Relaxation and metastability in a local search procedure for the random satisfiability problem

Guilhem Semerjian[1,*] and Rémi Monasson[1,2,†]
[1]*CNRS-Laboratoire de Physique Théorique de l'ENS, 24 rue Lhomond, 75005 Paris, France*
[2]*CNRS-Laboratoire de Physique Théorique, 3 rue de l'Université, 67000 Strasbourg, France*
(Received 15 January 2003; published 12 June 2003)

An analysis of the average properties of a local search procedure (RandomWalkSAT) for the satisfaction of random Boolean constraints is presented. Depending on the ratio α of constraints per variable, reaching a solution takes a time T_{res} growing linearly [$T_{res} \sim \tau_{res}(\alpha)N$, $\alpha < \alpha_d$] or exponentially ($T_{res} \sim \exp[N\zeta(\alpha)]$, $\alpha > \alpha_d$) with the size N of the instance. The relaxation time $\tau_{res}(\alpha)$ in the linear phase is calculated through a systematic expansion scheme based on a quantum formulation of the evolution operator. For $\alpha > \alpha_d$, the system is trapped in some metastable state, and resolution occurs from escape from this state through crossing of a large barrier. An annealed calculation of the height $\zeta(\alpha)$ of this barrier is proposed. The polynomial to exponential cross-over $\alpha_d \simeq 2.7$ is not related to the onset of clustering among solutions occurring at $\alpha \simeq 3.86$.

DOI: 10.1103/PhysRevE.67.066103 PACS number(s): 02.50.Ga, 05.40.−a, 89.20.Ff, 05.20.−y

I. INTRODUCTION

The study of combinatorial problems [1] with statistical physics techniques started almost twenty years ago [2]. Most of the efforts have been devoted to the calculation of the optimal solution of various problems (traveling salesman, matching, graph or number partitioning, satisfiability of Boolean constraints, vertex cover of graphs, etc.) as a function of the definition parameters of their inputs distributions. Central to these studies is the characterization of the properties of the extrema of correlated random variables, a question of considerable importance in probability theory [3]. From a computer science point of view, however, the main point is the characterization of solving times. Concepts and tools issued from the analysis of algorithms have allowed so far to understand the behavior and the efficiency of many algorithms of practical use [4,5], sorting for instance, but progress has been much slower in the analysis of search procedures for combinatorial problems. These are sophisticated algorithms hardly amenable to rigorous analysis with available techniques [6]. Statistical physics ideas and approximation techniques may then be of great relevance to help develop quantitative understanding and intuition about the operation of these algorithms. To some extent, the study of algorithms may be seen as a part of out-of-equilibrium statistical physics.

There exists a wide variety of algorithms for combinatorial problems [1]. Roughly speaking, two main classes may be identified. The first one includes complete algorithms, guaranteed to provide the optimal solution. They essentially proceed through an exhaustive albeit clever (making use of branch-and-bound procedures) search through the configuration space, and may require very large computational times, i.e., scaling exponentially with the size of the inputs to be treated. Recently, notions borrowed from statistical physics such as real-space renormalization and out-of-equilibrium growth processes made it possible to reach some understanding of the operation of complete algorithms for the satisfiability [7,8] and for the vertex cover [9,10] problems over random classes of inputs. Incomplete algorithms constitute another large class of solving procedures; they may be able to find the optimal solution very quickly, but may also run forever without ever finding it. An example is provided by local procedures, i.e., Monte Carlo dynamics, which attempt to find a solution from an arbitrary initial configuration through a sequence of stochastic local moves in the configuration space. When specialized to decision problems (for which the desired output is the answer YES or NO to a question related to the inputs, as: is there a way to color a given graph with seven colors only?), local search algorithms can sometimes be made one-sided error probabilistic algorithms [11]. When they stop, the answer is YES with certainty. If they run for a time t without halting, the probability (over the nondeterministic choices of local moves in the configuration space) that the correct answer is NO is bounded from below by a function $f(t,N)$ tends to 1 when $t \to \infty$. Obviously this function depends on the size N of the input: the larger the input, the larger the time $t_c(N)$ it takes to reach, say, a 99% confidence that the answer is NO, i.e., $f(t_c(N),N) = 0.99$. Determining the scaling (polynomial or exponential) of $t_c(N)$ with N is of capital importance to assess the efficiency of the algorithm.

In this paper, we study this question for a local search procedure, the RandomWalkSAT algorithm [12], and a decision problem defined over an easy-to-parametrize class of inputs, the random satisfiability problem [13]. Both procedure and problem are defined in Sec. II. We also recall the main results proven by mathematicians on these issues, and present an overview of the phenomenology of RandomWalkSAT. A major theoretical interest for studying the RandomWalkSAT algorithm is that it is, apparently, of purely dynamical nature. Detailed balance is indeed verified in a trivial way; the equilibrium measure is nonzero over solutions only, and the transition rates from a solution to any other configuration are null. The equilibrium measure is therefore of no use to understand long-time dynamics, a situation reminiscent of some models studied in out-of-equilibrium physics, e.g., the contact process for finite-size

*Electronic address: guilhem@lpt.ens.fr
†Electronic address: monasson@lpt.ens.fr

systems [14]. We are thus left with a study of the dynamical evolution of a spin system with disorder in the interactions (the instances of the combinatorial problem to be solved are random), a still largely open problem in statistical physics [15]. We show in Sec. III how the master equation for this evolution can be written as a Hamiltonian (in imaginary time) for 1/2 quantum spin systems, and use this representation to get exact results and develop systematic expansions for the quantities of interest, valid in some region of the parameter space. In Sec. IV, we present an approximate analysis of the RandomWalkSAT procedure in the whole parameter space. We show that depending on the value of the ratio α of the number of constraints per (Boolean) degree of freedom, resolution is either achieved in linear time or requires the escape from some metastable region in the configuration space, a slow process taking place over exponentially large times. Interestingly, the dynamics generated by RandomWalkSAT is very similar to the physical dynamics of (spin) glassy systems [15]. Some perspectives are presented in Sec. V. Note that a complementary study of the RandomWalkSAT procedure was very recently carried out by Barthel, Hartmann, and Weigt [16].

II. DEFINITIONS, KNOWN RESULTS AND PHENOMENOLOGY

A. The random *K*-Satisfiability problem

The 3-Satisfiability (3-SAT) decision problem is defined as follows. Consider a set of N boolean variables x_i, $i=1,\ldots,N$. A literal is either a variable x_i or its negation $\bar{x}_i$. A clause is the logical OR between three distinct literals. It is thus true as soon as one of the literals is true. A formula is the logical AND between M clauses, it is true if and only if all the clauses are true. A formula is said to be *satisfiable* if there is an assignment of the variables such that the formula is true, *unsatisfiable* otherwise.

3-SAT is an NP-complete problem [1]; it is believed that there is no algorithm capable of solving every instance of 3-SAT in a time bounded from above by a polynomial of the size of the instance. How well do existing and *a priori* exponential algorithms perform in practice? To answer these questions, computer scientists have devised a simple way of generating random instances of the 3-SAT problem, with a rich pattern of hardness [13]. Formulas are drawn in the following way. Repeat M times independently the same process; pick up a 3-uplet of distinct indices in $[1,N]$, uniformly on all possible three-uplets. For each of the three corresponding variables, choose the variable itself or its negation with equal probability (1/2), and construct a clause with the chosen literals. Repetition of this process M times gives a set of M independently chosen clauses, whose conjunction is the generated instance. The random generation of formulas makes the set of possible formulas a probability space, with a well-defined measure.

Numerical experiments indicate that a phase transition takes place when $N,M\to\infty$ at fixed ratio $\alpha=M/N$ of clauses per variables (in the thermodynamic limit). If α is smaller than some critical value $\alpha_c\simeq 4.3$, a randomly drawn instance admits at least one solution with high probability. Beyond this threshold, instances are almost never satisfiable. The existence of this transition has not been proven rigorously yet [17], but bounds on the threshold exist; the probability of satisfaction tends to 1 (respectively to 0) if $\alpha<3.42$ [18] (respectively if $\alpha>4.506$ [19]) All the above definitions can be extended to the K-SAT problem, where each clause is the disjunction of K, rather than 3, literals. 2-SAT is an easy (polynomial) problem, while K-SAT is an NP-complete problem for any $K\geqslant 3$. Location of the threshold is rigorously known for 2-SAT ($\alpha_c=1$) [20], but not for $K\geqslant 3$. We shall denote in the following averages on the random K-SAT ensemble by $[\cdot]$.

Statistical mechanics studies have pointed out the existence of another phase transition taking place in the satisfiable phase ($\alpha<\alpha_c$) with a location first estimated at $\alpha_s\simeq 3.95$ [21] and later at $\alpha_s\simeq 3.86$ [22]. This phase transition is related to the microscopic structure of the set of solutions. Define d, the Hamming distance between two solutions as the number of variables taking opposite values in these solutions. When $\alpha<\alpha_s$, there exist an exponentially large number of solutions (N), each pair of which is separated by a path in the solution space, that is, a sequence of solutions with a $O(1)$ Hamming distance between successive solutions along the path. The solution space is made up of a single cluster of solutions. For $\alpha_s<\alpha<\alpha_c$, the solution space breaks into an exponentially large number of clusters (N), separated by large voids without solutions. Two solutions in one cluster are linked through a path, while there is no path in the solution space linking two solutions in two different clusters. This clustering phenomenon, whose discovery was inspired from previous works in the context of information storage in neural networks [24], was subsequently found in various combinatorial problems [25,26], and rigorously demonstrated for the so-called random XOR-SAT problem [27]. It is a zero-temperature signature of the ergodicity breaking taking place in spin glasses [15,28]. Its precise relationship with dynamical properties, and in particular with the computational cost for finding a solution is not fully elucidated yet [21,23,29].

B. The RandomWalkSAT algorithm

The operation of RandomWalkSAT (also called RandomWalk) on an instance of the K-SAT problem is as follows [12].

(1) Choose randomly a configuration of the Boolean variables.

(2) If all clauses are satisfied, output ''satisfiable.''

(3) If not, choose randomly one of the unsatisfied clauses and one of the K variables of this clause. Flip the variable. Notice that the selected clause is now satisfied, but the flip operation may have violated other clauses that were previously satisfied.

(4) Go to step 2, until a limit on the number of flips fixed beforehand has been reached. Then output ''don't know.''

What is the output of the algorithm? Either ''satisfiable'' and a solution is exhibited, or ''don't know'' and no certainty

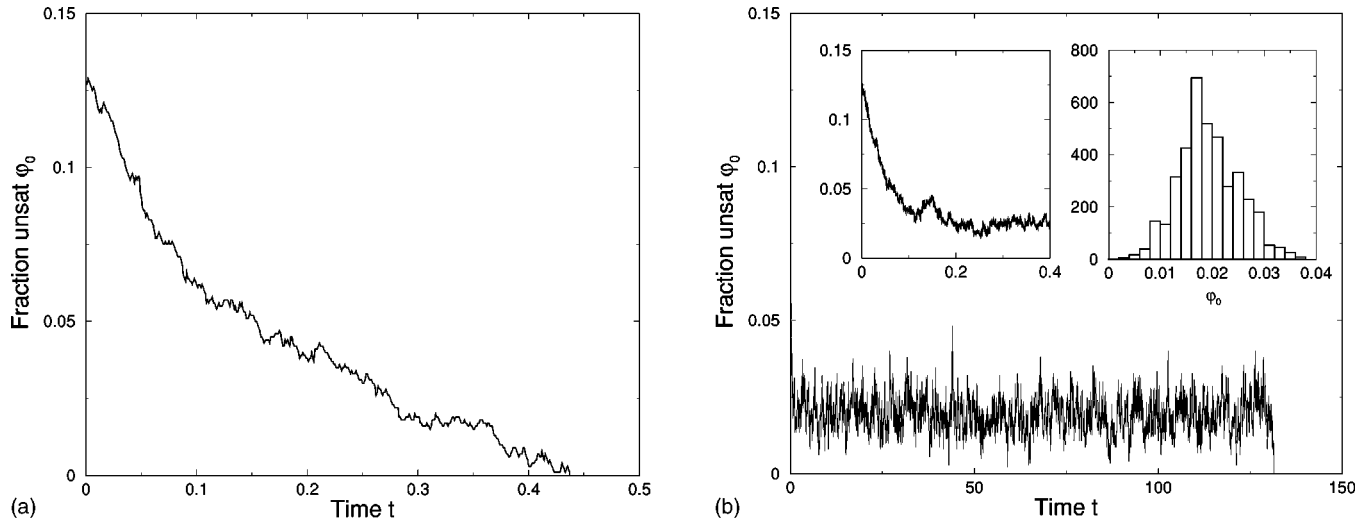

FIG. 1. Fraction φ_0 of unsatisfied clauses as a function of time t (number of flips over M) for two randomly drawn instances of 3-SAT with ratios $\alpha=2$ (a) and $\alpha=3$ (b) with $N=500$ variables. Note the difference of time scales between the two figures. Insets of (b): left, blow up of the initial relaxation of φ_0, taking place on the $O(1)$ time scale as in (a); right, histogram $p_{500}(\varphi_0)$ of the fluctuations of φ_0 on the plateau $1\leqslant t\leqslant 130$.

on the status of the formula is reached. Papadimitriou showed that RandomWalkSAT solves with high probability any satisfiable 2-SAT instance in a number of steps (flips) of the order of N^2 [12]. Recently, Schöning was able to prove the following very interesting result for 3-SAT [30]. Call ''trial'' a run of RandomWalkSAT consisting of the random choice of an initial configuration followed by $3\times N$ steps of the procedure. If none of T successive trials of RandomWalkSAT on a given instance has been successful (has provided a solution), then the probability that this instance is satisfiable is lower than $\exp[-T\times(3/4)^N]$. In other words, after $T\gg(4/3)^N$ trials of RandomWalkSAT, most of the configuration space has been ''probed''; if there were a solution, it would have been found. Though RandomWalkSAT is not a complete algorithm, the uncertainty on its output can be made as small as possible and it can be used to prove unsatisfiability (in a probabilistic sense).

Schöning's bound is true for any instance. Restriction to special input distributions allows one to improve his result. Alekhnovich and Ben-Sasson showed that instances drawn from the random 3-Satisfiability ensemble described above are solved in polynomial time with high probability when α is smaller than 1.63 [31]. It is remarkable that despite the quenched character of the disorder in this problem (the same clauses are seen various times in the course of the search), rigorous results on the dynamics of this spin model can be achieved.

C. Phenomenology of the operation of RandomWalkSAT

In this section, we briefly sketch the behavior of RandomWalkSAT, as seen from numerical experiments [32] and the analysis presented later in this paper. We find that there is a dynamical threshold α_d separating the following two regimes.

(1) For $\alpha<\alpha_d\simeq 2.7$ for 3-SAT, the algorithm finds a solution very quickly, namely, with a number of flips growing linearly with the number of variables N. Figure 1(a) shows the plot of the fraction φ_0 of unsatisfied clauses as a function of time t (number of flips divided by M) for one instance with ratio $\alpha=2$ and $N=500$ variables. The curves show a fast decrease from the initial value $[\varphi_0(t=0)=1/8$ in the large-N limit independent of $\alpha]$ down to zero on a time scale $t_{res}=O(1)$. Fluctuations become smaller and smaller as N grows. t_{res} is an increasing function of α. This *relaxation* regime corresponds to the one studied by Alekhnovich and Ben-Sasson, and $\alpha_d>1.63$ as expected [31]. Figure 2(a) symbolizes the behavior of the system in the relaxation regime.

(2) For instances in the $\alpha_d<\alpha<\alpha_c$ range, the initial relaxation phase taking place on a $t=O(1)$ time scale is not sufficient to reach a solution [Fig. 1(b)]. The fraction φ_0 of unsat clauses then fluctuates around some plateau value for a very long time. On the plateau, the system is trapped in a *metastable* state. The lifetime of this metastable state (trapping time) is so large that it is possible to define a (quasi) equilibrium probability distribution $p_N(\varphi_0)$ for the fraction φ_0 of unsat clauses. [Inset of Fig. 1(b)]. The distribution of fractions is well peaked around some average value (height of the plateau), with left and right tails decreasing exponentially fast with N, $p_N(\varphi_0)\sim\exp[N\overline{\zeta}(\varphi_0)]$ with $\overline{\zeta}\leqslant 0$ [Fig. 2(b)]. Eventually a large negative fluctuation will bring the system to a solution $(\varphi_0=0)$. Assuming that these fluctuations are independent random events occurring with probability $p_N(0)$ on an interval of time of order one, the solving time is a stochastic variable with exponential distribution. Its average is, to leading exponential order, the inverse of the probability that a solution is obtained on the $O(1)$ time scale: $[t_{res}]\sim\exp(N\zeta)$ with $\zeta=-\overline{\zeta}(0)$. Escape from the metastable state therefore occurs through barrier crossing and takes place on exponentially large-in-N time scales, as confirmed by numerical simulations for different sizes. Schöning's result [30] can be interpreted as a lower bound to

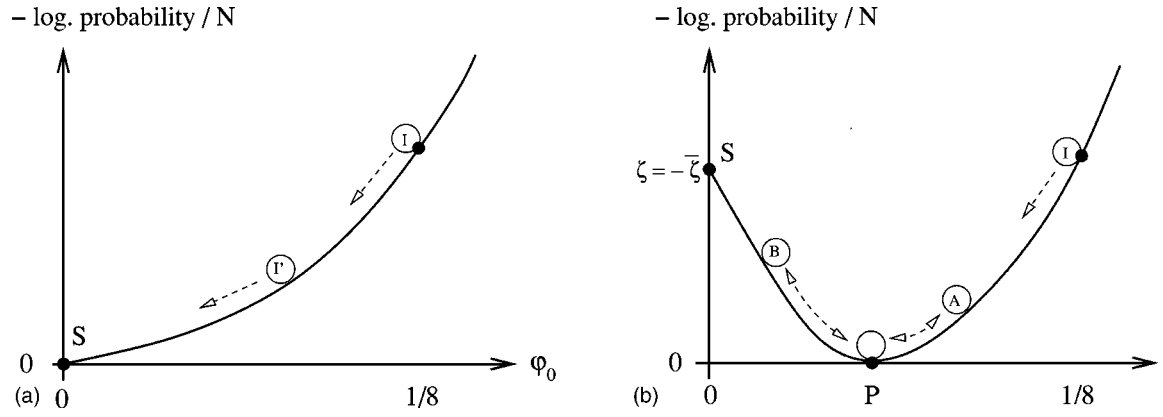

FIG. 2. A schematic picture of the operation of RandomWalkSAT for ratios α smaller (a) or larger (b) than the dynamical threshold α_d. Vertical axis is minus the logarithm of the probability (divided by N) that the system has a fraction φ_0 of unsat clauses after a large number of RandomWalkSAT flips. Its representative curve can be seen as an energy potential in which the configuration (represented by an empty ball) rolls down towards the more probable value of the order parameter φ_0, or up through stochastic fluctuations. The starting configuration violates $\varphi_0 = 1/8$ of the clauses (point I). At small ratios (a), the configuration rolls down to reach point S through a sequence of intermediary points (I'). The search for a solution is essentially a fast relaxation towards S [$O(N)$ time scale]. At large ratios (b), the ball first relaxes to the bottom of the well (point P with abscissa corresponding to the plateau height). Then slow negative (ball A) or positive (ball B) fluctuations of the fraction φ_0 take place on exponentially (N) long time scales. The time [$t_{res} \sim \exp(N\zeta)$] it takes to the system to reach a solution (point S) is, on the average, equal to the inverse probability $\exp(N\overline{\zeta})$ that a fluctuation drives the system to S.

the probability $\overline{\zeta}(0) > \ln(3/4)$, true for any instance.

The plateau energy, that is, the fraction of unsatisfied clauses reached by RandomWalkSAT on the linear time scale is plotted in Fig. 3. Notice that the ''dynamic'' critical value α_d above which the plateau energy is positive (RandomWalkSAT stops finding a solution in linear time) is strictly smaller than the ''static'' ratio α_c, where formulas go from satisfiable with high probability to unsatisfiable with high probability. In the intermediate range $\alpha_d < \alpha < \alpha_c$, instances are almost surely satisfiable, but RandomWalkSAT needs an exponentially large time to prove so. Interestingly, α_d and α_c coincides, for 2-SAT in agreement with Papadimitriou's result [12]. Furthermore, the dynamical transition is apparently not related to the onset of clustering taking place at α_s.

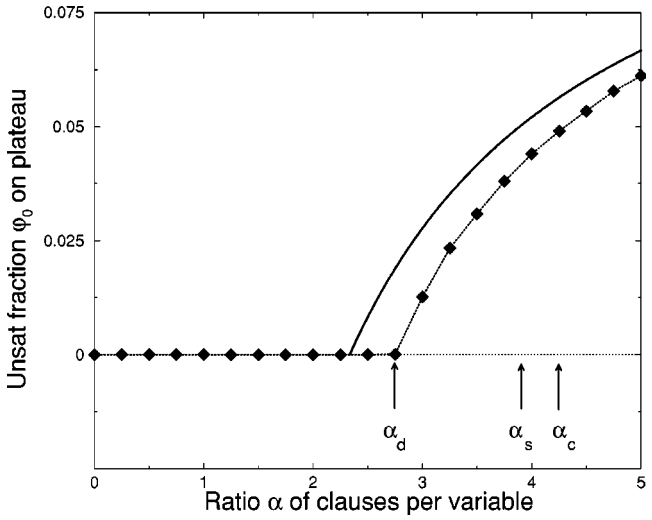

FIG. 3. Fraction φ_0 of unsatisfied clauses on the metastable plateau as a function of the ratio α of clauses per variable. Diamonds are the output of numerical experiment, and have been obtained through average of data from simulations at a given size N (number of variables) over 1 000 samples of 3-SAT, and extrapolation to infinite sizes (dotted line serves as a guide to the eye). The ratio at which φ_0 begins being positive, $\alpha_d \simeq 2.7$, is smaller than the thresholds $\alpha_s \simeq 3.9$ and $\alpha_c \simeq 4.3$ above which solutions gather into distinct clusters and instances have almost surely no solution, respectively. The full line represents the prediction of the Markovian approximations of Secs. IV C and IV E.

III. EXACT RESULTS: SPECIAL CASES AND EXPANSIONS

A. Evolution equations and quantum formalism

Boolean variables will be hereafter represented by Ising spins $S_i = 1$ (respectively -1) when the Boolean variable x_i is true (respectively false). A microscopic configuration $\mathbf{S}$ is specified by the states of all variables: $\mathbf{S} = (S_1, S_2, \ldots, S_N)$. We then define a 2^N-dimensional linear space with canonical basis $\{|\mathbf{S}\rangle\}$, orthonormal for the scalar product $\langle \mathbf{S}'|\mathbf{S}\rangle = \Pi_i \delta_{S_i', S_i}$. Let us denote Prob$[\mathbf{S}, T]$ the probability that the system configuration is $\mathbf{S}$ at time T, i.e., after T steps of the algorithm, and define [33]

$$|\mathbf{S}(T)\rangle = \sum_{\mathbf{S}} \text{Prob}[\mathbf{S}, T]|\mathbf{S}\rangle, \tag{1}$$

as the (time-dependent) vectorial state of the system. Knowledge of this vector gives access to the probability Prob$[\mathbf{S}, T] = \langle \mathbf{S}|\mathbf{S}(T)\rangle$ of being in a certain configuration $\mathbf{S}$.

RandomWalkSAT defines a Markov process on the set of configurations; during one step of the algorithm, the state vector of the system changes according to

$$|\mathbf{S}(T+1)\rangle = \hat{W}_d|\mathbf{S}(T)\rangle, \tag{2}$$

where the evolution operator in discrete time $\hat{W}_d$ reads as

$$\hat{W}_d = \sum_{\ell=1}^{M} \hat{F}_\ell \cdot \hat{U}_\ell \cdot \hat{E}^{-1}, \quad \hat{F}_\ell = \frac{1}{K} \sum_{i=1}^{N} C_{li}^2 \sigma_i^x,$$

$$\hat{U}_\ell = g_K\left(\sum_{j=1}^{N} C_{lj} \sigma_j^z \right), \quad \hat{E} = \sum_{\ell=1}^{M} \hat{U}_\ell . \tag{3}$$

In the above expression, we have made use of several different notations that we now explain. The $M \times N$ matrix $C_{\ell i}$ encodes the instance, $C_{\ell i}$ equals 1 if the ℓth clause contains the literal x_i, -1 if it contains the literal $\overline{x_i}$, and 0 otherwise. Since every clause contains K literals, $\Sigma_i C_{\ell i}^2 = K \ \forall \ell$. The Pauli operators are defined through $\sigma_i^z |\mathbf{S}\rangle = S_i |\mathbf{S}\rangle$ and $\sigma_i^x |\mathbf{S}\rangle = |\mathbf{S}^i\rangle$, where $\mathbf{S}^i$ is the configuration obtained from $\mathbf{S}$ by flipping the ith spin. It is a simple check that the argument of function g_K in $\hat{U}_\ell$ is a diagonal operator in the canonical basis $\{|\mathbf{S}\rangle\}$, with eigenvalues $x \equiv \Sigma_{i=1}^{N} C_{\ell i} S_i$ in $\{-K, -K+2, \ldots, K-2, K\}$, equal to $-K$ if and only if clause ℓ is unsatisfied. The function

$$g_K(x) = \delta_{x,-K} = \frac{1}{2^K K!} \prod_{p=0}^{K-1} (K - 2p - x) \tag{4}$$

is a polynomial of degree K in x equal to 1 if $x = -K$ and to zero for all the other possible eigenvalues. Operator $\hat{E}$ is diagonal in the canonical basis too, with eigenvalues equal to the numbers of unsatisfied clauses (also called energy) in each configuration. Therefore, $\hat{U}_\ell \cdot \hat{E}^{-1}$ in Eq. (3) acts as a filter retaining clause ℓ only if it is unsatisfied, and with a weight equal to the inverse number of unsatisfied clauses; only unsat clauses can be chosen at each time step, each of these with the same probability. $\hat{F}_\ell$ flips the spins of clause ℓ, each with probability $1/K$. Note that $\hat{W}_d$ conserves probabilities: $\langle O | \hat{W}_d = \langle O |$ where $\langle O | = \Sigma_{\mathbf{S}} \langle \mathbf{S}|$ is the superposition of all possible states.

In the thermodynamic limit, the evolution can be rewritten in continuous time, defining $t = T/M$,

$$\frac{d}{dt} |\mathbf{S}(t)\rangle = \hat{W} |\mathbf{S}(t)\rangle, \quad \hat{W} = M(\hat{W}_d - \hat{1}). \tag{5}$$

Formally, the solution of this equation is $|\mathbf{S}(t)\rangle = e^{t\hat{W}} |\mathbf{S}(0)\rangle$, where $|\mathbf{S}(0)\rangle = (1/2^N) \Sigma_{\mathbf{S}} |\mathbf{S}\rangle$ since the initial configuration is random. An important quantity to compute is the average fraction of unsatisfied clauses at time t, i.e., after $T = Mt$ steps of the algorithm, averaged both on the history of the algorithm and on the distribution of formulas:

$$\varphi_0(t) = \frac{1}{M} [\langle O | \hat{E} e^{t\hat{W}} | \mathbf{S}(0) \rangle]. \tag{6}$$

For the sake of analytical simplicity, we shall study a slightly different evolution operator in the next two sections, and use the parameter u to denote the time parameter of this modified evolution,

$$\hat{W}' = \hat{W} \cdot \frac{1}{M} \hat{E} = \sum_{\ell=1}^{M} (\hat{F}_\ell - \hat{1}) \cdot \hat{U}_\ell , \quad \frac{d}{du} |\mathbf{S}(u)\rangle = \hat{W}' |\mathbf{S}(u)\rangle. \tag{7}$$

Let us explain the meaning of this modification. The operator $\hat{W}'$ would have been obtained, starting from the following variant of the RandomWalkSAT stochastic process. At each step, choose a clause among the M ones. If it is satisfied, do nothing. If it is unsatisfied, flip one of its variables.

In the thermodynamic limit, the fraction of unsatisfied clauses φ_0 is expected to become a self-averaging quantity, i.e., to be peaked with high probability around its (t-dependent) mean value. Turning $\hat{W}$ into $\hat{W}'$ thus amounts to a local redefinition of time, independent of the instance of the problem. In definition (7), the operator $\hat{E}/M$ can be replaced with its mean value φ_0, leading to $(d/du) \equiv \varphi_0 d/dt$. The knowledge of the $\hat{W}'$—evolution of any observable in terms of u can then be rewritten in terms of t through

$$t(u) = \int_0^u du' \varphi_0(u'). \tag{8}$$

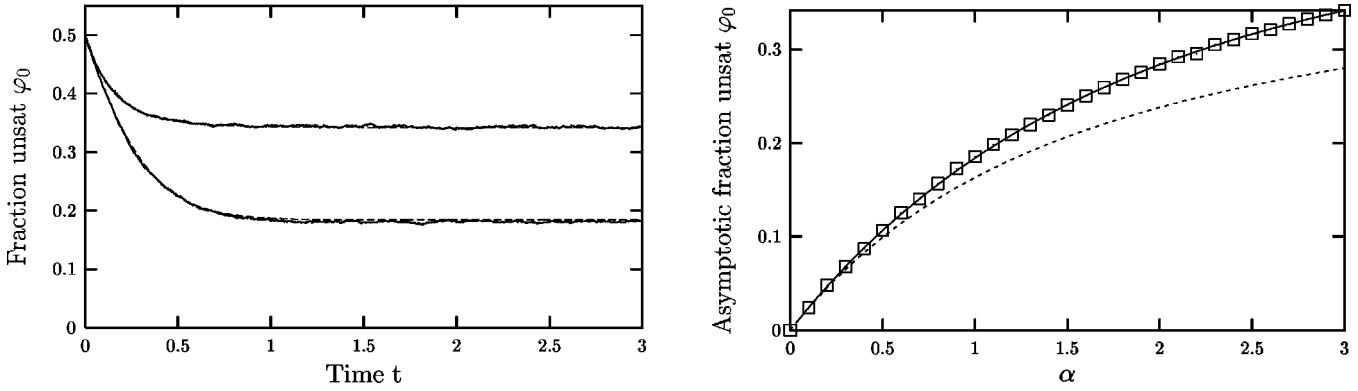

FIG. 4. A comparison between numerics and analytical results for $K = 1$. Left panel: fraction of unsat clauses $\varphi_0(t)$ as a function of time t for $\alpha = 3$ (top) and $\alpha = 1$ (bottom). Dashed line: analytical curve, solid line (almost superimposed): numerical results (single run with $N = 10^4$). Right panel: asymptotic fraction of unsat clauses as a function of α. Solid line: theoretical prediction (13), symbols: numerical results (averaged over ten runs, over $t \in [8,10]$, for $N = 10^4$), dashed line: ground-state energy φ_{GS}.

Examples of this time reparametrization will be shown below.

B. The $K=1$ case

Let us first study the simple case $K=1$. A clause is then a single literal, i.e., either a variable or its negation. If both a variable and its negation appear in the formula, it is obviously unsatisfiable. This is the case, with high probability in the thermodynamic limit, as soon as $\alpha>0$. Static properties of the 1-SAT model are known exactly [34], and its dynamics under the RandomWalkSAT evolution can be solved too.

An instance is described by a set of integers $\{m_i,n_i\}$, where m_i is the number of clauses in which the variable i appears, and $(m_i-n_i)/2$ is the number of times it has been chosen without negation. The evolution operator in terms of u—time is the sum of site operators,

$$\hat{W}_i'=\frac{1}{2}(m_i\sigma_i^x-m_i+n_i\sigma_i^z-n_i\sigma_i^x\sigma_i^z). \tag{9}$$

As the operators $\hat{W}_i'$ on different sites commute, the evolution operator can be diagonalized in each of the i subsets independently, and the vector state is a tensor product,

$$|\mathbf{S}(u)\rangle=\otimes_{i=1}^{N}\begin{cases}\frac{1}{2}(|+\rangle_i+|-\rangle_i) & \text{if } m_i=0\\ \frac{1}{2}\sum_{S=\pm1}\left(1+S\frac{n_i}{m_i}(1-e^{-m_iu})\right)|S\rangle_i & \text{if } m_i\neq0,\end{cases} \tag{10}$$

where $|\pm\rangle_i$ are the eigenvectors of σ_i^z with eigenvalues ±1. The fraction of unsatisfied clauses for a given instance and averaged over the choices of the algorithm reads as

$$\frac{1}{M}\langle O|\hat{E}|\mathbf{S}(u)\rangle=\frac{1}{2}-\frac{1}{2M}\sum_{i/m_i\neq0}\frac{n_i^2}{m_i}(1-e^{-m_iu}). \tag{11}$$

In the thermodynamic limit, the m_i's become independent random variables with identical Poisson distributions of parameter α; $(m_i-n_i)/2$ obeys a Binomial law of parameter 1/2 among m_i. Performing the quenched average over the formula, the fraction of unsatisfied clauses reads

$$\varphi_0(u)=\frac{1}{2}-\frac{1-\exp[\alpha(e^{-u}-1)]}{2\alpha}. \tag{12}$$

These exact results are compared with numerical simulations in Fig. 4. On the right panel has been drawn the asymptotic fraction of unsatisfied clauses obtained at large times, compared with the analytical prediction made above. The left panel shows the time evolution of the fraction of unsatisfied clauses as a function of time t, for two different values of α. The analytical curve has been obtained through the rescaling explained at the end of the preceding section; from the exact value of $\varphi_0(u)$ given above, the original time $t(u)$ has been obtained by numerical integration [Eq. (8)], and the plot $\varphi_0(t)$ is a parametric plot $\{t=t(u),\varphi_0=\varphi_0(u)\}$, parametrized by u.

Note that the asymptotic value of the energy reached is larger than the ground-state one [34],

$$\lim_{t\to\infty}\varphi_0(t)=\frac{1}{2}+\frac{e^{-\alpha}-1}{2\alpha}>\frac{1}{2}[1-e^{-\alpha}I_0(\alpha)-e^{-\alpha}I_1(\alpha)]$$
$$=\varphi_{GS}, \tag{13}$$

where I_n is the nth modified Bessel function. This is easily understood: if a formula contains x_1 once and $\overline{x_1}$ twice, the optimal value is x_1 false. But the algorithm does not stop in this configuration, as the clause x_1 is then violated. RandomWalkSAT will keep on flipping the variable and make the average energy higher than the optimal one.

C. The satisfiable phase of the $K=2$ case

We now turn to the 2-SAT case, where every clause involves two variables. It has been rigorously proven that there is a sharp threshold phenomenon taking place at $\alpha_c=1$ in this problem [20]. For lower values of α, the formulas are almost always satisfiable, beyond this threshold they are almost always unsatisfiable. The dynamical critical threshold of RandomWalkSAT has the same value, $\alpha_d=1$ [12]. Thus, there is no metastability in 2-SAT; if a solution of the formula exists it is almost surely found in polynomial time. Two questions of interest are: how fast will the number of unsatisfiable clauses decrease during the evolution of the process and how long will it take the algorithm to find a solution for $\alpha<1$? and what will be the typical energy after a long run of the algorithm when $\alpha>1$? In this and the next sections, we shall address the first of these two questions, and leave the second one to the next section.

As in the $K=1$ case, one can use the quantum formalism and write down the evolution operator W'. There appear both single site terms and couplings between pairs of sites,

$$\hat{W}'=\sum_i\hat{W}_i'+\sum_{i\neq j}\hat{W}_{ij}', \tag{14}$$

$$\hat{W}_i'=\frac{1}{8}(m_i\sigma_i^x-m_i+2n_i\sigma_i^z-n_i\sigma_i^x\sigma_i^z), \tag{15}$$

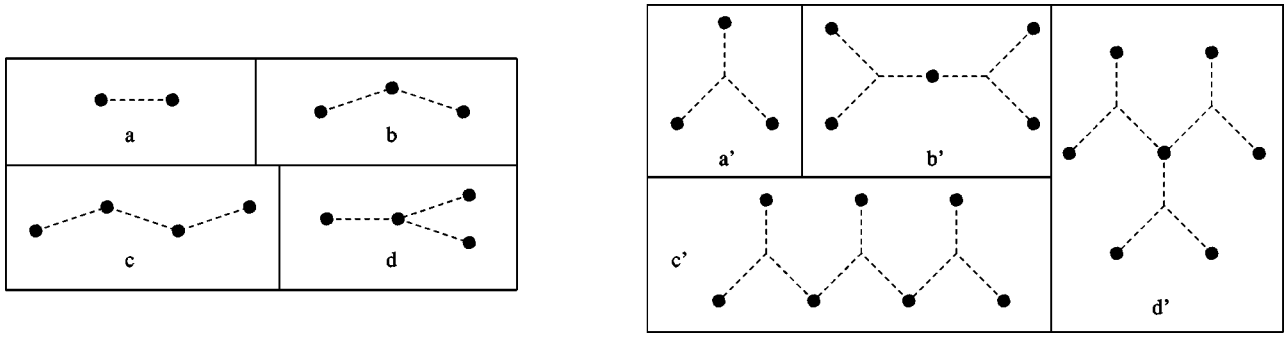

FIG. 5. Treelike clusters contributing to the expansions. Left panel: $K=2$. Right panel: $K=3$.

$$\hat{W}'_{ij}=\frac{1}{8}\left[a_{ij}\left(\frac{\sigma_i^x+\sigma_j^x}{2}-1\right)\sigma_i^z\sigma_j^z-b_{ij}\sigma_i^x\sigma_j^z\right] \tag{16}$$

with

$$n_i=\sum_l C_{li}, \quad m_i=\sum_l C_{li}^2,$$

$$a_{ij}=\sum_l C_{li}C_{lj}, \quad b_{ij}=\sum_l C_{li}^2C_{lj}. \tag{17}$$

Two-site operators $\hat{W}'_{ij}$ do not commute when they share a variable, and $\hat{W}'$ cannot be factorized as in the $K=1$ case. Yet, a cluster expansion in powers of α can be implemented. A detailed presentation of this cluster expansion in a closely related context has been given elsewhere [35]; the reader is referred to this previous work for more details. The method is presented below for a generic value of K. We call cluster and denote F_r a maximal set of variables connected by clauses. Any formula F can be decomposed as a conjunction of these clusters (note that despite the similarities in denominations, these clusters have nothing to do with the clustering phenomenon in configuration space). Consider now a quantity Q, depending on the realization of the disorder, namely, of formula F, and additive with respect to the cluster decomposition $Q(F)=\Sigma_{F_r}Q(F_r)$. Calculation of the average $[Q]$ over the random ensemble formulas of Q can be done as follows. First, consider the different possible topologies s of the clusters and compute $[Q]_s$, the quantity averaged over the choices of signs for a given topology, i.e., whether variables appear negated or not in a clause. Second, add up these contributions with combinatorial factors giving the frequency of appearance of the different topologies,

$$\frac{1}{M}[Q]=\sum_s \frac{1}{\alpha L_s}P_s[Q]_s. \tag{18}$$

The sum runs over the different topologies, L_s is the number of sites in such a cluster, and P_s is the probability of a given site belonging to an s-type cluster. In the thermodynamic limit, the finite-size clusters, which contribute to this sum are treelike: $P_s=(\alpha K!)^{m_s}e^{-L_s\alpha K}K_s$, where $m_s=(L_s-1)/(K-1)$ is the number of clauses of the cluster, and K_s is a symmetry factor. The smallest treelike clusters are shown in Fig. 5. Clauses are represented by dashed lines (for $K=2$) or stars (for $K\geqslant 3$). In principle, such an expansion is meaningful only below the percolation threshold of the underlying random hypergraph, $\alpha_p=1/[K(K-1)]$. Beyond α_p, a giant component appears, whereas the series (18) takes into account clusters of size $O(\ln N)$ only. However, rearranging the series as an expansion in powers of α (expanding out the exponential in P_s) allows one to extend its domain of validity beyond α_p, at least for quantities that are not singular at the percolation threshold.

This method can be employed here. The evolution operator $\hat{W}'$ is the sum of operators for each cluster $\hat{W}'(F_r)$. Operators attached to two distinct clusters commute with each other as they do not have common variables. The energy operator $\hat{E}$ can also be written as a sum over the clusters, thus the energy averaged over the history of the algorithm has the property of additivity over clusters if the time evolution is studied in terms of u. After averaging over the formulas, the fraction of unsatisfied clauses reads as

$$\varphi_0(u)=\sum_s \frac{1}{\alpha L_s}P_s[\langle E\rangle(u)]_s. \tag{19}$$

The evolution and energy operators for a cluster with n variables are $2^n\times 2^n$ matrices, $\hat{W}'_s$ and $\hat{E}_s$, respectively. With the help of a symbolic computation software package, one can easily study the small finite-size clusters by computing

$$[\langle E\rangle(u)]_s=[\langle O|\hat{E}_s e^{\hat{W}'_s u}|\mathbf{S}(0)\rangle]_s, \tag{20}$$

where $\langle O|$ and $|\mathbf{S}(0)\rangle$ are now 2^n-dimensional row and column vectors, and the average is over the choices of negating or not negating the variables in the clauses.

We have performed this task for clusters with up to four sites; their contributions are summarized in Table I. Up to order α^2, the expansion leads to

$$\varphi_0(u)=\frac{e^{-u}}{4}+\alpha\left(\frac{1}{4}e^{-u/2}-\frac{3}{8}e^{-u}+\frac{1}{8}e^{-2u}\right)+\alpha^2\left(-\frac{3}{8}e^{-u/2}\right.$$
$$\left.+\frac{19}{64}e^{-u}+\frac{1}{8}e^{-3u/2}-\frac{9}{32}e^{-2u}+\frac{3}{64}e^{-3u}\right)$$
$$+\alpha^2\left(\frac{\sqrt{3}-1}{32}e^{-[(3+\sqrt{3})/2]u}-\frac{\sqrt{3}+1}{32}e^{-[(3-\sqrt{3})/2]u}\right.$$

TABLE I. Contributions to the cluster expansions for $K=2$. See left panel of Fig. 5.

Type	L_s	K_s	$[\langle E\rangle(u)]_s$	$[t_{res}]_s$
a	2	1	$\frac{1}{4}e^{-u}$	1/4
b	3	3/2	$\frac{1}{16}e^{-2u}+\frac{5}{16}e^{-u}+\frac{1}{8}e^{-u/2}$	19/32
c	4	2	$\frac{1}{128}[e^{-3u}+10e^{-2u}+4e^{-3u/2}+53e^{-u}+20e^{-u/2}+2(2-\sqrt{2})[e^{-[(2+\sqrt{2})/2]u}+2(2+\sqrt{2})e^{-[(2-\sqrt{2})/2]u}]$	125/128
d	4	2/3	$\frac{3}{256}[e^{-3u}+10e^{-2u}+25e^{-u}+32e^{-u/2}-2(1+\sqrt{3})e^{-[(3-\sqrt{3})/2]u}-2(1-\sqrt{3})e^{-[(3+\sqrt{3})/2]u}]$	259/256

$$+\frac{2-\sqrt{2}}{16}e^{-[(2+\sqrt{2})/2]u}+\frac{2+\sqrt{2}}{16}e^{-[(2-\sqrt{2})/2]u}\Bigg)+O(\alpha^3). \quad (21)$$

The typical value $\varphi_0(t)$ of the fraction of unsatisfied clauses after $T=t\,M$ steps of the algorithm is obtained from Eq. (21) through the rescaling of time defined in Eq. (8). Our theory is compared in Fig. 6 to numerical simulations. The agreement is excellent at the beginning of the time evolution, and worsens at the end. A factor of explanation is that during the last steps of the algorithm, the fraction of unsatisfied clauses is low and thus for finite-size samples, the self-averaging hypothesis is violated.

Our approach is applicable to any value of K (note, however, that the size of the matrices to be diagonalized grows faster with the number of clauses in the clusters studied, that is, with the order in α in the expansion) but is restricted to the $\alpha<\alpha_d$ regime. The finite-size tree clusters considered at any (finite) order in the expansion are indeed solved in linear time by RandomWalkSAT. In Sec. IV, another kind of approximation will allow us to study the $\alpha>\alpha_d$ regime.

D. Cluster analysis of solving times

As the number of steps needed to solve a formula is an additive quantity over clusters, it can be calculated along the lines exposed above for the expansion of φ_0. Let us give two examples about how to compute the average time needed to solve a cluster of fixed topology.

The simplest example is a cluster made up of a single 2-SAT clause, $(x_1\vee x_2)$. With probability 3/4, the initial random condition is already a solution. If not, it will take one-time step to solve the cluster, as any of the two possible spin flips will lead to a solution. Thus the average solving time is 1/4. Similar analysis can be done on bigger clusters, even if the counting becomes increasingly tiresome. As a byproduct, some flaws of the RandomWalkSAT heuristic appear.

The second example is the cluster $(x_1\vee x_2)\wedge(\bar{x}_2\vee x_3)$. The eight configurations of the variables are represented in Fig. 7 with up or down spins corresponding to true or false boolean variables. The four solutions are drawn in the shaded box. Two configurations are turned into solutions in one flip (single outcoming arrow). Flipping a spin from each of the remaining two configurations (with two departing arrows) leads to a solution, or to the other configuration. Thus, the average solving time starting from one of these configurations is two time steps. Finally, the average solving time for the cluster is 3/4. To obtain the average over the signs of this solving time, one still has to make the same study for the case where the two clauses are not contradictory on the cen-

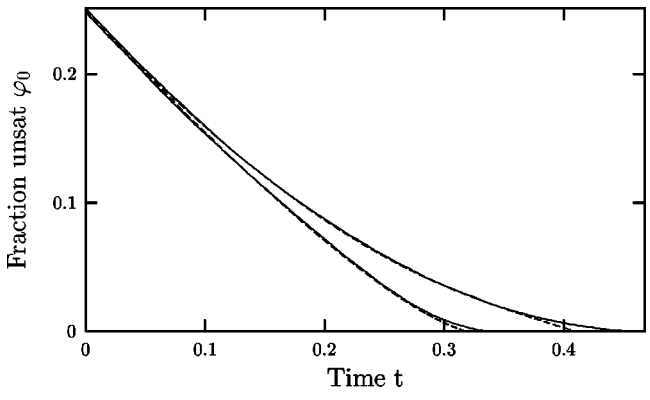

FIG. 6. Fraction of unsat clauses as a function of time for $K=2$, and $\alpha=0.6$ (top) and $\alpha=0.3$ (bottom). Dashed lines: cluster expansion prediction. Solid lines (almost superimposed): numerical results (averaged over ten runs for $N=10\,000$).

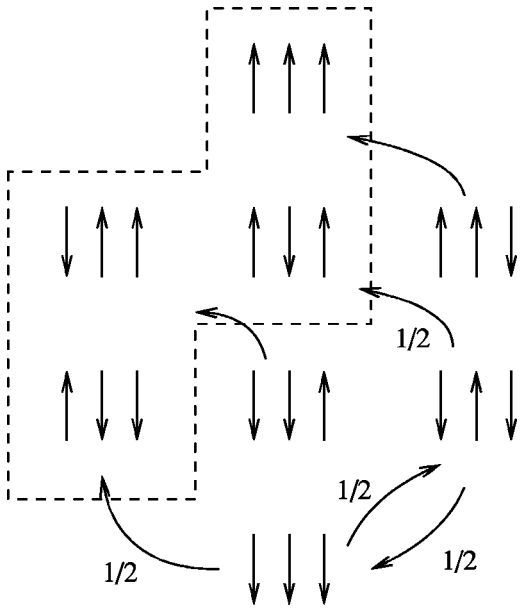

FIG. 7. An example of the behavior of RandomWalkSAT on a finite cluster $(x_1\vee x_2)\wedge(\bar{x}_2\vee x_3)$. Blocks of spins represent configurations of the Boolean variables, those in the shaded box are solutions. Arrows stand for possible transitions of the algorithm, see the text for details.

TABLE II. Contributions to the cluster expansion for the solving time, general K. See right panel of Fig. 5.

Type	L_s	$K_s(K!)^{m_s}/L_s$	$[t_{res}]_s$
a'	K	1	$\frac{1}{2^K}$
b'	$2K-1$	$\frac{K^2}{2}$	$\frac{1}{2^{2K}}\left[2^{K+1}+\frac{K+1}{K(K-1)}\right]$
c'	$3K-2$	$\frac{K^3(K-1)}{2}$	$\frac{1}{2^{3K}}\left[3\cdot 2^{2K}+2^{K+1}\frac{K+1}{K(K-1)}+\frac{4K^4+9K^3+9K^2+6K-4}{3K^2(K-1)(2K-1)(K^2-2)}\right]$
d'	$3K-2$	$\frac{K^3}{6}$	$\frac{1}{2^{3K}}\left[3\cdot 2^{2K}+2^K\frac{3(K+1)}{K(K-1)}-\frac{2K+1}{K^2(K-1)}\right]$

tral spin, hence the value 19/32 on type b of Table I.

Obviously, RandomWalkSAT can make ''bad'' choices on this very simple example, and ''oscillate'' a few time steps between the two configurations before finding a solution. One can imagine many local search heuristics that would do better than RandomWalkSAT on the example shown before. For instance, one could modify the algorithm so that once an unsatisfied clause has been chosen randomly, it prefers to flip a variable with a low number of neighbors, or one with the lowest number of contradicting clauses on it. The average solving time of any of these heuristics, as long as the information used to choose the variable to be flipped remains local, can be studied by such a cluster expansion. These simple enumerations could then provide an useful test ground for new heuristics.

For general K, the enumeration of the possible histories of clusters with three or less clauses leads to the quantities given in Table II,

$$t_{res}(\alpha,K)=\frac{1}{2^K}+\frac{K(K+1)}{K-1}\frac{1}{2^{2K+1}}\alpha + \frac{4K^6+K^5+6K^3-10K^2+2K}{3(K-1)(2K-1)(K^2-2)}\frac{1}{2^{3K+1}}\alpha^2 + O(\alpha^3). \quad (22)$$

This prediction is in a good agreement with numerical simulations, see Fig. 8 for results in the $K=3$ case.

The validity of expression (22) can be easily checked for $K=2$ from the findings of Sec. III C. Indeed, in terms of the time $t=T/M$, the fraction of unsat clauses $\varphi_0(t)$ vanishes after a finite time t_{res} given by

$$t_{res}=\lim_{u\to\infty} t(u)=\int_0^\infty du'\,\varphi_0(u'). \quad (23)$$

Integration of Eq. (21) coincides with prediction (22) for $t_{res}(\alpha,2)$.

IV. APPROXIMATE ANALYSIS OF RANDOMWALKSAT

In this section, an analysis of RandomWalkSAT, based on the Markovian approximation for the evolution equations (2) is proposed. It allows a quantitative description of the main features of RandomWalkSAT, namely, the asymptotic energy and the (exponentially small) probability of resolution in linear time for generic values of K and α ($>\alpha_d$). The approximation scheme is also applied to the analysis of RandomWalkSAT on the XOR-SAT problem [36].

A. Projected evolution

Consider an instance of the K-SAT problem. RandomWalkSAT defines a Markov process on the space of configurations of the Boolean variables, see Eq. (2), of cardinality 2^N. Call $\mathbf{S}(T)$ the configuration of the Boolean variables at a given instant T (number of flips) of the evolution of the algorithm. An observable $\mathcal{R}$ is a function of the configuration, e.g., the number of clauses violated by $\mathbf{S}$. The principle of the approach developed now is to track the evolution of $\mathcal{R}(T)\equiv\mathcal{R}(\mathbf{S}(T))$, that is, of one number (or a low-dimensional vector) instead of the whole configuration of spins. To do so, we make use of the projection operator formalism described below.

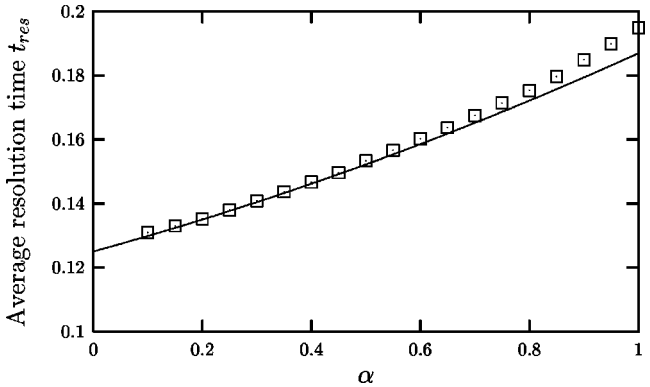

FIG. 8. Average solving time $t_{res}(\alpha,3)$ for 3-SAT. Symbols: numerical simulations, averaged over 1 000 runs for $N=10\,000$. Solid line: prediction from the cluster expansion (22).

Let us partition the configuration space into equivalence classes of microscopic configurations **S** associated with the same value of the macroscopic observable $\mathcal{R}(\mathbf{S})$. We call $\Omega(R)=\{\mathbf{S}|\mathcal{R}(\mathbf{S})=R\}$ these classes, and $|\Omega(R)|$ their cardinalities (number of configurations in these classes). Let us define the projection operator $\hat{\mathcal{P}}$ through its entries,

$$\langle \mathbf{S}_1|\hat{\mathcal{P}}|\mathbf{S}_2\rangle = \frac{1}{|\Omega(\mathcal{R}(\mathbf{S}_1))|}\delta(\mathcal{R}(\mathbf{S}_1)-\mathcal{R}(\mathbf{S}_2)), \qquad (24)$$

where δ denotes the (vectorial) Kronecker function. One can easily check that it is indeed a projector, $\hat{\mathcal{P}}^2=\hat{\mathcal{P}}$, that connects only configurations within the same class.

Now, consider the state vector $|\mathbf{S}(T)\rangle$, Eq. (1) and its projection $|\mathbf{P}(T)\rangle \equiv \hat{\mathcal{P}}|\mathbf{S}(T)\rangle$. Its components have the same value in each class, which is the average of $|\mathbf{S}(T)\rangle$ over the microscopic configurations in the class. Call $|\mathbf{Q}(T)\rangle = (1-\hat{\mathcal{P}})|\mathbf{S}(T)\rangle = |\mathbf{S}(T)\rangle - |\mathbf{P}(T)\rangle$. From the master equation (2), we obtain

$$|\mathbf{P}(T+1)\rangle = \hat{\mathcal{P}}\cdot\hat{W}_d|\mathbf{P}(T)\rangle + \hat{\mathcal{P}}\cdot\hat{W}_d|\mathbf{Q}(T)\rangle,$$

$$|\mathbf{Q}(T+1)\rangle = (1-\hat{\mathcal{P}})\cdot\hat{W}_d|\mathbf{P}(T)\rangle + (1-\hat{\mathcal{P}})\cdot\hat{W}_d|\mathbf{Q}(T)\rangle. \qquad (25)$$

The second equation can be formally ''integrated'' by iteration,

$$|\mathbf{Q}(T)\rangle = \sum_{T'=1}^{T}[(1-\hat{\mathcal{P}})\cdot\hat{W}_d]^{T'}|\mathbf{P}(T-T')\rangle, \qquad (26)$$

where the initial state vector $|\mathbf{S}(0)\rangle$ has been assumed to be uniform on each class, so that $|\mathbf{Q}(0)\rangle = 0$. Finally,

$$|\mathbf{P}(T+1)\rangle = \sum_{T'=0}^{T}\hat{\mathcal{P}}\cdot\hat{W}_d\cdot[(1-\hat{\mathcal{P}})\cdot\hat{W}_d]^{T'}|\mathbf{P}(T-T')\rangle, \qquad (27)$$

Equation (27) expresses that, once coarse gained by the action of the projection operation, the dynamics is not Markovian any longer. The principle of our approximation is precisely to omit all memory effects, by neglecting non-Markovian terms, i.e., $T' \geqslant 1$ contributions in Eq. (27), and averaging over disorder at each time step,

$$|\mathbf{P}(T+1)\rangle \simeq [\hat{\mathcal{P}}\cdot\hat{W}_d]|\mathbf{P}(T)\rangle. \qquad (28)$$

Obviously, the quality of the approximation depends on the observable $\mathcal{R}$. We shall see two examples in what follows.

B. Transition matrix for the number of unsatisfied clauses

A natural choice we study in this and the next two sections for the observable $\mathcal{R}$ is $\mathcal{M}_0$, which measures the numbers M_0 of clauses unsatisfied by each configuration of the variables. Defining the bra

$$\langle M_0| = \sum_{\mathbf{S}\in\Omega(M_0)}\langle \mathbf{S}|, \qquad (29)$$

and the probability $\mathrm{Prob}[M_0,T]=\langle M_0|\mathbf{P}(T)\rangle$ that the configuration of the variables is in class M_0 at time T, we obtain within the Markovian approximation (28),

$$\mathrm{Prob}[M_0',T+1] = \sum_{M_0} A_{M_0'M_0}\mathrm{Prob}[M_0,T] \qquad (30)$$

with $A_{M_0'M_0} \equiv N_{M_0'M_0}/D_{M_0}$ and

$$N_{M_0'M_0} = \sum_{j=1}^{N}\sum_{\mathbf{S}}\delta(M_0-\mathcal{M}_0(\mathbf{S}))\delta(M_0'-\mathcal{M}_0(\mathbf{S}^j))p_j(\mathbf{S}),$$

$$D_{M_0} = \sum_{\mathbf{S}}\delta(M_0-\mathcal{M}_0(\mathbf{S})), \qquad (31)$$

where $p_j(\mathbf{S})$ is the probability of flipping spin k when the system is in configuration **S**, i.e., the number of unsat clauses in which spin j appears divided by the number of unsat clauses. $\mathbf{S}^j$ denotes the configuration obtained from **S** by flipping spin j. The meaning of our Markovian approximation is clear: the transition rate from one value of the observable M_0 to another is the average of the microscopic transition rates from one microscopic configuration belonging to the first subset $\Omega(M_0)$ to another belonging to the second one, with a flat average on the starting subset. At time T, the only available information in the projected process is that the system is somewhere in the subset, and none of the corresponding microscopic configurations can be privileged.

To perform the average over the disorder, i.e., on the random distribution of formulas and compute $[A_{M_0'M_0}] = [N_{M_0'M_0}/D_{M_0}]$, we shall do the further approximation that the numerator and the denominator can be averaged separately. This 'annealed' hypothesis can be justified in some cases, see Sec. IV G. After some combinatorics, we find

$$[A_{M_0'M_0}] = \sum_{Z_u,Z_s}\frac{Z_uN}{KM_0}\binom{M_0}{Z_u}\binom{M-M_0}{Z_s}\left(1-\frac{K}{N}\right)^{M_0-Z_u} \times\left(\frac{K}{N}\right)^{Z_u}\left(1-\frac{K}{(2^K-1)N}\right)^{M-M_0-Z_s} \times\left(\frac{K}{(2^K-1)N}\right)^{Z_s}\delta(M_0'-M_0+Z_u-Z_s). \qquad (32)$$

Z_u is the number of unsatisfied clause which contains the variable to be flipped. All these clauses will become satisfied after the flip. The factor $Z_u/(KM_0)$ represents the probability of flip of the variable, the factor N coming from the sum over its index j. Z_s is the number of clauses satisfied prior to the flip and violated after. The meaning of the Binomial laws is transparent. Assume that the configuration violates M_0 clauses. In the absence of further information, the variable that is going to flip has probability K/N to be present in a given clause [there are $\binom{N}{K}$ possible K uplets over N variables, and $\binom{N-1}{K-1}$ that include a given variable]. Furthermore, a satisfied clause that contains the flipped variable has a probability $1/(2^K-1)$ to become unsatisfied later. Z_u (re-

spectively Z_s) is thus drawn from a binomial law with parameter K/N (respectively $K/[N(2^K-1)]$), over M_0 (resp. $M-M_0$) tries. This reasoning unveils the physical significance of our Markovian approximation; we neglect all correlations between flipped variables and clauses that inevitably arise as the algorithm runs beyond the description in terms of the macroscopic variable M_0.

C. Average evolution and the metastable plateau

The evolution equation for the average fraction $\varphi_0(T) = \Sigma_{M_0} M_0 \mathrm{Prob}[M_0,T]/M$ of unsatisfied clauses at time $T = tM$ is easily computed in the large-size limit from Eqs. (30) and (32). In particular, the average fraction of unsat clauses equals

$$\varphi_0(\alpha,K,t) = \frac{1}{2^K} + \frac{2^K-1}{\alpha K 2^K}\left(e^{-\alpha K(1-2^{-K})^{-1}t} - 1\right), \quad (33)$$

with $\varphi_0(0) = 1/2^K$. Two regimes appear. If the ratio α is smaller than the critical value

$$\alpha_d(K) = \frac{2^K-1}{K}, \quad (34)$$

the average fraction of unsat clauses φ_0 vanishes after a finite time t_{res}. Typically, the algorithm will find a solution after $t_{res} \times M$ steps (linear in N), and then stops. Predictions for α_d are in a good but not in a perfect agreement with estimates from numerical simulations, e.g., $\alpha_d = 7/3$ versus $\alpha_d \simeq 2.7-2.8$ for 3-SAT. The average solving time $t_{res}(\alpha,K)$ predicted within this approximation is given by the time at $\varphi_0(t)$ in Eq. (33) vanishes. It logarithmically diverges as α reaches the dynamical threshold at fixed K,

$$t_{res}(\alpha,K) \sim -\frac{1}{2^K}\ln(\alpha_d(K)-\alpha), \quad \alpha \to \alpha_d(K)^-. \quad (35)$$

On the contrary, when $\alpha > \alpha_d(K)$, φ_0 converges to a finite and positive value

$$\overline{\varphi_0}(\alpha,K) = \frac{1}{2^K}\left(1 - \frac{\alpha_d(K)}{\alpha}\right), \quad (36)$$

when $t \to \infty$ (Fig. 3). RandomWalkSAT is not able to find a solution and gets trapped at a positive level of unsatisfied clauses. This situation arises in the limit $T \propto N, N \to \infty$, and corresponds to the metastable plateau identified in Sec. II C.

D. Large deviations and escape from the metastable plateau

As explained above, when $\alpha > \alpha_d$, the system gets almost surely trapped in a metastable portion of the configuration space with a nonzero number of unsatisfied clauses. Numerical experiments indicate the existence of an exponentially small-in-N probability $\sim \exp(N\overline{\zeta}(\alpha))$ with $\overline{\zeta} < 0$ that this scenario is not correct, and that a solution is indeed found in linear time. We now make use of our Markovian hypothesis to derive an approximate expression for $\overline{\zeta}$.

Contrary to the preceding section, we now consider the large deviation of the process with respect to its typical behavior. This can be accessed through the study of the large deviation function $\pi(\varphi_0,t)$ of the fraction φ_0 [39],

$$\pi(\varphi_0,t) = \lim_{N\to\infty} \frac{1}{N} \ln \mathrm{Prob}[M_0 = M\varphi_0, T = tM]. \quad (37)$$

Introduction of the reduced time is a consequence of the following remark. $O(1)$ changes in the fraction φ_0, that is, $O(N)$ changes in the number M_0 of unsatisfied clauses are most likely to occur after a number of flips of the order of N. To compute the large deviation function π, we introduce the generating function of M_0,

$$G[y,T] = \sum_{M_0} \mathrm{Prob}[M_0,T]\exp(yM_0), \quad (38)$$

where y is a real-valued number. In the thermodynamic limit, G is expected to scale exponentially with N with a rate

$$g(y,t) \equiv \lim_{N\to\infty} \frac{1}{N} \ln G[y, T = tM] = \max_{\varphi_0}[\pi(\varphi_0,t) + \alpha y \varphi_0], \quad (39)$$

equal to the Legendre transform of π from insertion of definition (37) into Eq. (38). Using evolution equations (30) and (32), we obtain the following equation for g,

$$\frac{1}{\alpha}\frac{\partial g(y,t)}{\partial t} = -y + \frac{\alpha K}{2^K-1}(e^y-1) + K\left(e^{-y} - 1 - \frac{1}{2^K-1}(e^y-1)\right)\frac{\partial g(y,t)}{\partial y}. \quad (40)$$

along with the initial condition

$$g(y,0) = \alpha \ln\left(1 - \frac{1}{2^K} + \frac{e^y}{2^K}\right). \quad (41)$$

The average evolution studied in the preceding section can be found again from the location of the maximum of π or, equivalently, from the derivative of g in $y=0$: $\varphi_0(t) = (1/\alpha)\partial g/\partial y(0,t)$. The logarithm of the probability of that a solution is reached after a $O(N)$ time is given by

$$\overline{\zeta}(\alpha,K) = \pi(\varphi_0 = 0, t \to \infty) = \int_0^{\tilde{y}(\alpha)} dy z(y,\alpha), \quad (42)$$

where $z(y,\alpha) = [y - \alpha K(e^y-1)/(2^K-1)]/[K(e^{-y} - 1 - (e^y - 1)/(2^K-1))]$ and $\tilde{y}(\alpha)$ is the negative root of z.

Predictions for $\overline{\zeta}$ in the $K=3$ case are plotted in Fig. 9. They are compared to experimental measures of ζ, that is, the logarithm (divided by N) of the average solving times. It

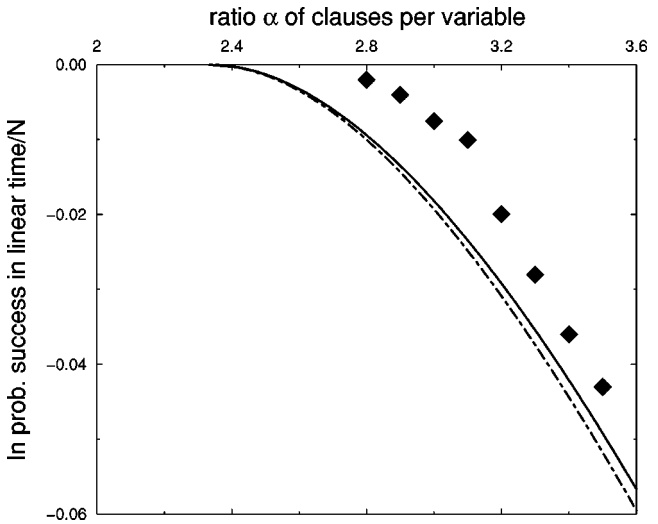

FIG. 9. Large deviations for the 3-SAT problem. The logarithm (in base e) $\overline{\zeta}$ of the probability of successful search (over the linear in N time scale) is plotted as a function of the ratio α of clauses per variables. Predictions for $\overline{\zeta}(\alpha,3)$ have been obtained within the approximations of Sec. IV D [Eq. (42), dot-dashed curve] and Sec. IV E [fourth order solution of Eq. (47), solid curve]. Diamonds corresponds to (minus) the logarithm ζ of the average solving times (averaged over 2 000 to 10 000 samples depending on the values of α,N, divided by N and extrapolated to $N\to\infty$) obtained from numerical simulations. Error bars are of the order of the size of the diamond symbol. Schöning's bound is $\overline{\zeta}\geqslant\ln(3/4)\simeq-0.288$.

is expected on intuitive grounds exposed in Sec. II C that ζ coincides with $-\overline{\zeta}$ (Fig. 2). Despite the roughness of our Markovian approximation, theoretical predictions are in qualitative agreement with numerical experiments.

E. Taking into account clause types

The calculation of Sec. IV B can be extended to other observables $\mathcal{R}$. In the following, we consider the case of a vectorial observable $\vec{\mathcal{M}}$, with $K+1$ components. Our starting point is the classification of clauses into types. A clause is said to be of type i, with $i=0,\ldots,K$, if the variables of the configuration $\mathbf{S}$ satisfy i among K of its literals. If $i=0$ the clause is unsatisfied while, as soon as $i\geqslant 1$, the clause is satisfied. Let us call $M_i(\mathbf{S})$ the number of clauses of type i, and $\vec{M}(\mathbf{S})=[M_0(\mathbf{S}),\ldots,M_K(\mathbf{S})]$ the vector made with these population sizes. Clearly, $\Sigma_i M_i(\mathbf{S})=M$ for any configuration. If $M_0(\mathbf{S})=0$ then $\mathbf{S}$ is a solution of the formula. Vector $\vec{M}$ is a natural characterization of the configuration of variables (and of the instance), and contains essential information about the operation of the algorithm. Indeed, the algorithm stops if the number M_0 of unsatisfied clauses vanishes. In addition, at each step of the algorithm, a single variable is flipped; clauses of type i become of type $i\pm 1$ if they include this variable, or remain of type i otherwise.

Within our Markovian annealed approximation, the probability $\mathrm{Prob}[\vec{M},T]$ that the configuration of the variables is in class $\Omega(\vec{M})$ at time T obeys the evolution equation,

$$\mathrm{Prob}[\vec{M}',T+1]=\sum_{\vec{M}}[A_{\vec{M}'\vec{M}}]\mathrm{Prob}[\vec{M},T], \tag{43}$$

with

$$[A_{\vec{M}'\vec{M}}]=\sum_{\vec{Z}}\frac{NZ_0}{KM_0}\delta(\vec{M}'-\vec{M}-\Delta\cdot\vec{Z})P(\vec{Z}|\vec{M}), \tag{44}$$

where $\vec{Z}=(Z_0,Z_1,Z_1^s,\ldots,Z_i,Z_i^s,\ldots,Z_{K-1},Z_{K-1}^s,Z_K)$ is a $(2K-1)$ dimensional vector. Component Z_i is the number of clauses of type i where the variable which is going to flip appears. In Z_i^s of these Z_i clauses, this variable was one of the i satisfying literals. It is not necessary to introduce components Z_0^s and Z_K^s for they have obvious values (respectively, equal to 0 and Z_K). Δ denotes a $(K+1)\times(2K-1)$ matrix such that $\Delta\cdot\vec{Z}$ gives the change in the observable $\vec{\mathcal{M}}$ when the variable is flipped. The ith line of $\Delta\cdot\vec{Z}$ reads $-Z_i+Z_{i+1}^s+(Z_{i-1}-Z_{i-1}^s)$. Clauses that contained the flipped variable and were of type i prior to the flip are no longer of this type after the flip (hence the term $-Z_i$), those which were of type $i+1$ and satisfied by the variable become of type i $(+Z_{i+1}^s)$, as those which were of type $i-1$ and unsatisfied by the flipping variable $(+Z_{i-1}-Z_{i-1}^s)$. The probability of $\vec{Z}$ conditioned on $\vec{M}$ is, as in the simpler case of Sec. IV B, a product of Binomial laws

$$P(\vec{Z}|\vec{M})=\prod_{i=0}^{K}\binom{M_i}{Z_i}\left(\frac{K}{N}\right)^{Z_i}\left(1-\frac{K}{N}\right)^{M_i-Z_i}\prod_{i=1}^{K-1}\binom{Z_i}{Z_i^s}\times\left(\frac{i}{K}\right)^{Z_i^s}\left(1-\frac{i}{K}\right)^{Z_i-Z_i^s}. \tag{45}$$

Repeating the procedures of Secs. IV C and IV D, we find the following.

(1) The average fraction of unsat clauses is calculated in Appendix A and reads as

$$\varphi_0(\alpha,K,t)=\frac{1}{2^K}+\frac{1}{\alpha K}\left[\frac{1}{(1+\tanh(\alpha t))^K}-1\right]. \tag{46}$$

The critical ratio separating polynomial from exponential search is found at the same value as in Sec. IV B, $\alpha_d(K)=(2^K-1)/K$. The average fraction of unsatisfied clauses on the plateau (φ_0 in the $t\to\infty$ limit) when $\alpha>\alpha_d(K)$ has also the same expression, cf. Eq. (36). Note, however, that the finite time evolution differs in the two calculations. When $\alpha<\alpha_d$, the solving time t_{res} is given by the vanishing of φ_0 in Eq. (46), and differs from its value found within the simpler approximation of Sec. IV B.

(2) The probability of easy (linear time) search is accessible from the large deviation function $\pi(\vec{\varphi},t)$ of the fractions φ_i of clauses of type $0\leqslant i\leqslant K$. Its Legendre transform $g(\vec{y},t)$ obeys the partial differential equation (PDE)

$$\frac{1}{\alpha}\frac{\partial g(\vec{y},t)}{\partial t} = -y_0 + y_1 + \sum_{i=0}^{K} ((K-i)e^{y_{i+1}-y_i} \times + i\, e^{y_{i-1}-y_i} - K)\frac{\partial g(\vec{y},t)}{\partial y_i} \tag{47}$$

along with the initial condition

$$g(\vec{y},0) = \alpha \ln\left[\sum_{i=0}^{K} \frac{1}{2^K}\binom{K}{i} e^{y_i}\right]. \tag{48}$$

The logarithm (divided by N) of the probability of the search to be successful on the $0(N)$ time scale, is given by

$$\overline{\zeta}(\alpha,K) = \max_{y_0} g(y_0, y_1 = y_2 = \cdots = y_K = 0, t). \tag{49}$$

We have not been able to calculate exactly $\overline{\zeta}$ for generic values of α and K, but have resorted to a polynomial expansion of g in powers of its arguments y_i. The expansion has been done up to order four with the help of a symbolic computation software package for $K=3$, and up to order two analytically for any K. Calculations are detailed in Appendix B. Predictions for $\overline{\zeta}(\alpha)$ in the $K=3$ case are plotted in Fig. 9.

F. The large *K* limit

A comparison between results of Secs. IV B and IV E shows that the output of the calculation quantitatively depends on the observable under study. However, we may expect some simplification to take place for large K. In this limit, if a clause gets unsatisfied twice, or more (but $\ll K$ times), it is very unlikely that each variable will be flipped more than once, and memory effects are lost. Therefore, the Markovian annealed approximation is expected to become correct. However, to avoid a trivial limit, the ratio α of clauses per variable must be rescaled accordingly. Inspection of the above result (34) indicates that the correct scaling is $\alpha, K \to \infty$ at fixed ratio $\alpha^* = \alpha K/2^K$. The dynamical threshold separating linear from exponential searches is located in

$$\alpha_d^* = 1. \tag{50}$$

As the critical threshold of K-SAT is known to scale as $\alpha_c(K) \sim 2^K \ln 2$ for large K [34,38], instances are always satisfiable on the reduced α^* scale.

For $\alpha^* < 1$, the initial fraction of unsatisfied clauses is $\simeq 1/2^K$, and decreases by $O(1)$ per unit of reduced time t, giving $t_{res} \sim 1/2^K$. For the same reason, the height φ_0, which is reached after a $O(1)$ relaxation time when $\alpha^* > 1$, is of the order of $1/2^K$. It is therefore natural to define the rescaled fraction of unsatisfied clauses through

$$\varphi_0^*(\alpha^*, t^*) = \lim_{K\to\infty} 2^K \varphi_0(\alpha^* 2^K/K, K, t^*/2^K), \tag{51}$$

from which we obtain the rescaled solving time $t_{res}^*(\alpha^*)$ for $\alpha^* < 1$ (vanishing of φ_0^*) and plateau height $\overline{\varphi_0^*}(\alpha^*)$ for $\alpha^* > 1$ (limit value of φ_0^* at large rescaled times). The two schemes of approximation given in Secs. IV B and IV E both yield

$$\varphi_0^*(\alpha^*, t^*) = 1 + \frac{1}{\alpha^*}(e^{-\alpha^* t^*} - 1), \tag{52}$$

$$t_{res}^*(\alpha^*) = -\frac{1}{\alpha^*}\ln(1-\alpha^*) = 1 + \frac{1}{2}\alpha^* + \frac{1}{3}(\alpha^*)^2 + O((\alpha^*)^3), \tag{53}$$

$$\overline{\varphi_0^*}(\alpha^*) = 1 - \frac{1}{\alpha^*} = \alpha^* - 1 + O((\alpha^* - 1)^2). \tag{54}$$

Note that the small α^* expansion (53) for the solving time coincides with the exact expansion obtained from Eq. (22) with the above rescaling of α and K. We conjecture that the equality holds for higher orders ($\geqslant 3$) in α^*, and that the above expressions for $\varphi_0^*(\alpha^*, t^*)$ and thus for $t_{res}^*(\alpha^*), \overline{\varphi_0^*}(\alpha^*)$ are correct.

The logarithm of the probability of fast search for $\alpha > \alpha_d(K)$ needs to be rescaled too,

$$\overline{\zeta}^*(\alpha^*) = \lim_{K\to\infty} K\overline{\zeta}(\alpha^* 2^K/K, K), \tag{55}$$

to acquire a well-defined limit when $K \to \infty$. The scalar approximation of Sec. IV D gives the asymptotic result

$$\overline{\zeta}^*(\alpha^*) = \int_0^{\tilde{y}(\alpha^*)} dy \frac{y - \alpha^*(e^y - 1)}{e^{-y} - 1} = -(\alpha^* - 1)^2 + O((\alpha^* - 1)^3), \tag{56}$$

where $\tilde{y}(\alpha^*)$ is the negative root of the numerator in the above integral. The quadratic resolution of the PDE arising from the study of Sec. IV E (cf. Appendix B) also leads to this result around $\alpha^* = 1$. Unfortunately, the exact results obtained in Sec. III are of no help to confirm identity (56).

G. The XOR-SAT case

XOR-SAT is a version of a satisfiability problem, much simpler than SAT from a computational complexity point of view [25,27,36,37]. One still draws K-uplets of variables, but each clause bears only one sign (instead of one for each variable in the K-SAT version), and the clause is said to be satisfied if the exclusive OR (XOR) of its boolean variables is equal to the sign of the clause. For a given clause, there are 2^{K-1} satisfiable assignments of the variables and also 2^{K-1} unsatisfiable assignments, in deep contrast with SAT where these numbers are, respectively, equal to $2^K - 1$ and 1. XOR-SAT may be recast as a linear algebra problem, where a set of M equations involving N Boolean variables must be satisfied modulo 2, and is therefore solvable in polynomial time by various methods, e.g., Gaussian elimination. Never-

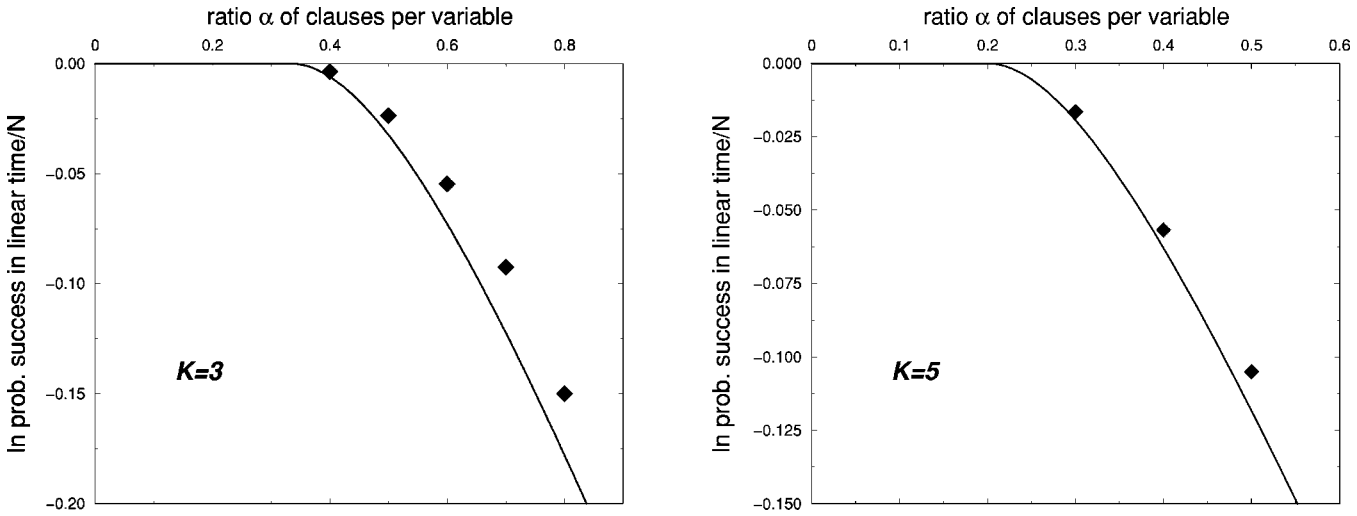

FIG. 10. Large deviations for the K-XOR-SAT problem for $K=3$ (left) and $K=5$ (right). The logarithm (in base e) $\overline{\zeta}$ of the probability of successful search (over the linear time scale) is plotted as a function of the ratio α of clauses per variables. Diamonds corresponds to (minus) the logarithm ζ of the average solving times (averaged over 10,000 samples, divided by N and extrapolated to $N\to\infty$) obtained from numerical simulations. Error bars are smaller than the size of the diamond symbol.

theless, it is legitimate to ask what the performance of local search methods as RandomWalkSAT are for this kind of computational problem.

A fundamental feature of XOR-SAT is that, whenever a spin is flipped, all clauses where this spin appears change status: the satisfied ones become unsatisfied and *vice versa*. There is thus no need to distinguish between clauses satisfied by a different number of literals, and the macroscopic observable we track is the number M_0 of unsatisfied clauses for configuration **S** as in Sec. IV B. It is an easy check that the transition matrix $[A]$ for XOR-SAT is given by the expression (32) where 2^K-1 is replaced with 1. Main results are the following.

(1) The average fraction of unsatisfied clauses $\varphi_0(t)$ reads as

$$\varphi_0(\alpha,K,t)=\frac{1}{2}+\frac{1}{2\alpha K}(e^{-2\alpha K t}-1), \tag{57}$$

and becomes asymptotically strictly positive if the ratio α of clauses per variables exceeds $\alpha_d(K)=1/K$, smaller than the clustering and critical ratios, $\alpha_s\simeq 0.818$ and $\alpha_c\simeq 0.918$ for $K=3$, respectively [25,27]. The overall picture of the algorithm behavior is identical to the SAT case.

(2) When $\alpha>\alpha_d(K)$, the average fraction of unsatisfied clauses on the plateau is given by

$$\overline{\varphi_0}(\alpha,K)=\frac{1}{2}\left(1-\frac{1}{\alpha K}\right). \tag{58}$$

(3) The partial differential equation for the generating function is

$$\frac{1}{\alpha}\frac{\partial}{\partial t}g(y,t)=-y+\alpha K(e^{y}-1)+K(e^{-y}-e^{y})\frac{\partial}{\partial y}g(y,t) \tag{59}$$

with $g(y,0)=\alpha\ln[(1+e^{y})/2]$. Resolution in the large time limit is straightforward, with the results shown in Fig. 10. The agreement with numerics is good, especially as K grows.

The relatively simple structure of XOR-SAT makes possible the test of some of the approximations we made. We show in Appendix C that the annealed hypothesis in the calculation of the evolution matrix A is justified in the thermodynamic limit. The validity of this approximation in the case of 3-SAT (Sec. IV B) is not established for finite K.

As for K-SAT, quantities of interest have a well-defined large α,K limit provided the ratio $\alpha^*=\alpha/\alpha_d(K)=\alpha K$ is kept fixed;

$$\varphi_0^*(\alpha^*,t^*)=\lim_{K\to\infty}\varphi_0(\alpha^*/K,K,t^*)$$
$$=\frac{1}{2}+\frac{1}{2\alpha^*}(e^{-2\alpha^* t^*}-1), \tag{60}$$

$$\overline{\zeta}^*(\alpha^*)=\lim_{K\to\infty}K\overline{\zeta}(\alpha^*/K,K)=\int_0^{\tilde{y}(\alpha^*)}dy z(y,\alpha^*)$$
$$=-\frac{1}{2}(\alpha^*-1)^2+O((\alpha^*-1)^3), \tag{61}$$

where $z(y,\alpha^*)=[y-\alpha^*(e^{y}-1)]/(e^{-y}-e^{y})$ and $\tilde{y}(\alpha^*)$ is the negative root of z.

V. CONCLUSION AND PERSPECTIVES

In this paper, we have studied the dynamics of a simple search procedure for the satisfaction of Boolean constraints, the RandomWalkSAT algorithm. We have shown using complementary techniques (expansions and approximations) that, for randomly drawn input instances, RandomWalkSAT may have two qualitatively distinct behaviors. Instances with

small ratios α of clauses per variable are almost surely solved in a time growing linearly with their size. On the contrary, for ratios above a threshold α_d, the dynamics gets trapped for an exponentially large time in a region of the configuration space with a positive fraction of unsatisfied clauses. A solution is finally reached through a large fluctuation from this metastable state.

The freezing taking place at α_d does not seem to be related to the onset of clustering between solutions [21]. Indeed, the value of α_d is expected to change with the local search rules. It would be interesting to pursue the study initiated in the present work to understand if and how the existence of this dynamical threshold is related to some property of the (static) energy landscape, as in mean-field models of spin glasses [15]. Another useful improvement would be to go beyond the Markovian approximation of Sec. IV. Unfortunately, keeping a finite (with respect to N) number of retarded terms in Eq. (27) should not be sufficient to achieve this goal. Improvements will require to take into account an extensive number of terms, or to extend the quantum formalism of Sec. III to the study of the metastable phase. Another possible direction of research would be to use projection operators on observables of extensive dimension [16,40]. Our Markovian approximation, expected to be exact in the large α, K limit, should be a starting point for a systematic expansion of the quantities of interest (plateau height, lifetime of the metastable regime, etc.). Finally, extension of our analysis to more sophisticated local search heuristics would be useful.

ACKNOWLEDGMENTS

We thank W. Barthel, A. K. Hartmann, and M. Weigt for oral communication of their results prior to publication [16], and S. Cocco, L. Cugliandolo, and A. Montanari for useful discussion and comments. The quantum formalism presented in Sec. III is inspired from a previous unpublished work of one of us with G. Biroli, whom we are very grateful to. We also thank C. Deroulers for his help in the perturbative resolution of Eq. (47) (Sec. III C). The present work was partly supported by the French Ministry of Research through the ACI Jeunes Chercheurs "Algorithmes d'optimization et systèmes désordonnés quantiques."

APPENDIX A: GENERATING FUNCTION FOR THE AVERAGE EVOLUTION

Defining $\vec{\varphi}=(\varphi_0,\ldots,\varphi_K)=\Sigma_{\vec{M}}\mathrm{Prob}[\vec{M},T]\vec{M}/M$ and the reduced time $t=T/M$, one get from Eq. (30),

$$\frac{d\vec{\varphi}}{dt}=\vec{v}+\alpha K W_r\cdot\vec{\varphi} \tag{A1}$$

with $\vec{v}$ a $(K+1)$-dimensional vector and W_r a $(K+1)\times(K+1)$ matrix defined as

$$\vec{v}=\begin{pmatrix}-1\\1\\0\\\vdots\\0\end{pmatrix},$$

$$W_r=\begin{pmatrix}-1 & \frac{1}{K} & 0 & 0 & 0 & \ldots & 0\\ 1 & -1 & \frac{2}{K} & 0 & 0 & \ldots & 0\\ 0 & \frac{K-1}{K} & -1 & \frac{3}{K} & 0 & \ldots & 0\\ \vdots & \vdots & \vdots & \vdots & \vdots & \ldots & \vdots\\ 0 & 0 & 0 & 0 & \ldots & \ldots & -1\end{pmatrix}. \tag{A2}$$

To solve Eq. (A1), it results convenient to introduce $\Phi(x,t)$, the polynomial in x whose coefficients are the fractions φ_i we want to determine,

$$\Phi(x,t)=\sum_{j=0}^{K}\varphi_j(t)x^j. \tag{A3}$$

The set of $K+1$ linear coupled differential (A1) reduces to a partial differential equation on $\Phi(x,t)$:

$$\frac{\partial\Phi}{\partial t}(x,t)=-1+x+\alpha K(x-1)\Phi(x,t)+\alpha(1-x^2)\frac{\partial\Phi}{\partial x}(x,t). \tag{A4}$$

At initial time, the variables are chosen randomly to be true or false, without any correlation with the formula studied. Thus, the number of satisfied literals in a clause obeys a binomial law with parameter 1/2: $\varphi_j(0)=(1/2^K)\binom{K}{j}$. The initial condition on Φ reads thus

$$\Phi(x,0)=\left(\frac{1+x}{2}\right)^K. \tag{A5}$$

Setting $\Psi(x,t)=\Phi(x,t)+1/(\alpha K)$, the constant terms in Eq. (A4) can be eliminated, with the resulting PDE for Ψ,

$$\frac{\partial\Psi}{\partial t}(x,t)=\alpha K(x-1)\Psi(x,t)+\alpha(1-x^2)\frac{\partial\Psi}{\partial x}(x,t). \tag{A6}$$

This can be transformed into a wave equation on $\chi(x,t)=(1+x)^{-K}\Psi(x,t)$:

$$\frac{\partial\chi}{\partial t}(x,t)=\alpha(1-x^2)\frac{\partial\chi}{\partial x}(x,t). \tag{A7}$$

This equation is solved in terms of an arbitrary function of a single argument,

$$\chi(x,t)=\omega\left(\frac{1}{\alpha}\tanh^{-1}(x)+t\right). \tag{A8}$$

Knowledge of $\Phi(x,0)$ for all x, Eq. (A5) allows us to determine unambiguously ω

$$\omega(u)=\frac{1}{2^K}+\frac{1}{\alpha K}[1+\tanh(\alpha u)]^{-K}. \tag{A9}$$

Going backwards, we obtain the expression of the generating function

$$\Phi(x,t)=\left(\frac{1+x}{2}\right)^K+\frac{1}{\alpha K}\left\{\left[\frac{1+x\tanh(\alpha t)}{1+\tanh(\alpha t)}\right]^K-1\right\} \tag{A10}$$

and of the fractions of clauses of type i through an expansion of the latter in powers of x,

$$\varphi_j(t)=\binom{K}{j}\left[\frac{1}{2^K}+\frac{1}{\alpha K}\frac{[\tanh(\alpha t)]^j}{[1+\tanh(\alpha t)]^K}-\frac{\delta_{j0}}{\alpha K}\right]. \tag{A11}$$

APPENDIX B: PERTURBATIVE RESOLUTION OF THE LARGE DEVIATION PDE

In this appendix, we sketch the resolution of PDE, Eq. (47), in the long-time limit, where the function g becomes independent of time. We expand it in powers of its arguments:

$$g(\vec{y},t\to\infty)=\sum_{i=0}^{K}a_i y_i+\frac{1}{2}\sum_{i,j=0}^{K}a_{ij}y_i y_j+\cdots. \tag{B1}$$

Substituting this expansion into Eq. (47), one obtains by identification of the monomials in y_i an infinite set of linear equations on the coefficients of g, which can be solved order by order. Constraint $\Sigma_i\varphi_i=1$ imposes a condition on g,

$$g(\vec{y}+c\vec{1})=\alpha c+g(\vec{y}), \tag{B2}$$

where $\vec{1}=(1,1,\ldots,1)$ and c is an arbitrary constant.

In the case $K=3$, we have solved the set of equations on the coefficients of g up to order four in the y_i's with the help of a symbolic computation software. To calculate $\overline{\zeta}$, we need to know g as a function of y_0 only, with $y_i=0$, $\forall i\geqslant 1$. We find

$$g(y_0)=\frac{3\alpha-7}{24}y_0+\frac{105\alpha-94}{1920}y_0^2+\frac{26\,460\alpha+10\,753}{193\,560}y_0^3+\frac{29\,645\alpha+66\,244}{1\,8923\,520}y_0^4+O(y_0^5). \tag{B3}$$

When $\alpha>\alpha_d(K=3)=7/3$, this function has a nontrivial extremum, in which g takes the value $\overline{\zeta}(\alpha)$.

We, now explain the resolution at quadratic order for a generic value of K. At linear order, following the calculation exposed in Appendix A,

$$a_i=\alpha\lim_{t\to\infty}\varphi_i(t)=\frac{1}{2^K}\binom{K}{i}\left(\alpha+\frac{1}{K}\right)-\frac{\delta_{i0}}{K}. \tag{B4}$$

In particular, $a_0=[\alpha-\alpha_d(K)]/2^K$. Then considering the monomials of second order in the expansion of the equation, a set of $(K+1)(K+2)/2$ linear equations determine the coefficients a_{ij}. As we shall not try to solve the equation at higher orders for this generic case, we need only a_{00}. Again, we introduce a generating function to turn the discrete algebraic problem into an analytic one. $f(s)\equiv\Sigma_{i,j}a_{ij}s^{i+j}$ obeys the ordinary differential equation,

$$2Kf(s)-(s+1)f'(s)=(s-1)\left[1-(1+\alpha K)\left(\frac{1+s^2}{2}\right)^{K-1}\right] \tag{B5}$$

with the condition $f(1)=0$ stemming from Eq. (B2). This equation can be easily solved, yielding

$$a_{00}=f(0)=\frac{2^K-1}{2^{2K}}\left(\alpha-\frac{(K-1)2^{2K}+2^K+2K-K2^{K+1}}{K(2K-1)(2^K-1)}\right). \tag{B6}$$

At this order of the expansion, the extremum of g in the subspace $y_i=0$, $\forall i\geqslant 1$ is reached in $y_0=-a_0/a_{00}$ and leads to $\overline{\zeta}(\alpha,K)=-a_0^2/(2a_{00})$.

APPENDIX C: VALIDITY OF THE ANNEALED HYPOTHESIS FOR THE XOR-SAT PROBLEM

We justify in this appendix the annealed average of the Markovian transition matrix in the XOR-SAT case. Our analysis is based on the Chebyshev inequality [41]: a positive integer valued random variable with a variance negligible with respect to the square of its average is sharply peaked around its mean value. Call $D_U=\Sigma_{\mathbf{S}}\delta(U-U(\mathbf{S}))$ the number of configurations with U unsatisfied clauses. The first moment of D_U over the distribution of XOR-SAT instances is easy to compute. After averaging, the 2^N configurations of the variables contribute equally to the sum; for each of these, the number U of unsatisfied clause has a binomial distribution of parameter 1/2 among the M clauses. In the thermodynamic limit, using Stirling's formula and denoting $\varphi_0=U/M$,

$$[D_U]\sim e^{Nf_1(\varphi_0,\alpha)},$$

$$f_1(\varphi_0,\alpha)=\ln 2+\alpha[-\varphi_0\ln\varphi_0-(1-\varphi_0)\ln(1-\varphi_0)-\ln 2] \tag{C1}$$

up to polynomial corrections. Suppose that $\varphi_0<1/2$. Call $\varphi_0^{(1)}(\alpha)$ the root of f_1 at fixed α. It is a growing function of α, vanishing for $\alpha\leqslant 1$, and monotically increasing to 1/2 as

α gets large. For $\varphi_0 > \varphi_0^{(1)}(\alpha)$, f_1 is positive and $[D_U]$ exponentially large. When $\varphi_0 < \varphi_0^{(1)}(\alpha)$, $[D_U]$ is exponentially small.

Consider now the second moment $[D_U^2]$ and its leading behavior $[D_U^2] \sim \exp[N f_2(\varphi_0, \alpha)]$. We introduce the generating function

$$\sum_U [D_U^2] e^{-2xU} = \int_0^{2\pi} \frac{d\theta}{2\pi} \sum_{\mathbf{S}_1, \mathbf{S}_2} [e^{-x[U(\mathbf{S}_1)+U(\mathbf{S}_2)]+i\theta[U(\mathbf{S}_1)-U(\mathbf{S}_2)]}]. \quad \text{(C2)}$$

The average on the right-hand side can be readily performed as the M clauses are drawn independently. The trace on the two configurations reduces to a sum on the Hamming distance between them. Evaluation of this sum and the integral over θ by the Laplace method yields

$$\operatorname*{ext}_{\varphi_0} [f_2(\varphi_0, \alpha) - 2\alpha x \varphi_0] = S(x, \alpha), \quad \text{(C3)}$$

where $S(x, \alpha)$ is the maximum over γ of

$$S(\gamma, x, \alpha) = \ln 2 - \gamma \ln \gamma - (1-\gamma)\ln(1-\gamma) - \alpha x + \alpha \ln[1 + p_e(\gamma)(\cosh x - 1)]. \quad \text{(C4)}$$

Here, $p_e(\gamma) = [1 + (1-2\gamma)^K]/2$ is the probability that a randomly drawn clause satisfies (or violates) two configurations at Hamming distance $d = \gamma N$. We are thus left with the problem of determining $S(x, \alpha)$ and of computing its Legendre transform with respects to x to obtain f_2. As the derivative of p_e in $\gamma = 1/2$ vanishes, this point is always an extremum of S. Two cases must be distinguished, depending on the value of φ_0, which fixes x [27,37,42]:

(1) If $\gamma = 1/2$ is the global maximum of S, then $f_2(\varphi_0, \alpha) = 2 f_1(\varphi_0, \alpha)$, in other words, $[D_U^2] \sim [D_U]^2$. In this case, it is possible to compute the polynomial corrections by expanding around the saddle point,

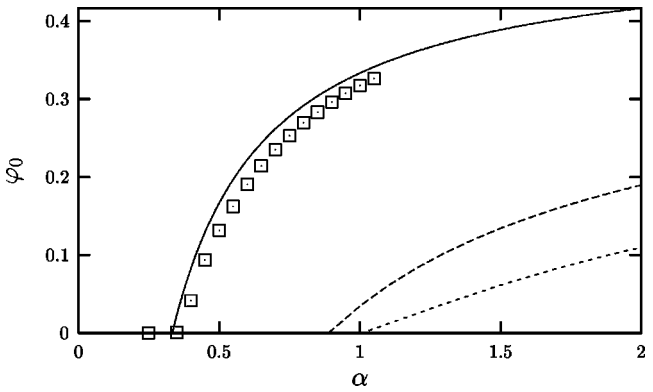

FIG. 11. Study of the moments of D_U for 3-XOR-SAT. Solid line: Markovian annealed prediction for the asymptotic fraction of unsat clauses φ_0. Long-dashed line: $\varphi_0^{(2)}$. Short-dashed line: $\varphi_0^{(1)}$. Symbols: asymptotic fraction of unsat clauses on the plateau, obtained through numerical simulations.

$$\frac{[D_U^2]}{[D_U]^2} \sim 1 + \frac{1}{N^{K-2}} \frac{\alpha^2 K!}{2} [1 - 4\varphi_0(1-\varphi_0)]. \quad \text{(C5)}$$

(2) If the global maximum of S is not in $\gamma = 1/2$, $f_2 > f_1$, thus $[D_U^2] \gg [D_U]^2$.

We have computed numerically for $K = 3$ the function $\varphi_0^{(2)}(\alpha)$ such that for $\varphi_0 > \varphi_0^{(2)}(\alpha)$, the global maxima of S is located in $\gamma = 1/2$. It is a growing function of α, vanishing when $\alpha < 0.889$ [37], and growing monotically to 1/2 when α diverges. The results are shown in Fig. 11 (all three curves reach $\varphi_0 = 1/2$ when $\alpha \to \infty$ without crossing each other). In the course of the algorithm operation, φ_0 decreases from its initial value (1/2) down to its plateau value, and remains confined to the region in the phase diagram where the second moment method applies: $\varphi_0 > \varphi_0^{(2)} > \varphi_0^{(1)}$. This proves that, within the Markovian approximation, the annealed average is correct: as the denominator of the transition matrix is peaked around its mean value, the numerator and denominator can be averaged separately. This analysis cannot be done in the case of K-SAT, for which the second moment fails as soon as $\alpha > 0$ [42].

[1] C.H. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity* (Prentice-Hall, Englewood Cliffs, 1982).

[2] For a recent review, see O. Martin, R. Monasson, and R. Zecchina, Theor. Comput. Sci. **265**, 3–67 (2001).

[3] M. Talagrand, Theor. Comput. Sci. **265**, 69–77 (2001).

[4] D.E. Knuth, *Selected Papers on Analysis of Algorithms*, Center for the Study of Language and Information, Stanford, CA, Report No. 102 (2000).

[5] R. Sedgewick and P. Flajolet, *Introduction to the Analysis of Algorithms* (Addison-Wesley, Boston, 2001).

[6] D. Achliotpas, Theor. Comput. Sci. **265**, 159 (2001).

[7] S. Cocco and R. Monasson, Phys. Rev. Lett. **86**, 1654 (2001); Eur. Phys. J. B **22**, 505 (2001).

[8] S. Cocco and R. Monasson, Phys. Rev. E **66**, 037101 (2002).

[9] M. Weigt and A. Hartmann, Phys. Rev. Lett. **86**, 1658 (2001).

[10] A. Montanari and R. Zecchina, Phys. Rev. Lett. **88**, 178701 (2002).

[11] R. Motwani and P. Raghavan, *Randomized Algorithms* (Cambridge University Press, Cambridge, 1995).

[12] C.H. Papadimitriou, in *Proceedings of the 32nd Annual IEEE Symposium on Foundations of Computer Science*, 1991, p. 163 (unpublished).

[13] B. Selman, H. Levesque, and D. Mitchell, in *Proceedings of the Tenth National Conference on Artificial Intelligence (AAAI-92)*, edited by W. Swartout (AAAI Press, Washington, D.C./ MIT Press, Cambridge, 1992).

[14] T.M. Ligget, *Stochastic Interacting Systems: Contact, Voter and Exclusion Processes* (Springer, Berlin, 1999).

[15] L.F. Cugliandolo, e-print cond-mat/0210312.

[16] W. Barthel, A. Hartmann, and M. Weigt, Phys. Rev. E **67,** 066104 (2003).
[17] E. Friedgut J. Amer. Math. Soc. **12**, 1017 (1999).
[18] A.C. Kaporis, L.M. Kirousis, and E.G. Lalas, in *European Symposium on Algorithms*, Lecture Notes in Computer Science, Vol. 2461 (Springer, Berlin, 2002), p. 574.
[19] O. Dubois, Y. Boufkhad, and J. Mandler, SODA **2002**, 126 (2002).
[20] For a review on random 2-SAT, see, F. De la Vega, Theor. Comput. Sci. **265**, 131 (2001).
[21] G. Biroli, R. Monasson, and M. Weigt, Eur. Phys. J. B **14**, 551 (2000).
[22] G. Parisi (unpublished).
[23] M. Mézard, G. Parisi, and R. Zecchina, Science (Washington, DC, U.S.) **297**, 812 (2002); M. Mézard and R. Zecchina, Phys. Rev. E **66**, 056126 (2002).
[24] For a recent review, see, A. Engel and C. Van den Broeck, *Statistical Mechanics of Learning* (Cambridge University Press, Cambridge, 2000).
[25] F. Ricci-Tersinghi, M. Weigt, and R. Zecchina, Phys. Rev. E **63**, 026702 (1999); S. Franz *et al.*, *ibid.* **87**, 127209 (2001).
[26] R. Mulet, A. Pagnani, M. Weigt, and R. Zecchina, Phys. Rev. Lett. **89**, 268701 (2002).
[27] O. Dubois and J. Mandler, Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science, Vancouver, 2002 (unpublished); S. Cocco, O. Dubois, J. Mandler, and R. Monasson, Phys. Rev. Lett. **90**, 047205 (2003); M. Mézard, F. Ricci-Tersenghi, and R. Zecchina (unpublished).
[28] C.M. Newman and D.L. Stein, Phys. Rev. E **55**, 5194 (1998); in *Mathematical Aspects of Spin Glasses and Neural Networks*, edited by A. Bovier and P. Picco (Birkhäuser, Boston, 1998), pp. 243–287.
[29] P. Svenson and M.G. Nordhal, Phys. Rev. E **59**, 3983 (1999).
[30] U. Schöning, Algorithmica **32**, 615 (2002).
[31] M. Alekhnovich and E. Ben-Sasson (unpublished).
[32] A.J. Parkes, Lect. Notes Comput. Sci. **2470**, 708 (2002).
[33] L.P. Kadanoff and J. Swift, Phys. Rev. **165**, 310 (1968).
[34] R. Monasson and R. Zecchina, Phys. Rev. E **56**, 1357 (1997).
[35] G. Semerjian and L.F. Cugliandolo, Phys. Rev. E **64**, 036115 (2001).
[36] N. Creignou and H. Daudé, Discrete Appl. Math. **96-97**, 41 (1999).
[37] N. Creignou, H. Daudé, and O. Dubois, e-print cs.DM/0106001.
[38] D. Achlioptas and Y. Peres (unpublished).
[39] R.B. Griffiths, C-Y. Weng, and J.S. Langer, Phys. Rev. **149**, 301 (1966).
[40] S.N. Laughton, A.C.C. Coolen, and D. Sherrington, J. Phys. A **29**, 763 (1996).
[41] N. Alon and J. Spencer, *The Probabilistic Method* (Wiley, New York, 2000).
[42] D. Achlioptas and C. Moore, in Proceedings of the 43rd Annual Symposium on Foundations of Computer Science (FOCS'02), Vancouver, 2002, pp. 779–788 (unpublished).