

Reweighted belief propagation and quiet planting for random K -SAT

Florent Krzakala

fk@espci.fr

*Laboratoire de Physique Statistique, Ecole Normale Supérieure, 45 rue d'Ulm
Université P. et M. Curie, Paris, France
ESPCI ParisTech, 10 rue Vauquelin, UMR 7083 Gulliver, Paris 75000, France.*

Marc Mézard

mezard@lptms.u-psud.fr

*Ecole Normale Supérieure, 45 rue d'Ulm, Paris France.
Université Paris-Sud & CNRS, LPTMS, UMR8626,
Bât. 100, Université Paris-Sud 91405 Orsay, France.*

Lenka Zdeborová

lenka.zdeborova@cea.fr

*Institut de Physique Théorique, IPhT,
CEA Saclay, and URA 2306, CNRS,
91191 Gif-sur-Yvette, France.*

Abstract

We study the random K -satisfiability problem using a partition function where each solution is reweighted according to the number of variables that satisfy every clause. We apply belief propagation and the related cavity method to the reweighted partition function. This allows us to obtain several new results on the properties of random K -satisfiability problem. In particular the reweighting allows to introduce a planted ensemble that generates instances that are, in some region of parameters, equivalent to random instances. We are hence able to generate at the same time a typical random SAT instance and one of its solutions. We study the relation between clustering and belief propagation fixed points and we give a direct evidence for the existence of purely entropic (rather than energetic) barriers between clusters in some region of parameters in the random K -satisfiability problem. We exhibit, in some large planted instances, solutions with a non-trivial whitening core; such solutions were known to exist but were so far never found on very large instances. Finally, we discuss algorithmic hardness of such planted instances and we determine a region of parameters in which planting leads to satisfiable benchmarks that, up to our knowledge, are the hardest known.

1. Motivation

The satisfiability of Boolean formulas is a fundamental problem in theoretical computer science. It was the first problem shown to be NP-complete [14, 22, 41], and it is of central relevance in various practical applications, including artificial intelligence, planning, hardware and electronic design, automation, verification and more. It can thus be thought of as the “Ising model” of computer science. Ensembles of randomly generated instances of the satisfiability problem have emerged in computer science as a way of evaluating algorithmic performance and addressing questions regarding the average case complexity.

An instance of the random “ K -SAT” problem consists of N Boolean variables and M clauses. Each clause contains a subset of K distinct variables chosen uniformly at

random, and each clause forbids one random assignment of these K variables out of the 2^K possible ones. The instance is *satisfiable* if there exists an assignment of variables that simultaneously satisfies all clauses. In this case we call such an assignment a *solution* to the instance. When the density of constraints, defined as $\alpha = M/N$, is increased, the formulas become less likely to be satisfiable. In the thermodynamic limit, i.e. when $N \rightarrow \infty$ at fixed α , there is a sharp transition from a phase, $\alpha < \alpha_c(N)$, in which the formulas are almost surely satisfiable to a phase, $\alpha > \alpha_c(N)$, where they are almost surely unsatisfiable [15]. Although the convergence of the sequence $\alpha_c(N)$ as $N \rightarrow \infty$ is still an open problem, it is widely believed that this sequence is convergent, and therefore that there is a phase transition at $\alpha_c = \lim_{N \rightarrow \infty} \alpha_c(N)$. It is also a well known empirical result that the hardest instances are found near to this threshold [11, 32, 17].

Random K -SAT has attracted the interest of mathematicians, computer scientists and statistical physicists. One very fruitful methodological direction to study random K -SAT is the belief propagation (BP) algorithm [43, 21], which is closely related to the cavity method that was developed in statistical physics for studies of mean field spin glasses [28, 27]. The results and insights coming from the cavity method are remarkable. The satisfiability threshold and other phase transitions in the structure of solutions have been described in [7, 31, 29, 18]. In particular, it was shown that for $K \geq 3$ the space of solutions for highly constrained but still satisfiable instances splits into exponentially many (in N) clusters which are far away from each other. Another important concept studied recently in random K -SAT [47, 35, 4] is the one of frozen variables: a given variable is frozen relatively to a given cluster of solutions if the variable is fixed to the same value in all the solutions of this cluster. Clearly any local search procedure which approaches a cluster must identify correctly the frozen variables.

The reader should note that, from a mathematical point of view, the majority of the results obtained with the cavity method are only conjectures since certain assumptions of the method were not proven yet. However, those assumptions are widely believed to be correct and hence these results, as well as the results of the present article, should be regarded as “presumably-exact conjectures” based on the so far heuristic, but sophisticated “cavity method”, and the related “replica method”. Notably a number of the predictions coming from this method (e.g. the existence of clustering of solutions) have been confirmed rigorously in random K -SAT [26, 3]. And an even larger number of these “presumably-exact conjectures” have been proven for other random constraint satisfaction problems such as coloring of random graphs, NAE-SAT or K -XOR-SAT since for those problems some methodological simplifications arise, see e.g. [12, 13, 33].

In spite of the large corpus of results/conjectures obtained in recent years via the cavity method, there are still many open questions in random K -SAT, even on this heuristic level. Let us mention the main ones.

- The relation between detailed properties of the energy landscape (where “energy” is defined as the number of unsatisfied clauses, and its domain are all the possible assignments) as predicted by the cavity method on one hand, and what is found by numerical simulations on the other hand is still not fully established. In the limit of large K , the situation is relatively well under control [18, 2]: phase transitions in the space of solutions can be identified with a threshold beyond which a large class of simple algorithms (or at least the methods used to analyze them) fail. How-

ever, for small K the algorithmic limits are strongly algorithm-dependent, and many algorithmically-relevant questions about the energy landscape are still largely open. In particular, relatively simple stochastic-local-search algorithms have been found to be efficient in the regimes of the density of constraints where clustering of solutions occurs [46].

- Another annoying observation is that so far no one has been able to find solutions of large (consider e.g. $N > 10^5$) random K -SAT formulas with frozen variables (i.e. variables which take the same value in all solutions) in spite of the fact that theory clearly predicts their existence [31, 49] at large enough density of constraints.
- Another aspect of this paradox is the fact that it was impossible so far to find numerically non-trivial solutions of the BP equations, while the theory shows that there are exponentially many of them (each associated with one cluster of solutions), and actually an algorithm like survey propagation is precisely based on the idea of counting the BP fixed points associated with frozen clusters [31].

In this work we shall address some of these open issues by developing a ‘quiet planting’ procedure for the random K -SAT problem. Planting is a way to generate a K -SAT instance with a known solution. Instead of generating the clauses randomly and then trying to find a solution, one first generates a special random configuration of variables, which we call the ‘planted configuration’ and then one generates the clauses in such a way that each clause is satisfied by the planted configuration. By construction the planted configuration is a solution of the instance so generated, one says that it has been ‘planted’ in the instance. Unfortunately, the naive planting that we have just described generates some instances that are very distinct from typical random K -SAT instances. A more careful procedure, called ‘quiet planting’, allows to generate K -SAT instances that have the same distribution as fully random instances (when the density of constraints is below the threshold α_c , to be defined later). Therefore quiet planting is a procedure that provides a random instance together with a solution. So far quiet planting had been developed for other constraint satisfaction problems [2, 19, 20, 50]. In this paper we study quiet planting in K -SAT. In order to do so, we introduce and study the *reweighted belief propagation* and the associated cavity method results. This study allows to answer several of the important open questions. From a mathematically rigorous point of view we provide new conjectures about the random K -SAT problem. More specifically, we address the following points:

- We find numerically some non-trivial BP fixed points by initializing BP in the quietly-planted configuration. This shows that such non-trivial fixed points do exist, as expected from the cavity method analysis. This answers a long-lasting question asked by random K -SAT practitioners who could not find these non-trivial BP fixed-points. They exist, but they are hard to find, and one needs a special procedure like quiet planting in order to see them.
- Clusters are often imprecisely described in literature as connected components of the “solution graph”. This is defined as a graph in the space of solutions where edges connect solutions that differ in no more than some number $d \geq 1$ of variables (the precise value of d depends on the problem or on the author). Hence to walk from one

such cluster to another one, it is necessary to violate a number of clauses $n(d) > 0$ on the way, this is called an ‘energetic’ barrier between clusters. It is known that clusters separated by energetic barriers exist (in the appropriate range of α), when K is large enough. Here we show the existence of different type of barriers, called “entropic barriers”. These occur whenever a connected component of the solution graph is made of two subsets of solutions connected by a very narrow path of solutions: The path from one subset to another does not require violation of any clause, but this path is rare and a random walker would need a very long time, typically a time that diverges exponentially when $N \rightarrow \infty$. Using the reweighted BP we explicitly show that entropic barriers do exist in random K -SAT. Entropic barriers had been known to exist in simple spin glass models (see for instance [5, 45, 39]), but it is the first time that they are found in K -SAT.

- An interesting problem which has been studied on several occasions in the literature [6, 36] is: how to plant a solution in a 3-SAT problem such that the resulting formula is nevertheless very hard. In this paper we clarify in which region the planted instances are hard, namely in the shaded region in Fig. 2. Qualitatively we reach the same conclusions as [6]. We give the exact boundary of the hard region (whereas the calculations in [6] were only approximate) and we provide more detailed heuristic arguments for this result. In the end we have a way to create the hardest known (at least to the authors) satisfiable 3-SAT formulas. Note that the hard ensemble discussed here is closely related to the one of [16].

2. Reweighting in random K -SAT

The satisfiability problem is defined over N Boolean variables $s_i \in \{0, 1\}$, and M clauses. Each clause contains K variables. If a variable i is negated in the clause a we set $J_{ia} = 1$, otherwise we set $J_{ia} = 0$. A clause a is satisfied if and only if

$$n_a = \sum_{i \in \partial a} (1 - \delta_{s_i, J_{ia}}) > 0, \quad (1)$$

where ∂a is the set of variables belonging to the clause a . Here we introduced n_a as the number of variables that are satisfying the clause a .

Let us first define the standard probability measure over solutions. We assume that there exists at least one solution, and we define the partition function as the number of solutions:

$$Z = \sum_{\{s_i\}} \prod_a C_a(\{s_i\}_{i \in \partial a}), \quad (2)$$

where $C_a = 1$ if clause a is satisfied and $C_a = 0$ otherwise. The sum is over all the 2^N assignments of variables. One introduces a ‘Boltzmann’-type measure as the uniform measure over all solutions:

$$\mu(\{s_i\}) = \frac{1}{Z} \prod_a C_a(\{s_i\}_{i \in \partial a}). \quad (3)$$

The cavity method (or BP) can be used to study heuristically the properties of this ‘standard’ measure.

Let us now define a reweighted measure. We introduce K parameters $\lambda_1, \dots, \lambda_K$ which are non-negative real numbers, and we define the ‘ λ -measure’ as:

$$\mu_\lambda(\{s_i\}) = \frac{1}{Z(\lambda)} \prod_a [C_a(\{s_i\}_{i \in \partial a}) \lambda_{n_a(\{s_i\}_{i \in \partial a})}] , \quad (4)$$

Where the reweighted partition function is defined as

$$Z(\lambda) = \sum_{\{s_i\}} \prod_a [C_a(\{s_i\}_{i \in \partial a}) \lambda_{n_a(\{s_i\}_{i \in \partial a})}] . \quad (5)$$

We denote the vector of $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_K)$. In order to have a unique measure associated with each vectors λ , we choose to fix the normalization

$$1 = \sum_{r=1}^K \binom{K}{r} \lambda_r . \quad (6)$$

(note that any normalization procedure that breaks the symmetry of (4) with respect to a global multiplication of all components of λ would be fine, we chose eq. (6) because it leads to convenient simplifications in some expressions that follow). The standard measure and partition function are recovered when all components of λ are equal: $\forall r \lambda_r = 1/(2^K - 1)$.

The main subject of this paper is to generalize the cavity-method-based conjectures on the phase diagram of random K -SAT to the generalized reweighted K -SAT problem with vector-parameter λ , and to discuss new interesting results which can be derived using the generalized measure (4). There are some properties that do not depend on the reweighting parameters λ (as long as $\lambda_r > 0$ for all $r = 1, \dots, K$), for instance the satisfiability threshold or the freezing transition, which is the density of constraints at which strictly all clusters start to contain frozen variables. But other physically important phase transition defined in [18, 35] such as the clustering or the condensation transitions do depend on reweighting parameters.

Note that special cases of the reweighted partition function were studied previously. In particular, [1] used the set of reweighting parameters

$$\lambda_r = \frac{\gamma^r}{\sum_{r=1}^K \binom{K}{r} \gamma^r} \quad (7)$$

as a tool to be able to use the second moment method to obtain a lower bound on the satisfiability threshold. Indeed, the success of this reweighting in the second moment calculations inspired our study. The same reweighting was later used in several other rigorous works in conjunction with the second moment method. Let us also mention the work of [36] where planting according to the reweighting (7) was studied numerically, together with the first and second moments calculations. However, it seems that the authors of [36] did not notice that for a special value of γ the generated instances are equivalent to random instances, which is one of the crucial points in the present paper. Reweighted planting was also studied in [6] as a way to create hard instances. We shall comment on all these results later in the light of our findings.

3. Belief propagation for the reweighted partition function

The reweighted partition function (5) can be computed via the belief propagation (BP) algorithm [43, 21]. The BP result is exact on trees, and it corresponds to the Bethe approximation on sparse random graphs. Considering the probability measure (4), we define $\psi_{s_i}^{a \rightarrow i}$ as the probability that the constraint a is satisfied, conditioned on the fact that the value of variable i is s_i . Similarly, $\chi_{s_j}^{j \rightarrow a}$ is the probability that the variable j takes value s_j conditioned on the fact that the constraint a has been removed from the graph. According to belief propagation these messages then satisfy the set of equations

$$\chi_{s_i}^{i \rightarrow a} = \frac{\prod_{b \in \partial i \setminus a} \psi_{s_i}^{b \rightarrow i}}{\prod_{b \in \partial i \setminus a} \psi_1^{b \rightarrow i} + \prod_{b \in \partial i \setminus a} \psi_0^{b \rightarrow i}}, \quad (8)$$

$$\psi_{J_{ia}}^{a \rightarrow i} = \frac{1}{Z^{a \rightarrow i}} \sum_{\{s_j\}_{j \in \partial a \setminus i}} \theta(K - 1 - \sum_{j \in \partial a \setminus i} \delta_{J_{ja}, s_j}) \lambda_{K-1-\sum_{j \in \partial a \setminus i} \delta_{s_j, J_{ja}}} \prod_{j \in \partial a \setminus i} \chi_{s_j}^{j \rightarrow a}, \quad (9)$$

$$\psi_{1-J_{ia}}^{a \rightarrow i} = \frac{1}{Z^{a \rightarrow i}} \sum_{\{s_j\}_{j \in \partial a \setminus i}} \lambda_{K-\sum_{j \in \partial a \setminus i} \delta_{s_j, J_{ja}}} \prod_{j \in \partial a \setminus i} \chi_{s_j}^{j \rightarrow a}, \quad (10)$$

where the normalization $Z^{a \rightarrow i}$ ensures that $\psi_{J_{ia}}^{a \rightarrow i} + \psi_{1-J_{ia}}^{a \rightarrow i} = 1$, and $\theta(x)$ is the Heaviside step function. Once these messages (i.e. the solution of these equations) have been found, the marginal probability that variable i takes value s_i , denoted by $\chi_{s_i}^i$, can be computed as:

$$\chi_{s_i}^i = \frac{\prod_{b \in \partial i} \psi_{s_i}^{b \rightarrow i}}{\prod_{b \in \partial i} \psi_1^{b \rightarrow i} + \prod_{b \in \partial i} \psi_0^{b \rightarrow i}}. \quad (11)$$

The log-partition function can then be computed using the Bethe formula. This gives the “free-entropy” density:

$$s(\lambda) = \frac{1}{N} \log Z(\lambda) = \frac{1}{N} \left(\sum_a \log Z^a + \sum_i \log Z^i - \sum_{ia} \log Z^{ia} \right), \quad (12)$$

where

$$Z^i = \prod_{a \in \partial i} \psi_1^{a \rightarrow i} + \prod_{a \in \partial i} \psi_0^{a \rightarrow i}, \quad (13)$$

$$Z^a = \sum_{\{s_i\}_{i \in \partial a}} \theta(K - \sum_{i \in \partial a} \delta_{J_{ia}, s_i}) \lambda_{K-\sum_{i \in \partial a} \delta_{s_i, J_{ia}}} \prod_{i \in \partial a} \chi_{s_i}^{i \rightarrow a}, \quad (14)$$

$$Z^{ia} = \chi_0^{i \rightarrow a} \psi_0^{a \rightarrow i} + \chi_1^{i \rightarrow a} \psi_1^{a \rightarrow i}. \quad (15)$$

4. Quiet planting in random K -SAT

A reweighting parameter set λ^* is defined as “magic” when the set of eqs. (8-10) has a solution (fixed point) for which the messages do not depend on the edge (ai) , i.e. $\chi_s^{i \rightarrow a} = \chi_s$, and $\psi_s^{a \rightarrow i} = \psi_s$ for all (ai) . Such a solution (fixed point) is referred to in the statistical physics literature as the “factorized” fixed point of belief propagation. We shall see that problems defined with magic reweighting have special properties.

4.1 Magic reweighting and quenched-annealed equivalence

Since in random K -SAT the negations are chosen at random there is a global symmetry between 0's and 1's. A factorized fixed point of BP hence also has to be symmetric, i.e. $\psi_s^{a \rightarrow i} = \chi_s^{i \rightarrow a} = 1/2$ for all (ai) and s . With the use of normalization (6) it is easy to see that Eqs. (9)-(10) permit such a fixed point if and only if λ satisfies

$$\frac{1}{2} = \sum_{r=1}^K \binom{K-1}{r-1} \lambda_r^*. \quad (16)$$

Note that, if we restrict to the power-law choice of the reweighting parameters (7) made in [1, 36], we recover from (16) the condition for the special value $\gamma = \gamma^*$ used in these works, defined by:

$$1 = (1 + \gamma^*)^{K-1} (1 - \gamma^*). \quad (17)$$

In what follows we will mostly speak about the general λ^* magic reweighting, but sometimes we focus for simplicity on the power-law choice of γ^* . We urge the reader not to confuse the two, as γ^* is only one special point in the space of magic reweightings.

The free-entropy density (12) in the case of a factorized BP fixed point is equal to

$$s_{\text{fBP}}(\lambda^*) = (1 - \alpha K) \log 2, \quad (18)$$

this is obtained simply by evaluating eq. (12) for $\psi_s^{a \rightarrow i} = \chi_s^{i \rightarrow a} = 1/2$.

We now define the quenched and annealed free-entropy density as

$$s_{\text{annealed}}(\lambda) \equiv \lim_{N \rightarrow \infty} \frac{1}{N} \log \mathbb{E}[Z(\lambda)], \quad (19)$$

$$s_{\text{quenched}}(\lambda) \equiv \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}[\log Z(\lambda)]. \quad (20)$$

where $Z(\lambda)$ is the partition function defined in (5) and the expectation is with respect to the measure (4). These names are motivated from properties of glassy material after a fast (quenched) or slow (annealed) decrease of temperature. It follows from the cavity method that under certain conditions the Bethe free-entropy density of the factorized fixed point is equal to the quenched free-entropy density $s_{\text{fBP}}(\lambda^*) = s_{\text{quenched}}(\lambda^*)$ in the thermodynamic limit $N \rightarrow \infty$. The condition for this equality to hold is that the density of constraints, α should be smaller than a threshold value where a thermodynamic phase transition called the “condensation transition” occurs. This transition is defined by the non-analyticity of the quenched free-entropy.

Let us now compute the annealed free-entropy density of the reweighted ensemble with arbitrary λ . We get

$$\mathbb{E}[Z(\lambda)] = 2^N \left[\sum_{r=1}^K \lambda_r \binom{K}{r} \frac{1}{2^K} \right]^{\alpha N} = 2^N 2^{-\alpha K N}, \quad (21)$$

where in the second equality we used the normalization condition (6). We see that the annealed free-entropy density equals the Bethe free-entropy density of the factorized BP solution (18) found for $\lambda = \lambda^*$, $s_{\text{annealed}} = s_{\text{fBP}}(\lambda^*)$. This is actually a general property, see [38]. Consequently, at low enough density of constraints, the quenched free-entropy density at λ^* is equal to the annealed free-entropy density. This last property is the crucial point that makes quiet planting in the sense of [2, 20, 50] possible.

4.2 Quiet planting in the reweighted ensemble

The definition of the λ -rewighted planted ensemble is the following:

1. Choose the geometry of the random K -SAT formula as before (M clauses where each contains a uniformly-chosen random K -tuple of variables).
2. Define a special configuration, called the *planted* configuration, by assigning randomly $1/2$ of the variables to the value 0 and the other $1/2$ to 1.
3. Define the probability $p(r)$ by:

$$p(r) = \binom{K}{r} \lambda_r \quad \text{for } r > 0, \quad p(0) = 0. \quad (22)$$

For each clause a , choose a random integer $r_a > 0$ from this distribution, and then choose at random one out of the $\binom{K}{r_a}$ configurations of negations for which the planted configuration satisfies the clause by r_a different variables.

In this planted ensemble, the probability that a randomly chosen clause is satisfied by r variables from the planted configuration is thus equal to $p(r)$ defined in (22).

First of all notice that since every planted configuration is consistent with the same number of instances the planted configuration has all the properties of a configuration sampled uniformly at random from the corresponding probability measure $\mu(\lambda)$, eq. (4). This is a very generic property of planting. Furthermore, it is easy to see that this planting procedure generates instances randomly but in general not uniformly over all instances. Every planted instance appears with probability proportional to its partition function $Z(\lambda)$. Hence the planted ensemble of instances is in general very different from the random ensemble. Only in the case when the annealed free-entropy density equals the quenched one, the fluctuations of $Z(\lambda)$ are small enough (in the large N limit) so that the planting procedure generates instances statistically equivalent to the random ones. This property called “quiet planting” means that every property that holds with high probability in the random ensemble holds also with high probability in the planted ensemble and vice versa, as discussed in [2, 20, 51, 40].

To summarize, when one uses magic-rewighting, with constraint densities smaller than the corresponding condensation phase transition [18], the annealed free-entropy density equals the quenched one and the planting is quiet. This means that the planting procedure has the following two properties: First it creates formulas of the satisfiability problem that are statistically equivalent to random formulas in the thermodynamic limit $N \rightarrow \infty$. Second, the planted configuration has all the properties of a configuration sampled randomly from the measure $\mu(\lambda^*)$. In physics language, the planted configuration is an “equilibrium configuration” of the measure $\mu(\lambda^*)$ (while it is not an equilibrium configuration of the uniform measure over all solutions). In the next sections we explore consequences of these simple but very useful properties.

4.3 Cavity equations in the “magically” reweighted ensemble

Iterating the BP equations on a given instance generated with quiet planting until a fixed point is reached is easy and provides an experimental access to the study of phase transitions

in this problem. However, the very notion of phase transition is an asymptotic one (it exists only in the thermodynamic limit $N \rightarrow \infty$), and it is thus very useful to get a control of this thermodynamic limit. This is done in the cavity method by studying an ensemble of planted problems. The corresponding distributional equations can be solved efficiently with population dynamics technique [27], these equations are referred to as the replica-symmetric cavity equations.

In the present case the object studied by the cavity method is the probability $P_{s_i}^t(\psi)$ that at the t -th iteration of BP equations a message arriving on a site i takes the value $\psi = (\psi_0, \psi_1)$, conditionally on the fact that the planted configuration on i is equal to s_i . It satisfies:

$$P_{s_i}^t(\psi) = \sum_{J_j} \prod_{j=1}^{K-1} \frac{\delta_{J_j,0} + \delta_{J_j,1}}{2} \sum_{\{s_j\}} \tilde{P}(\{s_j\}, s_i) \int \prod_{j=1}^{K-1} \sum_{l_j=1}^{\infty} q(l_j) \prod_{k_j=1}^{l_j} P_{s_j}^{t-1}(\psi^{k_j}) d\psi^{k_j} \delta(\psi - \mathcal{F}(\{\psi^{k_j}\})) \quad (23)$$

where the function $\mathcal{F}(\{\psi^{k_j}\})$ is obtained by combining (9-10) with (8)

$$\mathcal{F}_{J_i}(\{\psi^{k_j}\}) = \frac{1}{Z} \sum_{\{s_j\}_{j=1}^{K-1}} \theta(K-1 - \sum_{j=1}^{K-1} \delta_{J_j, s_j}) \lambda_{K-1 - \sum_{j=1}^{K-1} \delta_{s_j, J_j}} \prod_{j=1}^{K-1} \prod_{k_j=1}^{l_j} \psi_{s_j}^{k_j}, \quad (24)$$

$$\mathcal{F}_{1-J_i}(\{\psi^{k_j}\}) = \frac{1}{Z} \sum_{\{s_j\}_{j=1}^{K-1}} \lambda_{K - \sum_{j=1}^{K-1} \delta_{s_j, J_j}} \prod_{j=1}^{K-1} \prod_{k_j=1}^{l_j} \psi_{s_j}^{k_j}, \quad (25)$$

and $q(l_j)$ is the Poisson distribution with mean αK . The probability distribution $\tilde{P}(\{s_j\}, s_i)$ is obtained by drawing a random variable J_i uniformly from $\{0, 1\}$ and then a number r of variables s_j 's such that $s_j \neq J_j$, where r is drawn from the distribution $\tilde{p}(r)$:

$$\tilde{p}(r) = \frac{r p(r)}{\sum_{r'=0}^K r' p(r')} \quad \text{if } J_i \neq s_i, \quad (26)$$

$$\tilde{p}(r) = \frac{(K-r)p(r)}{\sum_{r'=0}^K (K-r')p(r')} \quad \text{if } J_i = s_i. \quad (27)$$

Where $p(r)$ is given by (22). The planted initialization then corresponds to $P_{s_i}^{\text{init}}(\psi) = \delta(\psi - u)$ where $u_{s_i} = 1$ and $u_{1-s_i} = 0$. The average Bethe free-entropy is expressed along the same lines.

5. Phase transitions in the reweighted ensemble, or how to use quiet planting

Following the argumentation in [19], the quietly-planted satisfiability instances have in general three sharp phase transitions as the density of constraints increases. The simplest way to locate these phase transitions is to directly perform the planting, and then iterate

till convergence the reweighted BP equations (8-10) on the resulting instance, with two possible initializations:

- “Random initialization”. One initializes messages $\chi^{i \rightarrow a}$ and $\psi^{a \rightarrow i}$ as random vectors drawn uniformly from the set of normalized vectors. In the cavity equations, $P_{s_i}^0(\psi)$ is the uniform density over the set $\psi_0 + \psi_1 = 1$.
- “Initialization in the planted configuration”. One initializes the messages in the planted configuration, i.e. $\chi_1^{i \rightarrow a} = 1$, $\chi_0^{i \rightarrow a} = 0$ if variable i was planted 1 or $\chi_1^{i \rightarrow a} = 0$, $\chi_0^{i \rightarrow a} = 1$ if variable i was planted 0. The distribution $P_{s_i}^0(\psi)$ in the cavity equations is the one induced by this choice, using Eqs. (9-10).

Let us summarize the general scenario found in quiet planting. The following results can be found empirically by performing the numerical experiment of iterating BP equations on large instances (which is easily done), but the real control of the phase transitions is done through the asymptotic analysis of the distributional equations of the cavity method written in the previous section.

In order to describe the various phases of the system, we use a vocabulary borrowed from statistical physics. We say that a system is “ferromagnetic” when the expectation of a variable has a positive correlation with the planted configuration (in the same way that a system of spins is ferromagnetic if the variables -the spins - have a positive correlation with a given spatial orientation, the orientation of the magnetic moment of the material). If we call τ the planted configuration, the probability measure $\mu(\{s_i\})$ is ferromagnetic if in expectation the configurations sampled from the measure have strictly positive overlap with the planted configuration, where the overlap q is defined as: $q = \lim_{N \rightarrow \infty} (1/N) \sum_{i=1}^N [\mu(s_i = \tau_i) - \mu(s_i \neq \tau_i)] > 0$. If this overlap is zero we call the measure “paramagnetic”.

As a function of the constraint density α we find four different phases in the quietly-planted satisfiability, identified by the behavior of the fixed points obtained from iterating BP with our two different initial conditions.

- For $\alpha < \alpha_d$ the system is in a purely paramagnetic phase. BP converges towards the factorized fixed point when started from the two initializations. The overlap between the BP fixed point and the planted configuration is zero. It turns out that α_d is the clustering threshold. In this phase the set of solutions forms one cluster, two random configurations from this cluster are completely uncorrelated, the planted configuration is simply one of them, and there is no way to tell which one.
- For $\alpha_d < \alpha < \alpha_c$ there is a paramagnetic phase with a subdominant (metastable) ferromagnet. BP started from the random initialization converges to the factorized fixed point, whereas BP started from the planted initialization converges to a fixed point with positive overlap with the planted configuration. The Bethe free-entropy of this second fixed point is smaller than the annealed free-entropy and hence the paramagnet is still the dominant solution. The Bethe free-entropy of the factorized fixed point is still equal to the quenched free-entropy. It turns out that α_c is the condensation threshold. In this phase there are exponentially many clusters of roughly the same size, the planted configuration belongs to one of them and there is no way to tell which one it belongs to.

- For $\alpha_c < \alpha < \alpha_l$ there is a ferromagnetic phase with a subdominant (metastable) paramagnet. Like in the previous phase, BP started from the random initialization still converges to the factorized fixed point, whereas BP started from the planted initialization converges to a fixed point with positive overlap with the planted configuration. But this time the Bethe free-entropy of the BP fixed point obtained from the planted initialization is larger than the Bethe free-entropy of the factorized fixed point. Hence the planted cluster dominates the measure. It is now the Bethe free-entropy of the BP fixed point obtained from the planted initialization that is equal to the quenched free-entropy. In this phase the planting is not quiet anymore, in the sense that the properties of planted and random instances are different. For instance random instances become typically unsatisfiable for $\alpha > \alpha_s > \alpha_c$ (α_s being the satisfiability threshold), whereas the planted ones are always satisfiable.
- For $\alpha_l < \alpha$ there is a ferromagnetic phase. In this phase the two initializations give the same result: in both cases BP converges to a fixed point which is correlated with the planted configuration. Hence it is easy to find a solution correlated with the planted configuration.

We emphasize here that the planting is “quiet” only for $\alpha \leq \alpha_c$, whereas the properties of the random and planted ensemble are different for $\alpha > \alpha_c$.

In the above discussion the phases have been identified by analyzing the behavior of BP fixed-points. It turns out that the phase transitions found in this way also correspond to phase transitions in the geometrical properties of the space of solutions of satisfiability found in previous works. We shall state this correspondance here without proof.

In the random K -SAT ensemble, reweighted with the parameter λ^* , α_d is the threshold that corresponds to the clustering (called dynamical 1RSB in statistical physics) transition of the measure $\mu(\lambda^*)$. For $\alpha < \alpha_d$ the set of solutions form one cluster. The planted configuration is one configuration of this cluster, it has no special property. The threshold α_c is the condensation (called static 1RSB in statistical physics) transition of the measure $\mu(\lambda^*)$. The phase transition at α_l is a point beyond which the paramagnetic phase ceases to exist. In the language of magnetic systems, α_l is called the spinodal threshold.

Another important property that follows from the analysis of quiet planting [19] is that the part of the phase space that is not correlated with the planted configuration remains unchanged in the whole region of α , i.e. also for $\alpha > \alpha_c$. Notably for $\alpha > \alpha_s$ (where α_s is the satisfiability threshold in the random SAT ensemble) only solutions correlated to the planted configuration exist in the planted ensemble. In the random K -SAT ensemble without planting, α_l does not have a physical interpretation, but at this point the BP iterations stop to converge.

The transition point α_l is investigated by computing the local stability of the factorized BP fixed point with respect to small random perturbations. This is done by linearizing equations (8-10) via expansion around the factorized fixed point and analyzing the largest eigenvalue x of the associated linear operator. The following analytical formula then locates the phase transition between a phase in which an infinitesimal random perturbation is damped to zero from a phase where this perturbation gets amplified in the direction of the

planted configuration:

$$\alpha_l = \frac{1}{K(K-1)x^2} \quad \text{where} \quad x = \frac{\sum_{r=1}^{K-2} \binom{K-2}{r} \lambda_r}{\sum_{r=1}^{K-1} \binom{K-1}{r} \lambda_r} + \frac{\sum_{r=2}^K \binom{K-2}{r-2} \lambda_r}{\sum_{r=1}^K \binom{K-1}{r-1} \lambda_r} - 1 \quad (28)$$

In order to locate the other phase transitions, α_d and α_c , we have solved the cavity equations (23)-(27) using the population dynamics [27].

In Fig. 1 we plot the full phase diagram for quietly planted 3-SAT. In this case the vector λ^* has only one free parameter, we chose it to be λ_2 , the other two components are then given from conditions (6) and (16). The phase transition is continuous, i.e. $\alpha_d = \alpha_c = \alpha_l$, for $\lambda_2^* > 0.118(3)$. It is discontinuous, i.e. $\alpha_d < \alpha_c < \alpha_l$, for λ_2^* smaller than this threshold. Within our numerical precision this tri-critical point in random 3-SAT agrees with the value of $\lambda_2^* = 0.1180$ corresponding to the γ^* reweighting (7). We did not find a theoretical reason for this.




Figure 1. The phase diagram of quietly-planted reweighted 3-SAT in the plane of the reweighting parameter λ_2^* and the density of constraints α . We remind that the other two reweighting parameters λ_1^* and λ_3^* are obtained using conditions (6) and (16). Note that the boundary value $\lambda_2^* = 0$ in some sense reduces the problem into planted 3-XOR-SAT, and $\lambda_2^* = 1/6$ to the planted NAE-SAT. Both these cases were studied in [30, 10] and phase transitions in these cases were known (the SAT/UNSAT transition for these problems is shown as a red cross). The SAT/UNSAT transition does not depend on λ_2^* , except for these two boundary cases. The black dot corresponds to the phase transition for the power-law reweighted ensemble with parameter γ^* , which is indistinguishable within our accuracy from the point where the transition goes from a first order (dash-dotted blue line) to a second order (full black line) one. The curve α_l is dashed in the region where it corresponds to a spinodal of the first order phase transition and not the second order phase transition. The shaded region marks the part of the phase diagram where planting with the corresponding λ^* creates extremely hard (as will be quantified in section 6.3) 3-SAT satisfiable instances.

For the general case of $K > 3$ we keep for simplicity to the power-law reweighting of (7) with the condition (17), as it was used in [1, 36]. Fig. 2 shows, in the case $K = 4$, the properties of the BP-fixed points, starting from the two possible initial conditions. The overlap of the resulting fixed points with the planted configurations allows to find the threshold α_d . The difference of the Bethe entropies of the paramagnetic and ferromagnetic fixed point allows to locate the first-order phase transition α_c .




Figure 2. Phase diagram of the γ^* planted 4-SAT, obtained by the population dynamics solution of the cavity equations, population size was $5 \cdot 10^4$. The overlap between the planted configurations and the BP fixed point when BP is initialized a) randomly - blue curve - showing a phase transition at α_d , b) in the planted configuration - green curve - showing a phase transition at α_l . A careful reader will remark that the blue curve jumps up in the overlap a bit before the constraint density α_l , this is due to a strong finite size correction associated to this phase transition. The red curve is the difference between the paramagnetic free-entropy (18) and the free-entropy corresponding to the BP fixed point obtained from the planted initialization. It becomes positive at the first order phase transition $\alpha = \alpha_c$: for $\alpha < \alpha_c$ the paramagnetic phase is the dominant one, for $\alpha > \alpha_c$ the ferromagnetic phase is the dominant one.

Table 1 summarizes the values of these phase transitions for different values of K . For comparison we also give the dynamical and condensation phase transitions in the canonical non-reweighted K -SAT that corresponds to $\hat{\lambda}_r = 1/(2^K - 1)$ for all $r = 1, \dots, K$.

6. Results unveiled by quiet planting

6.1 Non-trivial whitening and BP fixed points

In the standard K -SAT problem, without reweighting or planting, the iteration of BP equations starting from a generic random initial condition typically have two possible behaviors, depending on the value of α . Either they always converge to the same fixed-point, that we call the “trivial” fixed point, or they do not converge at all (as in 3-SAT for $\alpha > 3.86$). On the other hand, the statistical physics description of the clustered phase starts from the assumption that there should exist one different and non-trivial BP fixed point for each cluster of solutions, and then counts the number of non-trivial BP fixed points via

K	α_s	γ^*	$\alpha_d^{\gamma^*}$	$\alpha_c^{\gamma^*}$	$\alpha_l^{\gamma^*}$	$\alpha_d^{\hat{\lambda}}$	$\alpha_c^{\hat{\lambda}}$	$\alpha_{BP(\hat{\lambda})}^{\gamma^*}$	$\alpha_{2nd}^{\gamma^*}$	$\alpha_f^{\gamma^*}$
3	4.27	0.61803	2.991	2.991	2.991	3.86	3.86	3.93	2.574	4.29
4	9.93	0.83929	8.01	8.40	10.91	9.38	9.547	8.75	7.313	9.63
5	21.12	0.92756	17.2	19.30	35.40	19.16	20.80	17.7	17.617	19.4
6	43.4	0.96595	34.5	41.3	111.1	36.53	43.08	34.7	39.026	37.7

Table 1. For various values of K we give the satisfiability threshold α_s [24], the “quiet” value of the parameter γ^* , the corresponding dynamical, $\alpha_d^{\gamma^*}$, condensation, $\alpha_c^{\gamma^*}$, and spinodal, $\alpha_l^{\gamma^*}$, thresholds for the reweighted probability measure $\mu(\gamma^*)$ in random K -SAT. Recall that quiet planting generates typical random graphs for density of constraints smaller than $\alpha_c^{\gamma^*}$. For comparison we remind the dynamical and condensation thresholds for the K -SAT with no reweighting, i.e. for the measure $\mu(\hat{\lambda})$, from [18]. The column $\alpha_{BP(\hat{\lambda})}^{\gamma^*}$ gives the density of constraints beyond which even canonical BP (at $\hat{\lambda}$) converges to a non-trivial fixed point if initialized in the planted configuration. The $\alpha_{2nd}^{\gamma^*}$ column is taken from [37, 1], it is the rigorous lower bound on the satisfiability threshold that is obtained by computation of the reweighted second moment. The work of [37, 1] also implies that the annealed free-entropy density equals the quenched one in the range $\alpha < \alpha_{2nd}^{\gamma^*}$. Our arguments lead us to conjecture that actually the annealed free-entropy density equals the quenched one in the whole range $\alpha \leq \alpha_c^{\gamma^*}$. The last column $\alpha_f^{\gamma^*}$ is the constraint density beyond which the planted solution lies in a frozen cluster.

an augmented-BP-like approach. This counting finds an exponentially large number of BP fixed points, indicating an exponentially large number of clusters [31, 25].

However, on the algorithmic size, working with instances of large sizes (think of e.g. $N > 10^4$), even when we are able to find satisfiable assignments in the clustered region, the BP algorithm initialized in that assignment either converges back to the trivial fixed point, or (in 3-SAT) does not converge at all, in apparent contradiction with the basic assumption. This long-lasting open question had a correspondence in the coloring of random graphs. In the case of coloring it was resolved in [20, 50] where it was explained using the cavity method that the solutions found by slow simulated annealing indeed do not have a corresponding BP fixed point, whereas BP initialized in a random solution (obtained via quiet planting) leads indeed to a non-trivial BP fixed point. Whereas one could conjecture that the same reason probably applies to K -SAT there was no tool to check this explicitly. The quiet planting that we study in this paper permits such an explicit check for a first time.

To construct a large random K -SAT formula and a non-trivial BP fixed point we proceed as follows: A planted graph at λ^* is equivalent to non-planted random graph as long as the constraint density is $\alpha < \alpha_c(\lambda^*)$. This gives us a tool to generate both a typical instance of K -SAT and one of its solutions. We then initialize BP from the planted solution, and then iterate the canonical BP equations (non-reweighted, i.e. BP at $\hat{\lambda}$), we observe that this canonical BP converges to a non-trivial fixed point whenever $\alpha > \alpha_{BP(\hat{\lambda})}(\gamma^*)$. From Table 1 we see that $\alpha_{BP(\hat{\lambda})}(\gamma^*) < \alpha_c(\gamma^*)$ for $K \geq 5$. This shows that for $K \geq 5$ we have a non-empty interval of constraint densities where a non-trivial BP fixed point exist on a random formula. Note, however, that in this case (unlike for the coloring) the planted

configuration is not a typical configuration with respect to the uniform measure over all solutions $\mu(\hat{\lambda})$, therefore the corresponding non-trivial BP fixed point does not describe a random cluster relative to the original, flat, measure. Note also that the fact that for $K > 3$ one has $\alpha_{BP(\hat{\lambda})}(\gamma^*) < \alpha_d(\hat{\lambda})$ is a direct evidence for the presence of small subdominant clusters in the paramagnetic region where a single large cluster still dominates the flat $\mu(\hat{\lambda})$ measure.

Freezing of variables was argued to be an important ingredient in understanding the relation between the structure of the space of solutions and the algorithmic hardness of random constraint satisfaction problems [48]. A variable is frozen in a cluster if it takes the same value in all solutions belonging to that cluster. Within the assumptions of the cavity method one can investigate if a solution belongs to a frozen cluster or not by iterating the BP equations initialized in the solution and monitoring whether for some nodes i the BP messages stay completely polarized (i.e. $\chi_1^i = 1$ or $\chi_0^i = 1$). In fact a simpler version of BP can be written for such monitoring, it is called “warning propagation” or “whitening” in the literature [42, 8, 9, 23]. There is a long-standing open question in the field which concerns this whitening [48]. Whereas, for instance, the survey-propagation algorithm is based on the existence of solutions with a non-trivial whitening result, when one tries to run the whitening algorithm on solutions found by heuristic algorithms on large (think e.g. of $N > 10^4$) K -SAT instances, the whitening result was always observed to be trivial. In [50] it was shown for the coloring problem that random solutions do indeed have a non-trivial whitening, as expected from the theoretical calculations. This same work has also given theoretical reasons explaining why the solutions found by local algorithms do not have a corresponding non-trivial whitening. Now we show that random K -SAT formulas do have solutions with a non-trivial whitenings.

Importantly, warning propagation has no dependence on the reweighting λ , hence if a cluster is frozen at one value of λ it must be frozen for all λ . Table 1 gives the values of constraint density α beyond which the planted solution lies in a frozen cluster. We use the quiet planting procedure, which generates a typical instance in which the planted solution is also typical, as long as $\alpha < \alpha_c(\gamma^*)$. We see that, for $K \geq 6$, the planted solution lies in a frozen cluster on typical random instances of K -SAT (in the range $\alpha_f(\gamma^*) < \alpha < \alpha_c(\gamma^*)$).

6.2 Direct evidence for purely entropic barriers

Often researchers define clusters not via BP fixed points, but as connected components in an auxiliary graph where vertices are satisfying configurations and edges are drawn between vertices that differ in the value of only one (or some other sublinear in N number of) variable. It has been pointed out on many occasions that there might be important differences between these two definitions. The first definition is more meaningful from the physics point of view, because the BP fixed points are assumed to correspond to basins of attraction of local Monte Carlo Markov chains with detailed balance condition (such MCMC are used to model the behavior of true physical dynamics in materials). A basin of attraction is defined as a part of the configurational space in which the MCMC will stay blocked for a time that grows with the size of the system faster than a power law. Actually, it is often argued that the escape time should grow exponentially with system size N , based on a computation of a large deviation function (called the free energy): this




Figure 3. Non-trivial whitening in the random 6-SAT problem: here $N = 10^5$, the instance was created by planting, but for constraint density where planting provides instances equivalent to random ones. When $\alpha \geq 38$ a non trivial whitening core is obtained for the planted solution. No solution to such a large instance of random K -SAT problem with a non-trivial whitening core had ever been found so far.

function has a minimum corresponding to the cluster (the basin of attraction) in question and all the barriers around this minimum scale linearly with N . Since the MCMC acts as a kind of random walk in the landscape given by this free energy function, it takes with high probability an exponential time to find such a path (in the same way it would take exponential time to obtain a configuration on binary variables with more than 60% of ones if one started from a random string of N elements and flipped variables randomly, one by one).

The barrier in this free energy can have only two possible origins (in general it is a combination of the two)

- **Energetic barrier:** Consider all possible paths from one solution to another (one step being a flip of one variable), if each such path violates at least B clauses at some step, then we say there is an “energetic barrier” of size B between the two solutions.
- **Entropic barrier:** Consider an exponentially large subset of all solutions (covering an exponentially small fraction of all solutions), even if the energetic barrier between this subset and most of the other solutions is zero it can still be that a random walk will spend exponential time before escaping from the subset in question. Such a situation is called an entropic barrier.

The existence of clusters in the “energetic” sense, i.e. with the energetic barriers between cluster growing linearly with N , was proved for K -SAT with large K in [26, 3]. It is perfectly plausible that clusters with purely entropic barriers exist. Such a situation was described in simple spin glass models (see for instance [5, 45, 39]), but so far it was not found in a random constraint satisfaction problem.

Using quiet planting in K -SAT we show that in some region of parameters there are purely entropic clusters in random K -SAT. Consider a planted instance in the regime where

the following three conditions are met: there is a planted cluster, the planting is quiet $\alpha_c(\gamma^*) > \alpha > \alpha_d(\gamma^*)$ and simultaneously the canonical BP converges to a trivial fixed point $\alpha < \alpha_{BP(\hat{\lambda})}(\gamma^*)$ (for values see Table 1). Recall that the set of satisfying configurations (“solutions”) is independent of the reweighting λ . If we initialize the BP reweighted by γ^* in the planted configuration the iterations will converge to a non-trivial fixed point, meaning that there is a planted cluster in the measure reweighted $\mu(\gamma^*)$. If we run canonical non-reweighted BP the iterations will converge to the factorized fixed points (from both initializations - in the planted configurations and in the fixed point of the reweighted BP), meaning that there is no planted cluster in the flat measure $\mu(\hat{\lambda})$. Hence the dynamics of non-reweighted BP was able to escape from the region that formed a cluster in the measure $\mu(\gamma^*)$.

Relying on the conjectured relation between BP fixed points and basins of attraction of Monte Carlo Markov chains the above result translates into the following numerical experiment with a MCMC random walk in the space of solutions.

First, initialize in the planted configuration, consider a MCMC random walk that satisfies the detailed balance condition with respect to the measure $\mu(\lambda^*)$. This random walk will be restrained in proximity of the planted configuration and stay there for a very long time. We conjecture that this time is exponential in the size of the system based on the possible calculation of the barrier in the free energy and the relation between clusters and fixed points of BP equations. We do not attempt to have a full numerical confirmation of such exponential scaling: in the numerical experiment of Fig. 4 we content ourselves with a time long enough to see a saturation of the overlap with the initial configuration in the logarithmic time-scale.

In a second part of this numerical experiment, consider a Monte Carlo that satisfies the detailed balance with respect to the $\mu(\hat{\lambda})$ measure. This MCMC will be able to walk far away from the planted configuration in relatively short time. This means that it manages to find some paths that were very rare in the measure $\mu(\lambda^*)$ but are now easy to find in $\mu(\hat{\lambda})$. Using such paths, the second procedure escapes to a large distance from the planted configuration. These experiments are illustrated in Fig. 4. This provides a strong empirical evidence for the existence of purely entropic barriers.

6.3 How hard is λ^* -planted K -SAT

Hardness of the γ^* -planted (we remind that γ^* is only a special point in the space of λ^*) random 3-SAT instances was investigated in [36] for DPLL, WalkSAT and Survey Propagation algorithm. The authors found empirically a region where all the three algorithms fail or scale exponentially with the size of the system. Our results actually show that for simulated annealing or for belief propagation decimation random γ^* -planted 3-SAT is easy. Indeed, for $\alpha > \alpha_l(\gamma^*)$ the BP and MCMC dynamics is attracted close to the planted configuration and hence finding a solution nearby is easy. For $\alpha < \alpha_l(\gamma^*)$ the planted formula is equivalent to a random formula, and it is so sparse that MCMC will equilibrate in linear time and hence simulated annealing will find a solution in linear time, in this region also BP decimation works [34, 44]. On the other hand for $K \geq 4$ the γ^* -planted K -SAT has a hard region for values of $\alpha_s < \alpha < \alpha_l(\gamma^*)$. Indeed for $\alpha < \alpha_l(\gamma^*)$ the planted cluster is hidden to the MCMC and BP dynamics and for $\alpha > \alpha_s$ there are no solutions other than




Figure 4. Illustration of the existence of purely entropic barriers: a solution is planted using γ^* , $\alpha = 8.2$, $N = 10^5$, $K = 4$. Then a Monte-Carlo simulation satisfying detailed balance for the measure $\mu(\lambda^*)$ is performed starting from this planted configuration: the system is trapped in a cluster and does not exit it for a very long time. Switching then to a new Monte Carlo, which satisfied detailed balance for the measure $\mu(\hat{\lambda})$, i.e. the flat measure over all solutions, makes it very easy for the MCMC to exit this cluster and decorrelate from the initial solution. This gives a strong evidence for the presence of entropic barriers and the fact that the dynamics can be trapped even in absence of frozen variables or energetic barriers.

those belonging to the planted cluster. At this point we should mention that whereas we do not make any claim about hardness for an arbitrary algorithm, the class of algorithms for which the planted cluster is hard to find for $\alpha > \alpha_s$ seems to be quite large, it includes the local search algorithms, as well as message-passing ones. But for instance it would not include algorithms such as Gaussian elimination for constraint satisfaction problems that are linear over some field, such as the K -XOR-SAT problem.

The algorithmic hardness of λ^* -planted 3-SAT formulas was studied in [6] using a walk-SAT algorithm. We revisit the corresponding phase diagram in view of our results. We show that the conclusions of [6] were qualitatively correct, but the correct boundary of the “hard” region is different from what was estimated in [6] (as the statistical physics calculations in that work were only approximate). Most importantly we give further theoretical justification for why the λ^* -planted random 3-SAT formulas are (together with benchmarks from [16]) the hardest satisfiable formulas on the market of hard satisfiable benchmarks.

We remind that the space of solution in the λ^* -planted formulas can be split in two parts. The first part includes all the solutions (satisfying assignments) that would exist anyway in the non-planted random 3-SAT. The second part includes solutions correlated to the planted configuration. Algorithmic search for a solution belonging to the first part is just as hard (or as easy) as it is in the canonical random K -SAT. Notably for $\alpha > \alpha_s$ this space is empty (canonical random SAT is unsatisfiable). It follows from our results above that for $\alpha > \alpha_l$ it is easy to find a solution correlated with the planted configuration, to do so one can use λ^* -reweighted BP or λ^* -reweighted MCMC. We argue that, for $\alpha < \alpha_l$, finding solutions correlated to the planted one is hard. Finding a configuration (not necessarily a solution) correlated to the λ^* -planted configuration may be seen as an inference problem

and running λ^* -reweighted MCMC is then a Bayes-optimal algorithm for this inference problem. The cavity method predicts that a barrier of size linear in the system size would have to be overcome in order to find a configuration correlated to the planted one. Hence the running time of the MCMC would become exponentially large below the transition α_l (and above α_d). Concerning algorithms such as BP, the metastable state in which MCMC is blocked corresponds to a stable BP fixed point and one would have to initialize the BP messages extremely close to the planted configuration to find another fixed point. With random initializations this event would be exponentially rare. Clearly if it is exponentially hard to find any *configuration* correlated to the planted configuration, it will be even harder to find a *solution* correlated to the planted configuration. To conclude, the boundary of the hard region is given by $\alpha < \alpha_l$ in order not to find the planted configuration, and at the same time α sufficiently large for the canonical random K -SAT to be hard or better not to have any solutions at all for $\alpha > \alpha_s$. This region is depicted by shading in Fig. 1 for 3-SAT.

Let us also describe explicitly the relation between the benchmarks described here and those of [16]. Instances generated by the λ^* -planting with λ_2^* close to zero can be seen as a planted XOR-SAT formulas with “a bit of nonlinearity”. If moreover we restrict ourselves to the regular formulas where every variable appears in the same number of clauses then we are exactly in the setting of [16] who introduced planted regular XOR-SAT formulas with “a bit of nonlinearity” as the hardest known satisfiable benchmarks. Moreover the kind of nonlinearity that quiet planting adds by taking a nonzero value of λ_2^* is not easy to discover, even if the exact protocol of how the instances were created was known. Note that the regular instances are taken in order to lower the total number of variables (e.g. the leaves of random instances do not contribute to the overall hardness and can hence be omitted). The analysis presented in this paper can be used for regular instances straightforwardly.

7. Conclusions

In this paper, we have studied the reweighted measure over solutions of the random K -satisfiability problem, where each solution is reweighted according to the number of variables that satisfy every clause. This problem has been addressed both analytically, using the cavity method, and numerically, using belief propagation and Monte-Carlo Markov chains.

The main results of this study are the following: (i) The reweighting allows to introduce a planted ensemble that generates instances that are equivalent to random instances in some region of the phase diagram. In this case, we are thus able to generate simultaneously a typical SAT instance and one of its solutions. (ii) We have used this property to give a strong evidence for the existence of purely entropic (rather than energetic) barriers between clusters in some region of the phase diagram. This is a fundamental point demonstrating that the physical definition of clusters as basins of attraction (or belief propagation fixed points) is *not* equivalent to the geometric definition of clusters as disconnected components in the solution space. Such equivalence was often wrongly assumed in the literature and leads to confusions. (iii) We have used the quiet planting property to display explicitly, for the first time, solutions of large random K -SAT problems leading to a non-trivial whitening core: while such solutions were known to exist, they were so far never observed on very large instances. (iv) We discuss the algorithmic hardness of these planted instances and determine a region of parameter in which quiet planting leads to hard satisfiable benchmarks.

References

- [1] D. Achlioptas and Y. Peres. The threshold for random k -sat is $2^k \log 2 - O(k)$. *Journal of the AMS*, **17**:947–973, 2004.
- [2] Dimitris Achlioptas and Amin Coja-Oghlan. Algorithmic barriers from phase transitions. In *IEEE 49th Annual IEEE Symposium on Foundations of Computer Science, FOCS’08*, pages 793 – 802, 2008. arXiv:0803.2122v2.
- [3] Dimitris Achlioptas and Federico Ricci-Tersenghi. On the solution-space geometry of random constraint satisfaction problems. In *Proc. of 38th STOC*, pages 130–139, New York, NY, USA, 2006. ACM.
- [4] John Ardelius and Lenka Zdeborová. Exhaustive enumeration unveils clustering and freezing in random 3-sat. *Phys. Rev. E (Rapid Com.)*, **78**, 2008.
- [5] A. Barrat and M. Mézard. Phase space diffusion and low temperature aging. *J. Phys. I (France)*, **5**:941–947, 1995.
- [6] W. Barthel, A. K. Hartmann, M. Leone, F. Ricci-Tersenghi, M. Weigt, and R. Zecchina. Hiding solutions in random satisfiability problems: A statistical mechanics approach. *Phys. Rev. Lett.*, **88**:188701, 2002.
- [7] G. Biroli, R. Monasson, and M. Weigt. A variational description of the ground state structure in random satisfiability problems. *Eur. Phys. J. B*, **14**:551, 2000.
- [8] A. Braunstein, M. Mézard, M. Weigt, and R. Zecchina. Constraint satisfaction by survey propagation. In Allon Percus, Gabriel Istrate, and Cristopher Moore, editors, *Computational Complexity and Statistical Physics*, page 107. Oxford University Press, 2003.
- [9] A. Braunstein and R. Zecchina. Survey propagation as local equilibrium equations. *Journal of Statistical Mechanics: Theory and Experiment*, page P06007, 2004.
- [10] Tommaso Castellani, Vincenzo Napolano, Federico Ricci-Tersenghi, and Riccardo Zecchina. Bicoloring random hypergraphs. *J. Phys. A*, **36**:11037, 2003.
- [11] Peter Cheeseman, Bob Kanefsky, and William M. Taylor. Where the Really Hard Problems Are. In *Proc. 12th IJCAI*, pages 331–337, San Mateo, CA, USA, 1991. Morgan Kaufmann.
- [12] Amin Coja-Oghlan and Lenka Zdeborová. The condensation transition in random hypergraph 2-coloring. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 241–250. SIAM, 2012.
- [13] Amin Coja-Oglan and Konstantinos Panagiotou. Catching the k -naesat threshold. In *Proceedings of the 44th symposium on Theory of Computing*, pages 899–908. ACM, 2012.
- [14] Stephen A. Cook. The complexity of theorem-proving procedures. In *Proc. 3rd STOC*, pages 151–158, New York, NY, USA, 1971. ACM.

- [15] E. Friedgut. Sharp thresholds of graph properties, and the k -sat problem. *J. Amer. Math. Soc.*, **12**, 1999.
- [16] Harri Haanpää, Matti Järvisalo, Petteri Kaski, and Ilkka Niemelä. Hard satisfiable clause sets for benchmarking equivalence reasoning techniques. *Journal on Satisfiability, Boolean Modeling and Computations*, **2**:27–46, 2006.
- [17] S. Kirkpatrick and B. Selman. Critical behavior in the satisfiability of random boolean expression. *Science*, **264**:1297–1301, 1994.
- [18] F. Krzakala, A. Montanari, F. Ricci-Tersenghi, G. Semerjian, and L. Zdeborová. Gibbs states and the set of solutions of random constraint satisfaction problems. *Proc. Natl. Acad. Sci. U.S.A.*, **104**:10318, 2007.
- [19] Florent Krzakala and Lenka Zdeborová. Hiding quiet solutions in random constraint satisfaction problems. *Phys. Rev. Lett.*, **102**:238701, 2009.
- [20] Florent Krzakala and Lenka Zdeborová. Following gibbs states adiabatically the energy landscape of mean-field glassy systems. *EPL*, **90**:66002, 2010.
- [21] F. R. Kschischang, B. Frey, and H.-A. Loeliger. Factor graphs and the sum-product algorithm. *IEEE Trans. Inform. Theory*, **47**(2):498–519, 2001.
- [22] Leonid A Levin. Universal sequential search problems. *Problemy Peredachi Informatsii*, **9**(3):115–116, 1973.
- [23] Elitza N. Maneva, Elchanan Mossel, and Martin J. Wainwright. A new look at survey propagation and its generalizations. *J. ACM*, **54**(4), 2007.
- [24] Stephan Mertens, Marc Mézard, and Riccardo Zecchina. Threshold values of random k -sat from the cavity method. *Random Struct. Algorithms*, **28**(3):340–373, 2006.
- [25] M. Mézard and A. Montanari. *Physics, Information, Computation*. Oxford University Press, Oxford, 2009.
- [26] M. Mézard, T. Mora, and R. Zecchina. Clustering of solutions in the random satisfiability problem. *Physical Review Letters*, **94**:197205, 2005.
- [27] M. Mézard and G. Parisi. The bethe lattice spin glass revisited. *Eur. Phys. J. B*, **20**:217, 2001.
- [28] M. Mézard, G. Parisi, and M. A. Virasoro. *Spin-Glass Theory and Beyond*, **9** of *Lecture Notes in Physics*. World Scientific, Singapore, 1987.
- [29] M. Mézard, G. Parisi, and R. Zecchina. Analytic and algorithmic solution of random satisfiability problems. *Science*, **297**:812–815, 2002.
- [30] M. Mézard, F. Ricci-Tersenghi, and R. Zecchina. Alternative solutions to diluted p -spin models and XORSAT problems. *J. Stat. Phys.*, **111**:505, 2003.

- [31] M. Mézard and R. Zecchina. Random k -satisfiability problem: From an analytic solution to an efficient algorithm. *Phys. Rev. E*, **66**:056126, 2002.
- [32] David G. Mitchell, Bart Selman, and Hector J. Levesque. Hard and easy distributions for SAT problems. In *Proc. 10th AAAI*, pages 459–465, Menlo Park, California, 1992. AAAI Press.
- [33] Michael Molloy. The freezing threshold for k -colourings of a random graph. In *Proceedings of the 44th symposium on Theory of Computing*, pages 921–930. ACM, 2012.
- [34] A. Montanari, F. Ricci-Tersenghi, and G. Semerjian. Solving constraint satisfaction problems through belief propagation-guided decimation. *Proceedings of Allerton 2007*. Preprint: arXiv:0709.1667v1 [cs.AI], 2007.
- [35] A. Montanari, F. Ricci-Tersenghi, and G. Semerjian. Clusters of solutions and replica symmetry breaking in random k -satisfiability. *J. Stat. Mech.*, page P04004, 2008.
- [36] C. Moore, H. Jia, and D. Strain. Generating hard satisfiable formulas by hiding solutions deceptively. In *Proc. of the National Conference on Artificial Intelligence*, **20**, page 384, 2005. arXiv:cs/0503044v1 [cs.AI].
- [37] C. Moore and S. Mertens. *The Nature of Computation*. Oxford University Press, 2011.
- [38] T. Mora. *Géométrie et inférence dans l’optimisation et en théorie de l’information*. PhD thesis, Université Paris-Sud, 2007. <http://tel.archives-ouvertes.fr/tel-00175221/en/>.
- [39] T. Mora and L. Zdeborová. Random subcubes as a toy model for constraint satisfaction problems. *J. Stat. Phys.*, **131**(6):1121–1138, 2008. arXiv:0710.3804v2 [cs.CC].
- [40] Elchanan Mossel, Joe Neeman, and Allan Sly. Stochastic block models and reconstruction. arXiv:1202.1499v3 [math.PR], 2012.
- [41] C. H. Papadimitriou. *Computational complexity*. Addison-Wesley, 1994.
- [42] G. Parisi. On the survey-propagation equations in random constraint satisfiability problems. *J. Math. Phys.*, **49**, 2008. arXiv:cs.CC/0212009.
- [43] J. Pearl. Reverend bayes on inference engines: A distributed hierarchical approach. In *Proceedings American Association of Artificial Intelligence National Conference on AI*, pages 133–136, Pittsburgh, PA, USA, 1982.
- [44] F. Ricci-Tersenghi and G. Semerjian. On the cavity method for decimated random constraint satisfaction problems and the analysis of belief propagation guided decimation algorithms. *J. Stat. Mech.*, page P09001, 2009.
- [45] Felix Ritort. Glassiness in a model without energy barriers. *Phys. Rev. Lett.*, **75**:1190–1193, Aug 1995.
- [46] Sakari Seitz, Mikko Alava, and Pekka Orponen. Focused local search for random 3-satisfiability. *J. Stat. Mech.*, page P06006, 2005.

- [47] Guilhem Semerjian. On the freezing of variables in random constraint satisfaction problems. *J. Stat. Phys.*, **130**:251, 2008.
- [48] L. Zdeborová. Statistical physics of hard optimization problems. *Acta Physica Slovaca*, **59**(3):169–303, 2009.
- [49] L. Zdeborová and F. Krzakala. Phase transitions in the coloring of random graphs. *Phys. Rev. E*, **76**:031131, 2007.
- [50] Lenka Zdeborová and Florent Krzakala. Generalization of the cavity method for adiabatic evolution of gibbs states. *Phys. Rev. B*, **81**:224205, 2010.
- [51] Lenka Zdeborová and Florent Krzakala. Quiet planting in the locked constraint satisfaction problems. *SIAM Journal of Discrete Mathematics*, **25**:750–770, 2011.