Minimizing energy below the glass thresholds

Demian Battaglia,[1,*] Michal Kolář,[1,†] and Riccardo Zecchina[2,‡]
[1]*SISSA, Via Beirut 9, I-34100 Trieste, Italy*
[2]*ICTP, Strada Costiera 11, I-34100 Trieste, Italy*
(Received 26 February 2004; published 15 September 2004)

Focusing on the optimization version of the random *K*-satisfiability problem, the MAX-*K*-SAT problem, we study the performance of the finite energy version of the survey propagation algorithm. We show that a simple (linear time) backtrack decimation strategy is sufficient to reach configurations well below the lower bound for the dynamic threshold energy and very close to the analytic prediction for the optimal ground states. A comparative numerical study on one of the most efficient local search procedures is also given.

DOI: 10.1103/PhysRevE.70.036107 PACS number(s): 02.50.−r, 75.10.Nr, 02.60.Pn, 05.20.−y

I. INTRODUCTION

The problem of finding variable configurations that minimize the energy of a system with competitive interactions has been and still is a central one in the study of complex systems, like spin glasses in physics, protein folding and regulatory networks in biology, and optimization problems in computer science (see e.g., Refs. [1–5]).

Among the tools for numerical investigations of complex systems at low temperatures the simulated annealing (SA) algorithm [6] and its variants have played a major role. Such stochastic processes satisfy detailed balance and their behavior can be compared with static and dynamical mean-field calculations. However, in problems in which the interest is focused on zero temperature ground states and where the proliferation of metastable states causes an exponential slowdown in the equilibration rate, the applicability of SA-like algorithms is limited to relatively small system sizes.

In computer science the field of combinatorial optimization [7] deals precisely with the general issue of classifying the computational difficulty ("hardness") of minimization problems and of designing search algorithms. Similarly to statistical physics models, a generic combinatorial optimization problem is composed of many discrete variables—e.g., Boolean variables, finite sets of colors or Ising spins—which interact through constraints typically involving a small number of variables, that in turn sum up to give the global cost-energy function.

When the problem instances are extracted at random from nontrivial ensembles (that is ensembles which contain many instances that are hard to solve), computer science meets physics in a very direct way: many of the models considered to be of basic interest for computer science are nothing but spin glasses defined over finite connectivity random graphs, the well studied diluted spin glasses [8,9]. Their associated energy function counts the number of violated constraints in the original combinatorial problem (with ground states corresponding to optimal solutions). Understanding the onset of hardness of such systems is at the same time central to computer science and to $T=0$ statistical physics with surprisingly concrete engineering applications. For instance, among the most effective error correcting codes and data compression methods are the low density parity check algorithms [10–12], which indeed implement an energy minimization of a spin glass energy defined over a sparse random graph. In such problems, the choice of the graph ensemble is a part of the designing techniques, a fact that makes spin glass theory directly applicable.

The above example is however far from representing the general scenario for combinatorial problems: in many situations the probabilistic set up is not defined and, consequently, the notion of typical-case analysis does not play any obvious role. The study of the connection (if any) between worst-case and typical-case complexity is indeed an open one and very few general results are known [13]. Still, a precise understanding of nontrivial random problem instances promises to be important under many aspects. New algorithmic results as well as many mathematical issues have been put forward by the statistical physics studies, with examples ranging from phase transitions [14,15] and out-of-equilibrium analysis of randomized algorithms [16] to new classes of message-passing algorithms [17,18].

The physical scenario for the diluted spin glasses version of hard combinatorial problems predicts a trapping in metastable states for exponentially long times of local search dynamic process satisfying detailed balance. Depending on the models and on the details of the process—e.g., cooling rate for SA—the long time dynamics is dominated by different types of metastable states at different temperatures [19]. A common feature is that at zero temperature and for simulation times which are subexponential in the size of the problem there exists an extensive gap in energy which separates the blocking states from true ground states.

Such behavior can be tested on concrete random instances which therefore constitute a computational benchmark for more general algorithms. Of particular interest for computer science are randomized search processes which do not properly satisfy detailed balance and that are known (numerically) to be more efficient than SA-like algorithms in the search for ground states [20]. Whether the physical blocking scenario applies also to these artificial processes, which are not necessarily characterized by a proper Boltzmann distri-

*Electronic address: battagli@sissa.it
†Electronic address: kolarmi@sissa.it
‡Electronic address: zecchina@ctp.trieste.it

bution at long times, is a difficult open problem. The available numerical results and some approximate analytical calculations [21,22] seem to support the existence of a thermodynamical gap, a fact which is of up-most importance for optimization. For this reason (and independently from physics), during the last decade the problem of finding minimal energy configurations of random combinatorial problems similar to diluted spin-glasses—e.g., random *K*-satisfiability (K-SAT) or graph coloring—has become a very popular algorithmic benchmark in computer science [9].

In the last few years there has been a great progress in the study of spin glasses over random graphs which has shed new light on mean-field theory and has produced new algorithmic tools for the study of low energy states in large single problem instances. Quite surprisingly, problems which were considered to be algorithmically hard for local search algorithms, like for instance random K-SAT close to a phase boundary, turned out to be efficiently solved by the survey propagation (SP) algorithm arising from the replica symmetry-broken (RSB) cavity approach to diluted spin glasses. Such type of results calls for a rigorous theory of the functioning of SP (which is a nonlocal process) and bring new mathematical challenges of potential practical impact.

Scope of this paper is to display a set of numerical and algorithmic results which complete previously published results on the SP algorithm. We shall deal only with the random K-SAT problem even though we expect the algorithmic outcomes to be applicable to other similar problems like, for instance, the random graph coloring.

The paper is organized as follows. In Secs. II and III we briefly review the known results on random K-SAT together with the SP equations over single instances at finite pseudotemperature. We discuss as well in Sec. IV how the SP algorithm can be modified in order to study the region of parameters with finite ground state energy (not satisfiable or UNSAT phase), where not all constraints of the underlying random K-SAT problem can be satisfied simultaneously. In Sec. V we discuss then the performance of SP as an optimization device. At variance with the SAT phase in which many clusters of zero energy configurations coexist and where SP works efficiently without need of correcting variable assignments, in the UNSAT phase an efficient implementation of SP requires the introduction of—at least—a very simple form of backtracking procedure (similar to the one proposed in Ref. [23]). We show that a linear time backtrack is enough to reach energies compatible with those predicted by the analytic calculations in the infinite size limit in the relevant region of parameters. We give moreover numerical evidence for the existence of threshold states for one of the most efficient randomized local search algorithms for solving random K-SAT, namely WALKSAT [24]. We display a blocking mechanism at an energy level which is definitely above the lower bound for the dynamical threshold states predicted by the stability analysis of the 1-RSB cavity equations. Finally, for the deep UNSAT phase, we report on numerical data on convergence times for both WALKSAT and SA which are in agreement with the predicted existence of full RSB phases. Conclusions and perspectives are briefly discussed in Sec. VI.

II. BRIEF REVIEW OF RANDOM K-SAT

K-SAT is a *NP*-complete problem [25] (for $K>2$) which lies at the root of combinatorial optimization. It is very easy to state: Given N Boolean variables and M constraints taking the form of clauses, K-SAT consists in asking whether it exists an assignment of the variables that satisfies all constraints. Each clause contains exactly K variables, either directed or negated, and its truth value is given by the OR function. Since the same variable may appear directed or negated in different clauses, competitive interactions among clauses may set in.

As mentioned in the Introduction, in the last decade there has been a lot of interest on the random version of K-SAT: for each clause the variables are chosen uniformly at random (with no repetitions) and negated with probability 1/2.

In the large N limit, random K-SAT displays a very interesting threshold phenomenon. Taking as control parameter the ratio of number of clauses to number of variables, $\alpha = M/N$, there exists a phase transition at a finite value $\alpha_c(K)$ of this ratio. For $\alpha < \alpha_c(K)$ the generic problem is satisfiable (SAT), for $\alpha > \alpha_c(K)$ the generic problem is not satisfiable (UNSAT).

This phase transition has been seen numerically [26] and it is of special interest since extensive experiments [9] have shown that the instances which are algorithmically hard to solve are exactly those where α is close to α_c. Therefore, the study of the SAT-UNSAT phase transition is considered of crucial relevance for understanding the onset of computational complexity in typical instances [8]. A lot of work has been focused on the study of both the decision problem (i.e., determining with a YES/NO answer whether a satisfying assignment exists), and the optimization version in which one is interested in minimizing the number of violated clauses when the problem is UNSAT (random MAX-K-SAT problem).

On the analytical side, there exists a proof that the threshold phenomenon exists at large N [27], although the fact that the corresponding α_c has a limit when $N \to \infty$ has not yet been established rigorously. Upper bounds $\alpha_{\mathrm{UB}}(K)$ on α_c have been found using first moment methods [28] and variational interpolation methods [29], and lower bounds $\alpha_{\mathrm{LB}}(K)$ have been found using either explicit analysis of some algorithms [30], or some second moment methods [31]. For random MAX-K-SAT theoretical bounds are also known [32,33], as well as rigorous results on the running times of random walk and approximation algorithms [34–36].

Recently, the cavity method of statistical physics has been applied to K-SAT [15,17,37] and the thresholds have been computed with high accuracy. A lot of work is going on in order to provide a rigorous foundation to the cavity results and we refer to Ref. [37] for a more complete discussion of these aspects.

In what follows we shall concentrate on the $K=3$ case and we will be interested in analyzing the behavior of different algorithms in the region of parameter in which the random formulas are expected to be hard to solve or to minimize. The energy function which is used in the zero temperature statistical mechanics studies is taken proportional to the

number of violated clauses in a given problem so that a zero energy ground state corresponds to a satisfying assignment. The energy of a single clause is positive (equals 2 for technical reasons) if the clause is violated and zero if it is satisfied. The overall energy is obtained by summing over clauses and reads

$$E = 2\sum_a \frac{\Pi_{i=1}^3 (1 + J_{a,i} s_i^a)}{2}, \tag{1}$$

where s_i^a is the ith binary (spin) variable appearing in clause a and the coupling $J_{a,i}$ takes the value 1 (respectively, -1) if the corresponding variable appears not negated (respectively, negated) in clause a. For instance the clause $(x_1 \vee \overline{x_2} \vee x_3)$ has an energy $\frac{1}{4}(1+s_1)(1-s_2)(1+s_3)$ where the Boolean variables $x_i = \{0,1\}$ are connected to the spin variables by the transformation $s_i = (-1)^{x_i}$.

The phase diagram of the random 3-SAT problem as arising from the statistical physics studies can be very briefly summarized as follows.

For $\alpha < 3.86$, the $T=0$ phase is at zero energy (the problem is SAT). The entropy density is finite and the phase is replica symmetric (RS) and unfrozen. Roughly speaking, this means that there exists one giant cluster of nearby solutions and that the effective fields vanish linearly with the temperature.

For $3.86 < \alpha < 3.92$, there is a full RSB phase. The solution space breaks in clusters and the order parameter becomes a nested probability measure in the space of probability distribution describing cluster to cluster fluctuations. The phase is still SAT and unfrozen [38,39].

At $\alpha \simeq 3.92$ there is a discontinuous transition toward a clustered frozen phase [15,17]. Up to $\alpha = 4.15$ the phase is full RSB while above the 1-RSB solution becomes stable [40]. The complexity, that is the normalized logarithm of the number of clusters, is finite in this region. At finite energy there exist even more metastable states which act as dynamical traps. The 1-RSB metastable states become unstable at some energy density $E_G(\alpha)$ which constitutes a lower bound to the true dynamical threshold energy (see Sec. III for more details).

At $\alpha = 4.2667$ the ground state energy becomes positive and therefore the typical random 3-SAT problem becomes UNSAT. At the same point the complexity vanishes. The phase remains 1-RSB up to $\alpha = 4.39$ where an instability toward a zero complexity full RSB phase appears.

In the region $4.15 < \alpha < 4.39$, the 1-RSB ansatz for the ground state is stable against higher orders of RSB, but the 1-RSB predictions become unstable for energies larger than the Gardner energy. The instability line intersects with the 1-RSB ground state estimation at the two extremes of the interval, inside which it provides a lower bound to the true threshold energy (see Ref. [40] for a comprehensive discussion).

Further (preliminary) full-RSB corrections suggest that the true threshold states have energies very close to the lower bound and hence the interval $A = [4.15, 4.39]$ should be taken as the region where to take really hard benchmarks for algorithm testing. As displayed in Fig. 1, the actual value of the energy gap is very small close to the end points of A. In order to avoid systematic finite size errors, numerical simulations should be done close to the SAT-UNSAT point, i.e., far from the end point of A. Consistently with the fact that finite size fluctuations are relatively big $[O(\sqrt{N})]$, even close to α_c problem sizes of the order at least of $N = 10^5$ are necessary in order to observe a matching with the analytic predictions.

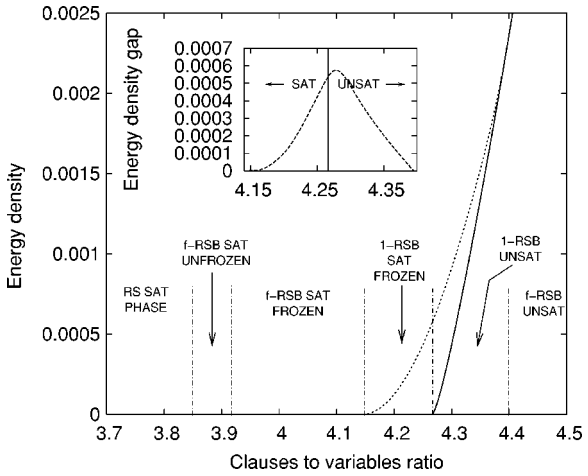

FIG. 1. The solid line is an estimation for the ground state energy, while the dashed curve represents the Gardner energy, providing a lower bound for the threshold states (numerical data adapted from Ref. [40]). In the inset we show that the difference between the Gardner and the ground state energy is strictly positive in the small 1-RSB stable region around the SAT/UNSAT transition critical point (indicated by the vertical line): it is expected that it is hard for heuristics based on local search to find assignments inside the closed area delimited by the energy gap curve.

III. BRIEF REVIEW OF SP EQUATIONS

The 1-RSB cavity equations which have been used to study the typical phase diagram of random K-SAT become the SP equations once reformulated to run over single problem instance [17]. This is done by avoiding the averaging process with respect to the underlying random graphs. Thanks to the self-averaging property of the random K-SAT free energy [41], the SP equations can be used both to rederive the phase diagram of the problem and, more important, to access detailed information of algorithmic relevance about a given problem instance. In particular, the SP equations provide information about the statistical behavior of the single variables in the stable and metastable states of given energy density.

The 1-RSB cavity equations are iterative equations (averaged over the disorder) for the probability distribution functions (PDF) of effective fields that describe their cluster-to-cluster fluctuations. The order parameter is a probability measure in the space of PDF's; it tells the probability that a randomly chosen variable has a certain associated PDF in states at a given energy density.

In SP and more in general in the cavity approach, one assumes to know PDF's of the fields of all variables in the

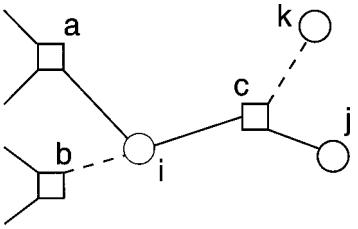

FIG. 2. Factor graph representation. Variables are represented by circles, and are connected by function nodes, represented by squares; if a variable appears negated in a clause, the connecting line is dashed.

temporary absence of one of them. Then one writes the induced PDF of the local field acting on this "cavity" variable in absence of some other variable interacting with it (i.e., the so-called Bethe lattice approximation for the problem). These relations define a closed set of equations for the PDF's that can be solved iteratively. The equations are exact if the cavity variables acting as inputs are uncorrelated, e.g., over trees, or are conjectured to be an asymptotically exact approximation over locally treelike structures [17] where the typical distance between randomly chosen variables diverges in the large N limit (as $\ln N$ for diluted random graphs). The full list of the cavity fields over the entire underlying graph, in the SP implementation, constitutes the order parameter. From the cavity fields one may determine the total field acting on each variable in all metastable states of given energy density and this information can be used for algorithmic purposes.

A clear formalism for the single sample analysis is given by the factor graph representation [42] of K-SAT: variables are represented by N circular "variable nodes" labeled with letters $i,j,k,\ldots$ whereas the K-body interactions are represented by M square "function nodes" (carrying the clause energies) labeled by $a,b,c,\ldots$ (see Fig. 2).

For random 3-SAT, function nodes have connectivity 3, variable nodes have a Poisson connectivity of average 3α and the overall graph is bipartite. The total energy is nothing but the sum of energies of all function nodes as given by Eq. (1).

Adopting the message-passing notation and strictly following [17], we call u-messages the contribution to the cavity fields coming from the different connected branches of the graph. In SP the messages along the links of the factor graph have a functional nature carrying information about distributions of u-messages over the states at a given value of the energy, fixed by a Lagrange multiplier y: we call these distributions of messages u-surveys. The SP equations can be written at any "temperature" (the inverse of the Lagrange multiplier y is actually a pseudotemperature, see Ref. [17]). However they acquire a particularly simple form in the limit $1/y \to 0$, which is the limit of interest for optimization purposes, at least in the SAT region.

In K-SAT, the u-surveys are parametrized by two real numbers and SP can be implemented very efficiently. Each edge $a \to i$, from a function node a to a variable node i, carries a u-survey $Q_{a\to i}(u)$. From these u-surveys one can compute the cavity fields $h_{i\to b}$ for every neighbor b, which in

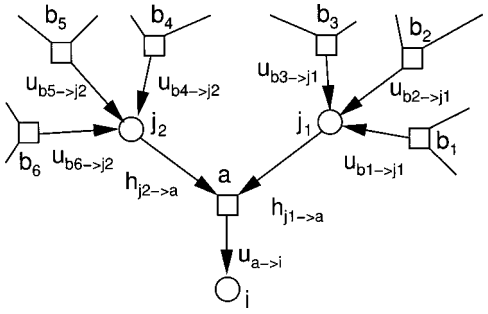

FIG. 3. Cavity fields and u-messages. The u-survey for the u-message $u_{a\to i}$ depends on the PDF's of the cavity fields $h_{j1\to a}$ and $h_{j2\to a}$. These are on the other side dependent on the u-surveys for the u-messages incoming to the variables j_1 and j_2.

turn determine new output u-surveys (see Fig. 3).

Very schematically, the SP equations can be implemented as follows. Let $V(i)$ be the set of function nodes connected to the variable i, $V(a)$ the set of variables connected to the function node a; let us denote by $V(i)\backslash a$ and $V(a)\backslash i$ the same sets deprived, respectively, of the clause a and of the variable i. Given then a random initialization of all the u-surveys $Q_{a\to i}(u)$, the function nodes are selected sequentially at random and the u-surveys are updated according to a complete set of coupled functional equations (see Fig. 3 for the notation):

$$P_{j\to a}(h_{j\to a}) = C_{j\to a}\int \mathcal{D}Q_{j,a}\,\delta\Big(h - \sum_{b\in V(j)\backslash a} u_{b\to j}\Big) \times \exp\Big(y\Big(\Big|\sum_{b\in V(j)\backslash a} u_{b\to j}\Big| - \sum_{b\in V(j)\backslash a} |u_{b\to j}|\Big)\Big), \tag{2}$$

$$Q_{a\to i}(u) = \int \mathcal{D}P_{a,i}\,\delta(u - \hat{u}_{a\to i}(\{h_{j\to a}\})), \tag{3}$$

where the $C_{i\to a}$'s are normalization constants, the function $\hat{u}_{a\to i}$ is

$$\hat{u}_{a\to i}(\{h_{j\to a}\}) = J_{a,i}\prod_{j\in V(a)\backslash i}\theta(J_{a,j}h_{j\to a}), \tag{4}$$

and the integration measures are given by

$$\mathcal{D}Q_{j,a} = \prod_{b\in V(j)\backslash a} Q_{b\to j}(u_{b\to j})du_{b\to j}, \tag{5}$$

$$\mathcal{D}P_{a,i} = \prod_{j\in V(a)\backslash i} P_{j\to a}(h_{j\to a})dh_{j\to a}. \tag{6}$$

Parametrizing the u-surveys as

$$Q_{a\to i}(u) = \eta^0_{a\to i}\delta(u) + \eta^+_{a\to i}\delta(u-1) + \eta^-_{a\to i}\delta(u+1), \tag{7}$$

where $\eta^0_{a\to i} = 1 - \eta^+_{a\to i} - \eta^-_{a\to i}$, the above set of equations (2) and (3) defines a nonlinear map over the η's.

Once a fixed point is reached, from the list of the u-surveys one may compute the normalized PDF of the local field acting on each variable,

$$P_i(H) = C_i \int \mathcal{D}\hat{Q}_i \, \delta\left(H - \sum_{b \in V(i)} u_{b\to i}\right) \times \exp\left(y\left(\left|\sum_{b\in V(i)} u_{b\to i}\right| - \sum_{b\in V(i)} |u_{b\to i}|\right)\right), \quad (8)$$

$$\mathcal{D}\hat{Q}_i = \prod_{b\in V(i)} Q_{b\to i}(u_{b\to i}) du_{b\to i}. \quad (9)$$

It should be remarked that $P_i(H)$ is in general different from the family of cavity fields PDF's $P_{i\to b}(h)$ computed by mean of (2).

From the knowledge of the cavity and local fields PDF's, one derives the (Bethe) free energy at the level of 1-RSB,

$$\Phi(y) = \frac{1}{N}\left(\sum_{a=1}^{M} \Phi_a^f(y) - \sum_{i=1}^{N} \Phi_i^v(y)(\Gamma_i - 1)\right), \quad (10)$$

where Γ_i is the connectivity of the variable i and

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

$$\Phi_i^v(y) = -\frac{1}{y}\ln\left\{\int \mathcal{D}\hat{Q}_i \exp\left[y\left(\left|\sum_{a\in V(i)} u_{a\to i}\right| - \sum_{a\in V(i)}|u_{a\to i}|\right)\right]\right\} = -\frac{1}{y}\ln(C_i). \quad (11)$$

Here, E_a is the energy contribution of the function node a. The maximum value of the free-energy functional provides a lower bound estimation of the ground state energy of the Hamiltonian (1) defined on the sample. In the SAT region the free-energy functional $\Phi(y)$ is always nonpositive and it is increasing in the limit $y\to\infty$; in the UNSAT region, on the contrary, it exhibits a positive maximum for $y=y^*$ (see [17]).

From the free-energy density of a given instance, it is straightforward to compute numerically its complexity $\Sigma(y) = \partial\Phi(y)/\partial(1/y)$ and its energy density $\epsilon(y) = \partial[y\Phi(y)]/\partial y$. We remind that the complexity is linked to the number of pure states (i.e., clusters of configurations) of energy E, by the defining relation $\mathcal{N}(E) = \exp[N\Sigma(E)]$. The energy level represented by the largest number of configurations, $e_{\rm th}$, is given by

$$\Sigma(e_{\rm th}) = \max_E[\Sigma(E)]. \quad (12)$$

Further RSB corrections may be needed to locate the precise value of $e_{\rm th}$, which is in any case lower bounded the largest energy of 1-RSB stable states, the so-called Gardner energy E_G. It is expected that local search strategies get trapped at energies close, but not necessarily equal, to the threshold energy (see Ref. [19] for a throrough discussion on the role of the isocomplexity states [43]). More elaborated strategies not properly satisfying detailed balance (e.g., WALKSAT for the K-SAT problem) could in principle overcome this type of barriers; however, the available numerical and analytical results suggest that also these more sophisticated randomized searches undergo an exponential slowdown, with different layers of states acting as dynamical traps, depending on the details of the heuristics.

IV. SP IN THE UNSAT REGION

In the SAT phase, where the $y\to\infty$ limit is taken, the convolutions (2) filter out completely any clause-violating truth value assignment. This feature is extremely useful for satisfiable formulas, but it becomes undesired when our sample is presumably unsatisfiable.

In the UNSAT region the SP equations require a finite value of the Lagrange multiplier y. The filtering action of the exponential reweighting term in (2) is then weakened and the messages computed by the SP equations can vehicle information pointing to states with a nonvanishing number of violated constraints.

A. The finite pseudotemperature recursive equations

The SP equations simplify considerably in the $y\to\infty$ limit and lead to extremely efficient algorithmic implementations, as discussed in great detail in Ref. [18]. In the case of finite pseudotemperature $1/y$ the same simplification cannot take place because of the presence of a nontrivial reweighting factor. Still, a relatively fast recursive procedure can be written. Let us consider a variable j having Γ_j neighboring function nodes and let us compute the cavity field PDF $P_{j\to a}(h)$ where $a\in V(j)$. We start by randomly picking up one function node in $V(j)\backslash a$, denoted as b_1, and we calculate the following "h-survey":

$$\widetilde{P}^{(1)}_{j\to a}(h) = \eta^0_{b_1\to i}\delta(h) + \eta^+_{b_1\to i}\delta(h-1) + \eta^-_{b_1\to i}\delta(h+1). \quad (13)$$

The function $\widetilde{P}^{(1)}_{j\to a}(h)$ would correspond to the true local field PDF of the variable j in the case in which b_1 was the only neighboring clause (as denoted by the upper index).

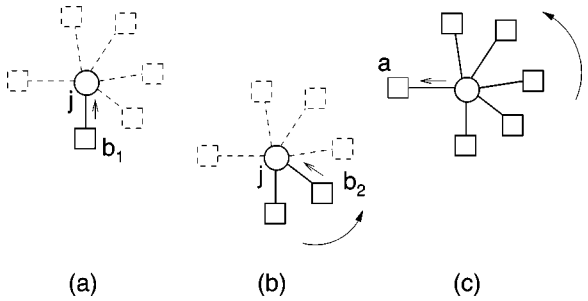

FIG. 4. Computing recursively a cavity PDF. (a) In order to find a single cavity PDF $P_{j\to a}(h)$, a single clause b_1 in $V(j)\backslash a$ is picked up at random and the u-survey $Q_{b_1\to j}$ is used to compute Eq. (13); (b) the contributions of all the other function nodes in $V(j)\backslash a$ are then added, clause by clause; (c) the PDF computed recursively after Γ_j-1 iterations coincides with $P_{j\to a}(h)$.

The following steps of the recursive procedure consist in adding the contributions of all the other function nodes in $V(j)\backslash a$, clause by clause (Fig. 4):

$$\widetilde{P}^{(\gamma)}_{j\to a}(h) = \eta^0_{b_\gamma\to j}\widetilde{P}^{(\gamma-1)}_{j\to a}(h) + \eta^+_{b_\gamma\to j}\widetilde{P}^{(\gamma-1)}_{j\to a}(h-1)\exp[-2y\hat{\theta}(-h)] + \eta^-_{b_\gamma\to j}\widetilde{P}^{(\gamma-1)}_{j\to a}(h+1)\exp[-2y\hat{\theta}(h)]. \quad (14)$$

Here $\widetilde{P}^{(\gamma)}_{j\to a}(h)$ is an unnormalized PDF and $\hat{\theta}(h)$ is a step function equal to 1 for $h\geqslant 0$ and zero otherwise. The recursion ends after $\gamma=\Gamma_j-1$ steps, when the influence of every clause in $V(j)\backslash a$ has been taken in account. The final cavity-field PDF $P_{j\to a}(h)$ can be found straightforwardly by computing the PDF $\widetilde{P}^{(\Gamma_j-1)}_{j\to a}(h)$ for all values of the field $-\Gamma_j+1<h<\Gamma_j-1$ and by normalizing it.

As already pointed out in Sec. III, the knowledge of $K-1$ input cavity-field PDF's can be used to obtain a single output u-survey. Let us compute for instance the u-survey $Q_{a\to i}(u)$ (see always Fig. 3 for the notation). In order to do that, we need first the cavity field PDF's $P_{j\to a}(h)$ for every $j\in V(a)\backslash i$. The parameters $\{\eta^0_{a\to i},\eta^+_{a\to i},\eta^-_{a\to i}\}$ are then updated according to the formulas

$$\eta^{J_{a,i}}_{a\to i} = \prod_{n=1}^{K-1} W^{J_{j_n,a}}_{j_n\to a}, \quad \eta^{-J_{a,i}}_{a\to i}=0, \quad \eta^0_{a\to i}=1-\eta^{J_{a,i}}_{a\to i}, \quad (15)$$

where we introduced the weight factors

$$W^+_{j\to a}=\sum_{h=1}^{\Gamma_j-1}P_{j\to a}(h), \quad W^-_{j\to a}=\sum_{h=-\Gamma_j+1}^{-1}P_{j\to a}(h). \quad (16)$$

It should be remarked that $Q_{a\to i}(u)$ depends only on one single nontrivial $\eta^{J_{a,i}}_{a\to i}$ (from now simply referred to as $\eta_{a\to i}$). We could say that a single kind of message can be produced, telling the receiver literal to assume the truth value "TRUE"; this message is transmitted along the edge $a\to i$ with a probability $\eta_{a\to i}$, corresponding to the probability that the only way of not violating the constraint a is to set appropriately the truth value of i.

Starting from a full collection of u-surveys at a given time, it is possible to realize a complete update of all the parameters $\{\eta_{a\to i}\}$ by systematical application of the recursions (13) and (14) and of the relation (15); from the new set of u-surveys, new cavity field PDF's can be computed and the procedure continues until when self-consistence of η's is reached. This procedure can be efficiently implemented numerically and allows us to determine the fixed point of the population-dynamics equations (2) and (3), for a general value of y.

B. The SP-Y algorithm

In the usual SP-inspired decimation [18], the computation of the local field PDF's $P_i(H)$ is used to decide a truth value assignment for the most biased variables. Indeed, it is reasonable that a spin tends to align itself with the most probable direction of the local field. A ranking can be realized by finding all the probability weights

$$W^+_j=\sum_{H=1}^{\Gamma_j}P_j(H), \quad W^-_j=\sum_{H=-\Gamma_j}^{-1}P_j(H), \quad (17)$$

and by sorting the variables according to the values of a bias function

$$b_{\mathrm{fix}}(j)=|W^+_j-W^-_j|. \quad (18)$$

The local field PDF's $P_j(H)$ can be naturally calculated resorting to the iterations (13) and (14): computation is done simply by sweeping over the whole set of neighboring function nodes $V(j)$, including also the contribution of the skipped edge $a\to j$. By fixing in the right direction the spin of the most biased variable, we actually reduce the original N variable problem to a new one with $N-1$ variables. New u-surveys are then computed. Doing that we have to take care of fixed variables: if i is fixed, its cavity field PDF's must be of the form

$$P_{i\to a}(h)=\delta(h-J_{a,i}s_i), \quad (19)$$

regardless of the recursions (13) and (14). The complete polarization reflects the knowledge of the truth value of the literals depending on the spin s_i.

The procedure of decimation continues until when a full truth assignment has been generated or until when convergence has been lost or a paramagnetic state has been reached; in the latter cases the original formula is simplified according to the partial truth assignment already generated and the simplified formula is passed to a specialized heuristic. Our choice of preference is the WALKSAT algorithm [24], which is by far more efficient than SA in the hard region of the 3-SAT problem, as we have checked exhaustively. Very briefly, the strategy of WALKSAT is the following one: at each time step the current assignment is changed by randomly alternating greedy moves (where the variable which maximizes the number of satisfied clauses is fixed) and random-walk steps (in which a variable belonging to a randomly chosen unsatisfied clause is selected and flipped). WALKSAT stops if either a satisfying assignment is found or if the maximum number of allowed spin flips (the "cutoff") is reached (see Ref. [44]

INPUT: a Boolean formula $\mathcal{F}$ in conjunctive normal form; a backtracking ratio r; optionally, a fixed inverse pseudo-temperature y_{in}

OUTPUT: a simplified Boolean formula $\mathcal{F}'$ in conjunctive normal form (ideally empty) and a partial truth value assignment for the variables of $\mathcal{F}$ (ideally a complete one)

0. For each edge $a \to i$ of the factor graph, randomly initialize the $\eta_{a\to i} \in \{0,1\}$
1. IF there is a fixed y_{in} as input, put $y^* = y_{in}$, ELSE after a fixed number of steps, determine by bisection the position of the free-energy maximum y^*
2. Compute all the fixed point u-surveys, using equations (13), (14), (15) and putting $y = y^*$
3. IF the population dynamics equations converge,
   - 3.1 FOR every unfixed variable i, compute the local field pdf using (13), (14)
   - 3.2 Extract a random number q in [0,1]
   - 3.3 IF $q \geq r$, Sort the variables according to the index function (18), and fix the most biased variable
   - 3.4 ELSE IF $q < r$, Sort the variables according to the index function (20) and unfix the highest ranked variable
   - 3.5 IF all the variables are fixed, RETURN the full truth value assignment and an empty sub-formula, ELSE, go to 1.
4. ELSE IF the population dynamics equations do not converge, simplify the formula by imposing the already assigned truth values, RETURN the partial solution and the obtained sub-formula

FIG. 5. The SP-Y simplification algorithm.

for another recently analyzed and very efficient heuristics).

When working at finite pseudotemperature, we have to take in account the possibility that some nonoptimal fixing is done in presence of thermal "noise." After several updates of the u-surveys some biases of fixed spins may become smaller than the value they had at the time when the corresponding spin was fixed. Certain local fields can even revert their orientation. Small or positive values of an index function like

$$b_{\text{backtrack}}(j) = -s_j(W_j^+ - W_j^-) \tag{20}$$

can track the appearance of such dangerous fixed spins and this information can be used to implement some "error removal" procedure; for instance, a simple strategy can be devised where both unfixing and fixing moves are performed at a fixed ratio $0 \leqslant r < 0.5$ (see Ref. [23] for another backtracking implementation).

The actual SP with finite y simplification procedure (SP-Y) will depend not only on the backtracking fraction r, but even more on the choice of the inverse pseudotemperature y. The simplest possibility is to keep it fixed during the simplification, but one may choose to dynamically update it, in order to stay as close as possible to the maximum y^* of the free energy functional $\Phi(y)$ (which corresponds to select the ground state in the 1-RSB framework, as we have seen in Sec. III).

The equations (10) and (11) can be rewritten in the following form, suitable for numerical computation:

$$\Phi_a^f(y) = -\frac{1}{y}\left[\ln\left(1 + (e^{-y} - 1)\prod_{i\in V(a)} W_{i\to a}^{J_{a,i}}\right) - \ln\left(\prod_{i\in V(a)} C_{i\to a}\right)\right], \tag{21}$$

$$\Phi_i^v(y) = -\frac{1}{y}\ln(C_i). \tag{22}$$

In Fig. 5 we give a summary of the simplification procedure in a standard pseudocode notation. The first release of the SP-Y code can be downloaded from Ref. [45].

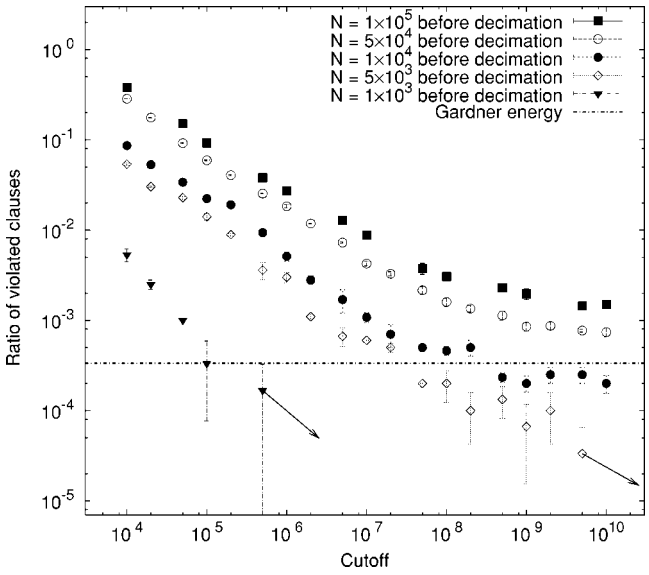

FIG. 6. Threshold energy effect in SAT region. The WALKSAT performance for various samples of different sizes and $\alpha=4.24$ is presented. With increasing size the curves appear to saturate above the Gardner energy. An arrow indicates that the next data point corresponds to a SAT assignment.

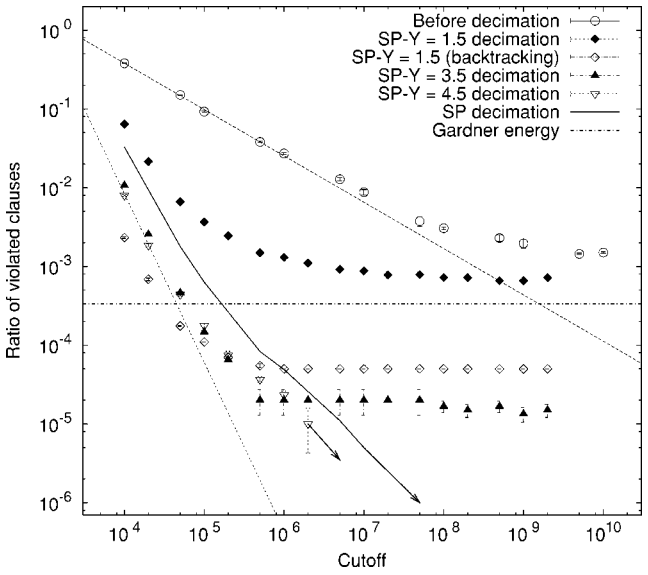

FIG. 7. Efficiency of SP-Y in the SAT region (single sample with $N=10^5$ variables and $\alpha=4.24$). After SP-Y simplification, WALKSAT is generally able to find solutions below the Gardner threshold; in some cases, it succeeds even in finding complete satisfying assignment. An arrow indicates that the next data point corresponds to a SAT assignment.

V. OPTIMIZING THE ENERGY BELOW THE THRESHOLD STATES

As we have already discussed in Sec. III, it is expected that, in the thermodynamical limit, any local search algorithm gets trapped in the vicinity of exponentially numerous threshold states with energy e_{th} and that any local heuristics is in general unable to find the optimal assignment in the thermodynamical limit. To verify this prediction, we conducted various experiments, both in the SAT and in the UNSAT phase, focusing on the comparison between the WALKSAT heuristics performance after and before different kinds of SP-Y simplification. In most of the situations, we decided to analyze carefully single large-sized samples instead of a larger number of smaller problems: we verified in fact that the sample-to-sample fluctuations tend to be unrelevant for size of order 10^4 and larger.

A. SAT region

The aim of the first set of experiments was to check the actual existence of the threshold effect. We ran WALKSAT over different formulas in the frozen-SAT region, with fixed $\alpha=4.24$ and sizes varying between $N=10^3$ and $N=10^5$, reaching a maximum cutoff of 10^{10} spin flips. The obtained results are plotted in Fig. 6; the Gardner energy is also reported for comparison with the data. Even if for small-size samples the local search algorithm is able to find a SAT assignment, for larger formulas $[N\sim O(10^4)]$ WALKSAT does not succeed in reaching the ground state, its relaxation profile suffers of critical slowdown, and saturates at some well defined level. This is actually expected, because the Gardner energy becomes $\mathcal{O}(1)$ only for $N\sim 10^4$ or larger, and for a smaller number of variables the threshold effect should be negligible when compared to finite size effects.

We remind that WALKSAT cannot be considered as an equilibrium stochastic process and that it is not possible to infer that its saturation level coincides with the sample threshold energy; we can anyway claim that WALKSAT is unable to explore the full energy landscape of the problem, and that the enormous number of nonoptimal valleys is unavoidably hiding the true ground states. Plateaus in the relaxation profiles of WALKSAT have indeed been already discussed in Refs. [21,22] and ascribed to metastable states acting as dynamical traps.

For the $N=10^4$ formula a trapping effect becomes clearly visible in our experiments, but the saturation plateau is below the Gardner lower bound. The finite-size fluctuations are still of the same order of the energy gap between the ground and the threshold states and the experimental conditions are distant from the thermodynamical limit. When the size is increased up to 10^5 variables, the saturation level moves finally between the full RSB lower bound and the 1-RSB upper bound for e_{th}.

The efficiency of the SP-Y simplification strategy against the glass threshold is discussed in Fig. 7. We simplified a single randomly generated formula ($N=10^5, \alpha=4.24$) at several fixed values of pseudotemperature. The solid line shows for comparison the WALKSAT results after a standard SP decimation (i.e., $y\to\infty$): the ground state, $E=0$, is reached as expected, after a rather small number of spin flips. The same happens after SP-Y simplifications performed at a large enough inverse pseudotemperature ($y>4$); one should remind indeed that in the SAT region the optimal value for y would be infinite, and that in that limit the SP-Y recursions reduce to the SP equations. After simplification with smaller y's, the WALKSAT cooling curves reach again a saturation

level, which is nevertheless below the Gardner energy, unless y is too small: the threshold states of the original formula have not been able to trap the local search, even if the ground state becomes inaccessible. As we have indeed already discussed, working at finite temperature increases the probability of violating a clause when doing a spin fixing, and this is particularly evident in the SAT region where every assignment that does not satisfy some constraint should be filtered out.

The procedure is intrinsically error prone, and it will allow us in general to reach only "good states," but not the true optimal solutions (the smaller the parameter y, the higher the saturation level will be). As we shall discuss in the next section, the use of backtracking partially cures the accumulation of errors at finite y: the saturation level can in fact be significantly lowered by keeping the same pseudotemperature and introducing a small fraction of backtrack moves during the simplification. In Fig. 7 the data for $y=1.5$ show the importance of backtracking. While the run of SP-Y without backtracking has led to a plateau above Gardner energy, with the introduction of backtrack moves we find energies well below the threshold.

B. UNSAT region

When entering the UNSAT region, the task of looking for the optimal state becomes harder. The expected presence of violated constraints in the optimal assignments really forces us to run the simplification at a finite pseudotemperature. Unfortunately, after many spin fixings, the recursions (13) and (14) stop to converge for some finite value of y before the maximum of the free energy is reached, most likely because the subproblem has entered a full RSB phase. At this point one should switch to a 2-RSB version of SP which we did not realize, yet. Alternatively, one could try to run directly the final heuristic search (hoping that the full RSB subsystem is not exponentially hard to optimize) or more simply one may continue the decimation process by selecting the largest y for which the computations converge. We decided to implement the latter choice until either convergence is lost independently from the value of y or a paramagnetic state is reached.

In our experiments we studied several 3-SAT sample problems belonging to the 1-RSB stable UNSAT phase. We employed WALKSAT as an example of standard well-performing heuristics. Although WALKSAT is not optimized for unsatisfiable problems, in the 1-RSB stable UNSAT region it performs still much better than any basic implementation of SA. We observed anyway that, even after 10^{10} spin flips, the WALKSAT best assignments were still quite distant from the Gardner energy, for various samples of different size and α. In Fig. 8 we show the results relative to many different SP-Y simplifications with various values of y and r for a single sample with $N=10^5$ and $\alpha=4.29$. The simplification produced always an improvement in the WALKSAT performance, but, in absence of backtracking, we were unable to go below the Gardner lower bound (although we touched it in some cases: in Fig. 8 we show the data for a simplification at fixed $y=2.5$; a simplification with runtime optimization of y reached the same level).

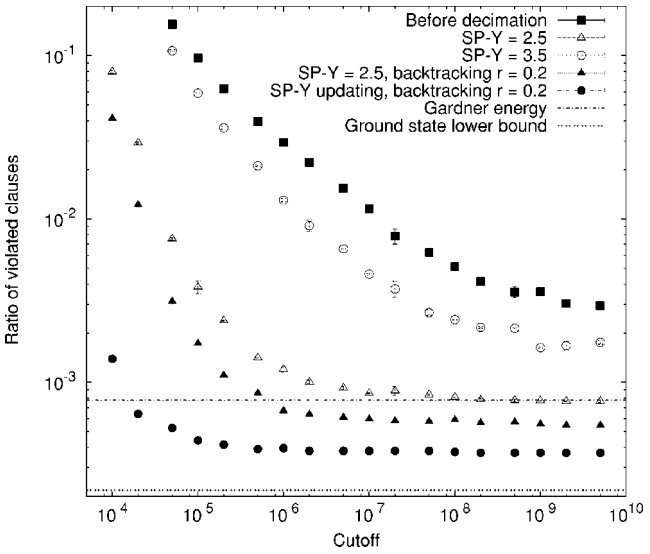

FIG. 8. SP-Y performance in the UNSAT region (single sample with $N=10^5$ variables and $\alpha=4.29$). Several simplification strategies are compared; the need for backtracking is readily visible, and its introduction allows us to reach energies closer to the ground state than to the Gardner lower bound.

The relative inefficiency of these first attempts of simplification was not due to the threshold effect alone, but also to an extreme sensitivity to the choice of y, as pointed out by a second set of experiments making use of backtracking. We performed first an extensive analysis of the simultaneous optimization of y and r, using smaller samples in order to produce more experimental points. After some trials, the fraction $r=0.2$ appeared to be the optimal one, at least for our implementation, and in the small region under investigation of the K-SAT phase diagram. The data in Fig. 9 refer to a formula with $N=10^4$ variables and $\alpha=4.35$. The dashed horizontal line shows the WALKSAT best energy obtained on the original formula after 10^9 spin flips. The WALKSAT performance was seriously degraded when simplifying at too small values of y, but the introduction of backtracking cured the problem, identifying and repairing most of the wrong assignments. The WALKSAT efficiency became actually almost independent from the choice of pseudotemperature, whereas in absence of error correction a time consuming parameter tuning was required for optimization.

Coming back to the analysis of the sample of Fig. 8, the backtracking simplifications allowed us to access states definitely below the Gardner lower bound. The combination of runtime y optimization and of error correction was even more effective: after a rather small number of spin flips, WALKSAT reached a saturation level strikingly closer to the ground state lower bound than to the Gardner energy. A further valuable effect of introduction of the backtracking was the increased efficiency of the formula simplification itself: in the backtracking experiments, SP-Y was able to determine a truth value for more than 80% of the variables before losing convergence, while without backtracking, the algorithm stopped on average after only 40% of fixings.

All the samples analyzed in the preceding sections were taken from the 1-RSB stable region of the 3-SAT problem,

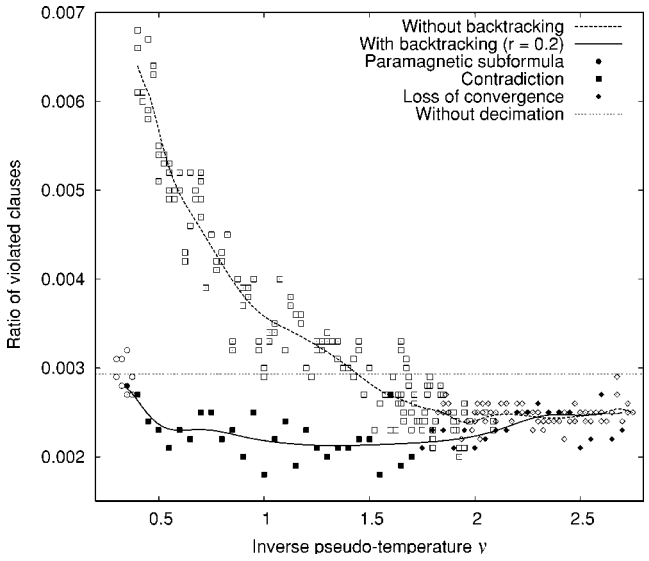

FIG. 9. Backtracking efficiency. Many SP-Y simplifications of a single sample with $N=10^4$ variables and $\alpha=4.35$ have been performed at fixed but different values of pseudotemperature; the introduction of a small fraction of backtracking moves eliminates essentially the need for a time consuming optimization of the parameter y. The empty points refer to simplifications without backtracking, the full points to simplifications with a backtracking ratio $r=0.2$. A diamond indicates that the simplification process stopped because of loss of convergence, a circle because of finding a completely unbiased paramagnetic state, and the squares indicate that the loss of convergence happened at an advanced stage where some clause-violating assignments have already been introduced by SP-Y.

where the equations (2) and (3) are considered to be exact. For $\alpha>4.39$, the phase becomes full RSB and SP loses convergence before the free energy $\Phi(y)$ reaches its maximum from the very first step of the decimation procedure. While a full RSB version of SP would most likely provide very good results, SP-Y still can be used in a suboptimal way by selecting the largest value of y for which convergence is reached. Numerical experiment shows that indeed the performance of SP-Y is in good agreement with the analytical expectations. However, it should be noticed that in this region the use of SP is not necessary. Although the performance of WALKSAT and SA can be improved by the SP simplification, the SA alone is already able of finding close-to-optimum assignments efficiently (as expected for a full RSB scenario) and behaves definitely better than WALKSAT.

VI. CONCLUSIONS

In this paper, we have displayed the performance of SP as an optimization device and shown that configurations well below the threshold states can be found efficiently. Similar results are expected to hold also for random satisfiable instances very close to the critical point for which the combined use of finite pseudotemperature and backtracking could give access to the SAT optima.

It would be of some interest to analyze further improvements of the decimation strategies as well as to consider more structured factor graphs within a variational framework, in which some correlations can be put under control.

A possible application of SP-Y-type algorithms can be found in information theory: lossy data compression based on low density parity check schemes leads to optimization problems which are indeed very similar to the one discussed in this paper.

ACKNOWLEDGMENTS

We thank A. Braunstein, M. Mézard, G. Parisi, and F. Ricci-Tersenghi for very fruitful discussions. This work was supported in part by the European Community's Human Potential Programme under Contract No. HPRN-CT-2002-00319, STIPCO.

[1] J.-P. Bouchaud, L. F. Cugliandolo, J. Kurchan, and M. Mézard, in *Spin Glasses and Random Fields*, edited by A. P. Young (World Scientific, Singapore, 1997).

[2] H. Frauenfelder, P. G. Wolynes, and R. H. Austin, Rev. Mod. Phys. **71**, S419 (1999); V. S. Pande, A. Y. Grosberg, and T. Tanaka, *ibid.* **72**, 259 (2000).

[3] M. Blatt, S. Wiseman, and E. Domany, Phys. Rev. Lett. **76**, 3251 (1996).

[4] A. K. Hartmann and H. Rieger, *Optimization Algorithms in Physics* (Wiley-VCH, Berlin, 2001).

[5] A. K. Hartmann and H. Rieger, *New and Advanced Optimization Algorithms in Physics and Computational Science* (Wiley-VCH, Berlin, 2004).

[6] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, Science **220**, 671 (1983).

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

[8] Theor. Comput. Sci. **265**, Issue 1–2 (2001), special issue on *NP-Hardness and Phase Transitions*, edited by O. Dubois, R. Monasson, B. Selman, and R. Zecchina.

[9] T. Hogg, B. A. Huberman, and C. Williams, Artif. Intell. **81**, 1 (1996).

[10] D. A. Spielman, Lect. Notes Comput. Sci. **1279**, 67 (1997).

[11] T. Richardson and R. Urbanke, in *Codes, Systems, and Graphical Models*, edited by B. Marcus and J. Rosenthal (Springer, New York, 2001).

[12] N. Sourlas, Nature (London) **339**, 693 (1989).

[13] N. Ajtai, Electronic Colloquium on Computational Complexity Report No. TR96–007, 1996.

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

[15] M. Mézard, G. Parisi, and R. Zecchina, Science **297**, 812 (2002) (Sciencexpress, published online 27 June 2002; 10.1126/science.1073287).

[16] S. Cocco, R. Monasson, A. Montanari, and G. Semerjian, e-print cs.CC/0302003.
[17] M. Mézard and R. Zecchina, Phys. Rev. E **66**, 056126 (2002).
[18] A. Braunstein, M. Mézard, and R. Zecchina, e-print cs.CC/0212002.
[19] A. Montanari and F. Ricci-Tersenghi, e-print cond-mat/0401649.
[20] R. Motwani and P. Raghavan, *Randomized Algorithms* (Cambridge University Press, Cambridge, 2000).
[21] G. Semerjian and R. Monasson, e-print cond-mat/0301272.
[22] W. Barthel, A. K. Hartmann, and M. Weigt, e-print cond-mat/0301271.
[23] G. Parisi, e-print cond-mat/0308510.
[24] B. Selman, H. Kautz, and B. Cohen, *Proc. AAAI-94* (AAAI Press, Seattle, WA, 1994), 337–343.
[25] M. R. Garey and D. S. Johnson, *Computers and Intractability* (Freeman, New York, 1979).
[26] S. Kirkpatrick and B. Selman, Science, **264**, 1297 (1994).
[27] E. Friedgut, J. Am. Math. Soc. **12**, 1017 (1999).
[28] O. Dubois, Y. Boufkhad, and J. Mandler, *Proeedings of the 11th ACM-SIAM Symposium on Discrete Algorithms* (San Francisco, CA, 2000);A. Kaporis, L. Kirousis, and E. Lalas, *Proceedings of the 4th European Symposium on Algorithms* (ESA, 2002), Lecture Notes in Computer Science (Springer, New York, in press).
[29] F. Guerra, Commun. Math. Phys. **233**, 1 (2003); S. Franz and M. Leone, J. Stat. Phys. **111**, 535 (2003).
[30] J. Franco, Theor. Comput. Sci. **265**, 147 (2001); D. Achlioptas and G. Sorkin, *41st Annual Symposium on Foundations of Computer Science* (IEEE Computer Society Press, Los Alamitos, CA, 2000).
[31] D. Achlioptas and C. Moore (unpublished).
[32] D. Achlioptas, U. Noar, and Y. Peres, extended abstract FOCS'03, pp. 362–370 (unpublished).
[33] H. Karlo and U. Zwick, *Proceedings of the 38th FOCS* (IEEE Computer Press, Miami Beach, FL, 1997), pp. 406–415.
[34] U. Schöning, Algorithmica **32**, 615 (2002).
[35] M. Alekhnovich and E. Ben-Sasson (unpublished).
[36] A. J. Parkes, Lect. Notes Comput. Sci. **2470**, 708 (2002).
[37] S. Mertens, M. Mézard, and R. Zecchina, e-print, cs.CC/0309020.
[38] G. Biroli, R. Monasson, and M. Weigt, Eur. Phys. J. B **14**, 551 (2000).
[39] G. Parisi, e-print cs.CC/0301015.
[40] A. Montanari, G. Parisi, and F. Ricci-Tersenghi, J. Phys. A **37**, 2073 (2004).
[41] A. Z. Broder, A. M. Frieze, and E. Upfal, *Proceedings of the 4th Annual ACM-SIAM Symposium on Discrete Algorithms* (SIAM Press, Austin, TX, 1993), p. 322.
[42] F. R. Kschischang, B. J. Frey, and H.-A. Loeliger, IEEE Trans. Inf. Theory **47**, 498 (2002).
[43] A. V. Lopatin and L. B. Ioffe, Phys. Rev. B **66**, 174202 (2002).
[44] S. Seitz and P. Orponen, Proceedings of the LICS'03 Workshop on Typical Case Complexity and Phase Transitions, Ottawa, Canada, 2003; *Electronic Notes in Discrete Mathematics* (Elsevier, Amsterdam, 2003), Vol. 16.
[45] Download site: www.ictp.trieste.it/~zecchina/SP