Random *K*-satisfiability problem: From an analytic solution to an efficient algorithm

Marc Mézard[1] and Riccardo Zecchina[2]
[1]*Laboratoire de Physique Théorique et Modèles Statistiques, CNRS and Université Paris Sud, Bâtiment 100, 91405 Orsay Cedex, France*
[2]*The Abdus Salam International Centre for Theoretical Physics, Statistical Mechanics and Interdisciplinary Applications Group, Strada Costiera 11, 34100 Trieste, Italy*
(Received 9 July 2002; published 26 November 2002)

We study the problem of satisfiability of randomly chosen clauses, each with K Boolean variables. Using the cavity method at zero temperature, we find the phase diagram for the $K=3$ case. We show the existence of an intermediate phase in the satisfiable region, where the proliferation of metastable states is at the origin of the slowdown of search algorithms. The fundamental order parameter introduced in the cavity method, which consists of surveys of local magnetic fields in the various possible states of the system, can be computed for one given sample. These surveys can be used to invent new types of algorithms for solving hard combinatorial optimizations problems. One such algorithm is shown here for the $K=3$ satisfiability problem, with very good performances.

DOI: 10.1103/PhysRevE.66.056126 PACS number(s): 02.50.−r, 75.10.Nr, 89.20.−a, 89.70.+c

I. INTRODUCTION

The K-satisfiability (K-SAT) problem deals with an ensemble of N Boolean variables, submitted to M constraints. Each constraint is in the form of an OR function of K variables in the ensemble (or their negations), and the problem is to know whether there exists one configuration of the variables (among the 2^N possible ones) which satisfies all constraints. The K-SAT problem for $K \geqslant 3$ is a central problem in combinatorial optimization: it was the first problem to be shown NP-complete [1,2], and an efficient algorithm for solving the K-SAT problem in its worst-case instances would immediately lead to other algorithms for solving efficiently thousands of different hard combinatorial problems.

At the core of the statistical physics of disordered systems is the spin-glass problem (SG), which also deals with Boolean variables (spins), interacting with random exchange couplings [3]. Each pair of interacting spins can be seen as a constraint, and finding the state of minimal energy in a spin-glass amounts to minimizing the number of violated constraints. Although the precise form of the constraints in SG and K-SAT differ, there exist deep similarities [4,5]; in both cases the difficulty comes from the existence of ''frustration'' [3], which forbids to find the global optimal state by a purely local optimization procedure. Links between combinatorial optimization and statistical physics have been known for long [3]. Two main categories of questions can be addressed. One type is algorithmic, for instance finding an algorithm which decides whether an instance is satisfiable or not. Another is more theoretical, and deals with large random instances, for which one wants to predict the typical behavior. Examples of use of statistical physics in each category are the simulated annealing algorithm [6] and the solution of the random assignment problem [7,8], or the direct mapping of certain graph partitioning problems to spin glasses [9]. Here we address the two types of questions in the $K=3$ satisfiability (3-SAT) problem.

The study of random K-SAT problems, where the clauses are chosen randomly, is also interesting from the viewpoint of optimization. In practice, algorithms that are used to solve real-world NP-complete problems display a huge variability of running times, ranging from linear to exponential, when the parameters (e.g., the number of clauses) are changed. A theory for the typical-case behavior of algorithms, on classes of random instances chosen from a given probability distribution, is therefore the natural complement to the worst-case analysis [10–15].

In random 3-SAT, numerical simulations have shown the existence of a phase transition when one varies the ratio $\alpha = M/N$ of the number of clauses to the number of variables. For $\alpha < \alpha_c$ the generic problem is satisfiable (SAT), for $\alpha > \alpha_c$ the generic problem is not satisfiable (UNSAT) [16]. Using the cavity method, first developed in spin-glass theory, we shall show the existence of this threshold and compute $\alpha_c \simeq 4.267$. We also find an intermediate region, the hard-SAT phase $\alpha_d = 3.921 < \alpha < \alpha_c$, where the generic problem is still SAT but the proliferation of metastable states makes it difficult for algorithms to find a solution. This proliferation is similar to the effect found in the theories of structural glasses using spin-glass models with multispin interactions [17], where it is known to lead to a dramatic slowdown of the relaxation. In this sense the difficulty to solve the 3-SAT problem in the intermediate region $\alpha_d < \alpha < \alpha_c$ is similar to the difficulty in equilibrating structural glasses.

This theoretical analysis is done using the cavity method, at a level equivalent to what is called one-step replica symmetry breaking in the replica language. This means that it assumes the existence of many states, but cannot handle a nontrivial correlation pattern between them. There are some arguments which point towards the correctness of this solution, although an exact proof looks somewhat remote at present.

In this cavity method with many states, the order parameter consists in the surveys of local magnetic fields acting on each spin. While for the theoretical analysis one averages over the random graph structure of the problem, it turns out that this order parameter can also be computed for one given sample, using a reasonably simple message passing proce-

dure which takes into account the multiplicity of states. This procedure provides a generalization of the belief propagation used in statistical inference [18]; it is shown here to converge in some difficult situation with many equilibrium states where ordinary belief propagation does not converge. The resulting surveys provide an interesting description of one given sample, where the various variables are found to play very different roles. This single sample analysis is very useful in order to find new algorithms for solving hard optimization problems. Here we show one such algorithm for the random 3-SAT problem, where the surveys are used to identify one spin and fix it. The problem is thus reduced and the surveys are computed again on the new system. This decimation procedure is shown to have very good performances, comparable to, or better than, the state of the art in this problem.

The paper presents a number of concepts and techniques, both analytical and numerical, which can be applied to a rather large class of combinatorial optimization problems. We have presented these concepts and techniques in a broad framework, in order to allow for future use on different problems. The concrete implementation is then done on the random 3-SAT problem. Some of the results discussed in this paper have been recently announced in Ref. [19].

The paper is organized as follows: in Sec. II we present the generic structure of the optimization problems in which we are interested. These can be represented as bipartite graphs called factor graphs. Section III recalls a general message passing procedure which can be used to study optimization or inference problems defined on these factor graphs. The basic ingredients of this procedure are messages which we call cavity biases which play a crucial role in the whole paper. Section IV defines the set of random graphs which we study, which are random hypergraphs with a fixed connectivity. Section V provides some background on the decomposition of the configuration space into states. It consists of a short introduction for nonphysicists, a specific definition of zero-temperature states for the random 3-SAT problem, and the definition of the *complexity* which is a crucial concept for a system with many states. Section VI provides an introduction to the zero-temperature cavity method, presented in the general setting of combinatorial problems on random factor graphs. This section summarizes, and puts in a more general context, some recent work which has developed the cavity method for finite connectivity problems, first at finite temperature [20], then at zero temperature [21]. It provides the whole formalism for the analytic study of the phase diagram. This formalism is applied to the random 3-SAT problem in Sec. VII, where all the results on the phase diagram are derived. We explain the survey propagation algorithm on a given sample in Sec. VIII, and the decimation algorithm for solving large random 3-SAT problems is presented in IX. Section X contains some concluding remarks. The Appendix A contains some technical details relative to the computation of the phase diagram done in Sec. VII. The Appendix B explains the computation of the free energy for one given sample.

II. FACTOR GRAPH REPRESENTATION

The models we are interested in involve Boolean variables which interact in groups, the energy being the sum of energies of all groups. We shall adopt the factor graph representation [22] familiar in computer science, but we shall keep to the representation of Boolean variables as Ising spins, more familiar to statistical physicists.

We consider a set of N Ising spins $\sigma_i \in \{\pm 1\}$ and we suppose that we have M groups of interacting variables, which are called function nodes. Each function node a involves a set of n_a spins. We denote by V_a the set of all these spins. The interaction is an arbitrary function of the spins in V_a, which depends on the problem one considers, and can also involve hidden variables.

The total energy of a configuration $\sigma_1, \ldots, \sigma_N$ is

$$E = \sum_{a=1}^{M} E_a, \tag{1}$$

and the goal in combinatorial optimization is to find a configuration of spins which minimizes E. A generalization of this problem, natural from the point of view of physics, and which connects with problems in statistical inference, consists in introducing an additional parameter β, an ''inverse temperature'' in the physics language, and in studying the Boltzmann probability distribution

$$P(\sigma_1, \ldots, \sigma_N) = \frac{1}{Z} \exp(-\beta E), \tag{2}$$

where Z is a normalization constant. As usual in physics, we shall denote by $\langle O \rangle$ the expectation of an observable O (which can be any function of the σ_i) with respect to this measure. In the large β (low-temperature) limit this measure concentrates onto the lowest energy configurations. At finite β one may be interested in computing for instance the expectation value of one spin variable σ_1 with the Boltzmann probability. Because we work with binary spins, this average determines the full marginal probability law of σ_1: this is precisely the quantity that one typically seeks in many inference problems, such as e.g., decoding procedures for error correcting codes.

The general problem can be represented by a graph consisting of two types of vertices, ''variable nodes'' associated with each spin, and ''function nodes.'' A function node a is connected by edges to all the variable nodes involved in S_a. Therefore, each variable node has connections towards all the function nodes in which it appears, and the graph is bipartite (see Fig. 1). Each spin σ_i is connected to n_i function nodes, we denote this set of function nodes by V_i. We call n_i the connectivity of spin i, n_a the connectivity of function node a. Throughout this paper, the variable nodes indices are taken in $i,j,k, \ldots$, while the function nodes indices are taken in a,b,c, etc.

Let us give here a few standard examples.

Spin glasses. All the interactions involve two spins, so all n_a are equal to 2; the energy of an interaction node a involving spins σ_i and σ_j is given by

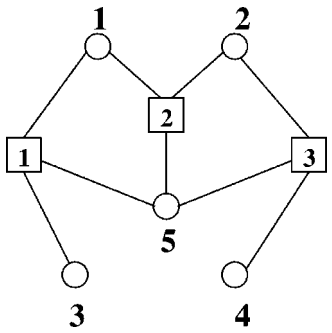

FIG. 1. An example of a factor graph with five variable nodes $i=1,\ldots,5$ and three function nodes $a=1,2,3$. In this case, each function node has connectivity 3, as in the 3-SAT problem. The connectivities of the five variable nodes are, respectively, 2,2,1,1,3.

$$E_a = -J_{ij}\sigma_i\sigma_j, \tag{3}$$

where the number J_{ij} is called the coupling constant. Generalized spin glasses with *p*-spin interactions,

$$E_a = -J_{i_1\cdots i_p}\sigma_{i_1}\cdots\sigma_{i_p}, \tag{4}$$

have also been studied a lot in statistical physics as models of structural glasses. They are the closest physical analogues of the satisfiability problems which we study here [23].

K-SAT: All interactions involve *K* spins, and the energy of an interaction node a involving spins $\sigma_{i_1},\ldots,\sigma_{i_K}$ is given by

$$E_a = 2\prod_{r=1}^{K}\frac{(1+J_a^r\sigma_{i_r})}{2}. \tag{5}$$

It depends on a set of K coupling constants $\mathbf{J_a}=(J_a^1,\ldots,J_a^K)$ which take values ± 1. This interaction node has a simple interpretation as a clause: the energy E_a is zero as soon as at least one of the spins σ_{i_r} is opposite to the corresponding coupling J_a^r. If all spins are equal to their couplings, the energy is equal to 2. The more conventional description of *K*-SAT uses Boolean variables: let us introduce x_i which is TRUE if and only if $\sigma_i=1$. The energy E_a depends on the OR of the K variables $y_{i_1},\ldots,y_{i_K}$, where y_{i_r} is the original x_{i_r} when $J_a^r=-1$ and is its negation when $J_a^r=+1$. The energy vanishes if $y_{i_1}\vee\cdots\vee y_{i_K}$ is true (the clause is then said to be satisfied), otherwise it is equal to 2 and the clause is unsatisfied. This arbitrary factor of 2 is introduced for future convenience.

One can also consider graphs involving mixtures of function nodes of different types, e.g., mixtures of $K=2$ clauses and $K=3$ clauses [5,24]. These are some examples of constraints satisfaction problems, but of course there exist many other instances of problems, much studied in computer science, which can be represented by such factor graphs.

In general, an instance of the problem (also called a sample in physics language) is given by a graph and the set of couplings needed to define each function node. In physics (e.g., spin glasses) one is interested in the configurations of low energy. In the SAT problem, one wants to know whether there exists a configuration of zero energy (in which case the instance is called SAT), or not (in which case the instance is UNSAT).

III. THE SUM-PRODUCT ALGORITHM

A popular method for studying the inference, i.e., the probability measure (2), is a message passing procedure called the ''sum-product'' algorithm [22,18]. When used at $\beta\to\infty$, the corresponding ''min-sum'' algorithm can also be used to get some information on the lowest energy configurations. This procedure is exact and fast on treelike graphs. In our case the sum-product algorithm amounts to sending some messages along the edges of the graph. We call cavity-field, and denote by $h_{i\to a}$, the message passed from a variable node i to a function node a. We call cavity-bias, and denote by $u_{a\to i}$, the message passed from a function node a to a variable node i.

The cavity-field $h_{i\to a}$ is given by the sum of cavity-biases converging to i from all function nodes b distinct from a:

$$h_{i\to a} = \sum_{b\in V(i)/a} u_{b\to i}. \tag{6}$$

The operation performed by a function node to compute the cavity-biases which it will send to its neighboring variable nodes is a partial summation: it computes the marginal probability law for that variable to which it sends the message. More precisely, let us consider a function node a of connectivity K and let us suppose, for notational simplicity, that it is connected to variables $\sigma_1,\ldots,\sigma_K$. The cavity-bias $u_{a\to 1}$ sent from the function node a to the variable node $i=1$ is a function of the cavity fields $h_{j\to a}$ sent from all other variables nodes $j\in\{2,\ldots,K\}$ towards node a. One considers the function of σ_1 defined by $\Sigma_{\sigma_2,\ldots,\sigma_K}\exp[-\beta E_a(\sigma_1,\sigma_2,\ldots,\sigma_K)+\beta(h_{2\to a}\sigma_2+\cdots+h_{K\to a}\sigma_K)]$. As $\sigma_1=\pm 1$, this function can be written for instance as the exponential of a linear form in σ_1. The cavity-bias $u_{a\to 1}$ sent from a to 1 is defined from

$$\begin{aligned}\sum_{\sigma_2,\ldots,\sigma_K}&\exp[-\beta E_a(\sigma_1,\sigma_2,\ldots,\sigma_K)\\ &+\beta(h_{2\to a}\sigma_2+\cdots+h_{K\to a}\sigma_K)]\\ &=\exp[\beta(w_{a\to 1}+\sigma_1 u_{a\to 1})].\end{aligned} \tag{7}$$

Besides the cavity-bias $u_{a\to 1}$, this equation also defines a ''free-energy shift'' $w_{a\to 1}$ which is not used in the sum-product algorithm but will become very important in our generalization later on. In physics words, $h_{i\to a}$ is the magnetic field on spin number i whenever the interaction a is turned off, and $u_{a\to i}$ is the contribution to the magnetic field on spin number i from the interaction a. Equation (6) indicates that the probability law of spin σ_i due to the interactions $a=2,\ldots,K$ is a product of independent laws due to each interaction, while the marginalization operation (7) is a

partial summation, hence the name sum-product. The algorithm is easily generalized to variables which are more complicated than Boolean.

The iteration of the above message passing algorithm, starting from a generic random initial condition, is known to converge whenever the underlying factor graph is a tree. Actually it converges in one sweep if one first computes the messages from the leaves of the tree. The resulting set of messages can be used to compute the probability distribution of one spin (or more generally of some subset of spins). One just needs to compute the local field H_i on spin σ_i:

$$H_i = \sum_{a \in V(i)} u_{a \to i}, \tag{8}$$

and the probability distribution of σ_i is

$$P(\sigma_i) = \frac{\exp(\beta H_i \sigma_i)}{2\cosh(\beta H_i)}. \tag{9}$$

One way to prove this result, using the physics language, is to show that the message passing algorithm minimizes the Bethe free energy of the spin system [18]. As the Bethe free energy is exact for treelike graphs, this provides the proof. If one is interested in the optimization problem ($\beta \to \infty$), one can show that the configuration $\sigma_i = \mathrm{sgn}\, H_i$ is the lowest energy configuration if there is a unique such configuration. If there are several lowest energy configurations, taking the $\beta \to \infty$ limit of $\langle \sigma_i \rangle$ with the measure (9) gives the average of the spin σ_i over all these configurations (but one needs to make this detour through the finite β problem in order to get this result).

In this work we shall mainly be interested in the optimization problem. Working directly at $\beta = \infty$ then simplifies the algorithm. The cavity-fields are generated as before by sums of cavity biases. For computing a cavity-bias, one performs a partial minimization, and the formula (7) simplifies to

$$\min_{\sigma_2, \ldots, \sigma_k} [E_a(\sigma_1, \sigma_2, \ldots, \sigma_k) - (h_{2 \to a}\sigma_2 + \cdots + h_{k \to a}\sigma_k)] = -(w_{a \to 1} + \sigma_1 u_{a \to 1}). \tag{10}$$

This equation defines the output messages w and u as functions of the input messages $h_{j \to a}$. In general, we shall write

$$w_{a \to 1} = \hat{w}_{\mathbf{J_a}}(h_{2 \to a}, \ldots, h_{K \to a}),$$

$$u_{a \to 1} = \hat{u}_{\mathbf{J_a}}(h_{2 \to a}, \ldots, h_{K \to a}), \tag{11}$$

which defines the functions $\hat{w}_{\mathbf{J_a}}$ and $\hat{u}_{\mathbf{J_a}}$ [the label $\mathbf{J_a}$ is here to explicitly remind that a given function node energy E_a will in general depend on some set of couplings—see Eqs. (4) and (5)—which we denote collectively as $\mathbf{J_a}$].

IV. RANDOM GRAPHS AND THERMODYNAMIC LIMIT

A. Definition of random graphs

In the rest of this paper we shall consider the random K-SAT problem, which is defined on some ensemble of random graphs which we now describe. To lighten the notation, we concentrate on the $K=3$ cases for which all the function nodes have connectivity $n_a = 3$, and we generate the random graphs as follows: For each triplet $i<j<k$ of variable nodes, a function node connecting them is present with a probability $6\alpha/N^2$, and it is absent with probability $1-6\alpha/N^2$. The average number of function nodes is then $M=\alpha N$. The graph model used is analogous to the $G(N,p)$ model of random graph theory [see, e.g., Ref. [25]], with $p = 6\alpha/N^2$.

For the problem which we consider, the energy E_a associated with a function node a also depends on some coupling constants $\mathbf{J_a}$ [see for instance Eq. (5)], which may be drawn randomly and independently for each function node from some probability distribution. For instance in random 3-SAT, each number J_a^1, J_a^2, J_a^3 takes values ± 1 with probability 1/2. In general we shall denote by $\mathcal{E}_J O$ the average of any quantity O over all the choices of the random graphs (with fixed N and M), and over the choices of couplings. In such a probabilistic setting, one is interested for instance in computing the average "ground state energy:" For each sample, an optimal configuration (one with minimal energy), is called the ground state, its energy is E_0, and one would like to compute $\mathcal{E}_J E_0$.

B. Thermodynamic limit

We shall be interested in the "*thermodynamic limit*" where M and N go to ∞, keeping the ratio $\alpha = M/N$ fixed. The connectivities of variable nodes become independent identically distributed (iid) random variables with a Poisson distribution $f_{3\alpha}(k)$ of mean 3α, since the probability of having k edges connected to a variable node is

$$\lim_{N \to \infty} \binom{(N-1)(N-2)/2}{k} (6\alpha/N^2)^k \times (1-6\alpha/N^2)^{(N-1)(N-2)/2-k} = \frac{(3\alpha)^k}{k!}\exp(-3\alpha) \equiv f_{3\alpha}(k). \tag{12}$$

The structure of the random graphs generated by this process for large N is interesting. Locally such a graph is treelike: the typical size of a loop in the graph scales like $\ln(N)$ for large N. On the other hand loops are definitely present, and they can induce frustration in the sense of having competing constraints [it has been argued that similar random graphs with a local treelike structure provide a natural setting for discussing the "Bethe approximation" of frustrated systems [26]]. The structure of the graph has one important consequence: consider one given function node, connected to three variable nodes (spins). If one deletes this function node, the typical distance between any two of these three spins (measured as the length of the shortest path on the

graph which connects them) is of order $\ln N$, and thus diverges in the thermodynamic limit: the spins are far apart. This property will be crucial in understanding the type of correlations existing between the spins, and in solving the model. Notice that the limit where $M,N\to\infty$ is also the one that is interesting from a computational complexity point of view.

For problems defined on random graphs with given N,M, the ground state energy fluctuates from sample to sample. It is often true, but it may be difficult to show, that the distribution of the ground state "energy density" E_0/N becomes more and more peaked when N increases, so that, in the thermodynamic limit, almost all samples have the same energy density, which can be computed as

$$\epsilon_0 = \lim_{N\to\infty} \mathcal{E}_J E_0/N. \tag{13}$$

For the random K-SAT problem it can be proved that the above condition holds [27].

One of our aims is to compute this limiting value ϵ_0 for a fixed value of $\alpha=M/N$. For the K-SAT problem it turns out that this value is equal to zero below a certain threshold α_c, and becomes >0 for $\alpha>\alpha_c$. In statistical physics it is very difficult to go beyond the estimate of the energy density: if one can compute ϵ_0, one knows that $E_0\sim\epsilon_0 N$ is the leading behavior of the energy for large N, but in general one cannot control the subleading part, and in principle it could be possible for instance that also in the small α phase of K-SAT where $\epsilon_0=0$, some finite contribution to E_0 (finite when $N\to\infty$) could make the problem typically UNSAT. Numerical simulations tend to show that this is not the case. Knowing ϵ_0 will then allow to get the phase diagram of the problem. But one is also interested in other properties of the generic samples in the thermodynamic limit, like the decomposition of the space of accessible configurations at a given energy, to which we now turn.

V. STATES AND CLUSTERING PROPERTY

A. A simple example of pure states: The ferromagnet

One of the main aims of statistical physics is to understand the building up of correlations between distant variables, when the basic interactions between them are short range. This is precisely the type of question that we need to address here: variables interact locally (the only direct interactions involve spins connected to the same function node). But we also need to control the correlations established between two spins belonging to the same function node, due to their indirect coupling through other nodes. As we saw, the geometry is such that this indirect interaction builds up through very long $(O(\ln N))$ paths.

Usually, in the statistical physics of systems with short range interactions, the correlation between distant variables displays a relatively simple variety of behaviors. The simplest one is when there is only one pure state in the system (typically a "paramagnetic phase"): then there exists a finite correlation length and the connected correlation function between two distant spins σ_i,σ_j decays exponentially with the distance d_{ij} at large distances

$$|\langle\sigma_i\sigma_j\rangle - \langle\sigma_i\rangle\langle\sigma_j\rangle| \simeq C\exp(-d_{ij}/\xi), \tag{14}$$

where ξ is the correlation length (C can be a constant, or involve power law corrections in the distance). This is called the clustering property. On the other hand, some systems can also have phase transitions, and display a low-temperatue phase with several pure state.

The archetypical case which we briefly describe here as a pedagogical example is the ferromagnetic $p=2$ spin system with energy given by Eq. (4) with $J_{ij}=1$: at low temperature the spins polarize in one of two pure states, related to each other by the global symmetry changing all σ_i to $-\sigma_i$. Let us call a configuration one assignment of the N spins, $\sigma_1,\ldots,\sigma_N$. The pure states are probability measures on the configuration space obtained using a slightly modified Boltzmann measure where one adds an external "symmetry breaking" magnetic field (in the present language one adds a function node of connectivity one connected to each variable node i, with energy $-B\sigma_i$). One computes $\lim_{B\to 0^\pm}\lim_{N\to\infty}\langle\sigma_i\rangle$, which defines the expectation value $\langle\sigma_i\rangle_\pm$ of spin i in each of the two states $+$ and $-$. It is a well known fact that the connected correlation functions within each state have the clustering property

$$|\langle\sigma_i\sigma_j\rangle_\pm - \langle\sigma_i\rangle_\pm\langle\sigma_j\rangle_\pm| \simeq C\exp(-d_{ij}/\xi). \tag{15}$$

This means that when the spins collectively polarize in the $+$ state, the correlations between distant spins vanishes. It is not true for the full Boltzmann measure: if one does not add the small symmetry breaking field B, but keeps to the Boltzmann measure (2), one gets for any observable $\langle O\rangle = 1/2(\langle O\rangle_+ + \langle O\rangle_-)$ (the fact that the two states enter with equal weight 1/2 is a consequence of the global symmetry of the original problem), and one easily shows that the corresponding correlations do not vanish at large distances.

The lesson we learn from statistical physics is that correlations decay at large distance within each pure state. In problems more complicated than the ferromagnet it may be difficult to identify the various pure states, especially when we do not have at hand a simple breaking of a symmetry. A large part of the work on spin glasses has been devoted to this problem and we shall not try to reproduce it here [see Ref. [3] for a review], nor to give a general definition of states at finite temperature in our problems.

B. States at zero temperature

Instead we shall focus on the zero-temperature limit ($\beta\to\infty$), where the situation is simpler. A state is defined in the thermodynamic limit as a cluster of configurations, all of equal energy, related to each other by single spin flip moves, and which are locally stable, in the sense that the energy cannot be decreased by any flip of a finite number of spins. In the random 3-SAT problem one can use an even simpler definition which allows to generalize the definition of states also to finite N problems: the condition of local stability can be substituted by a condition of stability with respect to any

sequence of one spin flips. The reason for this simplification, specific to the *K*-SAT problem, is that in this case the stability with respect to sequences of single spin flips insures stability with respect to collective flips of finite sets of spins.

C. Many states: Definition of the complexity

Experience with disordered and frustrated systems like glasses shows that there can exist many states, and the number of states typically grows exponentially with the number of variables. The number of states $\mathcal{N}(E)$ with energy E is written as

$$\mathcal{N}(E)=\exp[N\Sigma(\alpha,\epsilon)], \tag{16}$$

where the quantity $\Sigma(\alpha,\epsilon)$ is called the complexity. It is a function of $\alpha=M/N$ and $\epsilon=E/N$, and the form (16) is derived from the basic assumption that $\ln\mathcal{N}(E)$ is extensive. In general, whenever a problem has a nonzero complexity, one may expect that simple local algorithms will have great difficulty in finding the ground state, simply because the states proliferate (for large N) and the algorithm will easily get trapped into one state with energy above that of the ground state. We shall see in the next sections how the cavity method can handle such a situation.

VI. A PRIMER ON THE CAVITY METHOD AT ZERO TEMPERATURE

The cavity method was originally introduced in Ref. [28] to study spin glasses, but it gives a general framework for computing statistical properties of various frustrated systems, and is ideally adapted to systems with a locally treelike structure. It is always in principle equivalent to the replica method, which is a more compact and very appealing formalism; however, it possesses two advantages. On one hand, it proceeds through a standard probabilistic analysis, and makes explicit all the hypotheses involved in it. Roughly speaking, the cavity method assumes some properties about the correlations between variables in a system with N spins, and shows that these are self-consistently reproduced for a system with $N+1$ spin system. The problem in turning it into a rigorous methods is that these hypotheses only hold in the large N limit, not for N small. If one is able to have a good control of the correlations as a function of N, then the cavity method becomes also a choice method for rigorous probabilistic studies of frustrated systems [29]. On the other hand, in the cavity method, one considers explicitly the site dependence of the order parameter, and the averaging over ''disorder'' is performed at the end (this is in contrast with the replica approach where the disorder average is made from the very beginning). As we shall see here, this aspect allows to define some algorithm, inspired from the cavity method, which computes the order parameter on each site for one given sample.

In what follows we shall be interested in the zero-temperature version of the cavity method. As discussed in detail in Ref. [21], the formalism simplifies a lot in this limit. Here we shall mainly outline for completeness the basic aspects of the method, applied to the 3-SAT model where all function nodes involve exactly three spins (the generalization to more general problems is totally straightforward but would make the notation more cumbersome). We refer the interested reader to Ref. [21] for more details. We shall first present the method in its simple replica symmetric (RS) version where it assumes the presence of a single state, and we shall then turn to the more involved case in which many states exist but are uncorrelated, a situation called one-step replica symmetry breaking (1RSB) in the replica jargon.

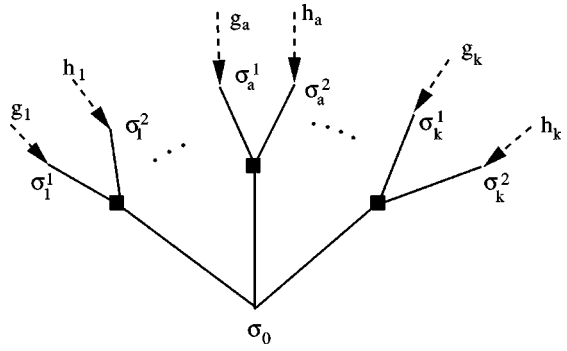

FIG. 2. When a new spin σ_0 is added to the system, it gets connected through k new function nodes to $2k$ other spins, σ_a^1 and σ_a^2. The cavity field on σ_a^1 is denoted by g_a, the one on σ_a^2 is denoted by h_a.

A. The cavity method with one single state (RS case)

1. Adding one spin

Consider a N spin system $\sigma_1,\ldots,\sigma_N$ and its interaction graph, and add to it a new spin σ_0. Then generate the new function nodes involving this new spin as follows: for each pair $1\leqslant i<j\leqslant N$, the function node $(0,i,j)$ is present with probability $6\alpha/N^2$. Therefore, we have added k new function nodes which we label by $a=1,\ldots,k$, where k is a random variable with probability distribution $f_{3\alpha}(k)$. Let us consider all the new function nodes which involves, besides σ_0, $2k$ other spins which we call σ_a^1 and σ_a^2 (see Fig. 2). Generically, on the original graph (i.e., before adding σ_0), these spins are far apart from each other. If there exists only one state, the clustering property implies that the correlations between these spins, before adding σ_0, vanish. Using the fact that σ_a^1 and σ_a^2 are binary variables, this decorrelation implies that the minimal energy of the original graph, for fixed values of the $2k$ spins σ_a^1 and σ_a^2, can be written as

$$E(\{\sigma_a^1,\sigma_a^2\})=A-\sum_{a=1}^{k}(g_a\sigma_a^1+h_a\sigma_a^2), \tag{17}$$

where the $2k$ local fields g_a and h_a are nothing but cavity-fields passed from each spin to the function node a, and A is a constant (independent of the local fields).

Looking at the function node a in the full graph including spin σ_0, we need to minimize the function $E_a(\sigma_0,\sigma_a^1,\sigma_a^2)-(g_a\sigma_a^1+h_a\sigma_a^2)$ with respect to σ_a^1,σ_a^2. This is precisely what one does in the message passing procedure of Sec. III,

and one can thus use Eq. (10) to get the minimal energy of the new graph with $N+1$ spins, for a given value of σ_0:

$$E(\sigma_0)=A-\sum_{a=1}^{k}\hat{w}_{\mathbf{J_a}}(g_a,h_a)-\sigma_0\sum_{a=1}^{k}\hat{u}_{\mathbf{J_a}}(g_a,h_a), \quad (18)$$

where $E_0=A-\Sigma_{a=1}^{k}(|h_a^1|+|h_a^2|)$ is the minimal energy of the N-spin system.

Equation (18) shows that the cavity field on the new spin σ_0 (the coefficient of $-\sigma_0$) can be written as

$$h_0=\sum_{a=1}^{k}\hat{u}_{\mathbf{J_a}}(g_a,h_a). \quad (19)$$

As we shall see, its is often useful to decompose this cavity field as a sum of cavity biases

$$h_0=\sum_{a=1}^{k}u_a; \quad u_a=\hat{u}_{\mathbf{J_a}}(g_a,h_a). \quad (20)$$

2. Self-consistency equation: The order parameter

Whenever one adds a new spin σ_0, one picks up a value of k, and a set of $2k$ fields g_a,h_a, which are iid variables taken from a probability distribution $\mathcal{P}(h)$. The cavity method assumes the existence of a thermodynamic limit $N\to\infty$ where the energy density E/N and the distribution of local fields $\mathcal{P}(h)$ have well defined limits. This means that the distribution of h_0 is the same as that of the $2k$ fields. Calling $\mathcal{Q}(u)$ the probability distribution of the u variables, this stability condition of the iteration (20) implies that

$$\mathcal{Q}(u)=\int dg\,dh\,\mathcal{P}(g)\mathcal{P}(h)\mathcal{E}_J\delta[u-\hat{u}_{\mathbf{J}}(g,h)],$$

$$\mathcal{P}(h)=\sum_{k=0}^{\infty}f_{3\alpha}(k)\int du_1\cdots du_k\mathcal{Q}(u_1)\cdots\mathcal{Q}(u_k)\times\delta\left(h-\sum_{a=1}^{k}u_a\right), \quad (21)$$

where $\mathcal{E}_J$ means an expectation value with respect to all the couplings J.

3. Computing the energy

One can easily compute the average shift in the ground state energy when adding one new spin. Looking at the addition process defined in Sec. VI A 1, we see that the energy of the original graph with N spins is $A-\Sigma_{a=1}^{k}(|g_a|+|h_a|)$, while that of the $N+1$ spin system is $A-\Sigma_{a=1}^{k}\hat{w}_{\mathbf{J}}(g_a,h_a)-|\Sigma_{a=1}^{k}\hat{u}_{\mathbf{J}}(g_a,h_a)|$. Therefore, the shift in energy when adding the new spin σ_0 is

$$\Delta E_1^{(0)}=\sum_{a=1}^{k}(-\hat{w}_{\mathbf{J_a}}(g_a,h_a)+|g_a|+|h_a|)-\left|\sum_{a=1}^{k}\hat{u}_{\mathbf{J_a}}(g_a,h_a)\right|. \quad (22)$$

Equation (21) gives an integral equation for the order parameter which is the probability distribution $\mathcal{P}(h)$ [or alternatively $\mathcal{Q}(u)$]. Let us now suppose that this equation has been solved (we shall see below how this can be done on each specific example), and show how the energy density can be deduced from this order parameter. We must compute the average of the energy shift for adding one spin (averaged over the choice of k and of the corresponding cavity fields):

$$\Delta E_1=\sum_{k=0}^{\infty}f_{3\alpha}(k)\mathcal{E}_J\int\prod_{a=1}^{k}[dg_a\mathcal{P}(g_a)dh_a\mathcal{P}(h_a)]\times\left(\sum_{a=1}^{k}[-\hat{w}_{\mathbf{J_a}}(g_a,h_a)+|g_a|+|h_a|]-\left|\sum_{a=1}^{k}\hat{u}_{\mathbf{J_a}}(g_a,h_a)\right|\right). \quad (23)$$

One might believe that, as the energy grows linearly in N at large N, this average energy shift would be equal to the energy density; however, there is a correction term due to the change in the number of function nodes per variable in the iteration $N\to N+1$. Indeed in the $N+1$ spin system we are generating function nodes with probability $6\alpha/N^2$ in a system with $N+1$ vertices and therefore we are slightly overgenerating function nodes. We need to cancel a fraction $1-N^2/(N+1)^2\simeq 2/N$ of them at random: the probability of deleting k' function nodes is

$$\binom{\alpha N}{k'}(2/N)^{k'}(1-2/N)^{\alpha N-k'}\simeq(2\alpha)^{k'}\exp(-2\alpha)/k'!=f_{2\alpha}(k') \quad (24)$$

and the average number of deleted function nodes is 2α. Each deleted function node contributes to the average energy change with a correction term

$$\Delta E_2=\mathcal{E}_J\int dh_1\,dh_2\,dh_3\,\mathcal{P}(h_1)\mathcal{P}(h_2)\mathcal{P}(h_3)\times\{\min_{\sigma_1,\sigma_2,\sigma_3}[E(\sigma_1,\sigma_2,\sigma_3)-h_1\sigma_1-h_2\sigma_2-h_3\sigma_3]+|h_1|+|h_2|+|h_3|\}. \quad (25)$$

The ground state energy density is finally given by

$$\epsilon_0=\Delta E_1-2\alpha\Delta E_2. \quad (26)$$

B. The cavity method with many states (one-step RSB)

1. Iteration within one state

Let us now see how the cavity method can be used to handle a situation in which there exist many states. As far as the clustering condition holds within each state, the iterative method can still be applied to each state. The problem is that the iteration induces some crossings of the energies of the states, and one needs to take this effect into account properly. We proceed as in the previous section by adding the new spin σ_0 connected to the $2k$ spins $\{\sigma_a^1,\sigma_a^2\}$. In each state α, one can reproduce the previous arguments: due to the vanishing of correlations, the energy of the state α, for fixed values of the $2k$ spins $\{\sigma_a^1\}$ and $\{\sigma_a^2\}$, can be written as

$$E^\alpha(\{\sigma_a^1,\sigma_a^2\})=A^\alpha-\sum_{a=1}^{k}(g_a^\alpha\sigma_a^1+h_a^\alpha\sigma_a^2). \tag{27}$$

We now have, for each state α, $2k$ local fields.

Within each state α, the optimization procedure on the $2k$ spins σ_a^1,σ_a^2 proceeds as before. The minimal energy of the new graph with $N+1$ spins, for a given value of σ_0, is

$$E^\alpha(\sigma_0)=A^\alpha-\sum_{a=1}^{k}\hat{w}_{\mathbf{J_a}}(g_a^\alpha,h_a^\alpha)-\sigma_0\sum_{a=1}^{k}\hat{u}_{\mathbf{J_a}}(g_a^\alpha,h_a^\alpha). \tag{28}$$

This shows that the local field on the new spin σ_0 in state α can be written as

$$h_0^\alpha=\sum_{a=1}^{k}\hat{u}_{\mathbf{J_a}}(g_a^\alpha,h_a^\alpha), \tag{29}$$

and the shift in energy of this state is [see Eq. (22)]

$$\Delta E^\alpha=\sum_{a=1}^{k}(-\hat{w}_{\mathbf{J_a}}(g_a^\alpha,h_a^\alpha)+|g_a^\alpha|+|h_a^\alpha|)-\left|\sum_{a=1}^{k}\hat{u}_{\mathbf{J_a}}(g_a^\alpha,h_a^\alpha)\right|. \tag{30}$$

2. Hypotheses on the states: u surveys

We suppose the existence of many states, with a complexity function $\Sigma(\alpha,E/N)$ defined as in Eq. (16) which is an increasing convex function. Let us consider all the states α with a given energy density $E/N=e$. We suppose that all the local fields $h_{j\to a}^\alpha$ on a given edge $j\to a$ are iid, taken from a probability distribution $P_{j\to a}^e(h)$ called an h survey. This probability distribution fluctuates from one edge to the next, so that the full order parameter, obtained by averaging over edges, is the functional probability distributions of these h surveys. The same hypotheses hold for the distribution of the cavity biases: all the $u_{a\to 0}$ on a given link are iid taken from a probability distribution $Q_{a\to 0}^e(u)$ called a u survey. Notice that the previous RS solution corresponds to having deterministic messages on each edge $P_{j\to a}^e(h)=\delta(h-h_{j\to a})$ and $Q_{a\to 0}^e(u)=\delta(u-u_{a\to 0})$. In the many state hypothesis, we have to generalize the messages, and the important quantities are the probability distributions (over the many states having a fixed energy density) of the cavity-biases going through a given link.

3. Iteration: level crossings and reweighting

One iteration step of the cavity procedure leads to an equation relating the probability distributions, before any averaging over the graph. In our iteration procedure, the h survey on the new site, $P_0^e(h)$, is related to the h surveys on the other $2k$ spins σ_a^1,σ_a^2 computed in the absence of σ_0. Let us denote by $P_a^e(g_a)$ the h survey incoming onto σ_a^1 and $P_a^{e\prime}(h_a)$ the h-survey incoming onto σ_a^2. The h survey $P_0^e(h)$ is given by

$$P_0^e(h)=C\int\prod_{a=1}^{k}[P_a^e(g_a)dg_aP_a^{e\prime}(h_a)dh_a]\times\delta\left[h-\sum_{a=1}^{k}\hat{u}_{\mathbf{J_a}}(g_a,h_a)\right]\times\exp\left[y\sum_{a=1}^{k}(\hat{w}_{\mathbf{J_a}}(g_a,h_a)-|g_a|-|h_a|)+y\left|\sum_{a=1}^{k}\hat{u}_{\mathbf{J_a}}(g_a,h_a)\right|\right], \tag{31}$$

where C is a normalization constant insuring that $P_0^e(h)$ is a normalized probability distribution. This equation provides the generalization to the RSB case of the simple iteration (20) of the previous section. Two complications have appeared: the simple messages (cavity fields and cavity biases) have become h surveys, i.e., probability distributions of simple messages, and a new term has appeared which is the exponential reweighting term. In this term, the parameter y is a number equal to the derivative of the complexity with respect to the energy:

$$y=\frac{\partial\Sigma}{\partial e}. \tag{32}$$

Let us now explain the origin of this new and crucial term. For a given state α, we add one new spin to the system and want to compute the new h survey. In this process, we have seen in Eq. (30) that there is an energy shift ΔE^α *which depends on the state, and is correlated to the value of the cavity-field* $h=\Sigma_{a=1}^k\hat{u}_{\mathbf{J}}(g_a^\alpha,h_a^\alpha)$. Let us call $S_0(h,\Delta E)$ the joint probability (when looking at all states) of the cavity bias and the free-energy shift, for this function node.

When we compute $P_0^e(u)$ at a fixed energy density $e=E/N$, we get a contribution from all states with energies before iteration equal to $E-\Delta E$. Therefore,

$$P_0^e(h) = C \int d(\Delta E) S_0(h,\Delta E) \exp\left[N\Sigma\left(\frac{E-\Delta E}{N}\right)\right]$$

$$\simeq C' \int d(\Delta E) S_0(h,\Delta E) \exp(-y\Delta E). \qquad (33)$$

The reweighting term in $\exp(-y\Delta E)$ is due to the level crossing, and the fact that the complexity $\Sigma(\alpha,\epsilon)$ is not constant, but increasing: states with a negative value of the energy shift are thus favored.

It is useful, as before, to decompose the iteration procedure (31) into two steps and introduce the u surveys. On any function node a, we merge two h surveys $P_a^e(g_a)$ and $P_a^e{}'(h_a)$ in order to build a u survey:

$$Q_a^e(u) = \int dg\, dh\, P_a^e(g) P_a^e{}'(h)\, \delta(u - \hat{u}_{\mathbf{J_a}}(g,h))$$

$$\times \exp\{y[\hat{w}_{\mathbf{J_a}}(g,h) - |g| - |h|]\}. \qquad (34)$$

Then we can combine all the u surveys incoming onto the new spin in order to build its h survey

$$P_0^e(h) = \int du_1, \ldots, du_k\, Q_1^e(u_1), \ldots, Q_k^e(u_k)$$

$$\times \exp\left(y \left| \sum_{a=1}^{k} u_a \right| \right) \delta\left(h - \sum_{a=1}^{k} u_a \right). \qquad (35)$$

Note that there is some degree of arbitrariness in the way one distributes the reweighting between the two iteration steps: different choices amount to different definitions of u surveys. The above one is the most natural one, and this is what we shall adopt from now on.

4. Order parameter and self-consistency: Population dynamics

Equations (34) and (35) are the main result giving the way to compute the messages sent along to a new added site. The assumed existence of a thermodynamic limit allows in principle to write a self-consistency of the iteration in a way similar to Eq. (21). In the present case, this is an equation for the functional $\mathcal{P}[P(h)]$ giving the probability, when one picks up an edge $i \to a$ at random, to observe on this edge a h survey $P_{i\to a}(h)$ equal to $P(h)$. Alternatively, one can use the functional $\mathcal{Q}[Q(u)]$ giving the probability, when one picks up an edge $a \to j$ at random, to observe on this edge a u survey $Q_{a\to j}(u)$ equal to $Q(u)$. In the following we shall rather work with the u surveys which turn out to have a simpler structure in practice, but obviously a fully equivalent description can be obtained working with h surveys.

These functional equations are the generalization to the RSB case of the Eqs. (21) for the RS case. One could write them explicitly, but they are not particularly illuminating, and we prefer to work directly with the iteration equations (34) and (35). These define a stochastic process; at each iteration, one performs the following operations.

(1) Pick up at random a number of neighbors k, with the probability $f_{3a}(k)$.

(2) Pick up at random k surveys of u, $Q_1(u_1), \ldots, Q_k(u_k)$ from the distribution $\mathcal{Q}[Q(u)]$.

(3) Compute a h survey $P_1(g)$ as the reweighted convolution

$$P_1(g) = C_1 \int du_1, \ldots, du_k\, Q_1(u_1), \ldots, Q_k(u_k)$$

$$\times \exp\left(y \left| \sum_{a=1}^{k} u_a \right| \right) \delta\left(g - \sum_{a=1}^{k} u_a \right). \qquad (36)$$

(4) Pick up at random a number of neighbors k', with the probability $f_{3\alpha}(k')$.

(5) Pick up at random k' surveys of u, $Q_{k+1}(u_1), \ldots, Q_{k+k'}(u_{k'})$ from the distribution $\mathcal{Q}[Q(u)]$;

(6) Compute a h survey $P_2(h)$ as the convolution

$$P_2(h) = C_2 \int du_1, \ldots, du_k\, Q_{k+1}(u_1), \ldots, Q_{k+k'}(u_{k'})$$

$$\times \exp\left(y \left| \sum_{a=1}^{k'} u_a \right| \right) \delta\left(h - \sum_{a=1}^{k'} u_a \right). \qquad (37)$$

(7) Pick up at random a set of couplings $\mathbf{J}$ characterizing a new function node, from the *a priori* distribution of couplings.

(8) Compute a new u survey, $Q_0(u)$ as

$$Q_0(u) = C_0 \int dg\, dh\, P_1(g) P_2(h)\, \delta(u - \hat{u}_{\mathbf{J}}(g,h))$$

$$\times \exp\{y[\hat{w}_{\mathbf{J}}(g,h) - |g| - |h|]\}, \qquad (38)$$

where C_0 is a normalization constant insuring that $Q(u)$ has an integral equal to one.

This iteration defines a stochastic process in the space of u surveys, which in turn defines a flow for $\mathcal{Q}[Q(u)]$, of which we would like to compute the fix point. Following Ref. [20], this is done in practice by a population dynamics algorithm: one uses a representative population of $\mathcal{N}$ u surveys from which the various $Q_\ell(u)$, $\ell \in \{1, \ldots, k+k'\}$ used in the iteration are extracted. After $Q_0(u)$ has been computed, one of the u surveys in the population, chosen randomly, is erased and substituted by $Q_0(u)$. After some transient, this population dynamics algorithm generates sets of u surveys which are sampled with a frequency proportional to the seeked $\mathcal{Q}[Q(u)]$.

The point of this stochastic process approach is to avoid trying to write explicitly the complicated functional equation satisfied by $\mathcal{Q}[Q(u)]$. This is one crucial place where the cavity method turns out to be superior to the replica method: with replicas one performs the average over disorder from the beginning, and one is forced to work directly with the functional $\mathcal{Q}[Q(u)]$ [30]. As this is very difficult, people have thus been constrained to look for approximate solutions of $\mathcal{Q}[Q(u)]$ where the functional is taken in a simple subspace, allowing for some explicit computations to be done.

5. Computing the energy and the complexity

Here we show how to generalize the computation of the energy of Sec. VI A 3 to the 1RSB case.

When adding one site 0, connected through k function nodes to $2k$ sites the energy shift in one given state is

$$\delta E = \sum_{\ell=1}^{k} [-\hat{w}_{\mathbf{J}_\ell}(h_\ell, g_\ell) + |h_\ell| + |g_\ell|] - \left| \sum_{\ell=1}^{k} \hat{u}_{\mathbf{J}_\ell}(h_\ell, g_\ell) \right|, \tag{39}$$

where h_ℓ, g_ℓ are the incoming fields onto the function node number ℓ. Let us call $P_\ell(h_\ell)$ and $P'_\ell(g_\ell)$ the corresponding field distributions (the h surveys). They induce a probability distribution $P_0(\delta E)$ of the energy change (39)

$$P_0(\delta E) = \int \prod_{\ell=1}^{k} [dh_\ell\, dg_\ell\, P_\ell(h_\ell) P'_\ell(g_\ell)] \times \delta\left(\delta E + \sum_{\ell=1}^{k} [\hat{w}_{\mathbf{J}_\ell}(h_\ell, g_\ell) - |h_\ell| - |g_\ell| + \left| \sum_{\ell=1}^{k} \hat{u}_{\mathbf{J}_\ell}(h_\ell, g_\ell) \right| \right). \tag{40}$$

Let us look at the corresponding change in complexity. The new system has $N+1$ variables and $M+k$ function nodes, and its number of states at energy E, $\exp[(N+1)\Sigma([M+k]/[N+1],[E/N+1])]$ is given by

$$\exp\left[(N+1)\Sigma\left(\frac{M+k}{N+1}, \frac{E}{N+1} \right) \right] = \int d(\delta E) P_0(\delta E) \exp\left[N\Sigma\left(\frac{M}{N}, \frac{E-\delta E}{N} \right) \right]. \tag{41}$$

This expression depends on the precise spin which has been added through the choice of the distributions P_ℓ and P'_ℓ, and of couplings $\mathbf{J}_\ell$, which appear in Eq. (39). As one expects $\Sigma(\alpha, \epsilon)$ to be self-averaging, one must average the logarithm of the expressions in Eq. (41) over the iteration of population dynamics algorithm. We denote this averaging by an overline. As δE is finite, one can expand in $\delta E/N$ in the thermodynamic limit, to get (calling as always $\alpha = M/N$ and $\epsilon = E/N$)

$$\Sigma(\alpha, \epsilon) - \epsilon \frac{\partial \Sigma}{\partial \epsilon} + 2\alpha \frac{\partial \Sigma}{\partial \alpha} = \overline{\ln\left(\int d(\delta E) P_0(\delta E) \exp[-y\, \delta E] \right)}, \tag{42}$$

where we used the fact that $\overline{k} = 3\alpha$. As in the RS case of Eq. (25), the derivative $\partial\Sigma/\partial\alpha$ can be computed by adding one function node to the system. For a generic function node a connected to the sites 1,2,3 and with interaction coupling $\mathbf{J}$, the probability distribution $P_a(\delta E)$ of the energy change is

$$P_a(\delta E) = \int dh_1 dh_2 dh_3 P_1(h_1) P_2(h_2) P_3(h_3) \times \delta[\, \delta E - \min_{\sigma_1, \sigma_2, \sigma_3} \{ -h_1\sigma_1 - h_2\sigma_2 - h_3\sigma_3 + \epsilon_{\mathbf{J}}(\sigma_1, \sigma_2, \sigma_3) \} - (|h_1| + |h_2| + |h_3|)]. \tag{43}$$

Let us look at the corresponding change in complexity. The new system has N variables and $M+1$ function nodes, so that

$$\exp\left[N\Sigma\left(\frac{M+1}{N}, \frac{E}{N} \right) \right] = \int d(\delta E) P_a(\delta E) \exp\left[N\Sigma\left(\frac{M}{N}, \frac{E-\delta E}{N} \right) \right]. \tag{44}$$

After averaging over the iterations, this gives

$$\frac{\partial \Sigma}{\partial \alpha} = \overline{\ln\left(\int d(\delta E) P_a(\delta E) \exp[-y\, \delta E] \right)}. \tag{45}$$

Combining the two expressions (42) and (45), it turns out that the quantity which is computed naturally in this scheme is the Legendre transform $\Phi(y)$, with respect to the energy density ϵ, of the complexity function $\Sigma(\alpha, \epsilon)$ [31,32]. This "zero-temperature free energy" is defined precisely as

$$\Sigma(\alpha, \epsilon) - y\epsilon = -y\Phi(y); \quad y \equiv \frac{d\Sigma}{d\epsilon} \tag{46}$$

and it can be computed from the population dynamics as

$$\Phi(y) = \Phi_1(y) - 2\alpha \Phi_2(y), \tag{47}$$

$$\Phi_1(y) = -\frac{1}{y} \overline{\ln\left(\int d(\delta E) P_0(\delta E) \exp[-y\, \delta E] \right)},$$

$$\Phi_2(y) = -\frac{1}{y} \overline{\ln\left(\int d(\delta E) P_a(\delta E) \exp[-y\, \delta E] \right)}, \tag{48}$$

where P_0 and P_a are given in Eqs. (40) and (43).

Technically, it turns out that $\Phi_1(y)$ is more easily computed through the normalization of the u surveys. When we compute a u survey $Q_0(u)$ as in Eq. (38), we can memorize the corresponding normalization constant C_0. Picking up k u surveys Q_ℓ at random in the population, and calling C_ℓ the corresponding normalizations, one gets

$$\Phi_1(y) = -\frac{1}{y} \overline{\ln\left(\int \prod_{\ell=1}^{k} \left[du_\ell \frac{Q_\ell(u_\ell)}{C_\ell} \right] \exp\left[y \left| \sum_{\ell=1}^{k} u_\ell \right| \right] \right)}. \tag{49}$$

VII. THE CAVITY METHOD APPLIED TO THE RANDOM 3-SAT PROBLEM

A. Known results on the phase diagram

Considering the random 3-SAT problem where the graph is generated at random and the various couplings take values ± 1 with probability 1/2, numerical experiments have provided a detailed study of the probability $P_N(\alpha,K)$ that a given F including $M=\alpha N$ clauses be satisfiable. For large sizes, there appears a remarkable behavior: P seems to reach unity for $\alpha<\alpha_c(K)$ and vanishes for $\alpha>\alpha_c(K)$ [13]. Such an abrupt threshold behavior, separating a SAT phase from an UNSAT one, has indeed been rigorously confirmed for 2-SAT, which is in P, with $\alpha_c(2)=1$ [33–35]. For larger $K \geqslant 3$, K-SAT is NP-complete and much less is known. The existence of a sharp transition has not been rigorously proven yet but estimates of the thresholds have been found. The present best numerical estimate for α_c at $K=3$ is 4.26 [36], and the rigorous bounds are [37–40] $3.26<\alpha_c<4.506$, while previous statistical mechanics analysis using the replica method, has found $\alpha_c(3)\sim 4.48$ [41] and $\alpha_c(3)\sim 4.396$ [42] in the framework of variational approximations.

The interest in random K-SAT arises from the fact that it has been observed numerically that hard random instances are generated when the problems are critically constrained, i.e., close to the SAT/UNSAT phase boundary [13,24]. The study of such hard instances represent a theoretical challenge towards a concrete understanding of complexity and the analysis of algorithms [15]. Moreover, hard random instances are also a testbed for the optimization of heuristic (incomplete) search procedures, which are widely used in practice.

B. The cavity analysis with one state

In the 3-SAT problem, the energy of a function node, as given by Eq. (5), is

$$E_a=(1+J_1^a\sigma_1)(1+J_2^a\sigma_2)(1+J_3^a\sigma_3)/4$$

and leads to

$$\hat{w}_{\mathbf{J_a}}(h_2,h_3)=|h_2|+|h_3|-\theta(J_2^a h_2)\,\theta(J_3^a h_3),$$

$$\hat{u}_{\mathbf{J_a}}(h_2,h_3)=-J_1^a\,\theta(J_2^a h_2)\,\theta(J_3^a h_3), \tag{50}$$

where the function $\theta(x)$ is defined as

$$\theta(x)=1 \quad \text{if} \quad x>0; \quad \theta(x)=0 \quad \text{if} \quad x\leqslant 0. \tag{51}$$

Let us consider the cavity iteration scheme, within the hypothesis of there being a single state. We thus use the general formalism presented in Sec. VI A. In the case of 3-SAT, Eq. (50) shows that the cavity bias on a given edge $u_{a\to j}$ takes values either in 0,1 if the corresponding coupling is negative, otherwise it takes values in $-1,0$. Therefore, the cavity fields are integers. [This is the reason for using the unusual factor 2 for each violated clause in the definition (5) of the energy]. Because the function $\hat{u}_{\mathbf{J_a}}$ is an odd function of one of the couplings in $\mathbf{J_a}$, and these couplings are random variables taking values ± 1 with probability 1/2, the distribution $\mathcal{Q}(u)$ of cavity biases must be of the form

$$\mathcal{Q}(u)=c_0\,\delta(u)+\frac{(1-c_0)}{2}[\,\delta(u-1)+\delta(u+1)]. \tag{52}$$

Plugging this expression into the self-consistency equations (21) leads to a relation between c_0 and the weight $p_0\equiv\mathcal{P}(h=0)$

$$p_0=\sum_k f_{3\alpha}(k)\sum_{q=0}^{[k/2]}\binom{k}{2q}c_0^{k-2q}\left(\frac{1-c_0}{2}\right)^{2q}\binom{2q}{q}$$
$$=\exp[-3\alpha(1-c_0)]I_0(3\alpha(1-c_0)), \tag{53}$$

where I_0 is the Bessel function and $[k/2]$ is the integer part of $k/2$. Let us now compute c_0. From Eq. (50) we find that a cavity-bias vanishes whenever at least one of the incoming fields (h_2 or h_3) is zero or has a sign opposite to the corresponding coupling. This shows that

$$c_0=1-\mathrm{Prob}[(h<0)\cap(g<0)]=1-\left(\frac{1-p_0}{2}\right)^2. \tag{54}$$

We obtain a closed set of Eqs. (54) and (53) which is easily solved. The distribution $\mathcal{P}(h)$ of cavity fields is then given by $\mathcal{P}(h)=\Sigma_r p_r\delta(h-r)$, where the weights p_r are equal to

$$p_r\equiv\mathcal{P}(h=r)=\mathcal{P}(h=-r)=\sum_{k=r}^{\infty}f_{3\alpha}(k)\sum_{q=0}^{[(k-r)/2]}\binom{k}{2q+r}$$
$$\times c_0^{k-2q-r}\left(\frac{1-c_0}{2}\right)^{2q+r}\binom{2q+r}{q} \tag{55}$$

and the energy is computed from Eqs. (23) and (25).

A nontrivial solution exists for $\alpha>4.667$, with a ground state energy that becomes positive at $\alpha=5.18$. The prediction of this hypothesis assuming a single state is a paramagnetic SAT phase with $c_0=p_0=1$, and energy $E_0=0$ for $\alpha<5.18$, and a frozen UNSAT glassy phase with $c_0<1$ and $E_0>0$ for $\alpha>5.18$. It is known [4,5] that this solution is wrong both quantitatively (the location of the transition point) and qualitatively (the structure of the order parameter).

The true transition is much more sophisticated, and the many state formalism corresponding to 1RSB is needed to unveil its structure.

C. The cavity analysis with many states

We introduce as before the h surveys and u surveys, and we use the population dynamics algorithm defined in Sec. VI B 4. It turns out to be more convenient to work only in terms of u surveys. The algorithm computes at each iteration a new u survey $Q_0(u)$ by taking $k+k'$ u surveys $Q_1(u),\ldots,Q_{k+k'}(u)$ in the population through

$$Q_0(u) = C_0 \int \int du_1 Q_1(u_1) \ldots du_k Q_k(u_k) \times dv_1 Q_{k+1}(v_1) \ldots dv_{k'} Q_{k+k'}(v_{k'}),$$

$$\delta(u - \hat{u}_{\mathbf{J}}(u_1 + \cdots + u_k, v_1 + \cdots + v_{k'})) \times \exp[y \hat{w}_{\mathbf{J}}(u_1 + \cdots + u_k, v_1 + \cdots + v_{k'})]. \quad (56)$$

Here, k and k' are two iid random numbers taken from the Poisson distribution $f_{3\alpha}(k)$ defined in Eq. (12), and $\mathbf{J}$ denotes a set of three iid random numbers J_1, J_2, J_3, each being equal to ± 1 with probability 1/2.

The functions $\hat{u}$ and $\hat{w}$ are defined in Eq. (50). A u survey always takes the simple form $Q_0(u) = (1-c)\delta(u) + c\,\delta(u + J_1)$; it is thus a probability distribution which can be characterized by a single number c, and therefore the iteration of the population dynamics is easily done numerically.

D. Solution of the self-consistency equations

Apart from the RS solution with $y = 0$ and $Q_j(u) = \delta(u - u_j)$, where the u_j are iid taken from a distribution $\mathcal{Q}(u)$, the numerical solution finds one other solution in the region $\alpha > \sim 4$. Generically, the u surveys found can be of three types:

$$Q_i(u) = \begin{cases} \delta(u) & \text{(trivial or type a)} \\ (1 - \eta_i e^{-y})\,\delta(u) + \eta_i e^{-y}\,\delta(u-1) & \text{(type } b_+\text{)} \\ (1 - \eta_i e^{-y})\,\delta(u) + \eta_i e^{-y}\,\delta(u+1) & \text{(type } b_-\text{)}. \end{cases} \quad (57)$$

The arbitrary factor e^{-y} has been introduced for convenience because numerical simulations show that the weight in $u = \pm 1$ of the nontrivial u-surveys scale proportionally to e^{-y} at large y.

The statistical symmetry of the problem due to the fact that the couplings take values ± 1 with probability 1/2 implies that the probability of finding in the population a "b_+" message is equal to that of finding a "b_-" message. We call $(1-t)/2$ these probabilities, and t the probability of finding a type a message. Nontrivial messages are fully characterized by the distribution $\rho(\eta)$ of the η_i variables. For any y, the full solution of the problem is given by the value of t and of the function $\rho(\eta)$. It can be obtained numerically by averaging over many iterations of the population dynamics. We shall now show how one can get some analytic control in the large y limit.

1. Existence of nontrivial u surveys

Looking at the iteration equation (56), the only way one can obtain a trivial u survey $Q_0(u) = \delta(u)$ is when either $J_2 \Sigma_{i=1}^{k} u_i \leqslant 0$ in the whole integration domain, or $J_3 \Sigma_{j=1}^{k'} v_j \leqslant 0$ in the whole integration domain, or both. The probability to have $J_2 \Sigma_{i=1}^{k} u_i \leqslant 0$ in the whole integration domain is

$$\text{Prob}\left[J_2 \sum_{i=1}^{k} u_i \leqslant 0 \right] = \sum_{k=0}^{\infty} f_{3\alpha}(k) \left[t^k + \frac{1}{2} \frac{(1-t)}{2} t^{k-1} 2 + 2 \frac{1}{2} \left(\frac{1-t}{2} \right)^2 t^{k-2} \binom{k}{2} + \cdots \right] = \exp\left[-3\alpha \frac{1-t}{2} \right]. \quad (58)$$

The first term t^k corresponds to all $Q_i(u)$, $i \in \{1, \ldots, k\}$, being of type a. The second term corresponds to [one of the $Q_i(u)$ of type b_+, all the other ones of type a, and $J_2 = -1$] or [one of the $Q_i(u)$ of type b_-, all the other ones of type a, and $J_2 = +1$]. The rest of the series is easily obtained similarly.

The probability t of having a trivial u survey is thus

$$t = 1 - \left(1 - \text{Prob}\left[J_2 \sum_{i=1}^{k} u_i \leqslant 0 \right] \right) \left(1 - \text{Prob}\left[J_3 \sum_{j=1}^{k'} v_j \leqslant 0 \right] \right) = 1 - \left(1 - \exp\left[-3\alpha \frac{1-t}{2} \right] \right)^2. \quad (59)$$

For α small the only solution is the paramagnetic one $t = 1$. A solution different from $t = 1$ appears above $\alpha_0 = 1.636\,94$ (which corresponds to $t = 0.048\,83$). (In fact there appears a pair of nontrivial solutions, the relevant one is the one with lowest t, as can be seen from its behavior at large α, and also checked in the population dynamics algorithm). It is interesting to observe that this threshold coincides with the point where the so called "unit clause literal" [27] algorithm ceases to converge. At large values of α, the fraction of u surveys of type a becomes very small, e.g., $t = 0.0051, 0.0011, 0.00025$, at $\alpha = 4, 5, 6$, respectively.

2. Expansion at large y: Location of the phase transition

Here we shall show that, at large y, the scaling in e^{-y} of the weights at $u = \pm 1$ for nontrivial u surveys is indeed consistent with the iteration of the population dynamics (56), and we deduce from this iteration a self-consistent equation for $\rho(\eta)$.

Let us study one given iteration of Eq. (56). The probability of having k cavity biases $Q_1, \ldots, Q_k$, among which m are of type b_+, n of type b_- and $k-m-n$ of type a is

$$C_{k,m,n} = f_{3\alpha}(k) \frac{k!}{m!n!(k-m-n)!} t^{k-m-n} \left(\frac{1-t}{2} \right)^{m+n}. \quad (60)$$

Without loss of generality, we can assume that the $m+n$ nontrivial u surveys are $Q_1, \ldots, Q_{m+n}$, and the couplings are $J_1 = J_2 = J_3 = 1$ for this iteration. We denote by η the vector of weights $\eta = (\eta_1, \ldots, \eta_{m+n})$. Since $\hat{u}_{\mathbf{J}}(u_1 + \cdots + u_k, v_1 + \cdots + v_{k'}) = 1$ if and only if $(J_1 \Sigma_i u_i > 0)$ and $(J_2 \Sigma_j v_j > 0)$, the new u survey $Q_0(u)$ depends only on the probabilities

$$f_q^{(m,n)}(\eta) \equiv \text{Prob}[u_1 + \cdots + u_k = q] \quad (61)$$

and is given by

$$Q_0(u)=C_0\Bigg\{\delta(u-1)\Bigg[\sum_{q=1}^{k}\sum_{q'=1}^{k'} f_q^{(m,n)}(\eta)f_{q'}^{(m',n')}(\eta') \times e^{y(|q|+|q'|-1)}\Bigg] + \delta(u)\Bigg[\Bigg(\sum_{q=1}^{k}\sum_{q'=-k}^{0}+\sum_{q=-k}^{0}\sum_{q'=1}^{k'}+\sum_{q=-k}^{0}\sum_{q'=-k}^{0}\Bigg) \times f_q^{(m,n)}(\eta)f_{q'}^{(m',n')}(\eta')e^{y(|q|+|q'|)}\Bigg]\Bigg\} \equiv C_0\{\delta(u-1)A_0e^{-y}+\delta(u)B_0\}. \tag{62}$$

Iterating the population dynamics, one finds that $\rho(\eta)$ satisfies the equation

$$\rho(\eta_0)=\frac{1}{1-t}\sum_{k=0}^{\infty}\sum_{m=1}^{k-1}\sum_{n=0}^{k-m}C_{k,m,n}\sum_{k'=0}^{\infty}\sum_{m'=1}^{k'-1}\sum_{n'=0}^{k'-m'} \times C_{k',m',n'}\int\prod_{\ell=1}^{m+n}[d\eta_\ell\,\rho(\eta_\ell)] \times\prod_{\ell'=1}^{m'+n'}[d\eta_{\ell'}\rho(\eta_{\ell'})]\delta\Bigg(\eta_0-\frac{A_0}{A_0e^{-y}+B_0}\Bigg). \tag{63}$$

This is an exact self-consistent equation for the distribution $\rho(\eta)$. It can be simplified at large y, and we show in Appendix A how to write a more tractable self-consistent equation in this limit. This equation is best written in terms of the probability distribution function $S(\phi)$ of the variable $\phi=\ln(1+\eta)$. One finds that $S(\phi)$ satisfies

$$S(\phi)=\int dx_+\,dx_-\,dx'_+\,dx'_-\,A(x_+)B(x_-)A(x'_+)B(x'_-) \tag{64}$$

$$\delta\Bigg(\phi-\ln\Bigg[1+\frac{(e^{x_+}-1)(e^{x_-}-1)}{(e^{x_+}-1)e^{x'_-}+(e^{x'_+}-1)e^{x_-}+e^{x_-}e^{x'_-}}\Bigg]\Bigg), \tag{65}$$

where $A(x)$ and $B(x)$ are two probability distributions related to $S(\phi)$ through its Fourier transform $\hat{S}(q)\equiv\int d\phi\exp(iq\phi)S(\phi)$,

$$A(x)\equiv\frac{1}{e^{3\alpha(1-t)/2}-1}\int\frac{dq}{2\pi}e^{-iqx}\Bigg(\exp\Bigg[\frac{3\alpha}{2}(1-t)\hat{S}(q)\Bigg]-1\Bigg),$$

$$B(x)\equiv\frac{1}{e^{3\alpha(1-t)/2}}\int\frac{dq}{2\pi}e^{-iqx} \times\exp\Bigg[3\alpha\Bigg(t-1+\frac{1}{2}(1-t)\hat{S}(q)\Bigg)\Bigg]. \tag{66}$$

Equations (65) and (66) are simple enough to be solved numerically to high accuracy.

Once $S(\phi)$ [and therefore $\rho(\eta)$] is known, one can deduce the value of the zero-temperature free-energy $\Phi(y)$. To leading order at large y, we show in Appendix A that $\Phi(y)=\Psi/y$, with

$$\Psi=-\int dx\,dz\,B(x)B(z)\ln(e^x+e^z-1) -3\alpha\int\prod_{i=1}^{2}dx_i\,dy_i\,B(x_i)B(z_i) \times\ln\Bigg[\prod_{i=1}^{2}(e^{x_i}+e^{z_i}-1)-\prod_{i=1}^{2}(e^{x_i}-1)\Bigg] +2\alpha\int\prod_{i=1}^{3}dx_i\,dy_i\,B(x_i)B(z_i) \times\ln\Bigg[\prod_{i=1}^{3}(e^{x_i}+e^{z_i}-1)-\prod_{i=1}^{3}(e^{x_i}-1)\Bigg]. \tag{67}$$

In order to compute the $S(\phi)$,$A(x)$,$B(x)$ functions, solutions of Eqs. (65) and (66), we have used the fact that these functions are probability distributions, and we have developed a population dynamics algorithm which follows a population of N variables $\phi_1,\ldots,\phi_N$, for a given value of α.

(1) Compute t, the solution of $t=1-[1-\exp[3\alpha(t-1)/2]]^2$.

(2) Initialize the ϕ_j as iid random positive variables, for instance with an exponential distribution of width 1. Initialize the "time variable" $\tau=1$.

(3) Upgrade the time $\tau\to\tau+1$.

(4) Generate an integer $k\geqslant 1$ with the distribution $\gamma^k/k!1/[\exp(\gamma)-1]$, where $\gamma=3\alpha(1-t)/2$. Pick up k integers $i_1,\ldots,i_k$ at random in $\{1,\ldots,N\}$, and compute the sum $\phi_{i_1}+\cdots+\phi_{i_k}$. The distribution of the variable $x_+(\tau)=\phi_{i_1}+\cdots+\phi_{i_k}$ is $A(x_+)$, related to $S(\phi)$ through Eq. (66).

(5) Generate a random variable $x_-(\tau)$, which will be distributed according to $B(x_-)$, as follows: with probability $\exp(-\gamma)$, one takes $x_-(\tau)=0$; with probability $1-\exp(-\gamma)$, one repeats the procedure of (4) and calls the output $x_-(\tau)=\phi_{i_1}+\cdots+\phi_{i_k}$.

(6) Repeat the steps (4) and (5) to generate two other variables $x'_+(\tau)$ and $x'_-(\tau)$.

(7) Compute $\chi(\tau)=e^{x_-(\tau)}/(e^{x_+(\tau)}-1)$, $\chi'(\tau)=e^{x'_-(\tau)}/(e^{x'_+(\tau)}-1)$, and $\phi_0(\tau)=\ln[1+1/(\chi(\tau)+\chi'(\tau)+\chi(\tau)\chi'(\tau))]$.

(8) Replace one randomly chosen variable in the population, ϕ_ℓ, by the new value $\phi_0(\tau)$.

Steps (3)–(8) must be repeated a large number of times, say $\mathcal{T}$. One can compute the average of any function of ϕ as

$$\int d\phi\, S(\phi) f(\phi) = \frac{4}{3\mathcal{T}} \sum_{\tau=\mathcal{T}/4+1}^{\mathcal{T}} f(\phi_0(\tau)). \qquad (68)$$

(One forgets the first $\mathcal{T}/4$ iterations in order to reach a stationary regime). The integrals involving a variable x distributed according to $B(x)$ can be estimated by summing functions of x_-. For instance, the first term in the zero-temperature free-energy Eq. (67) is evaluated as

$$-\int dx\, dz\, B(x)B(z)\ln(e^x+e^z-1)$$
$$= \frac{4}{3\mathcal{T}} \sum_{\tau=\mathcal{T}/4+1}^{\mathcal{T}} \ln(e^{x_-(\tau)}+e^{x'_-(\tau)}-1), \qquad (69)$$

and the two other terms are computed similarly, using the x_-, x'_- values from two (respectively three) successive time steps.

In practice, we used values of $N\sim 10^6$ and $\mathcal{T}\sim 1000N$, which was enough to achieve the precision given below in the results.

For $\alpha<\alpha_d=3.921$, the algorithm converges towards the solution $\phi_1=\cdots=\phi_N=0$. This is the paramagnetic solution where all the u surveys are trivial.

For $\alpha>\alpha_d=3.921$, we find a new solution with a nontrivial distribution $S(\phi)$. Computing the leading large y behavior of the zero-temperature free-energy function, $\Phi(y) \sim \Psi/y$, on this solution using Eq. (67), we find that Ψ is negative for $\alpha<\alpha_c=4.267$, while it is positive for $\alpha>\alpha_c$ [Ref. [19] reported $\alpha_c=4.256$, which was 3% too low because of some correlations in the random number generator used in solving Eqs. (65),(66)].

E. Phase diagram of the random 3-SAT problem

The previous large y analysis is in agreement with the direct numerical iteration of the population dynamics of Sec. VI B 4, but it allows one to get a much more precise determination of the thresholds. The results of this section have been obtained through the combined use of the numerical and analytical methods.

For $\alpha<\alpha_d$, the system is in the SAT phase, the solution is paramagnetic, it is easy to find a solution. Note that, although the t equation (59) in principle allows for the existence of nontrivial knowledge above $\alpha\simeq 1.64$, we have not found such a solution and the only one which remains is the paramagnetic one with $t=1$.

For $\alpha_d<\alpha<\alpha_c$, we find a monotonically increasing $\Phi(y)$ function, which reaches its maximum $\Phi\to 0$ at $y\to\infty$ (see Fig. 3).

Fitting $\Phi(y)$ by a function $\Sigma_{\ell=1}^{p}\nu_\ell \exp(-y\ell)/y$ with $p\in\{1,2,3\}$ gives a good and stable fit, from which the Legendre transform (46) is easily done. This allows one to reconstruct the complexity curve $\Sigma(\alpha,\epsilon)$ (see Fig. 4). It is found to be finite down to $\epsilon=0$. This implies that there is an exponentially large (in N) number of states with zero energy density.

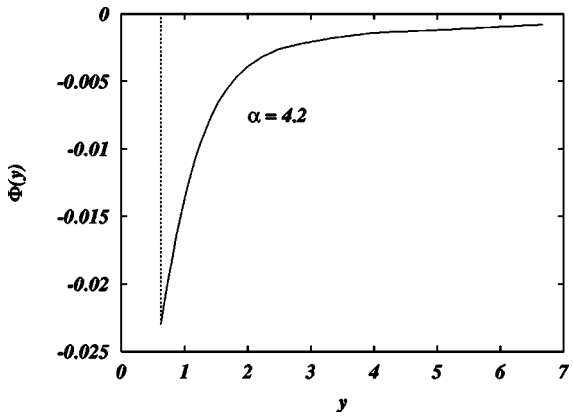

FIG. 3. Free energy $\Phi(y)$ versus the reweighting parameter y for the random 3-SAT problem at $\alpha=4.2$

The complexity of these ground states is plotted in Fig. 5. We call this phase the hard-SAT phase (HSP), since in this regime the typical sample is SAT, but the proliferation of states (most of which have strictly positive energy) makes it difficult to find a solution.

For $\alpha_c<\alpha$, the function $\Phi(y)$ has a maximum at a finite value y^*, and $\Phi(y^*)>0$. The complexity curve $\Sigma(\alpha,\epsilon)$ starts at a positive energy density $\epsilon_0=\Phi(y^*)$ (see Fig. 4). This energy density ϵ_0 is the minimal number of violated clauses per variable which will be found in almost all samples at large N. It is plotted in Fig. 5. We are in the UNSAT phase.

This figure also shows the value of ϵ where the complexity curve $\Sigma(\alpha,\epsilon)$ is maximum gives the energy where there exists the largest number of states. This is the ''threshold'' energy density e_{th} where a simple zero-temperature metropolis algorithm (ZTMA) will be trapped. This implies that ZTMA should find satisfying assignments only for $\alpha<\alpha_D$, in agreement with the numerical results of [43]. These predictions can be tested most clearly through their generalization to single instances which we discuss in the following section.

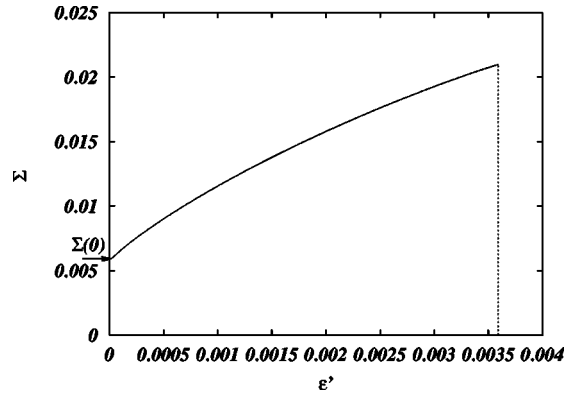

FIG. 4. Complexity Σ versus the fraction of violated clauses $\epsilon'=\epsilon/2$ for the random 3-SAT problem at $\alpha=4.2$, obtained from the Legendre transform of Fig. 3.

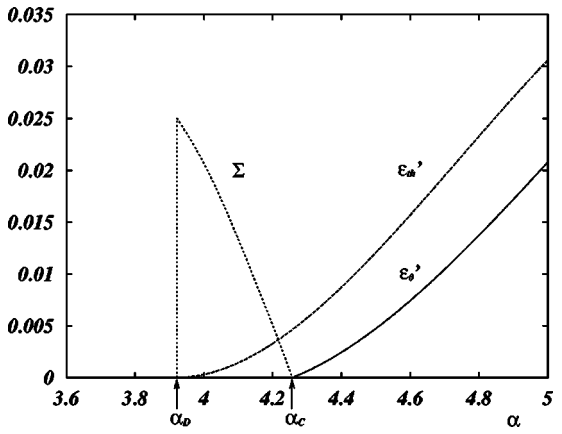

FIG. 5. The phase diagram of the random 3-SAT problem. Plotted is $\epsilon_0'=\epsilon_0/2$ (full line), the number of violated clauses per variable, versus the control parameter α which is the number of clauses per variable. The SAT-UNSAT transition occurs at $\alpha=\alpha_c\sim 4.267$. The dashed line is $\epsilon_{th}'=\epsilon_{th}/2$, the threshold energy (divided by two) per variable, where local algorithms get trapped. The dotted line is the complexity Σ of satisfiable states, equal to $1/N$ times the natural logarithm of their number.

Previous statistical mechanics attempts at finding this phase diagram culminated in powerful variational approximations using the replica method; see Ref. [41] for the first results and Ref. [42], which predicted approximate values for the SAT/UNSAT threshold—$\alpha_c\sim 4.48$ in the case of Ref. [41] and $\alpha_c\sim 4.39$ in Ref. [42]—with an intermediate phase appearing above $\alpha_s\simeq 3.96$ [41] and $\alpha_d\simeq 3.94$ with the Ansatz of Ref. [42].

The difference between the variational results and our cavity result is both quantitative and qualitative: in Ref. [41] the predicted nature of the intermediate phase is different with respect to ours while in Ref. [42] the structure of the order parameter is oversimplified. In the present approach (as well as in Ref. [42]) we work directly at zero temperature ($T=0$), which has the advantage that we do not need to study the subtle question of the limit $T\to 0$. The reason why this limit is subtle is due to the fact that some of the local fields, at low temperatures, vanish linearly in T, and thus contribute to the local magnetization $m=\tanh(\beta H)$; we call these fields evanescent fields. The local magnetization at $T=0$ is zero for a zero field, it is equal to 1 for a finite field, and it takes an intermediate value $m\in\,]-1,1[$ for an evanescent field. The variational approach of Ref. [41] focuses onto evanescent fields, and finds a *continuous* phase transition at $\alpha_s\simeq 3.96$ where the evanescent fields in different states start to cluster. However, as these are all evanescent fields, this means that the corresponding local magnetizations, in a given state, are not frozen to ± 1 but take some intermediate value, even in the $T\to 0$ limit. In our $T=0$ cavity approach (as well as in Ref. [42]), the HSP corresponds to a *discontinuous* transition at zero temperature, involving fields which are not evanescent, but are of order one [44]. This means that, in a given state, a finite number of local fields are nonzero integers, giving rise to magnetizations ± 1, as one could expect at zero temperature. This phenomenology cannot be found by considering evanescent fields. Its study with replicas would require using a more complicated Ansatz.

Note that our approach working directly at $T=0$ also has its limitations, for instance we are unable to determine precisely the self overlap (or the typical radius) of a state, or its internal entropy, precisely because we do not control the evanescent fields. We leave for future work the delicate study of the small T region of the phase diagram. Let us just notice here that a population dynamics study of this region with the 1RSB finite temperature Ansatz of Ref. [20] shows that the distributions of local fields tend to peak on integers when the temperature goes to zero in the HSP [45]. This is a strong argument in favor of the exactness of this 1RSB solution (i.e., the fact that we do not need to go to higher-order RSB), as argued in Ref. [21].

VIII. SURVEY PROPAGATION: CONFIGURATION SPACE ANALYSIS ON A SINGLE INSTANCE

The analysis of the iterative equations for the probability distributions of messages of the previous sections was done by using a population dynamics algorithm, which performs an average over the underlying random factor graphs. At each step of the iteration, a random choice of coupling constants, as well as neighboring nodes is performed with the proper probability distribution. While such an averaging step is central if one wants to estimate typical properties, the iterative equations make perfect sense on a single specific instance. The order parameters arising from the cavity equations, namely the u surveys, are histograms of probability distributions of cavity biases. They determine the bias of each spin in all metastable states of a given energy density for a given instance of the underlying factor graph. This is a very important piece of information which can be exploited to study specific problems and to invent new algorithms.

In the large N limit, we expect that the cavity assumptions hold for locally treelike factor graphs and we may use u-survey propagation to have access to the properties of optimal states of minimum energy. Here we shall develop one such application for random 3-SAT which is a concrete world-wide benchmark for search algorithms. However, the idea of exploiting the information on optimal states carried by the functional order parameter is rather general and we expect algorithmic applications in different fields.

Whenever the factor graph representing the problem does not lead to clustering within states, in practice whenever loops are short enough, one should think to the present approach as a first step of a sequence of possible approximations.

As is well known, the so called Cluster Variation Method [46] provides a systematic scheme that can be adopted to improve the approximate results given by the cavity approach [47]. The latter corresponds to the so called Bethe approximation which is the first step in the cluster variation scheme. Our present approach deals with the Bethe approximation in a frustrated case. While a great deal of work has been done concerning higher order cluster approximations for simple models, the corresponding analysis for frustrated systems such as spin-glasses or hard-combinatorial problems over nonlocally treelike graphs is largely unexplored.

A. The survey propagation (SP) algorithm

In the ordinary sum-product algorithms [22] as described in Sec. III, the messages arriving at a node are added up and then sent to a function node. Next, the function node transforms all input signals into a new message which is sent to the descendant variable node. At each time step, on the links of the factor graph there are signals traveling, just like in a communication network. SP works with the same principle, but now the messages traveling along the links of the factor graph are u surveys of usual messages over the various possible states of the system at a given value of the energy (or rather, in practice, at a given value of the reweighting parameter y). Of course this higher level of description is useful only when there are many states, which will be typically the case in hard optimization problems. One *a priori* drawback of the approach is that the messages are complicated, being functional probability distributions: a cavity bias is already a parametrization of a probability distribution (which turns out to be parametrized by a single variable in our case of binary spins, but could be more complicated in general); Here the u surveys are probability distribution functions of these cavity biases. In cases like K-SAT at $T=0$ in which the standard messages can take only few values (say r), a u survey is given by the probabilities of these values, i.e., by $r-1$ real numbers and the SP can be implemented easily. This is one big advantage of working directly at $T=0$, but we believe that the SP method could also be used more generally at finite temperature or with continuous variables, by using well adapted parametrizations of the cavity biases.

SP is defined for one given value of the reweighting parameter y and one given instance, with N variable nodes and M function nodes. Its basic ingredients are the u surveys. Each edge $a\to j$ from a function node to a variable node j carries a u survey $Q_{a\to j}(u)$. The algorithm finds these u surveys by a message passing procedure detailed below, and finds simultaneously all the h surveys $P_{i\to a}(h)$. Once these are known, one can compute the so called local field distributions and the zero-temperature free energy for this instance. The local field distribution $P_i(H)$ on a variable node i is the distribution, over all states selected by the reweighting parameter y, of the total local field H acting on spin σ_i [see Eq. (8)]. It is given by

$$P_i(H)=C_i\int\prod_{a\in V(i)}du_a\,Q_{a\to i}(u_a)\,\delta\times\left(H-\sum_{a\in V(i)}u_a\right)\exp\left(y\left|\sum_{a\in V(i)}u_a\right|\right) \tag{70}$$

C_i being the normalization constant.

We show in the Appendix B that the zero-temperature free-energy $\Phi(y)$ density of *this sample* can be computed as a sum of contributions $\Phi_a^f(y)$ for each function node a, corrected by the contributions $\Phi_i^v(y)$ from each variable node i, weighted by a factor n_i-1, where n_i is the connectivity of variable node i:

$$\Phi(y)=\frac{1}{N}\left(\sum_{a=1}^{M}\Phi_a^f(y)-\sum_{i=1}^{N}\Phi_i^v(y)(n_i-1)\right), \tag{71}$$

where

$$\Phi_a^f(y)=-\frac{1}{y}\ln\left\{\int\prod_{i\in V(a)}\left[\prod_{b\in V(i)-a}Q_{b\to i}(u_{b\to i})du_{b\to i}\right]\times\exp\left[-y\min_{\{\sigma_i,i\in V(a)\}}\left(E_a-\sum_{i\in V(a)}\left[\sum_{b\in V(i)-a}u_{b\to i}\right]\sigma_i\right)\right]\right\},$$

$$\Phi_i^v(y)=-\frac{1}{y}\ln\left\{\int\prod_{a\in V(i)}du_a\,Q_{a\to i}(u_a)\exp\left(y\left|\sum_{a\in V(i)}u_a\right|\right)\right\}=-\frac{1}{y}\ln(C_i). \tag{72}$$

The form (71) is the familiar one for the free energy in the Bethe approximation [18], and indeed one gets back the usual result using Eq. (72) in the $y\to 0$ limit. The generalization to $y\neq 0$ given in Eq. (72) adds the effect of the reweighting terms due to level crossings. The origin of these terms is exactly the same as in Sec. VI.

Let us now explain how SP works. We start with a general presentation of the algorithm, which applies to any optimization problem involving binary variables, characterized by a given factor graph. Some details of the implementation for the 3-SAT problem will be given below.

(1) All the u surveys $Q_{a\to i}(u)$ are initialized randomly.

(2) Function nodes are selected sequentially at random; for each such node a, we update the u surveys as follows (see Fig. 6):

(2.1) for each variable node i connected to the selected function node a, we compute the h survey $P_{i\to a}(h)$ as a reweighted convolution, see Fig. 6,

$$P_{i\to a}(h)=C_{i\to a}\int du_1\ldots du_k\,Q_{b_1\to i}(u_1)\ldots Q_{b_k\to i}(u_k)\times\delta\left(h-\sum_{a=1}^{k}u_a\right)\exp\left(y\left|\sum_{a=1}^{k}u_a\right|\right); \tag{73}$$

(2.2) successively, the u surveys on all edges $a\to i$ connected to a are updated using these h surveys:

$$Q_{a\to i}(u)=C_{a\to i}\int dg\,dh\,P_{j\to a}(g)P_{\ell\to a}(h)\times\delta(u-\hat{u}_J(g,h))\exp(y[\hat{w}_J(g,h)-|g|-|h|]). \tag{74}$$

($C_{i\to a}, C_{a\to i}$ are normalization constants).

(3) The iterative process of step (2) continues until convergence is reached. If the process converges, the corre-

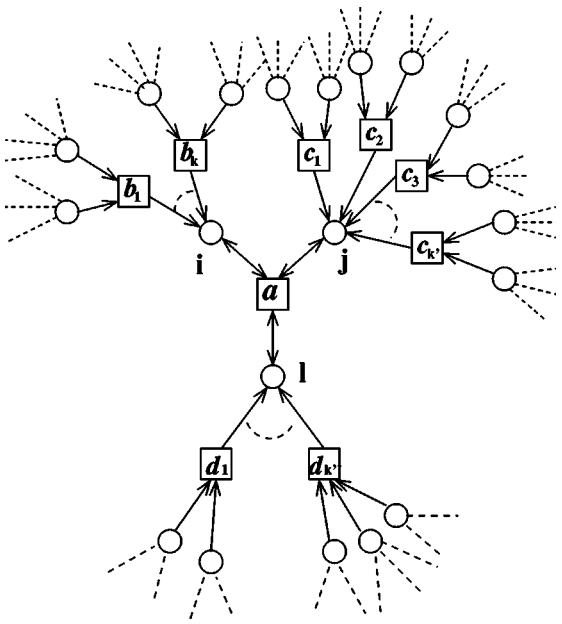

FIG. 6. Function node a and its neighboring graph.

sponding stable set of u surveys is used to compute the N local field distributions using Eq. (70), and the zero-temperature free energy $\Phi(y)$ given in Eqs. (71) and (72).

The above procedure can be repeated for different values of the reweighting y so that the complexity $\Sigma(y) = \partial\Phi(y)/\partial(1/y)$ and the energy density $\epsilon(y) = \partial(y\Phi(y))/\partial y$ of states can be estimated. The parametric plot of $\Sigma(y)$ versus $\epsilon(y)$, varying y, gives the complexity $\Sigma(\epsilon)$ of states of energy $E = N\epsilon$.

When it converges, SP allows to get an order parameter (the set of all the surveys), a zero-temperature free-energy density $\Phi(y)$, and a complexity curve $\Sigma(\epsilon)$ for one given instance. What is the meaning of these quantities in general is an open question. A given instance has a finite value of N, and therefore the notion of ''state'' is not easy to define. Roughly speaking, one can think that for large N, there might exist some effective finite N states, such that the number of spins to flip in order to reach one state from another one is large, leading to a separation of scales of the number of spins involved between the intrastate moves and the interstates moves. Such a situation would generally be difficult to handle for search algorithms, and this is where SP could be quite useful. In order to get a first understanding of these questions, we have experimented SP on single instances of the random 3-SAT problem for large values of N.

B. The case of random 3-SAT

We consider one instance of the 3-SAT problem, chosen randomly as in Sec. IV, with energy

$$E = 2\sum_{a=1}^{M} \prod_{j \in V(a)} \frac{1 + J_{a\to j}\sigma_j}{2}. \tag{75}$$

For 3-SAT the cavity-biases on a given link $a \to j$ takes values $u_{a\to j} = 0, -J_{a\to j}$; The corresponding survey is $Q_{a\to j}(u) = c_{a\to j}\delta(u) + (1 - c_{a\to j})\delta(u - J_{a\to j})$. The full set of u surveys is characterized by the $3M$ numbers $c_{a\to j}$ which are updated according to the SP algorithm described above, until convergence. The results of our numerical experiments are the following.

1. The paramagnetic phase

For $\alpha < \alpha_d$, SP converges toward the trivial paramagnetic solution $Q_{a\to i}(u) = \delta(u)$, for all $a \to i$ edges. Local field distributions are also trivial, $P_i(H) = \delta(H)\ \forall\ i$, and no information can be gained on the fine structure of ground states. In this region, there is a single state and the statistical properties of the zero energy configurations are totally driven by its entropy. A different formulation of the cavity approach, in which the proper $\beta \to \infty$ limit is taken and the evanescent fields are computed, could reveal some finer information for this paramagnetic phase, which however is known to be trivial from the algorithmic point of view.

2. The intermediate phase

For $\alpha_c > \alpha > \alpha_d$, that is in the glassy region, the random sequential updating of the iterative process converges to a unique nontrivial solution, provided y is large enough. In practice, we start from y large, like e.g., $y = 6$ (remember that the corrections to the $y \to \infty$ limit are exponentially small), run SP, and after finding a solution for the u surveys we decrease y (e.g., $y \to y - 0.2$) and rerun SP using the previous u surveys as a starting configuration for this new y. This speeds up the convergence. Below some value of y the nontrivial solution disappears abruptly and the algorithm converges to the paramagnetic solution.

In this region of α, the solution space, as well as the configurations of higher energy become divided into an exponential number of states. To compute the complexity, we measure the free energy $\Phi(y)$, see Fig. 7, and we perform the Legendre transform numerically.

The curve of the total complexity $N\Sigma$ versus the total energy $N\epsilon$ for one sample of random 3-SAT with $N = 10\,000$ and $M = 42\,000$ is given in Fig. 8. One finds $N\Sigma(E = 0) \sim 34$, meaning that the zero energy (SAT) states are predicted to be exponentially numerous, e^{34} at the leading exponential order [remember that each such state itself contains a large number of spin configurations [4]]. The threshold states have an energy of approximatively 44 violated clauses and their number is predicted to be about e^{216}. A cross check of such predictions is given by the behavior of ZTMA which cannot cross energetic barriers. It can be shown that for random 3-SAT zero-energy moves allow to explore configurations within each state and therefore, we expect such algorithms to get trapped in the most numerous ones (the threshold states). Indeed, extensive numerical simulations of ZTMA on many samples of different size (ranging from few hundreds to 10^5) and for different values of α confirm such scenario. As a representative example, we report that for the sample whose compexity is plotted in Fig. 8, repeated runs of ZTMA get stuck at an energy sharply peaked around 48 violated clauses, with a small residual dependence of the final energy on the simulation time (the final

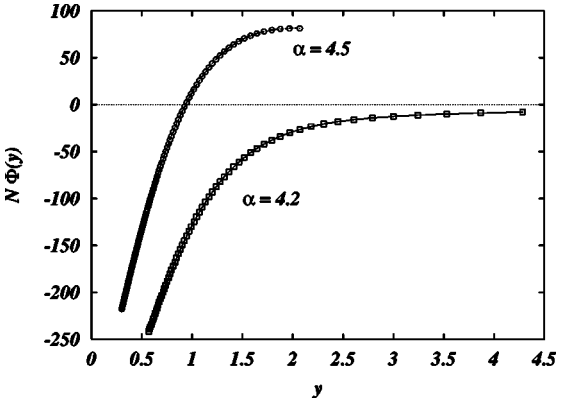

FIG. 7. Extensive zero temperature free energy $N\Phi(y)$ versus reweighting y for two specific instances of size $N=10\,000$ with $M=42\,000$ and $M=45\,000$ clauses, respectively. ($\alpha=4.2$ and $\alpha=4.5$).

energy found by ZTMA shows a power law behavior on the total number of spin flips).

We have checked that the functional order parameter given by the u surveys and the local field distributions carry precise information concerning the space of solutions for one given sample. Working with N not too large, some SAT configurations can be found efficiently by good algorithms like e.g., walkSAT-35 [48,49]. We have collected a large set (1000) of uncorrelated SAT configurations by running this algorithm many times with random initial conditions. In each such configuration ω, the spin σ_i takes a value σ_i^ω, and we have computed, for each given site i, the average of the whole set of SAT configurations ω. Next we have compared the above results with the predictions of SP as follows.

We first have selected the states with minimal energy, by picking the value of y which maximizes $F(y)$. Here this is $y=\infty$ and for practical computation it was enough to choose y sufficiently large (corrections are exponentially small). According to Eq. (70), the field $H_i=\Sigma_a u_{a\to i}$ in each state is an integer valued variable which can be computed from the u surveys. The total weights

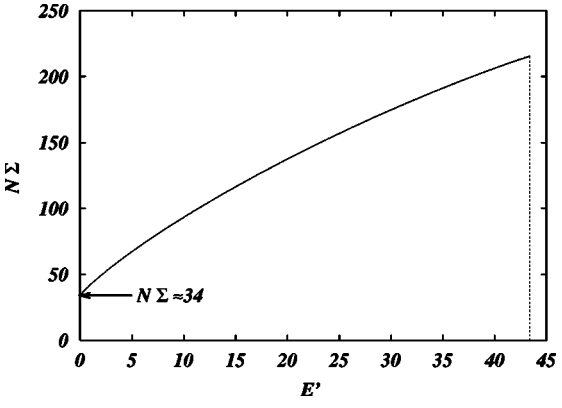

FIG. 8. Extensive complexity $N\Sigma$ versus the total number of violated clauses NE' ($=NE/2$), for the specific instance of size $N=10\,000$, $M=42\,000$ studied in Fig. 7. The complexity is obtained as the Legendre transform of the zero-temperature free energy.

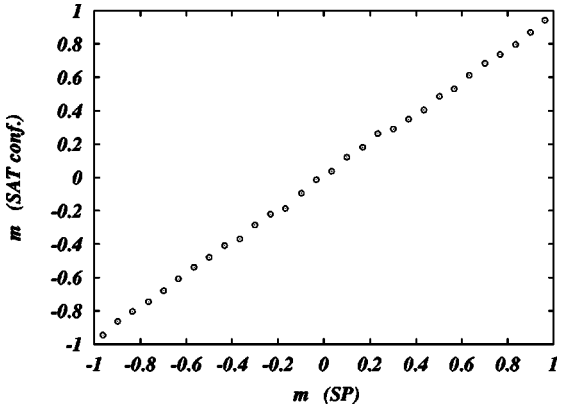

FIG. 9. The bias of the variables predicted by SP (with $y=8$) compared with the one measured analyzing SAT configurations from the same sample ($N=10\,000$, $M=42\,000$).

$$w_i^+=\int_{0^+}^{\infty} dH\, P_i(H); \quad w_i^-=\int_{-\infty}^{0^-} dH\, P_i(H) \tag{76}$$

of $P_i(H)$ on positive (negative) integers give the fractions of zero energy states where σ_i is fixed to 1 (to -1). As displayed in Fig. 9, we find a remarkable agreement between the local magnetizations $w_i^+-w_i^-$ predicted by SP and the local magnetizations measured by averaging over the ground states found by the walkSAT algorithm. In the figure we report data for $N=10\,000$ and $M=42\,000$: we have divided the local magnetization in 30 intervals and labeled spins according to the prediction of SP. Next on such a partitioning of spins we have taken the average over the configurations found by walkSAT. (The remarkable agreement of numerical and SP results indirectly shows that walkSAT is a good uniform sampler.)

The weight in $H=0$ of $P_i(H)$, $w_i^0=1-w_i^+-w_i^-$, measures the tendency of a variable to be under constrained: for instance, variables which belong to very few clauses have $w_i^0=1$.

3. The UNSAT phase

For $\alpha>\alpha_c$, SP predicts a positive ground state energy with zero complexity, whereas excited states remain exponentially numerous. Proper tuning of the reweighting, that is choosing y so that the complexity vanishes, allows to predict the ground state energy and to evaluate the probability distribution of effective fields for each variable. In this regime SP is found to converge only when the reweighting parameter is well chosen. For small values of y, SP converges to the paramagnetic solution or to the RS solution. For intermediate values of y, SP converges to the nontrivial solution whereas for larger values of y, SP stops converging. The range of y values for which SP converges to the nontrivial solution is sufficient to determine the free energy. An example is given in Fig. 7.

For large values of α, we expect multiple nested clustering phenomena to appear, that is continuous replica symmetry breaking [50]. This scenario could be analyzed by a fur-

ther generalization of SP which is beyond the scope of this work.

IX. SURVEY PROPAGATION AS A SOURCE OF NEW ALGORITHMS FOR HARD OPTIMIZATION PROBLEMS

The preceding section has shown how SP can give rather precise answers on the structure of the space of configurations and the ground state energy of the random 3-SAT problem. Here we shall stay within this problem and ask the following natural question: Given a random 3-SAT formula of size N, how can we take advantage of SP in order to find optimal configurations?

If SP could predict with very high accuracy the value of the ground state (g.s.) energy of a given formula, it could also predict its satisfiability. Then one could proceed in finding a satisfying assignment just by converting the decision algorithm into a search algorithm as follows [2]. A variable is selected and fixed to one value. We then use SP to evaluate the g.s. energy of the subproblem of size $N-1$ and decide whether it is still SAT or not. If the subproblem is SAT then we keep the assignment, otherwise the opposite value of the binary variable is chosen. The process is repeated untill all the variables have been exhausted (in at most $2N$ steps). If along this reduction process the subsystem becomes a paramagnet, then SP becomes ineffective and another search algorithm must be used on the subsystem (but for a paramagnet it is very easy to find the ground state).

The above scheme, however, suffers from finite size effects and from the imprecision in the determination of the ground state energy, a fact which is particularly important close to α_c. Moreover, it does not take advantage of the information provided by the u surveys.

A. Categories of variables in one specific instance

Somewhat coarse grained information contained in the u surveys, once SP has reached convergence, is that given by the total weights $w_i^{\pm}$ of the local field distribution which gives the fraction of states where the spin σ_i is positive (negative). Having computed these weights, we may distinguish three reference types of spins (of course all the intermediate cases will also be present): the *paramagnetic* ones with $w_i^0 \sim 1$, the *biased* ones with $w_i^+ \sim 1$ or $w_i^- \sim 1$ and the *balanced* ones with $w_i^+ \simeq w_i^-$ and w_i^0 small.

In order to characterize the differences between these various types of variables, we have performed a few numerical experiments and analyzed the effect of fixing one such spin on the structure of states of the subproblem of size $N-1$. As displayed in the complexity curve of Fig. 10, the three types of spins produce different effects, consistently with the interpretation of the order parameter. Fixing a biased spin does not alter the structure of the states and the complexity changes smoothly. Fixing a paramagnetic spin has an effect only on the internal entropy of the states (which we cannot measure) but leaves the energy unaltered. Interestingly enough, balanced spin have an enormous effect: the most balanced ones produce a decrease very close to $\ln 2$ in the complexity; indeed half of the states are eliminated by fixing one single balanced variable!

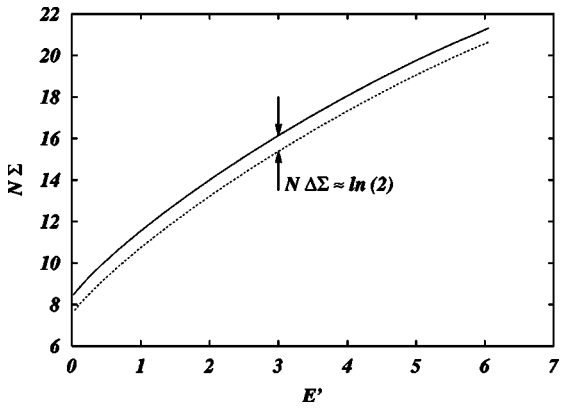

FIG. 10. Effect of fixing a single balanced spin on the complete complexity curve $N\Sigma$ versus the number of violated clauses (i.e., one half of the extensive energy) of an instance of 3-SAT with $N=1000$ and $M=4200$. The difference in the two curves is very close to $\ln 2$.

B. Survey inspired decimation (SID) algorithm

One strategy for using this information in order to produce an optimization algorithm is to fix as many variables as possible without altering the ground state energy, evaluated step by step as the size of the problem decreases. Eventually, either all variables have been fixed or (more likely) the remaining variables turn out to be paramagnetic [i.e., $P_i(H)=\delta(H), \ \forall \ i$], in which case a simple search process can be run to find the complete ground state configuration.

A straightforward implementation of the above ideas provides a simple algorithm that can be used to find solutions to random 3-SAT in the hard region $\alpha \in [\alpha_d, \alpha_c]$. We do not expect this implementation to be the most efficient one, in that no particular strategy has been worked out to optimize the decimation process. The scope of this first implementation consists in showing the potentiality of the novel algorithmic scheme and we leave for future work the design of optimized versions of the algorithm or applications in different contexts.

The overall idea underlying the search process is rather simple. At each time step a single variable is fixed according to the outcome of SP and the effect of such fixing is used to simplify the problem. The size of the problem reduces from N_t to N_t-1-S_t, where S_t is the number of variables which become fixed due to the simplification of the problem: satisfied clauses are eliminated, unsatisfied K clauses are transformed into $(K-1)$ clauses. $K=1$ clauses need to be satisfied and therefore their variables are properly fixed (unit clause propagation) leading to further spin elimination.

At the beginning of the process, randomly chosen balanced spins can be fixed in order to reduce the number of states. At each step one may compute the free energy to detect the onset of violated clauses. One may also evaluate the function $\Phi(y)$ to have an estimate of the complexity. Successively, biased spins are fixed. Whenever a paramagnetic state is found, or at some intermediate steps, a rapid search process like simulated annealing at a fixed cooling

rate or walkSAT is run on the subsystem. We may end up either by having found a solution or by having still few violated clauses. In the latter case we may simply restart. The sketch of the SID algorithm is as follows.

(1) Random initial condition for the cavity biases.

(2) Run SP and evaluate $\{P_i(H)\}$, or $\{w_i^+, w_i^-, w_i^0\}$, and $\Phi(y)$.

(3) Check for a paramagnetic state and in case (or at some intermediate step) run a fast local search process (e.g., simulated annealing or walkSAT). If a solution is found output ''SAT'' and stop.

(4) Select and fix the most biased variable (the one with the largest $|w_i^+ - w_i^-|$) and simplify the problem.

(5) If the problem is solved completely by unit clause propagation, then output ''SAT'' and stop. If no contradiction is found then continue the decimation process on the smaller problem (go to 1.) else (if a contradiction is reached) restart (go to 0.)

Extensive numerical experiments on random 3-SAT instances at $\alpha = 4.2$ with size up to $N = 10^5$ have shown a remarkable efficiency of SID. While the process of fixing a single variable takes some time [$O(N)$ operations] the number of assignments explored is very small. At $\alpha = 4.2$ typically a single run of SID (i.e., with no restarts) leads to a solution. Closer to the critical α, few restarts might be necessary in order to find a configuration of strictly zero energy. However, at each run the typical energy found by SID is very close to zero, well below the energy at which simulated annealing gets stuck. A detailed description of the numerical experiments will be given in a forthcoming paper [51]. We just mention that the largest public benchmarks of random 3-SAT [49] have been solved efficiently by SID.

In Fig. 11 we show the evolution of the complexity under SID. For a sample of size $N = 10\,000$ at $\alpha = 4.2$ we evaluate the complexity curve every 200 decimation steps until a paramagnetic state is reached. SID acts by eliminating clusters of solutions and hence reducing the complexity of the ground state down to the point where very few clusters remain.

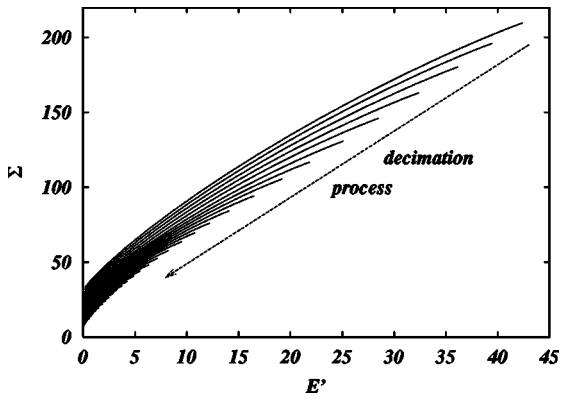

FIG. 11. Evolution of the complexity curve upon decimation: The Σ versus $E' = E/2$ curve is shown when 200,400,600, . . . spins have been fixed according to the SID algorithm ($N = 10\,000$, $M = 42\,000$ for this sample).

X. CONCLUSION

We have derived here two main results. The first one concerns the phase diagram of the random K-SAT problem, and establishes the existence of an intermediate phase where the problem is SAT but the solution is difficult to find because of the existence of many states. We would like to point out that the cavity method which we have used here is not rigorous: it relies on some hypotheses which can be true only for large systems and are thus difficult to prove (although this can be done in some cases [29]). However, the experience gained from similar problems, together with numerical results of this and previous papers, indicates that this solution is likely to be the correct one. If higher order replica symmetry breaking effects would show up, one can believe that in any case their quantitative influence on the results should be rather small. It should also be noted that the very same cavity strategy which we have used here has been tested on a variant of the K-SAT, the XORSAT problem, which can be also solved by exact methods [52,53]. In this case, the existence of the intermediate phase has been confirmed and all the predictions (qualitative and quantitative) from the cavity method have been checked rigorously [52].

The second result gives a new class of message passing algorithms for solving optimization problems in the regime where the proliferation of metastable states slows down a lot all the local algorithms. One such algorithm turns out to be quite powerful on the case of random 3-SAT problems. Clearly a lot of work needs to be done in order to develop such algorithms in various contexts and test them against traditional strategies. A more direct derivation and understanding of the algorithm, and in particular of the nontrivial reweighting term, would also be welcome.

Finally, we expect that the *single sample* SP *method* at finite temperature can become a useful tool in analyzing the fine structure of order parameters in disordered systems and other complex systems.

ACKNOWLEDGMENTS

We thank G. Parisi, B. Selman, and J. S. Yedidia for useful discussions and encouragement and S. Mertens for pointing out the inaccuracy in our previous estimate of α_c. R.Z. thanks the LPTMS, Orsay, for their kind hospitality.

APPENDIX A: LARGE y EXPANSION FOR THE 3-SAT PROBLEM

We give here some details on the solution of the population dynamics equation (56) of the 3-SAT problem at large y. We start from the self-consistent equation (63) for the distribution of the rescaled weights of the u surveys in $u = 1$, and we recall that $f_q^{(m,n)}(\eta)$ is defined as the probability that $u_1 + \cdots + u_k = q$, given that $Q_1(u_1), \ldots, Q_m(u_m)$ are of type b_+, and $Q_{m+1}(u_{m+1}), \ldots, Q_{m+n}(u_{m+n})$ are of type b_-. The quantities A_0, B_0 are expressed in terms of three numbers g_+, g_-, g_0:

$$g_+(\eta) \equiv \sum_{q=1}^{k} f_q^{(m,n)}(\eta) e^{y|q|},$$

$$g_-(\eta)\equiv\sum_{q=-k}^{-1} f_q^{(m,n)}(\eta)e^{y|q|},$$

$$g_0(\eta)\equiv f_0^{(m,n)}(\eta) \tag{A1}$$

as

$$A_0\equiv g_+(\eta)g_+(\eta'),$$

$$B_0\equiv g_+(\eta)[g_-(\eta')+g_0(\eta')]+g_+(\eta')[g_-(\eta)+g_0(\eta)] +[g_-(\eta')+g_0(\eta')][g_-(\eta)+g_0(\eta)]. \tag{A2}$$

The following step consists in introducing the joint probability distribution $P^{(m,n)}(g_+,g_-,g_0)$,

$$P^{(m,n)}(g_+,g_-,g_0)\equiv\int d\eta_1\cdots d\eta_{m+n}\rho(\eta_1)\cdots\rho(\eta_{m+n}) \times\delta\left(g_+-\sum_{q=1}^{k} f_q^{(m,n)}(\eta)e^{y|q|}\right) \times\delta\left(g_--\sum_{q=-k}^{-1} f_q^{(m,n)}(\eta)e^{y|q|}\right) \times\delta(g_0-f_0^{(m,n)}(\eta)). \tag{A3}$$

Equation (63) reads

$$\rho(w_0)=\frac{1}{1-t}\sum_{k=0}^{\infty}\sum_{m=1}^{k}\sum_{n=0}^{k-m} C_{k,m,n}\sum_{k'=0}^{\infty}\sum_{m'=1}^{k'}\sum_{n'=0}^{k'-m'} C_{k',m',n'} \times\int dg_+\,dg_-\,dg_0\,dg'_+\,dg'_-\,dg'_0 \times P^{(m,n)}(g_+,g_-,g_0)P^{(m',n')}(g'_+,g'_-,g'_0) \times\delta\left(w_0-\frac{A}{A+B}\right), \tag{A4}$$

where the coefficients $C_{k,m,n}$ are given in Eq. (60), and A,B are given in Eq. (62).

The scaling of the weights of the u surveys in $u=1$, w_0, with e^{-y} at large y is consistent with the self-consistency equation (63). In this limit we find $g_+,g_-,g_0\sim O(1)$, and these quantities simplify to (up to corrections of order e^{-y})

$$g_+(\eta)=(1+\eta_1)\cdots(1+\eta_m)-1,$$

$$g_-(\eta)=(1+\eta_{m+1})\cdots(1+\eta_{m+n})-1,$$

$$g_0(\eta)=1. \tag{A5}$$

Therefore, the three variables g_+,g_-,g_0 become uncorrelated random variables in the large y limit, with a distribution:

$$P^{(m,n)}(g_+,g_-,g_0)=\Omega^{(m)}(g_+)\Omega^{(n)}(g_-)\,\delta(g_0-1), \tag{A6}$$

where

$$\Omega^{(\ell)}(g)=\int d\eta_1\cdots d\eta_\ell\hat{\rho}(\eta_1)\cdots\hat{\rho}(\eta_\ell) \times\delta\left(g-\left[\prod_{i=1}^{\ell}(1+\eta_i)-1\right]\right). \tag{A7}$$

The equation for $\rho(\eta_0)$ reads

$$\rho(\eta_0)=\frac{1}{1-t}\sum_{k=0}^{\infty}\sum_{m=1}^{k}\sum_{n=0}^{k-m} C_{k,m,n}\sum_{k'=0}^{\infty}\sum_{m'=1}^{k'}\sum_{n'=0}^{k'-m'} C_{k',m',n'} \times\int dg_+\,dg_-\,dg'_+\,dg'_-\,\Omega^{(m)}(g_+) \times\Omega^{(n)}(g_-)\Omega^{(m')}(g'_+)\Omega^{(n')}(g'_-)\,\delta\left(\eta_0 -\frac{g_+g'_+}{g_+(1+g'_-)+g'_+(1+g_-)+(1+g_-)(1+g'_-)}\right). \tag{A8}$$

Clearly, the function $\Omega^{(\ell)}(g)$ can be seen as the ℓth convolution of a certain function after an appropriate change of variables. One is led to introduce the variables ϕ_i and x defined by

$$\phi_i=\ln(1+\eta_i),\quad x=\ln(1+g), \tag{A9}$$

and we call $S(\phi_i)$, $T^{(m)}(x)$ their probability distributions. Equation (A7) shows that

$$T^{(m)}(x)=\int d\phi_1\ldots d\phi_m\,S(\phi_1)\ldots S(\phi_m) \times\delta\left(x-\sum_{i=1}^{m}\phi_i\right), \tag{A10}$$

where $\phi_i\in[0,[\infty[$ and $x\in[0,m[\infty[$.

In order to simplify the self-consistency equation, we introduce the joint probability distribution

$$R(x_+,x_-)\equiv\frac{1}{\sqrt{1-t}}\sum_{k}\sum_{m=1}^{k}\sum_{n=0}^{k-m} C_{k,m,n}T^{(m)}(x_+)T^{(n)}(x_-), \tag{A11}$$

which is normalized: $\int dx_+\,dx_-\,R(x_+,x_-)=1$.

We may show that R factorizes by introducing its Fourier Transform: Using the coefficients (60), the triple series can be resummed and expressed in terms of the Fourier transform $\hat{S}(q)$ of $S(x)$:

$$\int dq_+\, dq_-\, R(x_+,x_-)e^{i(q_+x_+ + q_-x_-)}$$
$$= \frac{e^{-3\alpha}}{\sqrt{1-t}} \sum_k \frac{(3\alpha)^k}{k!} \left(\left[t + \frac{1-t}{2}\hat{S}(q_+) + \frac{1-t}{2}\hat{S}(q_-) \right]^k - \left[t + \frac{1-t}{2}\hat{S}(q_-) \right]^k \right)$$
$$= \frac{1}{\sqrt{1-t}} \exp\left[3\alpha \left(t - 1 + \frac{1-t}{2}\hat{S}(q_+) + \frac{1-t}{2}\hat{S}(q_-) \right) \right] - \frac{1}{\sqrt{1-t}} \exp\left[3\alpha \left(t - 1 + \frac{1-t}{2}\hat{S}(q_-) \right) \right]. \tag{A12}$$

Rearranging the above expression and taking the inverse transformation we find for R

$$R(x_+,x_-) = A(x_+)B(x_-), \tag{A13}$$

where

$$A(x_+) \equiv \frac{1}{e^{3\alpha(1-t)/2} - 1} \int \frac{dq_+}{2\pi} e^{-iq_+x_+} (e^{(3\alpha/2)(1-t)\hat{S}(q_+)} - 1),$$

$$B(x_-) \equiv \frac{1}{e^{3\alpha(1-t)/2}} \int \frac{dq_-}{2\pi} e^{-iq_-x_-} e^{(3\alpha/2)(1-t)\hat{S}(q_+)}. \tag{A14}$$

We may now write the self-consistency equation in a tractable form. Defining the variable ϕ_0 associated with η_0 as

$$\eta_0 = \frac{(e^{x_+}-1)(e^{x'_+}-1)}{(e^{x_+}-1)e^{x'_-} + (e^{x'_+}-1)e^{x_-} + e^{x_-}e^{x'_-}} = e^{\phi_0} - 1 \tag{A15}$$

we transform the equation for $\hat{\rho}(\eta_0)$ into an equation for $S(\phi_0)$

$$S(\phi_0) = \int dx_+\, dx_-\, dx'_+\, dx'_-\, R(x_+,x_-)R(x'_+,x'_-) \times \delta\left(\phi_0 - \ln\left[1 + \frac{(e^{x_+}-1)(e^{x'_+}-1)}{(e^{x_+}-1)e^{x'_-} + (e^{x'_+}-1)e^{x_-} + e^{x_-}e^{x'_-}} \right] \right). \tag{A16}$$

This is the equation that we have used in order to solve the problem numerically.

Let us mention however that a series of further simplifications may also be written. It is easy to verify that

$$S(\phi_0) = \int dz\, dz'\, C(z)C(z')\delta\left(\phi_0 - \ln\left[1 + \frac{1}{z + z' + zz'} \right] \right), \tag{A17}$$

where

$$C(z) \equiv \int dx_+\, A(x_+)\, dx_-\, B(x_-)\, \delta\left(z - \frac{e^{x_-}}{e^{x_+}-1} \right). \tag{A18}$$

Moreover, if we perform in the above the change of variable $1+z = e^{\zeta}$, defining the distribution

$$D(\zeta) = C(z)\frac{dz}{d\zeta}, \tag{A19}$$

we find

$$S(\phi_0) = \int d\eta\, E(\eta)\delta\left(\phi_0 - \ln\left[\frac{e^{\eta}}{e^{\eta}-1} \right] \right), \tag{A20}$$

where $E(\eta)$ is the convolution of D with itself

$$E(\eta) \equiv \int d\zeta\, d\zeta'\, D(\zeta)D(\zeta')\delta(\zeta + \zeta' - \eta). \tag{A21}$$

The transformations which we have written from $S \to A, B \to C \to D \to E \to S$ can all be done using one dimensional integrals, Fourier transforms, and changes of variables. Hence this provides an iterative mapping from the function $S(\phi)$ onto itself which can be handled efficiently numerically. Which form of the self-consistency equation to use is a matter of computational convenience. It turns out that, for our purpose, enough precision could be obtained from Eq. (A17) and we did not try to develop this alternative computation.

We now proceed with the computation of the energy $\Phi(y)$ defined in Eq. (48). We shall be interested in evaluating it at large y, using the solution for the distribution of u surveys that we have just found in this limit.

We start with the piece $\Phi_1(y)$. It consists of two terms

$$\Phi_1(y) = \frac{3\alpha}{y}\overline{\ln C_0} - \frac{1}{y}\overline{\ln A_0}, \tag{A22}$$

where the overline denotes the average over the population dynamics of Sec. VI B 4 and where C_0 and A_0 are given by

$$A_0 = \int \prod_{\ell=1}^{k} [du_\ell\, Q_\ell(u_\ell)] \exp\left[y \left| \sum_{\ell=1}^{k} u_\ell \right| \right], \tag{A23}$$

$$\frac{1}{C_0} = \int dg\, dh\, P_1(g)P_2(h)\exp[y\hat{w}_J(g,h)]. \tag{A24}$$

The term A_0 is easily written in terms of the variables g_+, g_-, g_0,

$$A_0 = \sum_q f_q^{(m,n)}(\eta)e^{y|q|} = g_+(\eta) + g_-(\eta) + g_0(\eta). \tag{A25}$$

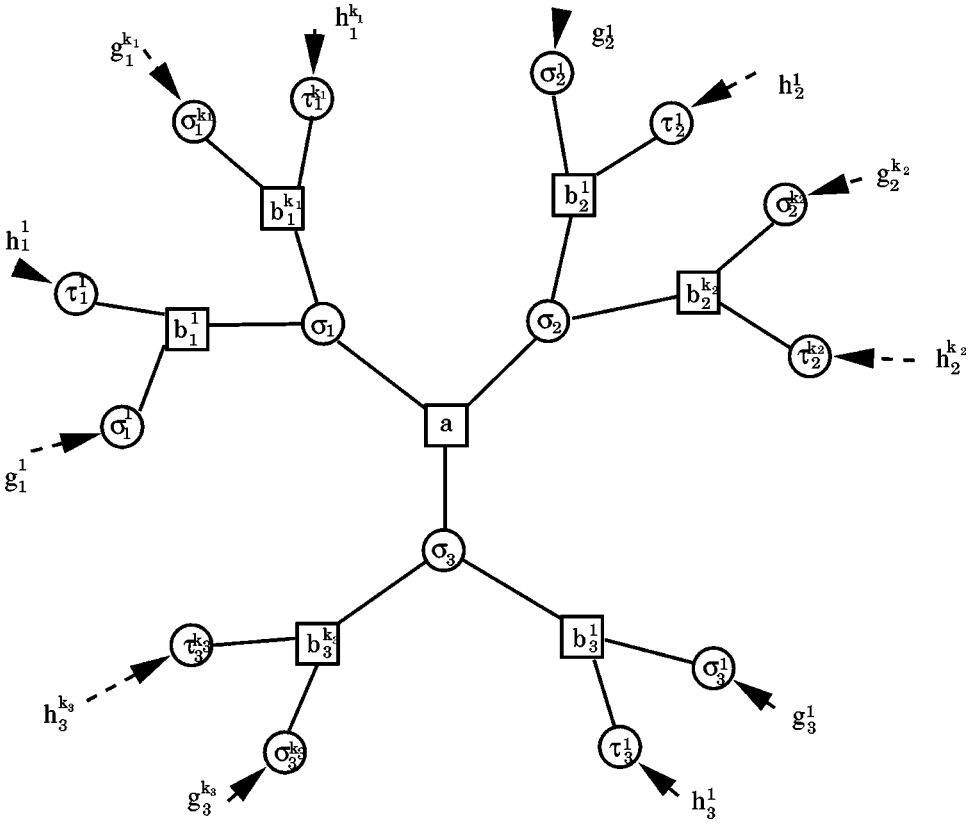

FIG. 12. The set of nodes with which one computes the free-energy shift $\Phi_a^f(y)$.

Averaging over the population dynamics, we get for large y

$$\frac{1}{y}\overline{\ln A_0} = \frac{1}{y}\sum_{k,m,n} C_{k,m,n}\int dg_0\, dg_+\, dg_-\, P^{(m,n)}(g_+,g_0,g_-) \times \ln[g_0+g_++g_-]$$
$$= \frac{1}{y}\sum_{k,m,n} C_{k,m,n}\int dx\, dz\, T^{(m)}(x)T^{(n)}(z) \times \ln(e^x+e^z-1). \quad \text{(A26)}$$

Treating separately the $m\geqslant 1$ piece, which can be resummed as in Eqs. (A11) and (A13), and the $m=0$ piece, the sum over k,m,n gives

$$\sum_{k,m,n} C_{k,m,n}T^{(m)}(x)T^{(n)}(z) = \sqrt{1-t}A(x)B(z)+\delta(x)\sum_{k,n} C_{k,0,n}T^{(n)}(z) = B(x)B(z). \quad \text{(A27)}$$

Putting this expression back into Eq. (A26), we finally find

$$\frac{1}{y}\overline{\ln A_0} = \frac{1}{y}\int dx\, dz\, B(x)B(z)\ln(e^x+e^z-1). \quad \text{(A28)}$$

We now turn to the second contribution, C_0, to Eq. (A24); before averaging over the iteration of the population dynamics, we have

$$\frac{1}{C_0} = \int \prod_{i=1}^{k}[du_i\, Q_i(u_i)]\prod_{j=1}^{k'}[dv_j\, Q_{k+j}(v_j)] \times \exp\left[y\hat{w}_{\mathbf{J}}\left(\sum_{i=1}^{k}u_i,\sum_{j=1}^{k'}v_j\right)\right]. \quad \text{(A29)}$$

Without loss of generality we assume $J_1=J_2=1$; we use the solution (57), and denote as before by m,n,m',n' the numbers of various u surveys appearing in Eq. (A29)

$k\to m$ type b_+, n type b_- and $k-m-n$ type a,

$k'\to m'$ type b_+, n' type b_- and k'-$m'-n'$ type a.

C_0 can then be written in terms of the g variables as

$$\frac{1}{C_0} = \int dg_0\, dg_+\, dg_-\, P^{(m,n)}(g_+,g_0,g_-) \times dg_0'\, dg_+'\, dg_-'\, P^{(m',n')}(g_+',g_0',g_-') \times (g_+g_+'e^{-y}+g_+[g_-'+g_0'] + g_+'[g_-+g_0]+[g_-'+g_0'][g_-+g_0]). \quad \text{(A30)}$$

For y large we can drop the term in $\exp(-y)$; averaging over the iteration of the population, we get

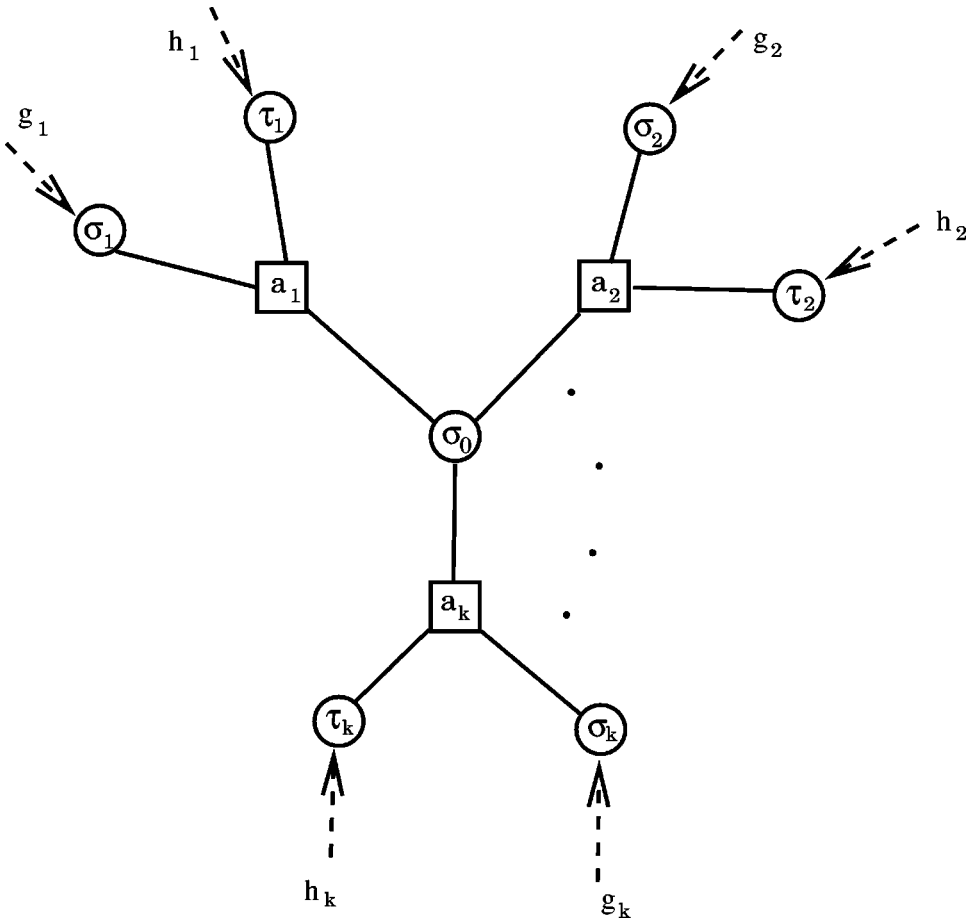

FIG. 13. The set of nodes with which one computes the free energy shift $\Phi_0^v(y)$.

$$
\begin{aligned}
-\overline{\ln C_0} = & \sum_{k,m,n} \sum_{k',m',n'} C_{k,m,n} C_{k',m',n'} \\
& \times \int dx\, dz\, T^{(m)}(x) T^{(n)}(z) \\
& \times \int dx'\, dz'\, T^{(m')}(x') T^{(n')}(z') \\
& \times \ln[(e^x-1)e^{z'}+(e^{x'}-1)e^z+e^z e^{z'}] \\
= & \int dx\, dz\, dx'\, dz'\, B(x) B(z) B(x') B(z') \\
& \times \ln[(e^x-1)e^{z'}+(e^{x'}-1)e^z+e^z e^{z'}].
\end{aligned}
\tag{A31}
$$

Finally we now compute the $\Phi_2(y)$ piece given in Eq. (48). For a generic clause with coupling J_1, J_2, J_3, involving the h surveys $P_1(h_1)$, $P_2(h_2)$, $P_3(h_3)$, we have

$$
\begin{aligned}
\exp[-y\Phi_2(y)] = & \int dh_1\, dh_2\, dh_3\, P_1(h_1) P_2(h_2) \\
& \times P_3(h_3) e^{-2y\theta(J_1 h_1)\theta(J_2 h_2)\theta(J_3 h_3)}.
\end{aligned}
\tag{A32}
$$

We write as before $P_1(h_1) = \int \Pi_{\ell=1}^{k_1} [du_\ell\, Q_l(u_\ell)]$, and suppose that this set of k_1 u surveys contains m_1 u surveys of type b_+, n_1 of type b_- and $k_1 - m_1 - n_1$ of type a, characterized by the weights η_1. The same decomposition is done for $P_2(h_2)$ [respectively $P_3(h_3)$], where the numbers of u surveys of various types are $m_2, n_2, k_2 - m_2 - n_2$ and the weights are η_2 (respectively $m_3, n_3, k_3 - m_3 - n_3, \eta_3$). For this iteration, the expression (A32) for $\Phi_2(y)$ can be reexpressed as

$$
\begin{aligned}
& \exp[-y\Phi_2(y)] \\
& = \sum_{q_1,q_2,q_3} f_{q_1}^{(m_1,n_1)}(\eta_1) f_{q_2}^{(m_2,n_2)}(\eta_2) f_{q_3}^{(m_3,n_3)}(\eta_3) \\
& \quad \times \exp[-2y\,\theta(q_1)\theta(q_2)\theta(q_3)] \\
& \simeq_{y\to\infty} \prod_{i=1}^{3} [g_+(\eta_i) + g_-(\eta_i) + g_0(\eta_i)] \\
& \quad - \prod_{i=1}^{3} [g_+(\eta_i)].
\end{aligned}
\tag{A33}
$$

The average over the population gives

$$\overline{\Phi_2(y)} = -\frac{1}{y} \sum_{k_1,m_1,n_1} \sum_{k_2,m_2,n_2} \sum_{k_3,m_3,n_3} C_{k_1,m_1,n_1} C_{k_2,m_2,n_2} \times C_{k_3,m_3,n_3} \prod_{i=1}^{3} \left\{ \int dx_i\, dz_i\, T^{(m_i)}(x_i) T^{(n_i)}(y_i) \times \ln\left[\prod_{i=1}^{3} (e^{x_i} + e^{z_i} - 1) - \prod_{i=1}^{3} (e^{x_i} - 1) \right] \right\}$$

$$= -\frac{1}{y} \int \prod_{i=1}^{3} dx_i\, dz_i\, B(x_i) B(z_i) \times \ln\left[\prod_{i=1}^{3} (e^{x_i} + e^{z_i} - 1) - \prod_{i=1}^{3} (e^{x_i} - 1) \right]. \quad \text{(A34)}$$

Grouping together the contributions (A22),(A28), (A31),(A34), we find the total zero-temperature free-energy density $\Phi(y) = \Phi_1(y) - 2\alpha\Phi_2(y)$ given in Eq. (67).

APPENDIX B: FREE ENERGY FOR ONE GIVEN SAMPLE

Let us explain here how to compute the zero-temperature free energy for one given sample. We start from the contribution of one given factor node a. We shall look at a somewhat large part of the graph containing a (see Fig. 12). We call $\sigma_1, \sigma_2, \sigma_3$ the three spins connected to it. One of these spins, σ_r is connected, beside a, to k_r other function nodes which we call $b_r^1, \ldots, b_r^{k_r}$. The function node b_r^s is connected, beside σ_r, to two other spins which we call σ_r^s and τ_r^s, and the cavity fields onto them are called g_r^s and h_r^s (see Fig. 12). In the absence of the spins $\sigma_1, \sigma_2, \sigma_3$ and of all the function nodes b_r^s, the ground state energy of the system would be

$$E_{init} = -\sum_{r=1}^{3} \sum_{s=1}^{k_r} (|g_r^s| + |h_r^s|). \quad \text{(B1)}$$

Adding the spins $\sigma_1, \sigma_2, \sigma_3$ and all the function nodes b_r^s, the ground state energy becomes

$$E_{fin} = \min_{\sigma_1,\sigma_2,\sigma_3} \left(E_a(\sigma_1,\sigma_2,\sigma_3) + \sum_{r=1}^{3} \sum_{s=1}^{k_r} \{ \min_{\sigma_r^s,\tau_r^s} [E_{b_r^s}(\sigma_r,\sigma_r^s,\tau_r^s) - g_r^s\sigma_r^s - h_r^s\tau_r^s] \} \right)$$

$$= \min_{\sigma_1,\sigma_2,\sigma_3} \left\{ E_a(\sigma_1,\sigma_2,\sigma_3) - \sum_{r=1}^{3} \sigma_r \sum_{s=1}^{k_r} \hat{u}_{\mathbf{J}}(g_r^s, h_r^s) \right\} - \sum_{r=1}^{3} \sum_{s=1}^{k_r} (\hat{w}_{\mathbf{J}}(g_r^s, h_r^s) - |g_r^s| - |h_r^s|). \quad \text{(B2)}$$

The zero-temperature free-energy shift induced by the addition of all these nodes is given by

$$e^{-y\Phi_a^f(y)} = \int e^{-y(E_{fin} - E_{init})}, \quad \text{(B3)}$$

where the integral is over all the g_r^s, h_r^s fields, each with a probability distribution given by its h survey. We can use the iteration equation (74) in order to simplify this complicated integral through the use of u surveys

$$e^{-y\Phi_a^f(y)} = \int \prod_{r=1}^{3} \prod_{s=1}^{k_r} \left[\frac{1}{C_{b_r^s \to r}} du_r^s\, Q_{b_r^s \to r}(u_r^s) \right] \times \exp\left(-y \min_{\sigma_1,\sigma_2,\sigma_3} \left\{ E_a(\sigma_1,\sigma_2,\sigma_3) - \sum_{r=1}^{3} \sigma_r \sum_{s=1}^{k_r} u_r^s \right\} \right). \quad \text{(B4)}$$

We now compute the contribution from one given variable node $i=0$. We use the notations of Fig. 13, calling a_r the k function nodes to which it is connected ($r \in \{1, \ldots, k\}$). The function node a_r is connected, beside σ_0, to the spins σ_r, τ_r, and we call g_r (respectively h_r) the cavity field on σ_r (respectively τ_r). In the absence of spin σ_0 and of the function nodes connected to it, the ground state energy of the system would be

$$E_{init} = -\sum_{r=1}^{k} (|g_r| + |h_r|). \quad \text{(B5)}$$

Adding the new spin and the function nodes a_r, the ground state energy becomes

$$E_{fin} = \min_{\sigma_0} \sum_{r=1}^{k} \{ \min_{\sigma_r,\tau_r} [E_{a_r}(\sigma_0,\sigma_r,\tau_r) - g_r\sigma_r - h_r\tau_r] \} = \min_{\sigma_0} \left(-\sigma_0 \sum_{r=1}^{k} \hat{u}_{\mathbf{J_r}}(g_r,h_r) \right) - \sum_{r=1}^{k} \hat{w}_{\mathbf{J_r}}(g_r,h_r) = -\left| \sum_{r=1}^{k} \hat{u}_{\mathbf{J_r}}(g_r,h_r) \right| - \sum_{r=1}^{k} \hat{w}_{\mathbf{J_r}}(g_r,h_r). \quad \text{(B6)}$$

The zero temperature free energy shift induced by the addition of these nodes is given by

$$e^{-y\Phi_0^v(y)} = \int e^{-y(E_{fin} - E_{init})}, \quad \text{(B7)}$$

where the integral is over all the g_r, h_r fields, each with a probability distribution given by its h survey. As usual, this can be simplified by the use of u-surveys derived from the iteration Eq. (74)

$$e^{-y\Phi_0^v(y)} = \int \prod_{r=1}^{k} \left[\frac{du_r}{C_{a_r \to 0}} Q_{a_r \to 0}(u_r) \right] e^{y|\Sigma u_r|}. \quad \text{(B8)}$$

A little thought shows that, when computing the total

zero-temperature free-energy $\Phi(y)=\Sigma_a \Phi_a^f(y)-\Sigma_i(n_i-1)\Phi_i^v(y)$, one is correctly counting each node once. In particular, in the limit of $y\to 0$, $\Phi(y)$ reduces to the sum of the energy of all factor nodes, as it should. The same reasoning shows that the "$1/C$" factors in Eqs. (B4) and (B8) actually cancel, so that one can forget these normalizations for the computation of $\Phi(y)$, as was done in the text in formulas (71) and (72).

[1] S. Cook, in *Proceedings of the 3rd Annual ACM Symposium on Theory of Computing* (ACM, New York, 1971), p. 151.

[2] M. Garey and D. S. Johnson, *Computers and Intractability; A guide to the Theory of NP-completeness* (Freeman, San Francisco, 1979); C. H. Papadimitriou, *Computational Complexity* (Addison-Wesley, Reading, MA, 1994).

[3] M. Mézard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1987).

[4] R. Monasson and R. Zecchina, Phys. Rev. Lett. **76**, 3881 (1996).

[5] R. Monasson and R. Zecchina, Phys. Rev. E **56**, 1357 (1997).

[6] S. Kirkpatrick, C. D. Gelatt Jr., and M. P. Vecchi, Science **220**, 671 (1983); V. Černy, J. Optim. Theory Appl. **45**, 41 (1985).

[7] M. Mézard and G. Parisi, J. Phys. (France) Lett. **46**, L771 (1985).

[8] D. Aldous, Random Struct. Algorithms **18**, 381 (2001).

[9] Y. Fu and P. W. Anderson, J. Phys. A **19**, 1605 (1986).

[10] L. A. Levin, SIAM J. Comput. **14**, 285 (1986).

[11] S. Ben-David, B. Chor, O. Goldreich, and M. Luby, J. Cell Sci. Suppl. **44**, 193 (1992).

[12] Y. Gurevich, J. Cell Sci. Suppl. **42**, 246 (1991).

[13] Artif. Intel. **81** (1996).

[14] S. A. Cook and D. G. Mitchell, *Satisfiability Problem: Theory and Applications*, edited by Du, Gu, and Pardalos, DIMACS Series in Discrete Mathematics and Theoretical Computer Science Vol. 35 (American Mathematical Society, 1997).

[15] Theor. Comput. Sci. **265** (2001), special issue on phase transitions in combinational problems, edited by O. Dubois, R. Monasson, B. Selman, and R. Zecchina.

[16] S. Kirkpatrick and B. Selman, Science **264**, 1297 (1994).

[17] For a review, see J.-P. Bouchaud, L. Cugliandolo, J. Kurchan, and M. Mézard, in *Spin Glasses and Random Fields*, edited by A. P. Young (World Scientific, Singapore, 1998).

[18] J. S. Yedidia, W. T. Freeman, and Y. Weiss, *Understanding Belief Propagation and its Generalizations* (IJCAI, 2001).

[19] M. Mézard, G. Parisi, and R. Zecchina, Science **297**, 812 (2002).

[20] M. Mézard and G. Parisi, Eur. Phys. J. B **20**, 217 (2001).

[21] M. Mézard and G. Parisi, *The Cavity Method at Zero Temperature* (to be published).

[22] F. R. Kschischang, B. J. Frey, and H.-A. Loeliger, IEEE Trans. Inf. Theory **47**, 498 (2002).

[23] F. Ricci-Tersenghi, M. Weigt, and R. Zecchina, Phys. Rev. E **63**, 026702 (2001).

[24] R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, and L. Troyansky, Nature (London) **400**, 133 (1999).

[25] B. Bollobàs, *Random Graphs* (Academic Press, New York, 1985).

[26] M. Mézard and G. Parisi, Europhys. Lett. **3**, 1067 (1987).

[27] A. Z. Broder, A. M. Frieze, and E. Upfal, in *Proceedings of the Fourth Annual ACM-SIAM Symposium on Discrete Algorithms* (ACM-SIAM, New York, 1993), pp. 322–330.

[28] M. Mézard, G. Parisi, and M. A. Virasoro, Europhys. Lett. **1**, 77 (1986).

[29] M. Talagrand, Probability Theory and Related Fields , **117**, 303 (2002).

[30] R. Monasson, J. Phys. A **31**, 513 (1998).

[31] R. Monasson, Phys. Rev. Lett. **75**, 2847 (1995).

[32] S. Franz and G. Parisi, J. Phys. I **5**, 1401 (1995).

[33] A. Goerdt, in Proceedings of the 17th International Symposium on Mathematical Foundations of Computer Science, Prague, 1992 (unpublished).

[34] V. Chvàtal and B. Reed, in Proceedings of the 33rd IEEE Symp. on Foundations of Computer, 1992 (unpublished), p. 620.

[35] B. Bollobàs, C. Borgs, J. T. Chayes, J. Han Kim, and D. B. Wilson, Rand Struct. Algorithms **18**, 301 (2001).

[36] J. A. Crawford and L. D. Auton, Artif. Intel. **81**, 31 (1996).

[37] A. Kaporis, L. Kirousis, and E. Lalas, in *Proceedings of the 4th European Symposium on Algorithms, ESA, 2002*, Lecture Notes in Computer Science (Springer-Verlag, Berlin, 2002).

[38] D. Achlioptas and G. Sorkin, *Proceedings of the 41st Annual Symposium of Foundations of Computer Science* (IEEE Computer Society Press, Los Alamitos, CA, 2000), p. 590.

[39] J. Franco, Theor. Comput. Sci. **265**, 147 (2001).

[40] O. Dubois, Y. Boufkhad, and J. Mandler, in *Proceedings of the 11th ACM-SIAM Symposium on Discrete Algorithms*, Vol. 124 (ACM-SIAM, San Francisco, CA, 2000).

[41] G. Biroli, R. Monasson, and M. A. Weigt, Eur. Phys. J. B **14**, 551 (2000).

[42] S. Franz, M. Leone, F. Ricci-Tersenghi, and R. Zecchina, Phys. Rev. Lett. **87**, 127209 (2001).

[43] P. Svenson and M. G. Nordahl, Phys. Rev. E **59**, 3983 (1999).

[44] Note that the existence of finite local fields is true for states of arbitrary energy. In particular, considering the zero energy states, selected by a reweighting $y\to\infty$, one sees from Eq. (57) that their nontrivial cavity fields scale as e^{-y}; however when computing the local fields as in Eq. (70), the extra reweighting exactly compensates this scaling, giving finite local fields also when $y\to\infty$.

[45] G. Parisi (unpublished).

[46] R. Kikuchi, Phys. Rev. **81**, 988 (1951).

[47] J. S. Yedidia, W. T. Freeman, and Y. Weiss, in *Advances in Neural Information Processing Systems 13*, edited by T. K. Leen, T. G. Dietterich, and V. Tresp (MIT Press, Cambridge, MA, 2001), pp. 689–695.

[48] B. Selman, H. Kautz, and B. Cohen, *Local Search Strategies for Satisfiability Testing*, Proceedings of DIMACS (American Mathematical Society, New York, 1993), p. 661.

[49] Satisfiability Library: www.satlib.org/

[50] L. Leuzzi and G. Parisi, J. Stat. Phys. **103**, 679 (2001); A. Crisanti, L. Leuzzi, and G. Parisi, J. Phys. A **35**, 481 (2002).
[51] A. Braunstein, M. Mézard, and R. Zecchina (unpublished).
[52] M. Mézard, F. Ricci-Tersenghi, and R. Zecchina, e-print cond-mat/0207140.
[53] S. Cocco, O. Dubois, J. Mandler, and R. Monasson, e-print cond-mat/0206239.