

Polynomial iterative algorithms for coloring and analyzing random graphs

A. Braunstein

*International Center for Theoretical Physics, Strada Costiera 11, P.O. Box 586, I-34100 Trieste, Italy
and SISSA, via Beirut 9, I-34100 Trieste, Italy*

R. Mulet

*International Centre for Theoretical Physics, Strada Costiera 11, P.O. Box 586, 34100 Trieste, Italy
and Henri-Poincaré-Chair of Complex Systems and Superconductivity Laboratory, Physics Faculty-IMRE, University of Havana,
La Habana CP 10400, Cuba*

A. Pagnani

*International Centre for Theoretical Physics, Strada Costiera 11, P.O. Box 586, 34100 Trieste, Italy
and Laboratoire de Physique Théorique et Modèles Statistiques, Bâtiment 100, Université Paris-Sud, F-91405 Orsay, France*

M. Weigt

*International Centre for Theoretical Physics, Strada Costiera 11, P.O. Box 586, 34100 Trieste, Italy
and Institute for Theoretical Physics, University of Göttingen, Bunsenstrasse 9, 37073 Göttingen, Germany*

R. Zecchina

International Centre for Theoretical Physics, Strada Costiera 11, P.O. Box 586, 34100 Trieste, Italy

(Received 24 April 2003; published 3 September 2003)

We study the graph coloring problem over random graphs of finite average connectivity c . Given a number q of available colors, we find that graphs with low connectivity admit almost always a proper coloring whereas graphs with high connectivity are uncolorable. Depending on q , we find with a one-step replica-symmetry breaking approximation the precise value of the critical average connectivity c_q . Moreover, we show that below c_q there exists a clustering phase $c \in [c_d, c_q]$ in which ground states spontaneously divide into an exponential number of clusters. Furthermore, we extended our considerations to the case of single instances showing consistent results. This leads us to propose a different algorithm that is able to color in polynomial time random graphs in the hard but colorable region, i.e., when $c \in [c_d, c_q]$.

DOI: 10.1103/PhysRevE.68.036702

PACS number(s): 02.70.-c, 89.20.Ff, 75.10.Nr, 05.70.Fh

I. INTRODUCTION

The graph coloring problem is a very basic problem in combinatorics [1] and in statistical physics [2]. Given a graph, or a lattice, and given a number q of available colors, the problem consists in finding a coloring of all vertices such that no edge has two ending vertices of the same color. The minimally needed number of colors is the *chromatic number* of the graph.

For planar graphs there exists a well-known theorem [3] showing that four colors are sufficient, and that a coloring can be found by an efficient algorithm, while for general graphs the problem is computationally hard to solve. In 1972 it was shown that graph coloring is NP complete [4] which means, roughly speaking, that the time required for determining the existence of a proper coloring grows exponentially with the graph size. On the other hand, if an efficient algorithm for solving coloring in its worst-case instances exists, the same algorithm up *polynomial reduction* can be applied to efficiently solve all other problems in the class NP (for a physicist's approach to complexity theory see Ref. [5]).

In modern computer science, graph coloring is taken as one of the most widely used benchmarks for the evaluation of algorithm performance [6]. The interest in coloring stems

from the fact that many real-world combinatorial optimization problems have component subproblems which can be easily represented as coloring problems. For instance, a classical application is the scheduling of registers in the central processing unit of computers [7]. All variables manipulated by the program are characterized by ranges of times during which their values are left unchanged. Any two variables that change during the same time interval cannot be stored in the same register. One may represent the overall computation by constructing a graph where each variable is associated with a vertex and edges are placed between any two vertices whose corresponding variables change during the same time interval. A proper coloring with a minimal number of colors of this graph provides an optimal scheduling for registers: two variables with the same color will not be connected by an edge and so can be assigned to the same register (since they change in different time intervals).

The q -coloring problem of random graphs represents a very active field of research in discrete mathematics which constitutes the natural evolution of the percolation theory initiated by Erdős and Rényi in the 1950s [8]. One point of contact between computer science and random graph theory arises from the observation that, for large random graphs, there exists a critical average connectivity beyond which the graphs become uncolorable by q colors with probability

tending to one as the graph size goes to infinity. This transition will be called the q -COL/UNCOL transition throughout this paper. The precise value of the critical connectivity depends, of course, on the number q of allowed colors and on the ensemble of random graphs under consideration. Graphs generated close to their critical connectivity are extraordinarily hard to color and therefore the study of critical instances is at the same time a well posed mathematical question as well as an algorithmic challenge for the understanding of the onset of computational complexity [9,10]. The notion of computational complexity refers to worst-case instances and therefore results for a given ensemble of problems might not be of direct relevance. However, on the more practical side, algorithms which are used to solve real-world problems display a huge variability of running times and a theory for their typical-case behavior, on classes of nontrivial random instances, constitutes the natural complement to the worst-case analysis. Similar to what happens for other very well-known combinatorial problems, e.g., the satisfiability problem of Boolean formulas, critical random instances of q coloring are a popular test bed for the performance of search algorithms [6].

From the physics side q coloring has a direct interpretation as a spin-glass model [11]. A proper coloring of a graph is simply a zero-energy ground-state configuration of a Potts antiferromagnet with q -state variables. For most lattices such a system is frustrated and displays all the equilibrium and out-of-equilibrium features of spin glasses (the ‘‘Potts glass’’).

Here we focus on the q -coloring problem (or Potts antiferromagnet) over random graphs of finite average connectivity, given by the $\mathcal{G}(N,p)$ ensemble: Graphs are composed of N vertices, every pair of them independently being connected by an edge with probability p , and being not directly connected with probability $1-p$. The relevant case of finite-connectivity graphs is described by $p=c/N$, with c staying constant in the thermodynamic limit $N\rightarrow\infty$. In this case, the expected number of edges becomes $M=c/N\binom{N}{2}\approx cN/2$, i.e., proportional to the vertex number. Each of the vertices is, on average, connected to c other vertices. This connectivity fluctuates according to a Poissonian, i.e., the probability of randomly selecting a vertex with exactly d neighbors is given by $e^{-c}c^d/d!$.

Two types of questions can be asked. One type is algorithmic, i.e., finding an algorithm that decides whether a given graph is colorable. The other type is more theoretical and amounts to asking whether a typical problem instance is colorable or not and what the typical structure of the solution space is. Here we address both questions using the so-called cavity method [12]. First, we provide a detailed description of the analytical calculations beyond the results presented in Ref. [13], where the questions of the coloring threshold and typical solution properties were addressed. Second, this analytical description is modified and applied to single-graph instances. This leads to an efficient graph coloring algorithm for the region slightly below the COL/UNCOL transition. In this region, known, complete, and stochastic algorithms are known to fail already for moderate system sizes.

Let us start with reviewing some known results on the

q -COL/UNCOL transition on random graphs. One of the first important finite-connectivity results was obtained by Luczak about one decade ago [14]. He proved that the threshold asymptotically grows like $c_q\sim 2q\ln q$ for large numbers of colors, a result, which up to a prefactor coincides with the outcome of a replica calculation on highly connected graphs [15] [$p=O(1)$ for large N]. For fixed number q of colors, all vertices with less than q neighbors, i.e., of *degree* smaller than q , can be colored for sure. The hardest to color structure is thus given by the maximal subgraph having minimal degree at least q , the so-called q core. Pittel, Spencer, and Wormald [16] showed that the emergence of a 2-core coincides with the percolation transition of random graphs at $c=1$ [8] and is continuous. For $q\geq 3$, however, the q core arises discontinuously, jumping from zero to a finite fraction of the full graph. For $q=3$ they found, e.g., that the core emerges at $c\approx 3.35$ and immediately contains about 27% of all vertices. Shortly after, it was realized that the existence of the core is necessary, but in no way sufficient for uncolorability [17]. In fact, the best lower bound for the 3-COL/UNCOL transition is 4.03 [18] and numerical results predict a threshold of about 4.7 [19]. The currently best rigorous upper bound is 4.99 [20]. It was obtained using a refined first-moment method. In statistical mechanics, the latter is known as the annealed approximation. More recently, a replica-symmetric analysis of the problem has been performed [21]. The resulting threshold 5.1 exceeds, however, the rigorous bound and one has to go beyond replica symmetry. At the level of one-step replica-symmetry breaking we are able to calculate a threshold value $c_3\approx 4.69$, which we believe to be exact. We also describe the solution space structure which undergoes a clustering transition at $c_d\approx 4.42$.

The remaining part of the paper is organized as follows. In Sec. II we present the replica-symmetric (RS) solution of the problem and discuss why it fails. In Sec. III the one-step replica-symmetry breaking (RSB) solution is presented. From this we derive the average graph connectivity for the q -COL/UNCOL transition and demonstrate the existence of a dynamic threshold associated with the appearance of solution clusters in configuration space. Then, in the following section we show that the previous ideas are valid even in the analysis of single-case instances. This allows us, in Sec. V, to propose an algorithm that colors, in the hard region, single instances in polynomial time. Finally, in Sec. VI some conclusions of the work are presented.

II. REPLICA-SYMMETRIC SOLUTION

As stated above, the question if a given graph is q colorable is equivalent to the question if there are zero-energy ground states of the antiferromagnetic q -state Potts model defined on the same graph. Denoting the set of all edges by E , the problem can thus be described by the Hamiltonian

$$H_G = \sum_{\{i,j\}\in E} \delta(\sigma_i, \sigma_j), \quad (1)$$

where $\{\sigma_i\}\in\{1,2,\dots,q\}$ are the usual Potts spins and $\delta(\dots,\dots)$ denotes the Kronecker symbol. This Hamil-

tonian obviously counts the number of edges being colored equally on both extremities, a proper coloring of the graph thus has energy zero. Since this Hamiltonian cannot take negative values, the combinatorial task of finding a coloring is translated to the physical task of finding a zero-energy ground state, i.e., to the statistical physics of the above model at zero temperature.

A. The cavity equation

In this paper we therefore apply the cavity method in a variant recently developed for finite-connectivity graphs directly at zero temperature [22–24]. This approach consists of a self-consistent iterative scheme which is believed to be exact over local treelike graphs, like $\mathcal{G}(N, c/N)$, the set we consider here. It includes the possibility of dealing with the existence of many pure states. One has to first evaluate the energy shift of the system due to the addition of a new spin σ_0 . Let us assume for a moment that the new spin is only connected to a single spin, say σ_1 , in the preexisting graph. Before adding the new site 0, the ground-state energy of the system with fixed σ_1 can be expressed as

$$E^N(\sigma_1) = A - \sum_{\tau=1}^q h_1^\tau \delta(\tau, \sigma_1), \quad (2)$$

where we have introduced the effective field $\vec{h}_1 = (h_1^1, \dots, h_1^q)$ and used the superscript N to stress that the previous quantity refers to the N -site systems. Note that a $(q-1)$ -dimensional field would be sufficient since one of the q fields above can be absorbed in A . We, however, prefer to work with q field components in order to keep evident the global color symmetry. When we connect σ_0 to σ_1 we can express the minimal energy of the $(N+1)$ -site graph at fixed σ_0, σ_1 as

$$E^{N+1}(\sigma_0, \sigma_1) = A - \sum_{\tau=1}^q h_1^\tau \delta(\tau, \sigma_1) + \delta(\sigma_0, \sigma_1). \quad (3)$$

The minimum energy for the $(N+1)$ -site system at fixed σ_0 is thus obtained by minimizing $E^{N+1}(\sigma_0, \sigma_1)$ with respect to σ_1 , it can be written as

$$\begin{aligned} E^{N+1}(\sigma_0) &= \min_{\sigma_1} E^{N+1}(\sigma_0, \sigma_1) \\ &\equiv A - \omega(\vec{h}_1) - \sum_{\tau=1}^q \hat{u}^\tau(\vec{h}_1) \delta(\tau, \sigma_0) \end{aligned} \quad (4)$$

with

$$\omega(\vec{h}) = -\min(-h^1, \dots, -h^q), \quad (5)$$

$$\hat{u}^\tau(\vec{h}) = -\min(-h^1, \dots, -h^\tau + 1, \dots, -h^q) - \omega(\vec{h}), \quad (6)$$

where we have introduced the *cavity biases* $\hat{u}(\vec{h})$. This choice of ω and \hat{u} is not unique but, according to the previous discussion, we have chosen the only manifestly color-

symmetric notation. The structure of the cavity biases is easily understood if we distinguish among two different cases: (i) $h^\tau > h^1, \dots, h^{\tau-1}, h^{\tau+1}, \dots, h^q$ for some τ : then $\hat{u}^\tau = -1$ and $\hat{u}^\sigma = 0$ for all $\sigma \neq \tau$; (ii) $h^{\tau_1} = h^{\tau_2} \geq h^1, \dots, h^q$ for some τ_1, τ_2 : then $\hat{u} = (\hat{u}^1, \dots, \hat{u}^q) = (0, 0, \dots, 0) = \vec{0}$.

Only vectors \vec{h} with nondegenerate maximal component give rise to nontrivial cavity bias \vec{u} in the direction of the minimal component. This is physically understandable: A unique maximal field component in \vec{h} fixes the corresponding color, which thus is forbidden to the newly added site. If there are two or more maximal field components, the color of the old site is not fixed, thus there cannot be any forbidden color for the new vertex. Each cavity bias in our problem thus belongs to the $(q+1)$ -element set $\{\vec{0}, \vec{e}_1, \dots, \vec{e}_q\}$, where \vec{e}_τ has all components 0 but the τ th equal to -1 .

If the new spin σ_0 is connected to k randomly chosen sites with fields $\vec{h}_1, \dots, \vec{h}_k$, the cavity bias has to be linearly superposed and the resulting cavity field on vertex 0 is given by $\vec{h}_0 = \sum_{i=1}^k \hat{u}(\vec{h}_i)$. With high probability (tending to one for large N) the k sites will be far from each other in the original graph: Although an extensive number of loops is surely present for $c > 1$ [8], these loops have lengths of the order of $\ln N$. Inside one Boltzmann state we can thus invoke the *clustering propriety* [11], resulting in a statistical independence of the k selected sites and their cavity fields h_i (for a more detailed discussion of this point see Refs. [23,24]). The simplest ansatz assumes that there exists just one such state (or a finite number, such as in ferromagnets at low temperature), which is equivalent to the Bethe-Peierls iterative scheme or the replica-symmetric ansatz in the replica method. Assuming furtheron the existence of a well-defined thermodynamic limit of the energy density E/N and of the probability distributions of local fields (for recent rigorous studies in this direction, see Refs. [25–27]), the distribution of the fields \vec{h}_0 of the newly added vertex becomes the equal to those of the k neighbors. It is consequently determined by the closed expression

$$\begin{aligned} P(\vec{h}) &= e^{-c} \sum_{k=0}^{\infty} \frac{c^k}{k!} \int d^q \vec{h}_1, \dots, d^q \vec{h}_k P(\vec{h}_1) \dots P(\vec{h}_k) \\ &\times \delta\left(\vec{h} - \sum_{i=1}^k \hat{u}_i(\vec{h}_i)\right), \end{aligned} \quad (7)$$

$$Q(\vec{u}) = \int d^q \vec{h} P(\vec{h}) \delta(\vec{u} - \hat{u}(\vec{h})), \quad (8)$$

where we have already used that the connectivities k are distributed according to a Poissonian of mean c . The previous equations can be combined in order to have a closed form for $Q(\vec{u})$:

$$\begin{aligned} Q(\vec{u}) &= e^{-c} \sum_{k=0}^{\infty} \frac{c^k}{k!} \int \prod_{i=1}^k d^q \vec{u}_i Q(\vec{u}_i) \\ &\times \delta(\vec{u} - \hat{u}(\vec{u}_1 + \vec{u}_2 + \dots + \vec{u}_k)). \end{aligned} \quad (9)$$

From the symmetry of our model under arbitrary permutations of the colors we conclude that

$$Q(\vec{e}_1) = Q(\vec{e}_2) = \dots = Q(\vec{e}_q) = \eta \quad \text{and} \quad Q(\vec{0}) = 1 - q, \quad (10)$$

i.e., we need a single real number η with $0 < \eta < 1/q$ to completely specify the probability distribution function $Q(\vec{u})$. Noting that the probability $P^{(k)}(\vec{h})$ for a site with k neighbors can be expressed by

$$P^{(k)}(\vec{h}) = \int \prod_{i=1}^k d^q \vec{u}_i Q(\vec{u}_i) \delta\left(h - \sum_{i=1}^k \vec{u}_i\right), \quad (11)$$

and recalling that $\vec{u}_i \in \{\vec{0}, \vec{e}_1, \dots, \vec{e}_q\}$ it is easy to rewrite this probability distribution in a compact multinomial form

$$P^{(k)}(\vec{h}) \equiv P^{(k)}(h^1, h^2, \dots, h^q) = \frac{k! \eta^{-\sum_{\tau=1}^q h^\tau (1-q\eta)^{k+\sum_{\tau=1}^q h^\tau}}}{\left(k + \sum_{\tau=1}^q h^\tau\right)! \prod_{\tau=1}^q (-h^\tau)!}, \quad (12)$$

with the convention that $1/n! = 0$ for $n < 0$. Note that $h^\tau \in (0, -1, \dots, -k)$ and that there are correlations among the different components of the cavity fields such that $P^{(k)}(h^1, \dots, h^q) \neq \prod_{\tau=1}^q P(h^\tau)$. Now we are ready to calculate the graph average over the Poissonian connectivity distribution of mean c ,

$$\begin{aligned} P(h^1, \dots, h^q) &= e^{-c} \sum_k \frac{c^k}{k!} P^{(k)}(h^1, \dots, h^q) \\ &= e^{-c} \prod_{\tau=1}^q \frac{(c\eta)^{-h^\tau}}{(-h^\tau)!} \equiv \prod_{\tau=1}^q \mathcal{P}_{c\eta}(h^\tau). \end{aligned} \quad (13)$$

It is interesting to note that after the average the correlations among the different colors disappear and P is the product of q Poissonian distributions with average $c\eta$. From Eq. (7) it is possible to derive a self-consistent equation for the order parameter noting that the probability η to obtain a nontrivial cavity bias—say \vec{e}_τ —is simply given by the probability that the τ th component of the local field is the nondegenerate smaller, so setting $\tau=1$,

$$\begin{aligned} \eta &= \sum_{h^1=0}^{\infty} \sum_{h^2=h^1+1}^{\infty} \dots \sum_{h^q=h^1+1}^{\infty} P(h^1, \dots, h^q) \\ &= e^{-c\eta} \sum_{n=0}^{\infty} \frac{(c\eta)^n}{n!} \left(1 - \frac{\Gamma(n+1, c\eta)}{\Gamma(n+1)}\right)^{q-1}, \end{aligned} \quad (14)$$

where $\Gamma(n, x)$ is the incomplete Gamma function defined from the following useful relation:




FIG. 1. Replica-symmetric order parameter η vs average connectivity c for $q=3$ from Eq. (14).

$$e^{-x} \sum_{k=n}^{\infty} \frac{x^k}{k!} = 1 - \frac{\Gamma(n, x)}{\Gamma(n)}. \quad (15)$$

The sum in Eq. (14) converges very fast. It is therefore easy to numerically construct a solution to this equation as a function of c . For $q > 2$ it turns out that η jumps discontinuously from zero to a finite value as shown in Fig. 1 where the order parameter η jumps at $c=5.141$ in the case of $q=3$.

This means that, up to $c=5.141$ and at the level of the replica-symmetric assumption, we only find the paramagnetic solution $\eta=0$. The solution $\eta > 0$ would account in a spontaneous breaking of this symmetry, there should be a finite number of pure states (similar to Neel order in Ising antiferromagnets).

B. The calculation of the energy

One can easily compute the average shift in the ground-state energy when a new spin is added to the N -site system and it is connected to k spins of the system. The energy of the original graph is given by $A - \sum_{i=1}^k \omega(\vec{h}_i)$ while the energy of the $(N+1)$ -site system is $A - \sum_{i=1}^k [\omega(\vec{h}_i) + \omega(\hat{u}(\vec{h}_i))]$. Therefore the average shift is given by

$$\begin{aligned} \Delta E_1 &= - \sum_{k=0}^{\infty} e^{-c} c^k / k! \int d^q \vec{u}_1 \dots d^q \vec{u}_k Q(\vec{u}_1) \dots dQ(\vec{u}_k) \\ &\quad \times \omega\left(\sum_{i=1}^k \vec{u}_i\right) \\ &= - \int d^q \vec{h} P(\vec{h}) \omega(\vec{h}). \end{aligned} \quad (16)$$

One might be tempted to conclude that Eq. (16) equals the energy density of the system, at least for N large enough, but this is not true. There is a correction term due to the change in the number of links per variable in the iteration $N \rightarrow N+1$. In fact, generating links with probability c/N in a $N+1$ system, instead of $c/(N+1)$, we are slightly overgenerating links. So, we need to calculate the average energy shift

in a system when two sites—say, spins σ_1 and σ_2 —are joined by an antiferromagnetic link.

Again, the energy of the original graph is $A - \omega(\vec{h}_1) - \omega(\vec{h}_2)$, while the energy after the two spins are joined is given by $A - \min_{\sigma_1, \sigma_2} [-h_1^{\sigma_1} - h_2^{\sigma_2} + \delta(\sigma_1, \sigma_2)]$. The difference between the two contributions can be written as

$$\begin{aligned} \Delta E_{\text{link}} &= \min_{\sigma_1} \{ -h_1^{\sigma_1} + \min_{\sigma_2} [-h_2^{\sigma_2} + \delta(\sigma_1, \sigma_2)] \} + \omega(\vec{h}_1) \\ &\quad + \omega(\vec{h}_2) \\ &= \min_{\sigma_1} [-h_1^{\sigma_1} - u^{\sigma_1}(\vec{h}_2) - \omega(\vec{h}_2)] + \omega(\vec{h}_1) + \omega(\vec{h}_2) \\ &= -\omega(\vec{h}_1 + \hat{u}(\vec{h}_2)) + \omega(\vec{h}_1). \end{aligned} \quad (17)$$

This allows us to express the average link-energy shift as

$$\Delta E_2 = \int d^q \vec{h}_1 d^q \vec{h}_2 P(\vec{h}_1) P(\vec{h}_2) [\omega(\vec{h}_1) - \omega(\vec{h}_1 + \hat{u}(\vec{h}_2))]. \quad (18)$$

It is interesting to observe that Eqs. (16) and (17) are *model independent* in the sense that the actual Hamiltonian is encoded into the functions $\omega(\vec{h})$ and $\hat{u}(\vec{h})$ defined by Eq. (5).

Using Eqs. (5) and (13) one shows easily that Eq. (16) reduces to

$$\begin{aligned} \Delta E_1 &= \sum_{h^1 \dots h^q} \mathcal{P}_{c\eta}(h^1) \dots \mathcal{P}_{c\eta}(h^q) \\ &\quad \times \min(-h^1, -h^2, \dots, -h^q) \\ &= - \sum_{\alpha=0}^{q-1} \binom{q}{q-\alpha} \sum_{h=0}^{-\infty} h \mathcal{P}_{c\eta}(h)^{q-\alpha} \left(\sum_{g < h}^{-\infty} \mathcal{P}_{c\eta}(g) \right)^\alpha. \end{aligned} \quad (19)$$

It is also not hard to prove that the average link-energy shift $\Delta E_2 = q\eta^2$. This result can be obtained either by direct computation of the integral or following a simple probabilistic argument: ΔE_{link} is different from zero whenever the two unlinked sites have the same color, but this happens with probability η^2 for each of the colors. Finally, we have the the following equation for the energy which is equivalent to the replica-symmetric approximation:

$$\begin{aligned} E &= N \left(\Delta E_1 - \frac{c}{2} \Delta E_2 \right) \\ &= - \sum_{\alpha=0}^{q-1} \binom{q}{q-\alpha} \sum_{h=0}^{-\infty} h \mathcal{P}_{c\eta}(h)^{q-\alpha} \left(\sum_{g < h}^{-\infty} \mathcal{P}_{c\eta}(g) \right)^\alpha \\ &\quad - \frac{c}{2} q \eta^2. \end{aligned} \quad (20)$$

The behavior of the energy for $q=3$ as a function of the average connectivity c is displayed in Fig. 2. Let us note that for average connectivity $5.141 < c < 5.497$ the energy is




FIG. 2. Energy density e vs average connectivity c for $q=3$ in the RS approximation from Eq. (20).

negative, a particularly baffling result if we consider that Hamiltonian (1) is at least positively defined. This phenomenon is analogous to what already observed for the RS approximation in random 3-SAT [39], and is a consequence of the approximation used. We will see in the following section how the one-step RSB (1-RSB) ansatz cures this pathology. However, before leaving the RS approximation we would like to compare our RS results with the RS approximation presented recently by van Mourik and Saad in Ref. [21] since their result clearly differs from ours. At the origin of the discrepancy is the fact that they work a population dynamics at very low but finite temperature finding a transition around $c=5.1$ but without negative energy region. Analogous to what is reported in Ref. [23], if one works directly at a zero temperature the distribution $P(\vec{h})$ must be concentrated around integer field components, but this is not true anymore at a temperature different from zero, as it happens in Ref. [21]. They find that, in the zero-temperature limit, there remain noninteger field components. In our opinion these are a direct hint to the existence of RSB. Instead of including these fields into an extended replica-symmetric approach, we directly switch to a replica-symmetry broken solution.

III. ONE-STEP RSB SOLUTION

The RS results show some evident pathologies and are at odd with numerical simulations [19,21] which predict a lower threshold around $c=4.7$, and with the rigorous upper bound $c=4.99$ [20]. What can be wrong in our analysis? The main assumption we have made is the statistical independence of the of the k cavity fields. Is it true that long distance among spins imply statistical independence? In general, the answer we obtain from statistical mechanics is “no”: the assumption holds only inside pure states (for a discussion on the notion of pure state, see Ref. [28] and references therein).

In this section we will focus on how the cavity method could be used to handle a situation in which there exist many different pure states. More precisely we assume that their number $\mathcal{N} \propto e^{N\Sigma}$ is exponentially large in N . The connectivity-dependent exponent Σ is called *complexity* and

denotes the entropy density of clusters. Note that it differs in general from the solution entropy since each cluster may contain as well an exponential number of solutions. The first basic assumption we made is that inside each pure state the clustering condition holds. Under this assumption the iteration can still be applied, but we have to take into account the reshuffling of energies of different states when new spins are added.

A. 1-RSB cavity equation

We proceed following the same steps of the preceding section. Let us take the new spin σ_0 and let us connect it to k spins $\sigma_1, \dots, \sigma_k$ in the same state α . Thanks to the fast decrease of correlations inside a pure state the energy of state α for fixed value of the k spins is

$$E_\alpha^N(\sigma_1, \dots, \sigma_k) = A_\alpha - \sum_{i=1}^k \sum_{\tau=1}^q h_{i,\alpha}^\tau \delta(\tau, \sigma_{i,\alpha}). \quad (21)$$

The optimization step within each pure state α runs still in close analogy to the RS computation: when we connect σ_0 to $\sigma_1, \dots, \sigma_k$, we express the minimal energy of the $(N+1)$ -site graph with fixed σ_0 ; by minimizing the $(N+1)$ -site system at fixed σ_0 is thus obtained by minimizing E_α^{N+1} with respect to the k spins:

$$E_\alpha^{N+1}(\sigma_0) = A_\alpha - \sum_{i=1}^k \omega(\vec{h}_{i,\alpha}) - \sum_{i=1}^k \sum_{\tau=1}^q \hat{u}^\tau(\vec{h}_{i,\alpha}) \delta(\tau, \sigma_0). \quad (22)$$

This last equation shows that the local field acting on the new spin σ_0 in the state α is

$$\vec{h}_{0,\alpha} = \sum_{i=1}^k \hat{u}(\vec{h}_{i,\alpha}) \quad (23)$$

and that the energy shift inside a state is

$$\Delta E_\alpha = - \sum_{i=1}^k \omega(\hat{u}(\vec{h}_{i,\alpha})). \quad (24)$$

All the previous equations are completely equivalent to those in the RS case except for the fact that now we have an α index labeling the different pure states. One natural question is how cavity fields and the related cavity biases are distributed for a given site among the different pure states. This leads us to the notion of *survey* [22–24], i.e., the site-dependent normalized histogram over the different states of both cavity biases and cavity fields:

$$\begin{aligned} Q_i(\vec{u}_i) &= \frac{1}{\mathcal{M}} \sum_{\alpha=1}^M \delta(\vec{u}_i - \vec{u}_{i,\alpha}), \\ P_i(\vec{h}_i) &= \frac{1}{\mathcal{M}} \sum_{\alpha=1}^M \delta(\vec{h}_i - \vec{h}_{i,\alpha}). \end{aligned} \quad (25)$$

In close analogy with what we have already done in the RS case, the existence of a well-defined thermodynamic limit

implies that there must exist unique functional probability distributions $\mathcal{Q}[Q(\vec{u})]$ and $\mathcal{P}[P(\vec{u})]$ for all the surveys. One may wonder how could we handle such a big functional space: Fortunately the Q surveys are described in terms of a single real number $0 \leq \eta \leq 1/q$, cf. Eq. (10), and scalar function $\rho(\eta)$ is enough for specifying their distribution:

$$\begin{aligned} \mathcal{Q}[Q(\vec{u})] &= \int d\eta \rho(\eta) \delta \left[Q(\vec{u}) - (1-q)\eta \delta(\vec{u}) \right. \\ &\quad \left. - \eta \sum_{\tau=1}^q \delta(\vec{u} - \vec{e}_\tau) \right], \end{aligned} \quad (26)$$

with $\delta[\dots]$ denoting a functional Dirac distribution. Assuming that the survey of site 0 is distributed equally to those of all its k neighbors, we can write

$$\begin{aligned} P_0(\vec{h}) &= e^{-c} \sum_{k=0}^{\infty} \frac{c^k}{k!} C_k \int d^q \vec{u}_1 Q_1(\vec{u}_1) \cdots d^q \vec{u}_k Q_k(\vec{u}_k) \\ &\quad \times \exp \left[y \omega \left(\sum_{i=1}^k \vec{u}_i \right) \right] \delta \left(\vec{h} - \sum_{i=1}^k \vec{u}_i \right), \end{aligned} \quad (27)$$

$$Q_0(\vec{u}) = \int d^q \vec{h} P_0(\vec{h}) \delta(\vec{u} - \hat{u}(\vec{h})). \quad (28)$$

Note the presence of the *reweighting factor* $\exp[y\omega(\sum_{i=1}^k \vec{u}_i)]$ that arises after conditioning the probability distributions of the \vec{h} 's to a given value of energy [23], the prefactors C_k are normalization constants depending on $Q_1(\vec{u}), \dots, Q_k(\vec{u})$. The reweighting parameter y is a number equal to the derivative of the complexity $\Sigma(e)$ of metastable states with respect to their energy density $e = E/N$:

$$y = \frac{\partial \Sigma}{\partial e}. \quad (29)$$

Intuitively, this reweighting factor can be understood as a penalty $e^{-y\Delta E_\alpha}$ one has to pay for positive energy shifts. Note that Eqs. (27) and (28) can be cast in the following form:

$$\begin{aligned} Q_0(\vec{u}_0) &= e^{-c} \sum_{k=0}^{\infty} \frac{c^k}{k!} C_k \int d^q \vec{u}_1 Q_1(\vec{u}_1) \cdots d^q \vec{u}_k Q_k(\vec{u}_k) \\ &\quad \times \exp \left[y \omega \left(\sum_{i=1}^k \vec{u}_i \right) \right] \delta(\vec{u}_0 - \hat{u}(\vec{u}_1 + \cdots + \vec{u}_k)) \\ &= e^{-c} \sum_{k=0}^{\infty} \frac{c^k}{k!} C_k \int d^q \vec{h} \tilde{P}(\vec{h}) e^{y\omega(\vec{h})} \delta(\vec{u}_0 - \hat{u}(\vec{h})). \end{aligned} \quad (30)$$

In the last line we have introduced the auxiliary distribution $\tilde{P}(\vec{h})$ which would result in Eq. (27) without reweighting

(i.e., by setting $y=0$). It has no direct physical meaning in this context, but it will be of great technical help in the following calculations.

Let us first concentrate on the *colorable phase*, where the ground states are proper q colorings and have zero energy. Consequently no positive energy shifts are allowed, so this phase is characterized by $y=\infty$. Let us first calculate the value of the normalization constants C_k in this limit. Note that $\omega(\vec{h})\leq 0$ for all allowed \vec{h} (each component of \vec{h} is nonpositive as \vec{h} results from a sum over \vec{u} s). This means that the only surviving terms in Eq. (30) are those with zero-energy shift $\omega(\vec{h})=0$, i.e., all fields must have at least one zero component, allowing for the selection of at least one color without violating an edge. Let us first specialize to the case $q=3$ for clarity, the generalization to arbitrary q is straightforward. Summing over \vec{u}_0 both sides of Eq. (30) we have

$$\frac{1}{C_k} = \tilde{P}(0,0,0) + 3 \sum_{h^1 < 0} \tilde{P}(h^1,0,0) + 3 \sum_{h^1, h^2 < 0} \tilde{P}(h^1, h^2, 0), \quad (31)$$

where the combinatorial factors 3 appearing on the right-hand sides are obtained by noting that $\tilde{P}(h,0,0) = \tilde{P}(0,h,0) = \tilde{P}(0,0,h)$ and that $\tilde{P}(h^1, h^2, 0) = \tilde{P}(h^1, 0, h^2) = \tilde{P}(0, h^1, h^2)$. Combining Eqs. (27), (28), and (10) we get

$$\tilde{P}(0,0,0) = \prod_{i=1}^k (1-3\eta_i), \quad (32)$$

$$\begin{aligned} \sum_{h^1 < 0} \tilde{P}(h^1, 0, 0) &= \prod_{i=1}^k (1-2\eta_i) - \tilde{P}(0,0,0) \\ &= \prod_{i=1}^k (1-2\eta_i) - \prod_{i=1}^k (1-3\eta_i), \end{aligned} \quad (33)$$

$$\begin{aligned} \sum_{h^1, h^2 < 0} \tilde{P}(h^1, h^2, 0) &= \prod_{i=1}^k (1-\eta_i) - 2 \sum_{h^1 < 0} \tilde{P}(h^1, 0, 0) \\ &\quad - \tilde{P}(0,0,0), \\ &= \prod_{i=1}^k (1-\eta_i) - 2 \prod_{i=1}^k (1-2\eta_i) \\ &\quad + \prod_{i=1}^k (1-3\eta_i). \end{aligned} \quad (34)$$

Plugging these relations into Eq. (31) we finally get

$$\frac{1}{C_k} = 3 \prod_{i=1}^k (1-\eta_i) - 3 \prod_{i=1}^k (1-2\eta_i) + \prod_{i=1}^k (1-3\eta_i). \quad (35)$$

Also note that in close analogy to the analysis that leads to Eq. (14), we can interpret Eq. (34) as the (unnormalized)

probability of having the survey in site 0 *pointing* in direction \vec{e}_3 . Therefore combining Eqs. (34) and (35) we obtain

$$\begin{aligned} \eta_0 &= \hat{f}_k(\eta_1, \dots, \eta_k) \\ &= \frac{\prod_{i=1}^k (1-\eta_i) - 2 \prod_{i=1}^k (1-2\eta_i) + \prod_{i=1}^k (1-3\eta_i)}{3 \prod_{i=1}^k (1-\eta_i) - 3 \prod_{i=1}^k (1-2\eta_i) + \prod_{i=1}^k (1-3\eta_i)}. \end{aligned} \quad (36)$$

At this point we are ready to write the one-RSB iterative equation for the Q surveys:

$$\begin{aligned} \rho(\eta) &= e^{-c} \sum_{k=0}^{\infty} \frac{c^k}{k!} \int d\eta_1 \rho(\eta_1) \cdots d\eta_k \rho(\eta_k) \\ &\quad \times \delta(\eta - \hat{f}_k(\eta_1, \dots, \eta_k)). \end{aligned} \quad (37)$$

Equation (36) can be easily generalized to an arbitrary number q of colors,

$$\hat{f}_k(\eta_1, \dots, \eta_k) = \frac{\sum_{l=0}^{q-1} (-1)^l \binom{q-1}{l} \prod_{i=1}^k [1-(l+1)\eta_i]}{\sum_{l=0}^{q-1} (-1)^l \binom{q}{l+1} \prod_{i=1}^k [1-(l+1)\eta_i]}. \quad (38)$$

The self-consistent equation (37) resembles a replica-symmetric equation and can be solved numerically using a population dynamic algorithm: (i) Start with an initial population $\eta_1, \dots, \eta_{\mathcal{M}}$ of size \mathcal{M} which can be easily chosen to be as large as 10^6 to generate high-precision data; (ii) randomly draw a number k from the Poisson distribution $e^{-c}c^k/k!$; (iii) randomly select $k+1$ indices i_0, i_1, \dots, i_k from $\{1, \dots, \mathcal{M}\}$; (iv) update the population by replacing η_{i_0} by $f_d(\eta_{i_1}, \dots, \eta_{i_k})$; and (v) go to (ii) until convergence of the algorithm is reached.

One obvious solution of Eqs. (37) and (38) is the paramagnetic solution $\delta(\eta)$. For small average connectivities c it is the only one. The appearance of a nontrivial solution coincides with a clustering transition of ground states into an exponentially large number of extensively separated clusters. In spin-glass theory, this transition is called dynamical. Still, $\rho(\eta)$ will contain a nontrivial peak in $\eta=0$ due to small disconnected subgraphs, dangling ends, low-connectivity vertices, etc. The shape of $\rho(\eta)$ in the case $q=3$ is displayed in Fig. 3 for connectivities c ranging from c_d to c_c .

The weight t of this peak can be computed self-consistently. Let us again consider first the case $q=3$. Keeping in mind that for $y \rightarrow \infty$ the field \vec{h} has at least one vanishing component, the only possibilities to obtain $\hat{u}(\vec{h}) = \vec{0}$ are given by $\vec{h} = \vec{0}$ or by a field \vec{h} with one single nonzero component. So the probability that the cavity field acting on a given site with k neighbors equals zero is given by the sum of the probabilities that all neighboring cavity fields are zero (equal to t^k) plus the probability that exactly one cavity bias among the k is nontrivial [equal to $k(1-t)t^{k-1}$]. The average over the Poissonian degree distribution leads to




FIG. 3. Probability distribution function $\rho(\eta)$ in the case $q=3$ for average connectivities $4.42 < c < 4.69$. Note also that a δ peak at $\eta=0$ is always present (but not displayed here).

$$t = e^{-c} \sum_{k=0}^{\infty} \frac{c^k}{k!} [t^k + k t^{k-1} (1-t)] = e^{-(1-t)c} [1 + (1-t)c]. \quad (39)$$

Generalizing Eq. (39) to a general number q of colors easily gives

$$t = e^{-(1-t)c} \sum_{l=0}^{q-2} \frac{(1-t)^l c^l}{l!}. \quad (40)$$

This equation is quite interesting, since a nontrivial solution forms a necessary condition for Eq. (37) to have a nontrivial solution. In fact, this equation was first found in Ref. [16], the fraction of edges belonging to the q core is given by $(1 - t_{min})$, with t_{min} being the smallest positive solution of Eq. (40). Thus, we also find that the existence of an extensive q core is necessary for a nontrivial $\rho(\eta)$ and forms a lower bound for the q -COL/UNCOL transition.

Unlike in the case of finite-connectivity p -spin glasses or, equivalently, random XOR-SAT problems [30–32], the existence of a solution $t < 1$ is not sufficient for a nontrivial $\rho(\eta)$ to exist. The latter appears suddenly at the dynamical transition c_d , which can be determined to high precision using the population dynamical algorithm. This solution does not imply uncolorability, but the set of solutions is separated into an exponentially large number of clusters. The number of these clusters, or more precisely its logarithm divided by the graph size N , is called the complexity Σ and can be calculated from $\rho(\eta)$.

B. The calculation of energy and complexity

More generally, we also expect a large number of metastable states at nonzero energy to exist. Hereafter we will assume that they are exponentially many, $\mathcal{N}(e) \propto \exp[N\Sigma(e)]$, where the complexity $\Sigma(e)$ is (despite the use of a capital letter) an intensive function of the energy density $e = E/N$. We can introduce a thermodynamic potential $\phi(y)$ [29] as

$$\phi(y) = -\frac{1}{yN} \ln \left(\int de e^{N\{-ye + \Sigma(e)\}} \right). \quad (41)$$

For large N , we calculate this integral by its saddle point:

$$\phi(y) = \min_e \left(e - \frac{1}{y} \Sigma(e) \right) = e_{sp} - \frac{1}{y} \Sigma(e_{sp}). \quad (42)$$

It is easily verified that potential ϕ calculated at the saddle point energy $e_{sp}(y)$ fulfills the usual Legendre relation:

$$\partial_y [y \phi(y, e_{sp}(y))] = e_{sp}, \quad (43)$$

$$y^2 \partial_y^2 \phi(y, e_{sp}(y)) = \Sigma'(e_{sp}). \quad (44)$$

Around the saddle point the complexity can be approximated, according to Eq. (42), by

$$\Sigma(e) \approx \Sigma(e_{sp}) + y(e - e_{sp}) = -y\phi(y) + ye. \quad (45)$$

We will now consider a cavity argument: let us denote by E_N the energy of a system composed of N sites, then the density of configurations is given by

$$d\mathcal{N}_N(E_N) \propto e^{-y\Phi_N(y) + yE_N} dE_N, \quad (46)$$

with $\Phi_N(y)$ denoting the extensive thermodynamic potential with limit $\Phi_N(y)/N \rightarrow \phi(y)$. Now we add a spin to the system. If we consider that the total energy is $E_{N+1} = E_N + \Delta E$, we can express the density of configurations in terms of E_N and ΔE :

$$d\mathcal{N}_{N+1}(E_N, \Delta E) \propto e^{y\Phi_N(y) + y(E_N + \Delta E)} dE_N P(\Delta E) d\Delta E. \quad (47)$$

Integrating over δE we get

$$d\mathcal{N}_{N+1}(E_{N+1}) = C e^{-y\Phi_{N+1}(y) + yE_{N+1}} dE_{N+1}, \quad (48)$$

$$C = \frac{1}{y} \int P(\Delta E) e^{y\Delta E} d\Delta E \equiv \frac{1}{y} \langle e^{y\Delta E} \rangle_{P(\Delta E)}. \quad (49)$$

Comparing the previous equations with Eq. (46) we can deduce that

$$\Phi_{N+1}(y) = \Phi_N(y) - \frac{1}{y} \ln \langle \exp(y\Delta E) \rangle_{P(\Delta E)}. \quad (50)$$

In the thermodynamic limit we can thus identify

$$\phi(y) = -\frac{1}{y} \ln \langle \exp(y\Delta E) \rangle_{P(\Delta E)}. \quad (51)$$

In close analogy with what we have already done in the RS case, and using Eq. (24), we can compute ϕ as a *site* contribution plus a *link* contribution in the 1-RSB scenario: site addition

$$\begin{aligned} \exp(-y\Delta\phi_1) &= \int d^q \vec{u}_{i_1} Q_{i_1}(\vec{u}_{i_1}) \cdots d^q \vec{u}_{i_k} Q_{i_k}(\vec{u}_{i_k}) \\ &\times \exp\left[y\omega\left(\sum_{j=1}^k \vec{u}_{i_j}\right)\right] = \frac{1}{C_k}, \end{aligned} \quad (52)$$

link addition

$$\begin{aligned} \exp(-y\Delta\phi_2) &= \int d^q \vec{h}_{i_1} P_{i_1}(\vec{h}_{i_1}) d^q \vec{h}_{i_2} P_{i_2}(\vec{h}_{i_2}) \\ &\times \exp[-y\omega(\vec{h}_{i_1}) + y\omega(\vec{h}_{i_1} + \hat{u}(\vec{h}_{i_2}))] \\ &= \int d^q \vec{h} P_{i_1}(\vec{h}) d^q \vec{u} Q_{i_2}(\vec{u}) \\ &\times \exp[-y(\omega(\vec{h}) - \omega(\vec{h} + \vec{u}))] \\ &= 1 + q\eta_{i_1}\eta_{i_2}(e^{-y} - 1). \end{aligned} \quad (53)$$

Note that in the limit $y \rightarrow 0$ and assuming $P_i = P$ for each site, we obtain the RS expressions. Once the functional distributions $\mathcal{Q}[Q(\vec{u})]$ and $\mathcal{P}[P(\vec{h})]$ are known we can eventually average the energy shifts $\Delta\phi_1, \Delta\phi_2$ in the usual linear combination:

$$\phi(y) = \overline{\Delta\phi_1} - \frac{c}{2} \overline{\Delta\phi_2}, \quad (54)$$

where the overlines denote the average over both disorder and functional distributions. One finally finds

$$\begin{aligned} \phi(y) &= -\frac{1}{y} \sum_{k=1}^{\infty} e^{-c} \frac{c^k}{k!} \int \prod_{i=0}^k \mathcal{D}Q_i \mathcal{Q}[Q_i] \\ &\times \ln\left\{ \int \prod_{i=0}^k d^q \vec{u}_i Q_i(\vec{u}_i) \exp\left[y\omega\left(\sum_{i=1}^k \vec{u}_i\right)\right] \right\} \\ &+ \frac{c}{2y} \int \mathcal{D}P_1 \mathcal{P}[P_1] \mathcal{D}P_2 \mathcal{P}[P_2] \\ &\times \ln\left(\int d^q \vec{h}_1 P_1(\vec{h}_1) d^q \vec{h}_2 P_2(\vec{h}_2) \right. \\ &\left. \times \exp[y\omega(\vec{h}_1) - y\omega(\vec{h}_1 + \hat{u}(\vec{h}_2))] \right). \end{aligned} \quad (55)$$

In the limit $y \rightarrow \infty$ these relations can be written in a more explicit form. Let us consider first the term $\Delta\phi_1$ in Eq. (52). Referring to Eq. (35) it easy to see that

$$\lim_{y \rightarrow \infty} e^{-y\Delta\phi_1} = \sum_{l=0}^{q-1} (-1)^l \binom{q}{l+1} \prod_{i=1}^k [1 - (l+1)\eta_i] \quad (56)$$

such that

$$\begin{aligned} \lim_{y \rightarrow \infty} -y\overline{\Delta\phi_1} &= \sum_{k=1}^{\infty} e^{-c} \frac{c^k}{k!} \int d\rho(\eta_1) \cdots d\rho(\eta_k) \\ &\times \ln\left[\sum_{l=0}^{q-1} (-1)^l \binom{q}{l+1} \prod_{i=1}^k [1 - (l+1)\eta_i] \right]. \end{aligned} \quad (57)$$

In order to compute the average link contribution $\overline{\Delta\phi_2(y)}$ we need to evaluate the large y limit of Eq. (53), which gives

$$\lim_{y \rightarrow \infty} -y\overline{\Delta\Phi_2(y)} = \int d\rho(\eta_1) d\rho(\eta_2) \ln(1 - q\eta_1\eta_2). \quad (58)$$

This equation has a good probabilistic interpretation complementary to that used in the derivation of ΔE_2 in the RS case. In fact the integrand of Eq. (53) is different from zero for $y \rightarrow \infty$ only when both sites i_1 and i_2 have a different color, and this happens with probability $(1 - q\eta_{i_1}\eta_{i_2})$ (note that $q\eta_{i_1}\eta_{i_2}$ is the probability that the two sites have the same color). It is now clear from Eq. (42) that taking $y \rightarrow \infty$ of $-y\Phi(y)$ gives us the complexity at least in the COL region where $e = 0$:

$$\begin{aligned} \Sigma(e=0) &= \sum_{k=1}^{\infty} e^{-c} \frac{c^k}{k!} \int d\eta_1 \rho(\eta_1) \cdots d\eta_k \rho(\eta_k) \\ &\times \ln\left[\sum_{l=0}^{q-1} (-1)^l \binom{q}{l+1} \prod_{i=1}^k [1 - (l+1)\eta_i] \right] \\ &- \frac{c}{2} \int d\eta_1 \rho(\eta_1) d\eta_2 \rho(\eta_2) \ln(1 - q\eta_1\eta_2). \end{aligned} \quad (59)$$

C. Results

The previous analysis results for the q -coloring problem in the existence of a dynamic transition, characterized by the sudden appearance of an exponential number of clusters that disconnect the solutions of the problem. This is represented in Fig. 4 for $q=3$ and 4, where the complexity is plotted as a function of the graph connectivity. Note that at a certain value average connectivity $c = c_d$ the complexity abruptly jumps from zero to a positive value. Then it decreases with growing c and disappears at c_q where the number of solutions become zero. It is not possible any more to find a zero-energy ground state for the system, i.e., the graph becomes uncolorable with q colors, and its chromatic number grows by one; see Fig. 5.

In Table I, we present the results for $q=3, 4$, and 5; for the dynamical transition we show the corresponding values of c_d of the entropy $s(c_d) = \ln q + c_d \ln(1 - 1/q)/2$ [33] and the complexity $\Sigma(c_d)$. For the q -COL/UNCOL transition, the critical connectivity c_q and the solution entropy are given. Like in random 3-satisfiability [34] and vertex covering [35], this entropy is found to be finite at the transition point.

In Fig. 6 we display the average complexity Σ as a function of the energy density e in the 1-RSB approximation.




FIG. 4. Top: Complexity $\Sigma(c)$ vs average connectivity for $q=3$ and $q=4$. Nonzero complexity appears discontinuously at the dynamical threshold c_d and goes down continuously to zero at the q -COL/UNCOL transition. The curves are calculated using the population-dynamical solution for $\rho(\eta)$ with population size $M=10^6$. Bottom: The full line shows the chromatic number of large random graphs vs their connectivity c . The symbols give results of small $N=10^3$, each averaged over 100 samples.

Recently Montanari and Ricci showed in Ref. [36] that in the p -spin spherical spin glass the 1-RSB scheme is correct only up to a certain critical energy density e_G , above which this solution becomes unstable and a FRSB calculation is required. It is possible that such a phenomenon might occur also in this case. The dynamical transition is not only characterized by a sudden clustering of ground states, at the same point an exponential number of metastable states of positive energy appears [24]. Such states (besides algorithm-dependent entropic barriers which may exist even below c_d)




FIG. 5. Average thermodynamic potential $\phi(y)$ vs y in the HARDCOL phase ($c=4.5$) and in the UNCOL phase ($c=5.0$). Note that $\phi(y)$ above the paramagnetic region ($\phi=0$) is a monotonously increasing function of y in the first case, while it displays a maximum at finite y in the second one.

TABLE I. Entropy and complexity at the static and dynamic thresholds for $q=3, 4$, and 5.

q	c_d	$s(c_d)$	$\Sigma(c_d)$	c_q	$s(c_q)$
3	4.42	0.203	0.0223	4.69	0.148
4	8.27	0.197	0.0553	8.90	0.106
5	12.67	0.196	0.0794	13.69	0.082

are expected to act as traps for local search algorithms causing an exponential slowing down of the search process. Well-known examples of search processes that are overwhelmed by the presence of excited states are simulated annealing or greedy algorithms based on local information.

To test this prediction, we have applied several of the best available solvers for coloring and SAT problems available in the net [6,38]. After some preliminary simulations we observed that the best results could be obtained with the small program [38] and concentrated our efforts on it. The simulation results, as shown in the lower half of Fig. 4, were obtained in the following way: First, a random graph ($N=10^3$) was generated and we tried to color it with a small number of colors (here $q=3$). If, after some cutoff time (we probed with 10 s, 1 min, and 2 min without substantial changes), the graph was not colored, we stopped and tried to color it with larger q . For each connectivity we averaged over 100 samples. As it can be clearly seen, the algorithm fails with q colors slightly below the dynamical transition, confirming our expectations. In Sec. IV we explain how the cavity approach helps to design an algorithm which is able to also deal with this problem.

D. The large- q asymptotic

From Eqs. (37) and (38) one can easily deduce the large- q asymptotics of $\rho(\eta)$. For average connectivities $c \gg q$ [the threshold c_q is expected to scale like $\mathcal{O}(q \ln q)$], f_k is dominated by the $l=0$ contributions in the numerator and in the denominator, leading to $\rho(\eta) = \delta(\eta - 1/q)$ in leading order.




FIG. 6. Average complexity Σ as a function of the energy density e for various average connectivities c . In this figure we only display the *physical* branches (see text).

Plugging this result into Eq. (59) one can easily calculate the COL/UNCOL threshold c_q by setting the complexity to zero. Taking care only of the dominant contribution we find

$$c_q = 2q \ln q + O(q \ln q). \quad (60)$$

This result coincides with the exact asymptotics found by Luczak [14]. Note, however, that the same dominant term can also be obtained from the vanishing of the replica-symmetric (paramagnetic) entropy $s(c)$ which is expected to be exact up to the COL/UNCOL transition. This means that, for $q \rightarrow \infty$, the threshold entropy goes to zero. This behavior could already be conjectured from the above table where the threshold entropies are given for small q . The derivation of the subdominant terms in Eq. (60) requires a much more detailed analysis and goes beyond the scope of this paper. It will be presented in a future publication together with analogous results for K -SAT [37].

IV. WORKING WITH SINGLE-GRAPH INSTANCES: SURVEY PROPAGATION

Up to now we have solved analytically the coloring problem averaged over the set of Erdős-Renyi graphs at given average connectivity. In this way we derived the q -dependent threshold connectivities of c_q at which the graph becomes almost surely uncolorable with q colors, i.e., the location of the COL/UNCOL transition. We have also demonstrated the existence of another threshold value c_d above which a clustering phenomenon takes place in the space of solutions.

However, one of the relevant consequences of this cavity approach is that it can be naturally implemented to study single-case instances, i.e., specific nonrandom graphs which have to have a locally treelike structure to fulfill the conditions of the cavity approach. In the average-case analysis at each step of the iteration, we selected *randomly* k sites from the \mathcal{M} possible ones to be used in Eq. (36), and we substituted another randomly chosen entry η_0 from the \mathcal{M} possible entries. From here on, we will assume that the iteration procedure used above is also valid for single instances—with one significant change: For the generation of survey for one vertex (or edge) we have to use its actual neighbors, the connections between sites are fixed once forever by the specific graph under consideration.

The survey-propagation algorithm

This algorithm works in a way similar to the sum product algorithm [41]. In the latter, to each vertex arrive u messages from $k-1$ neighbors, then this messages are transformed (become h fields) and sent as a new message through the link to the descendant k neighbors. So, at each time step, in the links of the graphs we will have messages traveling, like in a communication network. The survey-propagation (SP) algorithm, works with the same principle. The basic difference is that now the messages are replaced by u surveys of the messages (i.e., by probability distributions of messages). SP is defined for one given value of the reweighting parameter y that must be optimized to minimize the “free energy” of the system. To each edge $\{i, j\}$ of the graph we associate two u

surveys $Q_{i \rightarrow j}(\vec{u})$ and $Q_{j \rightarrow i}(\vec{u})$ of messages traveling in the two possible directions. The algorithm self-consistently determines these surveys by a message passing procedure to be described below, and finds consequently all the thermodynamic properties of the model defined on the specific graph. Let us now describe below how SP works in practice for the 3-coloring problem:

(1) Select a graph $G = (V, E)$.

(2) All the $Q_{i \rightarrow j}(\vec{u})$ with $\{i, j\} \in E$ are randomly initialized.

(3) Sequentially consider all sites i and randomly update the links $\{i, j\}$ to all neighbors j in the following way.

(a) For each neighbor j of i we calculate

$$P_{i|j}(\vec{h}) = C_{i|j} \int \left[\prod_{k \in V(i)/j} d^q \vec{u}_k Q_{k \rightarrow i}(\vec{u}_k) \right] \times \delta \left(\vec{h} - \sum_{k \in V(i)/j} \vec{u}_k \right) \exp \left\{ y \omega \left(\sum_{k \in V(i)/j} \vec{u}_k \right) \right\}, \quad (61)$$

where with the symbol $V(i)$ denotes all neighbors of i . The prefactor $C_{i|j}$ is chosen such that $P_{i|j}$ is properly normalized to one.

(b) From $P_{i|j}(\vec{h})$ we derive the new u surveys of all edges $\{i, j\}$:

$$Q_{i \rightarrow j}(\vec{u}) = \int d^q \vec{h} P_{i|j}(\vec{h}) \delta(\vec{u} - \hat{u}(\vec{h})). \quad (62)$$

(4) The iteration step 3 is repeated until convergence is reached.

It was already shown in Ref. [24] that the free energy of the system may be written as

$$\phi(y) = \frac{1}{N} \left[\sum_{\{i,j\} \in E} \phi_{i,j}^{link}(y) - \sum_i (n_i - 1) \phi_i^{node}(y) \right], \quad (63)$$

where n_i is the connectivity of the vertex i , and $\phi_{i,j}^{link}(y)$ and $\phi_i^{node}(y)$ represent the contributions of links and vertices which are given by

$$\phi_{i,j}^{link}(y) = -\frac{1}{y} \ln \left(\int d^q \vec{h} P_{i|j}(\vec{h}) d^q \vec{u} Q_{j \rightarrow i}(\vec{u}) \times \exp \{ -y [\omega(\vec{h}) - \omega(\vec{h} + \vec{u})] \} \right) \quad (64)$$

and

$$\phi_i^{node}(y) = -\frac{1}{y} \ln \left[\int \prod_{k \in V(i)} d^q \vec{u}_k Q_{k \rightarrow i}(\vec{u}_k) \times \exp \left\{ y \omega \left(\sum_{i=1}^k \vec{u}_i \right) \right\} \right]. \quad (65)$$

Repeating the above procedure for various values of y , Eqs. (64) and (65) do not only provide the values of $\phi(y)$, but also $\Sigma(y) = -y^2 \partial \phi(y) / \partial y$ and the energy density $e(y) = \partial(y \Phi(y)) / \partial y$ of states. The parametric plot of $\Sigma(y)$ vs




FIG. 7. Free energies ϕ as a function of y for three given samples of $N = 10\,000$ of connectivities $c = 4.60, 4.69,$ and 4.80 .

$\epsilon(y)$ gives the complexity of states as a function of their energy. For example, Fig. 7 shows the free energy $\phi(y)$ of single graphs with $N = 10\,000$ vertices as a function of y for three different values of the average connectivity c .

We observe that for high-enough connectivities the maximum of $\phi(y)$ is located at finite values of y . While decreasing c , the location of the maximum grows and approaches $y \rightarrow \infty$ at the coloring threshold. From these curves and by means of numerical derivatives, we may also calculate the complexity and energy. Figure 8 shows the two branches obtained in the parametric plot of $\Sigma(y)$ vs $e(y)$ for various connectivities c . While the physical meaning of the upper branches is not clear [23] we wanted to stress that they interpolate between the RS solution and the maximum complexity point.

From the previous figure we may extract two characteristic values of the energy: The first one is associated with the minimal number $e_0 N$ of miscolored edges in the graph, i.e., it gives the *ground-state energy* of the instance. The value of e_{gs} is determined as the positive point where the lower




FIG. 8. Complexity Σ as a function of ϵ for three given samples of random graph with average connectivities $c = 4.60, 4.69, 4.80$ and $N = 10\,000$ sites. At odds with Fig. 6 here we display both physical and unphysical branches.




FIG. 9. Density of miscolored links e_{gs} vs average connectivity of the graph c (lower dotted curve) and threshold energy density e_{th} vs c (upper continuous curve) in the 1-RSB approximation.

branch of the complexity curve intersects the energy axis or it equals zero if $\Sigma(e=0) > 0$ on the lower branch.

The other relevant energy value is the *threshold energy* e_{th} . It is determined by the point where the complexity reaches its maximum. It is therefore the point where, e.g., simulated annealing gets stuck. The same remark of Sec. III C holds here: this calculation should be probably improved along the line of Ref. [36] in order to take into account the FRSB instability at higher energy density as in the case of the p -spin spherical model.

From the practical side this is, of course, not the way to determine these values, it is much more desirable to look for the value of y at which $\phi(y)$ becomes maximal, cf. Eq. (44). Figure 9 shows a plot of these two energies as a function of the connectivities obtained using this single-instance algorithm.

Of course, the exact meaning of the numerical values of these quantities is an open question. In principle, they were defined for infinite systems, whereas our single-instance algorithm works for systems of finite sizes N . We expect that the numerical values give good approximations once we look at large values of N , where, e.g., the scales dividing distances of solutions inside one state from those between states are well separated. A more detailed discussion about this may be found in Refs. [12,24].

V. A POLYNOMIAL ALGORITHM TO COLOR GRAPHS

The survey propagation described above was very useful for the design of an efficient algorithm to find a solution of randomly generated 3-SAT formulas [22,24] in the hard but satisfiable phase. Here we will show that, with small modifications, the same idea can be extended to the q -coloring problem.

The relevant idea in this algorithm is to fix spins which are strongly biased toward (or away from) one color. Therefore, we have to first determine the distributions of local magnetic fields in the system using SP and select those which have the strongest bias. Once these are fixed the prob-

lem is reduced. We can rerun SP on the reduced instance, new spins may be biased and fixed. The procedure will be repeated until only paramagnetic spins remain. At this point SP cannot help any more, but surprisingly the decimated coloring problem becomes “easy.” Using any reasonable local solver known in the literature, we can proceed to construct a proper coloring.

In the case of q -COL the subject is technically slightly more complex than in K -SAT, since spins can be biased in q different directions and it is hard to decide what do we mean precisely by biased. In addition, by fixing the color of one vertex, all its neighbors have to have different colors, i.e., they are left with $q - 1$ colors. In the reduction process the problem, initially being a pure q -coloring problem, becomes a list coloring problem where each vertex has an own list of allowed colors. In this way the permutation symmetry of colors is broken, which requires a modification of the SP given above to nonsymmetric surveys.

In order to keep the presentation as simple as possible we concentrate our efforts on the 3-coloring problem and hence,

from now on, all the discussion will be associated for the case $q = 3$. The extension of the results to higher q is, however, straightforward although exponential in q .

As mentioned above, the first things we should do are a generalization of SP to non-color-symmetric situations, and to correctly define a biased spin. Let us start first noting that Eq. (26) may be written as

$$\mathcal{Q}[Q(\vec{u})] = \int d^q \vec{\eta} \rho(\vec{\eta}) \delta \left[Q(\vec{u}) - \eta^0 \delta(\vec{u}) - \sum_{\tau=1}^q \eta^\tau \delta(\vec{u} - \vec{e}_\tau) \right], \quad (66)$$

where we simply avoid to consider the color symmetry of the problem, and where we introduce $\eta^0 = (1 - \sum_{\tau=1}^3 \eta^\tau)$. Then, following the same lines of reasoning that lead from Eq. (26) to Eq. (36) we may deduce the following update of the surveys in the limit $y \rightarrow \infty$:

$$\eta_{i \rightarrow j}^r = \frac{\prod_{k \in V(i)/j} (1 - \eta_{k \rightarrow i}^r) - \sum_{p \neq r} \prod_{k \in V(i)/j} (\eta_{k \rightarrow i}^0 + \eta_{k \rightarrow i}^p) + \prod_{k \in V(i)/j} \eta_{k \rightarrow i}^0}{\sum_{p=1,2,3} \prod_{k \in V(i)/j} (1 - \eta_{k \rightarrow i}^p) - \sum_{p=1,2,3} \prod_{k \in V(i)/j} (\eta_{k \rightarrow i}^0 + \eta_{k \rightarrow i}^p) + \prod_{k \in V(i)/j} \eta_{k \rightarrow i}^0} \quad (67)$$

for $r \in \{1,2,3\}$. The value of $\eta_{i \rightarrow j}^0$ can be calculated by imposing the normalization condition. Using this update rule instead of the one proposed in the above version of SP, we directly work with a reweighting parameter $y = \infty$ which forbids any positive energy changes and thus characterizes proper colorings.

Having η_i^τ , for all the sites of the graph, we have to define the site dependent color polarizations

$$\Pi_i^r = \frac{\prod_{j \in V(i)} (1 - \eta_{j \rightarrow i}^r) - \sum_{p \neq r} \prod_{j \in V(i)} (\eta_{j \rightarrow i}^0 + \eta_{j \rightarrow i}^p) + \prod_{j \in V(i)} \eta_{j \rightarrow i}^0}{\sum_{p=1,2,3} \prod_{j \in V(i)} (1 - \eta_{j \rightarrow i}^p) - \sum_{p=1,2,3} \prod_{j \in V(i)} (\eta_{j \rightarrow i}^0 + \eta_{j \rightarrow i}^p) + \prod_{j \in V(i)} \eta_{j \rightarrow i}^0} \quad (68)$$

for $r = 1,2,3$. This equation is analogous to Eq. (67) but the products are extended to all neighbors. The polarization Π_i^r is the probability that vertex i is fixed to color r in a randomly selected cluster of solutions. Vertices that may change their color within one cluster are characterized by $\Pi_i^0 = (1 - \sum_{r=1}^3 \Pi_i^r)$. Once these polarizations are known, many strategies can be adopted for coloring the graph. We believe that the simplest and most intuitive one is the following

- (i) If one spin is very biased to one color, fix that spin and remove it from the graph. Forbid this color to all neighbors.
- (ii) If the bias of one spin toward some color is very low, forbid that color.

Forbidding color c to node i implies rewriting Eq. (67) using only two colors for that particular node. This can be achieved simply by taking Eqs. (67) and (68) but setting $\eta_{i \rightarrow k}^c = 0$ and $\eta_{k \rightarrow i}^c = 1$ for all $k \in V(i)$.

Furthermore, during the processes discussed above, it

turned out that many vertices get surrounded by neighbors with fixed colors. In that case, the spin can be fixed to one of their remaining allowed colors immediately, and it is removed from the graph.

In practice, we put a cutoff for the value of the bias to be used for the previous criteria. We use rule (i) every time a bias towards some color is greater than 0.8 and rule (ii) if the bias was lower than 0.15. There is no special reason for selecting specifically these values, but we found numerically a fast convergence to solvable paramagnetic problem instances. It could be useful to make a systematic analysis for improving this choice, and also to discuss other selection rules. However, this is not the objective of the present work. Here we just want to demonstrate that the algorithm works substantially better than every other local search algorithm we know, even without any parameter optimization.

Summarizing the discussion above, our algorithm follows the following steps:




FIG. 10. Probability of coloring a graph using our algorithm for different lattice sizes. From left to right: $N=4 \times 10^3$, 8×10^3 , 16×10^3 , 32×10^3 , and 64×10^3 .

- (1) Take the original graph and run SP in its infinite-version defined by Eq. (67).
- (2) Calculate the biases of all spins according to Eq. (68).
- (3) Select spins whose bias to one color is larger than 0.8, and fix and remove these spins from the graph. Forbid the color to all neighbors.
- (4) Select spins whose bias to one color is lower than 0.15 and forbid that color to these spins.
- (5) Take all spins where just one color is allowed, fix these spins, and remove them from the graph. Forbid the fixed color to all neighbors.
- (6) If the the graph is not completely paramagnetic: rerun SP and go to step (2).
- (7) Run any smart program that solves the coloring sub-problem.

Actually, we did not find any free program in the web which was able to easily handle large graphs for the coloring problem. The best we could find was the SMALLK program by Culberson [38], but even in the easy region it exploded in memory for graphs with sizes larger than $N=2000$. So step (7) above was changed into the following: (a) Transform the resulting graph into a satisfiability problem; (b) Run walkSAT [6] on this satisfiability problem.

An interesting point about the algorithm described above is the fact that we can fix a certain percentage of spins in every algorithmic step, without rerunning SP every time. This drastically reduces the computational time. How many spins we may fix depends in a nontrivial way on the system size and on the distance from the COL/UNCOL transition.

Figure 10 shows the success rate of our algorithm in 3-coloring random graphs in the hard region $c \in [4.42, 4.69]$. From left to right the sample sizes increase: $N=4 \times 10^3$, 8×10^3 , 16×10^3 , 32×10^3 , and 64×10^3 . In all the cases we fixed the 0.5% of the spins in every iteration step. Note that keeping this value fixed we find a clear improvement of the algorithm for sizes going from $N=4 \times 10^3$, 8×10^3 to $N=16 \times 10^3$ the performance is roughly the same for larger lattices suggesting that we should reduce the fraction of spins to fix. However, note that even within these conditions the algorithm works quite well in the hard

region of the system. Note that strong finite size effects are present, in fact the algorithm does not behave very well for small lattice sizes. Two reasons may explain this: First, there are short loops that disappear in the thermodynamic limit and second, there is some shift in the location of the COL/UNCOL transition towards higher connectivities for larger graphs. This point should be clarified in a forthcoming work.

Another relevant feature of the curve is the following: The closer our graph is to the critical point, the smaller is also the percentage of spins we may fix in one algorithmic step. However, extrapolating the results, the worst we can find is to fix only one single spin at the time. This would change the complexity of our algorithm from $N \ln N$ (resulting from sorting spins with respect to their biases) to N^2 , i.e., the algorithm remains polynomial.

VI. CONCLUSIONS AND OUTLOOK

In this work we presented a detailed derivation of the one-step replica-symmetry broken solution of the coloring problem on random graphs. The problem consists in finding a coloring of all vertices of the graph such that no two adjacent vertices carry equal colors. From the average case point of view, the one-step RSB approach allowed to determine the q -COL/UNCOL transition c_q for arbitrary color numbers q . This means that large random graphs of average connectivity below c_q have proper q colorings with high probability (approaching one for $N \rightarrow \infty$), whereas graphs with higher connectivity require more colors for a proper coloring. Moreover, we find the existence of a clustering transition in the colorable region. This transition is characterized by the appearance of an exponential number of states separated by large energetic barriers. The clustering transition is accompanied by the sudden appearance of an exponential number of metastable states that, intuitively, cause local the algorithm to get stuck.

We also extended our results to the study of single-case instances, i.e., specific realizations of random graphs, showing that the previous analysis remains valid. With this understanding we also implemented a different algorithm, based on the idea of a survey propagation that enabled us to solve the coloring problem in the hard clustering region in polynomial time. We present results for sizes as large as $N \approx 10^5$ vertices, which is far beyond the performance of other algorithms on random graphs.

There are many interesting directions in which we can extend this work. The first one concerns the survey-propagation algorithm. We were able to report quite encouraging results for the SP inspired graph reduction procedure if applied to the clustered, i.e., hard but colorable phase on random graphs. These graphs are characterized by a local treelike structure, loops are of length $O(\ln N)$. This structure allowed us to use the statistical independence of surveys restricting a randomly selected vertex inside each pure state. This assumption fails, however, if the graph has some non-trivial local structure as given by small loops, small highly connected subgraphs, etc. Before being of real practical value, SP should be extended to such situations, following, e.g., the lines used by Yedidia *et al.* in Ref. [42] in their generalization of belief propagation to local nontreelike graphs.

The second possible extension of our work concerns the interpretation of colorings as ground states of a Potts antiferromagnet, which is a model known to show glassy behavior at low temperatures (the so-called Potts glass), see, e.g., Ref. [11]. In the present work we have directly worked at zero temperature, but the extension to nonzero temperature is straightforward. In this context it is interesting to see that for $q=3$ a continuous full replica-symmetry breaking transition appears at the level of fields of $O(T)$ —before the one-step solution appears for fields of $O(1)$. So we expect that the one-step RSB transition in this model exists in a strict sense only at zero temperature, in temperature it is only a (sharp) crossover to glasslike behavior. This phenomenon disappears for larger q , but it is interesting in how far it can influence

the usual glassy phenomenology known from fully connected spin-glass models. Let us also point out that interestingly enough a similar scenario holds also in the random K -SAT [43] case. Using in addition the approach suitable for single-graph instances, one can, e.g., study inhomogeneities arising in the glassy phase [44] and thus go beyond the usual paradigm of disorder averaged results for randomly disordered models.

ACKNOWLEDGMENTS

We are grateful to A. Montanari, J. Culberson, B. Hayes, and F. Ricci-Tersenghi for their interest and many helpful discussions.

-
- [1] M.R. Garey and D.S. Johnson, *Computers and Intractability* (Freeman, New York, 1979).
- [2] F.Y. Wu, *Rev. Mod. Phys.* **54**, 235 (1982).
- [3] K. Appel and W. Haken, *Ill. J. Math.* **21**, 421 (1977); **21**, 491 (1977).
- [4] *Complexity of Computer Computation*, edited by R.E. Miller and J.W. Thatcher (Plenum Press, New York, 1972), p. 85.
- [5] S. Mertens, *Comput. Sci. Eng.* **4**, 31 (2002).
- [6] For a comprehensive collection of benchmark and tests for SAT solvers, see <http://www.satlib.org>
- [7] G.J. Chaitin, M.A. Auslander, A.K. Chandra, J. Cocke, M.E. Hopkins, and P. Markstein, *Comput. Speech Lang.* **6**, 47 (1981).
- [8] P. Erdős and A. Rényi, *Publ. Math. (Debrecen)* **6**, 290 (1959).
- [9] *Artif. Intell.* **81**, 1-2 (1996).
- [10] *Theor. Comput. Sci.* **265**, 1-2 (2001).
- [11] K. Binder and A.P. Young, *Rev. Mod. Phys.* **58**, 801 (1986); M. Mézard, G. Parisi, and M.A. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore 1987); K.H. Fisher and J.A. Hertz, *Spin Glasses* (Cambridge University Press, Cambridge, UK, 1991).
- [12] M. Mézard and G. Parisi, *Eur. Phys. J. B* **20**, 217 (2001).
- [13] R. Mulet, A. Pagnani, M. Weigt, and R. Zecchina, *Phys. Rev. Lett.* **89**, 268701 (2002).
- [14] T. Luczak, *Combinatorica* **11**, 45 (1991).
- [15] I. Kanter and H. Sompolinsky, *J. Phys. A* **20**, L673 (1987).
- [16] B. Pittel, J. Spencer, and N. Wormald, *J. Comb. Theory, Ser. B* **67**, 111 (1996).
- [17] M. Molloy, *Random Struct. Algorithms* **7**, 159 (1996).
- [18] D. Achlioptas and C. Moore, *Proceedings of the Symposium on Theory of Computing (STOC) 2002* (unpublished).
- [19] J. Culberson and Ian P. Gent, *Theor. Comput. Sci.* **265**, 227 (1991).
- [20] A.C. Kaporis, L.M. Kirousis, and Y.C. Stamatiou, *Electr. J. Comb.* **7**, R29 (2000).
- [21] J. van Mourik and D. Saad, *Phys. Rev. E* **66**, 056120 (2002).
- [22] M. Mézard, G. Parisi, and R. Zecchina, *Science* **297**, 812 (2002).
- [23] M. Mézard and G. Parisi, *J. Stat. Phys.* **111**, 1 (2003).
- [24] M. Mézard and R. Zecchina, *Phys. Rev. E* **66**, 056126 (2002).
- [25] F. Guerra and F.L. Toninelli, e-print cond-mat/0208579.
- [26] S. Franz and M. Leone, e-print cond-mat/0208280.
- [27] M. Talagrand, *Spin Glasses: A Challenge for Mathematicians. Mean-field Models and Cavity Method* (Springer-Verlag, Berlin, in press).
- [28] E. Marinari, G. Parisi, F. Ricci-Tersenghi, J. Ruiz-Lorenzo, and F. Zuliani, *J. Stat. Phys.* **98**, 973 (2000).
- [29] The educated reader will recognize that this construction is actually a rephrasing of the *real replica* argument due to Monasson [40] in which we consider the limiting case $\beta \rightarrow \infty$, $m \rightarrow 0$ such that $y = m\beta$ is finite.
- [30] F. Ricci-Tersenghi, M. Weigt, and R. Zecchina, *Phys. Rev. E* **63**, 026702 (2001).
- [31] S. Cocco, O. Dubois, J. Mandler, and R. Monasson, e-print cond-mat/0206239.
- [32] M. Mézard, F. Ricci-Tersenghi, and R. Zecchina, e-print cond-mat/0207140.
- [33] This formula is correct only in the case $O(T)$ field distributions are paramagnetic (see also Sec. VI). In the interval $c_d < c < c_q$ this is always true for $q > 3$. Values of $s(c_d)$ and $s(c_c)$ displayed in the table of Sec. I for $q=3$ are therefore to be considered as a lower bound to the true value.
- [34] R. Monasson and R. Zecchina, *Phys. Rev. Lett.* **75**, 2432 (1995).
- [35] M. Weigt and A.K. Hartmann, *Phys. Rev. Lett.* **84**, 6118 (2000).
- [36] A. Montanari and F. Ricci-Tersenghi, e-print cond-mat/0301591.
- [37] A. Braunstein, M. Mézard, M. Weigt, and R. Zecchina (unpublished).
- [38] Graph Coloring Page, <http://www.cs.ualberta.ca/~joe>
- [39] R. Monasson and R. Zecchina, *Phys. Rev. E* **56**, 1357 (1997).
- [40] R. Monasson, *Phys. Rev. Lett.* **75**, 2847 (1995).
- [41] F.R. Kschischang, B.J. Frey, and H.-A. Loeliger, *IEEE Trans. Inf. Theory* **47**, 498 (2002).
- [42] J.S. Yedidia, W.F. Freeman, and Y. Weiss, Mitsubishi Electrical Research Laboratories, Report TR-2002-35 (unpublished), available online at <http://www.merl.com>
- [43] S. Mertens, M. Mezard, and R. Zecchina (unpublished).
- [44] A. Montanari and F. Ricci-Tersenghi, *Phys. Rev. Lett.* **90**, 017203 (2003).