω under μ given that $\omega \in \Omega_F^\tau$ is ρ_F^τ. Thus, in general, the random-cluster model has two notions of Gibbs measures. For amenable quasi-transitive connected graphs, any Gibbs measure (of either of the two specifications) has at most one infinite cluster with probability 1; in this case, there is no distinction between free-DLR and wired-DLR Gibbs measures. In particular, $p_c^0(q) = p_c^1(q)$ for every amenable graph. For more general graphs, the number of infinite clusters may be infinite, in which case, this distinction is essential: for example, $\phi_{p,q}^1$ may not satisfy the free-DLR condition for certain graphs and values of p and q, but always satisfies the wired-DLR condition. Similarly, $\phi_{p,q}^0$ may not be a wired-DLR random-cluster measure (see Section 6.4 of [14] for further discussion). Also, every Gibbs measure (in either the free or wired DLR sense) stochastically dominates $\phi_{p,q}^0$ and is stochastically dominated by $\phi_{p,q}^1$.

Let $\rho = (\rho^+, \rho^-)$ be the upwards-downwards specification given by $\rho^+ := (\rho_F^{\mathrm{wired},\tau})_{F \subset E \text{ finite }, \tau \in \Omega^+}$ and $\rho^- := (\rho_F^{\mathrm{free},\tau})_{F \subset E \text{ finite }, \tau \in \Omega^-}$. Whenever $q \geq 1$, the FKG property of the random-cluster model implies that ρ is a monotone specification (see [16, Theorems 3.8 and 2.27]). It is clear from (3.1) that ρ is Γ-invariant. Since, by (3.1),

$$\left\{ \rho_e^{\mathrm{free},\tau}(\sigma_e = s),\, \rho_e^{\mathrm{wired},\tau}(\sigma_e = s) : \tau \in \Omega,\, e \in E,\, s \in \{0,1\} \right\} = \left\{ p, 1-p, \tfrac{p}{p+(1-p)q}, \tfrac{(1-p)q}{p+(1-p)q} \right\},$$

it is clear that ρ is irreducible and marginally finite.

Proof of Theorem 1.1. The free and wired random-cluster measures, $\phi_{p,q}^0$ and $\phi_{p,q}^1$, are precisely the measures μ^- and μ^+ obtained from the upwards-downwards specification ρ. Therefore, the first and second items of Theorem 1.1 are immediate consequences of the first and second items of Theorem 2.1, respectively.

We now turn to the third item of Theorem 1.1. Define $F_{e,r}$ be the set of edges whose distance (taken in the line graph of G) from e is at most r. For the third item of Theorem 1.1, it remains only to show that, when $p < p_c^1(q)$, there exist $C, c > 0$ such that

$$\left\| \rho_{F_{e,r}}^{\mathrm{wired},+}(\sigma_e \in \cdot) - \rho_{F_{e,r}}^{\mathrm{free},-}(\sigma_e \in \cdot) \right\|_{TV} \leq C e^{-cr} \qquad \text{for all } e \in E \text{ and } r \geq 0.$$

It is well-known that one may couple samples from $\rho_{F_{e,r}}^{\mathrm{wired},+}$ and $\rho_{F_{e,r}}^{\mathrm{free},-}$ so that they agree on e whenever the endpoints of e are disconnected from the boundary. Thus, if v is an endpoint of e,

$$\left\| \rho_{F_{e,r}}^{\mathrm{wired},+}(\sigma_e \in \cdot) - \rho_{F_{e,r}}^{\mathrm{free},-}(\sigma_e \in \cdot) \right\|_{TV} \leq \phi_{\mathbb{V}_{v,r-1},p,q}^1(e \leftrightarrow \partial \mathbb{V}_{v,r-1}).$$

The exponential decay of the right-hand side is exactly the content of [10, Theorem 1.2] which establishes the sharpness of the phase transition for the random-cluster model. $\square$

Proof of Theorem 1.3. In light of the above, Theorem 1.3 is a direct application of Corollary 2.3. $\square$

Proof of Theorem 1.2. We begin with the first item of the theorem, which we proceed to establish by contradiction. Thus, we assume towards a contradiction that $\phi_{p,q}^0 \neq \phi_{p,q}^1$ and

that $\phi^1_{p,q}$ is Γ-ffiid for some group Γ acting quasi-transitively by automorphisms. Since G is amenable, there exists at most one infinite connected component $\phi^i_{p,q}$-almost surely; thus, $\phi^i_{p,q}$ is both a free-DLR and a wired-DLR Gibbs state (see Remark 3.1 and [14, Proposition 6.19]). It thus does not matter which specification we work with; we will use the wired-DLR specification for concreteness, but denote it as ρ for notational clarity.

We begin by assuming that Γ acts transitively on E. We write ω for a generic random element of $\{0,1\}^E$. Since G is amenable, there exists a sequence $F'_n \subset \mathbb{V}$ of non-empty finite subsets such that $|\partial F'_n|/|F'_n| \to 0$ as $n \to \infty$. Letting F_n denote the set of edges spanned by F'_n, and ∂F_n denote the set of edges in $E \setminus F_n$ that share an endpoint with an edge in F_n, we have that $|\partial F_n|/|F_n| \to 0$ as $n \to \infty$. For $n \geq 1$, denote

$$Z_n := \frac{1}{|F_n|} \sum_{e \in F_n} \omega_e.$$

Denote $a_0 := \phi^0_{p,q}(\omega_e)$ and $a_1 := \phi^1_{p,q}(\omega_e)$ and note that $a_0 < a_1$. Since $\phi^1_{p,q}$ is Γ-ffiid, it follows that the convergence in the ergodic theorem occurs at an exponential rate for $\phi^1_{p,q}$ (this was shown in [3] for the case $G = \mathbb{Z}^d$, and the proof there goes through with no changes for an arbitrary quasi-transitive graph G). Hence, denoting $a := \frac{1}{2}(a_0 + a_1)$,

$$\phi^1_{p,q}(Z_n \leq a) \leq Ce^{-2c|F_n|} \qquad \text{for some } C, c > 0 \text{ and for all } n \geq 1.$$

By Markov's inequality and the fact that $\phi^1_{p,q}$ is a wired-DLR Gibbs measure,

$$\phi^1_{p,q}(T_n) \leq Ce^{-c|F_n|}, \qquad \text{where } T_n := \left\{ \tau \in \{0,1\}^E : \rho^\tau_{F_n}(Z_n \leq a) \geq e^{-c|F_n|} \right\}.$$

Let us show that the all 0 configuration $\mathbf{0}$ belongs to T_n for large n. By monotonicity of the specification ρ, T_n is a decreasing set for each n, and thus it suffices to show that T_n is non-empty for large n. Indeed, by Markov's inequality and the fact that $\phi^0_{p,q}$ is $\textit{also}$ a wired-DLR state,

$$\begin{aligned}
1 - \phi^0_{p,q}(T_n) &= \phi^0_{p,q}\left(\rho^\omega_{F_n}(Z_n \leq a) < e^{-c|F_n|} \right) \\
&\leq \phi^0_{p,q}\left(\rho^\omega_{F_n}(Z_n \leq a) \leq o(1) \right) \\
&= \phi^0_{p,q}\left(\rho^\omega_{F_n}(Z_n > a) \geq 1 - o(1) \right) \leq (1 + o(1)) \cdot \phi^0_{p,q}\left(\rho^\omega_{F_n}(Z_n > a) \right) \\
&\leq (1 + o(1)) \cdot \phi^0_{p,q}(Z_n > a) \leq (1 + o(1))\frac{a_0}{a}.
\end{aligned}$$

Since $a_0 < a$, we conclude that T_n is non-empty for large n. Let $\Omega^{\mathbf{0}}_{\partial F_n}$ denote the set of configurations that equal zero on ∂F_n. Observe that $\rho^\tau_{F_n} = \rho^{\mathbf{0}}_{F_n}$ for all n and $\tau \in \Omega^{\mathbf{0}}_{\partial F_n}$, as the free boundary conditions decouple the measure in F_n from the state of ω on $F_n^c \setminus \partial F_n$. Thus, $\Omega^{\mathbf{0}}_{\partial F_n} \subset T_n$ for large n, so that

$$\phi^1_{p,q}\left(\omega_{\partial F_n} = 0 \right) = \phi^1_{p,q}\left(\Omega^{\mathbf{0}}_{\partial F_n} \right) \leq \phi^1_{p,q}(T_n) \leq Ce^{-c|F_n|} \qquad \text{for large } n.$$

On the other hand, by finite energy, we have the lower bound

$$\phi^1_{p,q}\left(\omega_{\partial F_n} = 0 \right) \geq \left(\frac{p}{p + (1-p)q} \right)^{|\partial F_n|} \qquad \text{for all } n.$$

Since $\frac{|\partial F_n|}{|F_n|} \to 0$ as $n \to \infty$, we have reached a contradiction. This completes the first item in the case where Γ acts transitively on E.

To handle the case where Γ acts on E quasi-transitively, we let $(O_1, \ldots O_k)$ be the (finitely many) orbits of E under Γ. By possibly extracting a subsequence, we may assume

that $r_i = \lim_{n \to \infty} |O_i \cap F_n|/|F_n|$ exists for each i. We now define $a_j = \sum_{i=1}^{k} r_i \cdot \phi_{p,q}^j(\omega_{e_i})$, where e_i is an arbitrary element of O_i. By construction, $\phi_{p,q}^j(Z_n)$ converges to a_j, and the proof goes through, as above.

We now turn to the second item of the theorem. The proof given above adapts to this case with some modification which we now explain. First, we choose F_n differently. By definition of c-amenable, there exists a sequence $(F'_n, H'_n)_n$ of non-empty finite subsets of $\mathbb{V}$ such that H'_n is connected, $\partial F'_n \subset H'_n$ and $|H'_n|/|F'_n| \to 0$ as $n \to \infty$. Let H_n denote the set of edges incident to a vertex in H'_n, and let F_n denote the set of edges incident to a vertex in F'_n but not in H_n. Then H_n is a connected set of edges, which contains ∂F_n and is disjoint from F_n. We replace the occurrences of ∂F_n in the proof with H_n. The proof then goes through once we interchange the roles of free and wired, replace the all zero configuration $\mathbf{0}$ with the all one configuration $\mathbf{1}$, and replace Z_n by $(1 - Z_n)$. The reason for the H_n in this case (and not merely ∂F_n) is that, while the $\mathbf{0}$ configuration acts as a strong 'decoupling boundary condition' in the sense that $\rho_V^\tau = \rho_V^{\mathbf{0}}$ whenever $\tau_{\partial V} = 0$, the $\mathbf{1}$ configuration has a weaker decoupling property: in order to force 'true wired boundary condition' in the sense that $\rho_V^\tau = \rho_V^{\mathbf{1}}$, it is not sufficient to merely have $\tau_{\partial V} = 1$, but rather one needs $\tau_H = 1$ on a set H which connects ∂V from outside of V. This is the reason for our assumption of c-amenability. $\qquad \square$

3.2 The Potts model

In this section, we prove Theorem 1.4. Due to the relation between the Potts and random-cluster models (namely, the Edwards–Sokal coupling), the theorem follows (morally) from the results about the random-cluster (more specifically, the first and second items from Theorem 1.1 and the third item from Theorem 1.3). However, before proving Theorem 1.4, there is a technical issue we must face: Potts configurations belong to $\{1, \ldots, q\}^{\mathbb{V}}$, while random-cluster configurations belong to $\{0, 1\}^E$. Let Γ be a group acting on $\mathbb{V}$ by automorphisms of G. If $\mathcal{G}$ is the line graph of G, then Γ can be canonically embedded in the automorphism group of $\mathcal{G}$. In a slight abuse of notation, we allow $\gamma \in \Gamma$ to act on E through this identification. This allows us to discuss factors from processes on $\mathbb{V}$ to processes on E. The lemma below shows that we can produce any i.i.d. process on E using an i.i.d. process on $\mathbb{V}$ in a Γ-equivariant manner.

Lemma 3.2. *Let G be an infinite quasi-transitive graph and let Γ denote its full automorphism group. Then any i.i.d. process on E is a finitary Γ-factor of an i.i.d. process on $\mathbb{V}$ with bounded coding radius.*

Proof. Let $X = (X_e)_{e \in E}$ be an i.i.d. process, where each X_e takes values in a measurable space $(T, \mathcal{T})$. Let Δ be the maximal degree of G. We will show by direct construction that X is a Γ-factor of the i.i.d. process $(Y, Z) = (Y_v, Z_v)_{v \in \mathbb{V}}$, where Y_v and Z_v are independent, $Y_v = (Y_v^1, \ldots, Y_v^\Delta)$ is a collection of Δ i.i.d. random variables having the same distribution as X_e, and Z_v is a uniform random variable on $[0, 1]$.

Define
$$\psi \colon (T^\Delta \times [0, 1])^{\mathbb{V}} \to T^E$$
by
$$\psi(y^1, \ldots, y^\Delta, z)_{\{u,v\}} := \begin{cases} y_u^{|\{w \sim u \,:\, z_u \leq z_w \leq z_v\}|} & \text{if } z_u < z_v \\ y_v^{|\{w \sim v \,:\, z_v \leq z_w \leq z_u\}|} & \text{if } z_u \geq z_v \end{cases}.$$

In words, the value associated to an edge $\{u, v\}$ is obtained as follows: z induces an order on the vertices of G; the edge $\{u, v\}$ chooses its smaller endpoint with respect to this order – say, u – and takes on the value y_u^i for some $i \in \{1, \ldots, \Delta\}$. To determine the value of i, the set of edges that chose u is ordered according to the value of z at the *other* endpoint – the first edge takes y_u^1, the second y_u^2, etc.

If $z_u \neq z_v$ for any distinct $u, v \in \mathbb{V}$, then no value of y_u^i is assigned to more than one edge. Thus, we have $\psi(Y, Z) \overset{d}{=} X$. Since it is clear that ψ is Γ-equivariant and has coding radius at most 2, the lemma follows. $\qquad\square$

We also require the following simple lemma.

Lemma 3.3. *The composition of finitary codings with (stretched-)exponential tails is also a finitary coding with (stretched-)exponential tails.*

Proof. Let X, Y and Z be processes on $\mathbb{V}$. Let φ be a coding from Y to X and let φ' be a coding from Z to Y, both having (stretched-)exponential tails. Denote $\tilde{\varphi} = \varphi' \circ \varphi$. We denote the coding radii of φ, φ' and $\tilde{\varphi}$ at v by R_v, R_v' and $\tilde{R}_v$, respectively.

We first handle the case of exponential tails. Let $C, c > 0$ be such that $\mathbb{P}(R_v > r) \leq Ce^{-cr}$ and $\mathbb{P}(R_v' > r) \leq Ce^{-c'r}$ for all $r > 0$. Let Δ be the maximal degree of G and set $a := c/[2(c + \log \Delta)]$. Fix $v \in \mathbb{V}$ and define

$$S_{v,r} := \bigcap_{u \in \mathbb{V}_{v,ar}} \{R_u \leq (1-a)r\}.$$

By the union bound and the definition of a,

$$\mathbb{P}[S_{v,r}^c] \leq B(ar) \cdot Ce^{-c(1-a)r} \leq C\Delta^{ar} e^{-c(1-a)r} \leq Ce^{-cr/2}.$$

It straightforward to see that on the event $\{R_v' \leq ar\} \cap S_{v,r}$, we have that $\tilde{R}_v \leq r$. Thus,

$$\mathbb{P}[\tilde{R}_v > r] \leq \mathbb{P}[S_{r,v}^c] + \mathbb{P}[R_v' > ar] \leq Ce^{-cr/2} + Ce^{c'ar}.$$

This completes the proof in the case of exponential tails.

The case of stretched-exponential tails follows in a similar same way, where we let $C > 0$ and $0 < c, c' < 1$ be such that $\mathbb{P}(R_v > r) \leq Ce^{-r^c}$ and $\mathbb{P}(R_v' > r) \leq Ce^{-r^{c'}}$, and we set $a := r^{\frac{c}{2}-1}/\log \Delta$. $\qquad\square$

Proof of Theorem 1.4. Let Γ be the full automorphism group of G; we set $p = 1 - e^{-\beta}$, and assume that $\mu_{\beta,q}^0 = \mu_{\beta,q}^1 = \cdots = \mu_{\beta,q}^q$. Using the Edwards-Sokal coupling (see [16, Theorem 4.91]), this implies that $\phi_{p,q}^0 = \phi_{p,q}^1$. By Theorem 1.1, there exists an i.i.d. process Y on the edges, taking values in some measurable space $(T, \mathcal{T})$, and a finitary coding $\varphi : T^E \to \{0, 1\}^E$ from Y to $\phi_{p,q}$. Since φ is invariant under any automorphism of the line graph of G, we have that $\varphi \circ \gamma = \gamma \circ \varphi$ for any $\gamma \in \Gamma$ (since Γ is canonically embedded in the automorphism group of the line graph).

Next, we wish to construct the Edwards–Sokal coupling in a Γ-equivariant manner. Define

$$\Psi : [0, 1]^\mathbb{V} \times \{1, \ldots, q\}^\mathbb{V} \times \{0, 1\}^E \to \{1, \ldots, q\}^\mathbb{V}$$

by

$$\Psi(z, \sigma, \omega)_v := \sigma_u, \qquad \text{where } u = \operatorname{argmin}\{z_w : w \overset{\omega}{\leftrightarrow} v\},$$

where we recall that $w \overset{\omega}{\leftrightarrow} v$ indicates that there exists a path of ω-open edges connecting w and v. Heuristically, $\Psi(z, \sigma, \omega)_v$ outputs the color σ_u, where u is the vertex in the connected component of v which has the minimal z value. By construction, Ψ is Γ-equivariant. Let (Z, Σ) be an i.i.d. process on $\mathbb{V}$, where Z_v and Σ_v are independent and uniform on $[0, 1]$ and $\{1, \ldots, q\}$, respectively. Then the Edwards–Sokal coupling (see [16, Theorem 4.91]) implies that $\Psi(Z, \Sigma, \omega) \overset{d}{=} \mu_{\beta,q}$ whenever ω is sampled from $\phi_{p,q}$ independently of (Z, Σ). This implies that $\mu_{\beta,q}$ is ffid.

For the second item, we assume that $\beta < \beta_c^{\mathrm{w}}(q)$. This implies that $p < p_c(q)$, which implies (by the third item of Theorem 1.1) that the coding radius of φ has exponential

tails. Suppose now that Y and (Z, Σ) are independent. Then the composition $\Psi \circ (\mathrm{id}, \mathrm{id}, \varphi)$ is a coding from (Z, Σ, Y) to $\mu_{\beta,q}$. By Lemma 3.2, this implies that we can create a coding Ψ' from an i.i.d. process on $\mathbb{V}$ to the Potts model $\mu_{\beta,q}$. By Lemma 3.3, this coding has exponential tails, proving the second item of the theorem.

For the final item, we set $G = \mathbb{Z}^d$ and Γ to be the translation group of the lattice. The extra structure here allows us to skip the more complicated constructions above and do things 'by hand.' Let $\{e_1, \ldots, e_d\}$ denote the the standard basis of $\mathbb{Z}^d$. Any $e \in E$ has a unique representation $e = \{v, v + e_i\}$, where $v \in \mathbb{Z}^d$ and $1 \leq i \leq d$. Define $\tilde{\psi} : (T^d)^{\mathbb{V}} \to T^E$ by

$$\tilde{\psi}(y^1, \ldots, y^d)_e := y_v^i, \qquad \text{where } e = \{v, v + e_i\}.$$

We also define $\tilde{\Psi} : \{1, \ldots, q\}^{\mathbb{V}} \times \{0, 1\}^E$ by

$$\tilde{\Psi}(\sigma, \omega)_v := \sigma_u \qquad \text{where } u = \min\{w : w \overset{\omega}{\leftrightarrow} v\},$$

where the minimum over vertices is taken in the lexicographical order on $\mathbb{Z}^d$. Both $\tilde{\psi}$ and $\tilde{\Psi}$ are Γ-equivariant, as the lexicographical order is translation-invariant. Theorem 1.3 gives us a Γ-fv-ffiid coding $\tilde{\varphi}$ for $\phi_{p,q}$ with stretched-exponential tails. Then $\tilde{\Psi} \circ (\mathrm{id}, \tilde{\varphi} \circ \tilde{\psi})$ is a Γ-fv-ffiid coding for $\mu_{\beta,q}$. By Lemma 3.3, the coding radius of this map has stretched-exponential tails, completing the proof. $\qquad\square$

3.3 The loop $O(n)$ model

In this section, we prove Theorem 1.5.

We define the so-called *spin representation* of the loop $O(n)$ model as follows: set $\mathbb{V} := \mathbb{T}$ and $S := \{+, -\}$. For any $\sigma \in S^{\mathbb{V}}$, define the probability measure μ_V^τ defined by the formula

$$\mu_V^\tau(\sigma) := \frac{n^{k(\sigma)} x^{e(\sigma)}}{\mathbf{Z}_V^\tau} \cdot \mathbf{1}_{\{\sigma_{\mathbb{V} \setminus V} = \tau_{\mathbb{V} \setminus V}\}},$$

where $k(\sigma) + 1$ is the sum of the number of connected components of pluses and minuses in σ that intersect V or its neighborhood, $e(\sigma) := \sum_{u \sim v} \mathbf{1}_{\sigma_u \neq \sigma_v}$ is the number of edges $\{u, v\}$ that intersect V and have $\sigma_u \neq \sigma_v$, and $\mathbf{Z}_V^\tau$ is the unique constant making μ_V^τ a probability measure. Clearly, both $k(\sigma)$ and $e(\sigma)$ depend on V, but we omit it in the notation for brevity.

The spin representation is related to the original model in the following manner: if σ is distributed as μ_V^τ, then its 'domain walls', or the lines that separate $+$ from $-$, are distributed as $\nu_{\Omega,n,x}$ (see [9, Proposition 3] for more details).

Proof of Theorem 1.5. The family of measures $\{\mu_V^\tau\}_{V \subset \mathbb{V} \text{ finite}, \, \tau \in \Omega^+ \cup \Omega^-}$ defines an upwards-downwards specification. It is shown in [9, Theorem 4] that the spin representation is monotonic whenever $n \geq 1$ and $nx^2 \leq 1$. Thus, we can define the infinite-volume limits μ^+ and μ^-, which are the largest and smallest possible measures, respectively. The domain walls of μ^+ and μ^- are both distributed as $\nu_{n,x}$, the unique periodic Gibbs measure of the loop $O(n)$ model. The operation that maps a spin configuration to its domain walls has a finite coding radius, and therefore Lemma 3.2 allows us to transfer coding properties of μ^+ or μ^- to $\nu_{n,x}$.

The upwards-downwards specification satisfies the finite-energy property, and is therefore irreducible, and is clearly invariant under all automorphisms of $\mathbb{T}$. We also note that the automorphism group of the line graph of $\mathbb{H}$ (which is isomorphic to the Kagome lattice) is naturally identified with the automorphism group of $\mathbb{T}$. With this in mind, the first part of Theorem 1.5 follows from the first item of Theorem 2.1. For the second part of Theorem 1.5, we note that $\mu^+ = \mu^-$ when $n \in [1, 2]$ and $x = x_c(n)$, as was shown in [9] (see the last sentence in the paragraph after Theorem 5 in that paper), so that the result follows from the second item of Theorem 2.1. $\qquad\square$

3.4 Long-range Ising models

In this section, we prove Theorem 1.6. By Theorem 2.1, the long-rang Ising model satisfies the desired coding properties if μ^+ and μ^- are the limit measures of a monotone, irreducible Γ-invariant upwards-downwards specification. As was mentioned in the discussion above Theorem 1.6, the monotonicity is a classical consequence of Griffith's inequality. Irreducibility follows from the finite-energy property, using the fact that the coupling constants are summable. Finally, Γ-invariance of the model is a starightforward consequence of the definition of the model and the assumption that the coupling constants are Γ-invariant. $\qquad\qquad\square$

4 Construction of codings via coupling-from-the-past

In this section, we introduce a dynamics, stemming from a single-site heat-bath Glauber dynamics, on a general spin system with upwards and downwards specification. The dynamics is defined in any finite volume, given any starting state, and with one of two possible boundary conditions, corresponding to the largest and smallest configurations on the complement. Crucially, the construction ensures that this dynamics couples together all such choices simultaneously. In addition, the dynamics is monotone in the sense that it maintains the partial order on the spin space. The dynamics allows us to use coupling-from-the-past to manufacture an almost-sure limit, which will be the desired coding map. This procedure leads to a proof of Theorem 2.1.

4.1 Overview of the dynamics

We begin with an informal description of the dynamics, which play a central role in this section.

Consider the sequence of finite graphs $(\mathbb{V}_{v,r})_{r\in\mathbb{N}}$ and a monotonic irreducible upwards (or upwards-downwards) specification ρ. We will define a natural single-site dynamics on each $\mathbb{V}_{v,r}$, called the $+$ dynamics. A single step of this dynamics started at an arbitrary initial configuration $\omega^{(0)} \in \Omega_{\mathbb{V}_{v,r}}^+$ is defined by applying the following evolution:

- Order $\mathbb{V}$ in a Γ-invariant way and consider the order induced on $\mathbb{V}_{v,r} = \{v_1, \dots v_m\}$.

- Obtain $\omega^{(1)}$ from $\omega^{(0)}$ by resampling the value at v_1, i.e., $\omega^{(1)}$ is sampled from $\rho_{v_1}^{\omega^{(0)}}$.

- Repeat inductively, resampling ω_{v_k} using $\omega^{(k-1)}$, until all sites have been resampled.

- The final configuration $\omega^{(m)}$ is the new state.

When ρ is an upwards-downwards specification, the above can be applied to configurations in $\Omega_{\mathbb{V}_{v,r}}^-$, producing the $-$ dynamics.

In Section 4.2, we construct the $+$ dynamics on $\mathbb{V}_{v,r}$ so that all initial configurations ξ in $\Omega_{\mathbb{V}_{v,r}}^+$ are coupled at all times. The irreducibility assumption ensures the dynamics constructed above are ergodic for any fixed r. Thus, in the limit as the number of steps of the dynamics tends to infinity, the distribution converges to $\rho_{\mathbb{V}_{v,r}}^+$. Taking r to infinity as well (in a suitable manner) gives convergence in distribution to μ^+. The method of coupling-from-the-past allows us to move from distributional limits to stronger notions of convergence, and thus construct a coding for μ^+, as will be seen in Section 4.3. Finally, in Section 4.4 we consider the $+$ and $-$ dynamics simultaneously (in a properly coupled manner), and deduce that, under the appropriate assumptions, the coding radius R satisfies (2.1). This allows us to transfer quantitative control on the total variation distance between $\rho_{\mathbb{V}_{v,r}}^+$ and $\rho_{\mathbb{V}_{v,r}}^-$ to quantitative control on the coding radius; in particular, it shows that $\mu^+ = \mu^-$ is sufficient to prove that both measures are ffiid.

4.2 The coupled dynamics

Let $Y = (Y_v)_{v \in \mathbb{V}}$ be an i.i.d. process and suppose that, for each $v \in \mathbb{V}$,

$$Y_v = (Y_{v,n})_{n \in \mathbb{N}}$$

is a collection of i.i.d. random variables. Further suppose that $(\mathcal{A}, \pi)$ and $(\mathcal{B}, \theta)$ are two probability spaces and that, for each $v \in \mathbb{V}$ and $n \in \mathbb{N}$,

$$Y_{v,n} = (A_{v,n}, B_{v,n})$$

are two independent random variables sampled from π and θ, respectively. We denote $A_n := (A_{v,n})_{v \in \mathbb{V}}$ and $B_n := (B_{v,n})_{v \in \mathbb{V}}$. The dynamics we construct are functions of Y (specifically, the n-th step of the dynamics is a function of A_n and B_n), which thus yields a coding φ from Y to μ^+. We now explain how to choose $(\mathcal{A}, \pi)$, $(\mathcal{B}, \theta)$ and φ. To remain general, we not explicitly define $(\mathcal{A}, \pi)$, $(\mathcal{B}, \theta)$ and φ, but rather let them be arbitrary objects satisfying certain properties required for the proof. This gives us a framework which is sufficiently flexible to prove both Theorem 2.1 and Theorem 2.2. After each definition, we also provide constructions to ensure that the objects we require actually exist. In fact, these will be used for the proof of Theorem 2.1; more delicate versions of these constructions, in which $\mathcal{A}$ and $\mathcal{B}$ are finite, will be required for Theorem 2.2 (see Section 6).

As mentioned above, the dynamics we construct are a coupled version of single-site Glauber dynamics of the given upwards specification. We begin by selecting $(\mathcal{A}, \pi)$ and a measurable function

$$F : \Omega^+ \times \mathbb{V} \times \mathcal{A} \to S,$$

which is used to define a single-site update. Specifically, we require that

- The random variable $F(\omega, v, \cdot)$ matches the specification at v:

$$\pi\left(F(\omega, v, \cdot) = s\right) = \rho_v^\omega[\sigma_v = s] \qquad \text{for any } \omega \in \Omega^+ \text{ and } s \in S. \qquad (4.1)$$

- F is monotonic in ω:

$$F(\omega, v, a) \le F(\omega', v, a) \qquad \text{for any } a \in \mathcal{A} \text{ and } \omega, \omega' \in \Omega^+ \text{ such that } \omega \le \omega', \qquad (4.2)$$

- F is Γ-invariant:

$$F(\omega, v, a) = F(\gamma\omega, \gamma v, a) \qquad \text{for any } a \in \mathcal{A}, \ \omega \in \Omega^+ \text{ and } \gamma \in \Gamma. \qquad (4.3)$$

At this point, we place no additional restrictions on $\mathcal{A}$ (in Section 6, we will need $\mathcal{A}$ to be finite). This allows to give a simple construction for F: set $\mathcal{A} := [0, 1]$, $\pi := \mathrm{Leb}$, the Lebesgue measure on the interval, and, for any $\omega \in \Omega^+$ and $s \in S$,

$$a_*(\omega, v, s) := \rho_v^\omega[\sigma_v < s] \qquad \text{and} \qquad a^*(\omega, v, s) := \rho_v^\omega[\sigma_v \le s].$$

It is straightforward to check that, for any $\omega \in \Omega^+$, $s' \in S$ and $a \in (a_*(\omega, v, s'), a^*(\omega, v, s'))$,

$$\min\{s \in S : a^*(\omega, v, s) \ge a\} = \max\{s \in S : a_*(\omega, v, s) \le a\} = s'.$$

Therefore, choosing an arbitrarily $s_0 \in S$, we may now define F by

$$F(\omega, v, a) := \begin{cases} \min\{s \in S : a^*(\omega, v, s) \ge a\} = \max\{s \in S : a_*(\omega, v, s) \le a\} & \text{if } a \in \mathcal{A}' \\ s_0 & \text{otherwise} \end{cases}.$$

where

$$\mathcal{A}' := \bigcap_{\omega \in \Omega^+} \bigcap_{v \in \mathbb{V}} \bigcup_{s \in S} (a_*(\omega, v, s), a^*(\omega, v, s)).$$

This choice of $\mathcal{A}'$ ensures that $F(\omega, v, a)$ is well defined; we note that $\mathcal{A}'$ has Lebesgue measure one as its complement is the countable union of Lebesgue measure zero sets. Using that the upwards specification is monotone and Γ-invariant, it follows from the definition of F that (4.2) and (4.3) hold. To see that (4.1) holds, note that

$$\pi\left(F(\omega, v, \cdot) \le s\right) = \mathrm{Leb}(\{a \in [0,1] : a^*(\omega, v, s) \ge a\}) = a^*(\omega, v, s) = \rho_v^\omega[\sigma_v \le s].$$

For any $v \in \mathbb{V}$ and $a \in \mathcal{A}$, we define

$$F_{v,a} \colon \Omega^+ \to \Omega^+$$

by

$$F_{v,a}(\omega)_u := \begin{cases} F(\omega, v, a) & \text{if } u = v \\ \omega_u & \text{otherwise} \end{cases}, \qquad \omega \in \Omega^+, \ u \in \mathbb{V}.$$

Note that (4.3) implies that $F_{v,a}(\omega) = F_{\gamma v, a}(\gamma \omega)$ for $\gamma \in \Gamma$ and $\omega \in \Omega^+$.

We now describe how to choose the updating sites. In our dynamics, the set of updated sites are deterministic; we must, however, be careful as to the *order* of the chain of single-site updates which make up a single step of the dynamics. The most straightforward way to order the sites is to associate a uniform $[0,1]$ random variable to each, and use the inherited linear order. This approach is very useful, but is slightly too rigid to allow us to study the coding properties we are interested in – specifically, this will be an issue when we are looking for codings from a finite-valued i.i.d. process. Thus, we give a more abstract definition.

Let $(\mathcal{B}, \mathfrak{B})$ be a measurable space and let $\mathcal{O} \colon \mathcal{B}^{\mathbb{V}} \times \mathbb{V}^2 \to \{0,1\}$ be measurable and Γ-invariant - i.e.

$$\mathcal{O}(\gamma\eta, \gamma u, \gamma v) = \mathcal{O}(\eta, u, v) \qquad \text{for all } \gamma \in \Gamma.$$

We regard $\eta \in \mathcal{B}^{\mathbb{V}}$ as inducing via $\mathcal{O}$ an order $\preceq_\eta$ on $\mathbb{V}$, where $\mathcal{O}(\eta, u, v) = 1$ indicates that u precedes v in this order. Formally, $\preceq_\eta$ is a binary relation on $\mathbb{V}$, defined by

$$u \preceq_\eta v \qquad \text{if and only if} \qquad \mathcal{O}(\eta, u, v) = 1. \tag{4.4}$$

A general choice of $\mathcal{O}$ and η does not result in a linear ordering – or even a preorder, for that matter! For a probability measure θ on $\mathcal{B}$, we say that $\mathcal{O}$ is θ-compatible if $\preceq_\eta$ is almost surely a linear ordering, when $(\eta_v)_{v \in \mathbb{V}}$ are i.i.d. samples from θ. If $\mathcal{B} = [0,1]$ and θ is the uniform measure, we can choose $\mathcal{O}(\eta, u, v) = \mathbf{1}_{\eta_u \le \eta_v}$. This function is clearly θ-compatible, and recovers the simplest ordering described earlier.

Given a finite sequence $q = ((v_1, a_1), \ldots, (v_k, a_k)) \in (\mathbb{V} \times \mathcal{A})^k$, we denote

$$F_q := F_{v_1, a_1} \circ \cdots \circ F_{v_k, a_k}.$$

Given $\nu \in \mathcal{A}^{\mathbb{V}}$, $\eta \in \mathcal{B}^{\mathbb{V}}$ such that $\preceq_\eta$ is a linear order, a vertex $v \in \mathbb{V}$ and an integer $r \ge 0$, we define

$$q(\nu, \eta, v, r) := ((v_1, a_1), \ldots, (v_m, a_m)),$$

where

$$m = |\mathbb{V}_{v,r}|, \qquad \mathbb{V}_{v,r} = \{v_1, \ldots, v_m\}, \qquad v_1 \preceq_\eta \cdots \preceq_\eta v_m, \qquad a_i = \nu_{v_i}.$$

This gives rise to a coupled dynamics on Ω, namely,

$$F_{q(\nu, \eta, v, r)} \colon \Omega^+ \to \Omega^+.$$

For any $v \in \mathbb{V}$ and $r \in \mathbb{N}$, define

$$Q^{+,v,r} \colon \Omega \to \Omega^+_{\mathbb{V}_{v,r}}$$

to be the natural projection, i.e.

$$Q^{+,v,r}(\omega)_u := \begin{cases} \omega_u & \text{if } u \in \mathbb{V}_{v,r} \\ + & \text{otherwise} \end{cases}.$$

This allows us to define the random function from Ω to Ω^+:

$$\tilde{f}^{+,v,r}_n := F_{q(A_n,B_n,v,r)} \circ Q^{+,v,r}.$$

The function $\tilde{f}^{+,v,r}_n$ describes the nth round of updates in the coupled dynamics. We also define

$$f^{+,v,r}_n := \tilde{f}^{+,v,r}_1 \circ \cdots \circ \tilde{f}^{+,v,r}_n. \tag{4.5}$$

There are two important things to note about $f^{+,v,r}_n$. First, the order of composition is reverse from the usual convention. This will prove essential to our construction. For further discussion, see Section 4.3. Second, for any v, r and n, $f^{+,v,r}_n$ is a deterministic function of Y.

Having chosen our definitions carefully, we easily obtain the following.

Lemma 4.1. *For any $r \geq 0$ and $v \in \mathbb{V}$,*

$$f^{+,v,r}_n(+) \xrightarrow{(d)} \rho^+_{\mathbb{V}_{v,r}} \qquad \text{as } n \to \infty.$$

Proof. The consistency relations of the upwards specification ρ^+ imply that $\rho^+_{\mathbb{V}_{v,r}}$ is stationary with respect to $F_{u,A}$ for any $u \in \mathbb{V}_{v,r}$, where A is sampled from π. Since $Q^{+,v,r}$ is the identity map on $\Omega^+_{\mathbb{V}_{v,r}}$, $\rho^+_{\mathbb{V}_{v,r}}$ is also stationary with respect to $\tilde{f}^{+,v,r}_n$.

We now consider the countable-state Markov chain

$$g^{+,v,r}_n := \tilde{f}^{+,v,r}_n \circ \cdots \circ \tilde{f}^{+,v,r}_1,$$

given by composing the $\tilde{f}^{+,v,r}_n$ in the usual order. The chain is aperiodic (as $\tilde{f}^{+,v,r}_i(+)$ equals $+$ with positive probability) and irreducible (since ρ^+ is irreducible). Since there exists a stationary distribution, the Markov chain is ergodic on the states with positive $\rho^+_{\mathbb{V}_{v,r}}$ measure, and thus $g^{+,v,r}_n(+)$ converges in distribution to $\rho^+_{\mathbb{V}_{v,r}}$. Since $g^{+,v,r}_n(+)$ and $f^{+,v,r}_n(+)$ have the same distribution, we are done. $\qquad\square$

If we assume that S has a minimal element $-$ and ρ is an upwards-downwards specification, we may extend F to $\Omega^+ \cup \Omega^-$, define a $-$ projection $Q^{-,v,r}$, and thus create $\tilde{f}^{-,v,r}_n$ and $f^{-,v,r}_n$ in order to define the $-$ dynamics. Lemma 4.1 also applies to this dynamics.

4.3 Monotonicity and existence of factors

We can think of $f^{+,v,r}_n$ as a (random) function from Ω to $\Omega^+_{\mathbb{V}_{v,r}}$ which inherits several monotonicity properties from the upward specification ρ^+ (which must be stationary with respect to it).

Lemma 4.2. *The function $f^{+,v,r}_n(\omega)$ preserves the order in ω and is decreasing in r. That is, for any $n, r \geq 0$ and $v \in \mathbb{V}$,*

$$f^{+,v,r}_n(\omega) \preceq f^{+,v,r}_n(\omega') \qquad \text{for any } \omega, \omega' \in \Omega \text{ such that } \omega \leq \omega',$$

and

$$f_n^{+,v,r+1}(\omega) \preceq f_n^{+,v,r}(\omega) \qquad \text{for any } \omega \in \Omega.$$

In particular,

$$f_{n+1}^{+,v,r}(+) \leq f_n^{+,v,r}(+).$$

Proof. Thanks to (4.2), we know that $F_{u,a}(\omega) \leq F_{u,a}(\omega')$ for all $a \in \mathcal{A}$, $u \in \mathbb{V}_{v,r}$, and $\omega, \omega' \in \Omega^+$ such that $\omega \leq \omega'$. Furthermore, $Q^{+,v,r}$ also maintains the order on Ω for fixed v and r, and we conclude that $f_n^{+,v,r}$, a composition of monotone functions, is also monotonic.

We now turn to prove the second monotonicity statement. Since $Q^{+,v,r}$ is decreasing in r, it suffices to show that, almost surely,

$$F_{q(A_n,B_n,v,r+1)}(\omega) \leq F_{q(A_n,B_n,v,r)}(\omega) \qquad \text{for all } \omega \in \Omega_{\mathbb{V}_{v,r}}^+.$$

Fix $\nu \in \mathcal{A}^{\mathbb{V}}$ and $\eta \in \mathcal{B}^{\mathbb{V}}$ such that $\preceq_\eta$ is a total order on $\mathbb{V}$. Write $\mathbb{V}_{v,r} = \{v_1, \ldots, v_{B(r)}\}$ and $\mathbb{V}_{v,r+1} = \{u_1, \ldots, u_{B(r+1)}\}$, where $v_1 \preceq_\eta \cdots \preceq_\eta v_{B(r)}$ and $u_1 \preceq_\eta \cdots \preceq_\eta u_{B(r+1)}$. It is clear that $v_i = u_{i'}$ and $v_j = u_{j'}$ then $i \leq j$ if and only if $i' \leq j'$. Therefore, for some functions $G_i \colon \Omega^+ \to \Omega^+$ such that $G_i(\omega)_{\mathbb{V}_{v,r}} = \omega_{\mathbb{V}_{v,r}}$ for all $\omega \in \Omega^+$, we have

$$F_{q(\nu,\eta,v,r)} = \qquad F_{v_1,\nu_{v_1}} \qquad \circ F_{v_2,\nu_{v_2}} \qquad \circ \cdots \circ F_{v_{B(r)},\nu_{v_{B(r)}}}$$

$$F_{q(\nu,\eta,v,r+1)} = G_0 \circ F_{v_1,\nu_{v_1}} \circ G_1 \circ F_{v_2,\nu_{v_2}} \circ G_2 \circ \cdots \circ F_{v_{B(r)},\nu_{v_{B(r)}}} \circ G_{B(r)}.$$

Observe that, for any $\omega \in \Omega_{\mathbb{V}_{v,r}}^+$, we have $G_i(\omega) \leq \omega$ and $F_{w,\eta_w}(\omega) \in \Omega_{\mathbb{V}_{v,r}}^+$ for any $w \in \mathbb{V}_{v,r}$. Thus, given $\omega \in \Omega_{\mathbb{V}_{v,r}}^+$, removing every composition with G_i from the above sequence defining $F_{q(\nu,\eta,v,r+1)}(\omega)$ only increases the output, showing that $F_{q(\nu,\eta,v,r+1)}(\omega) \leq F_{q(\nu,\eta,v,r)}(\omega)$, as desired.

The final inequality now follows, since $+$ is the maximal element of Ω:

$$f_{n+1}^{+,v,r}(+) = f_n^{+,v,r}(\tilde{f}_{n+1}^{+,v,r}(+)) \leq f_n^{+,v,r}(+). \qquad \square$$

We deduce a simple but crucial consequence of the above lemma and the definition of $f_n^{+,v,r}$:

Corollary 4.3. *The random field*

$$\sigma^{+,v,r} := \lim_{n \to \infty} f_n^{+,v,r}(+)$$

is defined almost-surely and has the distribution $\rho_{\mathbb{V}_{v,r}}^+$. Furthermore, if μ^+ is a probability measure,

$$\sigma^+ := \lim_{r \to \infty} \sigma^{+,v,r} = \lim_{n,r \to \infty} f_n^{+,v,r}(+)$$

is also a well-defined random field, independent of v, with distribution μ^+.

Proof. The sequence $\{f_n^{+,v,r}(+)\}_n$ is decreasing, and must have an almost-sure limit, taking value in $\overline{S}^{\mathbb{V}}$ for some possibly larger $\overline{S} \supset S$ (if S is not finite, and thus not compact in the discrete topology, we cannot be sure that the limit is supported on Ω *a priori*). However, by Lemma 4.1, the distribution of the limiting random variable is known to be $\rho_{\mathbb{V}_{v,r}}^+$, which is supported on Ω^+, giving the desired result.

By Lemma 4.2, $\sigma^{+,v,r}$ is decreasing in r, meaning it, too, has an almost-sure limit in a possibly larger space as r goes to infinity. Thanks to the monotonicity of the specifications, the resulting limit is independent of v. Since μ^+ is defined by exhaustion, σ^+ must have distribution μ^+, meaning σ^+ is almost-surely supported in Ω. To complete the proof, we note that the array $\{f_n^{+,v,r}(+)\}_{n,r}$ is monotonically decreasing pointwise in both n and r, and thus extracting any diagonal sequence maintains the almost-sure convergence properties above. $\qquad \square$

Remark 4.4 (Coupling-from-the-past)**.** Let us momentarily reconsider the standard, "forward" dynamics given by $g_n^{+,v,r} := \tilde{f}_n^{+,v,r} \circ \cdots \circ \tilde{f}_1^{+,v,r}$. In distribution, the random variables $f_n^{+,v,r}(+)$ and $g_n^{+,v,r}(+)$ are identical for any fixed n. However, $g_n^{+,v,r}(+)$ cannot be monotonic in n - if our configuration $+$ evolved to be some different configuration ω at time n, there is no reason to believe that the next step in the dynamic is smaller than ω! In fact, $g_n^{+,v,r}(+)$ does not have an almost-sure limit, as it continues changing after every application of the g_n.

On the other hand, $f_n^{+,v,r}(\omega)$ is defined from the past, and can be thought of as evaluating the forward dynamics at time 0, with 'initial' conditions of ω at time $-n$. Since $f_n^{+,v,r}$ is a function from Ω to $\Omega^{+,v,r}$, we can sample $f_{n+1}^{+,v,r}$ given $f_n^{+,v,r}$ by taking ω to the (random) configuration $\tilde{f}_{n+1}^{+,v,r}(\omega)$, and then mapping it to the (deterministic, given the conditioning) configuration assigned to it by $f_n^{+,v,r}$. This concatenation construction is crucially important for the final inequality of Lemma 4.2, and is the conceptual justification for the existence of almost-sure limits in coupling-from-the-past.

Proof of Theorem 2.1, item 1. The existence of an i.i.d. coding follows by explicit construction: for any $v \in \mathbb{V}$, $\sigma_v^+ = \lim_{n,r\to\infty} f_n^{+,v,r}(+)_v$. Since μ^+ is a probability measure by assumption, Corollary 4.3 shows that σ^+ is distributed as μ^+; the Γ-invariance of the specifications implies that σ^+ is a deterministic and Γ-equivariant function of Y. $\qquad\square$

4.4 Finitary factors via quantitative bounds on coding radius

For this section, we assume that S is a finite spin space and that ρ is a monotone Γ-invariant irreducible upwards-downwards specification. In this case, S has both a maximal and minimal element, and both μ^+ and μ^- are probability measures.

As mentioned above, the construction in Section 4.2 extends to upwards-downwards specifications. Observe that $f_n^{-,v,r}$ enjoys similar properties as $f_n^{+,v,r}$, with the notable difference that $f_n^{-,v,r}(\omega)$ is *increasing* in r (it still preserves the order in ω). As in Corollary 4.3, $\sigma^{-,v,r}$ and σ^- are defined almost surely and are distributed as $\rho_{\overline{\mathbb{V}}_{v,r}}^-$ and μ^-, respectively.

We stress that the $+$ and $-$ dynamics are coupled as they are both defined through the same process Y. In particular, almost surely,

$$f_n^{-,v,r}(\omega) \le f_n^{+,v,r}(\omega') \quad \text{for any } v \in \mathbb{V}, \ r \in \mathbb{N} \text{ and } \omega, \omega' \in \Omega \text{ such that } \omega \le \omega'.$$

This implies that, almost surely,

$$\sigma^{-,v,r} \le \sigma^- \le \sigma^+ \le \sigma^{+,v,r} \quad \text{for any } v \in \mathbb{V}.$$

For finite spin spaces, we have the following lemma which relates the probability of disagreement under monotone couplings to total-variation bounds:

Lemma 4.5. *Let X and Y be random variables taking value in a totally ordered, finite spin space S. If $\mathbb{P}[X \le Y] = 1$, then*

$$\mathbb{P}[X \ne Y] \le (|S| - 1) \cdot \|X - Y\|_{TV}.$$

Proof. Identifying S with the set $\{0, \ldots, |S| - 1\}$, we see that, by Markov's inequality,

$$\mathbb{P}[X \ne Y] = \mathbb{P}[Y - X \ge 1] \le \mathbb{E}[Y - X] = \mathbb{E}_{\text{op}}[Y - X] \le (|S| - 1)\mathbb{P}_{\text{op}}[X \ne Y],$$

where $\mathbb{P}_{\text{op}}$ is an optimal coupling between X and Y. In an optimal coupling, the probability of two variables not matching is exactly the total variation distance, as required. $\qquad\square$

Proof of Theorem 2.1, item 2. We let $(\mathcal{A}, \pi)$ and F be as above. We let $(\mathcal{B}, \theta)$ be the Lebesgue measure space on $[0, 1]$ and set $\mathcal{O}(\eta, u, v) := \mathbf{1}_{\eta_u \leq \eta_v}$. With this choice, it is clear that $\sigma^{+,v,r}$ is measurable with respect to $Y_{\mathbb{V}_{v,r}}$. Recall that σ^+ describes a coding from Y to μ^+. Our goal is then to bound its coding radius R_v.

Define the random variable

$$\tilde{R}_v := \min\left\{ r \geq 0 : \sigma_v^{+,v,r} = \sigma_v^{-,v,r} \right\},$$

where we set $\tilde{R} = \infty$ if the two spins do not agree for any r. Thus, $\sigma_v^+ = \sigma_v^{+,v,r}$ for any $r \geq \tilde{R}_v$. Since $\sigma^{+,v,r}$ is independent of $(Y_u)_{u \notin \mathbb{V}_{v,r}}$ and $\tilde{R}_v$ is a stopping time with respect to the filtration of $(Y_{\mathbb{V}_{v,r}})_r$, it is clear that $R_v \leq \tilde{R}_v$.

By Lemma 4.5,

$$\mathbb{P}[\tilde{R}_v > r] = \mathbb{P}[\sigma_v^{+,v,r} \neq \sigma_v^{-,v,r}] \leq (|S| - 1) \cdot \left\| \rho_{\mathbb{V}_{v,r}}^+ (\sigma_v \in \cdot) - \rho_{\mathbb{V}_{v,r}}^- (\sigma_v \in \cdot) \right\|_{TV},$$

as required. In particular, if $\mu^+ = \mu^-$, it is clear the total variation distance must vanish as $r \to \infty$, so that R_v is almost-surely finite. $\qquad\square$

5 Weak spatial mixing implies exponential mixing in time

In this section, we prove bounds on the mixing-time of the dynamics considered in Section 4. Throughout this section, we will assume that S is finite and that we are given a monotone Γ-invariant irreducible upwards-downwards specification ρ. In particular, we use the coupled $+$ and $-$ dynamics as in Section 4.4.

A monotone Γ-invariant (upwards-downwards) specification ρ is said to satisfy *weak spatial mixing* with rate $c > 0$ if (2.3) holds for some C. Martinelli and Olivieri [22] show that when $G = \mathbb{Z}^d$ and ρ is a monotone Γ-invariant specification satisfying a *finite-range assumption* and a *finite-energy assumption*, weak spatial mixing implies that the mixing-time of the single-site Glauber dynamics (as considered in Section 4) has exponential tails (their setting is a continuous-time dynamics on $\mathbb{Z}^d$, but the proof easily adapts to our discrete-time dynamics on $\mathbb{Z}^d$). In our notation, this means that the total-variation distance between $\lim_{r \to \infty} f_n^{+,v,r}(+)_v$ and $\lim_{r \to \infty} f_n^{-,v,r}(-)_v$ is exponentially small in n. Using the order constructed in Section 4.2, we can see that the finite range assumption implies that $f_n^{+,v,n}(+)_v$ and $f_n^{-,v,n}(-)_v$ are exponentially close, as both $f_n^{+,v,r}(+)$ and $f_n^{-,v,r}(-)$ depend only on $\{A_{u,i}, B_{u,i}\}_{u \in \mathbb{V}_{v,Cr}, i \leq n}$ for some constant $C > 0$ depending on the range of the specification.

We extend the result of Martinelli–Olivieri in a number of directions. First, we allow an arbitrary quasi-transitive graph G of sub-exponential growth (though we require a slightly stronger quantitative bound on the rate of growth for the full conclusion). Second, we drop the finite-range and finite-energy assumptions, requiring only an irreducibility assumption. Third, we work with a monotone upwards-downwards specification, instead of a (usual) monotone specification. Lastly, we keep track not only of the amount of time required until mixing, but also the amount of *space* (in the graph G) required.

Theorem 5.1. *Let G be an infinite graph and Γ be a group acting quasi-transitively on $\mathbb{V}$ by automorphisms of G. Let S be a totally ordered finite spin space and ρ be a monotone Γ-invariant irreducible upwards-downwards specification that satisfies weak spatial mixing with rate $c > 0$. If G has sub-exponential growth, i.e., $B(r) = e^{o(r)}$ as $r \to \infty$, then*

$$\max_{v \in \mathbb{V}} \left\| f_n^{+,v,n}(+)_v - f_n^{-,v,n}(-)_v \right\|_{TV} \leq e^{-n^{1-o(1)}} \qquad \text{as } n \to \infty.$$

Moreover, if there exists $\beta < c \log 2$ such that $B(r) \leq e^{\frac{\beta r}{\log r}}$ for large r, then there exists

$c' > 0$ *such that*

$$\max_{v \in \mathbb{V}} \left\| f_n^{+,v,n}(+)_v - f_n^{-,v,n}(-)_v \right\|_{TV} \leq e^{-c'n} \qquad \text{for large } n.$$

For the proof, we require the following calculus lemma, whose proof we postpone to the end of the section.

Lemma 5.2. *Let $\psi \colon \mathbb{N} \to (0, \infty)$ be monotone decreasing to zero and let $b \colon \mathbb{N} \to [0, \infty)$ be sub-linear. Suppose that, for some $C, c > 0$,*

$$\psi(2n) \leq e^{b(s)}\psi(n)^2 + Ce^{-cs} \qquad \text{for all } n \geq s \geq 1.$$

Then $\psi(n)$ decays faster than any stretched-exponential, i.e., $\psi(n) \leq \exp(-n^{1-o(1)})$. In addition, if $b(n) \leq \frac{\beta n}{\log n}$ for some $\beta < c\log 2$ and all sufficiently large n, then $\psi(n)$ decays exponentially fast.

Proof of Theorem 5.1. For $n, r \geq 0$, define

$$\phi(n, r) := \max_{v \in \mathbb{V}} \mathbb{P}[f_n^{+,v,r}(+)_v \neq f_n^{-,v,r}(-)_v].$$

Throughout the proof, we repeatedly use Lemma 4.2 without explicit mention; in particular, we use the fact that $\phi(n, r)$ is decreasing in both n and r, as easily follows.

It suffices to show that $\psi(n) := \phi(n, n)$ has the desired decay rate. Note that the weak spatial mixing assumption implies that $\mu^+ = \mu^-$, which, together with irreducibility, implies that $\psi(n) \to 0$ as $n \to \infty$. Thus, the theorem will follow from Lemma 5.2 once we establish the following inequality:

$$\psi(2n) \leq \phi(2n, n + s) \leq 2B(s)\psi(n)^2 + 3C|S|e^{-cs} \qquad \text{for all } n \geq s \geq 0.$$

In fact, we show the slightly stronger inequality:

$$\phi(n+m, r+s) \leq 2B(s)\phi(n,r)\phi(m,r+s) + 2C|S|e^{-cs} + C|S|e^{-cr} \text{ for all } n,m,s \geq 0 \text{ and } r \geq s. \tag{5.1}$$

The earlier inequality follows by setting $n = m = r$.

Recall that $f_n^{\pm,v,r}$ is measurable with respect to the i.i.d. process Y. Let $\xi_v^+ \sim \rho_{\mathbb{V}_{v,r}}^+$ and $\xi_v^- \sim \rho_{\mathbb{V}_{v,r}}^-$ be a random variables, independent of Y, satisfying that $\xi_v^+ \geq \xi_v^-$ almost surely. Note that $f_n^{+,v,r}(\xi_v^+) \sim \rho_{\mathbb{V}_{v,r}}^+$ and $f_n^{-,v,r}(\xi_v^-) \sim \rho_{\mathbb{V}_{v,r}}^-$. By comparing $f_n^{+,v,r}(+)_v$ to $f_n^{+,v,r}(\xi_v^+)_v$ and $f_n^{-,v,r}(-)_v$ to $f_n^{-,v,r}(\xi_v^-)_v$ and using Lemma 4.5, we see that

$$\phi(n, r) \leq \phi^+(n, r) + \phi^-(n, r) + |S| \cdot \max_{v \in \mathbb{V}} \left\| \rho_{\mathbb{V}_{v,r}}^+(\sigma_v \in \cdot) - \rho_{\mathbb{V}_{v,r}}^-(\sigma_v \in \cdot) \right\|_{TV}, \tag{5.2}$$

where

$$\phi^+(n, r) := \max_{v \in \mathbb{V}} \mathbb{P}[f_n^{+,v,r}(+)_v \neq f_n^{+,v,r}(\xi_v^+)_v],$$
$$\phi^-(n, r) := \max_{v \in \mathbb{V}} \mathbb{P}[f_n^{-,v,r}(-)_v \neq f_n^{-,v,r}(\xi_v^-)_v].$$

Next, we now show that

$$\phi^{\pm}(n+m, r+s) \leq B(s)\phi(n,r)\phi(m,r+s) + |S| \cdot \max_{v \in \mathbb{V}} \left\| \rho_{\mathbb{V}_{v,s}}^+(\sigma_v \in \cdot) - \rho_{\mathbb{V}_{v,s}}^-(\sigma_v \in \cdot) \right\|_{TV}. \tag{5.3}$$

This statement will give (5.1) thanks to (2.3) and (5.2). We show (5.3) only for ϕ^+ as the proof for ϕ^- is similar. Recall from (4.5) that $f_{n+m}^{+,v,r+s}$ is the composition of $n + m$ independent copies of $\tilde{f}_1^{+,v,r+s}$. Letting $h_n^{+,v,r+s}$ be identical in distribution to $f_n^{+,v,r+s}$

and independent of Y and ξ_v^+, we see that $f_m^{+,v,r+s} \circ h_n^{+,v,r+s}$ has the same distribution as $f_{n+m}^{+,v,r+s}$. Therefore,

$$\phi^+(n+m,r+s) = \max_{v \in \mathbb{V}} \mathbb{P}\left[f_m^{+,v,r+s}(h_n^{+,v,r+s}(+))_v \neq f_m^{+,v,r+s}(h_n^{+,v,r+s}(\xi_v^+))_v \right],$$

Now, letting $E_{v,u}$ denote the event that $h_n^{+,v,r+s}(+)_u = h_n^{+,v,r+s}(\xi_v^+)_u$ and letting $E_v := \bigcap_{u \in \mathbb{V}_{v,s}} E_{v,u}$,

$$\phi^\pm(n+m,r+s) \leq \max_{v \in \mathbb{V}} \mathbb{P}(E_v^c) \cdot \mathbb{P}\left[f_m^{+,v,r+s}(+)_v \neq f_m^{+,v,r+s}(-)_v \right]$$
$$+ \max_{v \in \mathbb{V}} \mathbb{P}\left[f_m^{+,v,s}(h_n^{+,v,r+s}(\xi_v^+))_v \neq f_m^{-,v,s}(h_n^{+,v,r+s}(\xi_v^+))_v \right].$$

This follows by two different types of monotonocity: if the configurations did not couple by time n, we may assume they take on their maximal difference. If they do, we may assume the boundary conditions outside $\mathbb{V}_{v,s}$ take on the worst possible state.

For the first term, since $\mathbb{V}_{u,r} \subset \mathbb{V}_{v,r+s}$ for $u \in \mathbb{V}_{v,s}$, we have for any $u \in \mathbb{V}_{v,s}$ that

$$\mathbb{P}(E_{v,u}^c) \leq \mathbb{P}\left[h_n^{+,v,r+s}(+)_u \neq h_n^{-,v,r+s}(-)_u \right] \leq \mathbb{P}\left[h_n^{+,u,r}(+)_u \neq h_n^{-,u,r}(-)_u \right] \leq \phi(n,r),$$

so that

$$\mathbb{P}(E_v^c) \leq \sum_{u \in \mathbb{V}_{v,s}} \mathbb{P}(E_{v,u}^c) \leq B(s)\phi(n,r).$$

For the second term, we note that $s \leq r$ implies that $\rho_{\mathbb{V}_{v,s}}^- \leq_{\mathrm{st}} \rho_{\mathbb{V}_{v,r}}^- \leq_{\mathrm{st}} \rho_{\mathbb{V}_{v,r}}^+ \leq_{\mathrm{st}} \rho_{\mathbb{V}_{v,s}}^+$. Since $\xi_v^+ \sim \rho_{\mathbb{V}_{v,r}}^+$, we see that

$$f_m^{+,v,s}(h_n^{+,v,r+s}(\xi_v^+)) \leq f_m^{+,v,s}(h_n^{+,v,s}(\xi_v^+)) \leq_{\mathrm{st}} \rho_{\mathbb{V}_{v,s}}^+,$$
$$f_m^{-,v,s}(h_n^{+,v,r+s}(\xi_v^+)) \geq f_m^{-,v,s}(h_n^{-,v,s}(\xi_v^+)) \geq_{\mathrm{st}} \rho_{\mathbb{V}_{v,s}}^-.$$

Thus, Lemma 4.5 gives that

$$\mathbb{P}\left[f_m^{+,v,s}(h_n^{+,v,r+s}(\xi_v^+))_v \neq f_m^{-,v,s}(h_n^{+,v,r+s}(\xi_v^+))_v \right] \leq |S| \cdot \left\| \rho_{\mathbb{V}_{v,s}}^+(\sigma_v \in \cdot) - \rho_{\mathbb{V}_{v,s}}^-(\sigma_v \in \cdot) \right\|_{TV}.$$

Putting this together yields (5.3). $\qquad\square$

Proof of Lemma 5.2. Denote $a_n := -\log \psi(n)$ and observe that the main assumption implies that

$$a_{2n} \geq cs - \log(1+C) \qquad \text{for any } n \geq s \geq 1 \text{ such that} \qquad cs + b(s) \leq 2a_n. \qquad (5.4)$$

The restriction that $s \leq n$ is a nuisance; to rid ourselves of it, we note that either $a_n \geq cn/2$ for infinitely many n, or, whenever n is large, any solution to $cs + b(s) \leq 2a_n$ satisfies $s \leq n$. In the former case, it is not difficult to check that $a_n = \Omega(n)$, so that $\psi(n)$ decays exponentially. We may therefore assume that this is not the case.

We begin by showing that $\psi(n)$ decays faster than any stretched-exponential, that is, that a_n grows faster than n^δ for any $0 < \delta < 1$. Since $b(s)$ is sub-linear by assumption, for any fixed $\epsilon > 0$, we have $cs + b(s) \leq 2x$ for all $s \leq (2/c - \epsilon)x$ and large x. Since $a_n \to \infty$ as $n \to \infty$, it follows from (5.4) that $a_{2n} \geq c(2/c - \epsilon)a_n - \log(1+C) \geq (2 - \epsilon c - \epsilon)a_n$ for large n. We conclude that

$$\lim_{n \to \infty} a_{2^n} 2^{-\delta n} = \infty \qquad \text{for any } 0 < \delta < 1.$$

It is then straightforward to show that $a_n = n^{1-o(1)}$, establishing the first part of the lemma.

Towards showing an exponential bound under the additional assumption on the growth rate of $b(n)$, let $\beta < c \log 2$ be such that $b(n) \le \frac{\beta n}{\log n}$. Let $\beta/c < \alpha < \log 2$. We claim that

$$s \le \tfrac{2x}{c}(1 - \tfrac{\alpha}{\log x}) \implies cs + b(s) \le 2x \qquad \text{for large } x.$$

Indeed, since $s \mapsto cs + b(s)$ is increasing, this follows from

$$2x(1 - \tfrac{a}{\log x}) + \frac{\frac{2x\beta}{c}(1 - \frac{\alpha}{\log x})}{\log\left[\frac{2x}{c}(1 - \frac{\alpha}{\log x})\right]} = 2x(1 - \tfrac{a}{\log x}) \cdot \left(1 + \frac{\beta}{c \log\left[\frac{2x}{c}(1 - \frac{\alpha}{\log x})\right]}\right) \le 2x.$$

Thus, by (5.4), there exists N such that

$$a_{2n} \ge 2a_n(1 - \tfrac{\alpha}{\log a_n}) \qquad \text{for all } n \ge N.$$

Let $0 < \gamma < 1 - \frac{\alpha}{\log 2}$ and let $A > 0$ be small enough so that $a_{2N} \ge AN/(\log N)^{1-\gamma}$. We prove by induction that

$$a_n \ge \frac{An}{(\log n)^{1-\gamma}} \qquad \text{for all } n \in \{N, 2N, 4N, 8N, \dots\}.$$

Since $x \mapsto 2x(1 - \frac{\alpha}{\log x})$ is increasing,

$$a_{2n} \ge 2a_n\left(1 - \frac{\alpha}{\log a_n}\right) \ge \frac{2An}{(\log n)^{1-\gamma}} \cdot \left(1 - \frac{\alpha}{\log An - \log(\log n)^{1-\gamma}}\right)$$
$$= \frac{2An}{(\log 2n)^{1-\gamma}} \cdot \frac{(\log 2n)^{1-\gamma}}{(\log n)^{1-\gamma}} \cdot \left(1 - \frac{\alpha}{\log An - \log(\log n)^{1-\gamma}}\right).$$

The induction step now follows using that $(1 - \gamma) \log 2 > \alpha$ and that

$$\frac{(\log 2n)^{1-\gamma}}{(\log n)^{1-\gamma}} = \left(1 + \frac{\log 2}{\log n}\right)^{1-\gamma} \ge 1 + \frac{(1-\gamma)\log 2}{\log n}.$$

We conclude that

$$\psi(n) \le \exp\left(-\frac{An}{(\log n)^{1-\gamma}}\right) \qquad \text{for all } n \in \{N, 2N, 4N, 8N, \dots\}. \tag{5.5}$$

Towards upgrading this bound to the desired exponential bound, define

$$\ell(n) := e^{\frac{n}{(\log n)^{1-\gamma/2}}} \psi(n) + e^{-cn/4 + \sqrt{n}}.$$

Using the recursion assumption with $s = n$, we see that for large n,

$$\ell(2n) = e^{\frac{2n}{(\log 2n)^{1-\gamma/2}}} \psi(2n) + e^{-cn/2 + \sqrt{2n}}$$
$$\le e^{\frac{2n}{(\log 2n)^{1-\gamma/2}} + \frac{cn}{\log n}} \psi(n)^2 + e^{\frac{2n}{(\log 2n)^{1-\gamma/2}} - cn} + e^{-cn/2 + \sqrt{2n}}$$
$$\le e^{\frac{2n}{(\log n)^{1-\gamma/2}}} \psi(n)^2 + e^{-cn/2 + \sqrt{2n} + 1}$$
$$\le \ell(n)^2.$$

Since (5.5) implies that $\liminf_{n\to\infty} \ell(n) = 0$, there exists $M \ge 1$ such that $\ell(M) \le 1/e$ and $\ell(2n) \le \ell(n)^2$ for all $n \ge M$. Then $\ell(2^n M) \le e^{-2^n}$ for all $n \ge 0$. Since $\max_{n \le m \le 2n} \ell(m) \le e^{n/(\log n)^{1-\gamma/2}} \ell(n)$, it easily follows that $\ell(n)$ decays exponentially fast, and we conclude that $\psi(n)$ also decays exponentially fast. $\qquad\square$

6 Space-time finitary codings

The goal of this section is to prove Theorem 2.2. As such, the graph G will be an infinite quasi-transitive graph satisfying (2.2), S will be finite (and identified with $\{0, 1, \ldots, |S| - 1\}$), and ρ will be a monotone Γ-invariant irreducible marginally finite upwards-downwards specification such that $\mu^+ = \mu^-$.

We will reuse the dynamics of Section 4, but will now demand that the spaces $\mathcal{A}$ and $\mathcal{B}$ are both finite. Before we construct suitable versions of $(\mathcal{A}, \pi)$, F, $(\mathcal{B}, \theta)$ and $\mathcal{O}$, let us outline the properties required to obtain a space-time finitary coding.

For any $v \in \mathbb{V}$, define

$$T_v := \min \big\{ n : f_n^{+,v,n}(+)_v = f_n^{-,v,n}(-)_v \big\}.$$

The assumption $\mu^+ = \mu^-$ implies that $\sigma^+ = \sigma^-$ (as was shown in the proof of Theorem 2.1, item 2). Thus, from Corollary 4.3, we conclude that T_v is almost surely finite, and, in particular,

$$\sigma_v^+ = \sigma_v^- = f_{T_v}^{+,v,T_v}(+) = f_{T_v}^{-,v,T_v}(-), \qquad v \in \mathbb{Z}^d.$$

With a finite-valued construction of $(\mathcal{A}, \pi)$ and $(\mathcal{B}, \theta)$, this does not allow us to conclude that the coding φ is finitary (let alone space-time finitary). Indeed, T_v does not bound the coding radius as it is not necessarily a stopping time with respect to the filtration $(Y_{\mathbb{V}_{v,r}})$. This is because the ordering $\preceq_{B_n}$ restricted to $\mathbb{V}_{v,r}$ may depend on $\{B_{u,n}\}_{u \notin \mathbb{V}_{v,r}}$.

To deal with this issue, for $\eta \in \mathcal{B}^{\mathbb{V}}$, we define

$$R_{u,v}(\eta) := \min \big\{ r \geq 0 : \mathbf{1}_{\{u \preceq_\eta v\}} = \mathbf{1}_{\{u \preceq_{\eta'} v\}} \text{ for any } \eta' \text{ satisfying } \eta'_{\mathbb{V}_{u,r} \cup \mathbb{V}_{v,r}} = \eta_{\mathbb{V}_{u,r} \cup \mathbb{V}_{v,r}} \big\},$$

where we again set the variable to ∞ if the set is empty. In words $R_{u,v}(\eta)$ is the minimal radius around u and v needed to determine the relative $\preceq_\eta$-order between u and v. We now set

$$T_v^* := 2 \min \big\{ n : f_n^{+,v,n}(+)_v = f_n^{-,v,n}(-)_v \text{ and } R_{u,w}(B_i) \leq n \text{ for all } u, w \in \mathbb{V}_{v,n} \text{ and } 0 \leq i \leq n \big\}.$$

The factor of 2 is introduced to accommodate the fact that $\mathbb{V}_{u,n} \subset \mathbb{V}_{v,2n}$ for $u \in \mathbb{V}_{v,n}$. Then T_v^* is a bound on the coding radius at v. More importantly, it is a bound on the space-time coding radius, as one can easily see. Thus, the first part of Theorem 2.2 will follow once we can construct finite probability spaces $(\mathcal{A}, \pi)$ and $(\mathcal{B}, \theta)$ and their associated functions that imply T_v^* is almost surely finite. The second part will require quantitative bounds on the tails of T_v^*, which will require the use of the mixing time results of the previous section.

6.1 Constructing finite probability spaces

We begin by choosing $(\mathcal{A}, \pi)$ and F. We set

$$\mathcal{A} := \big\{ \rho_v^\omega[\sigma_v \leq s] \big\}_{\omega \in \Omega^+ \cup \Omega^-, \, v \in \mathbb{V}, \, s \in S}.$$

Since S is finite and ρ is marginally finite, we immediately see that $\mathcal{A}$ is finite as well. Thus, we may order the finite number of elements of $\mathcal{A}$ in increasing order, $0 \leq a_1 < \cdots < a_m = 1$, where $m := |\mathcal{A}|$. Letting $a_0 := 0$, we define π by

$$\pi(\{a_i\}) := a_i - a_{i-1}, \qquad 1 \leq i \leq m.$$

We then define F by

$$F(\omega, v, a_i) := \min\{s \in S : \rho_v^\omega[\sigma_v \leq s] \geq a_i\}, \qquad \omega \in \Omega^+ \cup \Omega^-, \, 1 \leq i \leq m.$$

It is straightforward to check that (4.1), (4.2) and (4.3) hold, where we recall that, in all those equations, we allow $\omega \in \Omega^+ \cup \Omega^-$.

We now turn to choosing $(\mathcal{B}, \theta)$ and $\mathcal{O}$. We set $\mathcal{B} := \{1, \ldots, D\}$ for some integer $D \geq 2$ and we set θ to be the uniform measure on $\mathcal{B}$. Given $\eta \in \mathcal{B}^{\mathbb{V}}$ and $v \in \mathbb{V}$, define $Z_v(\eta) = (Z_{v,n}(\eta))_{n \geq 0} \in \mathbb{N}^{\mathbb{N}}$ by

$$Z_{v,n}(\eta) := \sum_{u \in \mathbb{V}_{v,n} \setminus \mathbb{V}_{v,n-1}} \eta_u,$$

where it is understood that $Z_{v,0}(\eta) := \eta_v$. We now define

$$\mathcal{O}(\eta, u, v) := \mathbf{1}_{Z_u(\eta) \leq Z_v(\eta)},$$

where $\leq$ is used to indicate the lexicographical order on $\mathbb{N}^{\mathbb{N}}$. This creates a preorder $\preceq_\eta$ on $\mathbb{V}$. The following lemma shows that $\mathcal{O}$ is θ-compatible, i.e., that $\preceq_\eta$ is almost surely a total ordering, when $(\eta_v)_{v \in \mathbb{V}}$ are i.i.d. samples of θ.

Lemma 6.1. *Let G be an infinite quasi-transitive graph satisfying* (2.2). *Then $\mathcal{O}$ is θ-compatible and, letting $\eta = (\eta_v)_{v \in \mathbb{V}}$ be i.i.d. random variables sampled from θ,*

$$\mathbb{P}(R_{u,v}(\eta) > r) \leq D^{-r} \qquad \text{for any distinct } u, v \in \mathbb{V} \text{ and } r \geq 0.$$

Proof. Fix $u, v \in \mathbb{V}$ distinct. Consider the event

$$A_n := \bigcap_{i=1}^{n} \{Z_{u,i}(\eta) = Z_{v,i}(\eta)\}.$$

Observe that $\mathcal{O}$ is θ-compatible if and only if $\mathbb{P}(A_n) \to 0$ as $n \to \infty$. Observe also that $R_{u,v}(\eta) > n$ implies the occurrence A_n. Thus, the lemma will follow once we show that $\mathbb{P}(A_n \mid A_{n-1}) \leq \frac{1}{D}$ for all $n \geq 1$. By (2.2), there exists some $w_n \in (\mathbb{V}_{u,n} \setminus \mathbb{V}_{u,n-1}) \setminus \mathbb{V}_{v,n}$. Then

$$\mathbb{P}(A_n \mid \eta_{\mathbb{V} \setminus \{w_n\}}) \leq \max_{k \in \mathbb{Z}} \mathbb{P}(\eta_{w_n} = k) \leq \frac{1}{D}.$$

Since A_{n-1} is measurable with respect to $\eta_{\mathbb{V} \setminus \{w_n\}}$, it follows that $\mathbb{P}(A_n \mid A_{n-1}) \leq \frac{1}{D}$. $\square$

6.2 Proof of Theorem 2.2

The first item of the theorem will follow once we show that T_v^* is almost surely finite. By monotonicity, $f_n^{+,v,n}(+)_v = f_n^{-,v,n}(-)_v$ for all $n \geq T_v$. Thus, since T_v is almost surely finite, it suffices to show that $\mathbb{P}(E_n) \to 0$ as $n \to \infty$, where E_n is the event that there exists $0 \leq i \leq n$ and a pair of vertices $u, w \in \mathbb{V}_{v,n}$ for which $R_{u,w}(B_i) \geq n$. Indeed, taking D to be larger than $3\Delta^2$, where Δ is the degree of G, the union bound and Lemma 6.1 allow us to conclude that

$$\mathbb{P}[E_n] \leq (n+1)B(n)^2 D^{-n} \leq 10n\Delta^{2n}D^{-n} \leq 2^{-n} \qquad \text{for any sufficiently large } n.$$

This completes the proof of the first item of the theorem.

We now turn to the second item of the theorem. Since, on the complement of E_n, $T_v \leq n$ implies that $T_v^* \leq 2n$, we see that

$$\mathbb{P}[T_v^* > 2n] \leq \mathbb{P}[E_n] + \mathbb{P}[\{T_v^* > 2n\} \cap E_n^c] \leq \mathbb{P}[E_n] + \mathbb{P}[T_v > n].$$

Thus, for large n,

$$\mathbb{P}[T_v^* > 2n] \leq 2^{-n} + |S| \cdot \left\| f_n^{+,v,n}(+)_v - f_n^{-,v,n}(-)_v \right\|_{TV},$$

and the theorem follows from Theorem 5.1.

7 Codings from finite-valued i.i.d. processes

In this section, we prove Corollary 2.3. To this end, we require a result from [29] which allows to convert space-time finitary codings with exponential tails to fv-ffiid with stretched-exponential tails.

Suppose that $Y = (Y_{v,i})_{v \in \mathbb{Z}^d, i \geq 0}$ are i.i.d. random variables taking values in a finite set $\tilde{T}$. Let $\mathbf{0}$ be the origin of $\mathbb{Z}^d$, and $F = (F_n)_{n \geq 0}$ be a strictly increasing sequence of subsets of $\mathbb{Z}^d \times \mathbb{N}$ with $F_0 := \{(\mathbf{0}, 0)\}$, and consider the associated σ-algebras $\{\mathcal{F}_v^n\}_{v \in \mathbb{Z}^d, n \geq 0}$ defined by

$$\mathcal{F}_v^n := \sigma\big(\{Y_{v+u,i}\}_{(u,i) \in F_n}\big). \tag{7.1}$$

An $\mathbb{N}$-valued random field $\tau = (\tau_v)_{v \in \mathbb{Z}^d}$ is said to be a F-*stopping-process* for Y if, for every v, τ_v is an almost surely finite stopping time with respect to the filtration $(\mathcal{F}_v^n)_{n \geq 0}$. When we say that such a stopping-process is stationary, we shall mean that the same stopping rule is used at every vertex (rather than just meaning that its law is translation-invariant). Given a F-stopping-process, we denote by Y^τ the random field

$$Y^\tau := \big((Y_{v+u,i})_{(u,i) \in F_{\tau_v}}\big)_{v \in \mathbb{Z}^d}.$$

Note that $(Y^\tau)_v$ takes values in the finite-configuration space $\bigcup_{n \geq 0} \tilde{T}^{F_n}$. We say that F is *linear* if

$$\Delta_n := \max\big\{\max\{|u|, i\} : (u, i) \in F_n\big\} \leq \Delta n \qquad \text{for some } \Delta \geq 1 \text{ and all } n \geq 0. \tag{7.2}$$

Proposition 7.1 ([29, Proposition 10]). *Let $Y = (Y_{v,i})_{v \in \mathbb{Z}^d, i \geq 0}$ be a finite-valued i.i.d. process, let F be linear and let τ be a stationary F-stopping-process for Y. Suppose that τ_v has exponential tails and $\mathbb{E}|F_{\tau_v}| < M$ for some integer M. Then Y^τ is a translation-equivariant finitary factor of $((Y_{v,i})_{0 \leq i < M})_{v \in \mathbb{Z}^d}$ with stretched-exponential tails.*

Proof of Corollary 2.3. In the proof, all factors are translation-equivariant, i.e., Γ-factors where Γ is the group of translations.

Suppose first that $G = \mathbb{Z}^d$. In this case, μ^+ is a measure on $S^{\mathbb{Z}^d}$. By Theorem 2.2, there exists a space-time finitary coding φ from an i.i.d. process $Y = (Y_{v,i})_{v \in \mathbb{Z}^d, i \geq 0}$ to μ^+ whose space-time coding radius R^* has exponential tails. Let R_v^* denote the space-time coding radius of the vertex $v \in \mathbb{V}$. Towards applying Proposition 7.1, define $F_n := \{(u, i) : |u| \leq n, \ 0 \leq i \leq n\}$ and note that the random field $\tau = (R_v^*)_{v \in \mathbb{Z}^d}$ is a stationary F-stopping-process for Y. Since τ_v has exponential tails, it follows that $\mathbb{E}|F_{\tau_v}| < \infty$. Note that, by definition of the process Y^τ, there exists a deterministic function ψ such that $\varphi(Y)_v = \psi((Y^\tau)_v)$ for all v (ψ may, in some sense, be thought of as $\varphi(\cdot)_{\mathbf{0}}$). In particular, μ^+ is a finitary factor of Y^τ with coding radius 0. It thus suffices to show that Y^τ is fv-ffiid with stretched-exponential tails. Indeed, letting M be any integer larger than $\mathbb{E}|F_{\tau_v}|$, Proposition 7.1 yields that Y^τ is a finitary factor of $((Y_{v,i})_{0 \leq i < M})_{v \in \mathbb{Z}^d}$ with stretched-exponential tails. Since the latter process is a finite-valued i.i.d. process, this yields the required coding for μ^+, completing the proof in the case of $G = \mathbb{Z}^d$.

Suppose now that G is the line graph of $\mathbb{Z}^d$. In this case, μ^+ is a measure on $S^{E(\mathbb{Z}^d)}$ and the coding φ obtained from Theorem 2.2 is from an i.i.d. process $Y' = (Y'_{e,i})_{e \in E(\mathbb{Z}^d), i \geq 0}$. Converting the problem to one on the vertices of $\mathbb{Z}^d$ is simple: Let X' be sampled from μ^+, and define $X = (X_v)_{v \in \mathbb{Z}^d}$ and $Y = (Y_{v,i})_{v \in \mathbb{Z}^d, i \geq 0}$ by $X_v = (X'_{e_1}, \dots, X'_{e_d})$ and $Y_{v,i} = (Y'_{e_1,i}, \dots, Y'_{e_d,i})$, where $e_j = \{v, v + e_j\}$ for $1 \leq j \leq d$. Clearly, X is a space-time finitary factor of Y with space-time coding radius having exponential tails. The argument above now shows that, for some M, X is a a finitary factor $((Y_{v,i})_{0 \leq i < M})_{v \in \mathbb{Z}^d}$ with stretched-exponential tails. Going back to the prime processes, we obtain that X' is a finitary factor of $((Y'_{e,i})_{0 \leq i < M})_{e \in E(\mathbb{Z}^d)}$ with stretched-exponential tails, as required. We note that we have relied very much on the fact that Γ is the translation group (and not a larger group). $\qquad \square$

References

[1] Scot Adams, *Følner independence and the amenable Ising model*, Ergodic Theory and Dynamical Systems **12** (1992), no. 4, 633–657. MR1200333

[2] Jacob van den Berg and Jeffrey Steif, *On the existence and nonexistence of finitary codings for a class of random fields*, Annals of Probability (1999), 1501–1522. MR1733157

[3] GG Bosco, FP Machado, and Thomas Logan Ritchie, *Exponential rates of convergence in the ergodic theorem: A constructive approach*, Journal of Statistical Physics **139** (2010), no. 3, 367–374. MR2609444

[4] Robert M Burton and Michael Keane, *Density and uniqueness in percolation*, Communications in Mathematical Physics **121** (1989), no. 3, 501–505. MR0990777

[5] Emilio De Santis and Andrea Lissandrelli, *Developments in perfect simulation of Gibbs measures through a new result for the extinction of Galton-Watson-like processes*, Journal of Statistical Physics **147** (2012), no. 2, 231–251. MR2922768

[6] Emilio De Santis and Andrea Maffei, *Perfect simulation for the infinite random cluster model, Ising and Potts models at low or high temperature*, Probability Theory and Related Fields **164** (2016), no. 1-2, 109–131. MR3449387

[7] Roland L. Dobrushin, *The description of a random field by means of conditional probabilities and conditions of its regularity*, Theor. Probab. Appl. **13** (1968), 197–224. MR0231434

[8] H. Duminil-Copin, *Parafermionic observables and their applications to planar statistical physics models*, Ensaios Matematicos, vol. 25, Brazilian Mathematical Society, 2013. MR3184487

[9] Hugo Duminil-Copin, Alexander Glazman, Ron Peled, and Yinon Spinka, *Macroscopic loops in the loop $O(n)$ model at Nienhuis' critical point*, Journal of the European Mathematical Society **23** (2020), no. 1, 315–347. MR4186469

[10] Hugo Duminil-Copin, Aran Raoufi, and Vincent Tassion, *Sharp phase transition for the random-cluster and Potts models via decision trees*, Annals of Mathematics **189** (2019), no. 1, 75–99. MR3898174

[11] Sandro Gallo and Daniel Y Takahashi, *Attractive regular stochastic chains: perfect simulation and phase transition*, Ergodic Theory and Dynamical Systems **34** (2014), no. 5, 1567–1586. MR3255433

[12] A Galves, NL Garcia, and E Löcherbach, *Perfect simulation and finitary coding for multicolor systems with interactions of infinite range*, arXiv preprint arXiv:0809.3494 (2008).

[13] Antonio Galves, Eva Löcherbach, and Enza Orlandi, *Perfect simulation of infinite range Gibbs measures and coupling with their finite range approximations*, Journal of Statistical Physics **138** (2010), no. 1-3, 476–495. MR2594907

[14] Hans-Otto Georgii, Olle Häggström, and Christian Maes, *The random geometry of equilibrium phases*, Phase Transitions and Critical Phenomena, vol. 18, Elsevier, 2001, pp. 1–142. MR2014387

[15] Robert B Griffiths, *Correlations in Ising ferromagnets. I*, Journal of Mathematical Physics **8** (1967), no. 3, 478–483.

[16] Geoffrey R Grimmett, *The random-cluster model*, vol. 333, Springer Science & Business Media, 2006. MR2243761

[17] Olle Häggström, *The random-cluster model on a homogeneous tree*, Probability Theory and Related Fields **104** (1996), no. 2, 231–253. MR1373377

[18] Olle Häggström, Johan Jonasson, and Russell Lyons, *Coupling and Bernoullicity in random-cluster and Potts models*, Bernoulli **8** (2002), no. 3, 275–294. MR1913108

[19] Olle Häggström and Jeffrey Steif, *Propp–Wilson algorithms and finitary codings for high noise Markov random fields*, Combinatorics, Probability and Computing **9** (2000), no. 5, 425–439. MR1810150

[20] Michael Keane and Meir Smorodinsky, *Finitary isomorphisms of irreducible Markov shifts*, Israel Journal of Mathematics **34** (1979), no. 4, 281–286. MR0570887

[21] OE Lanford and David Ruelle, *Observables at infinity and states with short range correlations in statistical mechanics*, Communications in Mathematical Physics **13** (1969), no. 3, 194–215. MR0256687

[22] Fabio Martinelli and Enzo Olivieri, *Approach to equilibrium of Glauber dynamics in the one phase region*, Communications in Mathematical Physics **161** (1994), no. 3, 447–486. MR1269387

[23] Katalin Marton and Paul C Shields, *The positive-divergence and blowing-up properties*, Israel Journal of Mathematics **86** (1994), no. 1-3, 331–348. MR1276142

[24] Donald Ornstein, *Bernoulli shifts with the same entropy are isomorphic*, Advances in Mathematics **4** (1970), no. 3, 337–352. MR0257322

[25] Donald Ornstein and Benjamin Weiss, $\mathbf{Z}^d$-*actions and the Ising model*, unpublished.

[26] Ron Peled and Yinon Spinka, *Lectures on the spin and loop $O(n)$ models*, Sojourns in Probability Theory and Statistical Physics-I, celebrating Charles Newman's 70th birthday (2019), 246–320.

[27] Renfrey Burnard Potts, *Some generalized order-disorder transformations*, Proceedings of the Cambridge Philosophical Society, vol. 48, Cambridge Univ. Press, 1952, pp. 106–109. MR0047571

[28] James G. Propp and David B. Wilson, *Exact sampling with coupled Markov chains and applications to statistical mechanics*, Random Structures and Algorithms **9** (1996), no. 1-2, 223–252. MR1611693

[29] Yinon Spinka, *Finitary coding for the sub-critical Ising model with finite expected coding volume*, Electronic Journal of Probability **25** (2020), 1–27. MR4059186

[30] Yinon Spinka, *Finitary codings for spatial mixing Markov random fields*, The Annals of Probability **48** (2020), no. 3, 1557–1591. MR4112724

[31] Jeffrey E Steif, *Space-time Bernoullicity of the lower and upper stationary processes for attractive spin systems*, The Annals of Probability **19** (1991), no. 2, 609–635. MR1106279

[32] Fa-Yueh Wu, *The Potts model*, Reviews of Modern Physics **54** (1982), no. 1, 235. MR0641370