On the Concentration of the Number of Solutions of Random Satisfiability Formulas

Emmanuel Abbe,[1] Andrea Montanari[2]
[1] Department of Electrical Engineering and Program in Applied and Computational Mathematics, Princeton University; e-mail: eabbe@princeton.edu
[2] Department of Electrical Engineering and Department of Statistics, Stanford University; e-mail: montanari@stanford.edu

Received 13 December 2010; revised 17 May 2012; accepted 18 February 2013
Published online 12 April 2013 in Wiley Online Library (wileyonlinelibrary.com).
DOI 10.1002/rsa.20501

ABSTRACT: Let $Z(F)$ be the number of solutions of a random k-satisfiability formula F with n variables and clause density α. Assume that the probability that F is unsatisfiable is $O(1/\log(n)^{1+\delta})$ for some $\delta > 0$. We show that (possibly excluding a countable set of "exceptional" α's) the number of solutions concentrates, i.e., there exists a non-random function $\alpha \mapsto \phi_s(\alpha)$ such that, for any $\varepsilon > 0$, we have $Z(F) \in [2^{n(\phi_s-\varepsilon)}, 2^{n(\phi_s+\varepsilon)}]$ with high probability. In particular, the assumption holds for all $\alpha < 1$, which proves the above concentration claim in the whole satisfiability regime of random 2-SAT. We also extend these results to a broad class of constraint satisfaction problems. © 2013 Wiley Periodicals, Inc. Random Struct. Alg., 45, 362–382, 2014

Keywords: constraint satisfaction problems; satisfiability; counting; concentration; sharp threshold; interpolation method

1. INTRODUCTION AND MAIN RESULTS

Over the last twenty years, a considerable effort has been devoted to understanding the typical properties of random k-satisfiability (k-SAT) instances. This line of work was initially motivated by two surprising empirical discoveries. First of all, when the "clause density" (number of clauses per variable) crosses a critical threshold, the probability that a random instance is unsatisfiable increases sharply from 0 to 1. The critical clause density is usually referred to as the "satisfiability threshold" and depends of course on k. Second, the

Correspondence to: E. Abbe

typical running time of standard solvers peaks in proximity of the satisfiability threshold. A significant amount of work has been devoted to understanding and explaining these phenomena. In particular, an important motivation has been the hope to develop better heuristics to cope with empirically hard constraint satisfaction instances, that are generated via the satisfiability threshold.

Significant progress has been made along this path. In particular, it was estabilished early on by Friedgut that, indeed, the probability that a random instance is unsatisfiable has a sharp threshold as the number of variables increases [14]. This phenomenon is referred to as the "satisfiability phase transition." A key question is however left unresolved by Friedgut's theorem, namely whether the critical density[1] converges to a limit or not as the number of variables increases. Numerical simulations as well as heuristic arguments strongly point towards the first alternative [17, 21, 23]. In particular for 3-SAT, it is conjectured that the critical threshold is approximately 4.2. Despite such consensus no real strategy has been put forward to prove this outstanding conjecture. Upper and lower bounds are known to match up to a term that is of relative order $k\,2^{-k}$ as k increases [4]. Nevertheless, they are still far for any given k, and pushing the same arguments is unlikely to lead to matching upper and lower bounds. In particular, the "condensation" phenomenon studied in [18] is suggestive of fundamental obstructions to the second moment method.

On the other hand, a significantly more detailed picture has been conjectured, building on non-rigorous techniques from statistical physics such as the replica and cavity methods [18, 21–23]. In particular, not only an n-independent critical density is conjectured to exist, but explicit values (depending on k) were computed in [22]. It is instructive to recall some basic steps of this non-rigorous but highly sophisticated calculations. As usual in statistical mechanics, the starting point consists in defining an appropriate partition function. For a given formula F, and a truth assignment $x \in \{+1, -1\}^n$, let $U(x; F)$ count the number of clauses in F that are not satisfied by assignment x. For $\beta \geq 0$, the partition function is defined as

$$\mathcal{Z}(\beta; F) = \sum_{x \in \{+1,-1\}^n} \exp\{-\beta\, U(x; F)\}\,. \tag{1.1}$$

In particular if F is satisfiable $\mathcal{Z}(\infty; F) \equiv Z(F)$ is well defined and counts the number of satisfying assignments. Using the non-rigorous cavity method, the limit of $(1/n) \log \mathcal{Z}(\beta; F)$ as $n \to \infty$ is then computed (for random formulas with a clause density α) and found to be a non-random function $\Phi(\alpha; \beta)$. In particular, it is found that there exists a threshold density $\alpha_s(k)$ such that $\phi_s(\alpha) \equiv \lim_{\beta\to\infty} \Phi(\alpha; \beta)$ remains finite for $\alpha < \alpha_s(k)$ and diverges for $\alpha > \alpha_s(k)$. Since the number of satisfying assignments should be $Z(F) \doteq e^{n\phi_s}$, the value $\alpha_s(k)$ is identified with the satisfiability threshold.

Notice that the statistical mechanics approach is quite ambitious. Instead of establishing the existence of a threshold independent of the number of variables, it aims of computing exactly the exponential growth rate of the number of solutions. The present paper takes this philosophy seriously and, focusing on the crucial case $\beta = \infty$, tackles a basic conjecture at the heart of the statistical mechanics approach. To be precise, let $F(n, \alpha)$ denote a formula on n variables with clause density α (each clause having k literals) and hence let $Z(F(n, \alpha))$ be the number of solutions of $F(n, \alpha)$. The basic conjecture of the physics approach states

[1] This can be defined for instance as the density such that the probability that a random formula is satisfiable is equal to $1/2$.

that $Z(F(n,\alpha))$ concentrates on the exponential scale. Namely, for each $\alpha < \alpha_s(k)$ (the satisfiability threshold), there exists $\phi_s = \phi_s(\alpha)$ non-random such that, for any $\varepsilon > 0$, $2^{n(\phi_s-\varepsilon)} \le Z(F(n,\alpha)) < 2^{n(\phi_s+\varepsilon)}$ with high probability [18, 29]. In formula, there exists ϕ_s such that for any $\varepsilon > 0$

$$\lim_{n\to\infty} \mathbb{P}\{2^{n(\phi_s-\varepsilon)} \le Z(F(n,\alpha)) < 2^{n(\phi_s+\varepsilon)}\} = 1\,. \tag{1.2}$$

Notice that this conjecture would imply in particular the existence of an n-independent satisfiability threshold.

One fundamental difficulty in estabilishing (1.2) is that the concentration of $\log Z(F)$ cannot be proved using standard martingale methods. Such an argument typically requires to control the difference $|\log Z(F') - \log Z(F)|$ for F and F' differing in a single clause [20]. Unfortunately, adding a single clause can change the value of $Z(F)$ from exponentially large to $Z(F') = 0$ (we refer to the next section for further discussion on this point).

In this work, we prove the conjecture (1.2) for $k = 2$. Note that for $k = 2$, the satisfiability threshold is known to be 1 as proven in [7, 10, 15]. For arbitrary $k \ge 3$ we cannot establish the conjecture for all α below the satisfiable threshold. On the other hand, we are able to prove that (1.2) holds for any α such that $\mathbb{P}\{Z(F(n,\alpha)) = 0\}$ (the unsatisfiability probability) is upper bounded by $1/(\log n)^{1+\delta}$ for some $\delta > 0$ and all n large enough. In particular, this establishes the conjecture for $k \ge 3$ and $\alpha < 1$, which represents only a minor portion of the satisfiable phase since for random k-SAT the threshold grows like $2^k \log 2 - O(k)$ [5]. In fact, we expect that the condition on $\mathbb{P}\{Z(F(n,\alpha)) = 0\}$ holds up to the satisfiability threshold. Partial evidence on the basis of constructive satisfiability proofs is discussed below.

Let us briefly sketch the main ideas in the proof.

The first remark is that conjecture (1.2) amounts to saying that the monotone property $\mathcal{E}_n(\alpha,\phi) \equiv \{Z(F(n,\alpha)) \le 2^{n\phi}\}$ undergoes a sharp threshold in ϕ. Namely, its probability increases from close to 0 to close to 1 over a window in ϕ that shrinks to 0 as n goes to infinity. On the other hand, Friedgut's theorem implies quite straightfowardly that the same property $\mathcal{E}_n(\alpha,\phi)$ undergoes a sharp threshold in α (with a critical density that might depend on n); this was proved in [3], and in [25] for a generalized class of satisfiability models.

The real challenge to prove the conjecture is to show that this threshold phenomenon for fixed ϕ and α varying implies a threshold phenomenon with fixed α and ϕ varying as in Eq. (1.2). The difficulty originates from the limited control that Friedgut theorem implies about the thresholds, and namely the fact that a priori, these need not to converge as n increases. In this paper, we show that indeed the threshold (in α) for the property $\mathcal{E}_n(\alpha,\phi) \equiv \{Z(F(n,\alpha)) \le 2^{n\phi}\}$ does converge, under the mentioned hypothesis on $\mathbb{P}\{Z(F) = 0\}$.

To achieve this, we aim at showing that $(1/n)\mathbb{E}\log(1 + Z(F))$ converges in the limit $n \to \infty$. The limit value provides a candidate for the quantity ϕ_s entering Eq. (1.2).

A powerful tool to show similar convergence results in the context of spin glass theory is the interpolation method first developed by Francesco Guerra and Fabio Toninelli [16]. This method is normally applied to expected log-partition functions. In our case a naive application would make use of $(1/n)\mathbb{E}\log(Z(F))$. The interpolation method is then used to show a superadditivity (or sub-additivity) property on the sequence of interest. Despite its simplicity, the interpolation method appears to be surprisingly powerful and opened the way to a solution of several open problems so far (we refer to the next section for a brief overview).

In our context, the naive application fails, for an interesting reason. The difficulty is related to the very fact that we are considering a random constraint satisfaction problem,

and $Z(F(n,\alpha)) = 0$ with positive probability. Hence the expected log-number of solutions carries no information: $(1/n)\mathbb{E}\log(Z(F(n,\alpha))) = -\infty$. Our solution consists in considering the sequence $(1/n)\mathbb{E}\log(1 + Z(F(n,\alpha)))$. This allows to work with a well-defined quantity but one has to pay a price for this apparently innocuous modification. The interpolation method leads to a pseudo superadditivity property, namely $\mathbb{E}\log(1 + Z(F(n,\alpha))) \geq \mathbb{E}\log(1 + Z(F(n_1,\alpha))Z(F(n_2,\alpha)))$, when $n_1 + n_2 = n$.

The last part of the proof aims at extracting the consequences of the pseudo superadditivity property by controlling the effect of the "1+" term. In particular, we show that the pseudo superadditivity property can still be used to conclude convergence of $(1/n)\mathbb{E}\log(1 + Z(F(n,\alpha)))$, provided that $\mathbb{P}\{Z(F(n,\alpha)) = 0\}$ decays fast enough, i.e., as $O(1/\log(n)^{1+\delta})$.

This probability decay condition holds for random 2-SAT, which concludes the proof of conjecture (1.2) in this case. For other values of k, we obtain a conditional result. By comparison with 2-SAT, we obtain that $\mathbb{P}\{Z(F(n,\alpha)) = 0\} = O(1/n)$ for any $k \geq 2$ and $\alpha < 1$. Analysis of search algorithms can be used to get estimates on the unsatisfiability probability and extend the range of α's for which the decay condition holds beyond the interval $[0, 1)$. In particular, [8] could be used for that purpose. The most general result would be obtained by establishing a quantitative version of Friedgut's theorem to verify the decay condition.

Finally, we generalize in Section 4 the results obtained for k-SAT to a broad family of constraint satisfaction problems, including hypergraph 2-colorability, NAE k-SAT and k-XOR-SAT.

The rest of the paper is organized as follows. Section 2 discusses related work on the subject. Sections 3 and 4 state our main results for –respectively– k-satisfiability and a general class of constraint satisfaction problems. Proofs are presented in Sections 5 and 6.

2. RELATED WORK

From an algorithmic point of view, the problem of computing the number of solutions of a k-satisfiability formula is well known to be #P-complete for any $k \geq 2$; although 2-SAT is not NP-complete, #2-SAT (the problem of counting solutions for 2-SAT formulas) is #P-complete [28]. Even worse, there is no fully polynomial randomized approximation scheme (FPRAS) to approximate the number of solutions unless NP = RP [9].

Estimating the typical number of solutions for *random* satisfiability formulae has not been so far a major object of study within discrete mathematics. The only result in this direction is the paper [26] which computes the limit of $(1/n)\log \mathcal{Z}(\beta; F(n,\alpha))$ for all $\beta < \infty$ and all $\alpha \leq \alpha_{\mathrm{pl}}(k) = (2k^{-1}\log k)(1 + o_k(1))$. In the case of random 2-SAT, an unpublished result of A. Sharell cited in [11] claims a partial concentration result, which shows that the logarithm of the number of solutions concentrates around its expectation. However, it is mentioned in [11] that Sharell's result does not provide any information on the convergence of this expectation as $n \to \infty$.

On the other hand, as mentioned in the introduction, computing the log-partition function is the first step in all statistical mechanics calculations. It is worth mentioning that statistical mechanics analysis also lead to an intriguing picture of the geometry of the set of solutions of a random satisfiability instance [18, 21, 22]. While only a small subset of these results have been estabilished rigorously, they provided guidance and stimulus for exciting rigorous developments [2, 3].

A key tool in our analysis is the interpolation method first introduced in [16] for the Sherrington-Kirkpatrick model. This is a model for a spin-glass (i.e. a spin model with random couplings) on a complete graph. It was subsequently shown in [12, 13, 27] that the same ideas can be generalized to models on random sparse graphs. In particular, these papers prove[2] the existence of the limit of $(1/n)\mathbb{E}\log \mathcal{Z}(\beta; F(n,\alpha))$ for any $\beta < \infty$. Denoting by $U_*(F) \equiv \min_x U(x;F)$ the minimum number of unsatisfied clauses in the formula F, this also implies the existence of the limit $\lim_{n\to\infty} U_*(F(n,\alpha))/n$ by a standard argument.

The generalization to sparse random graphs opened the way to applications within coding theory [19, 24]. While the present paper was in preparation, applications of the interpolation method to other combinatorial models were developed independently by Bayati, Gamarnik and Tetali [6]. The two developments were concurrent.

Let us however stress that none of these papers deals with the key challenge posed by constraint satisfaction problems, namely that the number of solutions $Z(F(n,\alpha))$ vanishes with positive probability and hence $(1/n)\mathbb{E}\log Z(F(n,\alpha)) = -\infty$. For instance in the case of k-satisfiability, for any positive $\beta < \infty$, we have $2^n \geq \mathcal{Z}(\beta; F(n,\alpha)) \geq 2^n \exp\{-\beta(n\alpha)\}$ (because $0 \leq U(x;F) \leq n\alpha$) and therefore

$$\log 2 \geq (1/n)\mathbb{E}\log \mathcal{Z}(\beta; F(n,\alpha)) \geq \log 2 - \beta\alpha\,. \tag{2.1}$$

Hence superadditivity can be proved directly for the log-partition function, which simplifies the problem. Further, concentration on the exponential scale is immediate as well in those cases via bounded martingale methods. Consider again as an example the partition function (1.1). For the sake of simplicity, it is useful to consider the case in which a fixed number m of clauses is drawn uniformly at random, and call $F(n,m)$ the corresponding formula. It is then sufficient to bound the martingale differences, which follows from bounding the effect of adding one single clause. A simple calculation yields

$$\log \mathcal{Z}(\beta; F(n,m+1)) - \log \mathcal{Z}(\beta; F(n,m)) \\ = \log\{1 - (1-e^{-\beta})\mathbb{P}_{m,\beta}((m+1)\text{-th clause is not satisfied})\}\,,$$

where $\mathbb{P}_{m,\beta}$ denotes the probability distribution over truth assignments $x \in \{+1,-1\}^n$ given by $\mathbb{P}_{m,\beta}(x) \propto \exp\{-\beta U(x;F(n,m))\}$. From this it follows immediately that $|\log \mathcal{Z}(\beta; F(n,m+1)) - \log \mathcal{Z}(\beta; F(n,m))| \leq \beta$, hence the concentration of $\log \mathcal{Z}(\beta, F)$. In our case, these expansions do not apply since $\beta = \infty$.

The fact that the difference $|\log Z(F(n,m+1)) - \log Z(F(n,m))|$ is unbounded is intimately related to the structure of the set of solutions of $F(n,m)$. Consider an extreme case. If a fraction of all the variables takes the same value in all the solutions, then this quantity can be infinite. Indeed, the new clause can constrain k variables of this type, and therefore be violated by all solutions of $F(n,m)$. This leads to $Z(F(n,m+1) = 0$. This phenomenon is only possible at the threshold for that formula. Notice however that a much more frequent situation can lead to unbounded differences $|\log Z(F(n,m+1)) - \log Z(F(n,m))|$. Indeed, it is sufficient that the value of some variables is very biased when a solution is drawn uniformly at random. This phenomenon occurs at any positive α. Applying the martingale method would require controlling *a priori* the biases of all the variables, a task that is likely to be quite challenging.

[2]These paper consider the case of k even, but it was noticed early on by Elitza Maneva that the proof applies verbatimly to k odd as well.

3. RANDOM *K*-SAT

Definition 1. A k-clause is a disjunction of k Boolean variables or their negations. Let $C_k(n)$ be the set of all $N = \binom{n}{k}2^k$ possible k-clauses on n Boolean variables. We denote by $F_k(n,\alpha)$ a random formula which is formed by selecting independently each element in $C_k(n)$ with probability $p_k(n,\alpha) = \alpha n/N$, and by taking the conjunction of the selected clauses.

The number of clauses in the above model is a binomial random variable, which concentrates exponentially fast around its expectation αn. Some of our computations prove to be simpler within slightly different models, whereby the number of clauses is either Poisson or deterministic with the same mean αn. Standard monotonicity arguments can be used to show the equivalence of these models for our purposes and we will hence switch freely between these different models. Unless specified, the value of k will remain fixed throughout the paper and the k subscript is dropped.

Definition 2. We denote by $Z(F)$ the number of satisfying assignments (solutions) of a Boolean formula F and by

$$P_n(\alpha,\phi) := \mathbb{P}\{Z(F(n,\alpha)) < 2^{n\phi}\},$$

the probability that a random formula has fewer than $2^{n\phi}$ satisfying assignments.

Definition 3. We define

$$\alpha^* := \sup\{\alpha : P_n(\alpha,0) = O(1/(\log n)^{1+\delta}), \text{ for some fixed } \delta > 0\}.$$

Note that $P_n(\alpha,0) = \mathbb{P}\{Z(F(n,\alpha)) = 0\}$ is the probability that $F(n,\alpha)$ is unsatisfiable (UNSAT).

Remark 1. We expect the forthcoming results to hold when defining α^* to be $\sup\{\alpha : \sum_n P_n(\alpha,0)/n < \infty\}$. However, the above definition simplifies the proofs without much loss of generality.

Remark 2. For any $k \geq 2$, we have $\alpha^* \geq 1$. Indeed, considering the case of 2-SAT, [7, 10, 15] show that, for $\alpha < 1$, $\mathbb{P}\{Z(F_2(n,\alpha)) = 0\} = O(1/n)$ (cf. the last equality of the proof of Theorem 2.7 page 475 in [15]). Since for any $k \geq 3$, $\mathbb{P}\{Z(F_k(n,\alpha)) = 0\} \leq \mathbb{P}\{Z(F_2(n,\alpha)) = 0\}$, we conclude $\alpha^* \geq 1$ for $k \geq 3$ as well.

Unfortunately, the bounds on the satisfiability threshold based on the second moment method [4] do not imply any quantitative estimate on the probability that $F_k(n,\alpha)$ is UNSAT. It might be possible to prove such an estimate by a careful analysis of specific solution algorithms. In particular, a careful analysis of the recent algorithm [8] might lead to a proof of $\alpha^* \geq 2^k(1-\delta)\log k/k$, for k large enough (Coja-Oghlan, personal communication). We expect that α^* does coincide with the satisfiability threshold.

Our main result estabilishes the conjecture (1.2) for $\alpha < \alpha^*$, apart possibly for countably many values of α.

Theorem 1. *For every $k \geq 2$, there exist a countable set $\mathcal{C}$ and a function $\phi_s : [0, \alpha^*) \to [0, 1]$ where $\alpha^* \equiv \sup\{\alpha : \mathbb{P}\{Z(F_k(n, \alpha)) = 0\} = O(1/(\log n)^{1+\delta}),$ for some fixed $\delta > 0\}$ such that for every $\alpha \in [0, \alpha^*) \setminus \mathcal{C}$ and every $\varepsilon > 0$*

$$\lim_{n\to\infty} \mathbb{P}\{Z(F_k(n, \alpha)) < 2^{n(\phi_s(\alpha)-\varepsilon)}\} = 0,$$
$$\lim_{n\to\infty} \mathbb{P}\{Z(F_k(n, \alpha)) < 2^{n(\phi_s(\alpha)+\varepsilon)}\} = 1.$$

The function ϕ_s is decreasing with $\phi_s(0) = 1$ and its definition is given in Theorem 3 below. Hereafter, we will use the lighter notation introduced in Definition 2.

Remark 2 directly gives the following corollary.

Corollary 1. *For $k = 2$, i.e., for random 2-SAT, there exists a function $\phi_s : [0, 1) \to [0, 1]$ such that for almost every α in the satisfiable phase $[0, 1)$ and every $\varepsilon > 0$*

$$\lim_{n\to\infty} P_n(\alpha, \phi_s(\alpha) - \varepsilon) = 0,$$
$$\lim_{n\to\infty} P_n(\alpha, \phi_s(\alpha) + \varepsilon) = 1.$$

Theorem 1 is proved in Section 5 as a result of the following sharp threshold result, with a *n-independent* threshold.

Theorem 2. *For every $k \geq 2$, there exists a function $\phi_s : [0, \alpha^*) \to [0, 1]$ and a countable set $\mathcal{D}$ such that for every $\phi \in \phi_s([0, \alpha^*)) \setminus \mathcal{D}$, there exists $\alpha_s(\phi)$ such that for every $\varepsilon > 0$*

$$\lim_{n\to\infty} P_n(\alpha_s(\phi) - \varepsilon, \phi) = 0,$$
$$\lim_{n\to\infty} P_n(\alpha_s(\phi) + \varepsilon, \phi) = 1\,.$$

This theorem, proved in Section 5, is based on the subsequent results. First, the following sharp threshold result with a *n*-dependent threshold, which follows from Friedgut's theorem.

Lemma 1. *For every $k \geq 2$, for every $\phi \in [0, 1)$, there exists $\{\alpha_n(\phi)\}_{n\in\mathbb{Z}_+}$ such that for every $\varepsilon > 0$*

$$\lim_{n\to\infty} P_n(\alpha_n(\phi) - \varepsilon, \phi) = 0,$$
$$\lim_{n\to\infty} P_n(\alpha_n(\phi) + \varepsilon, \phi) = 1\,.$$

This lemma is proved in [3, Lemma 13]. An alternative proof is also available in [1]. Lemma 1 says that for any fixed $\phi \in [0, 1)$, the property $\{Z(F(n, \alpha)) < 2^{n\phi}\}$ has a sharp threshold in α. As for the $\phi = 0$ case, Lemma 1 is proved by showing that the monotone property $\{Z(F(n, \alpha)) < 2^{n\phi}\}$ cannot be approximated by a "local property" in the sense of Theorem 5.2 in [14], and hence must have a sharp threshold.

In order to prove Theorem 2 using Lemma 1, the key result is the following convergence theorem.

Theorem 3. *For every $k \geq 2$, for every $\alpha < \alpha^*$, the sequence*

$$\psi_n(\alpha) := \frac{1}{n}\mathbb{E}[\log Z(F(n, \alpha))|Z(F(n, \alpha)) \geq 1]$$

converges to a limit $\phi_s(\alpha)$.

To prove Theorem 3, which is done in Section 5, we first prove the following pseudo superadditivity property.

Lemma 2. *For every $k \geq 2$, for every $\alpha \geq 0$, let $Z_n := Z(F(n,\alpha))$, we then have for any $n_1, n_2 \geq 1$*

$$\mathbb{E}\log(1 + Z_{n_1+n_2}) \geq \mathbb{E}\log(1 + Z_{n_1} Z_{n_2}),$$

where on the right hand side Z_{n_1} and Z_{n_2} are understood to be the numbers of solutions of two independent random formulae on respectively n_1 and n_2 variables.

Lemma 2 is a specific case of Lemma 6 presented in the next section and proved in Section 6. It is based on the interpolation technique by Guerra and Toninelli [16], and Franz-Leone [12]. However, while in those cases one obtains superadditivity of $\mathbb{E}\log Z$, in the present case we get a weaker result because of the "1+" term. This problem comes from the fact that $Z = 0$ with positive probability, and therefore $\mathbb{E}\log Z$ is not defined. For α in the satisfiable phase Z_n is "typically" large, and one would expect that the effect of the "1+" term is negligible so that the pseudo superadditivity property can be approximated by a superadditivity property and hence imply a convergence result. This is made rigorous by requiring a decay condition on the probability of being UNSAT and using a technical lemma on the convergence of "approximately" superadditive sequences (Lemma 7 in Section 5).

4. A GENERAL FAMILY OF RANDOM CSP'S

In this section, we extend the results of the previous sections to a general family of random constraint satisfaction problems (CSP) over binary variables. An ensemble in this family is defined as follows (the definition is analogous to the one in [25]).

Definition 4. Let μ be a distribution over Boolean functions $\varphi : \{-1,+1\}^k \to \{0,1\}$, which we call the *clause type distribution*. Let n be an integer and $\alpha \in \mathbb{R}_+$. A random formula from the ensemble $F_k(n,\alpha,\mu)$ is drawn as follows. For each $a \in \{1,\ldots,m = \lfloor \alpha n \rfloor\}$ the a-th clause is drawn independently from previous ones. For clause a, k indices $i_1(a),\ldots,i_k(a)$ are drawn independently and uniformly at random in $[n]$ (i.e., indices are picked with replacement). Further $\varphi_a : \{-1,+1\}^k \to \{0,1\}$ is drawn under the distribution μ, producing the clause $\varphi_a(x_{i_1(a)},\ldots,x_{i_k(a)})$.

An assignment $x \in \{+1,-1\}^n$ is said to *satisfy* the formula $F_k(n,\alpha,\mu)$ if, for each $a \in [m]$, we have $\varphi_a(x_{i_1(a)},\ldots,x_{i_k(a)}) = 1$.

As in previous section, we will often drop the subscripts k in the following, $Z(F)$ denotes the number of satisfying assignments of formula F and we define

$$P_n(\alpha,\phi,\mu) := \mathbb{P}\{Z(F(n,\alpha,\mu)) < 2^{n\phi}\}.$$

Definition 5. Note that $P_n(\alpha,0,\mu) = \mathbb{P}\{Z(F(n,\alpha,\mu)) = 0\}$ is the probability that $F(n,\alpha,\mu)$ is UNSAT. We define

$$\alpha^*(\mu) := \sup\{\alpha :\ P_n(\alpha,0,\mu) = O(1/(\log n)^{1+\delta}) \text{ for some fixed } \delta > 0\}.$$

Definition 6. For $\varphi : \{-1,1\}^k \to \{0,1\}$ and $\theta \in [-1,1]$, let

$$\|\varphi\|_\theta^2 \equiv \sum_{x\in\{-1,1\}^k} \varphi(x)^2 v_\theta(x) \quad \text{and} \quad \|\varphi\| \equiv \|\varphi\|_0$$

where

$$v_\theta(x) \equiv \prod_{i=1}^{k} \frac{1+x_i\theta}{2}.$$

Note that $\|\varphi\|_\theta^2$ is the probability that $\varphi = 1$ under the measure v_θ, which assigns probabilities $(1-\theta)/2$ and $(1+\theta)/2$ to -1 and $+1$ respectively.

Our CSP ensemble is specified by the distribution μ over clause types, we now describe two set of hypotheses on this distribution.

H1. (a) *Dominance of balanced assignments.* For every $\theta \in [-1,1]$, $\mathbb{E}_\varphi \log \|\varphi\|_\theta \le \mathbb{E}_\varphi \log \|\varphi\|$, with equality only if $\theta = 0$. This condition implies that, in a typical random instance, most solutions have almost as many 1's as -1's.
(b) *Unsatisfiability of uniform assignments.* For every $s \in \{-1,+1\}$, there is at least one clause φ with $\mu(\varphi) > 0$ such that $\varphi(s,\dots,s) = 0$.
(c) *Balance property.* The distribution μ is supported on Boolean functions such that $\varphi(x_1,\dots,x_k) = \varphi(-x_1,\dots,-x_k)$. This condition implies that the odd Fourier coefficients of φ are zero.

Hypothesis (c) may not be necessary to establish the forthcoming results but is kept for consistency with [25].

H2. *Convexity of* Γ_l. Let $M_1(\{-1,1\}^l)$ be the set of probability measures on $\{-1,1\}^l$ and for a given $\nu \in M_1(\{-1,1\}^l)$, let $\{Z_i^{(1)},\dots,Z_i^{(l)}\}_{i=1}^k$ be i.i.d. under ν. We say that μ satisfies H2 if for a random clause type φ drawn under μ, the mapping

$$\Gamma_l : M_1(\{-1,1\}^l) \to \mathbb{R}$$

defined by

$$\Gamma_l(\nu) := \mathbb{E}_\varphi \mathbb{P}\{\varphi(Z^{(r)}) = 0, \forall 1 \le r \le l\} \tag{4.1}$$

is convex for any $l \ge 1$.

Notice that conditions H1.(a), H1.(b) and H1.(c) coincide respectively with conditions 4, 5 and 2 in [25]. Further, hypothesis H1.(a) and H1.(b) are satisfied by a number of interesting random CSP ensembles. Such examples include

- k-NAE-SAT, where $\varphi(x) = \varphi_s(x) = \mathbb{1}(x \notin \{-s,s\})$ and $\mu(\varphi_s) = 2^{-k}$ for each $s \in \{-1,1\}^k$;
- Hypergraph 2-coloring, where $\varphi(x) = \mathbb{1}(x \notin \{\underline{-1};\underline{+1}\})$ is the unique clause in the support of μ, with $\underline{-1} = (-1,\dots,-1)$ and $\underline{+1} = (+1,\dots,+1)$;
- k-XOR-SAT, where $\varphi(x) = \varphi_s(x) = \mathbb{1}(\prod_{i=1}^k x_i = s)$ and $\mu(\varphi_s) = 1/2$ for each $s \in \{-1,1\}$;
- k-SAT, where $\varphi(x) = \varphi_s(x) = \mathbb{1}(x \ne s)$ and $\mu(\varphi_s) = 2^{-k}$ for each $s \in \{-1,1\}^k$.

For the first three examples above, it is verified in [25] that hypotheses H1.(a) and H1.(b) are satisfied. Let us verify that this is also the case for k-SAT. Note that

$$\mathbb{E}_s\|\varphi\|_\theta^2 = 1 - \mathbb{E}_s\prod_{i=1}^{k}\frac{1-s_i\theta}{2} = 1 - 2^{-k} = \mathbb{E}_s\|\varphi\|^2,$$

hence

$$\mathbb{E}_s \log\|\varphi\|_\theta^2 \le \log\mathbb{E}_s\|\varphi\|_\theta^2 = \log\mathbb{E}_s\|\varphi\|^2 = \mathbb{E}_s\log\|\varphi\|^2.$$

This verifies condition H1.(a). Condition H1.(b) holds trivially. Condition H1.(c) holds for the first three CSP ensembles above as shown in [25], but H1.(c) does not hold for random k-SAT.

Hypothesis H2 is not straightforward to check, and it is not investigated in [25]. The next definition provides a family of clause type distributions satisfying it.

Definition 7 (k-factorizing distributions). A clause type distribution μ is said to *k-factorize* if it has the following structure. There exists an integer $J \ge 1$, such that any $\varphi \in \mathrm{supp}(\mu)$ is of the form

$$\varphi(x) = \mathbb{1}(x \notin \{s^{(1)}, \dots, s^{(J)}\}), \tag{4.2}$$

for some $s^{(1)}, \dots, s^{(J)} \in \{-1, 1\}^k$, and

$$\mu(\varphi) = \prod_{i=1}^{k}\bar{\mu}(s_i^{(1)}, \dots, s_i^{(J)}) \tag{4.3}$$

where $\bar{\mu}$ is a probability distribution on $\{-1, 1\}^J$. In other words, the vectors $(s_i^{(1)}, \dots, s_i^{(J)})$, for $i = 1, \dots, k$, can have correlated components but are mutually i.i.d. with distribution $\bar{\mu}$.

This definition can be generalized by letting J itself to be random, but we stick to the above case for the sake of simplicity.

The class of k-factorizing clause type distributions includes, among other problems:

- k-NAE-SAT: $\varphi(x) = \mathbb{1}(x \notin \{-s, s\})$ for $s \in \{-1, +1\}^k$ uniformly random. This is k-factorizing with $\bar{\mu}(-1, 1) = \bar{\mu}(1, -1) = 1/2$;
- Hypergraph 2-coloring: $\varphi(x) = \mathbb{1}(x \notin \{\underline{-1}, \underline{+1}\})$ with $\bar{\mu}(-1, 1) = 1$;
- k-SAT: $\varphi(x) = \mathbb{1}(x \notin \{s\})$ with $\bar{\mu}(1) = \bar{\mu}(-1) = 1/2$.

Condition H2 is satisfied by k-factorizing distributions as stated formally below.

Lemma 3. *If the clause type distribution k-factorizes, then the mapping Γ_l is convex for any $l \ge 1$.*

Note that k-XOR-SAT does not belong to this class of distributions, nevertheless, condition H2 holds in this case as well when k is even, as stated below.

Lemma 4. *The mapping Γ_l is convex for any $l \ge 1$ for k-XOR-SAT with k even.*

The proofs of Lemma 3 and Lemma 4 are differed to Section 6. We now state the equivalent of Theorem 1 for this general class of CSPs.

Theorem 4. *Assume μ to satisfy conditions H1 and H2. Then there exists a countable set $\mathcal{C}$ and a function $\alpha \mapsto \phi_s(\alpha)$ such that, for every $\alpha \in [0, \alpha^*(\mu)) \setminus \mathcal{C}$ and every $\varepsilon > 0$,*

$$\lim_{n\to\infty} P_n(\alpha, \phi_s(\alpha) - \varepsilon, \mu) = 0,$$
$$\lim_{n\to\infty} P_n(\alpha, \phi_s(\alpha) + \varepsilon, \mu) = 1.$$

Remark 3. By the same argument in Remark 2, we have $\alpha^* \geq 1$ for k-NAE-SAT. In fact, for each NAE-SAT formula, we can construct an associated SAT formula by forbidding only one of the two assignments s and $-s$ in each clause. Hence, the UNSAT probability for k-NAE-SAT is upper bounded by the UNSAT probability in k-SAT at the same value of α.

Remark 4. It is proved in [30] that k-hypergraph 2-coloring is satisfiable with probablity $1 - O(n^{-1/2})$ provided $\alpha \leq 2^k/(50k)$ (cf. Theorem 1, and Claim 1). It follows that in this case $\alpha^* \geq 2^k/(50k)$.

As in previous section, the proof of this theorem relies on the following two results.

Lemma 5. *For any μ satisfying H1 and $\phi \in [0, 1)$, there exists $\{\alpha_n(\phi)\}_{n\in\mathbb{Z}_+}$ such that for every $\varepsilon > 0$,*

$$\lim_{n\to\infty} P_n(\alpha_n(\phi) - \varepsilon, \phi, \mu) = 0,$$
$$\lim_{n\to\infty} P_n(\alpha_n(\phi) + \varepsilon, \phi, \mu) = 1.$$

This lemma is proved in [25, Lemma C.2].

Theorem 5. *Let*

$$\psi_n(\alpha) := \frac{1}{n}\mathbb{E}[\log Z(F(n, \alpha, \mu))|Z(F(n, \alpha, \mu)) \geq 1].$$

For any μ satisfying H2 and for any $\alpha < \alpha^*(\mu)$, $\psi_n(\alpha)$ converges to a limit $\phi_s(\alpha)$.

The proof of this theorem is based on the following pseudo superadditivity lemma.

Lemma 6. *For any α and μ satisfying H2, let $Z_n := Z(F(n, \alpha, \mu))$. For any $n_1, n_2 \geq 1$,*

$$\mathbb{E}\log(1 + Z_{n_1+n_2}) \geq \mathbb{E}\log(1 + Z_{n_1}Z_{n_2}),$$

where on the right hand side Z_{n_1} and Z_{n_2} are understood to be the numbers of solutions of two independent random formulas on respectively n_1 and n_2 variables.

The proof of Lemma 6 is differed to Section 6. The proofs of Theorem 5 and Theorem 4 follow the same analytical arguments as for the k-SAT case and hence we omit the details here and refer to the proofs of Theorem 3 and Theorem 1.

5. PROOFS OF THEOREM 1, THEOREM 2 AND THEOREM 3

We first need the following technical lemma.

Lemma 7. *Let $\Delta(n) = O(n/(\log n)^{1+\delta})$ for some $\delta > 0$, and $t(n) = o(n)$. Let $f(\cdot)$ be positive, such that $f(n)/n$ is bounded above and*

$$f(n_1 + n_2) + \Delta(n_1 + n_2) \geq f(n_1) + f(n_2), \quad \forall n_1, n_2 \geq t(n_1 + n_2).$$

Then $f(n)/n$ converges.

Remark 5. We expect this lemma to hold if $\Delta(n)$ is such that $\sum_n \frac{\Delta(n)}{n^2} < \infty$.

Proof of Lemma 7. Let $\varepsilon > 0$. Since $f(n)/n$ is bounded above, we have $S := \limsup_n f(n)/n < \infty$. Let n_0 be large enough such that $f(n)/n < S + \varepsilon$ for any $n \geq n_0$. Let r be an integer (to be chosen at our convenience but to be fixed independently of n) and let $\gamma(n)$ be such that $\gamma(n) \geq t(n)$ and $\gamma(n) = o(n)$. Define

$$m(n) = \inf\{m' \text{ such that } m' = r2^i \text{ for some } i \in \mathbb{Z}_+ \text{ and } m' \geq \gamma(n)\}$$

and note that $\gamma(n) \leq m(n) \leq \max(2\gamma(n), r)$. For any n, we have

$$n = \lfloor n/m(n) - 1 \rfloor m(n) + q(n),$$

where $q(n) \in [m(n), 2m(n)]$, with $\lfloor n/m(n) - 1 \rfloor \geq 0$ for n large enough. Hence, using the property of f, we have

$$f(n) \geq \lfloor n/m(n) - 1 \rfloor f(m(n)) + f(q(n)) - \sum_{i=1}^{\lfloor n/m(n)-1 \rfloor} \Delta(q(n) + im(n)),$$

and since f is positive,

$$\frac{f(n)}{n} \geq \frac{m(n)}{n} \lfloor n/m(n) - 1 \rfloor \frac{f(m(n))}{m(n)} - \frac{1}{n} \sum_{i=1}^{\lfloor n/m(n)-1 \rfloor} \Delta(q(n) + im(n)). \tag{5.1}$$

Since $\gamma(n) = o(n)$, for n large enough, we have

$$\frac{m(n)}{n} \lfloor n/m(n) - 1 \rfloor > 1 - \varepsilon.$$

We now show that for n large enough, we also have $\frac{f(m(n))}{m(n)} > S - \varepsilon$. First, note that we can take r large enough such that $r2^i/2 \geq \gamma(r2^i)$ for all $i \geq 0$, since $\gamma(n) = o(n)$. Hence, using the property of f, we have

$$\frac{f(r2^{i+1})}{r2^{i+1}} \geq \frac{f(r2^i)}{r2^i} - \frac{1}{r2^{i+1}} \Delta(r2^{i+1})$$

and by recursively using the last inequality

$$\frac{f(r2^i)}{r2^i} \geq \frac{f(r)}{r} - \sum_{j=1}^{i} \frac{1}{r2^j} \Delta(r2^j). \tag{5.2}$$

Since we can pick r at our convenience, note that if r is a power of 2,

$$\sum_{j=1}^{i} \frac{1}{r2^j} \Delta(r2^j) = \sum_{j=\log_2 r+1}^{\log_2 r+i} \frac{1}{2^j} \Delta(2^j),$$

which is, when r increases, tending to zero uniformly in i, provided that

$$\sum_{j=1}^{\infty} \frac{1}{2^j} \Delta(2^j) < \infty.$$

As previous condition follows from our hypothesis on Δ, and since we can always take r large enough to ensure that $f(r)/r > S - \varepsilon$, we can take r large enough such that, from (5.2), the following holds for any i

$$\frac{f(r2^i)}{r2^i} \geq S - \varepsilon$$

and since $m(n)$ is of the form $r2^i$, for any n

$$\frac{f(m(n))}{m(n)} \geq S - \varepsilon.$$

Finally, we need to show that the last term in (5.1) is vanishing, i.e., that

$$\frac{1}{n} \sum_{i=1}^{\lfloor n/m(n)-1 \rfloor} \Delta(q(n) + im(n)) \xrightarrow{n\to\infty} 0.$$

For this, we pick $\gamma(n)$ to be large enough. For $\Delta = O(n/(\log n)^{1+\delta})$, we have[3]

$$\frac{1}{n} \sum_{i=1}^{\lfloor n/m(n)-1 \rfloor} \Delta(q(n) + im(n)) \leq \frac{1}{n} \frac{n}{m(n)} \Delta(n), \tag{5.3}$$

and since $m(n) \geq \gamma(n)$, if $\gamma(n) = O(n/(\log n)^{1+\nu})$ with $\nu > \delta$, we conclude the proof. ∎

Proof of Theorem 3. The goal of the proof is to show that $\psi_n(\alpha)$ (or a related function) satisfies the hypothesis of Lemma 7.

Let $F_n = F(n, \alpha)$. For an event $\mathcal{A}$, we use the standard notation $\mathbb{E}[X, \mathcal{A}] \equiv \mathbb{E}[X\mathbb{1}_{\mathcal{A}}]$. Note that

$$\mathbb{E}\log(1 + Z(F_n)) = \mathbb{E}[\log(1 + Z(F_n)), Z(F_n) \geq 1]$$

and

$$\begin{aligned}&\mathbb{E}[\log(1 + Z(F_n)), Z(F_n) \geq 1] \\ &= \mathbb{E}[\log Z(F_n), Z(F_n) \geq 1] + \mathbb{E}[\log(1 + Z(F_n)^{-1}), Z(F_n) \geq 1].\end{aligned}$$

[3]In greater generality, one has to pick $\gamma(n)$ such that $\Delta(n)/\gamma(n) = o(1)$.

Let $c > 0$, we have

$$\mathbb{E}[\log(1 + Z(F_n)^{-1}), Z(F_n) \geq 1] \leq \mathbb{E}[Z(F_n)^{-1}, Z(F_n) \geq 1]$$
$$\leq \mathbb{E}[Z(F_n)^{-1}, Z(F_n) \geq 1, Z_{FV} \geq cn] + \mathbb{P}\{Z_{FV} < cn, Z(F_n) \geq 1\}$$

where Z_{FV} is the number of free variables in F_n (i.e. the number of variables that do not appear in any clause in F_n). Therefore, $\mathbb{E}[Z(F_n)^{-1}, Z(F_n) \geq 1, Z_{FV} \geq cn] \leq 2^{-cn}$. We have $\mathbb{P}\{Z_{FV} < cn, Z(F_n) \geq 1\} \leq \mathbb{P}\{Z_{FV} < cn\}$. Note that Z_{FV} is distributed as the number of empty bins when throwing $nk\alpha$ balls in n bins. Its expectation is $\mathbb{E}\{Z_{FV}\} = ne^{-k\alpha}(1+o(1))$. By simple martingale bounds, for any $c < e^{-k\alpha}$, there exists $c_2 > 0$ such that

$$\mathbb{P}\{Z_{FV} < cn\} \leq 2^{-c_2 n}.$$

Hence there exists $\xi > 0$ such that

$$\tau(n) := \mathbb{E}[\log(1 + Z(F_n)^{-1}), Z(F_n) \geq 1] = O(2^{-\xi n}). \tag{5.4}$$

On the other hand, we have (denoting by F_{n_1}, F_{n_2} two independent formulas and letting $n = n_1 + n_2$)

$$\mathbb{E}\log(1 + Z(F_{n_1})Z(F_{n_2})) = \mathbb{E}[\log(1 + Z(F_{n_1})Z(F_{n_2})), Z(F_{n_1})Z(F_{n_2}) \geq 1]$$

and

$$\mathbb{E}[\log(1 + Z(F_{n_1})Z(F_{n_2})), Z(F_{n_1})Z(F_{n_2}) \geq 1] \geq \mathbb{E}[\log(Z(F_{n_1})Z(F_{n_2})), Z(F_{n_1})Z(F_{n_2}) \geq 1].$$

Hence, using Lemma 2, we get

$$\mathbb{E}[\log Z(F_n), Z(F_n) \geq 1] + \tau(n) \geq \mathbb{E}[\log(Z(F_{n_1})Z(F_{n_2})), Z(F_{n_1})Z(F_{n_2}) \geq 1]$$

or equivalently

$$g(n) + \tau(n) \geq g(n_1) + g(n_2) - g(n_1)\varepsilon(n_2) - g(n_2)\varepsilon(n_1), \quad \forall n_1, n_2 \geq k \tag{5.5}$$

where

$$g(n) = \mathbb{E}[\log Z(F_n), Z(F_n) \geq 1],$$
$$\varepsilon(n) = \mathbb{P}\{Z(F_n) = 0\}.$$

Note that $0 \leq g(n) \leq n$. Therefore (5.5) implies

$$g(n) + \tau(n) \geq g(n_1) + g(n_2) - n_1\varepsilon(n_2) - n_2\varepsilon(n_1), \quad \forall n_1, n_2 \geq k. \tag{5.6}$$

Since $\alpha < \alpha^*$, we have that $\varepsilon(n) = O(1/(\log n)^{1+\delta})$, for some $\delta > 0$. We then restrict ourself to

$$n_1, n_2 \geq t(n) := n/(\log n)^{\eta}, \quad \text{with } \eta = \varepsilon/3.$$

This implies that $n^{1-\eta} \leq n_2$ and $(1-\eta)\log n \leq \log n_2$. So, for n_1, n_2 large enough, we have

$$n_1 \leq n_2(\log n_2)^{2\eta} \tag{5.7}$$
$$n_2 \leq n_1(\log n_1)^{2\eta}. \tag{5.8}$$

Going back to (5.6), we get

$$g(n)+\tau(n) \geq g(n_1)+g(n_2)-n_1(\log n_1)^{2\eta}\varepsilon(n_1)-n_2(\log n_2)^{2\eta}\varepsilon(n_2), \quad \forall n_1,n_2 \geq n/(\log n)^{\eta}$$

or equivalently

$$f(n)+\Delta(n) \geq f(n_1)+f(n_2), \quad \forall n_1,n_2 \geq t(n) \tag{5.9}$$

where

$$\begin{aligned} f(n) &= g(n)-n(\log n)^{2\eta}\varepsilon(n) \\ \Delta(n) &= n(\log n)^{2\eta}\varepsilon(n)+\tau(n) \\ t(n) &= n/(\log n)^{\eta}. \end{aligned}$$

Since $\varepsilon - 2\eta = \varepsilon/3 > 0$, we have

$$\Delta(n) \leq O\left(\frac{n}{(\log n)^{1+\varepsilon/3}}\right),$$

and by Lemma 7, $f(n)/n$ converges, hence $g(n)/n$ converges too. ∎

Proof of Theorem 2. From Theorem 3, for every $\alpha < \alpha^*$, $\psi_n(\alpha) = \frac{1}{n}\mathbb{E}[\log Z(F(n,\alpha,\mu)) | Z(F(n,\alpha,\mu)) \geq 1]$ converges to a limit $\phi_s(\alpha)$. Note that for $\alpha_1,\alpha_2 \in [0,\alpha^*)$ with $\alpha_1 \geq \alpha_2$, it must be (e.g., by a standard coupling argument) that $\psi_n(\alpha_1) \leq \psi(\alpha_2)$. Hence $\phi_s(\cdot)$ is a non-increasing function on $[0,\alpha^*)$ and from Froda's theorem it must have a countable number of plateaus and discontinuities. Let $\alpha_0 \in [0,\alpha^*)$ and denote $\phi_0 = \phi_s(\alpha_0)$. If $\alpha_n(\phi_0)$ does not converge, define

$$\begin{aligned} \underline{\alpha}_0 &= \liminf_{n\to\infty} \alpha_n(\phi_0), \\ \bar{\alpha}_0 &= \limsup_{n\to\infty} \alpha_n(\phi_0), \end{aligned}$$

and pick a sequence $\{n_i\}_{i=1}^{\infty}$ such that $n_i \nearrow \infty$ and

$$\lim_{i\to\infty} \alpha_{n_i}(\phi_0) = \underline{\alpha}_0,$$

and a sequence $\{m_i\}_{i=1}^{\infty}$ such that $m_i \nearrow \infty$ and

$$\lim_{i\to\infty} \alpha_{m_i}(\phi_0) = \bar{\alpha}_0.$$

Then, for any $\alpha \in (\underline{\alpha}_0, \bar{\alpha}_0)$, there exists $\varepsilon > 0$ such that

$$P_{m_i}(\alpha,\phi_0) \leq P_{m_i}(\alpha_{m_i}(\phi_0)-\varepsilon,\phi_0) \overset{i\nearrow\infty}{\to} 0 \tag{5.10}$$

and

$$P_{n_i}(\alpha,\phi_0) \geq P_{n_i}(\alpha_{n_i}(\phi_0)+\varepsilon,\phi_0) \overset{i\nearrow\infty}{\to} 1 \tag{5.11}$$

Moreover, if $\alpha < \alpha^*$,

$$P_{n_i}(\alpha,\phi_0) = \mathbb{P}\left\{\frac{1}{n_i}\log Z(F(n_i,\alpha)) < \phi_0 | Z(F(n_i,\alpha)) \geq 1\right\} + o(1),$$

hence

$$\lim_{i\to\infty} \mathbb{E}\left[\frac{1}{n_i}\log Z(F(n_i,\alpha))|Z(F(n_i,\alpha))\geq 1\right]\leq \phi_0,$$

i.e., since $\psi_n(\alpha)$ converges to $\phi_s(\alpha)$ from Theorem 3,

$$\phi_s(\alpha)\leq \phi_0.$$

Similarly, we have

$$\lim_{i\to\infty} \mathbb{E}\left[\frac{1}{m_i}\log Z(F(m_i,\alpha))|Z(F(m_i,\alpha))\geq 1\right]\geq \phi_0,$$

and

$$\phi_s(\alpha)\geq \phi_0.$$

Therefore, ϕ_0 is a plateau of $\phi_s(\cdot)$, and since $\phi_s(\cdot)$ has countably many plateaus, there are countably many $\phi_0 \in \phi_s([0,\alpha^*))$, for which $\alpha_n(\phi_0)$ does not converge. ∎

Proof of Theorem 1. From Theorem 3, there exists a function $\phi_s(\cdot)$, such that for any $\alpha \in [0,\alpha^*)$, we have $\phi_s(\alpha) = \lim_{n\to\infty}\psi_n(\alpha)$, where $\psi_n(\cdot)$ is defined in Theorem 3. Let $I := \phi_s([0,\alpha^*))$. From Theorem 2, there exists a countable set $\mathcal{C} \subseteq I$ and a function $A : I\backslash\mathcal{C} \to [0,\alpha^*)$ such that for any $\phi \in I\backslash\mathcal{C}$, we can define the limit $A(\phi) = \lim_{n\to\infty}\alpha_n(\phi)$. Note that for any $\phi \in I \setminus \mathcal{C}$, Lemma 1 implies $\phi_s(A(\phi)) = \phi$.

Now, for any $\alpha \in [0,\alpha^*)$ which is not a discontinuity point of ϕ_s (this holds except on a countable subset of $[0,\alpha^*)$), and for any $\varepsilon > 0$, there exists $\varepsilon' < \varepsilon$ such that $\phi_* := \phi_s(\alpha) - \varepsilon' \in I \setminus \mathcal{C}$ and hence $\alpha_n(\phi_*)$ tends to a limit A_*. Note that $A_* > \alpha$, since α is not a discontinuity point of ϕ_s and since $\phi_s(A_*) = \phi_*$. Therefore, there exists $\delta > 0$ such that

$$P_n(\alpha,\phi_s(\alpha)-\varepsilon)\leq P_n(\alpha,\phi_*)\leq P_n(\alpha_n(\phi_*)-\delta,\phi_*)$$

and we conclude by Lemma 1 that $P_n(\alpha,\phi_s(\alpha)-\varepsilon)\to 0$ when $n\to\infty$. With a similar argument, we conclude that $P_n(\alpha,\phi_s(\alpha)+\varepsilon)\to 1$ when $n\to\infty$. ∎

6. PROOFS OF LEMMA 3, LEMMA 4 AND LEMMA 6

Proof of Lemma 6. In this proof, we keep α fixed and split the n variables into two sets of n_1 and $n_2 = n - n_1$ variables, such as $\{1,\ldots,n_1\}$ and $\{n_1+1,\ldots,n\}$. For convenience, we now work with the interpolated Poisson model. We construct a random Boolean formula as follows: we first draw independently the integers M, M_1 and M_2 under Poisson distributions of parameters αnt, $\alpha n_1(1-t)$ and $\alpha n_2(1-t)$ respectively. We then draw independently M clauses from the full system, by picking for each clause the indices of the variables appearing in it independently and uniformly at random within the set of n variables and by picking φ under μ. We also draw independently M_i clauses from each sub-systems, by picking for each clause the indices of the variables appearing in it independently and uniformly at random within the set of n_i variables and by picking φ_i under μ. Finally, we take the conjunction of all clauses to create the formula $F_n(t)$.

Note that the claim of the lemma is equivalent to

$$\mathbb{E}\log(1+Z(F_n(1))) \geq \mathbb{E}\log(1+Z(F_n(0))), \tag{6.1}$$

which is proved by showing that

$$\frac{\partial}{\partial t}\mathbb{E}\log(1+Z(F_n(t))) \geq 0.$$

A straightforward calculation yields

$$\begin{aligned}
&\frac{\partial}{\partial t}\frac{1}{n}\mathbb{E}\log(1+Z(F_n(t))) \\
&\quad = \alpha\left[\mathbb{E}_{\varphi,I}\mathbb{E}\log(1+Z(F_n(t)\wedge\varphi(x_I))) - \mathbb{E}\log(1+Z(F_n(t)))\right] \\
&\qquad - \alpha\frac{n_1}{n}\left[\mathbb{E}_{\varphi_1,I_1}\mathbb{E}\log(1+Z(F_n(t)\wedge\varphi_1(x_{I_1}))) - \mathbb{E}\log(1+Z(F_n(t)))\right] \\
&\qquad - \alpha\frac{n_2}{n}\left[\mathbb{E}_{\varphi_2,I_2}\mathbb{E}\log(1+Z(F_n(t)\wedge\varphi_2(x_{I_2}))) - \mathbb{E}\log(1+Z(F_n(t)))\right],
\end{aligned}$$

where $\varphi,\varphi_1,\varphi_2 \overset{\text{iid}}{\sim} \mu$, $I \sim U^k$, $I_1 \sim U_1^k$, $I_2 \sim U_2^k$, all independent, and where U^k, respectively U_i^k, denotes the k-th product measure of U, respectively U_i (where U, resp. U_i, denotes the uniform measure on the n variables, resp. n_i variables). Hence, $x_I = (x_{i_1},\dots,x_{i_k})$ with $i_1,\dots,i_k$ i.i.d. uniform over the n variables.

We then have

$$\mathbb{E}_{\varphi,I}\mathbb{E}\log(1+Z(F_n(t)\wedge\varphi(x_I))) - \mathbb{E}\log(1+Z(F_n(t))) = \mathbb{E}_{\varphi,I}\mathbb{E}\log\langle\varphi(X_I)\rangle$$

where X is uniformly drawn within the augmented solution space $S(F_{n^*}(t)) = S(F_n(t))\cup\{*\}$, where $*$ is an assignment which returns true on any Boolean functions, and $\langle\cdot\rangle$ denotes the expectation with respect to X. Note that

$$\mathbb{E}_{\varphi,I}\mathbb{E}\log\langle\varphi(X_I)\rangle = -\mathbb{E}_{\varphi,I}\mathbb{E}\sum_{l=1}^{\infty}\frac{\langle\tilde{\varphi}(X_I)\rangle^l}{l}. \tag{6.2}$$

where $\tilde{\varphi} = 1-\varphi$. We now introduce the "replicas" $X^{(r)}$, which are i.i.d. copies of X. We then have

$$\langle\tilde{\varphi}(X_I)\rangle^l = \left\langle\prod_{r=1}^{l}\tilde{\varphi}(X_I^{(r)})\right\rangle, \quad \forall l \geq 1.$$

We are done if we can show that for any realization of the $X^{(r)}$'s and for any $l \geq 1$,

$$\mathbb{E}_{\varphi,I}\prod_{r=1}^{l}\tilde{\varphi}(X_I^{(r)}) - \frac{n_1}{n}\mathbb{E}_{\varphi,I_1}\prod_{r=1}^{l}\tilde{\varphi}(X_{I_1}^{(r)}) - \frac{n_2}{n}\mathbb{E}_{\varphi,I_2}\prod_{r=1}^{l}\tilde{\varphi}(X_{I_2}^{(r)}) \leq 0. \tag{6.3}$$

Note that

$$\mathbb{E}_{\varphi,I}\prod_{r=1}^{l}\tilde{\varphi}(X_I^{(r)}) = \mathbb{E}_{\varphi}\mathbb{E}_{\hat{P}}\prod_{r=1}^{l}\tilde{\varphi}(\xi^{(r)}) \tag{6.4}$$

where $\xi^{(1)}, \ldots, \xi^{(l)} \overset{\text{iid}}{\sim} \hat{P}$ and where $\hat{P}$ is the empirical distribution of $X_I^{(1)}, \ldots, X_I^{(l)}$, i.e. the distribution on $\{-1, 1\}^{kl}$ given by

$$\hat{P}(x_1^{(1)}, \ldots, x_k^{(1)}, \ldots, x_1^{(l)}, \ldots, x_k^{(l)}) = \prod_{i=1}^{k} \bar{P}(x_i^{(1)}, \ldots, x_i^{(l)})$$

with

$$\bar{P}(x_i^{(1)}, \ldots, x_i^{(l)}) = \frac{\#\{i \in \{1, \ldots, n\} : (X_i^{(1)}, \ldots, X_i^{(l)}) = (x_i^{(1)}, \ldots, x_i^{(l)})\}}{n}$$

and similarly

$$\mathbb{E}_{\varphi, I_s} \prod_{r=1}^{l} \tilde{\varphi}(X_{I_i}^{(r)}) = \mathbb{E}_{\varphi} \mathbb{E}_{\hat{P}_s} \prod_{r=1}^{l} \tilde{\varphi}(\xi_s^{(r)}), \quad s = 1, 2 \tag{6.5}$$

where $\xi_s^{(1)}, \ldots, \xi_s^{(l)} \overset{\text{iid}}{\sim} \hat{P}_s$ and where $\hat{P}_s$ is the empirical distribution of $X_{I_s}^{(1)}, \ldots, X_{I_s}^{(l)}$, i.e. the distribution on $\{-1, 1\}^{kl}$ given by

$$\hat{P}_s(x_1^{(1)}, \ldots, x_k^{(1)}, \ldots, x_1^{(l)}, \ldots, x_k^{(l)}) = \prod_{i=1}^{k} \bar{P}_s(x_i^{(1)}, \ldots, x_i^{(l)}), \quad s = 1, 2$$

with

$$\bar{P}_1(x_i^{(1)}, \ldots, x_i^{(l)}) = \frac{\#\{i \in \{1, \ldots, n_1\} : (X_i^{(1)}, \ldots, X_i^{(l)}) = (x_i^{(1)}, \ldots, x_i^{(l)})\}}{n_1},$$
$$\bar{P}_2(x_i^{(1)}, \ldots, x_i^{(l)}) = \frac{\#\{i \in \{n_1 + 1, \ldots, n\} : (X_i^{(1)}, \ldots, X_i^{(l)}) = (x_i^{(1)}, \ldots, x_i^{(l)})\}}{n_2}.$$

Since

$$\Gamma_l(\nu) := \mathbb{E}_{\varphi} \mathbb{P}\{\varphi(Z^{(r)}) = 0, \forall 1 \le r \le l\} = \mathbb{E}_{\varphi} \mathbb{E}_{Z^{(r)}} \prod_{r=1}^{l} (1 - \varphi(Z^{(r)})),$$

where $Z^{(r)}$ are Boolean random vectors of dimension k such that $Z_i = (Z_i^{(1)}, \ldots, Z_i^{(l)})$, $i = 1, \ldots, k$, are i.i.d. under ν, we have that (6.3) is equivalent to

$$\Gamma_l(\bar{P}) - \frac{n_1}{n} \Gamma_l(\bar{P}_1) - \frac{n_2}{n} \Gamma_l(\bar{P}_2) \le 0,$$

which holds by convexity of Γ_l (hypothesis H2), since

$$\bar{P} = \frac{n_1}{n} \bar{P}_1 + \frac{n_2}{n} \bar{P}_2.$$

∎

Proof of Lemma 3. We have

$$\begin{aligned}\Gamma_l(\nu) &= \mathbb{E}_\varphi \mathbb{P}\{\varphi(Z^{(r)}) = 0, \forall 1 \le r \le l\} \\ &= \mathbb{E}_{s^{(j)}} \sum_{z^{(1)},\dots,z^{(l)} \in \{s^{(1)},\dots,s^{(J)}\}} \prod_{i=1}^{k} \nu((z^{(1)})_i, \dots, (z^{(l)})_i) \\ &= \sum_{i_1,\dots,i_l \in \{1,\dots,J\}} [\mathbb{E}_{s_1^{(j)}} \nu(s_1^{(i_1)}, \dots, s_1^{(i_l)})]^k\end{aligned}$$

and Γ_l is convex for any $l \ge 1$. ∎

Proof of Lemma 4. We need to check the convexity of

$$\nu \mapsto \mathbb{E}_\varphi \mathbb{P}\{\varphi(Z^{(r)}) = 0, \forall 1 \le r \le l\} \tag{6.6}$$

where $Z^{(r)}$ are Boolean random vectors of dimension k such that $Z_i = (Z_i^{(1)}, \dots, Z_i^{(l)})$, $i = 1, \dots, k$, are i.i.d. with distribution ν,

$$\varphi_s(x) = \mathbb{1}\left(\prod_{i=1}^{k} x_i = s\right)$$

and

$$\mu(\varphi_1) = \mu(\varphi_{-1}) = 1/2.$$

Note that

$$\mathbb{P}\{\varphi(Z^{(r)}) = 0, \forall 1 \le r \le l\} = \mathbb{P}\left\{\prod_{i=1}^{k} Z_i = -s^l\right\} \tag{6.7}$$

where $-s^l$ denotes the vector $(-s, \dots, -s)$ with l components and where $\prod_{i=1}^k Z_i$ denotes the component-wise product of the vectors Z_i. Since the Z_i are i.i.d. under ν and valued in $\{-1, 1\}$, and since we are interested in their product, we now work with their Fourier transform. For any $Q \subseteq \{1, \dots, l\}$, let

$$f(Q) = f_{Z_1}(Q) = \mathbb{E} \prod_{r \in Q} Z_1^{(r)}.$$

Moreover

$$\mathbb{P}\{Z_1 = 1^l\} = \sum_{Q \in 2^{[l]}} f(Q)$$

and

$$\mathbb{P}\{Z_1 = -1^l\} = \sum_{Q \in 2^{[l]}} (-1)^{|Q|} f(Q).$$

Moreover,

$$f_{\prod_{i=1}^k Z_i}(Q) = f(Q)^k,$$

hence,

$$\mathbb{E}_s\mathbb{P}\left\{\prod_{i=1}^{k} Z_i = -s^l\right\} = 1/2 \sum_{Q\in 2^{[l]}} f(Q)^k + 1/2 \sum_{Q\in 2^{[l]}} (-1)^{|Q|} f(Q)^k$$
$$= \sum_{\substack{Q\in 2^{[l]} \\ |Q| \text{ even}}} f(Q)^k.$$

Since $f(Q)$ is linear in ν (it is the expectation of $\prod_{r\in Q} Z_1^{(r)}$ where $(Z_1^{(1)}, \ldots, Z_1^{(l)}) \sim \nu$), the above summation is clearly convex in ν if k is even. ∎

ACKNOWLEDGMENTS

The present research was initiated as far back as 2006, and remained dormant for a long period of time. Some results were presented at the "DIMACS Working Group on Message-Passing Algorithms" in October 2008. We were finally motivated to polish and publish the manuscript after discussions with Mohsen Bayati, David Gamarnik and Prasad Tetali regarding their recent paper [6]. It is a pleasure to thank them. We also thank Igal Sason and Emre Telatar for a careful reading of the manuscript and useful suggestions. This work was partially supported by a Terman fellowship, the NSF CAREER award CCF-0743978 and the NSF grants DMS-0806211, CCF-0915145.

REFERENCES

[1] E. Abbe and A. Montanari, On the concentration of the number of solutions of random satisfiability formulas (2010). arXiv:1006.3786v1 [cs.DM].

[2] D. Achlioptas and A. Coja-Oghlan, Algorithmic barriers from phase transitions, In 49th Annual Symposium on Foundations of Computer Science, Philadelphia, PA, October 2008, pp. 793–802.

[3] D. Achlioptas, A. Coja-Oghlan, and F. Ricci-Tersenghi, On the solution-space geometry of random constraint satisfaction problems, Random Struct Algorithms 38 (2011), 251–268.

[4] D. Achlioptas, A. Naor, and Y. Peres, Rigorous location of phase transitions in hard optimization problems, Nature 435 (2005), 759–764.

[5] D. Achlioptas and Y. Peres, The threshold for random k-sat is $2^k \log 2 - O(k)$, J Am Math Soc 17 (2004), 947–973.

[6] M. Bayati, D. Gamarnik, and P. Tetali, Combinatorial approach to the interpolation method and scaling limits in sparse random graphs, In Proceedings of the 42nd Annual ACM Symposium on Theory of Computing, STOC '10, Cambridge, MA, New York, NY, 2010, pp. 105–114.

[7] V. Chvátal 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, SFCS '92, IEEE Computer Society, Washington, DC, 1992, pp. 620–627.

[8] A. Coja-Oghlan, A better algorithm for random k-SAT, Proceedings of the 36th ICALP, In A. Susanne, M.-S. Alberto, M. Yossi, N. Sotiris, and T. Wolfgang (Editors), Lecture Notes in Computer Science, vol. 5555, Springer, Berlin/Heidelberg, 2009, pp. 292–303.

[9] M. Dyer, L. A. Goldberg, C. S. Greenhill, and M. Jerrum, The relative complexity of approximate counting problems, Algorithmica 38 (2003), 471–500.

[10] W. Fernandez de la Vega, On random 2-SAT, manuscript, 1992.

[11] W. Fernandez de la Vega, Random 2-SAT: results and problems, Theor Comput Sci 265 (2001), 131–146.

[12] S. Franz and M. Leone, Replica bounds for optimization problems and diluted spin systems, J Stat Phys 111 (2003), 535.

[13] S. Franz, M. Leone, and F. L. Toninelli, Replica bounds for diluted non-Poissonian spin systems, J Phys A 36 (2003), 10967.

[14] E. Friedgut, Sharp thresholds of graph properties, and the k-sat problem, J Amer Math Soc 12 (1999), 1017–1054; Appendix by J. Bourgain.

[15] A. Goerdt, A threshold for unsatisfiability, J Comput Syst Sci 53 (1996), 469–486.

[16] F. Guerra and F. L. Toninelli, The thermodynamic limit in mean field spin glasses, Commun Math Phys 230 (2002), 71–79.

[17] S. Kirkpatrick and B. Selman, Critical behavior in the satisfiability of random boolean expressions, Science 264 (1994), 1297–1301.

[18] F. Krzakala, A. Montanari, F. Ricci-Tersenghi, G. Semerjian, and L. Zdeborova, Gibbs states and the set of solutions of random constraint satisfaction problems, Proc Natl Acad Sci 104 (2007), 10318–10323.

[19] S. Kudekar and N. Macris, Sharp bounds for optimal decoding of low-density parity-check codes, IEEE Trans Inform Theory 55 (2009), 4635–4650.

[20] M. Ledoux, The concentration of measure phenomenon, American Mathematical Society, Providence, RI, 2001.

[21] M. Mézard and A. Montanari, Information, physics, and computation, Oxford University Press, Oxford, 2009.

[22] M. Mézard, G. Parisi, and R. Zecchina, Analytic and algorithmic solution of random satisfiability problems, Science 297 (2003), 812–815.

[23] R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, and L. Troyansky, Determining computational complexity from characteristic phase transitions, Nature 400 (1999), 133–137.

[24] A. Montanari, Tight bounds for LDPC and LDGM codes under MAP decoding, IEEE Trans Inform Theory 51 (2005), 3221–3246.

[25] A. Montanari, R. Restrepo, and P. Tetali, Reconstruction and clustering in random constraint satisfaction problems (2009). CoRR abs/0904.2751.

[26] A. Montanari and D. Shah, Counting good truth assignments of random k-sat formulae, SODA New Orleans, USA, January 2007, pp. 1255–1264.

[27] D. Panchenko and M. Talagrand, Bounds for diluted mean-field spin glass models, Prob Theor Rel Fields 130 (2004), 319–336.

[28] L. G. Valiant, The complexity of enumeration and reliability problems, SIAM J Comput 8 (1979), 410–421.

[29] A. Montanari, F. Ricci-Tersenghi, and G. Semerjian, Clusters of solutions and replica symmetry breaking in random k-satisfiability, J Stat Mech (2008), P04004.

[30] D. Achlioptas, J. Han Kim, M. Krivelevich, and P. Tetali, Two-coloring random hypergraphs, Random Struct Algorithms 20 (2002), 249–259.