

Two Solutions to Diluted p -Spin Models and XORSAT Problems

M. Mézard,¹ F. Ricci-Tersenghi,² and R. Zecchina³

Received October 1, 2002; accepted December 10, 2002

We derive analytical solutions for p -spin models with finite connectivity at zero temperature. These models are the statistical mechanics equivalent of p -XORSAT problems in theoretical computer science. We give a full characterization of the phase diagram: location of the phase transitions (static and dynamic), together with a description of the clustering phenomenon taking place in configurational space. We use two alternative methods: the cavity approach and a rigorous derivation.

KEY WORDS: Spin glass; satisfiability; leaf removal; cavity method.

1. INTRODUCTION

The very last years have seen a growth of interest in disordered models defined on Bethe-lattices-like topologies, that is finite connectivity random graphs (see, e.g., ref. 1). Appropriate generalizations of mean-field theory are exact on such structures allowing for an exact solution of spin-glass like models. The presence of large loops may induce frustration leading to highly non trivial properties at low enough temperatures. Interacting models defined over finite connectivity graphs provide a better approximation to finite-dimensional models than fully connected mean-field models, allowing for qualitatively new effects to be discussed. At zero temperature,

¹Laboratoire de Physique Théorique et Modèles Statistiques, Université Paris Sud, 91405 Orsay, France; e-mail: mezard@ipno.in2p3.fr

²Dipartimento di Fisica and INFN, Università di Roma "La Sapienza," Piazzale Aldo Moro 2, I-00185 Roma (Italy); e-mail: Federico.Ricci@roma1.infn.it

³International Center for Theoretical Physics, Strada Costiera 11, P.O. Box 586, I-34100 Trieste, Italy; e-mail: zecchina@ictp.trieste.it

spin glass like models over random graphs correspond to some random combinatorial optimization problems of central relevance in theoretical computer science.⁽²⁾

Quite in general, spin glass models show an interesting phase diagram in the (γ, T) plane (see, e.g., Fig. 2 in ref. 3), where γ is a parameter proportional to the mean connectivity of the underlying random graph and T is the temperature. The frozen phase is located at high γ and/or low T .

Open questions are, for example, the exact location of the critical lines (dynamic and static ones), the full characterization of the configurational space in the frozen phase (e.g., ground state energy and threshold energies), etc.

Here we focus on the simplest non trivial model that can be defined on a random graph with finite mean connectivity, namely the p -spin model. We concentrate on the zero temperature limit, which corresponds to the p -XORSAT problem in theoretical computer science.⁽⁴⁾

The decision problem of XORSAT, which is equivalent to stating if a system of linear equations in GF[2] has a solution, is known to be a polynomial problem, but the optimization problem, where one wants to minimize the number of unsatisfied equations, is NP-hard, and inapproximable.⁽⁵⁾

In this limit the model undergoes two relevant phase transitions.⁽⁶⁾ The first one takes place at γ_d (for $p = 3$, $\gamma_d = 0.818469$) and corresponds to a clustering phenomenon: For $\gamma < \gamma_d$ all the ground states (GS) form a unique cluster, while for $\gamma > \gamma_d$ they split into an exponentially large (in N) number of clusters, each one containing an exponential number of GS. This clustering phenomenon coincides with the formation, in the configurational space, of barriers (clusters are well defined only because of the presence of barriers) and of metastable states, which make any greedy search algorithm inefficient. This is why it is usually called *dynamic transition*.⁽⁷⁾ The second phase transition takes place at $\gamma_c > \gamma_d$ (for $p = 3$, $\gamma_c = 0.917935$) and marks the SAT/UNSAT transition, that is the point where frustration becomes manifest in the system and the GS energy becomes larger than zero.

We will derive the above scenario via two distinct and complementary methods. The first one is the *cavity method*. Its power relies in its generality, since it can be applied easily to more complex systems too, e.g., random k -SAT.^(8,9) Within this method the above scenario can be obtained using an Ansatz with a single step of replica symmetry breaking (1RSB). The second method is a *rigorous derivation* based on the “leaf removal” algorithm which is able to reduce the random (hyper)graph to its relevant *core*. On the core, any interesting quantity (e.g., the number of GS, cluster size and distance) can be easily calculated, since annealed averages coincide with quenched ones.

This rigorous derivation is of great importance also because this is one of the few cases⁽¹⁰⁾ where a highly non trivial scenario, previously obtained with a replica calculation,^(6, 11, 12) can be confirmed with rigorous methods. These results confirm the validity of the cavity approach, and may open the way towards the construction of mathematical bases for the Parisi's replica symmetry breaking theory.⁽¹³⁾

2. DEFINITION OF THE MODEL

The random p -XORSAT problem consists in finding an assignment to N boolean variables $x_i \in \{0, 1\}$, such that a set of $M = \gamma N$ parity checks on these variables are satisfied. Each parity check is of the kind

$$x_{i_1^m} + \cdots + x_{i_p^m} = y_m \pmod{2}, \quad m = 1, \dots, M \quad (1)$$

where, for each m , the p indices $i_1^m, \dots, i_p^m \in \{1, \dots, N\}$ are chosen randomly and uniformly among the $\binom{N}{p}$ possible p -uples of distinct indices, and the "coupling" y_m takes randomly value 0 or 1 with equal probability. The above set of constraints can be written in a more compact way as $\hat{A}\vec{x} = \vec{y} \pmod{2}$, where \hat{A} is a $M \times N$ random sparse matrix with exactly p ones per row and y is a random vector of 0s and 1s.

Once the mapping $\sigma = (-1)^x$ and $J = (-1)^y$ is performed, the XORSAT problem can be also studied as the zero-temperature limit of the following p -spin Hamiltonian giving the energy for a configuration of N Ising spins $\sigma_i \in \{-1, 1\}$:

$$\mathcal{H} = \sum_{m=1}^M (1 - J_m \sigma_{i_1^m} \cdots \sigma_{i_p^m}). \quad (2)$$

Unfrustrated ground states (GS) configurations have zero energy and correspond to solutions of the XORSAT problem, since they satisfy all the constraints: $\forall m, \sigma_{i_1^m} \cdots \sigma_{i_p^m} = J_m$.

3. $T=0$ PHASE DIAGRAM FROM THE ONE-STEP CAVITY METHOD

In this section we shall display the analysis of the phase diagram of the p -spin problem as it arises from the one-step cavity approach. We consider the cavity formalism directly at zero temperature as discussed by Mézard and Parisi^(1, 14) and developed further in ref. 9. We refer to those papers for a review of the method and the notations. Here we shall limit ourselves to the technical aspects of the analytical calculation for $p=3$ case, generalizations to $p > 3$ being straightforward.

The zero temperature p -spin model can be viewed as a relatively simple limit case of more general problems such as random k -SAT for which the cavity calculations have also been carried out recently.⁽⁹⁾ The main technical difference between random k -SAT like problems and the p -spin model consists in the fact that the site dependence of the functional order parameter simplifies dramatically in the p -spin problem below the static transition. This allows for a rigorous derivation of the cavity and replica results by alternative methods, as we shall thoroughly discuss in the subsequent sections.

In the cavity formalism⁽¹³⁾ one works with “cavity fields” h_i associated to the sites and “cavity biases” u_j associated to the hyperedges. The cavity field is the effective field on a variable once one of its interactions has been removed. Under a cavity iteration, cavity biases generate cavity fields and vice versa (see Fig. 1). The cavity field h is always the sum of the cavity biases u coming from all its interactions, but the one removed. The rule for generating u biases from h fields is in general more complex.

For $T = 0$ the formalism simplifies a lot:⁽¹⁴⁾ Cavity fields and cavity biases only take integer values and the cavity equations can be derived easily by implementing the energy minimization condition under the cavity iteration. Let us imagine to add a hyperedge connecting 3 spins, say spins σ_0, σ_1 and σ_2 among which spin σ_0 plays the role of cavity spin. We need to perform a partial minimization of the effective energy

$$\min_{\sigma_1, \sigma_2} [\epsilon(\sigma_0, \sigma_1, \sigma_2) - (h_1\sigma_1 + h_2\sigma_2)] = -w_J(h_1, h_2) - u_J(h_1, h_2)\sigma_0, \quad (3)$$

where $\epsilon(\sigma_0, \sigma_1, \sigma_2) = 1 - J\sigma_0\sigma_1\sigma_2$. The above relation defines the cavity biases w and u as functions of the “input” cavity fields h . After a little algebra one finds

$$\begin{aligned} w_J(h_1, h_2) &= |h_1| + |h_2| - |u_J(h_1, h_2)|, \\ u_J(h_1, h_2) &= \mathcal{L}(Jh_1h_2), \end{aligned} \quad (4)$$




Fig. 1. A pictorial view of the cavity iteration: h_1 and h_2 cavity fields are the sum of some cavity biases u , and in turn they generate a new cavity bias u according to Eq. (3).

where the function $\mathcal{S}(x)$ is defined as

$$\mathcal{S}(x) = \begin{cases} \text{sign}(x) & \text{if } x \neq 0, \\ 0 & \text{if } x = 0. \end{cases} \quad (5)$$

The free-energy of the system can be expressed either in terms of probability distributions of the cavity fields or of the cavity biases.^(1, 9, 14)

In a one step scenario the phase space breaks into many pure states and the order parameter of the model is a complete histogram, over the system, of probability distribution functions of fields, $\mathcal{P}[P(h)]$, and biases, $\mathcal{Q}[Q(u)]$. Such a rich structure of the order parameter can be understood by noticing that each spin may fluctuate from state to state and therefore the whole collection of single site probability distributions might be needed to capture such fluctuations. In the simple case of a single pure state, the so called replica symmetric (RS) phase, single site probability distributions become delta functions and the order parameter simplifies to a single global probability distribution.

Following the general scheme discussed in refs. 1, 9, 8, and 14, but with a more convenient normalization for the $Q(u)$, the self-consistency equation for the $\mathcal{Q}[Q(u)]$ reads

$$\begin{aligned} Q(u) &= E_J \int dh P^{(k)}(h) dg P^{(k')}(g) \delta(u - u_J(h, g)) \\ &\quad \text{with prob. } e^{-3\gamma} \frac{(3\gamma)^k}{k!} e^{-3\gamma} \frac{(3\gamma)^{k'}}{k'!}, \\ P^{(k)}(h) &= \frac{1}{A_k} \int du_1 Q_1(u_1) \cdots du_k Q_k(u_k) \delta\left(h - \sum_{i=1}^k u_i\right) \\ &\quad \times \exp\left[-y \left(\sum_{i=1}^k |u_i| - \left|\sum_{i=1}^k u_i\right|\right)\right], \\ A_k &= \int du_1 Q_1(u_1) \cdots du_k Q_k(u_k) \exp\left[-y \left(\sum_{i=1}^k |u_i| - \left|\sum_{i=1}^k u_i\right|\right)\right], \end{aligned} \quad (6)$$

where all the $Q_i(u)$ on the r.h.s. are chosen randomly from the distribution $\mathcal{Q}[Q(u)]$. The average E_J over the coupling signs $J = \pm 1$ forces all the distributions to be symmetric under $u \leftrightarrow -u$ or $h \leftrightarrow -h$. The parameter y is the so called reweighting coefficient ($y = \beta m$ where m is the Parisi breaking parameter) which takes into account level crossing of states under the cavity iterations.⁽¹⁴⁾ The parameter y must be chosen such as to maximize the free energy.

As the cavity biases take values in $\{0, \pm 1\}$, and thanks to above mentioned symmetry, each $Q_i(u)$ can be written, in full generality, as

$$Q_i(u) = \eta_i \delta(u) + \frac{1-\eta_i}{2} [\delta(u+1) + \delta(u-1)]. \quad (7)$$

Thus the self-consistency equation for $\mathcal{Q}[Q(u)]$ can be rewritten as a self-consistency equation for the probability distribution of $\eta_i, \rho(\eta)$.

Eventually, given the whole set of stationary $\{Q_i(u)\}$ or the stationary $\rho(\eta)$, the average ground state energy and the complexity can be deduced from the formulae of refs. 1, 9, 8, and 14.

3.1. Solution of the Self-Consistency Equation

3.1.1. RS Solution

We first notice that it is always possible to get back the simple replica symmetric solution by fixing $y = 0$ and assuming that the cavity biases are ‘‘certain,’’ $Q_j(u) = \delta(u - u_j)$, where the u_j are independent and identically distributed random variables taken from a distribution

$$\mathcal{Q}(u) = c_0 \delta(u) + \frac{(1-c_0)}{2} [\delta(u-1) + \delta(u+1)]. \quad (8)$$

Plugging the above form into Eq. (6), one finds for c_0

$$\begin{aligned} 1 - \sqrt{1-c_0} &= e^{-3\gamma} \sum_k \frac{(3\gamma)^k}{k!} \sum_{q=0}^{\lfloor k/2 \rfloor} \binom{k}{2q} \binom{2q}{q} \left(\frac{1-c_0}{2}\right)^{2q} c_0^{k-2q} \\ &= e^{-3\gamma(1-c_0)} I_0[3\gamma(1-c_0)], \end{aligned} \quad (9)$$

where $\lfloor x \rfloor$ is the integer part of x . However, the above equation leads to wrong predictions: a solution different from the trivial paramagnetic one, $Q_j(u) = \delta(u)$, appears at $\gamma = 1.16682$ with a negative energy. At $\gamma_{RS} = 1.29531$ the energy becomes positive, giving a lower bound for the true energy of the system.

3.1.2. 1RSB Solution and the Existence of Non Trivial Fields

The numerical solution of Eq. (6) indicates that there exists a non-trivial solution in the region $\gamma \gtrsim 0.82$ for sufficiently large values of the reweighting y . A careful look at the numerics shows that the probability distributions of η_i takes the form

$$\rho(\eta) = t \delta(\eta-1) + (1-t) \tilde{\rho}(\eta), \quad (10)$$

that is a fraction t of cavity biases are trivial. The non trivial cavity biases are characterized by a distribution $\tilde{\rho}$ which shrinks in the limit of large y , converging to delta function in $\eta = 0$. The $y \rightarrow \infty$ limit is particularly relevant in the region up to γ_c .

3.1.3. The $y \rightarrow \infty$ Limit: The Complexity and the Location of the Phase Transition

Looking at the self-consistency equations (6), the only way one can obtain a non-trivial distribution $Q(u)$ on the l.h.s. is when both $P^{(k)}(h) \neq \delta(h)$ and $P^{(k')}(g) \neq \delta(g)$. Moreover the probability that $P^{(k)}(h) = \delta(h)$ equals the probability of picking up k trivial distributions $Q(u)$, i.e., t^k . Putting everything in formulae, one has

$$\begin{aligned} 1-t &= e^{-6\gamma} \sum_{k, k'=0}^{\infty} \frac{(3\gamma)^k}{k!} \frac{(3\gamma)^{k'}}{k'!} (1-t^k)(1-t^{k'}) \\ &\Rightarrow 1-t = (1-e^{-3\gamma(1-t)})^2. \end{aligned} \quad (11)$$

For $\gamma < \gamma_d = 0.818469$ the only solution is $t = 1$ (the system is a paramagnet) whereas above γ_d a non-trivial solution appears.

For $y = \infty$, a direct inspection of the numerical results shows that the cavity biases spontaneously divide in two categories, such that $\rho(\eta) = t \delta(\eta-1) + (1-t) \delta(\eta)$. In terms of $\mathcal{Q}[Q(u)]$ it corresponds to having

$$Q(u) = \begin{cases} \delta(u) & \text{with prob. } t \\ \frac{\delta(u-1) + \delta(u+1)}{2} & \text{with prob. } 1-t \end{cases} \quad (12)$$

which indeed is a fixed point under the iteration process (6) for $y = \infty$, provided the fraction of trivial biases t satisfies Eq. (11).

Using the expressions of refs. 1, 9, 8, and 14, for very large y , the free energy can be written as $\Phi(y) = \frac{\psi}{y}$. As expected, one finds that, as long as $\psi < 0$ the maximum of $\Phi(y)$ is located in $y = \infty$ and corresponds to a zero ground state energy. Consequently the complexity or configurational entropy of zero-energy states, i.e., the normalised logarithm of the number of solutions clusters is given by

$$\Sigma = -\psi = \log(2) \left[1 - \frac{\lambda}{3} - e^{-\lambda} \left(1 + \frac{2}{3} \lambda \right) \right], \quad (13)$$

where $\lambda = 3\gamma(1-t)$ satisfies the self-consistency equation

$$\lambda = 3\gamma(1-e^{-\lambda})^2. \quad (14)$$

The critical point, i.e., the SAT/UNSAT threshold, $\gamma_c = 0.917935$ can be found as the γ value where the complexity becomes zero (see Fig. 2). For $\gamma > \gamma_c$ the free energy $\Phi(\gamma)$ has a positive maximum in a finite value of y , which corresponds to a positive ground state energy.

3.1.4. Expansion at Large γ : The Ground State Energy in the UNSAT Phase (MAX-3-XORSAT)

In order to study the ground state energy for $\gamma > \gamma_c$ we need to take care of the leading corrections in the limit $y \gg 1$. For finite y , the distribution in Eq. (12) is no longer stable and we need to study a more general distribution of biases which takes care of the appearance of a non-trivial contribution to the peak in $u=0$ arising from frustrated interactions. This more general $\mathcal{Q}[Q(u)]$ is such that a fraction t of messages is still completely trivial, $Q(u) = \delta(u)$, while non-trivial messages comes from distributions of the following kind

$$Q(u) = \xi e^{-2y} \delta(u) + \frac{1 - \xi e^{-2y}}{2} [\delta(u-1) + \delta(u+1)]. \quad (15)$$

The factor e^{-2y} has been introduced in order to have finite ξ in the limit of very large y . Moreover, from numerical solution of Eq. (6) we observe that ξ takes only integers values. Let us call a_m the fraction of non-trivial distributions having $\xi = m$. The generating function $a(z) = \sum_m a_m z^m$ satisfies the equation

$$a(z) = [Aa(z) + Bz + 1 - A - B]^2, \quad (16)$$

where $A = \frac{e^{-\lambda}}{1 - e^{-\lambda}}$ and $B = \frac{e^{-\lambda} \lambda^2 / 2}{1 - e^{-\lambda}}$.




Fig. 2. The complexity as a function of γ .

Using the distributions in Eq. (15) one can obtain the free energy density $\Phi(y)$ up to the first correction

$$\Phi(y) = \frac{\psi}{y} - \frac{\omega}{y} e^{-2y} + \mathcal{O}(e^{-4y}), \quad (17)$$

where

$$\omega = \frac{\lambda}{3} \left[1 - e^{-\lambda} \left(1 + \frac{3}{2} \lambda + \lambda^2 \right) + \langle \xi \rangle (1 - e^{-\lambda} (1 + 2\lambda)) \right]. \quad (18)$$

The mean value of ξ can be easily obtained differentiating Eq. (16) with respect to z and then putting $z = 1$,

$$\langle \xi \rangle = a'(1) = \frac{2B}{1-2A} = \frac{e^{-\lambda} \lambda^2}{1 - e^{-\lambda} (1 + 2\lambda)}, \quad (19)$$

and thus we have

$$\omega = \frac{\lambda}{3} \left(1 - e^{-\lambda} \left(1 + \frac{3}{2} \lambda \right) \right). \quad (20)$$

Summarizing the statistical mechanics analysis, for any $\gamma > \gamma_d = 0.818469$, one can solve Eq. (14) for λ , deduce the large y behaviour of $\Phi(y)$ from Eq. (17) and maximize $\Phi(y)$ with respect to y . We find a critical value of γ , $\gamma_c = 0.917935$, where ψ changes sign. For $\gamma < \gamma_c$, $\psi < 0$ and therefore the maximum of $\Phi(y)$ is found at $y = \infty$. The distribution of cavity biases is given by Eq. (12), and the maximum value of Φ is 0, showing that all hyperedges are satisfied (apart from maybe a vanishing fraction at large N).

At $\gamma = \gamma_c$, ψ changes sign and, for $\gamma > \gamma_c$, $\Phi(y)$ has a maximum at a finite y , which shows that the ground state energy becomes strictly positive: It is no longer possible to satisfy simultaneously all the constraints.

The value of the energy for γ slightly above γ_c can be computed from the large y expansion. Moreover, such an expansion allows us to compute the complexity $\Sigma(E)$ of states of given energy E by a Legendre transformation of the free energy. The complexity function $\Sigma(E)$ is obtained by solving $E = \partial_y(y\Phi)$ and $\Sigma = y^2 \partial_y \Phi$.

From Eq. (17) for Φ we get

$$E = 2\omega e^{-2y} + \mathcal{O}(e^{-4y}), \quad (21)$$

$$\Sigma = -\psi + (2y + 1) \omega e^{-2y} + \mathcal{O}(e^{-4y}). \quad (22)$$




Fig. 3. Ground state energy for γ values above the critical point $\gamma_c = 0.917935$. Numerical data seem to converge to the analytic prediction. Finite-size corrections roughly decrease as $1/N$.

For $\gamma_d < \gamma < \gamma_c$, the constant ψ is negative and one finds a complexity curve which starts positive at $E = 0$

$$\Sigma(E) \simeq -\psi - \frac{E}{2} \left[\log \left(\frac{E}{2\omega} \right) - 1 \right]. \quad (23)$$

In particular, the number of lowest lying states, which have an energy $E = 0$, scales with the number of N of spins as $\exp(-N\psi)$.

For $\gamma > \gamma_c$, the expression (23) for the complexity still holds, but ψ is positive. The regime of energies close to 0 where $\Sigma(E)$ is negative corresponds to a region where the average number of states is exponentially small in N . Therefore there are no states in this region in the typical sample. States appear above the ground state energy E_0 which is the point where $\Sigma(E)$ vanishes, and corresponds to the maximum of $\Phi(\gamma)$. In Fig. 3 we show the analytic prediction for the ground state energy E_0 (lowest curve) together with numerical results from exact optimization on small systems. Numerical data are compatible with the analytic solution, which has been obtained expanding around the critical point.

4. RIGOROUS DERIVATION OF THRESHOLDS AND CLUSTERING

We now show how the results of the previous Section can be rederived in a rigorous way. We will exploit concepts from graph theory and all the calculations will be simple annealed averages, which are rigorous. All the formulas will be written for generic p , and the particular case $p = 3$ will be considered in order to make connection with calculations in the Section 3.

The physical idea behind the graph theoretical derivation is the following. In a random hypergraph there are many variables with connectivities 0 and 1, whose cavity fields are null. A small fluctuation in the number of these variables, induce very large fluctuations in physical observables, like, e.g., in the entropy. Thus the idea is to remove all these spins and to study the properties of the residual hypergraph, the core. We find that, on the core, sample-to-sample fluctuations are negligible and this allow us to study its properties by means of very simple annealed averages.

The plan of this section is the following: (1) definition of some graph theoretical concepts, like random hypergraph and hyperloop; (2) introduction of the “leaf removal” algorithm and solution to its dynamics (estimation of the γ_d threshold); (3) statistical description of the hypergraph core (the part left by the application of leaf removal algorithm); (4) calculation of γ_c , the SAT/UNSAT threshold; (5) derivation of GS clustering properties.

4.1. Random Hypergraphs and Hyperloops

In the Hamiltonian (2) disorder enters in 2 ways: in the sign of the couplings $J_m = \pm 1$ and in the M random p -uples of indices $\{i_1^m, \dots, i_p^m\}_{m=1, \dots, M}$, which define the interactions topology. This topology has finite connectivity (each variable appears on average in $p\gamma$ interactions) and locally tree-like (a Husimi tree for $p > 2$).

This topology can be represented as a *hypergraph* \mathcal{G} made of a set of N vertices (corresponding to the variables in the problem) and a set of M hyperedges (corresponding to the constraints in the problem), each one connecting p vertices. The disorder ensemble thus corresponds to all the possible ways one can place $M = \gamma N$ hyperedges among N vertices, each hyperedge connecting p vertices and carrying a random sign $J_m = \pm 1$.

Analogously to what happens with loops in usual graphs ($p = 2$), in a disordered model defined on a hypergraph ($p > 2$) frustration is induced by the presence of hyperloops,^(6, 11) which are also called hypercycles in the literature.⁽¹⁵⁾ The definition of a hyperloop can be given both in terms of the hypergraph \mathcal{G} or in terms of the matrix \hat{A} .

A hyperloop is a sub-hypergraph $\mathcal{C} \subset \mathcal{G}$, i.e., a set of hyperedges belonging to \mathcal{G} , such that every vertex has even degree (connectivity) in \mathcal{C} .

In terms of the matrix \hat{A} it corresponds to a set of rows \mathcal{R} such that, for every column, the sum modulo 2 of the elements is zero, i.e., $\sum_{i \in \mathcal{R}} A_{ij} \bmod 2 = 0 \forall j$.

The presence of hyperloops is directly related to the presence of frustration in the system: If the product of the signs of hyperloop interactions is negative, $\prod_{m \in \mathcal{C}} J_m = -1$, then not all such interactions can be satisfied at

the same time. The critical point γ_c , where hyperloops percolate, is a $T = 0$ phase boundary for the p -spin glass models defined by Hamiltonian (2): For $\gamma < \gamma_c$ all the interactions can be satisfied and the GS energy is zero, while for $\gamma > \gamma_c$ the system is in a frustrated spin glass phase and GS of zero energy no longer exist.

The critical point γ_c corresponds to the SAT/UNSAT threshold for the random p -XORSAT problem. In terms of the random linear system $\hat{A}\vec{x} = \vec{y} \bmod 2$, as long as $\gamma < \gamma_c$, solutions to the system will exist with probability 1 in the large N limit for any y .

4.2. “Leaf Removal” Algorithm

Given a hypergraph the leaf removal algorithm proceeds as follows:⁽¹⁶⁾ As long as there is a vertex of degree 1 remove its unique hyperedge. A single step of the algorithm is illustrated in Fig. 4 for a graph ($p = 2$) and for a hypergraph ($p = 3$). Very similar algorithms have been recently studied in refs. 17 and 18.

During the whole process the remaining hypergraph is still a random one, since no correlation can arise among the hyperedges if it was not present at the beginning. When there are no more vertices of degree 1 in the hypergraph the process stops and we call *core* the resulting hypergraph, cleared of all isolated vertices.

The leaf removal algorithm is not able to break up any hyperloop, since each vertex in the hyperloop has at least degree 2. The γ value where the core size becomes different from zero, let us call it γ_d , is certainly smaller than the percolation point of hyperloops γ_c (for $p = 2$ these two values coincide).




Fig. 4. A single step of the “leaf removal” algorithm on a graph (top) and on a hypergraph (bottom).

The evolution of a hypergraph under the application of the leaf removal algorithm can be described in terms of the probability, $f_k(t)$, of finding a vertex of degree k after having removed tN hyperedges where the “time” t ranges from 0 to γ . The initial condition is $f_k(0) = e^{-p\gamma} \frac{(p\gamma)^k}{k!}$ and the evolution equations read (see ref. 18 for a detailed derivation of similar equations)

$$\begin{aligned}\frac{\partial f_0(t)}{\partial t} &= (p-1) \frac{f_1(t)}{m(t)} + 1, \\ \frac{\partial f_1(t)}{\partial t} &= (p-1) \frac{2f_2(t) - f_1(t)}{m(t)} - 1, \\ \frac{\partial f_k(t)}{\partial t} &= (p-1) \frac{(k+1) f_{k+1}(t) - k f_k(t)}{m(t)} \quad \forall k \geq 2,\end{aligned}\tag{24}$$

where $m(t) = \sum_k k f_k(t) = p(\gamma - t)$, since the mean degree linearly decreases with time (we remove one interaction per step) and vanishes at $t = \gamma$.

Thanks to the simplicity of the leaf removal process, the degree distribution always remains Poissonian for degrees larger than 1, with a time dependent average $\lambda(t)$,

$$f_k(t) = e^{-\lambda(t)} \frac{\lambda(t)^k}{k!} \quad \forall k \geq 2.\tag{25}$$

The solution to Eqs. (24) reads

$$\lambda(t) = p[\gamma(\gamma - t)^{p-1}]^{\frac{1}{p}},\tag{26}$$

$$f_1(t) = \lambda(t) \left[e^{-\lambda(t)} - 1 + \left(\frac{\lambda(t)}{p\gamma} \right)^{\frac{1}{p-1}} \right],\tag{27}$$

$$f_0(t) = 1 - \sum_{k=1}^{\infty} f_k(t).\tag{28}$$

The leaf removal algorithm stops when there are no more vertices of degree 1, so one can predict the resulting core by fixing $\lambda(t) = \lambda^*$, where λ^* is the largest zero of the equation $f_1 = 0$ or equivalently

$$e^{-\lambda^*} - 1 + \left(\frac{\lambda^*}{p\gamma} \right)^{\frac{1}{p-1}} = 0.\tag{29}$$

More precisely λ^* is the first zero of Eq. (27) one finds decreasing λ , starting from the initial value of $\lambda(0) = p\gamma$, but this always coincides with the largest zero. Note that once we define $m = [\lambda^*/(p\gamma)]^{1/(p-1)}$, Eq. (29) can be rewritten as

$$1 - m = \exp(-p\gamma m^{p-1}),\tag{30}$$

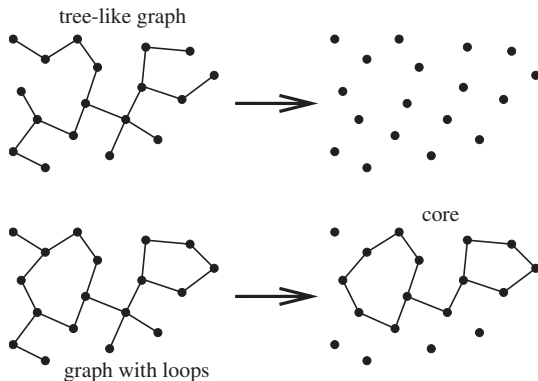


Fig. 5. On graphs ($p = 2$) the leaf removal algorithm is not able to break loops, which thus remain in the residual core.

which is nothing but the equation for the magnetization in the ferromagnetic state,⁽⁶⁾ equivalently the equation for the backbone size in any cluster. Note that Eq. (29) with $p = 3$ is identical to Eq. (14), which indeed determines the mean connectivity of the sub-hypergraph made of hyperedges with non-trivial biases.

In the $p = 2$ case the leaf removal algorithm is able to delete all the edges only for tree-like graphs. As soon as there are loops in the graph, a core containing these loops arises (see Fig. 5). In a random graph the leaf removal transition coincides with the percolation one at $\gamma_p = 1/2$. The shape of the function $f_1(\lambda)$ is shown in Fig. 6. For $\gamma \leq \gamma_p$, there is only one zero in $\lambda^* = 0$; While, for $\gamma > \gamma_p$, $\lambda^* > 0$ and a core arises, whose size grows as $(\gamma - \gamma_p)^2$ near the critical point.

For $p > 2$ the percolation transition, taking place at $\gamma_p = \frac{1}{p(p-1)}$, does not affect at all the leaf removal algorithm which is able to delete all the




Fig. 6. The function $f_1(\lambda)/\lambda$ for $p = 2$.




Fig. 7. For $p > 2$ the leaf removal algorithm is able to break loops (but not hyperloops!).

hyperedges, even those forming loops (but not those forming hyperloops!), far beyond γ_p (see Fig. 7).

The leaf removal transition takes place at γ_d , which is defined as the first γ value where a second solution to Eq. (30) appears. For $p = 3$ we have $\gamma_d = 0.818469$. The transition is first order and, at the critical point, the core already occupies a finite fraction of the system. In Fig. 8 we show the function $f_1(\lambda)$ for $p = 3$. It is clear (see inset of Fig. 8) that when $\lambda^*(\gamma)$ becomes different from zero it directly jumps to a finite value: $\lambda^*(\gamma_d) = 1.25643$ for $p = 3$.

4.3. Statistical Description of the Core

Once the leaf removal process has come to an end the distribution of connectivities on the core (normalized with respect to the original number of vertices) is a truncated Poissonian

$$P_c(k) = \begin{cases} 0 & \text{for } k = 0, 1 \\ e^{-\lambda^*(\gamma)} \frac{\lambda^*(\gamma)^k}{k!} & \text{for } k \geq 2 \end{cases} \quad (31)$$




Fig. 8. The function $f_1(\lambda)/\lambda$ for $p = 3$. Inset: function $\lambda^*(\gamma)$ for $p = 3$.

The number of vertices N_c and the number of hyperedges M_c in the core can be expressed in terms of p , γ and $\lambda^*(\gamma)$ as

$$N_c(\gamma) = N \sum_{k=2}^{\infty} f_k(\lambda^*) = N[1 - (1 + \lambda^*) e^{-\lambda^*}], \quad (32)$$

$$M_c(\gamma) = M - Nt^* = N \left[\frac{1}{\gamma} \left(\frac{\lambda^*}{p} \right)^p \right]^{\frac{1}{p-1}} = \gamma N (1 - e^{-\lambda^*})^p = N \frac{\lambda^*}{p} (1 - e^{-\lambda^*}). \quad (33)$$

The first of these equations has a simple interpretation: The number of vertices in the core is nothing but the number of vertices with a degree larger than 1, after the application of the leaf removal algorithm. The second equation states that the number of hyperedges left is the initial one minus the number of step the leaf removal algorithm has been run (during each step only one hyperedge is deleted). The running time t^* is the solution to Eq. (26) with λ^* on the left hand side. The last two, and more compact, expressions for M_c have been obtained with the use of Eq. (29). The lower curves in Fig. 9 show the normalized number of vertices N_c/N and number of interactions M_c/N in the core as a function of γ , for $p = 3$.

It is natural now to study the residual problem on the core, $\hat{A}_c \vec{x}_c = \vec{y}_c \bmod 2$, where \hat{A}_c is the $M_c \times N_c$ sparse random matrix obtained from \hat{A} deleting all the rows corresponding to removed interactions and all empty columns. In the next subsection we will derive a general result that, when applied to the problem on the core, gives a necessary and sufficient condition for the existence of solutions to $\hat{A}_c \vec{x}_c = \vec{y}_c \bmod 2$. Then we will show that, from a solution in the core, a solution for the original system can always be constructed.




Fig. 9. From bottom to top (on the left): For $p = 3$, normalized number of hyperedges and vertices in the core, and fraction of frozen sites, i.e., magnetization (or backbone) in a state.

4.4. Calculation of the γ_c Threshold

Let us call $\mathcal{N}_{\mathcal{J}, N, M}$ the number of GS for a given disorder realization \mathcal{J} (i.e., a given hypergraph and coupling signs) with N variables and M interactions. We will show that, in the large N limit, if the hypergraph does not contain any vertex of degree less than 2, $\mathcal{N}_{\mathcal{J}, N, M}$ is a self averaging quantity, that is it does not fluctuate changing \mathcal{J} .

In order to show self-averageness we will prove that, on hypergraphs ($p > 2$) with minimum degree at least 2, the following equalities hold

$$\overline{\mathcal{N}_{\mathcal{J}, N, M}} = 2^{N-M}, \quad \lim_{N \rightarrow \infty} \frac{\overline{\mathcal{N}_{\mathcal{J}, N, M}^2} - (\overline{\mathcal{N}_{\mathcal{J}, N, M}})^2}{(\overline{\mathcal{N}_{\mathcal{J}, N, M}})^2} = 0, \quad (34)$$

where the overline stands for the average over the disorder ensemble, that is over the ways of choosing M hyperedges among $\binom{N}{p}$ and the ways of giving them a sign $J_m = \pm 1$. The above equalities state that the probability distribution of $\mathcal{N}_{\mathcal{J}, N, M}$ over the disorder ensemble is a delta function, and thus the quenched average equals the annealed one

$$\overline{\log \mathcal{N}_{\mathcal{J}, N, M}} = \log \overline{\mathcal{N}_{\mathcal{J}, N, M}} = \log(2) (N - M). \quad (35)$$

Given the definition

$$\mathcal{N}_{\mathcal{J}, N, M} = \sum_{\vec{\sigma}} \prod_{m=1}^M \delta(\sigma_{i_1^m} \cdots \sigma_{i_p^m} = J_m), \quad (36)$$

the first moment is trivially given by

$$\overline{\mathcal{N}_{\mathcal{J}, N, M}} = \sum_{\vec{\sigma}} \prod_{m=1}^M \delta(J_m = \sigma_{i_1^m} \cdots \sigma_{i_p^m}) = 2^{N-M}, \quad (37)$$

since, for every given spin configuration and topology, the probability that coupling signs satisfy all the M interactions is exactly 2^{-M} .

The second moment is given by

$$\begin{aligned} \overline{\mathcal{N}_{\mathcal{J}, N, M}^2} &= \overline{\sum_{\vec{\sigma}} \prod_{m=1}^M \delta(\sigma_{i_1^m} \cdots \sigma_{i_p^m} = J_m) \delta(\sigma'_{i_1^m} \cdots \sigma'_{i_p^m} = J_m)} \\ &= \overline{\sum_{\vec{\sigma}} \prod_{m=1}^M \delta(J_m = \sigma_{i_1^m} \cdots \sigma_{i_p^m}) \sum_{\vec{\sigma}'} \prod_{m=1}^M \delta(\sigma'_{i_1^m} \cdots \sigma'_{i_p^m} = \sigma_{i_1^m} \cdots \sigma_{i_p^m})} \\ &= 2^{N-M} \overline{\sum_{\vec{\tau}} \prod_{m=1}^M \delta(\tau_{i_1^m} \cdots \tau_{i_p^m} = 1)}, \end{aligned} \quad (38)$$

where $\tau_i = \sigma_i \sigma'_i$ and the last expression is nothing but the annealed average of the partition function at $T = 0$ for a system where all the coupling signs have been set to 1, i.e., a ferromagnetic model. Such an average can be computed by standard saddle point integration and the final result is

$$F_{N,M} = \lim_{N \rightarrow \infty} \frac{1}{N} \log \sum_{\vec{\tau}} \prod_{m=1}^M \delta(\tau_{i_1}^m \cdots \tau_{i_p}^m = 1) = \sum_k P(k) \log(x_+^k + x_-^k), \quad (39)$$

where $P(k)$ is the distribution of connectivities in the hypergraph and x_+, x_- solve the following equations

$$x_+ + x_- = \left[\sum_k \frac{kP(k)}{\langle k \rangle} \frac{x_+^{k-1} + x_-^{k-1}}{x_+^k + x_-^k} \right]^{p-1}, \quad (40)$$

$$x_+ - x_- = \left[\sum_k \frac{kP(k)}{\langle k \rangle} \frac{x_+^{k-1} - x_-^{k-1}}{x_+^k + x_-^k} \right]^{p-1}. \quad (41)$$

Here $\langle k \rangle = \sum_k kP(k) = p \frac{M}{N}$ is the mean connectivity. When more than one solution to Eqs. (40) and (41) exist, the one maximizing Eq. (39) must be chosen. The value of x_+ (resp. x_-) is proportional to the fraction of variables taking values 1 (resp. -1) in the set of configurations which maximize the sum in Eq. (39). Then the typical magnetization of this model is given by $m = \frac{x_+ - x_-}{x_+ + x_-}$.

Solutions to Eqs. (40) and (41) can be classified depending on the value of magnetization m . In full generality there are 3 solutions: a first symmetric one ($x_+ = x_-$) with $m = 0$, a second one with large magnetization and a third one with an intermediate value of m . For some choices of $P(k)$ (e.g., a Poissonian) solutions with $m > 0$ may exist only for $\frac{M}{N}$ large enough. The solution with intermediate magnetization always corresponds to a minimum of $F_{N,M}$ and can be in general neglected.

The symmetric solution $x_+ = x_- = 2^{-1/p}$ always exists and gives $F_{N,M} = \log(2) (1 - \frac{M}{N})$. For $p > 2$ and $P(0) = P(1) = 0$, i.e., for hypergraphs with minimum degree 2, the solution with large magnetization also exist for any γ value and has $x_+ = 1, x_- = 0$ and $F_{N,M} = 0$. As expected, the intermediate solution, when it exists, has $F_{N,M} < 0$.

Then, for $p > 2$ and $P(0) = P(1) = 0$, we can conclude that the average in the last term of Eq. (38) equals $e^{NF_{N,M}} = 2^{N-M}$ (the coefficient can be easily calculated and is exactly 1). Thus, equalities in Eq. (34) hold and the number of GS is a self-averaging quantity.

Since the core generated by the leaf removal algorithm has minimum degree 2, we may apply the above result, and find that the SAT/UNSAT threshold is given by the condition

$$N_c(\gamma_c) = M_c(\gamma_c) \Rightarrow 1 - (1 + \lambda_c) e^{-\lambda_c} = \frac{\lambda_c}{p} (1 - e^{-\lambda_c}), \quad (42)$$

where $\lambda_c = \lambda^*(\gamma_c)$. For $\gamma \leq \gamma_c$ there are $2^{N_c - M_c}$ solutions (i.e., unfrustrated GS) in the core, while for $\gamma > \gamma_c$ there is none. For $p = 3$, solution to Eq. (42) gives $\lambda_c = 2.14913$ and $\gamma_c = 0.917935$.

For any given solution in the core, a solution for the whole original system can be easily reconstructed. Indeed, we reintroduce in the system the interactions removed during the leaf removal process, but in a reversed order (i.e., the last removed is the first to be reintroduced). At each step, together with one interaction, at least one variable is reintroduced in the system (the variable having degree 1 when that interaction was removed) and this variable must be set such as to satisfy the interaction. Very often more than one variable per step is reintroduced, allowing for multiple and equivalent choices. This redundancy is what makes the total number of solutions larger than the number of solutions in the core (see below).

In the table below we report the thresholds γ_d and γ_c for some p values.

p	γ_d	γ_c
2	1/2	1/2
3	0.818469	0.917935
4	0.772278	0.976770
5	0.701780	0.992438
6	0.637080	0.997380

4.5. Clustering of Ground States

Let us come back to the problem of clustering solutions below the SAT/UNSAT threshold ($\gamma \leq \gamma_c$). In this region the system is not frustrated and then a gauge transformation setting all coupling signs to 1 can always be found: Given an unfrustrated GS $\vec{\sigma}^0$ a possible gauge transformation is $\sigma'_i = \sigma_i \sigma_i^0$ and $J'_m = J_m \sigma_{i_1}^0 \cdots \sigma_{i_p}^0 = 1$. Thanks to this, in the rest of the paper we will consider only a ferromagnetic system ($J_m = 1 \forall m$), which corresponds to the linear system $\hat{A} \vec{x} = \vec{0} \pmod{2}$.

The solutions to the linear system $\hat{A}\vec{x} = \vec{0} \pmod 2$ form a group: The sum of 2 solutions is still a solution and the null element is the solution $\vec{x} = \vec{0}$. The symmetry group is telling us that if one looks at the configurational space sitting on a reference GS, the set of GS will look the same, whatever the reference GS is. An immediate consequence of this symmetry is that, if GS form clusters, these clusters must be all of the same size.

For $\gamma \leq \gamma_c$, hyperloops are absent and the total number of GS (or solutions) is always given by 2^{N-M} , i.e., their entropy is $S(\gamma) = \log(2)(1 - \gamma)$. Let us divide the N variables in 2 sets: \vec{x}_c represents the N_c variables in the core, and \vec{x}_{nc} the $N - N_c$ variables in the non-core part of the hypergraph, that is variables corresponding to vertices remained isolated at the end of the leaf removal process. Thus also the entropy can be divided in 2 parts. One part is given by the solutions in the core, that is by the possible assignments of \vec{x}_c ,

$$S_c(\gamma) = \log(2) \frac{N_c(\gamma) - M_c(\gamma)}{N}, \tag{43}$$

which is non-negative for $\gamma_d \leq \gamma \leq \gamma_c$. The other part is given by the possible multiple assignments of \vec{x}_{nc} during the reconstruction process

$$S_{nc}(\gamma) = S(\gamma) - S_c(\gamma). \tag{44}$$

This separation of the entropy in 2 parts is physically relevant, and we will show here that it corresponds to the proper clustering of the solutions.

The physical picture we have in mind is sketched in Fig. 10. For $\gamma_d \leq \gamma \leq \gamma_c$, the solutions of $\hat{A}\vec{x} = \vec{y} \pmod 2$, or equivalently the ground states

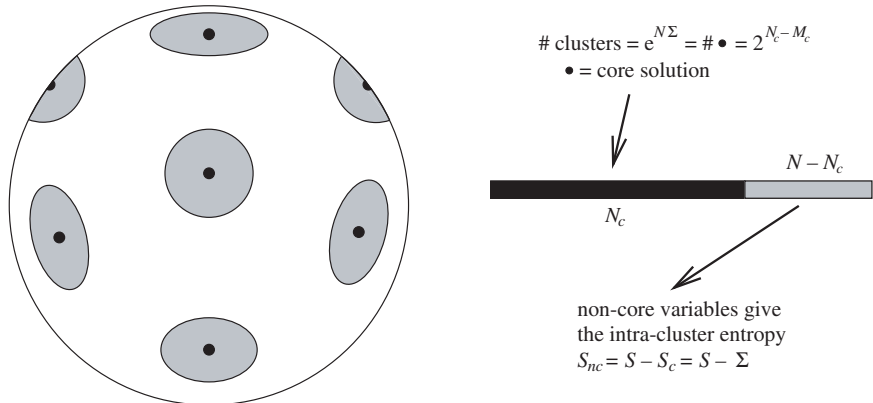


Fig. 10. Schematic picture of the clustering of solutions for $\gamma_d < \gamma < \gamma_c$.




Fig. 11. Total entropy $S(\gamma)$ and configurational entropy $\Sigma(\gamma)$ for $p = 3$.

of (2), spontaneously form clusters. By definition, two solutions having a finite Hamming distance d , i.e., $d/N \rightarrow 0$ for $N \rightarrow \infty$, are in the same cluster, while two solutions in different clusters must have an extensive distance, that is $d/N \sim \mathcal{O}(1)$ for large N .

In virtue of the property stated at the beginning of this subsection, all the clusters have the same size. Their number is $e^{N\Sigma(\gamma)}$, where $\Sigma(\gamma)$ is called complexity or configurational entropy. We will show that the number of clusters equals the number of solution in the core, that is

$$\Sigma(\gamma) = S_c(\gamma). \quad (45)$$

The intra-cluster entropy, i.e., the normalized logarithm of the cluster size, is then given by the non-core entropy $S_{nc}(\gamma) = S(\gamma) - S_c(\gamma) = S(\gamma) - \Sigma(\gamma)$. For $p = 3$ these entropies are shown in Fig. 11.

The proof of Eq. (45) is given in 2 steps. First we show that all the solution assignments of the core variables \vec{x}_c are “well separated”, that is the distance among any pair of them is extensive. This is what gives rise to the clustering, with a number of clusters which is at least as large as the number of core solutions ($\Sigma \geq S_c$). Then we show that, for any fixed \vec{x}_c , all possible assignments of non-core variables \vec{x}_{nc} belong to the same cluster, and so $\Sigma = S_c$.

The first step is accomplished by calculating the probability distribution of the distance among any two solutions in the core. Thanks to the group property, we can restrict the calculation fixing one solution to the null vector $\vec{0}$. For simplicity we have performed an annealed average

$$S(d, \gamma) = \lim_{N_c \rightarrow \infty} \frac{1}{N_c} \log \sum_{\vec{\sigma}} \delta \left(\sum_i \sigma_i = N_c - 2d \right) \prod_{m=1}^{M_c} \delta(\sigma_{i_1}^m \cdots \sigma_{i_p}^m = 1), \quad (46)$$




Fig. 12. Entropy of distances among solutions in the core for $p = 3$ (in the annealed approximation).

which gives an upper bound to the exact result. The expression for this entropy is given by Eq. (39), where x_+ comes from the solution of Eq. (40), keeping the ratio $\frac{d}{N_c} = \frac{x_-}{x_+ + x_-}$ fixed.

In Fig. 12 we plot the resulting entropy as a function of the distance d , for $p = 3$ and some values of γ . For $\gamma_d < \gamma < \gamma_c$ the entropy is negative for $0 < d < d_{\min}(\gamma)$, and so $d_{\min}(\gamma)$ is a lower bound on the minimum distance among any two solutions in the core. This minimum distance is shown for $p = 3$ in Fig. 13.

Then all the e^{NS_c} core solutions are well separated, and can be represented as the centers of the clusters (see Fig. 10). It remains to be proven that, for any fixed \vec{x}_c , the solution assignments of \vec{x}_{nc} form a single cluster. Thus no further clustering is present and the picture of Fig. 10 is correct.

This last proof is given in the Appendix, and it is based on an algorithm which allows one to change the value to any variable in \vec{x}_{nc} by simply




Fig. 13. Lower bound for the minimum distance among any 2 solutions in the core for $p = 3$.

adjusting other $\mathcal{O}(1)$ variables in \vec{x}_{nc} . This shows that all the solutions in one cluster are connected in the following sense. One solution can be reached from any other one by a sequence of moves, where each move involves flipping only a finite number of spins.

5. CONCLUSION AND DISCUSSION

In this work we have solved, with two alternative methods, the p -XORSAT model, which corresponds to the zero-temperature limit of the diluted p -spin model.

Increasing the γ parameter (number of interactions per variable) the model undergoes two phase transitions. At γ_d , solutions to the p -XORSAT problem (i.e., ground states for the p -spin model) spontaneously form an exponentially large number of clusters, thus giving a finite configurational entropy. At γ_c , frustration percolates throughout the system, and consequently the number of clusters (and solutions) goes to zero, and the ground state energy becomes positive. γ_c corresponds to the SAT/UNSAT threshold. These exact results perfectly agree with previous replica calculations^(6, 11, 12) and may suggest new approaches for finding mathematical bases to Parisi's theory of spin glasses.⁽¹³⁾

The use of the cavity method combined with a rigorous derivation based on the topological properties of the interaction hypergraph, allow us to establish some interesting links among distributions of cavity fields on a given variable and the position of the corresponding vertex in the hypergraph. In particular all the variables with a non-trivial distribution of cavity fields belong to the 'frozen' part of the hypergraph (see Appendix), that is to the core and to the part that can be uniquely fixed, once an assignment to core variables has been chosen. The "frozen" part is exactly the backbone of a cluster (variables which take the same value for all the solutions in the cluster) and its size is given by the largest solution to Eq. (30). The rest of the hypergraph, the "floppy" part, only contains paramagnetic variables, that is variables always having a null cavity field.

APPENDIX A

In this appendix we show that assignments of non-core variables \vec{x}_{nc} are not clustered. To this end, we define an algorithm which allows one to flip any non-core variable, by simply adjusting other $\mathcal{O}(1)$ non-core variables. With this algorithm one can move through all the \vec{x}_{nc} assignments by doing *finite* steps, thus proving that non-core solutions form a single cluster.

Let us fix the core variables \bar{x}_c to any solution, and call them “frozen.” All the variables, belonging to at least one equation where the other $p-1$ variables are already frozen, must be frozen too (see e.g., the dashed triangle in Fig. 14, where the dashed blobs represent the frozen core). In this way one is able to freeze a number of variables $m(\gamma) N$, where $m(\gamma)$ turns out to coincide with the largest solution of Eq. (30), that is with the magnetization in the ferromagnetic state or the backbone in a generic cluster. For $p=3$ the function $m(\gamma)$ is shown in Fig. 9 (upper curve).

After having fixed all the variables one could, one is left with the “floppy” part of the hypergraph. The typical situation is sketched in Fig. 14, where the dashed part is frozen (hereafter we refer only to the $p=3$ case for the sake of clarity). All the interactions involving both frozen and floppy variables (those which form the boundary between the frozen and the floppy part of the hypergraph) must contain 2 floppy and 1 frozen variables, otherwise (2 frozen and 1 floppy) that interaction would become frozen as well and would not longer be on the boundary.

The numbers in Fig. 14 have been assigned during a slightly different leaf removal process with the following rule. Starting with the original hypergraph, the number “1” is given to all the vertices of degree less than 2 (isolated vertices and leafs) and their hyperedges are deleted. Then, in the new hypergraph, the number “2” is given to all vertices of degree less than 2 and their hyperedges deleted. And so on. We call these numbers the *depth* of a vertex. Vertices of depth 1 represent the ‘external boundary’ or the “surface” of the hypergraph.

The evolution of this “collective” leaf removal process can be described in terms of the same function $f_1(t)$ used previously. At each time step a depth is assigned to a fraction $f_1(t)$ of vertices and then the time is




Fig. 14. The bold tree-like structure is a possible seaweed (see text) in order to flip the variable on the circled vertex and still keep all the interactions satisfied. Note that the seaweed passes through at most 2 vertices on the same interaction.

increased by $\Delta t = f_1(t)$, in order to take into account the deletion of hyperedges leaving from the just numbered vertices. For very large times and depths, $f_1(t)$ is very small and can be approximated by $f_1(t) \simeq (t - t^*) \partial_t f_1(t^*)$, where t^* is such that $f_1(t^*) = 0$. In this regime we have that

$$\Delta t = f_1(t), \quad (\text{A.1})$$

$$\Delta f_1(t) = \left. \frac{\partial f_1(t)}{\partial t} \right|_{t^*} \Delta t = \left. \frac{\partial f_1(t)}{\partial t} \right|_{t^*} f_1(t), \quad (\text{A.2})$$

and so $f_1(t + \Delta t) = f_1(t)[1 + \partial_t f_1(t^*)]$. Then the probability of having a (large) depth h satisfies the equation $\mathcal{P}(h+1) \simeq \mathcal{P}(h)\mu$, where

$$\mu(\gamma) = 1 + \left. \frac{\partial f_1(t)}{\partial t} \right|_{t^*} = 1 + \left. \frac{\partial f_1(\lambda)}{\partial \lambda} \frac{\partial \lambda(t)}{\partial t} \right|_{\lambda^*(\gamma)}. \quad (\text{A.3})$$

Since the probability of having depth h drops exponentially for large h as $\mathcal{P}(h) \propto \mu^h$, the largest depth assigned with this process is $\mathcal{O}(\log N)$. For any $\gamma \neq \gamma_d$ we have that $\mu(\gamma) < 1$, since $\lambda(t)$ is a decreasing function of t and $\partial_\lambda f_1(\lambda)$ is positive in the largest root λ^* , unless $\gamma = \gamma_d$.

Once depths have been assigned, there is an algorithm (described below) which allows one to change the value to any floppy variable, by adjusting, at the same time, only $\mathcal{O}(1)$ other floppy variables. Such a new configuration will be a finite distance far apart, and, by definition, will belong to the same cluster. In this way one can change the configuration of the floppy (and non-core) variables to any admissible one, and these configurations will form a unique cluster.

The physical idea behind the algorithm for flipping any floppy variable, keeping all the interactions satisfied, and adjusting only a finite number of other variables, is the following. Suppose, as in Fig. 14, that we flip the variable of depth 5. Then the interactions it participates to will be unsatisfied, and we have to move this “excess energy” by flipping other variables, along the shortest way, towards the boundaries of the hypergraph, that is the vertices of depth 1, where it can be freely released. The only delicate point is the definition of the “path to the boundary,” which has to contain a finite number of vertices. In Fig. 14 we show a possible way to release the excess energy generated by flipping variable of depth 5. Flipping all the variables belonging to the tree-like bold structure

will keep all the interactions satisfied, since every interaction contains an even number of vertices belonging to the bold structure.

We will call this tree-like structure a *seaweed*, since it has a root, corresponding to the vertex of maximum depth, and the number of its branches grows approaching the surface. Now we give the rules for constructing a seaweed, such that its size is finite.

Let us start with some nomenclature: We say that a hyperedge e is “below” a vertex v , and analogously v is “above” e , if the depth of v is the smallest among the depths of all the vertices in e .

Thanks to the way depths have been assigned, each vertex may have at most 1 hyperedge below. This property can be easily proved, remembering that to any given vertex v the depth is assigned only when its connectivity is 0 or 1. At this time, all the other hyperedges of v have been removed, since we have assigned smaller depths to its neighbours. The only hyperedge which can be below v is the last one. Moreover, if the depth is assigned to v when its connectivity is 0 (isolated vertex), the vertex v will have no hyperedges below, and we will call it a *root*.

The construction of the seaweed starts from the vertex corresponding to the variable that we want to flip (let us call it *seed*). In this way we are sure that such a vertex will be in the structure, and the corresponding variable flipped. The seaweed is built up recursively, that is we give the rules for growing a single branch, both upwards (i.e., towards the surface) and downwards (i.e., towards a root), and then these rules must be applied to any branch of the seaweed, until it reaches the surface of the hypergraph or a root vertex. The branches are such that along an upwards (downwards) direction the depth strictly decreases (increases). Rare exceptions to this property will be illustrated below.

When a branch passes through a hyperedge it will visit only 2 vertices in this hyperedge, such that, when all the variables belonging to the seaweed will be flipped, the interaction will remain satisfied.

Suppose the seed vertex has connectivity k . Then we start k different branches, 1 downwards entering the only hyperedge below the seed vertex and $k-1$ upwards entering the other hyperedges.

Any upwards branch entering a hyperedge e through vertex v has to be continued with the vertex above e . If there are many vertices of the same minimum depth in e , any of them can be chosen equivalently. With this rule we are ensuring that the new vertex added to the upwards branch is of smaller depth than v .

Any downwards branch entering hyperedge e through vertex v has to be continued with the vertex of maximum depth in e . If there are many vertices of the same maximum depth in e , any of them can be chosen equivalently. With this rule we can ensure that the new vertex added

to the downwards branch will be deeper than v , since v is of minimum depth in e .

Any growing branch reaching a vertex v of connectivity k has to be continued with $k-1$ branches, in order to satisfy the rule that all the hyperedges of v must be visited by a branch. If the just reached vertex is on the surface (i.e., it has depth 1 and connectivity 1) the branch ends there. On the contrary, reaching a vertex of connectivity larger than 2, the growing branch generates new branches. More in particular, if the branch is an upwards one it will generate only upwards branches (since it is coming from the only hyperedge below v). While, if it is the downwards one, it may generate at most one downwards branch (all the rest being upwards ones). This is a consequence of the property that every vertex may have at most 1 hyperedge below it.

In two cases the unique downwards branch ends in a vertex v , which is thus the root of the seaweed: (1) v is a root vertex, that is it has no hyperedges below it (2) vertex v is above hyperedge e , but v is not the only vertex of minimum depth in e . In this case the branch entering e through v becomes an upwards one, and makes a single step without decreasing the depth (this is the only exception to the rule on the monotonicity of the depth along a branch stated above).

Since each branch of the seaweed is grown independently, it may be that at the end of the process some vertices result in more than one branch. This is not a problem: The rule says that every vertex which has been included an odd number of times in the seaweed must be in it; While those entering an even number of times must be left out. The net result is a decrease in the total number of vertices in the structure. The seaweed can eventually break up in more than a single connected component. All the components, but that containing the seed, can be removed from the seaweed.

The choice of growing the branches always along vertices of maximum and minimum depths is dictated by the need of reaching a root vertex and the surface of the hypergraph as soon as possible, thus making the seaweed as small as possible. It is worth noticing that the probability that a vertex is a root increases for larger depths.

The last point to be proven is that the typical distance, ℓ , measured along any branch, among the root of the seaweed and the surface, is finite (and not order $\log N$). This property together with the fact that the branching ratio is proportional to the connectivity, which is finite too, implies that the number of vertices in the seaweed, which is roughly proportional to $(3\gamma-1)^\ell$, is finite. On the contrary if ℓ would be of order $\log N$, the volume of the seaweed would diverge for large N .

In order to show that ℓ is finite, even when the root depth is as large as possible (i.e., order $\log N$), we need to know the probability that a vertex

has depth h . This probability distribution function, $\mathcal{P}(h)$, can be calculated exactly, but its expression is too involved to be presented here. We only report some features relevant for our purposes. It depends on the connectivity of the vertex, $\mathcal{P}_k(h)$, and for $k=0$ or $k=1$ it is trivially given by $\mathcal{P}_{0,1}(h) = \delta(h-1)$. For any $k \geq 2$, it decreases exponentially fast for large h , and the probability of reaching a vertex (not on the surface) of depth h is $Q(h) = \sum_{k \geq 2} k f_k(0) \mathcal{P}_k(h) \propto \mu(\gamma)^h$ for large h . For the present calculation the exact shape of $Q(h)$ at small depths is irrelevant, and we only care about its tail, so we can hereafter use $Q(h) = \mu^h$ for all h .

We show now that, with such a distribution of depths, even starting from a root of depth $\mathcal{O}(\log N)$, an upwards branch needs only a finite number of steps to reach the surface (for simplicity we fix to 0, instead of 1, the surface depth). The probability of going in a single step from depth h_1 to depth h_2 is

$$w(h_1 \rightarrow h_2) = \frac{1-\mu}{1-\mu^{h_1}} \mu^{h_2}, \quad (\text{A.4})$$

which has the correct normalization $\sum_{h_2=0}^{h_1-1} w(h_1 \rightarrow h_2) = 1$. The probability of going from depth h to depth 0 in m steps is then

$$\begin{aligned} W_h(m) &= \sum_{h_1=h_2+1}^{h-1} \sum_{h_2=h_3+1}^{h_1-1} \cdots \sum_{h_{m-1}=1}^{h_{m-2}-1} w(h \rightarrow h_1) w(h_1 \rightarrow h_2) \cdots \\ &\quad w(h_{m-2} \rightarrow h_{m-1}) w(h_{m-1} \rightarrow 0) \\ &= \frac{1-\mu}{1-\mu^h} \frac{(1-\mu)^{m-1}}{(m-1)!} \sum' \prod_{j=1}^{m-1} \frac{\mu^{i_j}}{1-\mu^{i_j}} \\ &< \frac{1-\mu}{1-\mu^h} \frac{(1-\mu)^{m-1}}{(m-1)!} \left[\sum_{i=1}^{h-1} \frac{\mu^i}{1-\mu^i} \right]^{m-1}, \end{aligned} \quad (\text{A.5})$$

where the primed sum is over the $m-1$ intermediate depths, taking different values between 1 and $h-1$, and the inequality follows since in the last term we have included also configurations with indices taking equal values. So $W_h(m)$ is upper bounded by a Poissonian distribution with a mean number of steps

$$\ell(h) = (1-\mu) \sum_{i=1}^{h-1} \frac{\mu^i}{1-\mu^i}. \quad (\text{A.6})$$

As expected, ℓ is an increasing function of h . In the limit of a very deep root, $h \rightarrow \infty$, the series converges for any $\mu < 1$ (i.e., $\gamma > \gamma_d$), and thus $\ell(\infty)$ is still finite.

ACKNOWLEDGMENTS

We thank Silvio Franz, Michele Leone and Martin Weigt for discussions. F.R.T. thanks ICTP for kind hospitality during the completion of this work.

While writing this work, we became aware of an independent work by S. Cocco, O. Dubois, J. Mandler, and R. Monasson on similar issues.⁽¹⁹⁾

REFERENCES

1. M. Mézard and G. Parisi, *Eur. Phys. J. B* **20**:217 (2001).
2. O. C. Martin, R. Monasson, and R. Zecchina, *Theoret. Comput. Sci.* **265**:3 (2001).
3. S. Franz, M. Mézard, F. Ricci-Tersenghi, M. Weigt, and R. Zecchina, *Europhys. Lett.* **55**:465 (2001).
4. T. J. Schaefer, in *Proc. 10th STOC* (San Diego, CA, USA, ACM, 1978), p. 216. N. Creignou, H. Daudé, and O. Dubois, preprint arXiv:cs.DM/0106001.
5. J. Hastad, *J. ACM* **48**, 798 (2001).
6. F. Ricci-Tersenghi, M. Weigt, and R. Zecchina, *Phys. Rev. E* **63**:026702 (2001).
7. A. Braunstein, M. Leone, F. Ricci-Tersenghi, and R. Zecchina, *J. Phys. A* **35**:7559 (2002), preprint arXiv: cond-mat/0203613, for the role of the dynamical transition on global algorithms.
8. M. Mézard, G. Parisi, and R. Zecchina, *Science* **297**:812 (2002).
9. M. Mézard and R. Zecchina, *Phys. Rev. E* **66**:056126 (2002), arXiv:cond-mat/0207194.
10. For a very recent review on mathematical results on spin glass models see A. Bovier and I. Kurkova, preprint arXiv:cond-mat/0206562, and references therein.
11. M. Leone, F. Ricci-Tersenghi, and R. Zecchina, *J. Phys. A* **34**:4615 (2001).
12. S. Franz, M. Leone, F. Ricci-Tersenghi, and R. Zecchina, *Phys. Rev. Lett.* **87**:127209 (2001).
13. M. Mézard, G. Parisi, and M. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1987).
14. M. Mézard and G. Parisi, *J. Stat. Phys.* **111**:1 (2003), arXiv:cond-mat/0207121.
15. V. F. Kolchin, *Random Graphs* (Cambridge University Press, 1999).
16. B. Pittel, J. Spencer, and N. Wormald, *J. Comb. Theory B* **67**:111 (1996).
17. M. Bauer and O. Golinelli, *Eur. Phys. J. B* **24**:339 (2001).
18. M. Weigt, *Eur. Phys. J. B* **28**:369 (2002), preprint arXiv:cond-mat/0203281.
19. S. Cocco, O. Dubois, J. Mandler, and R. Monasson, *Phys. Rev. Lett.* **90**:047205 (2003), preprint arXiv:cond-mat/0206239.