One-step replica symmetry breaking of random regular NAE-SAT

Danny Nam *, Allan Sly*, and Youngtak Sohn†
*Department of Mathematics
Princeton University, NJ 08544
Email:{dhnam,asly}@math.princeton.edu
†Department of Statistics
Stanford University, CA 94305
Email: youngtak@stanford.edu

2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS) | 978-1-6654-2055-6/22/$31.00 ©2022 IEEE | DOI: 10.1109/FOCS52979.2021.00039

***Abstract*—In a broad class of sparse random constraint satisfaction problems (CSP), deep heuristics from statistical physics predict that there is a *condensation phase transition* before the satisfiability threshold, governed by *one-step replica symmetry breaking* (1RSB). In fact, in random regular k-NAE-SAT, which is one of such random CSPs, it was verified [1] that its free energy is well-defined and the explicit value follows the 1RSB prediction. However, for any model of sparse random CSP, it has been unknown whether the solution space indeed *condensates* on O(1) clusters according to the 1RSB prediction. In this paper, we give an affirmative answer to this question for the random regular k-NAE-SAT model. Namely, we prove that with probability close to one, most of the solutions lie inside a bounded number of solution clusters whose sizes are comparable to the scale of the free energy. Furthermore, we establish that the overlap between two independently drawn solutions concentrates precisely at two values. This is the defining property of the one-step replica symmetry breaking class which we establish for the first time in a sparse random CSP. Our proof is based on a detailed moment analysis of a spin system, which has an infinite spin space that encodes the structure of solution clusters. We develop new techniques to study the partition function as well as enhance previous approaches which were only applicable to spin systems with finitely many spins. We believe that our method is applicable to a broad range of random CSPs in the 1RSB universality class.**

***Keywords*-random constraint satisfaction problems, random regular NAE-SAT model, one-step replica symmetry breaking, condensation phase transition**

I. Introduction

A random constraint satisfaction problem (rCSP) is defined by a collection of variables whose configuration should satisfy a set of randomly chosen constraints. Namely, there are n variables $\underline{x} = \{x_i\}_{i=1}^n \in \mathfrak{X}^n$ taking values in a finite alphabet set $\mathfrak{X}$, and they are subject to $m \equiv \alpha n$ randomly drawn constraints. The major interest is to understand the structure of the solution space of rCSPs as $n, m \to \infty$ while α being fixed. Indeed, statistical physicists developed a deep but non-rigorous theory to study these problems and conjectured that in a wide class of rCSPs, there is a fascinating series of phase transitions as α varies ([2], [3]; cf. [4] and Chapter 19 of [5] for a survey). As we detail below, the present paper focuses on investigating the solution space structure when α is in the *condensation regime*, for a rCSP model called the *random regular* k-NAE-SAT.

The canonical rCSP is random k-SAT, a random Boolean CNF formula formed by taking the AND of clauses, each of which is the OR of k variables or their negations. A *not-all-equal-satisfiability* (NAE-SAT) formula, has the same form as k-SAT but asks that both $\underline{x}$ an assignment of the variables and $\neg\underline{x}$ its negation evaluate to true in the formula. We call such formula k-NAE-SAT if the clauses appearing in the CNF formula have exactly k literals, and it is called d-regular if each variable appears precisely in d clauses. One can then choose a d-regular k-NAE-SAT problem of n variables uniformly at random, which gives the *random d-regular* k-NAE-SAT problem, with clause density $\alpha = d/k$. Compared to the k-SAT problem, the NAE-SAT problem possesses extra symmetries that make it more tractable from a mathematical perspective. Nevertheless, it is predicted to belong to the same universality class of rCSPs as random k-SAT and random graph coloring, and hence is expected to share the most interesting qualitative behaviors with them.

Let $Z \equiv Z_n$ denote the number of solutions for a given random d-regular k-NAE-SAT instance. Physicists predict that for each fixed α, there exists $\mathsf{f}(\alpha)$ such that

$$\frac{1}{n}\log Z \longrightarrow \mathsf{f}(\alpha) \quad \text{in probability.}$$

A direct computation of the first moment $\mathbb{E}Z$ gives that

$$\mathbb{E}Z = 2^n\left(1-2^{-k+1}\right)^m = e^{n\mathsf{f}^{\mathrm{RS}}(\alpha)}, \quad \text{where}$$
$$\mathsf{f}^{\mathrm{RS}}(\alpha) \equiv \log 2 + \alpha\log\left(1-2^{-k+1}\right),$$

(the superscript RS refers to the *replica-symmetric* free energy) and we see that $\mathsf{f} \le \mathsf{f}^{\mathrm{RS}}$, by Markov's inequality. The previous works of Ding-Sly-Sun [7] and Sly-Sun-Zhang [1] established some of the physics conjectures on the description of Z and f given in [8], [3], [9], which can be summarized as follows.

2575-8454/21/$31.00 ©2021 IEEE
DOI 10.1109/FOCS52979.2021.00039

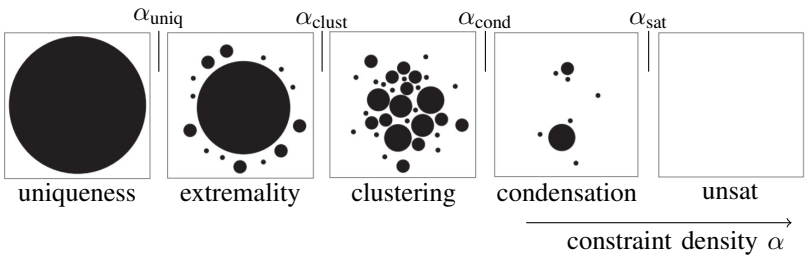

Figure 1: *Figure adapted from* [3], [6]. A pictorial description of the conjectured phase diagram of random regular k-NAE-SAT. In the condensation regime $(\alpha_{\text{cond}}, \alpha_{\text{sat}})$, there remains a bounded number of clusters containing most of the solutions.

- ([7]) For large enough k, there exists the *satisfiability threshold* $\alpha_{\text{sat}} \equiv \alpha_{\text{sat}}(k) > 0$ such that

$$\lim_{n\to\infty} \mathbb{P}(Z > 0) = \begin{cases} 1 & \text{for } \alpha \in (0, \alpha_{\text{sat}}); \\ 0 & \text{for } \alpha > \alpha_{\text{sat}}. \end{cases}$$

- ([1]) For large enough k, there exists the *condensation threshold* $\alpha_{\text{cond}} \equiv \alpha_{\text{cond}}(k) \in (0, \alpha_{\text{sat}})$ such that

$$\mathsf{f}(\alpha) = \begin{cases} \mathsf{f}^{\text{RS}}(\alpha) & \text{for } \alpha \le \alpha_{\text{cond}}; \\ \mathsf{f}^{\text{1RSB}}(\alpha) & \text{for } \alpha > \alpha_{\text{cond}}, \end{cases} \tag{1}$$

  where $\mathsf{f}^{\text{1RSB}} \equiv \mathsf{f}^{\text{1RSB}}(\alpha)$ is the 1RSB *free energy*. Moreover, $\mathsf{f}^{\text{RS}} > \mathsf{f}^{\text{1RSB}}$ on $(\alpha_{\text{cond}}, \alpha_{\text{sat}})$. For the explicit formula and derivation of f^{1RSB}, we refer to Section 1.6 of [1] for a concise overview.

Furthermore, the physics predictions say that the solution space of the random regular k-NAE-SAT is *condensed* when $\alpha \in (\alpha_{\text{cond}}, \alpha_{\text{sat}})$ into a finite number of *clusters* (Figure 1). Here, *clusters* are defined by the connected components of the solution space, where we connect two solutions if they differ by one variable. Our first main result verifies the prediction for all $k \ge k_0$, where k_0 is a universal constant. It is the first to provide a rigorous *cluster-level* description of the solution space of a sparse rCSP in the condensation regime.

Theorem I.1. *Let $k \ge k_0$ and $\alpha \in (\alpha_{\text{cond}}, \alpha_{\text{sat}})$ such that $d \equiv \alpha k \in \mathbb{N}$. For all $\varepsilon > 0$ and $M \in \mathbb{N}$, there exist constants $K \equiv K(\varepsilon, \alpha, k) \in \mathbb{N}$ and $C \equiv C(M, \varepsilon, \alpha, k) > 0$ such that with probability at least $1-\varepsilon$, the random d-regular k-NAE-SAT instance satisfies the following:*

(a) *The number of solutions is no greater than $\exp(n\mathsf{f}^{\text{1RSB}}(\alpha) - c_\star \log n + C)$, where f^{1RSB} is the* 1RSB *free energy and $c_\star \equiv c_\star(\alpha, k) > 0$ is a fixed constant (see (6) for the definition).*

(b) *The K largest solution clusters, $\mathcal{C}_1, \dots, \mathcal{C}_K$, occupy at least $1 - \varepsilon$ fraction of the solution space.*

(c) *There are at least $\exp(n\mathsf{f}^{\text{1RSB}}(\alpha) - c_\star \log n - C)$ many solutions in $\mathcal{C}_1, \dots, \mathcal{C}_M$, the M largest clusters.*

We now briefly discuss the principles underlying the condensation predictions which are helpful in understanding the main theorem. As shown in Figure 1, the solution space of the random regular k-NAE-SAT is predicted to be *clustered* into exponentially many clusters with each of them occupying an exponentially small mass when $\alpha \in (\alpha_{\text{clust}}, \alpha_{\text{cond}})$. As α gets larger than $\alpha_{\text{cond}}(> \alpha_{\text{clust}})$ (the *condensation regime*), the solution space becomes *condensed*, which causes the failure of the first moment analysis as seen in (1). When $\alpha \in (\alpha_{\text{cond}}, \alpha_{\text{sat}})$, the number of clusters that contribute the most to $\mathbb{E}Z$ is exponentially small in n, meaning that those clusters are no longer present in a typical instance of the rCSP. Thus, the leading order of Z is given by the largest clusters that can typically exist (which are thus smaller than the main contributors to $\mathbb{E}Z$), and the number of such clusters is believed to be bounded. Moreover, it is expected that the sizes of those clusters are comparable to the 1RSB free energy.

Theorem I.1 verifies that the solution space indeed becomes *condensed* in the condensation regime, while the previous works [10], [1] obtained the evidence of the condensation phenomenon in the level of free energy. Furthermore, it is believed that the nature of the condensation is governed by *one-step replica symmetry breaking*, which we detail in the following subsection.

Compared to the previous related works [6], [11], [7], [1] in similar settings, we interpret the partition function from a different perspective in order to acquire information on the *number of clusters of particular sizes*. Our approach requires a detailed analysis of an auxiliary spin system with an *infinite* spin space, and one of our major accomplishments is to develop new ideas and generalize existing theories to understand such a system.

A. One-step replica symmetry breaking

In the condensation regime $\alpha \in (\alpha_{\text{cond}}, \alpha_{\text{sat}})$, the random regular k-NAE-SAT model is believed to possess a single layer of hierarchy of clusters in the solution space. Roughly speaking, the prediction is that within a cluster, we can move from one solution to another

by flipping one (or a small number of) variable(s) at once, in such a way that the intermediate steps all belong to the same cluster. Thus, the solutions are fairly *well-connected* inside each cluster so that no additional hierarchical structure occurs in it. Such behaviors are conjectured in various other models, called 1RSB universality class, such as random graph coloring and random k-SAT. However, we remark that there are also other models such as maximum independent set (or high-fugacity hard-core model) in random graphs with small degrees [12] and Sherrington-Kirkpatrick model (on the complete graph) [13], which are expected (or proven [14]) to undergo *full* RSB, meaning that there are infinitely many levels of hierarchy inside the solution clusters.

One way to characterize 1RSB is to look at the *overlap* between two uniformly and independently drawn solutions. In the condensation regime, since there are a bounded number of clusters containing most of the mass, with a non-trivial probability the two solutions belong to the same cluster. According to the description of 1RSB, there is no additional structure inside each cluster, and hence the Hamming distance between two independently selected solutions is expected to be concentrated precisely at *two values*, depending on whether they came from the same cluster or not.

Our second result verifies that this is indeed the case for the random regular k-NAE-SAT with large k, establishing for the first time a rigorous characterization of 1RSB in sparse rCSPs.

Definition I.2. For $\underline{x}^1, \underline{x}^2 \in \{0,1\}^n$, let $\underline{y}^i = 2\underline{x}^i - \mathbf{1}$. The overlap $\rho(\underline{x}^1, \underline{x}^2)$ is defined by

$$\rho(\underline{x}^1, \underline{x}^2) \equiv \frac{1}{n}\underline{y}^1 \cdot \underline{y}^2 = \frac{1}{n}\sum_{i=1}^{n} y_i^1 y_i^2.$$

Theorem I.3. *Let $k \geq k_0$, $\alpha \in (\alpha_{\text{cond}}, \alpha_{\text{sat}})$ such that $d \equiv \alpha k \in \mathbb{N}$. There exists an explicit constant $p^\star \equiv p^\star(\alpha, k) \in (0,1)$ such that the following holds: for all $\varepsilon > 0$, there exists a constant $\delta = \delta(\varepsilon) > 0$ such that with probability at least $1-\varepsilon$, the random d-regular k-NAE-SAT instance $\mathscr{G}$ satisfies the following. Let $\underline{x}^1, \underline{x}^2 \in \{0,1\}^n$ be independent, uniformly chosen satisfying assignments given $\mathscr{G}$. Then, the absolute value $\rho_{\text{abs}} \equiv |\rho|$ of their overlap $\rho \equiv \rho(\underline{x}^1, \underline{x}^2)$ satisfies*

(a) $\mathbb{P}\big(\rho_{\text{abs}} \leq n^{-1/3} \big| \mathscr{G}\big) \geq \delta$;
(b) $\mathbb{P}\big(\big|\rho_{\text{abs}} - p^\star\big| \leq n^{-1/3} \big| \mathscr{G}\big) \geq \delta$;
(c) $\mathbb{P}\big(\min\{\rho_{\text{abs}}, |\rho_{\text{abs}} - p^\star|\} \geq n^{-1/3} \big| \mathscr{G}\big) \leq n^{-1/3}$.

We remark that in (b), ρ can take either $p^\star + O(n^{-1/3})$ or $-p^\star + O(n^{-1/3})$ with asymptotically equal probability as $n \to \infty$. This is due to the symmetric nature of the NAE-SAT, where $-\underline{x}$ is also a solution if $\underline{x}$ is. Thus, the clusters of solutions come in pairs as well: if $\mathcal{C}$ is a cluster, then so is $-\mathcal{C} := \{-\underline{x} : \underline{x} \in \mathcal{C}\}$.

According to the physics predictions [3], the relative sizes of the largest clusters in the rCSPs with 1RSB in the condensation regime are conjectured to converge to a Poisson-Dirichlet process. Although we provide a cluster-level illustration of the solution space and show that it follows the 1RSB prediction, our method is not strong enough to study the limiting distributions of the cluster sizes, and the conjecture is left as an important open problem in the field.

Remark I.4. Although our definition of a cluster in Theorem I.1 is a connected component of the solution space, our proof shows that Theorem I.1 also holds for a slightly different definition of clusters, where we merge the connected components if they differ in a small, say $\log n$, number of variables. The conjectured description of cluster sizes according to Poisson-Dirichlet process actually corresponds to the latter definition of clusters.

B. Related works

Earlier works on rCSPs focused on determining their satisfiability thresholds and verifying the sharpness of SAT-UNSAT transitions. For rCSP models that are known not to exhibit RSB, such goals were established. These models include random 2-SAT [15], [16], random 1-IN-k-SAT [17], k-XOR-SAT [18]–[20], and random linear equations [21]. On the other hand, for the models which are predicted to display the condensation phenomenon, intensive studies have been conducted to estimate their satisfiability threshold, as shown in [22]–[24] (random k-SAT), [25]–[27] (random k-NAE-SAT), and [28]–[31] (random graph coloring).

The satisfiability thresholds for rCSPs with RSB have been rigorously determined in several models (random regular k-NAE-SAT [7], maximum independent set [11], random regular k-SAT [24] and random k-SAT [6]), where they looked at the number of *clusters* instead of the number of solutions and carried out a demanding second moment method. Although determining the location of colorability threshold is left open, the condensation threshold for random graph coloring was settled in [10], where they conducted a technically challenging analysis based on a clever "planting" technique, and the results were further generalized to other models in [32]. Similarly, [33] identified the condensation threshold for random regular k-SAT, where each variable appears $d/2$-times positive and $d/2$-times negative.

Further theory was developed in [1] to establish the 1RSB free energy prediction for random regular k-NAE-SAT in the condensation regime. However, [1] was not able to present a cluster-level description of an rCSP instance, nor to explain the nature of the condensation phenomenon. Our main contribution is to illustrate the solution space of the random regular NAE-SAT instance

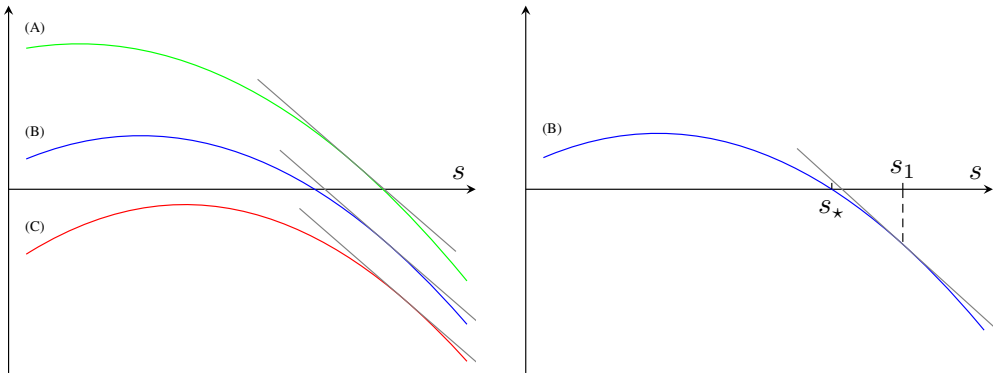

Figure 2: A description of $\Sigma(s;\alpha)$ in s for different values of α. In the left, the curves correspond to the different values of α, with (A) $\alpha \in (\alpha_{\text{clust}}, \alpha_{\text{cond}})$, (B) $\alpha \in (\alpha_{\text{cond}}, \alpha_{\text{sat}})$, and (C) $\alpha > \alpha_{\text{sat}}$, with the gray lines depicting the locations of s_1. In the right, curve (B) is shown with the values s_1 and $s_\star$.

at the cluster-level and to verify that its condensation is governed by 1RSB.

Lastly, the recent work [34] studied the random k-MAX-NAE-SAT beyond α_{sat}, where they verified that the 1RSB description breaks down before $\alpha \asymp k^{-3}4^k$. Indeed, the *Gardner transition* from 1RSB to FRSB is expected at $\alpha_{\text{Ga}} \asymp k^{-3}4^k > \alpha_{\text{sat}}$ [35], [36], and [34] provides evidence of this phenomenon.

C. Heuristic description of condensation

We briefly overview what happens in an rCSP in 1RSB class as the clause density $\alpha = d/k$ varies, as well as a heuristic illustration of condensation.

When α is fairly small, most of the solutions lie inside a single well-connected cluster (in the sense explained in Section I-A). As α becomes larger than α_{clust}, the *clustering threshold*, the solution space becomes shattered into exponentially many clusters, each containing exponentially many solutions yet exponentially small compared to the whole solution space. In this regime, define $\Sigma(s) \equiv \Sigma(s;\alpha)$, the *cluster complexity function*, as

$$\exp(n\Sigma(s)) \equiv \text{expected number of clusters of size } e^{ns}.$$

Indeed, the number of size–e^{ns} clusters is believed to be concentrated around its mean $e^{n\Sigma(s)}$. Thus, the expected number of solutions can be written as

$$\begin{aligned}\mathbb{E}Z &= \sum_s \exp(n\{s+\Sigma(s)\}) \\ &\doteq \exp(n\cdot\max\{s+\Sigma(s) : s \geq 0\}),\end{aligned}$$

where $\doteq$ denotes the equality up to the leading exponential order. The function $\Sigma(s;\alpha)$ is believed to be smooth and concave in s for each fixed α, and indeed physicists predict an explicit formula for $\Sigma(s)$ via the 1RSB *cavity method* [3], [5]. Hence, if this is the case, we have that

$$\mathbb{E}Z \doteq \exp(n\{s_1+\Sigma(s_1)\}),$$

where $s_1 \equiv s_1(\alpha) > 0$ is the unique solution of $\Sigma'(s_1;\alpha) = -1$. However, if $\Sigma(s_1;\alpha) < 0$, meaning that the expected number of size–e^{ns_1} clusters are exponentially small, those clusters are unlikely to exist in a typical instance and hence the main contribution to Z is given by

$$Z \doteq \exp(n\{s_\star+\Sigma(s_\star)\}),$$

where $s_\star$ is defined as

$$s_\star \equiv s_\star(\alpha) \equiv \arg\max_s\{s+\Sigma(s) : \Sigma(s) \geq 0\}. \quad (2)$$

This is the regime where the *condensation phenomenon* occurs, and hence the *condensation threshold* α_{cond} is defined by

$$\begin{aligned}\alpha_{\text{cond}} &\equiv \max\{\alpha : \Sigma(s_1(\alpha);\alpha) \geq 0\} \\ &= \max\{\alpha : s_\star(\alpha) \geq s_1(\alpha)\}.\end{aligned}$$

Thus, for $\alpha > \alpha_{\text{cond}}$, typically we have $Z \ll \mathbb{E}Z$, which causes the failure of the second moment method applied to Z.

For α beyond the *satisfiability threshold* α_{sat}, the problem becomes unsatisfiable ($Z = 0$) with high probability, where α_{sat} is given by

$$\alpha_{\text{sat}} \equiv \min\{\alpha : \Sigma(s;\alpha) \leq 0 \text{ for all } s\}.$$

An illustration of the above discussion is given in Figure 2. We can also see that when $\alpha \in (\alpha_{\text{cond}}, \alpha_{\text{sat}})$, $\Sigma(s_\star(\alpha);\alpha) = 0$, implying that the primary contribution to Z should come from a bounded number of clusters of size roughly $e^{ns_\star}$, whereas if $\alpha < \alpha_{\text{cond}}$ the leading term consists of the clusters of size roughly e^{ns_1} whose numbers are exponentially large. Indeed, in the latter case Z becomes concentrated around $\mathbb{E}Z$ [25]–[27] and the overlap is expected to concentrate around one point as opposed to two points, stated in Theorem I.3.

As $k \to \infty$, asymptotic values of the thresholds are known to be

$$\alpha_{\text{cond}} = \left(2^{k-1} - 1\right) \log 2 + o_k(1),$$
$$\alpha_{\text{sat}} = \left(2^{k-1} - \frac{1}{2} - \frac{1}{4 \log 2}\right) \log 2 + o_k(1).$$

The known upper bound for α_{clust} [37] tells us that it is relatively much smaller than α_{cond} and α_{sat} if k is large. Moreover, α_{clust} is believed to coincide with the *reconstruction threshold*, where we refer the readers to [38], [3], [39], [40] for further information.

D. Tilted cluster partition function and encoding clusters

The main object of study in the present paper shares the same spirit as [1], and its derivation is based on the ideas discussed in Section I-C. We consider the *tilted cluster partition function* $\overline{\mathbf{Z}}_\lambda$, defined as

$$\overline{\mathbf{Z}}_\lambda \equiv \sum_{\Upsilon} |\Upsilon|^\lambda, \tag{3}$$

where the sum is taken over all clusters Υ. If we compute $\mathbb{E}\overline{\mathbf{Z}}_\lambda$ for $\lambda \equiv \lambda(\alpha) \equiv -\Sigma'(s_\star; \alpha)$ (with $s_\star$ as in (2)), then we see that the main contribution comes from the clusters of size $e^{ns_\star}$, following the same reasoning as Section I-C. Thus, we expect to have $\overline{\mathbf{Z}}_\lambda \doteq \mathbb{E}\overline{\mathbf{Z}}_\lambda$, and indeed [1] carried out challenging moment computations in a similar setting to obtain the 1RSB free energy $\mathsf{f}^{\mathsf{1rsb}}$ for random regular k-NAE-SAT.

The next issue is to obtain a combinatorial representation of a cluster. We follow the *coarsening algorithm*, which is an inductive process starting from a solution $\underline{x}$ that sets a variable in $\underline{x}$ to be $\mathtt{f}$ (free) one by one, if no clause is violated when the variable is flipped (that is, $0 \to 1$ or $1 \to 0$). Then, it was observed in [7], [1] that the resulting *frozen configuration* $\underline{y} \equiv \underline{y}(\underline{x}) \in \{0, 1, \mathtt{f}\}^n$ obtained by such a procedure serves as a good representation for a cluster.

To study the size of a cluster, we adapt the framework from [1] to count the number of ways to assign $0/1$-values to free variables in a frozen configuration, which we detail as follows. In the regime of our interest, an important observation is that most of the variables in a solution $\underline{x}$ are *frozen* (so that those variables cannot be flipped in the solution space), while a small constant fraction of them are *free*. Thus, in a frozen configuration $\underline{y} \in \{0, 1, \mathtt{f}\}^n$, the connected structure among the free variables (and their neighboring clauses) would mostly look like trees that are not too large. Heuristically, they can be thought of subcritical branching processes, so the maximal connected *free component* will have size $O(\log n)$. In [1], they utilize the idea of *belief propagation* from statistical physics to effectively count the number of NAE-SAT assignments on a given tree of free variables.

The previous work [1] studied the *truncated* partition function $\overline{\mathbf{Z}}_{\lambda,L}$, which only counts the contributions from the clusters whose free components are *trees* of size at most some finite threshold L. Again based on the branching process heuristics, there is always a constant probability for a subcritical branching process to be larger than L, and hence we may expect that

$$\overline{\mathbf{Z}}_{\lambda,L} \doteq e^{-\delta n} \overline{\mathbf{Z}}_\lambda,$$

where $\delta(L) \to 0$ as L tends to infinity. Thus, they investigated the moments of $\overline{\mathbf{Z}}_{\lambda,L}$ and let $L \to \infty$ to deduce the conclusion on the free energy of the original model. Imposing the finite-size truncation played a crucial role in their work, since it makes the space of *free trees* to be finite so that some of the important methods from the earlier works [6], [11], [7] are applicable without significant changes. However, to obtain Theorem I.1, working with the truncated model is insufficient, since we cannot afford the cost of $e^{-\delta n}$ for any small $\delta > 0$. In the following subsection, we describe a brief overview on the ideas to overcome such difficulties along with an outline of the proof.

E. Proof ideas

The major difficulties in understanding the solution space in the cluster-level can be summarized as follows.

1) In addition to investigating $\overline{\mathbf{Z}}_\lambda$, we need to study the contributions from clusters of sizes in a constant window $[e^{ns}, e^{ns+1})$:
$$\overline{\mathbf{Z}}_{\lambda,s} := \sum_{\Upsilon} |\Upsilon|^\lambda \mathbb{1}\{|\Upsilon| \in [e^{ns}, e^{ns+1})\}. \tag{4}$$
2) As mentioned above, it is required to work with the full space of *free trees* which is infinite.

The ideas to overcome the difficulties above is explained in Section I-E1 below.

In fact, we first compute the first and second moments of $\overline{\mathbf{Z}}_{\lambda,s}$ for $\lambda = \lambda(\alpha) \equiv -\Sigma'(s_\star; \alpha)$ and s sufficently close to the free energy $\mathsf{f}^{\text{1RSB}}(\alpha)$. Let $\overline{\mathbf{N}}_s$ denote the number of clusters whose size is in the interval $[e^{ns}, e^{ns+1})$:

$$\overline{\mathbf{N}}_s := \sum_{\Upsilon} \mathbb{1}\{|\Upsilon| \in [e^{ns}, e^{ns+1})\}.$$

Since $e^{-\lambda}\overline{\mathbf{Z}}_{\lambda,s} \le e^{n\lambda s}\overline{\mathbf{N}}_s \le \overline{\mathbf{Z}}_{\lambda,s}$, a successful computation of first and second moments of $\overline{\mathbf{Z}}_{\lambda,s}$ will give us information on $\overline{\mathbf{N}}_s$ based on the second moment method.

Indeed, we show that for $|s - \mathsf{f}^{\text{1RSB}}(\alpha)| = O(n^{-2/3})$, $\mathbb{E}\overline{\mathbf{Z}}_{\lambda(\alpha),s}$ equals up to constant

$$\mathbb{E}\overline{\mathbf{Z}}_{\lambda(\alpha),s} \asymp \frac{1}{\sqrt{n}} \exp\left(n\lambda(\alpha)\mathsf{f}^{\text{1RSB}}(\alpha)\right). \tag{5}$$

Thus, in order for $\mathbb{E}\overline{\mathbf{N}}_s \asymp 1$ which we expect from the discussion in Section I-C, we must take s to be

$$s_\circ \equiv s_\circ(n, \alpha, C) \equiv \mathsf{f}^{\text{1RSB}}(\alpha) - \frac{c_\star \log n}{n} + \frac{C}{n}, \quad \text{where}$$

$$c_\star \equiv \left(2\lambda(\alpha)\right)^{-1}, \tag{6}$$

and $C \in \mathbb{R}$ is a constant which does not depend on n. Moreover, we show that the second moment can be bounded by

$$\mathbb{E}\overline{\mathbf{N}}_{s_\circ} \lesssim_k e^{-2\lambda(\alpha)C} + e^{-\lambda(\alpha)C}. \tag{7}$$

Therefore, the estimates (5), (7) and the Paley-Zygmund inequality imply

$$\liminf_{n\to\infty} \mathbb{P}\left(\overline{\mathbf{N}}_{s_\circ} \geq 1\right) > 0 \tag{8}$$

for any $C \in \mathbb{R}$. The equations (5), (7), (8) will serve as the building blocks of the proof of Theorem I.1.

In the rest of the Section, we describe the high-level ideas on how to compute the first and second moment of $\overline{\mathbf{Z}}_{\lambda(\alpha),s_\circ}$, and to establish Theorem I.1 and I.3 from such computations.

1) Moment computations: The previous approaches in [6], [11], [7], [1] to study the moments of $\overline{\mathbf{Z}}_\lambda$ were to decompose the quantity into the contributions from different types of "local neighborhood profile" of configurations. However, in our case which has infinitely many types of free components, such methods do not give a good enough understanding on $\overline{\mathbf{Z}}_\lambda$, since the Stirling approximations which were crucial in the earlier works are no longer precise.

Instead, we focus on computing the *cost* of containing each type of free component inside a cluster. One of the interesting observations we make is that conditioned on the "boundary profile" of non-free variables and clauses, the profile of free components is essentially given as the result of *independently throwing in each type of free component with a prescribed probability*.

In an informal manner, we describe this crucial observation for the first moment analysis. Denote B by the boundary profile, which is a collection of empirical measures of certain "types" of non-free variables and clauses, and denote $n_\mathfrak{f}$ by the number of the free component $\mathfrak{f}$. Let $\overline{\mathbf{Z}}_\lambda[B,(n_\mathfrak{f})_\mathfrak{f}]$ be the contribution to $\overline{\mathbf{Z}}_\lambda$ (cf. (3)) from the clusters having the boundary profile B and the number of free components $(n_\mathfrak{f})_\mathfrak{f}$. Then, we show that

$$\mathbb{E}\overline{\mathbf{Z}}_\lambda\left[B,(n_\mathfrak{f})_\mathfrak{f}\right] \asymp g(n,B)\cdot\prod_{\mathfrak{f}}\left[\frac{1}{n_\mathfrak{f}!}\left(\left(\frac{e}{n}\right)^{\gamma(\mathfrak{f})} J_\mathfrak{f} w_\mathfrak{f}^\lambda\right)^{n_\mathfrak{f}}\right], \tag{9}$$

where $g(n,B)$ is a certain explicit function of n and B, $\gamma(\mathfrak{f})+1$ equals the number of tree-excess edges of $\mathfrak{f}$, $w_\mathfrak{f}$ is the number of ways to fill NAE-SAT solutions into the free variables of $\mathfrak{f}$, and $J_\mathfrak{f}$ is a certain embedding number of $\mathfrak{f}$.

From (9), we observe that conditioned on B, the distribution of $(n_\mathfrak{f})_\mathfrak{f}$ is a multinomial. Moreover, fixing B and s amounts to a linear constraint on the number of free components $(n_\mathfrak{f})_\mathfrak{f}$. Thus, calculating $\mathbb{E}\overline{\mathbf{Z}}_{\lambda,s}[B]$, where $\overline{\mathbf{Z}}_{\lambda,s}[B]$ is the contribution to $\overline{\mathbf{Z}}_{\lambda,s}$ (cf. (4)) from clusters having the boundary profile B, is equivalent to calculating the probability of a large deviation event. We thus introduce an exponential tilting factor and appeal to local central limit theorem to compute $\mathbb{E}\overline{\mathbf{Z}}_{\lambda,s}[B]$ for $\lambda=\lambda(\alpha)$ and s close to $\mathsf{f}^{\text{1RSB}}(\alpha)$, which is a technique often used in large deviations theory [41].

However, to sum up the $\mathbb{E}\mathbf{Z}_{\lambda,s}[B]$ for different boundary profiles B, we need to show that the negative definiteness of the free energy of B, i.e. the leading exponent of $\frac{1}{n}\log\mathbb{E}\mathbf{Z}_{\lambda,s}[B]$. In order to so, we use the *resampling method* which we describe below.

2) The resampling method: The resampling method was first introduced in [1] to show negative definiteness of the free energy around its maximizer. The main idea behind the method can be summarized as follows. Given a NAE-SAT instance $\mathscr{G}$ and a frozen configuration $\underline{y}\in\{0,1,\mathfrak{f}\}^n$, sample small, say ε, fraction of variables Y. We sample $v\in Y$ far away from each other so that each free tree containing $v\in Y$ do not intersect. Next, resample the spins around Y conditioned on the configuration outside of depth 1 neighborhood of Y. Then, the empirical profile should become closer to the *optimal profile*, which is obtained by solving a fixed point equation of a certain tree recursion called *belief propagation*. The main issue is to quantify the improvement coming from this *local* update procedure, and it turns out that it is closely related to a convex *tree optimization*.

However, the techniques from [1] are limited to the analysis of spin systems with bounded number of spins. In the *untruncated* partition function $\overline{\mathbf{Z}}_\lambda$, the large trees inevitably appear and we can no longer sample Y so that the free trees around Y are guaranteed to never intersect. In order to overcome this issue, we first show that the frequency of large free trees decays exponentially in the number of variables. We then appeal to this rareness of the large free trees to show that if we sample $|Y|=\lfloor\varepsilon n\rfloor$ vertices uniformly at random, the free trees around Y do not intersect with good enough probability. Then, we perform the resampling procedure $O(\frac{1}{\varepsilon})$ times. This step is the most technically challenging part out of the whole proof.

3) Achieving probability 1: One may hope to have $\mathbb{E}\overline{\mathbf{N}}_{s_\circ}^2 \approx (\mathbb{E}\overline{\mathbf{N}}_{s_\circ})^2$ to show that the right hand side of (8) can be pushed near 1, but this is indeed false in the case of random regular NAE-SAT. One of the main reasons is that the existence of short cycles in the graph causes multiplicative fluctuations of $\overline{\mathbf{N}}_{s_\circ}$. Therefore, our approach is to show that if we rescale $\overline{\mathbf{N}}_{s_\circ}$ according to the effects of short cycles, then the resulting rescaled partition function concentrates, that is, $\mathbb{E}[\widetilde{\mathbf{N}}_{s_\circ}^2] \approx (\mathbb{E}\widetilde{\mathbf{N}}_{s_\circ})^2$ (to be precise, this will only be true when C is negative with a huge magnitude,

due to the intrinsic correlations coming from the largest clusters). Furthermore, we argue that the fluctuations coming from the short cycles are not too big, and hence can be absorbed by $\overline{\mathbf{N}}_{s_\circ}$ if $\mathbb{E}\overline{\mathbf{N}}_{s_\circ}$ is large. To this end, we develop a new argument that combines the ideas of *small subgraph conditioning* [42], [43] and the *Doob martingale approach* [11], [7], [1], which are not effective in our model if used alone.

The *small subgraph conditioning method* ([42], [43]; for a survey, see Chapter 9.3 of [44]) is proven to be useful in many settings [45]–[47] to derive a precise distributional limit of partition functions. Indeed, in [46], this method was applied to the proper coloring model of bipartite random regular graphs, where they determined the limiting distribution of the number of colorings. However, this method relies much on algebraic identities specific to the model which may not robust and is not clear in the current model.

Another technique that inspired our proof is the *Doob martingale approach* introduced in [11], [7]. This method rather directly controls the multiplicative fluctuations of $\mathbf{N}$, by investigating the Doob martingale increments of $\log \mathbf{N}$. It has proven to be useful in the study of the models like random regular NAE-SAT, as seen in [1]. However, in the spin systems with infinitely many spins like our model, some of the key estimates in the argument become false, due to the existence of rare spins which appear with probability $o(1)$.

Our approach blends the two techniques in a novel way to back up each other's limitations. We derive the algebraic identities required for the small subgraph conditioning not in a combinatorial manner, but by a modified Doob martingale approach for the *truncated* partition function $\overline{\mathbf{Z}}_{\lambda,L}$ which has a finite spin space. Then, we take $L \to \infty$ limit on these algebraic identities, and show that they converge to the corresponding formulas for the *untruncated* partition function $\mathbb{E}\overline{\mathbf{Z}}_{\lambda}$. This step requires a refined knowledge on the first and second moments of $\mathbf{Z}_{\lambda,s}$ including the constant coefficient in front of the leading exponential term, which was not needed in the earlier works [7], [1]. We then appeal to the small subgraph conditioning method to deduce the conclusion based on those identities.

4) Concentration of the overlap: Observe that for two uniformly and independently drawn solutions $\underline{x}^1, \underline{x}^2$ from a random regular k-NAE-SAT instance, Theorem I.1 shows that they can be contained either in the same cluster or in different ones, each with strictly positive probability. If they are from the same cluster, the set of frozen variables in both solutions will be the same. Moreover, from the moment computations, the number of free trees will concentrate around an explicit value. Since the $0/1$-values for the free variables are assigned independently for each free trees, we show that the absolute value of the overlap concentrates on a single value $p^\star$. On the other hand, if the two solutions are from different clusters, the results from the second moment computation show that the corresponding two frozen configurations are near-independent and from the $0/1$ symmetry of NAE-SAT model, we will conclude that the overlap concentrates around 0.

The actual proof is more complicated than the description above, since we need to take account of the free components containing a cycle. Based on our methods, we develop a coupling argument between the clusters containing cyclic free components and those without cyclic components, which requires an extended analysis on the moment computations.

The full version of this paper is available as an online preprint.

(https://arxiv.org/abs/2011.14270).

ACKNOWLEDGMENT

We thank Amir Dembo, Nike Sun and Yumeng Zhang for helpful discussions. DN is supported by a Samsung Scholarship. AS is supported by NSF grants DMS-1352013 and DMS-1855527, Simons Investigator grant and a MacArthur Fellowship. YS is partially supported by NSF grants DMS-1613091 and DMS-1954337.

REFERENCES

[1] A. Sly, N. Sun, and Y. Zhang, "The number of solutions for random regular NAE-SAT," in *Proceedings of the 57th Symposium on Foundations of Computer Science*, ser. FOCS '16, 2016, pp. 724–731.

[2] M. Mézard, G. Parisi, and R. Zecchina, "Analytic and algorithmic solution of random satisfiability problems," *Science*, vol. 297, no. 5582, pp. 812–815, 2002. [Online]. Available: https://science.sciencemag.org/content/297/5582/812

[3] F. Krzakała, A. Montanari, F. Ricci-Tersenghi, G. Semerjian, and L. Zdeborová, "Gibbs states and the set of solutions of random constraint satisfaction problems," *Proceedings of the National Academy of Sciences*, vol. 104, no. 25, pp. 10 318–10 323, 2007. [Online]. Available: https://www.pnas.org/content/104/25/10318

[4] D. Achlioptas, A. Naor, and Y. Peres, "Rigorous location of phase transitions in hard optimization problems," *Nature*, vol. 435, no. 7043, pp. 759–764, 2005.

[5] M. Mézard and A. Montanari, *Information, physics, and computation*, ser. Oxford Graduate Texts. Oxford University Press, Oxford, 2009. [Online]. Available: https://doi.org/10.1093/acprof:oso/9780198570837.001.0001

[6] J. Ding, A. Sly, and N. Sun, "Proof of the satisfiability conjecture for large k," in *Proceedings of the Forty-seventh Annual ACM Symposium on Theory of Computing*, ser. STOC '15. New York, NY, USA: ACM, 2015, pp. 59–68. [Online]. Available: http://doi.acm.org/10.1145/2746539.2746619

[7] ——, "Satisfiability threshold for random regular NAE-SAT," *Commun. Math. Phys.*, vol. 341, no. 2, pp. 435–489, 2016.

[8] L. Zdeborová and F. Krzakała, "Phase transitions in the coloring of random graphs," *Phys. Rev. E*, vol. 76, p. 031131, 2007.

[9] A. Montanari, F. Ricci-Tersenghi, and G. Semerjian, "Clusters of solutions and replica symmetry breaking in random k-satisfiability," *J. Stat. Mech. Theory E*, vol. 2008, no. 04, p. P04004, apr 2008. [Online]. Available: https://doi.org/10.1088%2F1742-5468%2F2008%2F04%2Fp04004

[10] V. Bapst, A. Coja-Oghlan, S. Hetterich, F. Raß mann, and D. Vilenchik, "The condensation phase transition in random graph coloring," *Comm. Math. Phys.*, vol. 341, no. 2, pp. 543–606, 2016. [Online]. Available: https://doi.org/10.1007/s00220-015-2464-z

[11] J. Ding, A. Sly, and N. Sun, "Maximum independent sets on random regular graphs," *Acta Math.*, vol. 217, no. 2, pp. 263–340, 2016. [Online]. Available: https://doi.org/10.1007/s11511-017-0145-9

[12] J. Barbier, F. Krzakała, L. Zdeborová, and P. Zhang, "The hard-core model on random graphs revisited," *Journal of Physics: Conference Series*, vol. 473, p. 012021, dec 2013. [Online]. Available: https://doi.org/10.1088%2F1742-6596%2F473%2F1%2F012021

[13] M. Talagrand, "The Parisi formula," *Ann. of Math. (2)*, vol. 163, no. 1, pp. 221–263, 2006. [Online]. Available: https://doi.org/10.4007/annals.2006.163.221

[14] A. Auffinger, W.-K. Chen, and Q. Zeng, "The SK model is infinite step replica symmetry breaking at zero temperature," *Communications on Pure and Applied Mathematics*, vol. 73, no. 5, pp. 921–943, 2020. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/cpa.21886

[15] V. Chvatal and B. Reed, "Mick gets some (the odds are on his side) (satisfiability)," in *Proceedings of the 33rd Annual Symposium on Foundations of Computer Science*, ser. SFCS '92. Washington, DC, USA: IEEE Computer Society, 1992, pp. 620–627. [Online]. Available: https://doi.org/10.1109/SFCS.1992.267789

[16] B. Bollobás, C. Borgs, J. T. Chayes, J. H. Kim, and D. B. Wilson, "The scaling window of the 2-SAT transition," *Random Structures Algorithms*, vol. 18, no. 3, pp. 201–256, 2001. [Online]. Available: https://doi.org/10.1002/rsa.1006

[17] D. Achlioptas, A. Chtcherba, G. Istrate, and C. Moore, "The phase transition in 1-in-k SAT and NAE 3-sat," in *Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms*, ser. SODA '01. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 2001, pp. 721–722. [Online]. Available: http://dl.acm.org/citation.cfm?id=365411.365760

[18] O. Dubois and J. Mandler, "The 3-XORSAT threshold," in *Proceedings of the 43rd Symposium on Foundations of Computer Science*, ser. FOCS '02. Washington, DC, USA: IEEE Computer Society, 2002, pp. 769–778. [Online]. Available: http://dl.acm.org/citation.cfm?id=645413.652160

[19] M. Dietzfelbinger, A. Goerdt, M. Mitzenmacher, A. Montanari, R. Pagh, and M. Rink, "Tight thresholds for cuckoo hashing via XORSAT," in *Automata, Languages and Programming*, S. Abramsky, C. Gavoille, C. Kirchner, F. Meyer auf der Heide, and P. G. Spirakis, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 213–225.

[20] B. Pittel and G. B. Sorkin, "The satisfiability threshold for k-XORSAT," *Combin. Probab. Comput.*, vol. 25, no. 2, pp. 236–268, 2016. [Online]. Available: https://doi.org/10.1017/S0963548315000097

[21] P. Ayre, A. Coja-Oghlan, P. Gao, and N. Müller, "The satisfiability threshold for random linear equations," *arXiv preprint arXiv:1710.07497*, 2017.

[22] L. M. Kirousis, E. Kranakis, D. Krizanc, and Y. C. Stamatiou, "Approximating the unsatisfiability threshold of random formulas," *Random Structures Algorithms*, vol. 12, no. 3, pp. 253–269, 1998. [Online]. Available: https://doi.org/10.1002/(SICI)1098-2418(199805)12:3⟨253::AID-RSA3⟩3.3.CO;2-H

[23] D. Achlioptas and Y. Peres, "The threshold for random k-SAT is $2^k \log 2 - O(k)$," *J. Amer. Math. Soc.*, vol. 17, no. 4, pp. 947–973, 2004. [Online]. Available: https://doi.org/10.1090/S0894-0347-04-00464-3

[24] A. Coja-Oghlan and K. Panagiotou, "The asymptotic k-SAT threshold," *Adv. Math.*, vol. 288, pp. 985–1068, 2016. [Online]. Available: https://doi.org/10.1016/j.aim.2015.11.007

[25] D. Achlioptas and C. Moore, "Random k-SAT: two moments suffice to cross a sharp threshold," *SIAM J. Comput.*, vol. 36, no. 3, pp. 740–762, 2006. [Online]. Available: https://doi.org/10.1137/S0097539703434231

[26] A. Coja-Oghlan and L. Zdeborová, "The condensation transition in random hypergraph 2-coloring," in *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*, ser. SODA '12. ACM, New York, 2012, pp. 241–250.

[27] A. Coja-Oghlan and K. Panagiotou, "Catching the k-NAESAT threshold [extended abstract]," in *STOC'12—Proceedings of the 2012 ACM Symposium on Theory of Computing*. ACM, New York, 2012, pp. 899–907. [Online]. Available: https://doi.org/10.1145/2213977.2214058

[28] D. Achlioptas and A. Naor, "The two possible values of the chromatic number of a random graph," *Ann. of Math. (2)*, vol. 162, no. 3, pp. 1335–1351, 2005. [Online]. Available: https://doi.org/10.4007/annals.2005.162.1335

[29] A. Coja-Oghlan, "Upper-bounding the k-colorability threshold by counting covers," *Electron. J. Combin.*, vol. 20, no. 3, pp. Paper 32, 28, 2013.

[30] A. Coja-Oghlan and D. Vilenchik, "Chasing the k-colorability threshold," in *2013 IEEE 54th Annual Symposium on Foundations of Computer Science—FOCS '13*. IEEE Computer Soc., Los Alamitos, CA, 2013, pp. 380–389. [Online]. Available: https://doi.org/10.1109/FOCS.2013.48

[31] A. Coja-Oghlan, C. Efthymiou, and S. Hetterich, "On the chromatic number of random regular graphs," *J. Combin. Theory Ser. B*, vol. 116, pp. 367–439, 2016. [Online]. Available: https://doi.org/10.1016/j.jctb.2015.09.006

[32] A. Coja-Oghlan, F. Krzakała, W. Perkins, and L. Zdeborová, "Information-theoretic thresholds from the cavity method," *Adv. Math.*, vol. 333, pp. 694–795, 2018. [Online]. Available: https://doi.org/10.1016/j.aim.2018.05.029

[33] V. Bapst and A. Coja-Oghlan, "The condensation phase transition in the regular k-SAT model," in *Approximation, randomization, and combinatorial optimization. Algorithms and techniques*, ser. LIPIcs. Leibniz Int. Proc. Inform. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2016, vol. 60, pp. Art. No. 22, 18.

[34] Z. Bartha, N. Sun, and Y. Zhang, "Breaking of 1RSB in random MAX-NAE-SAT," *arXiv preprint, arXiv:1904.08891*, 2019.

[35] A. Montanari and F. Ricci-Tersenghi, "On the nature of the low-temperature phase in discontinuous mean-field spin glasses," *The European Physical Journal B - Condensed Matter and Complex Systems*, vol. 33, no. 3, pp. 339–346, Jun 2003.

[36] F. Krzakała, A. Pagnani, and M. Weigt, "Threshold values, stability analysis, and high-q asymptotics for the coloring problem on random graphs," *Phys. Rev. E*, vol. 70, p. 046705, Oct 2004.

[37] M. Molloy and R. Restrepo, "Frozen variables in random boolean constraint satisfaction problems," in *Proceedings of the Twenty-fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, ser. SODA '13. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 2013, pp. 1306–1318. [Online]. Available: http://dl.acm.org/citation.cfm?id=2627817.2627912

[38] A. Gerschenfeld and A. Montanari, "Reconstruction for models on random graphs," in *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*, ser. FOCS '07. Washington, DC, USA: IEEE Computer Society, 2007, pp. 194–204. [Online]. Available: http://dx.doi.org/10.1109/FOCS.2007.58

[39] A. Montanari, R. Restrepo, and P. Tetali, "Reconstruction and clustering in random constraint satisfaction problems," *SIAM J. Discrete Math.*, vol. 25, no. 2, pp. 771–808, 2011. [Online]. Available: https://doi.org/10.1137/090755862

[40] L. Budzynski and G. Semerjian, "The asymptotics of the clustering transition for random constraint satisfaction problems," *Journal of Statistical Physics*, vol. 181, no. 5, pp. 1490–1522, 2020. [Online]. Available: https://doi.org/10.1007/s10955-020-02635-8

[41] A. Dembo and O. Zeitouni, *Large deviations techniques and applications*, ser. Stochastic Modelling and Applied Probability. Springer-Verlag, Berlin, 2010, vol. 38. [Online]. Available: https://doi-org.stanford.idm.oclc.org/10.1007/978-3-642-03311-7

[42] R. W. Robinson and N. C. Wormald, "Almost all cubic graphs are Hamiltonian," *Random Structures Algorithms*, vol. 3, no. 2, pp. 117–125, 1992. [Online]. Available: https://doi.org/10.1002/rsa.3240030202

[43] ——, "Almost all regular graphs are Hamiltonian," *Random Structures Algorithms*, vol. 5, no. 2, pp. 363–374, 1994. [Online]. Available: https://doi.org/10.1002/rsa.3240050209

[44] S. Janson, T. Ł uczak, and A. Rucinski, *Random graphs*, ser. Wiley-Interscience Series in Discrete Mathematics and Optimization. Wiley-Interscience, New York, 2000. [Online]. Available: https://doi.org/10.1002/9781118032718

[45] A. Sly, "Computational transition at the uniqueness threshold," in *Proceedings of the 2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, ser. FOCS '10. Washington, DC, USA: IEEE Computer Society, 2010, pp. 287–296.

[46] A. Galanis, D. Štefankovič, and E. Vigoda, "Inapproximability for antiferromagnetic spin systems in the tree nonuniqueness region," *J. ACM*, vol. 62, no. 6, pp. Art. 50, 60, 2015. [Online]. Available: https://doi.org/10.1145/2785964

[47] ——, "Inapproximability of the partition function for the antiferromagnetic ising and hard-core models," *Combinatorics, Probability and Computing*, vol. 25, no. 4, pp. 500–559, 2016.