Broadcasting with Random Matrices

Charilaos Efthymiou ✉
Computer Science, University of Warwick, Coventry, UK

Kostas Zampetakis ✉
Computer Science, University of Warwick, Coventry, UK

Abstract

Motivated by the theory of *spin-glasses* in physics, we study the so-called *reconstruction problem* on the tree, and on the sparse random graph $\boldsymbol{G}(n, d/n)$. Both cases reduce naturally to analysing broadcasting models, where each edge has its own broadcasting matrix, and this matrix is drawn independently from a predefined distribution.

We establish the *reconstruction threshold* for the cases where the broadcasting matrices give rise to symmetric, 2-spin Gibbs distributions. This threshold seems to be a natural extension of the well-known *Kesten-Stigum bound* that manifests in the classic version of the reconstruction problem. Our results determine, as a special case, the reconstruction threshold for the prominent *Edwards–Anderson model* of spin-glasses, on the tree.

Also, we extend our analysis to the setting of the Galton-Watson random tree, and the (sparse) random graph $\boldsymbol{G}(n, d/n)$, where we establish the corresponding thresholds. Interestingly, for the Edwards–Anderson model on the random graph, we show that the *replica symmetry breaking* phase transition, established by Guerra and and Toninelli in [21], coincides with the reconstruction threshold.

Compared to classical Gibbs distributions, spin-glasses have several unique features. In that respect, their study calls for new ideas, e.g. we introduce novel estimators for the reconstruction problem. The main technical challenge in the analysis of such systems, is the presence of (too) many levels of randomness, which we manage to circumvent by utilising recently proposed tools coming from the analysis of Markov chains.

2012 ACM Subject Classification Theory of computation → Design and analysis of algorithms; Theory of computation → Randomness, geometry and discrete structures; Mathematics of computing → Discrete mathematics

Keywords and phrases spin-system, spin-glass, sparse random graph, reconstruction, phase transitions

Digital Object Identifier 10.4230/LIPIcs.ICALP.2023.55

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: `https://arxiv.org/abs/2302.11657`

Funding *Charilaos Efthymiou*: EPSRC New Investigator Award (grant no. EP/V050842/1) and Centre of Discrete Mathematics and Applications (DIMAP), The University of Warwick.
Kostas Zampetakis: EPSRC New Investigator Award (grant no. EP/V050842/1) and Centre of Discrete Mathematics and Applications (DIMAP), The University of Warwick.

Acknowledgements We are grateful to the anonymous reviewers for their thorough review of our submission, and for their insightful comments and suggestions.

1 Introduction

Motivated by the theory of *spin-glasses* in physics, we study the so-called *reconstruction problem* with respect to the related distributions, on the tree, and on the sparse random graph $\boldsymbol{G}(n, d/n)$.

50th International Colloquium on Automata, Languages, and Programming (ICALP 2023).
Editors: Kousha Etessami, Uriel Feige, and Gabriele Puppis; Article No. 55; pp. 55:1–55:14
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Spin-glasses are *disordered* magnetic materials that are studied by physicists (not necessarily the theoretical ones). It has been noted that even though they are a type of magnet, actually, "they are not very good at being magnets". Metallic spin-glasses are "unremarkable conductors", and the insulating spin-glasses are "fairly useless as practical insulators ... ", e.g. see [30].

However, the research on spin-glasses has provided tools to analyse some exciting, and extremely challenging, problems in mathematics, physics, but also *real world* ones. Through their study, we have garnered a deep understanding of the nature of complex systems. A case in point is the pioneering work of Giorgio Parisi in '70s on the so-called *Sherrington-Kirkpatrick* spin-glass, which introduces the formulation of the renowned *replica symmetry breaking* [27]. Parisi's ideas were highly influential in physics community, and later, in mathematics, and computer science. The theory of replica symmetry breaking was among the groundbreaking ideas which got Parisi the Nobel Prise in Physics in 2021.

Perhaps one of the most successful, and extensively studied, spin-glass models, is the famous *Edwards-Anderson* model (EA-model for short), introduced back in '70s by Sam Edwards and Philip Anderson in [16]. Few months after the work of Edwards and Anderson, David Sherrington and Scott Kirkpatrick, in [28], introduced their own model of spin-glasses, the well-known in computer science literature, *Sherrington-Kirkpatrick* model (or SK-model for short). As it turns out, the SK-model corresponds to the *mean field* version of the EA-model.

Given a fixed graph $G = (V, E)$, the Edwards-Anderson model with *inverse* temperature $\beta > 0$, is the *random* Gibbs distribution $\boldsymbol{\mu}$ on the configuration space $\{\pm 1\}^V$ defined as follows: let $\{\boldsymbol{J}_e : e \in E\}$ be independent identically distributed (i.i.d.) *standard Gaussians*. Then each configuration $\sigma \in \{\pm 1\}^V$ receives probability mass $\boldsymbol{\mu}(\sigma)$, defined by

$$\boldsymbol{\mu}(\sigma) \propto \exp\left(\beta \cdot \sum_{\{u,w\}\in E} \mathbf{1}\{\sigma(u) = \sigma(w)\} \cdot \boldsymbol{J}_{\{u,w\}}\right) , \tag{1}$$

where $\propto$ stands for "proportional to". We usually refer to $\{\boldsymbol{J}_e\}_{e\in E}$ as the *coupling parameters*. Let us comment here that, traditionally, the Gibbs distribution is defined by replacing the indicator $\mathbf{1}\{\sigma(u) = \sigma(w)\}$ in (1), with the product $\sigma(u)\sigma(w)$, in the physics literature. However, the two formulations are equivalent, as a simple transformation converts one to the other (see the full version). We also note that there is a simpler version of the Edwards-Anderson model, in which coupling parameters take independently ± 1 values, uniformly at random.

Apart from its mathematical elegance, and theoretical importance, the Edwards-Anderson model, and the related spin-glass distributions, arise also in applications such as neural networks (e.g. the so-called Hopfield model), protein folding, and conformational dynamics. We refer the interested reader to [30], and references therein.

In this work, we largely study the Edwards-Anderson model on trees, and the (locally tree-like) *random graph* $\boldsymbol{G}(n, d/n)$ with constant expected degree d. This is the random graph on n vertices, such that each edge appears independently with probability d/n.

Since the Edwards-Anderson model on $\boldsymbol{G}(n, d/n)$ shares essential features with random *Constraint Satisfaction Problems* (r-CSPs for short), it is not surprising that has been studied extensively in terms of phase transitions, in physics, e.g. [19, 25], mathematics, e.g. [21, 12], but also in computer science, e.g. for sampling algorithms [17, 2].

In contrast to the standard Gibbs distributions on trees, e.g. the Ising model, the Hard-core model, and the Potts model, the Edwards-Anderson model, despite being the most basic distribution for spin-glasses, has not been sufficiently studied. As a result, several fundamental questions about it still remain open. Here, we consider the tree *reconstruction problem* for the Edwards-Anderson model (and some natural extensions).

The reconstruction problem studies the effect of the configuration at a vertex, r, on that of the vertices at distance h from r, as $h \to \infty$. Specifically, we want to distinguish the region of parameters where the effect is vanishing, from that where the effect is non-vanishing. Typically, the two regions are specified in terms of a *sharp threshold*, i.e., we have an abrupt transition from one region to the other as we vary the parameters of the model. We usually call this phenomenon *reconstruction threshold*, and it has been the subject of intense study, e.g. [26, 1, 22, 7, 29, 10]. In the context of r-CSPs, the onset of reconstruction has been linked to an abrupt deterioration of the performance of algorithms (both searching and counting), e.g. see [1].

In this work, among other results, we establish precisely the reconstruction threshold for the Edwards-Anderson model on the Δ-ary tree, the Galton-Watson tree with general offspring distribution, and the random graph $\boldsymbol{G}(n, d/n)$. Furthermore, as far as the Edwards-Anderson model on $\boldsymbol{G}(n, d/n)$ is concerned, we combine our results with [21, 12], to conclude that the reconstruction threshold coincides with the so-called *Replica Symmetry Breaking* phase transition.

Interestingly, for the Δ-ary tree, we establish the reconstruction threshold, not only for the Edwards-Anderson model, but also for the general version of the Gibbs distribution $\boldsymbol{\mu}$ defined in (1). That is, the coupling parameters are i.i.d. following a *general distribution*, not necessary the standard Normal.

It turns out that the corresponding reconstruction problems on the Galton-Watson tree with Poisson(d) offspring, and on the sparse random graph $\boldsymbol{G}(n, d/n)$, are not too different from each other. Connections have been established between these two Gibbs distributions, e.g. see [4, 15, 11, 14]. We relate the two reconstruction results, i.e., for the tree and the graph, by exploiting the idea of planted-model (Teacher-Student model [31]) and the notion of mutual *contiguity* [12]. In that respect, our basic analysis involves the complete Δ-tree, and the Galton-Watson tree, while, subsequently, we extend these results to the random graph $\boldsymbol{G}(n, d/n)$.

We study the reconstruction problem on trees by means of the broadcasting models. These are abstractions of *noisy transmission* of information over the edges of the tree, i.e., the edges act as noisy channels. To our knowledge, the study of the broadcasting models, and the closely related reconstruction problem, dates back to '60s with the seminal work of Kesten and Stigum [24].

Establishing the reconstruction threshold for the Edwards-Anderson model on the Δ-ary tree, as well as the generalisation of this distribution, turns out to be a challenging problem. The difficulty of these models stems from the manifestation of local *frustration phenomena*, i.e., mixed ferromagnetic and antiferromagnetic interaction in the same neighbourhood, but also from the "many levels of randomness" we need to deal with in their analysis.

To this end, we make an extensive use of various potentials in order to simplify the analysis. To establish non-reconstruction, we employ some newly introduced techniques in the area of Markov chains and Spectral Independence [3, 9], that combine potential functions to analyse tree recursions. To establish reconstruction, we use a carefully crafted potential as an estimator for the root configuration. We call this estimator *flip-majority vote*.

1.1 Broadcasting, Reconstruction and the Kesten-Stigum bound

Consider the Δ-ary tree $T = (V, E)$, of height $h > 0$. Let r be the root of the tree T. Broadcasting on T, is a stochastic process which abstracts noisy transmission of information over the edges of the tree.

There is a finite set of spins $\mathcal{A}$, and an $\mathcal{A} \times \mathcal{A}$ stochastic matrix M, which we call the *broadcasting matrix*, or *transition matrix.* With the broadcasting we obtain a configuration $\boldsymbol{\sigma} \in \mathcal{A}^V$ by working recursively as follows: assume that the configuration at the root r is obtained according to some predefined distribution over $\mathcal{A}$. If for the non-leaf vertex u in T we have $\boldsymbol{\sigma}(u) = i$, then for each vertex w, child of u, we have $\boldsymbol{\sigma}(w) = j$ with probability $M(i, j)$, independently of the other children, i.e.,

$$\Pr[\boldsymbol{\sigma}(w) = j \mid \boldsymbol{\sigma}(u) = i] = M(i, j) \ .$$

Here we assume that $\boldsymbol{\sigma}(r)$ is distributed uniformly at random in $\mathcal{A}$.

A natural problem to study in this setting is the so-called *reconstruction problem.* Suppose that μ_h is the marginal distribution of the configuration of the vertices at distance h from the root. The reconstruction problem amounts to studying the influence of the configuration at the root of the tree to the marginal μ_h. Specifically, we want to compare the two distributions $\mu_h(\cdot \mid \boldsymbol{\sigma}(r) = i)$, and $\mu_h(\cdot \mid \boldsymbol{\sigma}(r) = j)$ for different $i, j \in \mathcal{A}$, i.e., μ_h conditional on the configuration at the root being i and j, respectively. The comparison is by means of the total variation distance, i.e.,

$$||\mu_h(\cdot \mid \boldsymbol{\sigma}(r) = i) - \mu_h(\cdot \mid \boldsymbol{\sigma}(r) = j)||_{\mathrm{TV}} \ .$$

Typically, we focus on the behaviour of the quantity above, as h grows.

▶ **Definition 1.** *We say that the distribution μ exhibits* reconstruction *if there exist spins $i, j \in \mathcal{A}$ such that*

$$\limsup_{h \to \infty} ||\mu_h(\cdot \mid \boldsymbol{\sigma}(r) = i) - \mu_h(\cdot \mid \boldsymbol{\sigma}(r) = j)||_{\mathrm{TV}} > 0 \ .$$

On the other hand, if for all $i, j \in \mathcal{A}$ the above limit is zero, then we have non-reconstruction*.*

The broadcasting process we describe above gives rise to well-known Gibbs distributions on T such as the *Ising model*, the *Potts model* etc. In terms of the Gibbs distributions on the tree, the reconstruction problem can be formulated as to whether the free-measure on the tree is *extremal*, or not. The extremality here is considered with respect to whether the Gibbs distribution can be expressed as a convex combination of two, or more measures, e.g. see [20]. It is interesting to compare the extremality condition with various spatial mixing conditions of the Gibbs distribution. Perhaps the most interesting case is to compare it with the Gibbs tree *uniqueness.* Then, it is standard to show that the extremality is a *weaker* condition than uniqueness.

The reconstruction problem has been studied since 1960s. Perhaps the most general result in the area is the so-called *Kesten-Stigum bound* [24], or KS-bound (for short). Let $\varDelta_{\mathrm{KS}} = \varDelta_{\mathrm{KS}}(M)$ be such that

$$\varDelta_{\mathrm{KS}} = \lambda_2^{-2}(M) \ , \tag{2}$$

where $\lambda_2(M)$ is the second largest, in magnitude, eigenvalue of the transition matrix M. The result of [24] implies that if $\varDelta > \varDelta_{\mathrm{KS}}$, then we have reconstruction.

In light of the above, a natural question is whether the condition $\varDelta < \varDelta_{\mathrm{KS}}$ implies that we have non-reconstruction. In general, the answer to this question is no, e.g. see [5, 29]. However, for several important distributions, including the Ising model, the KS-bound is tight, in the sense that the condition $\varDelta < \varDelta_{\mathrm{KS}}$ indeed implies non-reconstruction, see [7, 18, 22].

1.2 Broadcasting with random matrices

Here, we consider the natural problem of broadcasting on a tree, where the transition matrix is *random*. In this setting, as before, we consider the Δ-ary tree $T = (V, E)$, of height $h > 0$, rooted at r. Also, we have a finite set of spins $\mathcal{A}$. Rather than using the same matrix for every edge of the tree, each edge has its own matrix, which is an independent sample from a predefined distribution ψ.

More formally, every $\mathcal{A} \times \mathcal{A}$ stochastic matrix can be viewed as a point in the $|\mathcal{A}|^2$ Euclidean space. We endow the set of all $\mathcal{A} \times \mathcal{A}$ stochastic matrices with the σ-algebra induced by the Borel algebra. Then, ψ is a distribution over the set of these matrices.

Once we have a matrix for each edge of T, the broadcasting proceeds with the same rules as in the deterministic case. If for the non-leaf vertex u in T we have $\boldsymbol{\sigma}(u) = i$, then the vertex w, child of u, gets $\boldsymbol{\sigma}(w) = j$ with probability $\boldsymbol{M}_e(i, j)$, independently of the other children of u, i.e.,

$$\Pr[\boldsymbol{\sigma}(w) = j \mid \boldsymbol{\sigma}(u) = i] = \boldsymbol{M}_e(i, j) \ ,$$

where $e = \{u, w\}$.

The above setting gives rise to a *random* probability measure on the set of configurations $\mathcal{A}^V$ which we denote as $\boldsymbol{\mu} = \boldsymbol{\mu}_{T,\psi}$. Hence, the configuration $\boldsymbol{\sigma} \in \mathcal{A}^V$ we get from the broadcasting, consists of *two-levels of randomness*. The first level is due to the fact that the measure $\boldsymbol{\mu}$ is induced by the random instances of the broadcasting matrices $\{\boldsymbol{M}_e\}_{e \in E}$. Once these matrices have been fixed, the second level of randomness emerges from the random choices of the broadcasting process. The above formulation gives rise to well-studied Gibbs distributions, such as the Edwards–Anderson model of spin-glasses, by choosing appropriately the distribution ψ.

In this new setting, we study the reconstruction problem. Here, the definition of reconstruction differs slightly from Definition 1 above. Denote with $\boldsymbol{\mu}_h$ the marginal of $\boldsymbol{\mu}$ on the vertices at distance h from the root of the tree T. Then, the reconstruction problem is defined as follows:

▶ **Definition 2.** *For a distribution ψ on $\mathcal{A} \times \mathcal{A}$ stochastic matrices , we say that the random measure $\boldsymbol{\mu} = \boldsymbol{\mu}_{T,\psi}$ exhibits* reconstruction *if there exist spins $i, j \in \mathcal{A}$ such that*

$$\limsup_{h \to \infty} \mathbb{E}\left[||\boldsymbol{\mu}_h(\cdot \mid \boldsymbol{\sigma}(r) = i) - \boldsymbol{\mu}_h(\cdot \mid \boldsymbol{\sigma}(r) = j)||_{\mathrm{TV}}\right] > 0 \ ,$$

where the expectation is with respect to the randomness of $\boldsymbol{\mu}$.

On the other hand, if for all $i, j \in \mathcal{A}$ the above limit is zero, then we have non-reconstruction*.*

We consider the reconstruction problem in terms of the KS-bound, i.e., we examine whether it is tight, or not. Before addressing this question, we need to specify what the parameter Δ_{KS} might be in this setting.

It turns out that a natural candidate for Δ_{KS} can be defined as follows:

Let $\boldsymbol{M}$ be a matrix sampled from the distribution ψ, and define

$$\Xi = \mathbb{E}\left[\boldsymbol{M} \otimes \boldsymbol{M}\right] \ , \tag{3}$$

i.e., the matrix Ξ is the expectation of the tensor product of the matrix $\boldsymbol{M}$ with itself. Let $\mathbf{1} \in \mathbb{R}^{\mathcal{A}}$ denote the vector whose entries are all equal to one. Also, write

$$\mathcal{E} = \left\{ z \in \mathbb{R}^{\mathcal{A}} \otimes \mathbb{R}^{\mathcal{A}} : \forall y \in \mathbb{R}^{\mathcal{A}} \langle z, \mathbf{1} \otimes y \rangle = \langle z, y \otimes \mathbf{1} \rangle = 0 \right\} \ ,$$

where $\langle \cdot, \cdot \rangle$ is the standard inner product operation. Then, we define $\Delta_{\rm KS}(\psi)$ to be such that

$$\Delta_{\rm KS}(\psi) = \left(\max_{x \in \mathcal{E}: ||x||=1} \langle \Xi x, x \rangle \right)^{-1} . \tag{4}$$

The above quantity, $\Delta_{\rm KS}$, arises in the study of phases transitions in random CSPs [12]. Specifically, it signifies an *upper bound* on the density of the so-called *Replica Symmetric* phase, of symmetric Gibbs distributions. The value $\Delta_{\rm KS}$ is derived in [12] by means of a stability analysis of the so-called *free-energy* functional. Note that the above definition for $\Delta_{\rm KS}(\psi)$ applies to any set of spins $\mathcal{A}$, and any distribution ψ on $\mathcal{A} \times \mathcal{A}$ matrices.

Here, we prove that the above is indeed the analogue of KS-bound for *symmetric*, 2-spin distributions $\boldsymbol{\mu}$. That is, for any value of the parameter $\beta > 0$, and for any distribution ψ over the broadcasting matrices whose support is comprised of symmetric 2×2 matrices, we prove that the Δ-ary tree T exhibits reconstruction when $\Delta > \Delta_{\rm KS}(\psi)$, while we have non-reconstruction when $\Delta < \Delta_{\rm KS}(\psi)$.

Furthermore, we go beyond the basic case of the Δ-ary tree. Firstly, we extend our results to the cases where the underlying graph is the *Galton-Watson* random tree with general offspring distribution. Secondly, we exploit the notion of contiguity of measures to derive non-reconstruction results for the Edwards-Anderson model on the random graph $\boldsymbol{G}(n, d/n)$.

2 Results

We start the presentation of our results on the 2-spin, symmetric distributions, by considering the Δ-ary tree. Specifically, for integers $\Delta > 0$ and $h > 0$, let $T = (V, E)$ be the Δ-ary tree of height h, rooted at vertex r. We let $\mathcal{A} = \{\pm 1\}$ be the set of spins.

Assume that each edge of the tree is equipped with its own broadcasting matrix, each matrix drawn *independently* from the distribution induced by the following experiment: We have two parameters, a real number $\beta > 0$, and a distribution ϕ on the real numbers $\mathbb{R}$, i.e., we have the probability space $(\mathbb{R}, \mathcal{F}, \phi)$ where $\mathcal{F}$ is the σ-algebra induced by the Borel algebra. We generate a matrix $\boldsymbol{M}$ following the two steps below:

Step 1 Draw $\boldsymbol{J} \in \mathbb{R}$ from the distribution ϕ.

Step 2 Generate the $\mathcal{A} \times \mathcal{A}$ matrix $\boldsymbol{M}$ such that

$$\boldsymbol{M} = \frac{1}{\exp(\beta \boldsymbol{J}) + 1} \begin{bmatrix} \exp(\beta \boldsymbol{J}) & 1 \\ 1 & \exp(\beta \boldsymbol{J}) \end{bmatrix} . \tag{5}$$

Note that our broadcasting matrices are always *symmetric*.

The above broadcasting process gives rise to configurations in $\mathcal{A}^V$ following the Gibbs distribution $\boldsymbol{\mu}_{\beta,\phi}$ specified as follows: Let $\{\boldsymbol{J}_e\}_{e \in E}$ be independent, identically distributed (i.i.d.) random variables such that each one of them is distributed as in ϕ (this is the same distribution used to generate matrix $\boldsymbol{M}$). Each $\sigma \in \mathcal{A}^V$ is assigned probability mass $\boldsymbol{\mu}_{\beta,\phi}(\sigma)$ defined by

$$\boldsymbol{\mu}_{\beta,\phi}(\sigma) \propto \exp\left(\beta \textstyle\sum_{\{w,u\} \in E} \mathbf{1}\{\sigma(u) = \sigma(w)\} \cdot \boldsymbol{J}_{\{u,w\}} \right) , \tag{6}$$

where $\propto$ stands for "proportional to".

At this point, it is immediate that by choosing ϕ to be the *standard Gaussian* distribution, we retrieve the Edwards-Anderson model in (1). Note however, that (6) above generates a whole family of "spin-glass" distributions with the EA-model being a special case.

The definition of the distribution of the broadcasting matrix in (5) allows us to derive an explicit formula for the quantity Δ_{KS} in (4). Specifically, for $\boldsymbol{J}$ distributed according to ϕ, it is not hard to prove (see the full version) that

$$\Delta_{\mathrm{KS}}(\beta,\phi)=\left(\mathbb{E}\left[\left(\frac{1-\exp(\beta\boldsymbol{J})}{1+\exp(\beta\boldsymbol{J})}\right)^2\right]\right)^{-1}, \tag{7}$$

where the expectation is with respect to the random variable $\boldsymbol{J}$. In light of the above, we prove the following result for the general Gibbs distribution.

▶ **Theorem 3.** *For a real number $\beta > q0$, and a distribution ϕ on the real numbers $\mathbb{R}$ let $\Delta_{\mathrm{KS}} = \Delta_{\mathrm{KS}}(\beta,\phi)$ be defined as in (7).*

For any integer $\Delta > \Delta_{\mathrm{KS}}$, the Gibbs distribution $\boldsymbol{\mu}_{\beta,\phi}$, defined as in (6), on the Δ-ary tree exhibits reconstruction. On the other hand, if $\Delta < \Delta_{\mathrm{KS}}$ the distribution $\boldsymbol{\mu}_{\beta,\phi}$ exhibits non-reconstruction.

The proof of Theorem 3 appears in the full version. Let us state the implications of Theorem 3 for the Edwards-Anderson model on the Δ-ary tree.

▶ **Corollary 4.** *For $\beta > 0$ and the standard Gaussian $\boldsymbol{J}$, let*

$$\Delta_{\mathrm{EA}}(\beta)=\left(\mathbb{E}\left[\left(\frac{1-\exp(\beta\boldsymbol{J})}{1+\exp(\beta\boldsymbol{J})}\right)^2\right]\right)^{-1},$$

where the expectation is with respect to $\boldsymbol{J}$.

For any integer $\Delta > \Delta_{\mathrm{EA}}(\beta)$, the distribution $\boldsymbol{\mu}_\beta$, the Edwards-Anderson model with inverse temperature β on the Δ-ary tree, exhibits reconstruction. On the other hand, if $\Delta < \Delta_{\mathrm{EA}}(\beta)$ the distribution $\boldsymbol{\mu}_\beta$ exhibits non-reconstruction.

2.1 The case of the Galton-Watson tree

As a further step, we study the reconstruction problem on the Galton-Watson tree. Even though this is a very interesting problem on its own, we make use of our results for the Galton-Watson tree to derive subsequent results for $\boldsymbol{G}(n,d/n)$, see Section 2.2.

Let $\zeta : \mathbb{Z}_{\geq 0} \to [0,1]$ be a distribution over the non-negative integers. Then, the rooted tree $\boldsymbol{T}$ is a Galton-Watson tree with offspring distribution ζ, if the number of children for each vertex in $\boldsymbol{T}$ is distributed according to ζ, *independently* from the other vertices.

Note that broadcasting with random matrices over the Galton-Watson tree $\boldsymbol{T}$, gives rise to configurations that consist of *three* levels of randomness. One of the challenges we circumvent with our analysis, is to disentangle all of three levels of randomness, and make clear the contribution of each one of them. Before getting there, we need to clarify what we mean by (non-)reconstruction in the current setting.

▶ **Definition 5.** *Consider the distributions ϕ over $\mathbb{R}$ and ζ over $\mathbb{Z}_{\geq 0}$, and a real number $\beta \geq 0$. Let the Galton-Watson tree $\boldsymbol{T}$ with offspring distribution ζ, while let the measure $\boldsymbol{\mu} = \boldsymbol{\mu}_{\beta,\phi}$ be defined as in (6), on the tree $\boldsymbol{T}$. We say that $\boldsymbol{\mu}$ exhibits* reconstruction *if*

$$\limsup_{h\to\infty} \mathbb{E}_{\boldsymbol{T}}\left[\ \mathbb{E}_{\boldsymbol{\mu}}\left[||\boldsymbol{\mu}_h(\cdot \mid \boldsymbol{\sigma}(r)=+1)-\boldsymbol{\mu}_h(\cdot \mid \boldsymbol{\sigma}(r)=-1)||_{\mathrm{TV}} \ \mid \boldsymbol{T}\right]\ \right] > 0\ .$$

On the other hand, if the above limit is zero, then we have non-reconstruction*.*

For the above, recall that $\boldsymbol{\mu}_h$ is the marginal of $\boldsymbol{\mu}$ on the set of vertices at distance h from the root. Note that if $\boldsymbol{T}$ has no vertex at level h, then the total variation distance above is, degenerately, equal to zero. We use the double expectation in Definition 5 for the sake of clarity: we can just replace it by a single expectation with respect to both the random tree $\boldsymbol{T}$, and the random measure $\boldsymbol{\mu}$.

As far as the reconstruction problem on the Galton-Watson trees is concerned, we have the following result.

▶ **Theorem 6.** *For any real numbers $d > 0, \beta > 0$, for any distribution ϕ on $\mathbb{R}$, for any distribution ζ on $\mathbb{Z}_{\geq 0}$ with expectation d, and* bounded second moment*, let $\boldsymbol{T}$ be the Galton-Watson tree with offspring distribution ζ. Let also $\boldsymbol{\mu}_{\beta,\phi}$ be the Gibbs distribution defined as in* (6)*, on the tree $\boldsymbol{T}$. Finally, let $\Delta_{\mathrm{KS}} = \Delta_{\mathrm{KS}}(\beta, \phi)$ be defined as in* (7)*.*

The distribution $\boldsymbol{\mu}_{\beta,\phi}$ exhibits reconstruction if $d > \Delta_{\mathrm{KS}}$. On the other hand, if $d < \Delta_{\mathrm{KS}}$, the distribution $\boldsymbol{\mu}_{\beta,\phi}$ exhibits non-reconstruction.

Let us now state the implications of Theorem 6 for the Edwards-Anderson model on the Galton-Watson tree.

▶ **Corollary 7.** *For $\beta > 0$, consider the quantity $\Delta_{\mathrm{EA}}(\beta)$ defined in Corollary 4. For any real number $d > 0$, and any distribution $\zeta : \mathbb{Z}_{\geq 0} \to [0, 1]$ with expectation d, and* bounded second moment*, let $\boldsymbol{T}$ be the Galton-Watson tree with offspring distribution ζ.*

Then, for $\boldsymbol{\mu}_\beta$ the Edwards-Anderson model with inverse temperature β, on the tree $\boldsymbol{T}$, the following is true. The distribution $\boldsymbol{\mu}_\beta$ exhibits reconstruction if $d > \Delta_{\mathrm{EA}}(\beta)$. On the other hand, if $d < \Delta_{\mathrm{EA}}(\beta)$, the distribution $\boldsymbol{\mu}_\beta$ exhibits non-reconstruction.

2.2 The Edwards-Anderson model on $\boldsymbol{G}(n, d/n)$

For integer $n \geq 1$, and real $p \in [0, 1]$, let $\boldsymbol{G} = \boldsymbol{G}(n, p)$ be the random graph on $V_n = \{x_1, \ldots, x_n\}$, whose edge set $E(\boldsymbol{G})$ is obtained by including each edge with probability, p *independently.*

The *Edwards-Anderson model* on $\boldsymbol{G}$ at inverse temperature $\beta > 0$, is defined as follows: for $\boldsymbol{J} = \{\boldsymbol{J}_e\}_{e \in E(\boldsymbol{G})}$ a family of independent *standard Gaussians*, we let

$$\mu_{\boldsymbol{G},\boldsymbol{J},\beta}(\sigma) = \tfrac{1}{Z_\beta(\boldsymbol{G},\boldsymbol{J})} \exp\left(\beta \textstyle\sum_{x \sim y} \mathbf{1}\{\sigma(y) = \sigma(x)\} \cdot \boldsymbol{J}_{\{x,y\}}\right) \ , \tag{8}$$

where

$$Z_\beta(\boldsymbol{G}, \boldsymbol{J}) = \textstyle\sum_{\tau \in \{\pm 1\}^{V_n}} \exp\left(\beta \textstyle\sum_{x \sim y} \mathbf{1}\{\tau(y) = \tau(x)\} \cdot \boldsymbol{J}_{\{x,y\}}\right) \ .$$

Here we assume that $p = \frac{d}{n}$, where $d > 0$ is a fixed number. Typically, we study this distribution as $n \to \infty$. The natural question we ask here is how does the model change as we vary d. According to the physics predictions, for any β there exists a *condensation threshold*, denoted as $d_{\mathrm{cond}}(\beta)$, where the function

$$d \mapsto \lim_{n \to \infty} \frac{1}{n} \mathbb{E}[\ln Z_\beta(\boldsymbol{G}, \boldsymbol{J})]$$

is non-analytic [19]. This conjecture was proved by Guerra and Toninelli [21]. The regime $d < d_{\mathrm{cond}}(\beta)$ is called the *replica symmetric phase.* This region has several interesting properties; here we consider one that seems to be most relevant to our discussion. For any $d < d_{\mathrm{cond}}(\beta)$ the distribution $\mu_{\boldsymbol{G},\boldsymbol{J},\beta}$ satisfies the following property: for $\boldsymbol{\sigma}$ distributed as in $\mu_{\boldsymbol{G},\boldsymbol{J},\beta}$, for two randomly chosen vertices $\mathbf{x}$ and $\mathbf{y}$, the configurations $\boldsymbol{\sigma}(\mathbf{x})$ and $\boldsymbol{\sigma}(\mathbf{y})$ are asymptotically independent. Formally, the above can be expressed as follows: for $d < d_{\text{cond}}(\beta)$ and any $i, j \in \{\pm 1\}$, we have that

$$\limsup_{n\to\infty} \frac{1}{n^2} \sum_{x,y\in V_n} \mathbb{E}\left[\langle \mathbf{1}\{\boldsymbol{\sigma}(x)=i\} \times \mathbf{1}\{\boldsymbol{\sigma}(y)=j\}\rangle - \langle \mathbf{1}\{\boldsymbol{\sigma}(x)=i\}\rangle \times \langle \mathbf{1}\{\boldsymbol{\sigma}(y)=j\}\rangle\right] = 0 \ ,$$

where $\langle\cdot\rangle$ denotes expectation with respect to the Gibbs distribution $\mu_{\boldsymbol{G},\boldsymbol{J},\beta}$. Note that the above holds not only for pairs of vertices, but also for sets of k vertices, for any fixed integer $k > 0$. Using our notation, the work by Guerra and Toninelli [21] implies the following result.

▶ **Theorem 8** ([21])**.** *For any $\beta > 0$, for the distribution $\mu_{\boldsymbol{G},\boldsymbol{J},\beta}$ defined as in* (8)*, we have that*

$$d_{\text{cond}}(\beta) = \left(\mathbb{E}\left[\left(\tfrac{1-\exp(\beta\boldsymbol{J})}{1+\exp(\beta\boldsymbol{J})}\right)^2\right]\right)^{-1} \ ,$$

where $\boldsymbol{J}$ is a standard Gaussian random variable.

Interestingly, one obtains the above by combining our Theorem 6 and using results from [12, 13]. Our main focus is on the reconstruction threshold for the Edwards-Anderson model on $\boldsymbol{G}$. The reconstruction for $\mu_{\boldsymbol{G},\boldsymbol{J},\beta}(\cdot)$ is defined in a slightly different way than what we have for the random tree.

▶ **Definition 9.** *For $d > 0$, for $\beta > 0$, consider the Gibbs distribution $\mu_{\boldsymbol{G},\boldsymbol{J},\beta}$ as this is defined in* (8)*. We say that the measure $\mu = \mu_{\boldsymbol{G},\boldsymbol{J},\beta}$ exhibits* reconstruction *if*

$$\limsup_{h\to\infty} \lim_{n\to\infty} \frac{1}{n} \sum_{x\in V_n} \mathbb{E}\left[||\mu_{x,h}(\cdot \mid \boldsymbol{\sigma}(x)=+1) - \mu_{x,h}(\cdot \mid \boldsymbol{\sigma}(x)=-1)||_{\text{TV}}\ \right] > 0 \ ,$$

where $\mu_{x,h}$ denote the Gibbs marginal at the vertices at distance h from vertex x. On the other hand, if the above limit is zero, then we have non-reconstruction*.*

Perhaps, it is interesting to notice the order with which we take the double limit in the above definition. We let the reconstruction threshold, denoted as d_{recon}, to be the infimum over $d > 0$ such that

$$\limsup_{h\to\infty} \lim_{n\to\infty} \frac{1}{n} \sum_{x\in V_n} \mathbb{E}\left[||\mu_h(\cdot \mid \boldsymbol{\sigma}(x)=+1) - \mu_h(\cdot \mid \boldsymbol{\sigma}(x)=-1)||_{\text{TV}}\ \right] > 0 \ .$$

The region of values of d such that $d < d_{\text{recon}}$ is called the *non-reconstruction phase*. It is immediate from Definition 9 that, for any $d < d_{\text{recon}}$, we have that non-reconstruction.

In the following result, we prove that the replica symmetric phase coincides with the non-reconstruction phase of the Edwards-Anderson model on $\boldsymbol{G}$.

▶ **Theorem 10.** *For any $\beta > 0$, for the distribution $\mu_{\boldsymbol{G},\boldsymbol{J},\beta}$ defined as in* (8)*, we have that $d_{\text{recon}}(\beta) = d_{\text{cond}}(\beta)$.*

The above follows from Theorems 8, 7 and [12, Corollary 1.5].

Notation

For the graph $G = (V, E)$ and the Gibbs distribution μ on the set of configurations $\{\pm 1\}^V$. For a configuration σ, we let $\sigma(\Lambda)$ denote the configuration that σ specifies on the set of vertices Λ. We let μ_Λ denote the marginal of μ at the set Λ. We let $\mu(\cdot \mid \Lambda, \sigma)$, denote the distribution μ conditional on the configuration at Λ being σ. Also, we interpret the conditional marginal $\mu_\Lambda(\cdot \mid \Lambda', \sigma)$, for $\Lambda' \subseteq V$, in the natural way.

3 Approach

A major challenge in our setting is that we have to deal with multiple levels of randomness, i.e., we have two levels of randomness in the case of the $\varDelta$-ary tree, while the levels increase with the Galton-Watson trees. To circumvent this problem, we follow an analysis that allows us to disentangle the different sources of randomness in our models. In this section, we provide a high-level description of our approach. We restrict our discussion on the $\varDelta$-ary tree.

Non-reconstruction

Consider the $\varDelta$-ary tree $T = (V, E)$ rooted at r. Suppose that we have a distribution μ as in (6) on T, while assume that each edge $e \in E$ has its own coupling parameter J_e. Assume, for the moment, that the coupling parameters at the edges are fixed, e.g. the reader may assume that are arbitrary real numbers. That is, each J_e can be either positive, or negative. Hence, one might consider the aforementioned distribution as a *non-homogenous* Ising model which involves both ferromagnetic and anti-ferromagnetic interactions. Let us focus on non-reconstruction. We derive an upper bound on

$$||\mu_h(\cdot \mid \boldsymbol{\sigma}(r) = +1) - \mu_h(\cdot \mid \boldsymbol{\sigma}(r) = -1)||_{\mathrm{TV}} \quad ,$$

which is expressed in terms of the *influence* between neighbouring vertices. The notion of influence between vertices is the same as the one developed in the context of *Spectral Independence* technique for establishing rapid mixing of Glauber dynamics [3, 9]. These influences are used in the context of the so-called *down-up* coupling to establish non-reconstruction. This is a coupling approach from [6], which also relies on ideas in [29].

Let us be more specific. For the probability measure μ we consider, let R_r be the *ratio of Gibbs marginals* at the root r defined by

$$R_r = \frac{\mu_r(+1)}{\mu_r(-1)} \quad . \tag{9}$$

Recall that $\mu_r(\cdot)$ denotes the marginal of the Gibbs distribution $\mu(\cdot)$ at the root r. For a vertex $u \in V$, we let T_u be the subtree of T that includes u, and all its descendants. Also, we let R_u be the ratio of marginals at vertex u, where the Gibbs distribution is, now, with respect to the subtree T_u.

Suppose that the vertices $w_1, \dots, w_\varDelta$ are the children of the root r. Our focus is on expressing $\log R_r$ recursively, as a function of $\log R_{w_1}, \dots, \log R_{w_\varDelta}$. Note that we study the logarithm of the ratios involved, which can be viewed as applying the potential function $\log(\cdot)$ to the tree recursions. We have that $\log(R_r) = H(\log R_{w_1}, \dots, \log R_{w_\varDelta})$ where

$$H(x_1, x_2, \dots, x_\varDelta) = \sum_{i=1}^{\varDelta} \log\left(\frac{\exp\left(x_i + \beta J_{\{r,w_i\}}\right) + 1}{\exp(x_i) + \exp\left(\beta J_{\{r,w_i\}}\right)}\right) \quad . \tag{10}$$

Note that $J_{\{r,w_i\}}$ is the coupling parameter that corresponds to the edge between the root r with its child w_i. All the above extends naturally in the case where we impose boundary conditions. That is, for a region $K \subseteq V$, and $\tau \in \{\pm 1\}^K$, we define the ratio of marginals $R_r^{K,\tau}$ at the root, where now the ratio is between the conditional marginals $\mu_r(+1 \mid K, \tau)$ and $\mu_r(-1 \mid K, \tau)$. The recursive function H for the conditional ratios is exactly the same as the one above.

Our interest is on the *gradient* of the function H. Specifically, for every $i \in [\Delta]$, we let

$$\Gamma_{\{r,w_i\}} = \sup_{x_1,\dots,x_\Delta} \left| \frac{\partial}{\partial x_i} H(x_1, x_2, \dots, x_\Delta) \right| \ . \tag{11}$$

It turns out that, in our case, $\Gamma_{\{r,w_i\}}$ has a simple form

$$\Gamma_{\{r,w_i\}} = \frac{\left|1 - \exp\left(\beta J_{\{r,w\}}\right)\right|}{1 + \exp\left(\beta J_{\{r,w_i\}}\right)} \ .$$

Utilising the idea of down-up coupling from [6], we prove the following:

$$||\mu_h(\cdot \mid \boldsymbol{\sigma}(r) = +1) - \mu_h(\cdot \mid \boldsymbol{\sigma}(r) = -1)||_{TV} \leq \sqrt{\sum_{v \in \Lambda} \prod_{e \in \text{path}(r,v)} \Gamma_e^2} \ , \tag{12}$$

where $\Lambda = \Lambda(h)$ denotes the set of vertices at distance h from the root r. Note that the above provides a bound for the total variation distance of the the marginals for fixed, i.e., non-random, couplings $\{J_e\}_{e \in E}$. Inequality (12), extends naturally when we study reconstruction for the distribution $\boldsymbol{\mu}$ defined in (6), i.e., when the coupling parameters $\boldsymbol{J}_e$ are i.i.d. samples from a distribution ϕ. Indeed, averaging yields

$$\mathbb{E}\left[\left(||\boldsymbol{\mu}_h(\cdot \mid \boldsymbol{\sigma}(r) = +1) - \boldsymbol{\mu}_h(\cdot \mid \boldsymbol{\sigma}(r) = -1)||_{TV}\right)^2\right] \leq \sum_{v \in \Lambda} \prod_{e \in \text{path}(r,v)} \mathbb{E}\left[\boldsymbol{\Gamma}_e^2\right] \ , \tag{13}$$

where we have $\boldsymbol{\Gamma}_e = \frac{|1-\exp(\beta \boldsymbol{J}_e)|}{1+\exp(\beta \boldsymbol{J}_e)}$, for each $e \in E$. Note that the above holds, since each $\boldsymbol{\Gamma}_e$ depends only on $\boldsymbol{J}_e$, while the coupling parameters $\boldsymbol{J}_e$ are assumed to be independent with each other.

At this point, and since the $\boldsymbol{J}_e$'s are identically distributed, we further observe that for any $e \in E$, we have that

$$\Delta_{\text{KS}}(\beta, \phi) = \left(\mathbb{E}\left[\boldsymbol{\Gamma}_e^2\right]\right)^{-1} \ .$$

Since the underlying tree T is Δ-ary, it is immediate to see that for $\Delta < \Delta_{\text{KS}}(\beta, \phi)$, the r.h.s. of (13) tends to zero as $h \to \infty$. From this point on, it is standard to prove non-reconstruction.

Our analysis allows to deal with the randomness of the spin-glass measure $\boldsymbol{\mu}$ by utilising the bound in (12). That is, the upper bound on the total variation distance has a nice *product* form of the quantities Γ_e, which, in turn, expresses the dependence of the total variation distance on the edge couplings $\{J_e\}_{e \in E}$. This product form of the bound, behaves rather nicely when we need to take averages over the randomness of the coupling parameters $\{\boldsymbol{J}_e\}_{e \in E}$ of the the spin-glass measure $\boldsymbol{\mu}$.

Reconstruction

In the reconstruction regime, the configuration at the root has a non-vanishing effect on the configuration of the vertices at distance h, regardless of the height h. Specifically, the corresponding leaf configurations from the measure conditioned on root's spin being $+1$, and -1, are so different with each other, that discrepancies cannot be attributed to random fluctuations. Therefore, a question that naturally arises is how can we take advantage of the discrepancies so that we infer the spin of the root.

For the standard ferromagnetic Ising, several approaches have been developed to establish reconstruction (see [18], [8], [23]). Here, we build on an elegant argument in [18]. The authors in this work, show that a simple *majority vote* of the leaf spins, conveys information sufficient to reconstruct root's spin, The majority vote on the leaves is defined by

$$M_h = \sum_{u \in \Lambda} \sigma(u) \ . \tag{14}$$

The estimation rule is to infer that the spin at the root is $\mathrm{sgn}\{M_h\}$, i.e., the sign of M_h. Impressively, it turns out that this estimator is optimal, i.e., it coincides with the *maximum likelihood* one. For the Δ-ary tree, one establishes reconstruction for the ferromagnetic Ising model by employing a second moment argument on the estimator M_h.

For the distributions we consider here, the above estimator is far from sufficient. This is due to various facts. Firstly, we allow for mixed couplings on the edges, i.e., certain edges can be ferromagnetic, and others can be anti-ferromagnetic. Secondly, the strength of the interaction, i.e., the magnitude of $\boldsymbol{J}_e$'s, is expected to vary from one edge to the other. To this end, we introduce a new estimator, and we establish reconstruction by building on the second moment argument from [18]. The starting point towards deriving this estimator, comes from just considering the standard anti-antiferromagnetic Ising. The statistic from (14), clearly does not work for this distribution. However, there is an easy remedy, by taking into account the parity of the height h, i.e., if h is an even, or an odd number. We infer that the spin at the root is equal to $\mathrm{sgn}\left\{\widehat{M_h}\right\}$, where

$$\widehat{M_h} = (-1)^h \sum_{u\in\Lambda} \sigma(u) \ .$$

For the spin-glass distributions we consider here, we need to get the above idea even further. Firstly, in order to accommodate the *mixed* ferromagnetic and anti-ferromagnetic couplings on the edges of the tree. It seems meaningful to use the estimator $\mathrm{sgn}\left\{\widetilde{M_h}\right\}$ for the root configuration, where

$$\widetilde{M_h} = \sum_{u\in\Lambda} \sigma(u) \prod_{e\in\mathrm{path}(r,u)} \mathrm{sign}\{\boldsymbol{J}_e\} \ ,$$

with $\mathrm{path}(r,u)$ denoting the set of edges along the unique path connecting r to u. So that in $\widetilde{M_h}$, for each leaf we essentially examine the parity of the number of antiferromagnetic couplings along the path that connects it to the root. Unfortunately, for the above estimator, our second moment argument does not seem to work all that well.

The estimator we end up using, is a *reweighted* version of $\widetilde{M_h}$, which we call the "flip majority" vote, and is defined by

$$F_h = \sum_{u\in\Lambda} \sigma(u) \prod_{e\in\mathrm{path}(r,u)} \tfrac{1-\exp(\beta \boldsymbol{J}_e)}{1+\exp(\beta \boldsymbol{J}_e)} \ .$$

Note that the absolute value of the weight for the edge e, above, coincides with the quantity $\boldsymbol{\Gamma}_e$ in (13). Naturally, the estimation rule is to infer that the root spin is $\mathrm{sgn}\left\{F_h\right\}$.

References

1. Dimitris Achlioptas and Amin Coja-Oghlan. Algorithmic barriers from phase transitions. In *2008 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 793–802. IEEE, 2008.
2. Ahmed El Alaoui, Andrea Montanari, and Mark Sellke. Sampling from the Sherrington-Kirkpatrick Gibbs measure via algorithmic stochastic localization. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, Denver, CO, USA, October 31 - November 3, 2022*, pages 323–334. IEEE, 2022. doi:10.1109/FOCS54457.2022.00038.
3. Nima Anari, Kuikui Liu, and Shayan Oveis Gharan. Spectral independence in high-dimensional expanders and applications to the hardcore model. *SIAM Journal on Computing*, 0(0):FOCS20–1–FOCS20–37, 2021. doi:10.1137/20M1367696.

4 Victor Bapst, Amin Coja-Oghlan, and Charilaos Efthymiou. Planting colourings silently. *Combinatorics, probability and computing*, 26(3):338–366, 2017.

5 Nayantara Bhatnagar, Allan Sly, and Prasad Tetali. Reconstruction threshold for the hardcore model. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques: 13th International Workshop, APPROX 2010, and 14th International Workshop, RANDOM 2010, Barcelona, Spain, September 1-3, 2010. Proceedings*, pages 434–447. Springer, 2010.

6 Nayantara Bhatnagar, Juan Vera, Eric Vigoda, and Dror Weitz. Reconstruction for colorings on trees. *SIAM Journal on Discrete Mathematics*, 25(2):809–826, 2011.

7 Pavel M Bleher, Jean Ruiz, and Valentin A Zagrebnov. On the purity of the limiting Gibbs state for the Ising model on the Bethe lattice. *Journal of Statistical Physics*, 79:473–482, 1995.

8 Christian Borgs, Jennifer Chayes, Elchanan Mossel, and Sébastien Roch. The Kesten-Stigum reconstruction bound is tight for roughly symmetric binary channels. In *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*, pages 518–530. IEEE, 2006.

9 Zongchen Chen, Kuikui Liu, and Eric Vigoda. Rapid mixing of Glauber dynamics up to uniqueness via contraction. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1307–1318. IEEE, 2020.

10 Amin Coja-Oghlan and Charilaos Efthymiou. On independent sets in random graphs. *Random Structures & Algorithms*, 47(3):436–486, 2015.

11 Amin Coja-Oghlan, Charilaos Efthymiou, and Nor Jaafari. Local convergence of random graph colorings. *Combinatorica*, 38(2):341–380, 2018.

12 Amin Coja-Oghlan, Charilaos Efthymiou, Nor Jaafari, Mihyun Kang, and Tobias Kapetanopoulos. Charting the replica symmetric phase. *Communications in Mathematical Physics*, 359:603–698, 2018.

13 Amin Coja-Oghlan, Andreas Galanis, Leslie Ann Goldberg, Jean Bernoulli Ravelomanana, Daniel Stefankovic, and Eric Vigoda. Metastability of the Potts Ferromagnet on Random Regular Graphs. In Mikolaj Bojanczyk, Emanuela Merelli, and David P. Woodruff, editors, *49th International Colloquium on Automata, Languages, and Programming, ICALP 2022, July 4-8, 2022, Paris, France*, volume 229 of *LIPIcs*, pages 45:1–45:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.ICALP.2022.45.

14 Amin Coja-Oghlan, Tobias Kapetanopoulos, and Noela Müller. The replica symmetric phase of random constraint satisfaction problems. *Combinatorics, Probability and Computing*, 29(3):346–422, 2020.

15 Amin Coja-Oghlan, Florent Krzakala, Will Perkins, and Lenka Zdeborová. Information-theoretic thresholds from the cavity method. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 146–157, 2017.

16 Samuel Frederick Edwards and Phil W Anderson. Theory of spin glasses. *Journal of Physics F: Metal Physics*, 5(5):965, 1975.

17 Charilaos Efthymiou. On Sampling Symmetric Gibbs Distributions on Sparse Random Graphs and Hypergraphs. In Mikolaj Bojanczyk, Emanuela Merelli, and David P. Woodruff, editors, *49th International Colloquium on Automata, Languages, and Programming, ICALP 2022, July 4-8, 2022, Paris, France*, volume 229 of *LIPIcs*, pages 57:1–57:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.ICALP.2022.57.

18 William Evans, Claire Kenyon, Yuval Peres, and Leonard J Schulman. Broadcasting on trees and the Ising model. *Annals of Applied Probability*, pages 410–433, 2000.

19 Silvio Franz, Michele Leone, Federico Ricci-Tersenghi, and Riccardo Zecchina. Exact solutions for diluted spin glasses and optimization problems. *Physical review letters*, 87(12):127209, 2001.

20 Hans-Otto Georgii. *Gibbs measures and phase transitions*, volume 9. Walter de Gruyter, 2011.

21 Francesco Guerra and Fabio Lucio Toninelli. The high temperature region of the Viana-Bray diluted spin glass model. *Journal of statistical physics*, 115:531–555, 2004.

22 Yasunari Higuchi. Remarks on the limiting Gibbs states on a (d+ 1)-tree. *Publications of the Research Institute for Mathematical Sciences*, 13(2):335–348, 1977.
23 Dmitry Ioffe. On the extremality of the disordered state for the Ising model on the Bethe lattice. *Letters in Mathematical Physics*, 37:137–143, 1996.
24 Harry Kesten and Bernt P Stigum. Additional limit theorems for indecomposable multidimensional Galton-Watson processes. *The Annals of Mathematical Statistics*, 37(6):1463–1481, 1966.
25 Marc Mézard and Andrea Montanari. Reconstruction on trees and spin glass transition. *Journal of statistical physics*, 124:1317–1350, 2006.
26 Michael Molloy. The freezing threshold for k-colourings of a random graph. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 921–930, 2012.
27 Giorgio Parisi. Infinite number of order parameters for spin-glasses. *Physical Review Letters*, 43(23):1754, 1979.
28 David Sherrington and Scott Kirkpatrick. Solvable model of a spin-glass. *Physical review letters*, 35(26):1792, 1975.
29 Allan Sly. Reconstruction of random colourings. *Communications in Mathematical Physics*, 288(3):943–961, 2009.
30 Daniel L Stein and Charles M Newman. *Spin glasses and complexity*, volume 4. Princeton University Press, 2013.
31 Lenka Zdeborová and Florent Krzakala. Statistical physics of inference: Thresholds and algorithms. *Advances in Physics*, 65(5):453–552, 2016.