

Simplest random K -satisfiability problem

Federico Ricci-Tersenghi,^{1,*} Martin Weigt,^{2,†} and Riccardo Zecchina^{1,‡}

¹*The Abdus Salam International Centre for Theoretical Physics, Condensed Matter Group, Strada Costiera 11, P.O. Box 586, I-34100 Trieste, Italy*

²*Institute for Theoretical Physics, University of Göttingen, Bunsenstrasse 9, D-37073 Göttingen, Germany*

(Received 16 August 2000; published 23 January 2001)

We study a simple and exactly solvable model for the generation of random satisfiability problems. These consist of γN random boolean constraints which are to be satisfied simultaneously by N logical variables. In statistical-mechanics language, the considered model can be seen as a diluted p -spin model at zero temperature. While such problems become extraordinarily hard to solve by local search methods in a large region of the parameter space, still at least one solution may be superimposed by construction. The statistical properties of the model can be studied exactly by the replica method and each single instance can be analyzed in polynomial time by a simple global solution method. The geometrical and topological structures responsible for dynamic and static phase transitions as well as for the onset of computational complexity in the local search method are thoroughly analyzed. Numerical analysis on very large samples allows for a precise characterization of the critical scaling behavior.

DOI: 10.1103/PhysRevE.63.026702

PACS number(s): 89.20.Ff, 75.10.Nr

I. INTRODUCTION

Complexity theory [1], as arising from Cook's theorem of 1971 [2], deals with the issue of classifying combinatorial optimization problems according to the computational cost required for their solution. The hard problems are grouped in a class named NP, where NP stands for "non-deterministic polynomial time." These problems are such that a potential solution can be checked rapidly whereas finding one solution may require a time growing exponentially with system size in the worst case. In turn, the hardest problems in NP belong to a subclass called NP-complete which is at the root of computational complexity. The completeness property refers to the fact that if an efficient algorithm for solving just one of these problems could be found, then one would have an efficient algorithm for solving all problems in NP. By now, a huge number of NP-complete problems have been identified [1], and the lack of an efficient algorithm corroborates the widespread conjecture that no such algorithm exists, or more formally that $NP \neq P$ where P includes all problems solvable in polynomial time.

Complexity theory is based on a worst-case analysis and therefore does not depend on the properties of the particular instances of the problems under consideration. In order to deepen the understanding of typical-case complexity rather than the worst-case one and to improve and test algorithms for real world applications, computer scientists have recently focused their attention on the study of random instances of hard computational problems, seeking for a link between the onset of computational complexity and some intrinsic (i.e., algorithm independent) properties of the model. Analytical and numerical results have accumulated [3–8] showing that the computationally hard instances appear with a significant probability only when generated near "phase boundaries,"

i.e., when problems are critically constrained. This phenomenon is known as the easy-hard transition.

Randomized search algorithms provide efficient heuristics for quickly finding solutions provided they exist. At the phase boundary, however, there appears an exponential critical slowing down which makes the search inefficient for any practical purpose. Understanding the behavior of search processes at the easy-hard transition point constitutes an important theoretical challenge which can be viewed as the problem of building a generalized off-equilibrium theory for stochastic processes which do not satisfy detailed balance. No static probability measure describing the asymptotic statistical behavior of the search processes is guaranteed to exist. Moreover, the hardest random instances of combinatorial optimization problems provide a natural test bed for the optimization of heuristic search algorithms which are widely used in practice.

How to generate hard and solvable instances is far from obvious and very few examples of such generators are known [9]. In most cases, e.g., in the random Boolean satisfiability problem (K -SAT [6,7,10]), for a short definition see Ref. [11], hard instances can only be found in a very narrow region of the parameters space. In this region the probability that a random instance of the problem has no solution at all is finite. Then, heuristic (incomplete) search algorithms have no way to disentangle, in a given finite time, the unsatisfiable instances, from those which are simply very hard to solve.

In this paper, we shall discuss a very simple and exactly solvable model for the generation of random combinatorial problems. On one hand, these become extraordinarily hard to solve by local search methods in a large region of the parameter space and yet at least one solution may be superimposed by construction. On the other hand, the model may be solved in polynomial time by a simple global method and therefore belongs to the class P .

At variance with respect to the famous random 2-SAT problem [12,10], which is in P and can be solved efficiently by local search methods [13] also at the phase boundary, the model we consider undergoes an easy-hard transition very

*Electronic address: riccife@ictp.trieste.it

†Electronic address: weigt@theorie.physik.uni-goettingen.de

‡Electronic address: zecchina@ictp.trieste.it

similar (even harder) to the one observed in 3-SAT as far as local search methods are concerned. However, the exact mapping of the model on a minimization problem over uniform random hypergraphs makes the problem analytically tractable. It is also solvable in polynomial time by a global method which allows for the numerical study of very large systems. Therefore some of the open questions which arise from the analysis of 3-SAT and which are common to the present model can be answered exactly.

In the context of statistical physics the model provides a simple model for the glass transition, in which the crystalline state can be viewed as the superimposed solution and the structure of the excited states is responsible for the off-equilibrium behavior and the associated structural glass transition. These aspects will be the subject of a forthcoming paper. The limit of infinite connectivity provides one of the most studied models in the context of spin glass theory, see, for instance [14–19].

II. MODEL

In order to unveil the different aspects of the model, to be referred to as hyper-SAT (hSAT), we give explicitly its definition both as a satisfiability problem and as a minimization problem over hypergraphs.

Here we discuss the hSAT model with $K=3$ variables per constraint, which can be viewed as a perfectly balanced version of the famous random 3-SAT problem. The case $K=2$ does not present any interesting computational features as far as hardness is concerned because it can be solved efficiently both by local and global methods. Generalizations to $K>3$ are straightforward.

Given a set of N Boolean variables $\{x_i=0,1\}_{i=1,\dots,N}$, we construct an instance of 3-hSAT as follows. First we define the following elementary constraints (4-clauses sets with 50% satisfying assignments)

$$\begin{aligned} C(ijk|+1) &= (x_i \vee x_j \vee x_k) \wedge (\bar{x}_i \vee \bar{x}_j \vee \bar{x}_k) \wedge (\bar{x}_i \vee x_j \vee \bar{x}_k) \\ &\quad \wedge (\bar{x}_i \vee \bar{x}_j \vee x_k), \\ C(ijk|-1) &= (\bar{x}_i \vee \bar{x}_j \vee \bar{x}_k) \wedge (\bar{x}_i \vee x_j \vee x_k) \wedge (x_i \vee \bar{x}_j \vee x_k) \\ &\quad \wedge (x_i \vee x_j \vee \bar{x}_k), \end{aligned} \quad (1)$$

where \wedge and \vee stand for the logical AND and OR operations, respectively, and the overbar is the logical negation. Then, by randomly choosing a set E of M triples $\{i,j,k\}$ among the N possible indices and M associated unbiased and independent random variables $J_{ijk} = \pm 1$, we construct a Boolean expression in conjunctive normal form (CNF) as

$$F = \bigwedge_{\{i,j,k\} \in E} C(ijk|J_{ijk}). \quad (2)$$

A logical assignment of the $\{x_i\}$ s satisfying all clauses, that is evaluating F to true, is called a solution of the 3-hSAT problem. If no such assignment exists, F is said to be unsatisfiable.

A slightly different choice of J_{ijk} allows one to construct hSAT formulas which are random but guaranteed to be satisfiable. To every Boolean variable we associate indepen-

dently drawn random variables $\varepsilon_i = \pm 1$, and define $J_{ijk} = \varepsilon_i \varepsilon_j \varepsilon_k$ for all $\{i,j,k\} \in E$. For this choice, CNF formula (2) is satisfied by $\{x_i | x_i = +1 \text{ if } \varepsilon_i = +1, x_i = 0 \text{ if } \varepsilon_i = -1\}$. As we shall discuss in great detail, these formulas provide a uniform ensemble of hard satisfiable instances for local search methods. We refer to this version of the model as the satisfiable hSAT. Indeed, the random signs of J_{ijk} can be removed in this satisfiable case by negating all Boolean variables x_i associated to negative ε_i . The resulting model has $J_{ijk} = +1$ for all $\{i,j,k\} \in E$, and the forced satisfying solution is $x_i = 1, \forall i = 1, \dots, N$. The use of the $\{\varepsilon_i\}$ is a way of hiding the latter solution by a random gauge transformation without changing the properties of the model. The impossibility of efficiently inverting the gauge transformation by local methods is a consequence of the branching process arising from the presence of $K=3$ variables in each constraint. For any $K>3$ the same result would hold whereas for $K=2$ the problem trivializes.

The hSAT model can be easily described as a minimization problem of a cost-energy function over a random hypergraph. Given a random hypergraph $\mathcal{G}_{N,M} = (V, E)$, where V is the set of N vertices and E is the set of M hyperedges joining triples of vertices, the energy function to be minimized reads

$$H_J[\mathbf{S}] = M - \sum_{\{i,j,k\} \in E} J_{ijk} S_i S_j S_k, \quad (3)$$

where each vertex i bears a binary ‘‘spin’’ variable $S_i = \pm 1$, and the weights J_{ijk} associated to the random bonds can be either ± 1 at random, in the so-called frustrated case, or simply equal to 1 in the unfrustrated model.

Once the mapping $S_i = 1$ if $x_i = 1$ and $S_i = -1$ if $x_i = 0$ is established, one can easily notice that the energy function in Eq. (3) simply counts twice the number of violated clauses in the previously defined CNF formulas with the same set of J 's. The frustrated and the unfrustrated cases correspond to the hSAT and to the satisfiable hSAT formulas, respectively.

The computational issue consists in finding a configuration of spin variables which minimizes H . If all the terms $J_{ijk} S_i S_j S_k$ appearing in the energy are simultaneously maximized (‘‘satisfied’’) the energy vanishes. This is always possible in the unfrustrated case just by setting $S_i = 1, \forall i$. In the frustrated case there exist a critical value of the average connectivity above which the various terms start to be in conflict, that is frustration becomes effective in the model. In random hypergraphs the control parameter is the average density of bonds, $\gamma = M/N$ (or, for the CNF formula, the density of clauses $\alpha = 4\gamma$). For sufficiently small densities, the graph consists of many small connected clusters of size up to $O(\ln N)$. If γ increases up to the percolation value $\gamma_p = 1/6$, there appears a giant cluster containing a finite fraction of the N sites in the limit of large N . However, this cluster can also *a priori* have a treelike structure, for which the randomness of the couplings $J_{ijk} = \pm 1$ can be eliminated by a proper gauge transformation, $S_i \rightarrow \pm S'_i$, of the spin variables. As we shall see, there exist two other thresholds of the bond density at which more complicated and interesting dynamical and structural changes take place.

In spite of apparent similarities, hSAT and the random Boolean satisfiability problem (K -SAT [11]) differ in some basic aspects.

In K -SAT the fluctuations of the frequencies of appearance of the variables in the clauses lead to both single and two body interactions in the associated energy function [20] which force the minima in some specific random directions and which rule out the existence of a purely dynamical threshold (see below). Algorithms may take advantage of such information and both heuristic as well as complete algorithms show a performance which indeed depends on the criterion used to fix the variables. For example, rigorous lower bounds to the critical threshold have been improved recently by exploiting this opportunity in a simple tractable way [21]. On the same footing, the efficiency of the most popular heuristic and complete search algorithms, namely walk-sat [22] and TABLEAU [23], is again based on strategies which exploit the above structure. Note that the above improvements cannot be applied to the hSAT model where formulas are completely balanced.

Moreover, in K -SAT the mapping of the problem over directed random graphs is rather involved and the exact analytical solution is still lacking, while in hSAT the connection to random hypergraphs is clear and makes the analysis tractable.

Finally, restriction of K -SAT to satisfiable instances (for instance, by selecting at random clauses which are satisfied by a previously fixed assignment of variables) does not provide a uniform ensemble of hard satisfiable problems even when restricted to local search methods [24].

Given the mapping over random hypergraphs, the satisfiability problem for hSAT can be solved in $O(N^3)$ steps by simply noticing that the problem of satisfying all constrains is nothing but the problem of solving an associated linear system modulo 2, i.e., in $GF[2]$. Upon introducing the two sets of binary variables $\{a_i\} \in \{0,1\}^N$ and $\{b_{ijk}\} \in \{0,1\}^M$ such that $(-1)^{a_i} = S_i$ and $(-1)^{b_{ijk}} = J_{ijk}$, the hSAT decision problem becomes simply the problem of determining the existence of a solution in $GF[2]$ to the random linear system $a_i + a_j + a_k = b_{ijk} \pmod{2}$, with ijk running over all triples.

Finally, we notice that in the high γ UNSAT (or frustrated) region the optimization problem of minimizing the number of violated constrains, the so-called MAX-hSAT, is indeed computationally very hard both for complete and incomplete algorithms and no global method for finding ground states is available.

III. OUTLINE OF MAIN RESULTS

For the sake of clarity, we anticipate here the main results leaving for the following sections a thorough discussion of the analytical and numerical studies.

The frustrated hSAT model presents two clear transitions. The first one appears at $\gamma_d = 0.818$ and it is of purely dynamical nature. There the typical formula still remains satisfiable with probability one, but an exponential number of local energy minima appear at positive energies. Deterministic algorithms, like greedy search or zero temperature dynamics, try to decrease the energy in every step and thus get

stuck at least at this threshold. Randomized algorithms may escape from these minima, but they undergo a slowing down from an exponentially fast convergence to a polynomially slow one, i.e., at γ_d the typical time for finding a solution diverges as a power of the number of variables. The dynamical transition at γ_d seems to be accompanied by a dynamical glassy transition due to replica symmetry breaking (RSB) effects connected with the appearance of an exponential number of local minima. An approximate variational calculation (see Ref. [25] for a discussion on the method) involving RSB gives $\gamma_{d,RSB}^{\text{var}} \approx 0.83$ which is in good agreement with the value of γ_d where local minima appear.

The second transition appears at $\gamma_c = 0.918$ and corresponds to the so-called SAT/UNSAT transition (below γ_c the typical problem is satisfiable whereas above γ_c it becomes unsatisfiable). At this point the structure of the global energy minima changes abruptly. The ground states have strictly positive energy, thus no satisfying assignments for the hSAT formula exist any more. While the number of these configurations is always exponentially large (the ground state entropy is always finite), at γ_c a finite fraction of the variables, the so-called backbone component, becomes totally constraint, i.e., the backbone variables take the same value in all minima [26]. An important difference of the SAT/UNSAT transition in hSAT compared to K -SAT [25] is the nonexistence of any precursor. For $\gamma < \gamma_c$ and large N , all variables S_i take equally often the values $+1$ and -1 in the ground states (they have zero local magnetization), even those which become backbone elements when γ_c is reached by adding new 4-clauses sets. The lack of any precursor comes from the nonexistence of single- or two-body interactions in Eq. (3).

The unfrustrated or satisfiable hSAT problem has by construction at least one solution which we find to be superimposed without affecting the statistical features of the model for $\gamma < \gamma_c$ in the limit of large N , including the dynamical transition at $\gamma_d = 0.818$. It is impossible to get any information on the superimposed solution by looking at the full solution space because it is completely hidden by the exponential number of ground states. Randomly chosen satisfying assignments do not show any correlation. At exactly the same $\gamma_c = 0.918$ as in the frustrated model, there appears a transition from a SAT phase with exponentially many unbiased solutions to another SAT phase where the solutions are strongly concentrated around the superimposed solution. The latter one is now hidden by the presence of exponentially many local energy minima with positive cost. These minima have exactly the statistical properties of the global minima of the corresponding frustrated hSAT problem, that is the hSAT problem defined over the same hypergraph but with randomized signs of the couplings J_{ijk} . More specifically, the energy, the entropy, and the backbone component size coincide. Due to their finite entropy, an algorithm will thus hit many of these local minima before it reaches the satisfying ground state. As one can see from Fig. 5, finding this solution by backtracking, e.g., with the Davis–Putnam (DP) procedure [27], is nevertheless easier than proving the unsatisfiability of hSAT (or identifying ground states in the frustrated version). This results stems from the missing in-

formation on the true ground state energy of hSAT above γ_c . The solution time is, however, found to be clearly exponential in both cases. In the $\gamma \geq \gamma_c$ region, the model provides a uniform ensemble of hard SAT instances for local methods which can be used to test and optimize algorithms.

IV. STATISTICAL MECHANICS ANALYSIS: THE REPLICA RESULTS

In our analytical approach, we exploit the well-known analogies between combinatorial optimization problems and statistical mechanics. In both cases, the system is characterized by some cost-energy function, as it is given, e.g., by Eq. (3) for hSAT. In equilibrium statistical mechanics, any configuration $\mathbf{S} = \{S_i\}_{i=1, \dots, N}$ is realized with probability $\exp\{-\beta H[\mathbf{S}]\}/Z$ where $\beta = 1/T$ is the inverse temperature and Z the partition function. If the temperature is lowered, the probability becomes more and more concentrated on the global energy minima and finally, for $T=0$, only the ground states keep nonzero weights.

In order to compute the average free energy, we resort to the replica symmetric (RS) functional replica method developed for diluted spin glasses which is known to provide exact results for ferromagnetic models: To circumvent the difficulty of computing the average value of $\ln Z$, we compute the n th moment of Z for integer-valued n and perform an analytical continuation to real n to exploit the identity $\langle\langle Z^n \rangle\rangle = 1 + n\langle\langle \ln Z \rangle\rangle + O(n^2)$. The n th moment of Z is obtained by replicating n times the sum over the spin configuration and averaging over the disorder [28]

$$\langle\langle Z^n \rangle\rangle = \sum_{s^1, s^2, \dots, s^n} \left\langle \left\langle \exp\left(-\beta \sum_{a=1}^n H_J[\mathbf{S}^a]\right) \right\rangle \right\rangle, \quad (4)$$

which in turn may be viewed as a generating function in the variable $\exp(-\beta)$.

In order to compute the expectation values that appear in Eq. (4), one notices that each single term $\exp(-\beta \sum_{a=1}^n H_J[\mathbf{S}^a])$ factorizes over the sets of different triples of indices due to the absence of any correlation in the probability distribution of the J_{ijk} . It follows

$$\langle\langle Z^n \rangle\rangle = \sum_{s_i^1, s_i^2, \dots, s_i^n} \exp\left\{-\beta \gamma N n - \gamma N + \frac{\gamma}{N^2} \sum_{ijk} e^\beta \sum_a s_i^a s_j^a s_k^a + O(1)\right\}. \quad (5)$$

The averaged term in Eq. (5) depends on the $n \times N$ spins only through the 2^n occupation fractions $c(\vec{\sigma})$ labeled by the vectors $\vec{\sigma}$ with n binary components; $c(\vec{\sigma})N$ equals the number of labels i such that $S_i^a = \sigma^a$, $\forall a = 1, \dots, n$. Therefore the final expression of the n th moment of Z to the leading order in N (i.e., by resorting to a saddle-point integration), can be written as $\langle\langle Z^n \rangle\rangle \approx \exp(NF[c])$ where $F[c]$ is the maximum over all possible $c(\vec{\sigma})$'s of the functional [28]

$$-\beta F[c] = -\gamma(1 + \beta n) - \sum_{\vec{\sigma}} c(\vec{\sigma}) \ln c(\vec{\sigma}) + \gamma \sum_{\sigma, \rho, \tau} c(\vec{\sigma}) c(\vec{\rho}) c(\vec{\tau}) \exp\left(\beta \sum_a \sigma^a \rho^a \tau^a\right). \quad (6)$$

The saddle-point equation $\partial(-\beta F)/\partial c(\vec{\sigma}) = \Lambda - 1$ reads

$$c(\vec{\sigma}) = \exp\left\{-\Lambda + 3\gamma \sum_{\rho, \tau} c(\vec{\rho}) c(\vec{\tau}) \exp\left(\beta \sum_a \sigma^a \rho^a \tau^a\right)\right\}, \quad (7)$$

where the Lagrange multiplier Λ enforces the normalization constraint $\sum_{\vec{\sigma}} c(\vec{\sigma}) = 1$, and goes to 3γ for $n \rightarrow 0$. In Eq. (6), one may easily identify two terms, one model dependent and the other $[-\sum_{\vec{\sigma}} c(\vec{\sigma}) \ln c(\vec{\sigma})]$ simply describing the degeneracy (the so-called combinatorial entropy) with which each term of the generating function appears given the representation in terms of the occupation fractions. In the limit of interest $T \rightarrow 0$ and in the replica symmetric subspace, the freezing of the spin variables is properly described by a rescaling of the local magnetizations of the form $m = \tanh(\beta h)$. The probability distribution $P(h)$ is therefore introduced through the generating functional

$$c(\vec{\sigma}) = \int_{-\infty}^{\infty} dh P(h) \frac{e^{\beta h \sum_a \sigma^a}}{[2 \cosh(\beta h)]^n}, \quad (8)$$

where h is nothing but an effective field in which the spins are immersed. c depends on $\vec{\sigma}$ only via $s = \sum_a \sigma^a$. In this representation, the free-energy reads

$$-\beta F[P(h)] = \gamma \int dh_1 dh_2 dh_3 P(h_1) P(h_2) P(h_3) \times \ln \frac{2 \cosh[\beta(h_1 + h_2)][e^{\beta h_3} + e^{-2\beta - \beta h_3}] + 2 \cosh[\beta(h_1 - h_2)][e^{-\beta h_3} + e^{-2\beta + \beta h_3}]}{[2 \cosh(\beta h_1)][2 \cosh(\beta h_2)][2 \cosh(\beta h_3)]} + \int \frac{dh dK}{2\pi} e^{ihK} P_{FT}(K) [1 - \ln P_{FT}(K)] \ln[2 \cosh(\beta h)], \quad (9)$$

where $P_{FT}(K)$ is the Fourier transform of $P(h)$. The associated saddle-point equation reads

$$\int dh P(h) e^{\beta h s} = \exp \left\{ -3\gamma + 3\gamma \int dh_1 dh_2 P(h_1) \times P(h_2) G(h_1, h_2) \right\}, \quad (10)$$

where

$$G(h_1, h_2) = \left(\frac{\cosh[\beta(h_1 + h_2)] + e^{-2\beta} \cosh[\beta(h_1 - h_2)]}{\cosh[\beta(h_1 - h_2)] + e^{-2\beta} \cosh[\beta(h_1 + h_2)]} \right)^{s/2}. \quad (11)$$

In the case of satisfiable hSAT, at $\beta \rightarrow \infty (T=0)$ and in the version having no random gauge ($J_{ijk} = +1, \forall \{i, j, k\} \in E$), the spins turn out to be subject to an effective local field h which fluctuates from site to site according to the following simple probability distribution

$$P(h) = \sum_{\ell \geq 0} p_\gamma^{(\ell)} \delta(h - \ell) \quad (12)$$

with the saddle-point conditions

$$p_\gamma^{(\ell)} = (3\gamma)^\ell \frac{(1 - p_\gamma^{(0)})^{2\ell} p_\gamma^{(0)}}{\ell!}, \quad (13)$$

$$p_\gamma^{(0)} = \exp\{-3\gamma(1 - p_\gamma^{(0)})^2\}.$$

The above structure is not surprising for a ferromagnetic model since $1 - p_\gamma^{(0)}$ is nothing but the fraction of sites which have nonvanishing field and are therefore totally magnetized. The saddle-point equations simplify once rewritten in terms of the probability distribution $P(m)$ of the local magnetizations $m_i = 0, 1$ which takes the particularly simple form

$$P(m) = p_\gamma^{(0)} \delta_{m,0} + (1 - p_\gamma^{(0)}) \delta_{m,1}, \quad (14)$$

where $\delta_{\cdot, \cdot}$ is the Kronecker symbol. Thus a fraction $1 - p_\gamma^{(0)}$ of all logical variables is frozen to $+1$ in all ground states, whereas the others take both truth values with the same frequency. The self-consistent equation for $p_\gamma^{(0)}$ in Eq. (13) can be rewritten as

$$p_\gamma^{(0)} = \sum_{c=0}^{\infty} e^{-3\gamma} \frac{(3\gamma)^c}{c!} [1 - (1 - p_\gamma^{(0)})^2]^c, \quad (15)$$

and can be justified by simple probabilistic arguments: A variable is frozen if and only if it is contained in at least one hyperedge $\{i, j, k\} \in E$ where also the two neighbors are frozen. Thus a variable is unfrozen, $m_i = 0$, if and only if every adjacent hyperedge contains at least one more unfrozen variable. For a spin of connectivity c , this happens according to Eq. (14) with probability $[1 - (1 - p_\gamma^{(0)})^2]^c$. The average over the Poisson-distribution $e^{-3\gamma} (3\gamma)^c / c!$ of connectivities c results in the total probability for a variable to be unfrozen,

so Eq. (15) follows. As an additional result of replica theory, we derive the ground-state entropy

$$s(\gamma) = \lim_{N \rightarrow \infty} \frac{1}{N} \ln \mathcal{N}_{gs} = \ln(2) (p_\gamma^{(0)} (1 - \ln p_\gamma^{(0)}) - \gamma [1 - (1 - p_\gamma^{(0)})^3]). \quad (16)$$

For small γ , Eq. (15) has only the trivial solution $p_\gamma^{(0)} = 1$ where all variables are unfrozen, i.e., $m_i = 0$ for all i . No internal structure is found in the set of satisfying assignments and, choosing randomly two of them, they have Hamming distance $0.5N + O(\sqrt{N})$. To leading order in N , the M 4-clauses sets act independently, each dividing the number of satisfying assignments by two, i.e., $\mathcal{N}_{gs} = 2^{N(1-\gamma)}$. This is a clear sign that the structure of the hypergraph is still treelike.

At $\gamma_d = 0.818$, a new solution of Eq. (15) appears discontinuously, having a fraction $(1 - p_\gamma^{(0)}) = 0.712$ of completely magnetized variables. This transition can be seen as a percolation transition of fully magnetized triples of connected variables. The entropy of this solution remains, however, smaller than the entropy $1 - \gamma$ of the paramagnetic solution, thus the total solution space is still correctly described by $m_i = 0$ for all $i = 1, \dots, N$. The appearance of the new solution signals, however, a structural change in the set of solutions which breaks into an exponential number of clusters. The cluster containing the imposed solution $x_i = +1$ is described by the new metastable solution.

Another important difference to the low- γ phase is an exponential number of local minima of the energy function (3) showing up at γ_d . These have positive energies, and the corresponding logical assignments do not satisfy the hSAT formula. Algorithms which decrease the energy in every time step by local variable changes, e.g., zero-temperature Glauber dynamics or greedy algorithms, get almost surely trapped in these states and do not find a zero-energy ground state for $\gamma > \gamma_d$. Randomized algorithms may escape from these minima, but as found numerically, this causes a polynomial slowing down.

By increasing γ above γ_d , the number of ground-state clusters decreases further. At $\gamma_c = 0.918$ all but one ground-state clusters disappear, and the nontrivial solution of Eq. (15) becomes the stable one. So only the cluster including the imposed solution survives, it still contains $2^{0.082N}$ solutions, but 88.3% of all variables are fixed to $+1$, thus forming the backbone which appears discontinuously. As is known from Ref. [7], the existence of an extensive backbone is closely related to the exponential computational hardness of a problem. The remaining 11.7% of unfrozen variables change their values from ground state to ground state. They are contained in small disconnected components or dangling ends of the hypergraph.

The behavior of frustrated hSAT is similar, as given both by numerical analysis and by RS or variational RSB calculations. We find that the solution $x_i = +1$ (and its corresponding cluster) in satisfiable hSAT are just superimposed to the solution structure of random hSAT. Thus the statistical properties of the solutions do not change for $\gamma < \gamma_c$, including also the clustering of solutions above γ_d . At $\gamma_c = 0.918$

the model undergoes a SAT/UNSAT transition, and the solution entropy jumps from 0.082 down to minus infinity. The variational RSB calculation gives a value for the dynamical critical connectivity $\gamma_{d,\text{RSB}}^{\text{var}} \approx 0.83$ which is close to the exact value 0.818. This result gives evidence for the validity of the variational approach in the region where local minima first appear, i.e., where the result does not depend strongly on the specific functional ansatz made for the RSB probability distributions. For the SAT/UNSAT static transition the predictions of the variational RSB analysis can be strongly affected by the restriction of the functional space which does not necessarily match the geometrical structure (clustering) of the space of solution. However, in the case of hSAT the results are still in good agreement, we find $\gamma_{c,\text{RSB}}^{\text{var}} \approx 0.935$.

V. CONNECTION WITH GRAPH THEORY

In the hSAT model, we are able to extract exact results, without the need of RSB, by identifying the topological structures in the underlying hypergraph which are responsible of the SAT/UNSAT transition (or of frustration and glassiness). The presence (or the absence) of such topological structures in the hypergraph drastically changes the statistical mechanical properties of the model. The different phase transitions can be viewed as different kinds of percolation in the random graph theory language [29].

We have already seen that the formation of a locally stable ferromagnetic state in unfrustrated hSAT at $\gamma_d = 0.818$ can be understood in term of percolation arguments. The same arguments reveal that at γ_d many metastables states appear in both versions of the model, giving rise to a dynamical transition.

In order to understand what happens at the critical point γ_c we need to introduce the notion of hyperloops, that is the most natural generalization of the usual loop to hypergraphs with multivertex links. Given a random graph $\mathcal{G} = (V, E)$, where V is the set of vertices and E is the set of (hyper)links, a hyperloop can be defined as a nonzero set of (hyper)links, $R \subset E$, such that the degree of the subgraph $\mathcal{L} = (V, R)$ is even, i.e., every vertex belongs to an even number of (hyper)links (including zero). In Fig. 1 (left) we show the smallest hyperloop in a $K=3$ random hypergraph. Note that in random hypergraph typical hyperloops are very large and the one shown in Fig. 1 (left) is extremely rare for N large.

In a similar way we can identify those vertices which are totally constrained. A set of (hyper)links, $T^{(i)} \subset E$, constrains completely the spin at site i if in the subgraph $\mathcal{F} = (V, T^{(i)})$ the vertex i has an odd degree and the remaining vertices an even one. In Fig. 1 (right) we show the smallest of such structures.

In a zero-energy configuration (SAT assignment) we have $S_i S_j S_k = J_{ijk}$, $\forall \{i, j, k\} \in E$. Then, given any hyperloop R , we conclude

$$\prod_{\{i,j,k\} \in R} J_{ijk} = \prod_{\{i,j,k\} \in R} S_i S_j S_k = 1, \quad (17)$$

where the second equality comes from the fact that in the second product every spin appears an even number of times.




FIG. 1. The simplest hyperloop (left) and the hyperloop with one totally constrained vertex of odd degree (right) [30]. Triangles represent the interaction between the three spins located at the vertices. The black dot represents the constrained spin residing on the odd-degree vertex of the hyperloop.

In frustrated hSAT the couplings are randomly fixed to ± 1 and, consequently, the first product in Eq. (17) is equal to -1 with probability $1/2$. Then we can conclude that as soon as one hyperloop arises in the hypergraph half the formulas become unsatisfiable. In general, given a hypergraph with N_{hl} hyperloops, the fraction of SAT formulas (with that given hypergraph) is $2^{-N_{hl}}$. Still one needs to average this fraction over the random hypergraph in order to obtain the right fraction of SAT formulas.

We have numerically found that at the critical value $\gamma_c = 0.918$ the percolation of hyperloops takes place, that is, in the large N limit, the average number of hyperloops $N_{hl}(\gamma)$ is zero for $\gamma < \gamma_c$ and $\mathcal{O}(N)$ for $\gamma > \gamma_c$. This is the direct explanation of the SAT/UNSAT transition in terms of hypergraph topology.

In the unfrustrated model $J_{ijk} = 1$ and Eq. (17) is always satisfied. However, the mean number of hyperloops $N_{hl}(\gamma)$ is related to the entropy of satisfying assignments through $s(\gamma) = [1 - \gamma + N_{hl}(\gamma)/N] \ln 2$. The derivation of this equality is straightforward if we consider the linear system modulo 2 of M equations in N variables, introduced at the end of Sec. II. In terms of the linear system hyperloops represent combinations of equations giving a trivial one (e.g., $0=0$) which does not fix any degree of freedom. The entropy, which is proportional to the number of degree of freedom, is then given by $S(\gamma) = \ln(2)[N - M - N_{hl}(\gamma)]/N$.

Considering now a totally constrained spin at site l and a SAT assignment, we have that

$$\prod_{\{i,j,k\} \in T^{(l)}} J_{ijk} = \prod_{\{i,j,k\} \in T^{(l)}} S_i S_j S_k = S_l. \quad (18)$$

Then, in every SAT formula, hyperloops with one odd-degree vertex (to be denoted by the label $hl-1$) fix one spin variable to a complicated function of the couplings. We have

numerically checked, that such structures arise at γ_c , like hyperloops, but in a discontinuous way.

In satisfiable hSAT we have $J_{ijk}=1$, so that any independent hyperloop with one odd degree vertex fixes one spin to 1 [31]. Then the magnetization of the model is equal to the mean density of such loops, $m(\gamma)=\rho_{hl-1}(\gamma)=N_{hl-1}(\gamma)/N$. Because of the discontinuous nature of the transition the limits $\lim_{\gamma \rightarrow \gamma_c^-} \rho_{hl-1}(\gamma)=0$ and $\lim_{\gamma \rightarrow \gamma_c^+} \rho_{hl-1}(\gamma)=m_c=0.883$ do not coincide.

In frustrated hSAT Eq. (18) fixes the variables belonging to the backbone. Then one would be tempted to relate the backbone size to the density $\rho_{hl-1}(\gamma)$ of hyperloops with one frozen vertex (which is true) and to estimate the backbone size at the critical point to be 88.3% (which is not true). Indeed at the critical point there is a coexistence of SAT and UNSAT formulas (see the next section) and for $\gamma > \gamma_c$ all the formulas become UNSAT in the large N limit. Then Eq. (18) can be applied only for $\gamma < \gamma_c$ where the density ρ_{hl-1} goes to zero when $N \rightarrow \infty$. While the appearance of the backbone is necessarily related to the presence of hyperloops with frozen vertices, the estimation of its size is nontrivial. A very rough estimate can be obtained assuming that at the critical point half the formulas are SAT (according to the numerical results presented in the next section) and that the backbone size is 0 for UNSAT formulas and 0.88 for SAT ones. Under these very crude hypothesis the backbone size would be 0.44, which is not too far from the numerical result (see the next section).

VI. NUMERICAL RESULTS

We have performed extensive numerical experiments on both versions of hSAT in order to confirm analytical predictions and to compute quantities which are not accessible analytically. Beside the $GF[2]$ polynomial method, we have also used two local algorithms, namely the Davis-Putnam (DP) complete backtrack search [27] and the incomplete walk-SAT randomized heuristic search [22], to check the hardness of the problem for local search. The existence of at least one solution in the satisfiable hSAT allowed us to run walk-SAT in the whole range of γ , the halting criterion being always finding a SAT assignment.

The first set of results concerns the numerical determinations of the critical points of hSAT obtained by the polynomial method over large samples.

For the frustrated case, the fraction of satisfiable instances drops down to zero at $\gamma_c=0.918$. In Fig. 2 we show this fraction as a function of γ , which has been obtained, for any size N , counting the number of hyperloops in 10^4 different random hypergraphs. For any given random hypergraph the fraction of SAT formulas is given by $2^{-N_{hl}}$, where $N_{hl}(\gamma)$ is the number of hyperloops. The same set of simulations run on the satisfiable hSAT show that at exactly the same γ_c , the model undergoes a discontinuous ferromagnetic transition.

At $\gamma_d=0.818$ a dynamical transition takes place in both versions of hSAT. There appears an exponentially large number of positive energy local minima strongly affecting nonrandomized dynamics, which is not able to overcome en-




FIG. 2. The probability that a formula is SAT as a function of the coupling density. Inset: The energy reached by a deterministic rule becomes different from zero at the dynamical critical point.

ergy barriers. We can easily detect the dynamical transition by adopting the following deterministic algorithm as a probe and by checking where it stops converging to solutions. The algorithm exploits the only local source of correlations among variables, that is, fluctuations in connectivity. At each step, the algorithm chooses the variable with the highest connectivity, fixes its value at random, and it simplifies the formula (“unit clause propagation” [27]). As can be seen in the inset of Fig. 2 the energy reached running the above process on very large formulas ($N=10^2, 10^3, 10^4$) starts to deviate from zero at a value which is highly compatible with the analytical prediction $\gamma_d=0.818$. Unfortunately the mathematical analysis of this kind of algorithm appears to be beyond our present skills due to the correlations induced into the simplified formulas by the particular choice of variables. For a simple random (connectivity independent) choice of the variable the algorithm can be analyzed along the lines of Ref. [32] and a convergence can be proven up to $\gamma=2/3$, which is also a rigorous lower bound to the true critical density γ_c . A rigorous upper bound is easily established by noticing that the probability for the satisfiability of a formula at fixed γ is bounded by the number of satisfying assignments, averaged over all formulas of length γN . It follows $\gamma_c < 2 \ln 2$ (which is the so-called annealed bound known in the statistical mechanics of disordered media).

We have performed standard finite size scaling analysis in order to determine the size of the critical window $w(N)$ and the ν exponent defined by $w(N) \propto N^{-1/\nu}$ for large N .

In a growing random hypergraph as soon as the first hyperloop arises the fraction of SAT formulas drops down to 0.5. We have measured the mean γ value where this event takes place, $\gamma_c(N)$. Such value scales as $\gamma_c(N) - \gamma_c \propto N^{-1}$, i.e., its critical exponent is $\nu=1$ as expected for a discontinuous transition (see lower inset in Fig. 3).

However in the model there is also another source of pure statistical (not critical) fluctuations [33]. These fluctuations come from the fact that almost every formula can be modified by the addition (or deletion) of order \sqrt{N} clauses without changing its satisfiability. Therefore in the large N limit these




FIG. 3. Scaling function for the SAT probability. Lower inset: The γ value where the first hyperloop arises scales as N^{-1} . Upper inset: The critical width undergoes a crossover from $\nu=1$ to $\nu=2$. The fitting curve is $3.4/N + 0.74/\sqrt{N}$, while the line is the asymptote $0.74/\sqrt{N}$.

purely statistical fluctuations will dominate the critical ones, leading to an exponent $\nu=2$ in the scaling of the SAT probability. In the upper inset in Fig. 3 we show the width of the critical region [34] as a function of N , together with the best fit of the kind $Ax^{-B} + Cx^{-1/2}$. Notably the best fitting value for B is perfectly compatible with 1, giving more evidence to the crossover from critical fluctuations ($\nu=1$) to statistical ones ($\nu=2$).

In the main part of Fig. 3 we show the scaling function for the SAT probability. Note that the value at criticality is equal to 0.5 up to the numerical precision. Slight deviations from perfect scaling appear in the $\gamma > \gamma_c$ region. However, scaling relations hold only close to the critical point and our data perfectly collapse in all the ranges where the SAT probability is between 0.2 and 0.8.

The different kind of transition taking place at γ_c in the two versions of hSAT is reflected in the behavior of their ground-state entropies $s(\gamma)$ shown in Fig. 4. For $\gamma \leq \gamma_c$ both entropies coincide and they have the analytical expression $s(\gamma) = \ln(2)(1-\gamma)$ up to γ_c . For $\gamma > \gamma_c$, while the entropy of satisfiable hSAT decreases exponentially fast with γ (the solutions are more and more concentrated around the superimposed one), in the frustrated version the entropy decreases more slowly with γ , indicating that the number of UNSAT assignments minimizing the energy remains exponentially large up to $\gamma \gg \gamma_c$.

At the SAT/UNSAT transition the solution space acquires a backbone structure, with a finite fraction of the variables that take the same value in all the solutions. Above the critical threshold a similar structure characterizes the ground states. In the inset of Fig. 4 we report the results of exhaustive ground-states enumeration on small systems, giving the average size of the backbone and the average energy. Increasing the system size, the average energy per hyperlink converges to zero for $\gamma < \gamma_c$ and it becomes positive at γ_c in a continuous way. The appearance of the backbone becomes sharper, increasing the system size and, in the thermody-




FIG. 4. The lowest lines are the analytical expressions for the entropy of the unfrustrated model. The numerical estimation (not reported) perfectly coincides. Dashed parts correspond to metastable states. The rest of the data (entropy in the main body and energy and backbone size in the inset) come from exhaustive enumeration of the ground states in the frustrated model and of first excited states in the unfrustrated one (only $N=40,60$) and they coincide.

namical limit ($N \rightarrow \infty$), we expect it to be zero for $\gamma < \gamma_c$ and finite for $\gamma \geq \gamma_c$, consistently with a random first order phase transition predicted by the replica theory. As can be seen in the inset of Fig. 4 the backbone size does not depend strongly on the system size in the UNSAT phase. As discussed in Ref. [7] the presence of a finite backbone is conjectured to be the source of computational hardness in finding solutions at the SAT/UNSAT transition for both complete and randomized local algorithms.

In the $\gamma > \gamma_c$ region the backbone size shows clear oscillations, due to finite size effects. At fixed energy the backbone size is a nondecreasing function of γ , but it typically decreases when the energy jumps to a higher value. For finite systems such jumps, which are of order $1/N$, are particularly evident and induce observable fluctuations in the backbone. We expect these fluctuation to disappear in the thermodynamic limit.

In satisfiable hSAT, once we consider only the lowest local minima configurations just above the zero energy solutions (the so-called excited states) we find that they share completely the same statistical properties with the ground states of the corresponding frustrated hSAT, i.e., the model defined over the same random hypergraph with randomized couplings. We have performed a set of DP runs in satisfiable hSAT similar to the ones used previously, with the additional requirement of not considering the superimposed solution. The backbone size, the average energy, and the entropy of the excited states just above the solution are identical to those measured on the ground states of the frustrated version (see Fig. 4). These results, together with some preliminary analytical findings [35], show in detail why a model without quenched frustration behaves and can be modeled as having random sign interactions, i.e., like a spin glass model. Such a




FIG. 5. The computational costs for finding a solution or proving unsatisfiability with the Davis–Putnam algorithm strongly increase approaching the critical point. For $\gamma \geq \gamma_c$ they grow exponentially with the problem size N . Inset: The same computational costs for the walk-SAT algorithm, which can be run for every γ in the satisfiable model ($N = 25, 50, 75, 100$).

mapping is believed to play a particularly important role in spin glass theory of structural glasses, in which the only source of frustration is geometrical (i.e., dynamical). Once the Boltzmann temperature T is introduced in the model, the critical points of hSAT can be thought of as zero temperature limits of critical lines in the (T, γ) plane. In spite of the absence of any static frustration and of the existence of a pure “crystalline” state (the spin configurations corresponding to the satisfying assignment), the spin system undergoes several dynamical and static transitions as the temperature is lowered. Both the crystalline state and the first excited states are never reached in any subexponential time and the system stays for very large times in the metastable states (the same happens in the frustrated version).

In Fig. 5 we report data concerning the computational costs for finding a solution in the satisfiable hSAT and for proving satisfiability for the frustrated hSAT [36]. For both algorithms (DP and walk-SAT) and in the whole range of γ , we have measured the logarithm of the running time averaged over thousands of samples of different sizes. The choice of analyzing the averaged logarithm instead of the logarithm of the average is dictated by the presence of fat tails in the running time distributions, even in the $\gamma < \gamma_c$ region. The averaged logarithm provides directly the information on the most probable cost.

The main body of the figure displays the DP computational costs for proving satisfiability in hSAT and for finding the satisfying assignment in the satisfiable hSAT (given the same underlying hypergraph structure). Both costs show a sharp easy-hard transition at γ_c , where an enormous increase in the typical running times take place. For $\gamma < \gamma_c$ both costs obviously coincide and they increase as a power law of N , the only effect of γ_d being a change of the exponent from 1 to a large value which eventually diverges at γ_c . For $\gamma > \gamma_c$, the computational costs remain very high, i.e., $\langle \ln[\pi(\gamma)] \rangle \propto \sigma(\gamma)N$, with an exponent σ slowly decreasing as $1/\gamma$ [37].

In the inset of Fig. 5 we show the average logarithm of the running times needed by walk-SAT for finding a solution in the satisfiable hSAT model. Analogously to DP the walk-SAT costs undergo an easy-hard transition at γ_c . Interestingly enough, above γ_c the computational cost for finding solutions remains quite high and does not decrease as in DP, where the additional constraints act as a pruning strategy in the search process. In the hard satisfiable region standard heuristic algorithms, like walk-SAT, get stuck in local minima and they are not able to exploit the large number of constraints in order to reduce the searching space. In particular, the large scale structure [$O(N)$] of the hyperloops makes them difficult to detect in polynomial time by a local search process which is dominated by the exponential branching process arising at each step when the tentative choices for the variables are made. However, having at hand a model on which new heuristic algorithms can be tested, such a searching optimization can hopefully be pushed far forward.

A thorough analysis of the dependence of computational costs on N gives the following overall picture. For $\gamma \leq \gamma_d$ the cost is a linear function of N . For $\gamma \in [\gamma_d, \gamma_c]$ the typical cost increases as a power law of N , with an exponent which should diverge in γ_c . For $\gamma \geq \gamma_c$ the costs are exponential in N .

VII. CONCLUSIONS

In this paper we have studied a model for the generation of random combinatorial problems, called hyper-SAT. In the context of theoretical computer science such a model is simply the completely balanced version of the famous K -SAT model, while in statistical physics it corresponds to a diluted p -spin model at zero temperature. We have studied two versions of the model, a frustrated one and an unfrustrated one.

Increasing the density of interactions, $\gamma = M/N$, the model undergoes two transitions. The first one is of purely dynamical nature whereas the second one is static. Such phase transitions have a straightforward interpretation in terms of the structure of the underlying hypergraphs, leading to a very simple connection between theoretical computer science and graph theory, and statistical physics of random systems.

The locations of phase boundaries can be computed exactly within the RS replica formalism, leading to $\gamma_d = 0.818$ and $\gamma_c = 0.918$. We expect the replica results to be computable also by more rigorous probabilistic methods.

Exploiting a global solution method which is polynomial in the problem size, we have been able to study very large problems, determining with high precision critical points and critical exponent, and a crossover from critical fluctuations to statistical ones has been measured.

We have found that the computational costs for finding a solution to a typical problem or to prove that it is unsatisfiable using only local search methods undergoes easy-hard transitions at γ_d and γ_c . The growth of the costs with the problem size N is linear up to γ_d , is polynomial in N in the range $\gamma_d \leq \gamma \leq \gamma_c$, and finally it becomes exponential in N above γ_c . The above scenario has been checked for both complete and incomplete local algorithms, thanks to the existence of a halting criterion in the unfrustrated version of

hyper-SAT where at least one solution is guaranteed to exist. The use of this model as a benchmark for heuristic algorithms may result in a good improvement of their performances in the phase where many local minima are present.

hSAT can be viewed as an intermediate model between 2-SAT and K -SAT ($K > 2$), which is exactly solvable and in which the presence of hidden solutions can be kept under some control. Hopefully, some of the results and methods of

our analysis of hSAT can be extended to NP-complete problems.

ACKNOWLEDGMENTS

We thank J. Berg, S. Franz, M. Leone, and D.B. Wilson for very useful discussions. M.W. acknowledges the hospitality of the ICTP where the majority of this work was done.

-
- [1] M. Garey and D.S. Johnson, *Computers and Intractability; A Guide to the Theory of NP-Completeness* (Freeman, San Francisco, 1979); C. Papadimitriou, *Computational Complexity* (Addison-Wesley, Reading, MA 1994).
- [2] S.A. Cook, in *Proceedings of the Third Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, 1971), p. 151.
- [3] Y.-T. Fu and P.W. Anderson, in *Lectures in the Sciences of Complexity*, edited by D. Stein (Addison-Wesley, Reading, MA, 1989), p. 815.
- [4] D. Mitchell, B. Selman, and H. Levesque, in *Proceedings of the Tenth National Conference on Artificial Intelligence (AAAI-92)* (AAAI Press, Menlo Park, CA, 1992) p. 459.
- [5] S. Kirkpatrick and B. Selman, *Science* **264**, 1297 (1994).
- [6] *Artif. Intel.* **81**, issue 1-2 (1996).
- [7] R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, and L. Troyansky, *Nature (London)* **400**, 133 (1999).
- [8] M. Weigt and A.K. Hartmann, *Phys. Rev. Lett.* **84**, 6118 (2000).
- [9] D. Achlioptas, C. Gomes, D. Liang, H. Kautz, and B. Selman (unpublished).
- [10] B. Bollobás, C. Borgs, J.T. Chayes, J.H. Kim, and D.B. Wilson, e-print arXiv:math.CO/9909031.
- [11] In this context, the study of Boolean satisfiability (K -SAT) problems has played a major role, allowing for both theoretical and numerical analysis. An instance of random K -SAT consists of a set of $M = \alpha N$ random logical clauses over N Boolean variables. Each clause contains exactly K variables which are connected by logical OR operations and appear negated with probability $1/2$. The important computational question is whether there exists an assignment to the variables that simultaneously satisfies all clauses (“constraints”) for a given instance. When α crosses a critical value $\alpha_c(K)$ and for $N \gg 1$, the probability of finding solutions vanishes abruptly, i.e., $\alpha_c(K)$ identifies the so-called satisfiable to unsatisfiable (SAT/UNSAT) transition. At the same point and for $K \geq 3$, hard computational instances cluster, leading to an exponential median search cost for the most efficient known algorithms. The random K -SAT problem therefore provides a good model for the study of the onset of true intractability of “typical” instances of NP-complete problems. Moreover, recent results [7] have pointed out a clear connection between typical-case computational complexity and the type of threshold phenomena taking place at $\alpha_c(K)$.
- [12] A. Goerdt, in *Proceedings of the 17th International Symposium on Mathematical Foundations of Computer Science* in Lecture Notes in Computer Science Vol. 629 (Springer, Berlin, 1992), p. 264; *J. Comput. Syst. Sci.* **53**, 469 (1996); V. Chvátal and B. Reed, in *Proceedings of the 33rd IEEE Symposium on Foundations of Computer Science* (IEEE, New York, 1992), p. 620.
- [13] B. Aspöqvist, M.F. Plass, and R.E. Tarjan, *Inf. Process. Lett.* **8**, 121 (1979).
- [14] M. Mézard, G. Parisi, and M.A. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1987).
- [15] D. Gross and M. Mézard, *Nucl. Phys. B* **240**, 431 (1984).
- [16] E. Gardner, *Nucl. Phys. B* **257**, 747 (1985).
- [17] E. Gardner and B. Derrida, *J. Phys. A* **22**, 1983 (1989).
- [18] T.R. Kirkpatrick and D. Thirumalai, *Phys. Rev. Lett.* **58**, 2091 (1987).
- [19] A. Barrat and R. Zecchina, *Phys. Rev. E* **59**, 1299 (1999).
- [20] Similarly to the hSAT model, also random K -SAT can be mapped onto a minimization problem of a spin cost function. Keeping the same definition of the spin variables one can construct a function which counts the number of violated clauses. In the $K=3$ case such a function reads
- $$H[\Delta, S] = \frac{\alpha}{8} N - \sum_{i=1}^N H_i S_i + \sum_{i < j} T_{ij} S_i S_j - \sum_{i < j < k} J_{ijk} S_i S_j S_k, \quad (19)$$
- where the couplings are given by $H_i = \frac{1}{8} \sum_{\ell=1}^M \Delta_{\ell, i}$, $T_{ij} = \frac{1}{8} \sum_{\ell=1}^M \Delta_{\ell, i} \Delta_{\ell, j}$; and $J_{ijk} = \frac{1}{8} \sum_{\ell=1}^M \Delta_{\ell, i} \Delta_{\ell, j} \Delta_{\ell, k}$. The random parameters $\Delta_{\ell, i}$ take the value 1 (respectively -1) if the variable x_i (respectively \bar{x}_i) appears in the ℓ th clause, and 0 otherwise. Comparison of the above expression with Eq. (3) clarifies the difference with hSAT in which only the 3-spin interaction term appears.
- [21] D. Achlioptas (unpublished).
- [22] B. Selman, H.A. Kautz and B. Cohen, *Cliques, Coloring, and Satisfiability*, edited by D.S. Johnson and M.A. Trick, DIMACS Series in Discrete Mathematics and Theoretical Computer Science Vol. 26 (AMS, Providence, RI, 1996).
- [23] J.M. Crawford and L.D. Auton, in *Proceedings of the Eleventh National Conference on Artificial Intelligence (AAAI-93)* (AAI Press, Menlo Park, CA, 1993), p. 21.
- [24] The Boolean expressions which are generated in this way are simply those which have the highest probability given the satisfying assignment and therefore are those with maximum entropy, i.e., the easiest to satisfy.
- [25] G. Biroli, R. Monasson, and M. Weigt, *Eur. Phys. J. B* **14**, 551 (2000).
- [26] However, the entropy is not trivially given by $\ln(2)$ times the fraction of variables which do not belong to the backbone.
- [27] M. Davis and H. Putnam, *J. Assoc. Comput. Mach.* **7**, 201 (1960).

- [28] R. Monasson and R. Zecchina, Phys. Rev. Lett. **76**, 3881 (1996); Phys. Rev. E **56**, 1357 (1997).
- [29] B. Bollobás, *Random Graphs* (Academic Press, London, 1985).
- [30] Note that in random hypergraphs the probability that any pair of vertices belongs to more than one hyperlink is proportional to $1/N$ for large N , we have not considered this possibility when constructing the structures in Fig. 1.
- [31] Taking also the gauge into account we have $J_{ijk} = \varepsilon_i \varepsilon_j \varepsilon_k$ and the constraint becomes $S_{\ell} = \varepsilon_{\ell}$.
- [32] M.-T. Chao and J. Franco, Inf. Sci. (N.Y.) **51**, 289 (1990).
- [33] D.B. Wilson, e-print arXiv:math/0005136.
- [34] The size of the critical window has been estimated as the inverse of the derivative at γ_c of the curves in Fig. 2. Other methods give compatible results.
- [35] J. Berg, S. Franz, M. Leone, F. Ricci-Tersenghi, and R. Zecchina (unpublished).
- [36] In any SAT phase proving satisfiability is equivalent to finding a ground state whereas in an UNSAT phase the proof of unsatisfiability does not provide any information on the precise assignments of the variables in the ground states.
- [37] Similarly to what happens in Max-KSAT. See, for instance, P. Beame, R. Karp, T. Pitassi, and M. Saks, in *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing* (ACM, New York, 1998), p. 561.