

The Phase Diagram of 1-in-3 Satisfiability Problem

Jack Raymond¹, Andrea Sportiello², Lenka Zdeborová³

¹ *NCRG; Aston University, Aston Triangle, Birmingham, B4 7EJ*

² *Università degli Studi di Milano, via Celoria 16, I-20133 Milano*

³ *CNRS; Univ. Paris-Sud, UMR 8626, Orsay CEDEX, France 91405, LPTMS*

(Dated: May 23, 2007)

We study the typical case properties of the 1-in-3 satisfiability problem, the boolean satisfaction problem where a clause is satisfied by exactly one literal, in an enlarged random ensemble parametrized by average connectivity and probability of negation of a variable in a clause. Random 1-in-3 Satisfiability and Exact 3-Cover are special cases of this ensemble. We interpolate between these cases from a region where satisfiability can be typically decided for all connectivities in polynomial time to a region where deciding satisfiability is hard, in some interval of connectivities. We derive several rigorous results in the first region, and develop the one-step-replica-symmetry-breaking cavity analysis in the second one. We discuss the prediction for the transition between the almost surely satisfiable and the almost surely unsatisfiable phase, and other structural properties of the phase diagram, in light of cavity method results.

PACS numbers: 89.20.Ff, 75.10.Nr, 05.70.Fh, 02.70.-r

I. INTRODUCTION AND MOTIVATION

Classification of the average-case computational complexity of the constraint satisfaction problems is a major task in theoretical computer science. Many problems were successfully analyzed by rigorous probabilistic methods. However, the average-case complexity remains an open question for most of the well known NP-complete problems [1, 2], for example the K -satisfiability, Vertex Coloring, and also 1-in- K satisfiability, on the commonly studied random ensembles (sparse random regular or Erdős-Rényi graphs).

In recent years heuristic methods from statistical physics [3, 4, 5] have allowed us to understand some average-case properties of large random instances [6]. The aim of these studies was not to prove average NP-completeness [7], it was rather to understand why the problems appear hard for some local algorithms, in some intervals of ensemble parametrization. These efforts culminated in designing a new polynomial algorithm, *survey propagation* [8, 9], which empirically outspeeds all the previously known heuristics. Rigorous understanding of this algorithm is, however, still missing.

The fact that lies behind this success is the intrinsic similarity of the combinatorial optimization problems to physical systems called spin glasses [10]. The organization of solutions is analogous to the structure of the free energy landscape of the physical models. Several phases can be located in the parameter space, with abrupt transitions between the different phases. An example is the SAT/UNSAT transition, i.e. the transition from the SATisfiable (SAT) phase where almost every instance is satisfiable (ground state of energy zero) to the UNSATisfiable (UNSAT) phase where almost every instance is unsatisfiable (positive ground-state energy). Another is the glassy transition where the phase space splits into many clusters and metastable states, and where many dynamical procedures (a physical dynamics or an algorithm) are unable to find the ground state. This connection between the structure of solution and the average algorithmic performance was the main motivation for detailed studies of the phase diagram of the K -satisfiability [5, 8, 9, 11], Vertex Coloring [12, 13] and many other problems.

The presented study of the phase diagram of the 1-in- K satisfiability (sometimes called *exact satisfiability* [14]) problem adds one more item to this list. But prolonging a list is not the main motivation for this work. 1-in- K SAT on the symmetric ensemble (probability that a variable in a clause is negated is 1/2) is one of the few NP-complete problem which has been proven to be on average polynomial (Easy) [15]. On the other hand for the positive ensemble (no negations, equivalent to Exact Cover) no such proof exists, nor is there a heuristic algorithm with empirically polynomial time performance in the vicinity of the SAT/UNSAT transition. However, by analogy with K -SAT and coloring, we may expect polynomial time performance in this region using survey propagation.

Our main motivation is to interpolate between the symmetric and positive ensemble to show how the phase space changes. For this reason we introduce a ϵ -1-in- K SAT problem, and study the phase diagram in parameters (γ, ϵ) , where ϵ stays for the probability that a variable in a clause is negated, and γ is the average connectivity of a variable. To our knowledge, this general ensemble is considered here for the first time. We generalize the rigorous probabilistic analysis to the general ϵ case. Then we use the replica-

symmetric and one-step-replica-symmetry-breaking cavity methods [3, 4] to understand more features of the problem in the whole space of parameters.

Our motivation is similar to the one which led to the introduction of the $(2+p)$ -SAT problem [16, 17, 18], where the instances are a mixture of 2-SAT and 3-SAT clauses on Erdős-Rényi graphs. A parameter p interpolates between the ensemble of random 3-SAT formulas, which are known to be computationally hard [8, 9] in a region including the SAT/UNSAT transition, and random 2-SAT formulas, for which an anytime polynomial algorithm exists. A statistical physics approach has been applied to study the $(2+p)$ -SAT problem, however only the replica-symmetric solution was investigated. Analogical interpolation between P and NP-complete cases of some other problems have been investigated in [19, 20].

A. The model

A *factor graph* $G = (V_v, V_c; E)$ is a bipartite graph, where the two species of vertices are called respectively *variables* $i \in V_v$ and *clauses* $a \in V_c$. It is a common graphical object used in computer science in order to encode the geometrical framework of a problem in combinatorial optimization, as it often allows us to shorten and clarify the “rules of the game”.

This is the case also for 1-in- K -SAT. Indeed, cases of both K -SAT and 1-in- K -SAT are examples of boolean satisfiability problems, and thus formally inscribed in a framework of boolean logic expressions: we deal with M boolean clauses over N variables, which should be simultaneously satisfied (i.e. evaluated to True), in order to consider the N -tuple of assignments to the variables a *solution* of the problem instance. Each clause a involves K out of the $2N$ literals $x_1, \dots, x_N, \bar{x}_1, \dots, \bar{x}_N$ (not x_i and \bar{x}_i simultaneously). While a K -SAT clause is satisfied if *at least* one of the involved literals is True, a 1-in- K -SAT clause is satisfied if *exactly* one of the involved literals is True.

For both problems, a factor graph G (whose clause-nodes a have degree K) and a function $J : E(G) \rightarrow \pm 1$ fully encode an instance: if clause a involves literal x_i or \bar{x}_i we will have an edge $(i, a) \in E(G)$, and $J_{ai} = +1$ or -1 respectively. It is customary to draw edges with $J = +1$ as solid lines, and edges with $J = -1$ as dashed lines.

If we use the common identification with “spin variables” s_i

$$\begin{aligned} x_i = \text{True} & \longleftrightarrow s_i = +1 \\ x_i = \text{False} & \longleftrightarrow s_i = -1 \end{aligned}$$

the function $E_{\{J_a\}}(s)$ corresponding to a 1-in- K -SAT clause is

$$E_{\{J_a\}}(s) = \begin{cases} \text{True} & s_{i_1} J_{ai_1} + \dots + s_{i_K} J_{ai_K} = 2 - K \\ \text{False} & \text{otherwise} \end{cases} \quad (1)$$

and we say that s is a solution of the given instance if $\bigwedge_a E_{\{J_a\}}(s) = \text{True}$.

1-in- K -SAT is polynomial in the case $K = 2$ (coinciding with 2-XOR-SAT, or 2-Coloring), while it is NP-complete for $K \geq 3$, even in the restriction to all J 's positive (unlike K -SAT).

For what concerns average-case complexity, two Erdős-Rényi-like random ensembles of instances are commonly considered. In both cases we have N variables and every possible clause is present with probability p such that the average number of clauses is $M = N\gamma/K$, and variables have Poissonian degree with average γ . Then we distinguish

Positive Poisson ensemble: The edge parameters J_{ai} are all $+1$. In this case we use a shorthand for the energy function $E_{J_a=(+,+,+)} = E_a$.

Symmetric Poisson ensemble: The edge parameters J_{ai} are random independent in $\{\pm 1\}$ with equal probability.

The positive version of the problem is the one which corresponds to Exact Cover, in the case of incidence matrices whose columns have K nonzero entries.

In this paper we study the generalization to the ensemble in which the J 's are taking value ± 1 independently, $\epsilon \in [0, 1/2]$ being the probability of having $J = -1$. We call this generalization ϵ -1-in- K SAT, in order not to confuse with the 1-in- K SAT by which is often meant only the symmetric ensemble. We describe the phase diagram of this problem in the parameters (ϵ, γ) .

B. Main results and the paper organization

Throughout this paper we present methods and results relevant to the problem of ϵ -1-in- K -SAT with $K = 3$ only. Generalisation to instances of larger clause length requires, in most cases, only small changes in methodology.


In section II and appendices A-B we derive algorithmic and probabilistic bounds, both rigorous, for the SAT/UNSAT threshold in the ϵ -1-in-3-SAT problem. The most remarkable result of those sections is that the bound is tight for $\epsilon \in [0.2726, 1/2]$, so in that interval the SAT/UNSAT threshold is rigorously known. This generalizes the result of [15] for the symmetric ensemble $\epsilon = 1/2$.

In section III we develop the Replica-Symmetric (RS) solution. First we write the replica-symmetric equations III A, then we discuss the zero-temperature limit III B. We analyze the hard-fields solution in III C, and the soft-fields solution in III D. However, as we show in III E, this solution cannot be correct (cease to be stable) above a certain connectivity not larger than the expected SAT/UNSAT transition. At this connectivity the belief propagation algorithm would fail to converge. In fact, there even exists a region in the phase diagram, where the RS solution is not stable, and yet the Short Clause Heuristics (SCH) algorithm is proven to work in on average polynomial time. To our knowledge this does not happen in any of the previously studied models, and is a point worth further investigation.

In section IV we work out the one-step-Replica-Symmetry-Breaking (1RSB) solution. In this case we assume the existence of many disconnected clusters of solutions, and many metastable states, which can actually trap most of the traditional algorithms. We write the general equations in section IV A, then we concentrate on the zero-temperature zero-energy case IV B, which leads to the survey propagation equations. The zero-temperature positive-energy case is studied in appendix C. In appendix D we check the local stability of the 1RSB solution.

The main result of the 1RSB analysis is the prediction for the SAT/UNSAT threshold, fig. 1. For $\epsilon < 0.07$ the 1RSB approach is stable around the SAT/UNSAT line, so the threshold is likely to be exact. Whereas for $0.07 < \epsilon < 0.2726$ the 1RSB result is unstable, and a more involved analysis would be required to locate exactly the SAT/UNSAT threshold (the 1RSB result is expected to be an upper bound). The presence of a nontrivial 1RSB solution in the small- ϵ region suggests the presence of a Hard-SAT region. Details of these results are discussed in section IV C.

FIG. 1: The phase diagram of ϵ -1-in-3 SAT problem, for what concerns the SAT/UNSAT transition. The parameters ϵ and γ describe the probability of negations and the average variable connectivity. For $\epsilon > 0.2726$, the threshold is rigorously $\gamma^*(\epsilon) = 1/(4\epsilon(1-\epsilon))$ (drawn as a solid line), since the unit-clause upper bound and short-clause-heuristic lower bound coincide in that region. For $\epsilon < 0.2726$, the dot-dashed, dashed and dotted line denote respectively the SCH lower bound, and the UC and first-moment-method (1MM) upper bounds. The solid line is our one-step replica-symmetry-breaking (1RSB) prediction for the SAT/UNSAT threshold. For $0 \leq \epsilon < 0.07$ the 1RSB result is stable (gray shading) and so the threshold is likely to be exact. For $0.07 < \epsilon < 0.2726$ the 1RSB result is unstable, and so the threshold is just approximate (expected to be an upper bound).



II. RIGOROUS BOUNDS ON THE SAT/UNSAT THRESHOLD

A. Unit-Clause Propagation Analysis

Unit-Clause (UC) algorithms are a class of randomized algorithms for boolean satisfiability problems, which when applied to a specific instance, seek a solution or a certificate of unsatisfiability by assigning variables to ± 1 (or “True/False”) whilst maximizing the amount of logical deductions coming from uniquely determined constraints (*unit clauses*). In the absence of immediate deductions, variables are fixed by some heuristic rule, and these free steps determine a branching process on the space of feasible configurations. Our analysis of unit-clause propagation is elucidated in appendix A, while for a more general description consider [21].

Algorithms based on unit-clause propagation have been already analysed for the problems of symmetric and positive 1-in- K SAT. For the positive ensemble ($\epsilon = 0$) the best known lower bound to the SAT/UNSAT transition is $\gamma = 1.638$ [22], and no upper bound is known from unit-clause algorithms. For the symmetric ensemble ($\epsilon = 1/2$), the method allows us to determine that the exact SAT/UNSAT transition is $\gamma = 1$ [15]. Here we extend these results to compute the upper bound $\gamma_{uc}(\epsilon)$ and the lower bound $\gamma_{sch}(\epsilon)$ for general probability of negation, which describe regions that are almost surely (a.s.) Easy SAT or Easy UNSAT for an instance of the ϵ -1-in-3-SAT, sampled from the random (ϵ, γ) ensemble (cfr. section IA).

In appendix A1 we demonstrate the upper bound $\gamma_{uc}(\epsilon)$ to the connectivity above which an Easy UNSAT phase exists. Whenever $\gamma > \gamma_{uc}(\epsilon)$, the instance is a.s. proven to be UNSAT by a randomized linear-time decimation algorithm in which one tests, for all variables i , if both fixing $s_i = +1$ or $s_i = -1$ lead to contradictions through unit-clause implications alone. This line has the analytic form $\gamma_{uc}(\epsilon) = 1/(4\epsilon(1 - \epsilon))$.

In appendix A2 we obtain the lower bound $\gamma_{sch}(\epsilon)$ to the connectivity below which an Easy SAT phase exists. Now, we perform an extensive number of free choices, and thus we should specify our heuristic rule. It turns out that, among those tested, the Short Clause Heuristics (SCH; assigning a variable in one of the shortest clauses remaining) is the one attaining the best bound on the whole interval of ϵ . If $\gamma < \gamma_{sch}(\epsilon)$, by fixing variables according to SCH one can find a solution with finite probability on any run. Restarting the procedure many times allows us to find a solution in on average linear time. This extends the idea employed for Exact Cover in [22]. At all ϵ we find lower bounds by numerical integration (see above figure), including $\gamma_{sch}(0) = 1.639$, consistent with the analysis of [22].

Finally in appendix A3 we prove analytically that on the interval $\epsilon \in [0.2726, 1/2]$ the curves $\gamma_{uc}(\epsilon)$ and $\gamma_{sch}(\epsilon)$ coincide. The result includes the symmetric ensemble, for which it was originally proven in [15]. This fact indicates that there exists a region of the phase diagram in which typical instances of the (ϵ, γ) ensemble are easily solved, except at the exactly determined SAT/UNSAT transition line.

B. Upper bounds for small ϵ

For $\epsilon \rightarrow 0$ we have $\gamma_{uc} \sim \epsilon^{-1}$ and so we would like to find a better upper bound by some different method. An improvement is obtained through the First Moment Method (1MM) on the 2-core of the graph. Restriction to the 2-core makes the bound tighter, as it reduces instance-to-instance fluctuations. This provides a line $\gamma_{1mm}(\epsilon)$, which is finite everywhere, and thus beats γ_{uc} in some interval of small ϵ (details are in appendix B1). The best known upper bound for the positive ensemble ($\epsilon = 0$) is $\gamma = 1.932$, obtained by a refinement of the first moment method [23].

Still, the first moment method is only probabilistic, and does not allow us to find a certificate in polynomial time for a given instance. Such a task is achieved at finite γ , also in the region of small ϵ and $\epsilon = 0$, through the embedding of 1-in-3-SAT into an instance of 3-XOR-SAT. While $E_a^{1-in-3}(s_1, s_2, s_3) = \text{True}$ on the three configurations $(s_1, s_2, s_3) = (+, -, -)$, $(-, +, -)$ and $(-, -, +)$, the function $E_a^{3-XOR}(s_1, s_2, s_3)$ also allows for $(s_1, s_2, s_3) = (+, +, +)$. So all the constraints are linear relations, and the problem is formally solved by Gaussian elimination. This gives an upper bound for an “Easy-UNSAT” phase, independently of ϵ at $\gamma = 3\alpha^* = 2.754$ where α^* is the SAT/UNSAT threshold (clause-to-variable ratio) in 3-XOR-SAT [24, 25] (cfr. appendix B2 for details). Note for comparison that in random 3-SAT at given finite α (however large) in the UNSAT phase, there is no polynomial algorithm which can find a.s. a certificate for a typical instance, and intuition strongly suggests that such an algorithm can not exist [26].

III. RS CAVITY APPROACH

The cavity method is developed within a statistical mechanics formulation. For this purpose we choose an integer-valued “energy function” for a single clause, $E_{\{J_a\}}(s)$, to be associated to the original boolean-valued function $E_{\{J_a\}}^{\text{bool}}(s)$ in (1); just as $E_{\{J_a\}}(s) = 0$ or 1 respectively if $E_{\{J_a\}}^{\text{bool}}(s) = \text{True}$ or False . We thus have a Hamiltonian

$$\mathcal{H}(s) = \sum_a 2E_{\{J_a\}}(s), \quad (2)$$

which counts (twice) the number of contradictions. The factor 2 is a useful convention, so that all the cavity parameters to be introduced below (*cavity fields* and *biases*) will be integer in the “zero-temperature” limit.

The introduction of the cost function above allows us to define a Gibbs weight $e^{-\beta\mathcal{H}(s)}$, where β is some parameter (*inverse temperature*), so that a single contradiction causes a dump of a factor $e^{-2\beta}$ in the measure of the configuration. As customary in statistical mechanics, one introduces a partition function

$$Z(\beta) = \sum_s e^{-\beta\mathcal{H}(s)} \quad (3)$$

and a set of observables (say, probabilities of having patterns A)

$$\text{prob}(A) := Z^{-1} \sum_{s: A \text{ happens}} e^{-\beta\mathcal{H}(s)}. \quad (4)$$

Within this framework the cavity method translates certain obvious recurrence relations for interaction structures on a factorized graph (tree) to approximate self-consistent equations for local expectation values on a graph which is only “locally tree-like” (e.g. a sparse Erdős-Rényi graph at large N , where loops are expected to arise at lengths of order $\ln N$).

The replica-symmetric (RS) assumption is used at a certain point. It consists in assuming that there is a single pure state describing the equilibrium behaviour of the ensemble. In turns, it will allow to neglect certain connected correlation functions. In this section we develop the cavity method under this hypothesis, while extensions are discussed in later sections.

We will not review in detail all the derivations of the equations, instead we just introduce, in section III A, some notations on the “easy” case of interaction on a tree, and give without proof the further formulas which are valid in the various contexts. A heuristic consideration of the complications arising on a random graph with long loops can be found in [3, 4].

A. RS cavity equations

Consider a problem defined on a factor graph G , such that, for a certain edge (i, a) , G is composed of factorized components attached to the vertices in a neighbourhood of radius 1 of (i, a) . Call $(\partial i \setminus a)$ the set of other clauses, besides a , neighbouring i , and $(\partial a \setminus i)$ the set of other variables, besides i , neighboring a . This description motivates a factor graph of the form on fig. 2 left, where “gray bubbles” stands for some other parts of the graph, and there are no paths connecting distinct bubbles, except through the explicitly drawn neighborhood of (i, a) .

For a variable $j \in (\partial a \setminus i)$, call $Z_s^{j \rightarrow a}$ the quantity $Z \cdot \text{prob}(s_j = s)$ on the system consisting of the (gray-bubble) subgraph attached to vertex j , see fig. 2 right, Z being the partition function of this subsystem. Similarly for a clause $b \in (\partial i \setminus a)$, call $Y_s^{b \rightarrow i}$ the quantity $Z \cdot \text{prob}(s_i = s)$ on the system consisting on the subgraph attached to node i through edge (b, i) , see fig. 2 right. Then we have the composition relations (say $\partial a \setminus i = \{j, k\}$ for our clause of degree 3)

$$Z_{s_i}^{i \rightarrow a} = \prod_{b \in \partial i \setminus a} Y_{s_i}^{b \rightarrow i}, \quad (5a)$$

$$Y_{s_i}^{a \rightarrow i} = \sum_{s_j, s_k} e^{-2\beta E_{\{J_a\}}(s_i, s_j, s_k)} Z_{s_j}^{j \rightarrow a} Z_{s_k}^{k \rightarrow a}. \quad (5b)$$




FIG. 2: Left: Factor graph representing the 1-in-3 SAT problem, in a neighbourhood of an edge (i, a) . Right: Definition of the cavity partition functions.

At this stage we note that, in rewriting the cavity partition sums in terms of probabilities, the *belief propagation* equations [27, 28], familiar to computer scientists, are attained.

As usual in physics we reparametrize the pairs Z_+, Z_- as different natural quantities, a *free energy* $F = 1/\beta \ln(Z_+ + Z_-)$ and a *magnetic field* $h = 1/(2\beta) \ln(Z_+/Z_-)$, the field that, if applied to the variable in substitution of the whole system, would cause the same average magnetization. We define the *cavity fields* and *cavity biases* in the following way


$$e^{2\beta h_{i \rightarrow a}} := \frac{Z_+^{i \rightarrow a}}{Z_-^{i \rightarrow a}}, \quad e^{2\beta u_{a \rightarrow i}} := \frac{Y_+^{a \rightarrow i}}{Y_-^{a \rightarrow i}}. \quad (6)$$

The recursion equations for h and u then follows from (5)

$$h_{i \rightarrow a} = \sum_{b \in \partial i \setminus a} u_{b \rightarrow i}, \quad (7a)$$

$$u_{a \rightarrow i} = \frac{1}{2\beta} \ln \frac{\sum_{s_j, s_k} \exp[\beta(h_{j \rightarrow a} s_j + h_{k \rightarrow a} s_k - 2E_{\{J_a\}}(+1, s_j, s_k))]}{\sum_{s_j, s_k} \exp[\beta(h_{j \rightarrow a} s_j + h_{k \rightarrow a} s_k - 2E_{\{J_a\}}(-1, s_j, s_k))]} . \quad (7b)$$

One can think of u 's and h 's as messages being attached to the edges of the graph, and oriented (the u 's towards the variable, the h 's towards the clause). Then, the update functions (7) are represented on the variable- and clause-nodes, as in the figure below.



Similarly, we handle the free energies. First define the accessory quantities

$$Z^i = \prod_{a \in \partial i} Y_+^{a \rightarrow i} + \prod_{a \in \partial i} Y_-^{a \rightarrow i}, \quad (8a)$$

$$Y^a = \sum_{s_i, s_j, s_k} e^{-2\beta E_{\{J_a\}}(s_i, s_j, s_k)} Z_{s_i}^{i \rightarrow a} Z_{s_j}^{j \rightarrow a} Z_{s_k}^{k \rightarrow a}. \quad (8b)$$

Then we write the free-energy shift $\Delta F^{a \cup \partial a}$ after adding a clause a and connecting all its neighbors $i \in \partial a$, and free-energy shift ΔF^i after connecting all the components incident on variable i

$$e^{-\beta \Delta F^{a \cup \partial a}} = \frac{Y^a}{\prod_{i \in \partial a} \prod_{b \in \partial i \setminus a} (Y_+^{b \rightarrow i} + Y_-^{b \rightarrow i})}, \quad (9a)$$

$$= \frac{\sum_{s_i, s_j, s_k} \exp\{\beta[h_{i \rightarrow a} s_i + h_{j \rightarrow a} s_j + h_{k \rightarrow a} s_k - 2E_{\{J_a\}}(s_i, s_j, s_k)]\}}{\prod_{i \in \partial a} \prod_{b \in \partial i \setminus a} 2 \cosh(\beta u_{b \rightarrow i})},$$

$$e^{-\beta \Delta F^i} = \frac{Z^i}{\prod_{a \in \partial i} (Y_+^{a \rightarrow i} + Y_-^{a \rightarrow i})} = \frac{2 \cosh(\beta \sum_{a \in \partial i} u_{a \rightarrow i})}{\prod_{a \in \partial i} 2 \cosh(\beta u_{a \rightarrow i})}. \quad (9b)$$

Finally, for the free energy, $F = 1/\beta \ln Z(G)$, of a tree graph one gets

$$F(\beta) = \sum_a \Delta F^{a \cup \partial a} - \sum_i (d_i - 1) \Delta F^i, \quad (10)$$

where d_i is degree of variable i . Writing this in terms of fields we note the cancellation of the factors $2 \cosh(\beta u_{a \rightarrow i})$ between the denominators of (9a) and (9b). Furthermore, in the numerator of (9b), the combination $h_i := \sum_{a \in \partial i} u_{a \rightarrow i}$ appears. This is the “total field” parameter for the magnetization of variable i in cavity approximation.

The free energy as a function of inverse temperature β on a given graph allows us to determine, by Legendre transform, the number $\exp(S(E))$ of configurations of given energy E (number of violated clauses). Both S and E are extensive, i.e. of order N , and corrections decreasing with $1/N$ are understood.

$$E(\beta) = \frac{\partial(\beta F(\beta))}{\partial \beta}; \quad S(E) = (E - F) \beta(E). \quad (11)$$

The main insight here is that we can think of the set of cavity fields as parameterization for the local “magnetizations” of variables i (i.e. probability of being $s_i = +1$, for two-state variables), in a system in which the interaction of i with a neighboring clause a has been modified (*cavity system*). If the clause a has been “switched off”, the nodes in the neighborhood of a now become well-separated on the graph which effectively describes the cost function. An assumption of decorrelation of variables, which is exactly true for variables in disconnected components, and approximatively valid for variables sufficiently far away on the graph, out of a critical temperature and *within a pure thermodynamic phase*, provides us self-consistent equations for the cavity fields. The equations are exactly the same as those we wrote for factorized graphs and hold in the leading order in the system size N , in particular eq. (10), see [3]. The cavity assumption can be self-consistently checked as we will describe in section III E.

B. The zero-temperature limit: Hard and soft fields

In the limit of zero temperature, $\beta \rightarrow \infty$, the update of cavity biases (7) simplifies significantly to

$$h_{i \rightarrow a} = \sum_{b \in \partial i \setminus a} u_{b \rightarrow i}, \quad (12a)$$

$$u_{a \rightarrow i} = \frac{1}{2} \left[\max_{s_j, s_k} (h_{j \rightarrow a} s_j + h_{k \rightarrow a} s_k - 2E_{\{J_a\}}(+1, s_j, s_k)) \right. \\ \left. - \max_{s_j, s_k} (h_{j \rightarrow a} s_j + h_{k \rightarrow a} s_k - 2E_{\{J_a\}}(-1, s_j, s_k)) \right]. \quad (12b)$$

It is immediately seen that, as $E_{\{J_a\}}(s)$ is evaluated over $\{0, 1\}$, it is self-consistent to assume that $h \in \mathbb{Z}$ and $u \in \{-1, 0, +1\}$.

In fact the only characteristic property of $E_{\{J_a\}}(s)$ we need to have is that $E_{\{J_a\}}(s) = 0$ if and only if s satisfies clause a , and that $E_{\{J_a\}}(s_1, s_2, s_3) - E_{\{J_a\}}(-s_1, s_2, s_3) \in \{-1, 0, +1\}$, (a kind of discrete “Lipschitz” condition), which clearly holds for our choice of Hamiltonian (2).

The only other choice of Hamiltonian for 1-in-3-SAT sharing this property is

$$\mathcal{H}'(s) = 2 \sum_a E'_{\{J_a\}}(s), \quad (13)$$

where $E'_a(s)$ coincides with $E_a(s)$ except that on $(+, +, +)$, where it is valued $E' = 2$ instead of 1 (because 2 flips are required in order to satisfy the clause).

The fact that $h, u \in \mathbb{Z}$ is much more than self-consistent, it is necessary, even for the “true” cavity fields (the ones that we would find from the evaluation of global partition functions, instead of the ones being solution of the cavity equations), and approximatively true in a whole region of large β (it suffices that $\beta \gg \ln N$). Let us concentrate first on $h_{i \rightarrow a}$ and say that $Z_+^{i \rightarrow a} = \sum_n g_+(n) e^{-2\beta n}$, where the integer coefficients $g_+(n)$ count the configurations with n violated clauses, in the proper cavity system labeled by $(i \rightarrow a)$. There will be a certain value n_+ corresponding to the first non-vanishing coefficient $g_+(n)$. Identical definitions are assumed for $+\leftrightarrow -$. Then we have

$$h_{i \rightarrow a} = \frac{1}{2\beta} \ln \frac{Z_+^{i \rightarrow a}}{Z_-^{i \rightarrow a}} = (n_- - n_+) + \frac{1}{2\beta} \ln \frac{g_+(n_+)}{g_-(n_-)} + \mathcal{O}(e^{-2\beta}). \quad (14)$$

So, at all orders in a purely algebraic expansion in powers of $1/\beta$, we only have two terms: a first one, $(n_- - n_+)$, the *hard field*, is constrained to be integer; and secondly the coefficient in the second term, the *soft field*, which being the logarithm of the ratio of two (potentially large in N) integers, is simply taken as a value over \mathbb{R} .

In particular the ground-state energy, $\beta \rightarrow \infty$ limit of eq. (11), can be computed using only the hard fields. On the other hand, to compute the ground-state entropy the soft fields are necessary (the general relation follows from (11), but is rather lengthy).

We will see in the next section that working only with the hard fields has huge computational advantages, however, as they do not contain all the information, we come back to the soft fields in section III D.

C. The hard-fields analysis: Warning Propagation

In this section we return to equations (12), and neglect for this moment the $1/\beta$ part of the field. The corresponding equations are called *warning propagation*, and the discrete set of possible values for the biases takes an interpretation in terms of “kinds of warnings”:

$$\begin{array}{l|l} u_{a \rightarrow i} = -1 & \text{Clause } a \text{ tells to variable } i: \text{ “I think you should be } -1\text{”} \\ u_{a \rightarrow i} = 0 & \text{Clause } a \text{ tells to variable } i: \text{ “I can deal with any value you take”} \\ u_{a \rightarrow i} = +1 & \text{Clause } a \text{ tells to variable } i: \text{ “I think you should be } +1\text{”} \end{array}$$

The analogous interpretation for fields h is

$$\begin{array}{l|l} h_{i \rightarrow a} < 0 & \text{Variable } i \text{ tells to clause } a: \text{ “I would prefer to be } -1\text{”} \\ h_{i \rightarrow a} = 0 & \text{Variable } i \text{ tells to clause } a: \text{ “I don’t have any strong preferences”} \\ h_{i \rightarrow a} > 0 & \text{Variable } i \text{ tells to clause } a: \text{ “I would prefer to be } +1\text{”} \end{array}$$

From which the prescriptions (12) on how to update the “warnings” over the graph also become intuitive.

We now determine statistics over the ensemble of random pairs $(G, \{J_a\})$. It turns out that, although fields h can take infinite values, by virtue of the Poisson ensemble the equations are closed under a finite number of parameters: We define probabilities $p_-/p_+/p_0$ that cavity fields h are negative/positive/zero, and similar probabilities $q_-/q_+/q_0$ that biases u are $-1/+1/0$. Then, the statistical average of equations (12), seen as defining a dynamics of time evolution for the distributions of the fields, gives the following dynamical map over $\mathbf{q} = (q_+, q_-)$ (the auxiliary vector $\mathbf{p} = (p_+, p_-)$ is also defined)

$$\mathbf{q} = M\mathbf{q}'; \quad \mathbf{q}' = (p_-^2, 2p_+ - 2p_+^2); \quad (15a)$$

$$\mathbf{p} = M\mathbf{p}'; \quad \mathbf{p}' = (f(\gamma q_+, \gamma q_-), f(\gamma q_-, \gamma q_+)); \quad (15b)$$

where we recall that γ is the mean degree of variable, while matrix M describes the probability of negations

$$M = \begin{pmatrix} 1 - \epsilon & \epsilon \\ \epsilon & 1 - \epsilon \end{pmatrix} \quad (16)$$


Finally the function $f(r, s)$ gives the probability that the difference of two Poissonian-distributed integers (resp. with rate r and s) is positive

$$f(r, s) := \sum_{m=0}^{\infty} \sum_{n=m+1}^{\infty} \text{Pois}_r(n) \text{Pois}_s(m); \quad \text{Pois}_\alpha(n) := e^{-\alpha} \frac{\alpha^n}{n!}. \quad (17)$$

The “paramagnetic” state $\mathbf{q} = \mathbf{0}$ is everywhere a solution of (15). It is however numerically unstable above the line $\gamma_{\text{uc}} = 1/(4\epsilon(1 - \epsilon))$ (coinciding with the unit-clause upper bound). A non-paramagnetic $\mathbf{q} \neq \mathbf{0}$ solution appears continuously above this line and is stable. Conversely, for $\epsilon < \epsilon^*$, the non-paramagnetic solution appears discontinuously at connectivity γ_{RS}^* , and is a stable local attractor.

The line $\gamma_{\text{RS}}^*(\epsilon)$, and even the “triple point” ϵ^* at which $\gamma_{\text{RS}}^*(\epsilon)$ touches γ_{uc} , depend on the choice of Hamiltonian \mathcal{H} (2) or \mathcal{H}' (13), and are plotted in fig. 3. Since the ground-state energy $E(\beta \rightarrow \infty)$ (11) is zero if and only if $\mathbf{q} = \mathbf{0}$ we conclude that the line γ_{RS}^* for $\epsilon < \epsilon^*$, and γ_{UC} for $\epsilon \geq \epsilon^*$ is the replica-symmetric prediction for the SAT/UNSAT threshold.

FIG. 3: Results of the replica-symmetric cavity analysis and its stability. The continuous curves with no bullets are the rigorous bounds left for comparison. The line with diamonds corresponds to the replica-symmetric prediction for the SAT/UNSAT threshold derived from the Hamiltonian \mathcal{H} , with $E_a(+, +, +) = 1$, while the line with triangles to the one derived from the Hamiltonian \mathcal{H}' , with $E'_a(+, +, +) = 2$. The dotted line is the soft-fields instability of the replica-symmetric solution, it is separating the stable region (below) from the unstable one (above). For $\epsilon > 0.33 \pm 0.02$ the stability line seems to coincide with the SAT/UNSAT threshold.



There are striking hints towards the badness of the replica-symmetric solution. For example, the predicted critical value $\gamma_{\text{RS}}^*(\epsilon = 0)$ is different from the numerical one, and even larger than the rigorous upper bound. But there is also an inconsistency internal to the method. It is not possible that the satisfiability line in the phase diagram depends on the finite-temperature Hamiltonian used in the cavity equations, but the lines γ_{RS}^* coming from \mathcal{H} and \mathcal{H}' are different. Another argument comes from the cavity prediction of the ground-state energy $E_{\text{min}} = E(\beta \rightarrow \infty)$ (11) which is zero if and only if $\mathbf{q} = \mathbf{0}$. The *discontinuous* appearance of a new fixed point leads to a discontinuity in γ of $E_{\text{min}}(\gamma, \epsilon)$, but this is impossible, as we have the Lipschitz condition

$$\frac{\partial}{\partial \gamma} E_{\text{min}}(\gamma, \epsilon) \in [0, \frac{1}{3}] \quad (18)$$

coming from the fact that adding randomly M' clauses to an instance whose minimum energy is E_{min} can only give an instance with minimum energy in the range $[E_{\text{min}}, E_{\text{min}} + M']$. In section III E we will explain in detail why and where exactly the replica-symmetric solution breaks down.

D. The soft-fields analysis

In the region of the phase diagram where the RS hard-fields analysis predicts zero ground-state energy (SAT region) all the hard fields and biases are zero ($g_+(0), g_-(0) > 0$ in the language of equation (14)). We denote with capital letters $U_{a \rightarrow i}$ and $H_{i \rightarrow a}$ the soft fields, i.e.

$$u_{a \rightarrow i} = 0 + \frac{U_{a \rightarrow i}}{\beta}; \quad h_{i \rightarrow a} = 0 + \frac{H_{i \rightarrow a}}{\beta}. \quad (19)$$

Their update is deduced analyzing the general cavity equations (7)

$$H_{i \rightarrow a} = \sum_{b \in \partial i \setminus a} U_{b \rightarrow i}; \quad (20)$$

$$J_{ai} U_{a \rightarrow i} = -\frac{1}{2} \ln (e^{2J_{aj} H_{j \rightarrow a}} + e^{2J_{ak} H_{k \rightarrow a}}); \quad (21)$$

Defining $q(U)$ and $p(H)$ as the probability distributions of U and H over the graph, we have the self-consistent equations

$$p(H) = (1 - \epsilon)\tilde{p}(H) + \epsilon\tilde{p}(-H); \quad (22)$$

$$\tilde{p}(H) = \sum_{k=0}^{\infty} e^{-\gamma} \frac{\gamma^k}{k!} \int \prod_{i=1}^k [dU_i q(U_i)] \delta\left(H - \sum_{i=1}^k U_i\right); \quad (23)$$

being the analogue of (15b), while for the (15a) one has

$$q(U) = (1 - \epsilon)\tilde{q}(U) + \epsilon\tilde{q}(-U) \quad (24)$$

$$\tilde{q}(U) = \int dH_j dH_k p(H_j) p(H_k) \delta\left(U + \frac{1}{2} \ln(e^{2H_j} + e^{2H_k})\right) \quad (25)$$

These equations are already beyond the possibilities of an analytical treatment, and can be solved only by a population-dynamics technique [3].

In this “paramagnetic” region the expression for the RS ground-state entropy (logarithm of the number of SAT configurations) simplifies. Thus, knowing the distributions $q(w)$ and $p(g)$ we can compute the average entropy. In appendix B1 we compute the average number of solutions $\langle \mathcal{N} \rangle$, annealed entropy $\ln(\langle \mathcal{N} \rangle)$, whereas the RS computation leads to quenched average of the entropy $\langle \ln \mathcal{N} \rangle$. The same quantity may be computed for a given large sparse graph from the message passing procedure (21). However, the result will be valid only in the region where the RS assumption is valid, see section III E.

E. The RS stability

The replica-symmetric solution will turn out to be incorrect if the assumption of having a single pure phase is proven to fail. As we know, a necessary condition for this is that fields incoming to a given node are uncorrelated. This property can be tested on the RS solution: if the spin-glass susceptibility diverges, then all the fields (and in particular the pairs of incoming ones) are strongly co-fluctuating, and the RS assumption is inconsistent. The (nonlinear) spin-glass susceptibility is defined as [29, 30]

$$\chi_{\text{SG}}(G) = \frac{1}{N} \sum_{i,j \in V(G)} \langle s_i s_j \rangle_c^2; \quad \chi_{\text{SG}} = \sum_{d=0}^{\infty} (2\gamma)^d \mathbb{E}(\langle s_0 s_d \rangle_c^2). \quad (26)$$

On the left is the definition for a fixed graph G , $\langle s_i s_j \rangle_c$ is the connected correlation function between nodes i and j . On the right we consider the average over instances, in the thermodynamic limit, where sites s_0 and s_d are at distance d . The factor $(2\gamma)^d$ stands for the average number of neighbours at distance d , when $d \ll \ln N$. Assuming that the limit for large d of the summands in (26) exists (with the limit $N \rightarrow \infty$ performed first), we relate it to the *stability parameter*:

$$\bar{\lambda} = \lim_{d \rightarrow \infty} (2\gamma) \left(\mathbb{E}(\langle s_0 s_d \rangle_c^2) \right)^{\frac{1}{d}}. \quad (27)$$


Then the series in (26) is essentially geometric, and converges if and only if $\bar{\lambda} < 1$.

Using the fluctuation-dissipation theorem we relate the correlation $\langle s_0 s_d \rangle_c$ to the variation of magnetization in s_0 , caused by an infinitesimal magnetic field in s_d . Then, one relates this quantity to cavity fields, i.e., up to a factor \mathcal{C} independent from d ,

$$\mathbb{E}(\langle s_0 s_d \rangle_c^2) = \mathcal{C} \cdot \sum_{\substack{a \in \partial d \\ b \in \partial 0}} \mathbb{E} \left[\left(\frac{\partial u_{a \rightarrow d}}{\partial u_{b \rightarrow 0}} \right)^2 \right]. \quad (28)$$

Finally, using the fact that we perform the large- N limit first, the variation above is dominated by the direct influence through the length- d path between the two nodes, and this induces a “chain” relation: if the path involves the clause and variable nodes $(a_d, d, a_{d-1}, d-1, \dots, a_0, 0)$ we have

$$\mathbb{E}(\langle s_0 s_d \rangle_c^2) = \mathcal{C} \cdot \mathbb{E} \left[\prod_{\ell=1}^d \left(\frac{\partial u_{a_\ell \rightarrow \ell}}{\partial u_{a_{\ell-1} \rightarrow \ell-1}} \right)^2 \right]. \quad (29)$$



Inside the paramagnetic phase and in the zero-temperature limit (but keeping soft fields of section III D), from (21), we have for a path like the upper one in the above figure

$$J_{ai} \frac{\partial u_{a \rightarrow i}}{\partial u_{b \rightarrow j}} = J_{ai} \frac{\partial U_{a \rightarrow i}}{\partial U_{b \rightarrow j}} = \frac{-J_{aj} e^{2J_{aj} H_{j \rightarrow a}}}{e^{2J_{aj} H_{j \rightarrow a}} + e^{2J_{ak} H_{k \rightarrow a}}} . \quad (30)$$

Instead of directly computing the stability parameter $\bar{\lambda}$, it is equivalent, and numerically easier, to associate every bias $u_{a \rightarrow i}$ in the population-dynamics algorithm with a positive number $v_{a \rightarrow i}$. We update this number together with the fields according to

$$v_{a \rightarrow i} = \sum_{b \in \partial j \setminus a} \left(\frac{\partial u_{a \rightarrow i}}{\partial u_{b \rightarrow j}} \right)^2 v_{b \rightarrow j} + \sum_{c \in \partial k \setminus a} \left(\frac{\partial u_{a \rightarrow i}}{\partial u_{c \rightarrow k}} \right)^2 v_{c \rightarrow k} , \quad (31)$$

(we recall that, when performing population-dynamics technique, the labels “ $a \rightarrow i$ ” do not have any spatial meaning: the population is just a collection, and messages are randomly combined at each step). After equilibration, the numbers v will change on average geometrically, with a factor $\bar{\lambda}$.

Numerically, we see that for a given ϵ the stability parameter grows with connectivity γ . On fig. 3 we see the line above which the replica-symmetric solution is unstable (i.e. $\bar{\lambda} > 1$). This line coincides with the unit-clause upper bound within the errors from about $\epsilon > 0.33 \pm 0.02$. In particular all the region where the RS results are contradictory is unstable. It is furthermore remarkable (and unexpected) that there exist a region in which RS solution is not stable, and yet the short clause heuristics a.s. finds a solution in polynomial time.

IV. 1RSB CAVITY APPROACH AND ITS IMPLICATIONS FOR THE PHASE DIAGRAM

The understanding of the role of ergodicity in the validity of the replica-symmetric cavity assumption allows us to recast the cavity method as a more powerful tool, for the case in which there are (exponentially) many phases. The assumptions underlying this process go under the jargon term “1RSB” type of symmetry breaking [3, 4]. In the 1RSB approach we assume that exponentially many pure thermodynamical states (phases) exists, and that the neighbors of a node in absence of this node are uncorrelated only within each of these states. This happens because the cluster property (the small correlation of observables far from each other) holds only within a pure phase. The name replica symmetry breaking is due to historical reasons, since the mechanism was first proposed by Parisi [31], while using the “replica trick” in analysing a spin-glass model.

The necessary but conceptually impossible handling of the multiplicity of pure phases may be replaced by a “survey” over these phases. The only memory of the original structure is through the free energies F_α of the various phases α . As the phases have to be weighted with a Boltzmann weight in F_α , a *reweighting term* has to be introduced in the “survey” equations, as we show in section IV A. In the zero-temperature limit the analysis leads to what are now called *survey propagation equations*, which are developed fully in section IV B.

The solution to these equations is described in section IV C 1, in section IV C 2 are results of the stability analysis of this solution. Checking the validity of the 1RSB cavity assumption (1RSB stability analysis) gives us a hint if this solution could be the final correct one. This has been done for the K -SAT [11, 32] and Coloring problems [12] on random graphs. The results in those cases supported strongly the conclusion that the SAT/UNSAT thresholds computed with the 1RSB cavity method were exact. The 1RSB stability analysis is technically involved, we summarize the main steps in appendix D.

A. General 1RSB cavity equations

We define a complexity function $\Sigma(F)$ as the logarithm of number of states, and hence it is computable by Legendre transform

$$-\beta m \Phi(m, \beta) = -\beta m F(\beta) + \Sigma(F); \quad \frac{\partial \Sigma(F)}{\partial F} = \beta m; \quad (32)$$

where the parameter m plays role of a second temperature, for free energies of states instead of energies of configurations, and is called the Parisi parameter of the replica symmetry breaking. The function $\Phi(m, \beta)$ is called the “replicated free energy”.

Instead of one field and one bias on every edge, now we need to keep one field and one bias for every edge *and* every state. Or equivalently, as we assumed that there is a huge number of states, a distribution of fields and biases on every edge. The self-consistent equation for this distribution is [3, 4]

$$\begin{aligned} \mathcal{P}^{a \rightarrow i}(u_{a \rightarrow i}) &= Z_{a \rightarrow i}^{-1} \int \prod_{b \in \partial j \setminus a} du_{b \rightarrow j} \mathcal{P}^{b \rightarrow j}(u_{b \rightarrow j}) \int \prod_{b \in \partial k \setminus a} du_{b \rightarrow k} \mathcal{P}^{b \rightarrow k}(u_{b \rightarrow k}) \\ &\cdot \delta(u_{a \rightarrow i} - \mathcal{F}(\{u_{b \rightarrow j}\}, \{u_{b \rightarrow k}\})) \cdot \exp(-\beta m \Delta F^{a \rightarrow i}). \end{aligned} \quad (33)$$

The function \mathcal{F} is the single-phase update of biases given by equations (7), and the last term is the reweighting of states, where $\Delta F^{a \rightarrow i}$ is the free-energy shift after adding clause a and all its neighbors except i . Referring to our calculations in equation (9a), this free-energy shift is

$$\begin{aligned} e^{-\beta \Delta F^{a \rightarrow i}} &= \frac{Y_{+1}^{a \rightarrow i} + Y_{-1}^{a \rightarrow i}}{\prod_{j \in \partial a \setminus i} \prod_{b \in \partial j \setminus a} (Y_{+1}^{b \rightarrow j} + Y_{-1}^{b \rightarrow j})} \\ &= \frac{\sum_{s_i, s_j, s_k} \exp[\beta(h_{j \rightarrow a} s_j + h_{k \rightarrow a} s_k - E_{\{J_a\}}(s_i, s_j, s_k))]}{\prod_{j \in \partial a \setminus i} \prod_{b \in \partial j \setminus a} (2 \cosh \beta u_{b \rightarrow j})}. \end{aligned} \quad (34)$$

Then, in analogy with (10), the replicated free energy Φ is calculated as

$$\Phi(m, \beta) = \sum_a \Delta \Phi^{a \cup \partial a} - \sum_i (d_i - 1) \Delta \Phi^i, \quad (35)$$

where d_i is the degree of the node i . The replicated free-energy shifts are

$$e^{-\beta m \Delta \Phi^{a \cup \partial a}} = \int e^{-\beta m \Delta F^{a \cup \partial a}}; \quad e^{-\beta m \Delta \Phi^i} = \int e^{-\beta m \Delta F^i}. \quad (36)$$

The integrals are making an average over the distributions of the fields incoming to the cavity, similarly as in equation (33), i.e. $\int f(u_1, \dots, u_k) = \int du_1 \mathcal{P}^{(1)}(u_1) \dots \int du_k \mathcal{P}^{(k)}(u_k) f(u_1, \dots, u_k)$.

Since we are interested mainly in the ground-state properties of the ϵ -1-in-3 SAT problem, we need to take the zero-temperature limit. There are two standard ways of doing this

The energetic $T \rightarrow 0$ limit [4]: We take the limit $\beta \rightarrow \infty$, $m \rightarrow 0$ with $y = \beta m$ fixed and finite. Then

$$-y \Phi(y) = -yE + \Sigma(E). \quad (37)$$

Here we neglected the entropic contribution and we can obtain complexity as a function of energy. This is the 1RSB analog of the RS analysis with hard fields (warnings).

The entropic $T \rightarrow 0$ limit [33]: We take the limit $\beta \rightarrow \infty$ at energy fixed to zero, $E = 0$. Then

$$m \tilde{\Phi}(m) = mS + \Sigma(S); \quad (38)$$

where $-\beta \Phi(\beta, m) \rightarrow \tilde{\Phi}(m)$. Here we are fixed to zero energy, but on the other hand we can compute complexity (number of states) as a function of the state internal entropy. This is the 1RSB analog of the RS analysis with soft fields (beliefs).

In this paper we work out only the simpler energetic limit, the same as in [9] for K -SAT, or in [13] for Coloring. We will see how this analysis alone already gives us a large amount of information about the phase diagram.

The reweighting (34) becomes in the zero-temperature energetic limit

$$\Delta E^{a \rightarrow i} = - \max_{s_i, s_j, s_k} (h_{j \rightarrow a} s_j + h_{k \rightarrow a} s_k - E_{\{J_a\}}(s_i, s_j, s_k)) + \sum_{b \in \partial j \setminus a} |u_{b \rightarrow j}| + \sum_{b \in \partial k \setminus a} |u_{b \rightarrow k}|. \quad (39)$$

Since the fields h and biases u are integers, by relations (12), $\Delta E^{a \rightarrow i}$ also takes nonnegative integer values. In fact it counts the number of contradictions in one warning-propagation update.

We want to determine whether a typical ϵ -1-in-3 SAT instance has any satisfying configuration. We do this in section IV B, by taking the $y \rightarrow \infty$ limit. The reweighting term $\exp(-y \Delta E^{a \rightarrow i})$ then guarantees that we keep only the cases without contradictions, $\Delta E^{a \rightarrow i} = 0$. Conversely, in order to compute the ground-state energy in the UNSAT region, or the complexity at energies higher than zero, we need to keep y finite. We undertake this in appendix C.

B. The zero-energy case, survey propagation

In the limit $y \rightarrow \infty$ we fix the energy to be zero. The energy shift $\Delta E^{a \rightarrow i}$ is zero if and only if:

- The $\{u_{b \rightarrow j}\}_{b \in \partial j \setminus a}$ are all nonnegative or all nonpositive,
- The $\{u_{b \rightarrow k}\}_{b \in \partial k \setminus a}$ are all nonnegative or all nonpositive,
- The terms $J_{aj} \sum_{b \in \partial j \setminus a} u_{b \rightarrow j}$ and $J_{ak} \sum_{b \in \partial k \setminus a} u_{b \rightarrow k}$ are not both positive.


We begin by simplifying the form of eq. (33). In the zero-temperature energetic limit we can write the distribution of fields and biases over states on every edge as

$$\mathcal{P}^{a \rightarrow i}(u_{a \rightarrow i}) = q_-^{a \rightarrow i} \delta(u_{a \rightarrow i} + 1) + q_+^{a \rightarrow i} \delta(u_{a \rightarrow i} - 1) + q_0^{a \rightarrow i} \delta(u_{a \rightarrow i}); \quad (40)$$

$$\mathcal{P}^{i \rightarrow a}(h_{i \rightarrow a}) = p_-^{i \rightarrow a} \mu_-(h_{i \rightarrow a}) + p_+^{i \rightarrow a} \mu_+(h_{i \rightarrow a}) + p_0^{i \rightarrow a} \delta(h_{i \rightarrow a}); \quad (41)$$

where $q_-^{a \rightarrow i} + q_+^{a \rightarrow i} + q_0^{a \rightarrow i} = p_-^{i \rightarrow a} + p_+^{i \rightarrow a} + p_0^{i \rightarrow a} = 1$, and $\mu_{\pm}(h)$ are normalized measures with support over \mathbb{Z}^{\pm} .

So, to every oriented edge we associate a triple of numbers $q = (q_-, q_0, q_+)$ or a triple $p = (p_-, p_0, p_+)$ (resp. if oriented towards the variable or the clause). In analogy with the self-consistent equations (7) for fields and biases we can write self-consistent equations for probabilities (surveys) q and p .



Considering only the combinations with $\Delta E^{a \rightarrow i} = 0$, the surveys of fields are given by incoming surveys of biases as

$$p_+^{i \rightarrow a} + p_0^{i \rightarrow a} = \mathcal{N}_{i \rightarrow a}^{-1} \prod_{b \in \partial i \setminus a} (q_+^{a \rightarrow i} + q_0^{a \rightarrow i}); \quad (42a)$$

$$p_-^{i \rightarrow a} + p_0^{i \rightarrow a} = \mathcal{N}_{i \rightarrow a}^{-1} \prod_{b \in \partial i \setminus a} (q_-^{b \rightarrow i} + q_0^{b \rightarrow i}); \quad (42b)$$

$$p_0^{i \rightarrow a} = \mathcal{N}_{i \rightarrow a}^{-1} \prod_{b \in \partial i \setminus a} q_0^{b \rightarrow i}; \quad (42c)$$

where $\mathcal{N}_{i \rightarrow a}$ is the normalization factor (in the update, we have three equations for three independent unknown, p_{\pm} and \mathcal{N}). And the surveys of biases are given by the incoming surveys of fields

$$q_{J_{ai}}^{a \rightarrow i} = \mathcal{N}_{a \rightarrow i}^{-1} p_{-J_{aj}}^{j \rightarrow a} p_{-J_{ak}}^{k \rightarrow a}; \quad (43a)$$

$$q_{-J_{ai}}^{a \rightarrow i} = \mathcal{N}_{a \rightarrow i}^{-1} (p_{J_{aj}}^{j \rightarrow a} (1 - p_{J_{ak}}^{k \rightarrow a}) + (1 - p_{J_{aj}}^{j \rightarrow a}) p_{J_{ak}}^{k \rightarrow a}); \quad (43b)$$

$$q_0^{a \rightarrow i} = \mathcal{N}_{a \rightarrow i}^{-1} (p_{-J_{aj}}^{j \rightarrow a} p_0^{k \rightarrow a} + p_0^{j \rightarrow a} p_{-J_{ak}}^{k \rightarrow a} + p_0^{j \rightarrow a} p_0^{k \rightarrow a}); \quad (43c)$$

where $\mathcal{N}_{a \rightarrow i}$ is the normalization factor, and the lower indexes of q 's and p 's are multiplied by -1 when variable is negated in the clause.

These equations describe survey propagation [9], and can be used inside an algorithm to find a solution to a typical instance of ϵ -1-in-3 SAT, hopefully also in a region of parameters (ϵ, γ) where short-clause-like heuristics or belief-propagation methods fail. They might also be used to compute quantities averaged over instances in our random ensemble, particularly the complexity function which determines the SAT/UNSAT transition.

In the zero-temperature limit, the replicated free energy (35) becomes

$$-y\Phi(y) = \sum_a \ln \left(\int e^{-y\Delta E^{a \cup \partial a}} \right) - \sum_i (d_i - 1) \ln \left(\int e^{-y\Delta E^i} \right), \quad (44)$$

where from (9) and (36) we get the energy shifts

$$\Delta E^{a \cup \partial a} = - \max_{s_i, s_j, s_k} (h_{i \rightarrow a} s_i + h_{j \rightarrow a} s_j + h_{k \rightarrow a} s_k - E_{\{J_a\}}(s_i, s_j, s_k)) + \sum_{i \in \partial a} \sum_{b \in \partial i \setminus a} |u_{b \rightarrow i}|; \quad (45)$$

$$\Delta E^i = - \left| \sum_{a \in \partial i} u_{a \rightarrow i} \right| + \sum_{a \in \partial i} |u_{a \rightarrow i}|; \quad (46)$$

again both these energy shifts are non-negative integers.

Furthermore in the $y \rightarrow \infty$ limit, we distinguish only if $\Delta E = 0$, then $\exp(-y\Delta E) = 1$, or if $\Delta E > 0$, then $\exp(-y\Delta E) = 0$. From eq. (44) we get for the complexity at zero energy

$$\Sigma(E = 0) = \sum_a \ln (\text{prob}(\Delta E^{a \cup \partial a} = 0)) - \sum_i (d_i - 1) \ln (\text{prob}(\Delta E^i = 0)), \quad (47)$$

where, calling $\mathcal{P}_0^i := \prod_{a \in \partial i} q_0^{a \rightarrow i}$ and $\mathcal{P}_{\pm}^i := \prod_{a \in \partial i} (q_{\pm}^{a \rightarrow i} + q_0^{a \rightarrow i})$,

$$\text{prob}(\Delta E^i = 0) = \mathcal{P}_+^i + \mathcal{P}_-^i - \mathcal{P}_0^i; \quad (48a)$$

$$\begin{aligned} \text{prob}(\Delta E^{a \cup \partial a} = 0) &= \prod_{i \in \partial a} (\mathcal{P}_{J_{ai}}^{i \rightarrow a} + \mathcal{P}_{-J_{ai}}^{i \rightarrow a} - \mathcal{P}_0^{i \rightarrow a}) - \prod_{i \in \partial a} (\mathcal{P}_{-J_{ai}}^{i \rightarrow a} - \mathcal{P}_0^{i \rightarrow a}) - \prod_{i \in \partial a} (\mathcal{P}_{J_{ai}}^{i \rightarrow a} - \mathcal{P}_0^{i \rightarrow a}) \\ &\quad - \sum_{i \in \partial a} \mathcal{P}_{-J_{ai}}^{i \rightarrow a} \prod_{j \in \partial a \setminus i} (\mathcal{P}_{J_{aj}}^{j \rightarrow a} - \mathcal{P}_0^{j \rightarrow a}). \end{aligned} \quad (48b)$$

The second equation collects the contributions from all combinations of arriving fields except the ‘‘contradictory’’ ones $(+, +, +)$, $(-, -, -)$, $(+, +, 0)$ and $(+, +, -)$ (plus permutations of the latters). Note at this point that all these equations (42)-(48) correctly do not depend on the choice of Hamiltonian (2) or (13).

C. 1RSB results for the phase diagram

In this section we give results of 1RSB cavity analysis for the ϵ -1-in-3 SAT problem. In the first two subsections we concentrate on the positive 1-in-3 SAT ($\epsilon = 0$). In the third one IV C 3 we show results for general probability of negation.

1. Complexity as a function of connectivity

To compute numerically the average value of complexity from eq. (47) we first need to find a fixed point of the survey propagation equations (42) and (43). We do that using the population-dynamics algorithm [3]. The result is in fig. 4.




FIG. 4: Average complexity density (logarithm of number of states divided by size of the graph) as a function of mean degree γ for the positive 1-in-3 SAT problem. At $\gamma_{\text{sp}} = 1.822$ a nontrivial solution of survey propagation equations appears, with positive complexity. At $\gamma = \gamma^* = 1.8789 \pm 0.0002$ the complexity becomes negative: this is the SAT/UNSAT transition. At $\gamma_p = 1.992$ the solution at zero energy ceases to exist. The inset magnifies the region where the complexity crosses zero, together with the error bar for the SAT/UNSAT transition.

Below mean degree $\gamma_{\text{sp}} = 1.822 \pm 0.001$ the survey propagation equations (42, 43) have only the trivial paramagnetic solution, with $p_0 = q_0 = 1$ and $q_{\pm} = p_{\pm} = 0$ for all edges. At γ_{sp} a solution of survey propagation equations with positive complexity appears discontinuously. The emergence of this transition far below the numerically known SAT/UNSAT threshold suggests that, in a whole interval of parameters near to the threshold the phase space restricted to solutions is clustered into many pure states, a *Hard-SAT phase* [9] exists. Furthermore, in that interval, there are also many metastable states, entropically relevant, and local minima with positive energy cost separated by macroscopic barriers. This means that local algorithms, like decimation heuristics or variants of annealed Monte Carlo, get trapped and are unable to find any ground state in polynomial time. Nonetheless, a decimation procedure based on the stationary distribution of survey propagation equations is expected to work beyond this threshold.

Note that γ_{sp} is referred to as a “dynamical threshold” in [9, 13], we stress that it is not this point which is connected to a real dynamical transition [34]. Neither is it the point where the local algorithm ceases to work in polynomial time. We come back to comment about this point in the discussion, section V.

At mean degree $\gamma^* = 1.8789 \pm 0.0002$ the complexity becomes negative. Instead of having a.s. in each instance an exponential number of clusters which contain at least one solution, the fraction of instances having any cluster which contains at least one solution (i.e. the fraction of SAT instances) becomes exponentially small. So this point identifies the SAT/UNSAT transition.

We are aware of two works where results from numerical simulation for this SAT/UNSAT threshold are given. In [22] they conclude the value of threshold is $\gamma^* = 1.86 \pm 0.03$. In [23] they give $\gamma^* = 1.875 \pm 0.015$. In fact, the latter do not give an error bar, so we guessed it from their fig. 4. Our result agrees with these estimations, and as it is based on an analytical method we reduce the error bar by one order of magnitude with a very small numerical effort.

At mean degree $\gamma_p = 1.992 \pm 0.001$ the solution at zero energy ceases to exist. In the $y \rightarrow \infty$ limit the population dynamics converges to a solution which shows a finite fraction of surveys of type $(p_{\pm}, p_0, p_{\mp}) = (0, 0, 1)$. Then, with finite probability we would find two such surveys creating a contradiction, the normalization in (43) then would be zero. We call this situation a “hard contradiction”.

Note that such a phenomenon does not occur in K -SAT or Coloring problems. Cavity equations deal with the messages incoming to a clause from all neighbours but one. In both K -SAT and Coloring (and in many other problems, like NAE- K -SAT, Vertex Covering, and so on), there is no way of making a clause unsatisfied if one of the neighbouring variables is not restricted, and indeed a γ_p threshold has never appeared in the cavity analysis for these systems.

In order to obtain a non-singular solution above connectivity γ_p , we need to work with the equations at finite y , which is able to account for the positive energy contributions. The results of the finite y analysis are shown in appendix C 1.

2. The stability analysis




FIG. 5: The stability parameter of the second kind $\ln(\mu_{II}(d))$ (D3) for positive 1-in-3 SAT for different connectivities as a function of length of the chain d . When the slope is negative, the 1RSB at this connectivity is stable against bug proliferation, and vice versa. This happens for $\gamma > \gamma_{II} = 1.838 \pm 0.002$. Right: The stability parameter of the first kind μ_I (D9) as a function of connectivity for positive 1-in-3 SAT. The stability parameter is smaller than 1 for $\gamma < \gamma_I = 1.948 \pm 0.002$, for these connectivities the 1RSB equations are stable against noise propagation.

In appendix D we introduce two stability parameters μ_I (D9) and μ_{II} (D3). Their meaning is analogous to that of the replica-symmetric stability parameter $\bar{\lambda}$ (27). The 1RSB solution is stable if and only if both $\mu_I < 1$ and $\mu_{II} < 1$.

The results for stability parameter of the second kind μ_{II} (D3) for positive 1-in-3 SAT are shown in fig. 5 (left), for finite d . Extrapolation to $d \rightarrow \infty$ is done by linear fit, which looks reasonable from the data points. So our criterion is that, if the slope in the logarithmic plot is positive, the limit value μ_{II} is larger than 1, and vice versa. We estimate that 1RSB is “type-II” stable for $\gamma > \gamma_{II} = 1.838 \pm 0.002$.

More directly, for the stability parameter of the first kind μ_I (D9), the results are shown in fig. 5, right. So we get that 1RSB is “type-I” stable for $\gamma < \gamma_I = 1.948 \pm 0.002$ (the generous error estimate is due to potential biases caused by finite population sizes).

The 1RSB solution may be correct only if both the stability parameters are smaller than one. For positive 1-in-3 SAT this happens for connectivities in the range $1.838 \simeq \gamma_{II} < \gamma < \gamma_I \simeq 1.948$. So, in particular, the SAT/UNSAT threshold γ^* is in the range of stability, and its value is to be considered exact. Note that such a situation, in which the SAT/UNSAT threshold falls into a narrow stability region, is quite common, and it has been seen also in K -SAT [11] and Coloring [12].

3. 1RSB results for general probability of negation

We applied the techniques of section IV C to the problem at finite ϵ . As expected, all the critical connectivities describe curves which are continuous at $\epsilon = 0$. We thus show, in figure 6 (left), the curves $\gamma^*(\epsilon)$, $\gamma_{sp}(\epsilon)$ and $\gamma_p(\epsilon)$, and, in the magnification on the right of fig. 6, also $\gamma_I(\epsilon)$ and $\gamma_{II}(\epsilon)$.

As ϵ approaches about 0.20, the interesting interval $\gamma_{sp} < \gamma < \gamma_p$ becomes very narrow (fig. 6, right), and the complexity value very small, $\Sigma \approx 10^{-5}$, three orders of magnitude smaller than the analogous values for $\epsilon = 0$. Above $\epsilon = 0.20$ we do not have sufficient numerical resolution to examine this region at all.

In fig. 6 (right) we plot the four curves $\gamma_{sp}(\epsilon)$, $\gamma_p(\epsilon)$, $\gamma_I(\epsilon)$ and $\gamma_{II}(\epsilon)$, shifted by the curve $\gamma^*(\epsilon)$, which is used as a reference. This allows us to appreciate that the differences $(\gamma^* - \gamma_{sp})(\epsilon)$ and $(\gamma_p - \gamma^*)(\epsilon)$ seem to vanish linearly at about $\epsilon_p = 0.21 \pm 0.01$, these two linear fits extrapolating to the same value of ϵ with reasonable confidence.

Above ϵ_p , as soon as a nontrivial solution of (energetic finite- y) 1RSB cavity equations exists, it has immediately a complexity $\Sigma(E)$ of the qualitative shape for the connectivities above γ_p , (for example see fig. 8 in appendix C). We are thus led to conclude that in this interval the line $\gamma_p(\epsilon)$ should be taken as the 1RSB prediction for the SAT/UNSAT line $\gamma^*(\epsilon)$. At about $\epsilon = 0.26 \pm 0.01$ the curve $\gamma_p(\epsilon)$ joins the unit-clause upper bound.

We should add that, above $\epsilon \simeq 0.07$, both stability criteria fail along the curve $\gamma^*(\epsilon)$ (and then, along $\gamma_p(\epsilon)$, for $\epsilon > \epsilon_p$), so that the 1RSB prediction for the SAT/UNSAT transition is not expected to be exact, but only an upper bound, in the range $0.07 < \epsilon < 0.2726$ [35, 36].

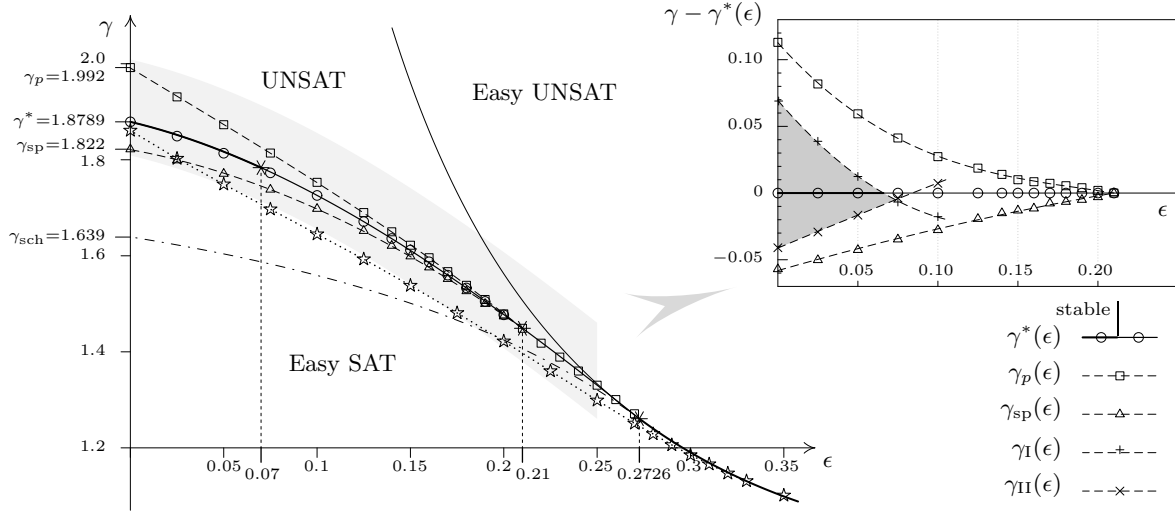


FIG. 6: Left: plot of the three curves $\gamma^*(\epsilon)$, $\gamma_{sp}(\epsilon)$ and $\gamma_p(\epsilon)$ described in the text. We left for comparison the SCH lower bound (dot-dashed curve) and the RS instability line (dotted curve with stars data points). Right: The same data, with connectivity plotted with respect to the SAT/UNSAT threshold prediction $\gamma^*(\epsilon)$. Also the stability lines $\gamma_I(\epsilon)$ and $\gamma_{II}(\epsilon)$ are shown, and the interval of stability for the SAT/UNSAT curve is $\epsilon \in [0, 0.07 \pm 0.01]$. In the inset, all the error bars are approximately as big as the point size.

V. DISCUSSION AND CONCLUSIONS

We studied the average-case behaviour of 1-in-3 SAT in the random ϵ -1-in-3 SAT ensemble, where $\epsilon \in [0, 1/2]$ is the probability of negation. This generalizes the random (symmetric) 1-in-3 SAT problem ($\epsilon = 1/2$) and random positive 1-in-3 SAT problem ($\epsilon = 0$), which is a special ensemble of Exact Cover.

Our main result is the phase diagram in fig. 1, and, magnified, in fig. 6 above. It fills the conceptual gap between the symmetric problem, of polynomial average-case complexity both in the SAT and UNSAT regions, and the positive problem, which shows a Hard-complexity phase around the SAT/UNSAT threshold.

Concerning the SAT/UNSAT transition curve, $\gamma^*(\epsilon)$, we computed upper bounds coming from unit-clause technique (UC), and from first moment method with restriction to the 2-core (1MM), and the lower bound coming from short clause heuristics (SCH). The UC and SCH bounds have been proven to coincide on the interval $\epsilon \in [0.2726, 1/2]$, and thus determine the corresponding portion of the SAT/UNSAT line rigorously.

All the other results are obtained with the non-rigorous cavity method. The results of the replica-symmetric calculations do not give us a better result for the SAT/UNSAT threshold than the UC and SCH bounds, since its prediction have to be rejected above the RS stability line γ_{RS} , fig. 3.

It is remarkable that a region of the phase diagram exists, where the replica symmetry is broken, while the short clause heuristics is proven to succeed a.s. in polynomial time. For what we know, such a feature has not been proven in any of the previously studied models, while it has been often observed empirically that some local algorithms, e.g. the Walk-SAT [37, 38, 39, 40], works in linear time inside a phase with replica-symmetry breaking. We are tempted to say that this result actually proves that the onset of a non-trivial replica symmetry broken solution does not imply to the onset of computational hardness (unfortunately the cavity results are not rigorous and the term “computational hardness” would have to be defined properly to be allowed to speak about a proof). However, we hope that this could be used to study in a new way the nature of the replica-symmetry-broken phase. On the other hand, and as claimed before, we are persuaded that a stable 1RSB solution suggest an existence of a Hard-SAT phase near to the SAT/UNSAT transition (nearer than the stability threshold). For quantitative study of this point for

the coloring problem see [41]. The analysis of [41] should be repeated for 1-in- K SAT problem in future works.

Our main insights for the region with $\epsilon < 0.2726$ comes from the one-step-Replica-Symmetry-Breaking (1RSB) calculations by analysis of the energetic zero-temperature limit of the 1RSB cavity equations (33), in which we keep only the weights of the hard fields instead of whole probability distribution.

For $\epsilon < 0.21$ we can locate, on the curves for the zero-energy complexity function $\Sigma(\gamma)$ fig. 4, the connectivities $\gamma^*(\epsilon)$ at which Σ vanishes, this corresponding to the SAT/UNSAT transition. The same computation shows the existence of a nontrivial solution above $\gamma_{\text{sp}}(\epsilon) < \gamma^*(\epsilon)$, thus predicting a whole interval of Hard-SAT phase with many pure states. However for exact location of the “dynamical transition” we would need to keep the information about the soft fields and compute when a nontrivial solution of eq. (33) appears for the entropically dominating clusters, see [34]. Note here also that the inequality $\gamma_{\text{sp}} < \gamma_{\text{RS}}$ for small ϵ is due to the discontinuity of the transition towards nontrivial 1RSB phase, that can not be seen by the local RS stability analysis. Analysis of [34] would also show that a nontrivial 1RSB solution exists everywhere in the region of $\gamma > \gamma_{\text{RS}}$. This analysis for 1-in- K SAT might be a direction for future work.

Above the line $\gamma_p(\epsilon)$, fig. 6, the system shows a transition to a phase where the 1RSB solution at zero energy ceases to exist, while it still exists above some value $E_{\text{min}}(\gamma, \epsilon)$. This is due to the presence of hard contradictions, a phenomenon specific of strongly constrained problems, like 1-in- K SAT, and to our knowledge it is a newly observed fact. Interpretation of this transition may be that the SAT formulas start to be subexponentially rare at the connectivity γ_p .

For $\epsilon > 0.21$ the nontrivial 1RSB solution at zero energy never exists, and we can trace only the curve γ_p . This would be the 1RSB prediction for the SAT/UNSAT transition. This result suggests that for $\epsilon > 0.21$ a Hard-SAT region might actually be absent. Specifying what sort of replica-symmetry-broken solution is connected with the breakdown of local algorithms, like decimation heuristics or variants of annealed Monte Carlo, is an important direction of future research.

We have checked the local stability of 1RSB solution towards 2RSB. The result is that 1RSB is stable only in a small region between the lines γ_{I} and γ_{II} in fig. 6. This means, among other things, that the 1RSB location of the SAT/UNSAT transition for $\epsilon < 0.07$ is likely to be exact. In particular, this is true for the positive 1-in-3 SAT threshold, $\gamma^* = 1.8789$. For $0.07 < \epsilon < 0.2726$ the 1RSB result is unstable so the exact location of the SAT/UNSAT transition in that region remains an open question. We can only conjecture that our 1RSB result is an upper bound, in analogy with proofs for other models in [35, 36].

Furthermore, it would be interesting to compare our results with the behaviour of the structurally affine $(2+p)$ -SAT problem [16, 17] for which the 1RSB analysis is still missing.

Finally, as we mentioned in several places above, the 1RSB cavity approach allows for algorithmic implementations. We started studying this aspect also together with Elitza Maneva and Talya Meltzer, and the results will be published elsewhere.

Acknowledgments

We thank the organizers of the “Complex Systems” July ’06 Les Houches Summer School (Session 85), where this work has been started. This work has been supported in part by the EU through the network MTR 2002-00319 “STIPCO” and the FP6 IST consortium “EVERGROW”. A.S. also thanks the LPTMS of Orsay – Paris Sud for support and kind hospitality in large stages of the work.

We thank also Elitza Maneva and Talya Meltzer for collaboration in some still unpublished algorithmic parts of this work. We appreciated a lot discussions with Florent Krzakala, Rémi Monasson and last but not least we thank Marc Mézard for his precious hints.

APPENDIX A: UPPER AND LOWER BOUNDS FROM UNIT-CLAUSE PROPAGATION AND DECIMATION HEURISTICS

Consider an instance drawn from an appropriate ensemble, and subject to a decimation algorithm whereby in each time set a variable is set to ± 1 . Call X the discrete decimation time (number of variables set, among the N total) and $C_i(X)$ the number of clauses remaining of length $i \geq 2$. Thus the initial conditions for the instance are $X = 0$ and $C_i(X) = N \frac{2}{3} \delta_{i,3}$. Assume for now these quantities are sufficient to describe the instance in the absence of clauses of “length 1” (*unit clauses*). If we assume a variable is fixed (decimated) from such an instance, the remaining instance involves a smaller number of literals,

so it is simplified in some respect. More importantly, some of the clauses are shortened, and may even be reduced to unit clauses. The unit clauses, being only 1 literal, allows no ambiguity in the values their variables must take in order for the instance to be SAT. The initial fixing of one variable by this process, forces the value of some other variables, which may again propagate, i.e. a branching process. So, a single binary choice could decrease the number of variables by a considerable amount, as a result of a cascade of these unit-clause implications.

The justification in considering the instance at all times described by $\{C_i(X)\}$ and X is the following. For sufficiently simple decimation rules, the distribution of remaining variables within clauses will be uniform and random at all X , and if N is sufficiently large the values of $C_i(X)$ are self-averaging. Furthermore, at fixed clause length, the fraction of clauses with a given number of negations is the one expected from an independent Bernoulli process: among clauses of length i , at all times there is a fraction $\binom{i}{h}\epsilon^h(1-\epsilon)^{i-h}$ of clauses with h negations. All these elements are necessary to allow a sufficiently concise dynamical description to make the progress in the following sections. Among the various possible heuristics – which determine the values set in the absence of unit clauses, one typically is interested in the (suboptimal) subset of heuristics which preserve these decorrelation properties of the Poissonian ensemble, so that a statistical analysis is achievable.

It is useful to consider the algorithm as partitioned into rounds, which consist of a single application of the heuristic rule (free step), followed by the cascade of unit-clause propagations (forced steps). In expectation, the number of variables fixed throughout a round of unit-clause propagation is described by a transition matrix depending on the clause distribution $\{C_i(X)\}$, which are constant to leading order during any “subcritical” round (defined below). The unit clauses generated in the first free step go on to generate other unit clauses and so forth, this can be described by a geometric series in the transition matrix $\mathcal{M}(X)$. Calling $\mathbf{p} = (p_T, p_F)$ and $\mathbf{m} = (m_T, m_F)$ respectively the expectations for the numbers of variables fixed to (True, False) at the first level of the cascade (p), and on the whole cascade (m), we have

$$\mathbf{m} = \mathbf{p} + \mathcal{M}(x)\mathbf{p} + \mathcal{M}^2(x)\mathbf{p} + \dots = (I - \mathcal{M}(x))^{-1}\mathbf{p}. \quad (\text{A1})$$

The matrix inverse above is justified from the fact that, for consistency of the approximations, we require the round to be subcritical, thus we must consider the range of parameters where the modulus of the largest eigenvalue of \mathcal{M} is smaller than 1. This is also responsible for the approximation that $\mathcal{M}(x)$ remains invariant (up to order $\frac{1}{N}$) during the cascade.

The transition matrix has two components. A first one comes from K -clauses which are “broken” into a bunch of $K - 1$ unit-clauses, because the fixed variable already satisfies the original one, and all the other $K - 1$ variables thus should take the value not satisfying the clause. A second one comes from 2-clauses (if any) which were just “shortened” because although the fixed variables are not satisfying them, the left over variables must, by definition creating unit clauses. Since there are $C_i(x)$ clauses of length i and $N - X$ variables left in the instance, these two terms are combined in the expression

$$\mathcal{M}(x) = \frac{2}{N - X} \left[(C_2(X) + 3C_3(X)) \begin{pmatrix} \epsilon(1 - \epsilon) & \epsilon^2 \\ (1 - \epsilon)^2 & \epsilon(1 - \epsilon) \end{pmatrix} + C_2(X) \begin{pmatrix} \epsilon(1 - \epsilon) & (1 - \epsilon)^2 \\ \epsilon^2 & \epsilon(1 - \epsilon) \end{pmatrix} \right]. \quad (\text{A2})$$

Here we identify that the unit-clause cascades on a large graph are a simple uncorrelated process, governed by the spectrum of a certain finite-size “transition matrix” (in our case, 2×2). If all the eigenvalues have $|\lambda_i| < 1$, the process is *subcritical*: the typical size of the cascades is $\sim 1/\min_i(1 - |\lambda_i|)$, and their average size concentrates. Conversely, when the gap $1 - |\lambda_i|$ vanishes, a single cascade could visit a finite fraction of a graph even in the large N limit, and thus could lead to a contradiction.

Also note that both the upper- and lower-bound arguments later derived are complemented by an analysis of the concentration properties of the process [21], and by a non-rigorous argument on the approximate decorrelation of distinct random restarts on a fixed instance. The vulgate version reads: “If a random algorithm succeeds with finite probability p on its first run, after $\sim n$ independent runs the probability of success will be $\sim 1 - \exp(-n \ln(1 - p) + \dots)$ ”, where the dots stand for some function of the correlations, small in n/N , caused by working with a fixed finite instance. We do not discuss here these complex technical points.

1. Upper bound

If, for a variable i selected randomly from the initial instance, both the cascades initiated by $s(i) = +1$ and $s(i) = -1$ percolate, there is a finite probability that they result in a certificate of contradiction.

Thus the upper bound for the SAT/UNSAT transition comes from the requirement that the cascades are on the edge of criticality already at time $X = 0$. At this point we have $C_2 = 0$ and $C_3 = N\gamma/3$, so that we get

$$\mathcal{M}(0) = 2\gamma \begin{pmatrix} \epsilon(1-\epsilon) & \epsilon^2 \\ (1-\epsilon)^2 & \epsilon(1-\epsilon) \end{pmatrix} \longrightarrow (\lambda_1, \lambda_2) = (4\gamma\epsilon(1-\epsilon), 0). \quad (\text{A3})$$

From this we see that a random instance is a.s. (randomized linear time) provable to be unsatisfiable for γ larger than the percolation threshold

$$\gamma_{\text{uc}}(\epsilon) = \frac{1}{4\epsilon(1-\epsilon)}. \quad (\text{A4})$$

2. Lower bound

The differential equations studied here are a generalization of those found by Kalapala and Moore [22] for the case of positive 1-in- K -SAT.

The heuristic determines the nature of the free step in our rounds. The two rules examined here in are random heuristic (RH[p]) and short clause heuristic (SCH). In RH[p] a variable is chosen at random, with uniform probability, from the remaining unassigned variables and set to true with probability p . In SCH, a random 2-clause is selected (if any exists) and a random literal set to satisfy the clause (hence the other literal is set to not satisfy it). If at some time no 2-clauses exist, a RH choice is performed, but fact will not be relevant in our statistical analysis: criticality of the cascade process will always arise after a time interval throughout which an extensive number of short-length clauses have been present (except at $X = \mathcal{O}(1)$).

If at some time X , p_T variables are set to True and p_F to False in expectation, then C_i changes accordingly. If an i -clause contains the variable just fixed, it is reduced to an $(i-1)$ -clause, and similarly an $(i+1)$ -clause can be reduced to an i -clause. Call $\boldsymbol{\epsilon} = (\epsilon, 1-\epsilon)$, $\boldsymbol{p} = (p_T, p_F)$, and $\mathbf{1} = (1, 1)$. Still in expectation:

$$C_i(X + \mathbf{1} \cdot \boldsymbol{p}) = (C_i(X) - a\delta_{i,2}(1 - \delta_{C_2(X),0})) \left(1 - \frac{i}{N-X}(\mathbf{1} \cdot \boldsymbol{p}) \right) + \frac{i+1}{N-X}(\boldsymbol{\epsilon} \cdot \boldsymbol{p}) C_{i+1}(X); \quad (\text{A5})$$

where $a = 0$ or 1 respectively in the case of RH[p] and SCH. The two heuristics are distinguished in that to initiate the cascades for RH[p] we have $\boldsymbol{p}_{\text{RH}[p]} = (p, 1-p)$, and for SCH we have $\boldsymbol{p}_{\text{SCH}} = (1, 1)$. The SCH value can be understood since setting one random literal in a two clause implies setting the other to the opposite value, thus setting variables to either ± 1 is equally likely in expectation.

A round can be described by incorporating the variables set in the forced steps. Suppose that during a subcritical round m_T variables are set to True and m_F to False in expectation (including the free step), and call \boldsymbol{m} the vector (m_T, m_F) . To leading order in $N-X$ the variation is

$$C_i(X + \mathbf{1} \cdot \boldsymbol{m}) = (C_i(X) - a\delta_{i,2}(1 - \delta_{C_2(X),0})) \left(1 - \frac{i}{N-X}(\mathbf{1} \cdot \boldsymbol{m}) \right) + \frac{i+1}{N-X}(\boldsymbol{\epsilon} \cdot \boldsymbol{m}) C_{i+1}(X), \quad (\text{A6})$$

A final simplification in the clause dynamics is to summarize the behavior by continuous variables $x = X/N$ and $c_i = C_i/N$. In the hypothesis of subcriticality, \boldsymbol{m}/N is infinitesimal, and we attain a differential equation description


$$\frac{d}{dx} c_i(x) = -a\delta_{i,2}\theta(c_2(x)) \left\langle \frac{\mathbf{1}}{\mathbf{1} \cdot \boldsymbol{m}} \right\rangle + \frac{1}{1-x} \left(-i c_i(x) + (i+1) \left\langle \frac{\boldsymbol{\epsilon} \cdot \boldsymbol{m}}{\mathbf{1} \cdot \boldsymbol{m}} \right\rangle c_{i+1}(x) \right), \quad (\text{A7})$$

For both RH[p] and SCH rules, the equation for $c_3(x)$ gives

$$c_3(x) = \frac{\gamma}{3}(1-x)^k, \quad (\text{A8})$$

Instead, for $c_2(x)$ the equation is nonlinear. Indeed we get that $\langle m_T \rangle$ and $\langle m_F \rangle$ are given by the combination of equations (A1, A2, A8), and thus depend on the unknown function $c_2(x)$ (besides, of course, x ,

FIG. 7: On the left, profiles of $\lambda(x)$ along decimation time x , for RH[1], at various ϵ and at the corresponding critical value of γ . In all the cases, the functions $\lambda(x)$ are concave (up to the limit value $\epsilon = 1/2$, where $\lambda(x) = 1 - x$). For ϵ larger or smaller than the tricritical value $\epsilon^* = 0.272633$, the maximum of $\lambda(x)$ is achieved respectively at $x = 0$ or at $x > 0$. On the right, critical curves $\gamma_{\text{sch}}(\epsilon)$ and $\gamma_{\text{RH}[1]}(\epsilon)$, obtained through short-clause (SCH) and random heuristic at optimal parameter $p = 1$ (RH[1]). For a comparison with other values and curves here out of range, cfr. figures 1, and 3.



ϵ and γ). Using this expression within (A7) allows us to determine $c_2(x)$ by numerical integration, and thence $\lambda_{\text{max}}(x)$.

The best choice for the parameter p in RH[p] is the one which creates the smallest cascades, i.e. the one “more orthogonal” to the principal eigenvalue $\epsilon = (\epsilon, 1 - \epsilon)$, but compatible with the probabilistic interpretation of \mathbf{p} . Thus, in the whole interval $\epsilon \in [0, 1/2]$, the choice $p = 1$ is optimal.

Here we thus show the results for RH[1] and for SCH. The latter is always at least as good as the former, and gives a lower bound of $\gamma_{\text{sch}} = 1.6393$, while RH[1] attains $\gamma_{\text{rh}} = 1.6031$, for the case $\epsilon = 0$. Kalapala and Moore calculated these quantities for positive 1-in- K SAT, with compatible results for the $K = 3$ case (up to maybe a misprint exchanging RH[p] with RH[$1 - p$]).

3. Exact SAT/UNSAT threshold for $\epsilon > 0.2726$

This section proves the coincidence of the curves $\gamma_{\text{sch}}(\epsilon)$ and $\gamma_{\text{uc}}(\epsilon)$ for $\epsilon > 0.2726$. It was shown in the previous section that whenever the cascades remain subcritical we are in the Easy-SAT phase. The criteria for the cascades to be subcritical at $x = 0$ is precisely $\gamma < \gamma_{\text{uc}}(\epsilon)$. It thus suffices to show that the maximum (over the decimation time x) of the $\max_i |\lambda_i(x)|$, is attained for $x = 0$. This is indeed what happens in the interval $\epsilon \in [0.2726, 1/2]$.

Building on the previous section we will see that, for ϵ -1-in-3-SAT and our heuristics, $\lambda(x)$ is a concave function. So, the interval on which $\gamma_{\text{uc}}(\epsilon)$ and $\gamma_{\text{sch}}(\epsilon)$ coincide is the one in which

$$\left. \frac{d\lambda(x; \epsilon, \gamma = \gamma_{\text{uc}}(\epsilon))}{dx} \right|_{x=0} \leq 0, \quad (\text{A9})$$

the endpoint being determined by the corresponding equality.

It is possible to calculate the characteristic polynomial (and differentiate w.r.t. x). The expressions thereby found can, however, only be evaluated exactly at $x = 0$. At this value we have expressions for $c_i(x)$ and their derivatives in terms of the initial conditions and \mathbf{m} .

As we increase connectivity towards the critical limit, $4\gamma\epsilon(1 - \epsilon) = 1$, a further simplification is in the eigenvectors of \mathcal{M} which become ϵ and its orthogonal, the latter having null eigenvalue. Regardless of \mathbf{p} (which must have a component along ϵ in order to have positive components), we have

$$\mathcal{F}(x, \epsilon) := \frac{\epsilon \cdot \mathbf{m}}{\mathbf{1} \cdot \mathbf{m}}; \quad \mathcal{F}(0, \epsilon) = \frac{\epsilon \cdot \epsilon}{\mathbf{1} \cdot \epsilon} = 1 - 2\epsilon(1 - \epsilon). \quad (\text{A10})$$

Finally, the condition (A9) becomes

$$\left(1 + \frac{1}{4} \left(\frac{\epsilon}{1 - \epsilon} + \frac{1 - \epsilon}{\epsilon} \right)^2 \right) (1 - 2\epsilon(1 - \epsilon)) - 2 \leq 0. \quad (\text{A11})$$

So that finally, after the change of variable $x = 2\epsilon(1 - \epsilon)$, one gets the equation for the endpoint of the interval

$$2x^3 - 2x^2 + 3x - 1 = 0, \quad (\text{A12})$$

whose only real solution gives $\epsilon = 0.272633$, or its symmetric point.

To show that the properties at $x = 0$ are sufficient to determine γ_{sch} it is necessary to show that whenever criteria (A11) is met, and $\lambda(0) < 1$, the algorithm is subcritical at all x . If we want a true analytic proof, besides the numerical evidence of figure 7 (left), a method is to find a function $\hat{\lambda}(x)$ such that

$$\lambda(x) \leq \hat{\lambda}(x) \leq \lambda(0), \quad (\text{A13})$$

hence establishing the result.

Since we find that $\lambda(x)$ is a monotonically increasing function of $c_2(x)$, an upper bound $\hat{c}_2(x) > c_2(x)$ implies an upper bound in $\lambda(x)$ also which we take to be $\hat{\lambda}(x)$. The bound function \hat{c}_2 is defined by replacing the complicated function $\mathcal{F}(x, \epsilon)$ by the constant value $\mathcal{F}(0, \epsilon)$ in the expression (A7), which are then exactly solvable for all x as

$$\hat{c}_2(x) = \gamma x(1-x)^2 \mathcal{F}(0, \epsilon) = \gamma x(1-x)^2 (1 - 2\epsilon(1-\epsilon)). \quad (\text{A14})$$

For $\text{RH}[\frac{1}{2}]$ and certain other heuristics this approximation can be shown to produce an upper bound for $c_2(x)$, and yet be exact at $x = 0$ in both absolute value and derivative.

This then allows an exact expression for $\frac{d\hat{\lambda}(x)}{dx}$ to be written in terms of x . Though the dependency on x remains complicated it can be established that

$$\frac{d\hat{\lambda}(x)}{dx} < 0 \quad \text{and} \quad \hat{\lambda}(0) < 1, \quad (\text{A15})$$

exactly in the same interval of ϵ in which (A11) holds. These fact proves that our “local” analysis at $x = 0$ was sufficient for the purpose of identifying the maximum over x of $\lambda(x)$ in this interval.

As a final remark we note that the proof of this exact bound is indirectly reliant on the convexity of the curves for all ϵ (figure 7 left). Interestingly we found that for ϵ -1-in-3-SAT with $k > 3$ the curves are not convex for some ϵ ; the gradient at $x = 0$ may be negative and yet the global maximum in lambda appears at $x > 0$. On first inspection a rigorous bound appears more challenging to obtain in these cases.

APPENDIX B: OTHER BOUNDS

1. Upper bound from first moment method

Here we obtain the statistical properties of the 2-core of random bipartite graphs, in the Erdős-Rényi ensemble described in section IA, with $K = 3$. Assuming $\gamma > 1/2$, the percolation transition, we solve self-consistently for the probability that a given branch of the graph is not percolating.

We use “giant” or “small” for synonymous of “of size of order N ” or “of size of order 1” respectively. Indeed, for graphs in our ensemble, a.s. there is a single giant 2-core component. Each edge is either attached on both sides to a small tree; is attached to a small tree on one of the two extrema, and to the giant 2-core on the other one (i.e. it is in the leaf-part of the giant component); or is connected to the 2-core through both extrema. Only in this last case is it in the 2-core of the graph.

Consider an incoming edge from a clause on the original graph. Call q the probability that the part of the graph “downstream” is a tree. The edge will be attached to a variable, participating in k other clauses (k Poissonian distributed of rate γ). For each clause there will be two incoming edges, which must also be connected to finite trees. Self-consistency will thus require that

$$q = \sum_k e^{-\gamma} \frac{\gamma^k}{k!} q^{2k} = e^{-\gamma(1-q^2)}. \quad (\text{B1})$$

The two functions $q(\gamma)$ and $\gamma(q)$, inverse of each other, are both monotonic on our domains $q \in [0, 1]$, $\gamma \in [1/2, +\infty)$. In particular, $\gamma(q)$ has an algebraic form

$$\gamma(q) = -\frac{\ln q}{1-q^2} \quad (\text{B2})$$

so that using q as a parameter instead of γ will simplify our equations.

The probability that an incoming branch from a variable is connected to a tree is q^2 , as it is the probability that both outgoing branches from the neighbouring clause are connected to a tree. Thus, the average number of variables of coordination $k \geq 2$, $\langle N_k \rangle$, is proportional to a Poissonian with rate $\gamma(1 - q^2) = -\ln q$. Then, there will be $\langle M_3 \rangle = (1 - q)^3 M$ clauses remaining of degree 3, and $\langle M_2 \rangle = 3q(1 - q)^2 M$ clauses reducing to 2-sat clauses, all the others are decimated by the leaf removal. The number of edges is $E = \sum_k k N_k = 3M_3 + 2M_2$, so on average $\langle E \rangle = N(-\ln q)(1 - q)$. All these averaged quantities are concentrated.

Consider the ensemble of configurations s , in “spin” notations as in the rest of the paper. Call x_k the fraction of variables of degree k which take value $+1$: the space of configurations is thus described by the infinite vector $\{x_k\}_{k \geq 2}$, with each x in $[0, 1]$, and a vector $\{x_k\}$ comes with an entropy

$$S_{\text{var}}(\vec{x}) = \sum_k N_k h(x_k), \quad (\text{B3})$$

where we use the common two-state entropy function $h(x) = -x \ln x - (1 - x) \ln(1 - x)$. Denote by x the fraction of incoming edges from variables assigned value $+1$

$$x = \frac{1}{E} \sum_k k N_k x_k, \quad (\text{B4})$$

and by p the fraction of edges (ai) such that $J_{ai}s_i = +1$, and hence $p = (1 - \epsilon)x + \epsilon(1 - x) = \epsilon + (1 - 2\epsilon)x$.

The probability that a 1-in-3-clause is satisfied is thus $3p(1 - p)^2$, and the probability that a 2-sat clause is satisfied is $1 - p^2$. So we get for the entropy term coming from clauses

$$S_{\text{cla}}(p(\vec{x})) = M_3 \ln(3p(1 - p)^2) + M_2 \ln(1 - p^2). \quad (\text{B5})$$

The upper bound on the SAT/UNSAT transition is achieved by the line, in the (q, ϵ) plane (with range $[0, 1] \times [0, 1/2]$), where the total (intensive) entropy $S(q, \epsilon)$ vanishes:

$$S(q, \epsilon) = \max_{\{x_k\}} S(q, \epsilon; \vec{x}); \quad S(q, \epsilon; \vec{x}) = S_{\text{var}}(\vec{x}) + S_{\text{cla}}(p(\vec{x})). \quad (\text{B6})$$

This variational problem is infinite dimensional, thus at first sight intractable. Instead, stationarity with respect to x_k produces

$$\frac{1}{k} \ln \frac{x_k}{1 - x_k} = (1 - 2\epsilon) \left(\frac{(1 - q)}{3(1 + q)} \frac{1 - 3p}{p(1 - p)} - \frac{q}{1 + q} \frac{2p}{1 - p^2} \right) =: y(p). \quad (\text{B7})$$

Remarkably, a single parameter y describes the family of possibly stationary vectors $\{x_k\}$, this being a residue of the original independence of the Poisson ensemble.

$$x_k(y) = \frac{1}{1 + e^{-ky}}. \quad (\text{B8})$$

Then, we can get self-consistently p from y , through the x_k 's

$$p(y) = \epsilon + (1 - 2\epsilon) \frac{1}{E} \sum_k k N_k \frac{1}{1 + e^{-ky}} = \epsilon + (1 - 2\epsilon) \frac{q}{1 - q} \sum_k \frac{(-\ln q)^k}{k!} \frac{1}{1 + e^{-(k+1)y}}. \quad (\text{B9})$$

So, the expression for the entropy $S(q, \epsilon)$ is given by the function

$$S(q, \epsilon) = q \sum_k \frac{(-\ln q)^k}{k!} h \left(\frac{1}{1 + e^{ky}} \right) + \frac{(-\ln q)(1 - q)}{3(1 + q)} ((1 - q) \ln(3p(1 - p)^2) + 3q \ln(1 - p^2)), \quad (\text{B10})$$

where the values of p and y are determined by the (only) solution of the nonlinear system of two equations (B7) and (B9). The set of points $(\epsilon, \gamma(q))$ where the function $S(q, \epsilon)$ vanishes describes a curve which appears in the figures 1 and 3.

Finally, we remark that a better upper bound can be achieved if one realizes that a further removal procedure is allowed: if a variable is connected only to 2-sat clauses, and with edges all of the same sign, then one can safely fix it, satisfying all the neighboring clauses. The new reduced instance is SAT if and only if the original one is, but the number of solutions is potentially smaller: as this decreases fluctuations, the bound is improved.

For the case $\epsilon = 0$, this program is achieved in [23], although in the restriction to the variational space of x_k all equal, and leads in that case to the (better) bound $\gamma_c \leq 1.932$.

2. Algorithmic upper bound through embedding into 3-XOR-SAT

An instance of 1-in-3-SAT is SAT only if the corresponding 3-XOR-SAT instance is SAT, where 3-XOR clauses allows also for the extra “spurious” configuration $(J_1\sigma_1, J_2\sigma_2, J_3\sigma_3) = (+, +, +)$.

Random Erdős-Rényi graphs with $K = 3$ have a finite core (under hypergraph leaf removal: if a variable has degree 1, one removes it together with the incident clause) beyond a “dynamical” threshold $\alpha_d = 0.818$. In a range $\alpha_d < \alpha < \alpha_c = 0.918$ there is an exponential number of solutions in the XOR-SAT problem, even if restricted to the core. However, beyond the critical value α_c there are no longer solutions (other than the single trivial one, with all $\sigma_i = +1$, in the case of “positive” instances with all $J_{ai} = +1$). These results can be found, for example, in [24, 25].

The last situation can be detected in polynomial time, by Gaussian elimination on the adjacency matrix: if the rank equals the number of variables (more generally, if it is smaller by at most $\mathcal{O}(\ln N)$), the solutions on the core can be checked in polynomial time. As a.s. all the variables are forced to be $+1$, a fraction of order 1 of the clauses in the core will be proven to be satisfiable only by the “spurious” configuration $(+, +, +)$. This provides a certificate of unsatisfiability for the original instance in the random ϵ -1-in-3-SAT ensemble.

So, at all ϵ , for $\gamma > 3\alpha_c = 2.754$, one gets a.s. a certificate of unsatisfiability in randomized cubic time (an upper bound for matrix triangulation). This proves that the Easy-UNSAT phase starts from a finite γ at all ϵ .

The method strongly relies on the fact that a XOR-SAT core exists in the instance. Unfortunately, this is not the case for the customary reductions of SAT to ϵ -1-in- K -SAT, even at large α (as the former constraints are much more sloppy than the latter, the reduction makes use of auxiliary leaf structures), so the method does not extend to a randomized polynomial-time algorithm for finding a certificate in large- α 3-SAT instances, in agreement with the widespread conjecture that such an algorithm can not exist [26].

APPENDIX C: 1RSB AT POSITIVE ENERGY (FINITE y)

In this appendix we describe the 1RSB solution in the energetic zero-temperature limit, but at finite value of parameter y . That allows us to compute the dependence of complexity Σ on energy E for a given probability of negation ϵ and connectivity γ .

The survey propagation equations (42), (43) at finite y becomes

$$p_+^{i \rightarrow a} + p_0^{i \rightarrow a} = \mathcal{N}_{i \rightarrow a}^{-1} \left(\sum_r e^{-yr} \text{prob}(\Delta \tilde{E}^{i \rightarrow a} = r, \sum_{b \in \partial i \setminus a} u_{b \rightarrow i} \geq 0) \right); \quad (\text{C1a})$$

$$p_-^{i \rightarrow a} + p_0^{i \rightarrow a} = \mathcal{N}_{i \rightarrow a}^{-1} \left(\sum_r e^{-yr} \text{prob}(\Delta \tilde{E}^{i \rightarrow a} = r, \sum_{b \in \partial i \setminus a} u_{b \rightarrow i} \leq 0) \right); \quad (\text{C1b})$$

$$p_0^{i \rightarrow a} = \mathcal{N}_{i \rightarrow a}^{-1} \left(\sum_r e^{-yr} \text{prob}(\Delta \tilde{E}^{i \rightarrow a} = r, \sum_{b \in \partial i \setminus a} u_{b \rightarrow i} = 0) \right); \quad (\text{C1c})$$

where $\mathcal{N}_{i \rightarrow a}$ is the normalization factor, and $\Delta \tilde{E}^{i \rightarrow a} = \sum_{b \in \partial i \setminus a} |u_{b \rightarrow i}| - |\sum_{b \in \partial i \setminus a} u_{b \rightarrow i}|$. When using Hamiltonian \mathcal{H} as in (2) (as we did), the probabilities of biases are given by the incoming probabilities of fields as

$$q_{J_{ai}}^{a \rightarrow i} = \mathcal{N}_{a \rightarrow i}^{-1} p_{-J_{aj}}^{j \rightarrow a} p_{-J_{ak}}^{k \rightarrow a}; \quad (\text{C2a})$$

$$q_{-J_{ai}}^{a \rightarrow i} = \mathcal{N}_{a \rightarrow i}^{-1} (p_{J_{aj}}^{j \rightarrow a} (1 - p_{J_{ak}}^{k \rightarrow a}) + (1 - p_{J_{aj}}^{j \rightarrow a}) p_{J_{ak}}^{k \rightarrow a}); \quad (\text{C2b})$$

$$q_0^{a \rightarrow i} = \mathcal{N}_{a \rightarrow i}^{-1} [(p_{-J_{aj}}^{j \rightarrow a} p_0^{k \rightarrow a} + p_0^{j \rightarrow a} p_{-J_{ak}}^{k \rightarrow a} + p_0^{j \rightarrow a} p_0^{k \rightarrow a}) + p_{J_{aj}}^{j \rightarrow a} p_{J_{ak}}^{k \rightarrow a} e^{-y}]; \quad (\text{C2c})$$

where $\mathcal{N}_{a \rightarrow i}$ is again the normalization. If we were using \mathcal{H}' of equation (13) instead, the last summand in (C2c), proportional to e^{-y} , would have appeared in (C2b).

Similarly, for the replicated free energy and energy we have

$$-y\Phi(y) = \sum_a \ln \left(\sum_r e^{-yr} \text{prob}(\Delta E^{a \cup \partial a} = r) \right) - \sum_i (d_i - 1) \ln \left(\sum_r e^{-yr} \text{prob}(\Delta E^i = r) \right); \quad (\text{C3})$$

$$E(y) = \sum_a \frac{\sum_r r e^{-yr} \text{prob}(\Delta E^{a \cup \partial a} = r)}{\sum_r e^{-yr} \text{prob}(\Delta E^{a \cup \partial a} = r)} - \sum_i (d_i - 1) \frac{\sum_r r e^{-yr} \text{prob}(\Delta E^i = r)}{\sum_r e^{-yr} \text{prob}(\Delta E^i = r)}; \quad (\text{C4})$$

γ	y_{ex}	y_{max}	E_{max}	Σ_{max}	y_{gs}	E_{gs}	y_{min}	E_{min}	Σ_{min}
1.850	2.9	3.35	0.000410	0.00414		0			(+0.00247)
1.879	2.1	2.83	0.000969	0.00441		0			(-0.00015)
1.900	1.6	2.64	0.00143	0.00471	5.31	0.000282			(-0.00215)
1.935	1.0	2.40	0.00236	0.00480	4.31	0.000939			(-0.00570)
1.970	0.6	2.25	0.00336	0.00517	3.83	0.00176			(-0.00954)
2.000	0.4	2.09	0.00428	0.00534	3.60	0.00255	7.37	0.000149	-0.0111
2.050	0.1	2.04	0.00573	0.00611	3.27	0.00406	6.37	0.00107	-0.0121
2.100	0.1	1.94	0.00613	0.00815	3.04	0.00580	5.76	0.00232	-0.0130
2.150	0.1	1.86	0.00766	0.0104	2.88	0.00632	5.27	0.00377	-0.0138
2.200	0.1	1.67	0.00993	0.0126	2.72	0.00674	5.32	0.00549	-0.0155

TABLE I: Values at special points in the complexity vs. energy curves of figure 8. The parametrizations $E(y)$ and $\Sigma(y)$ at various gamma exist for $y \gtrsim y_{\text{ex}}(\gamma)$ (the estimate is an upper bound, being the first value for which a non-paramagnetic RS solution emerged). All energies E and complexities Σ are extensive, and factors $1/N$ in the table entries are understood. The triplets $(y, E, \Sigma)_{\text{max}}$ and (if any) $(y, E, \Sigma)_{\text{min}}$ are the two turning points, while $(y, E, \Sigma = 0)_{\text{gs}}$ is the point where the complexity vanishes, along the physical branch of the curve. In the column for Σ_{min} , the first 5 values are instead $\Sigma(E = 0)$, as $E_{\text{min}} = 0$ without turning point in that case.

The probabilities $\text{prob}(\Delta E = r)$ have to be computed algorithmically (sum over all combinations). We found a closed formulas only when $r = 0$ (no contradictions), equations (42) and (48).

1. Complexity as a function of energy

Again with the population-dynamics algorithm we solve eqs. (C1)-(C4) and from the solution obtain the replicated potential $\Phi(y)$ (C3). The function $\Sigma(E)$, plotted in fig. 8, is obtained by fitting function $\Phi(y)$ and computing the Legendre transform (37) of the fit:

$$E(y) = \frac{\partial(y\Phi(y))}{\partial y}; \quad \Sigma(y) = y^2 \frac{\partial\Phi(y)}{\partial y}. \quad (\text{C5})$$

A non-paramagnetic solution of (C1,C2) exists only above a value $y_{\text{ex}}(\gamma)$, which decreases monotonically, from its asymptotics $y_{\text{ex}} \rightarrow +\infty$ for $\gamma \searrow \gamma_{\text{sp}}$ to $y_{\text{ex}} \rightarrow 0$ for $\gamma \rightarrow \infty$. The function $\Sigma(E)$ is parametrically identified from the two $\Sigma(y)$ and $E(y)$ above. Assuming the latter are regular functions and noting that $\partial\Sigma(E)/\partial E = y$, one also finds that $\Sigma(E)$ is regular, and convex or concave respectively if parameter y grows towards right or left, up to (possibly) special points where the curve changes concavity (*turning points*). Note that only the concave parts have a physical meaning, while the convex parts are the “non-physical” portion of the formal solution.

It turns out that, for $\gamma_{\text{sp}} < \gamma < \gamma_p$, there is a single turning point, from convex at small y to concave at large y , the values E_{max} and y_{max} labeling the values of E and y at this point. Instead, for $\gamma > \gamma_p$, a second turning point, from concave to convex, arises at higher values of y , thus defining the values E_{min} and y_{min} .

The labels of “max” and “min” stand for the fact that there exist phases with energy E (a number approximatively $\exp(\Sigma(E))$, if $\Sigma(E) > 0$, or in a fraction of instances of order $\exp(\Sigma(E))$, if $\Sigma(E) < 0$), only for energies in a range $E_{\text{min}} < E < E_{\text{max}}$ (where $E_{\text{min}} = 0$ for $\gamma < \gamma_p$), while we should interpret that there are no pure phases with energy E out of the range above, up to a subexponential fraction of instances.

We find that, for all γ , $\Sigma(E_{\text{max}}) > 0$, while for $\gamma > \gamma_p$, $\Sigma(E_{\text{min}}) < 0$. An intermediate value $E_{\text{min}} < E_{\text{gs}} < E_{\text{max}}$, corresponding to the one at which the complexity vanishes, exists for all $\gamma > \gamma^*$. It is the value of minimum energy of a configuration in a typical instance sampled from the corresponding ensemble, so it is important for the statistical properties of the “optimization” problem in the UNSAT phase.

In the table I we show the numerical values of the quantities described above, for a range of relevant γ 's, obtained by population dynamics.

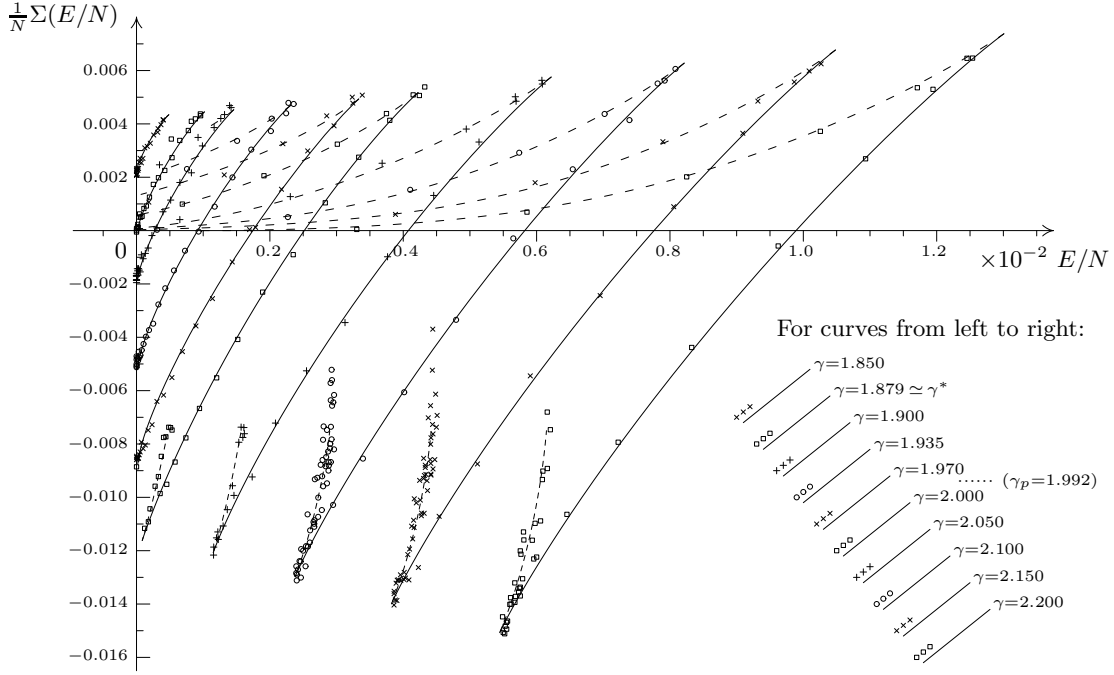


FIG. 8: Complexity as a function of energy for positive 1-in-3 SAT, for several different connectivities. $\gamma = 1.850$ is in the SAT region, $\gamma = 1.879$ is near the SAT/UNSAT transition, $1.900 \leq \gamma \leq 1.970$ is in the UNSAT region and $2.000 \leq \gamma \leq 2.200$ is in the UNSAT region where a solution at zero energy does not exist.




FIG. 9: The values of $E_{\max}(\gamma)$, $E_{\text{gs}}(\gamma)$ and $E_{\min}(\gamma)$, and a quadratic fit which extrapolates to $E_{\max}(\gamma_{\text{sp}}) = E_{\text{gs}}(\gamma^*) = E_{\min}(\gamma_p) = 0$.

Let us mention that, in our interpretation, the possibility of hard contradictions, peculiar to 1-in-3 SAT and other highly constrained NP-complete problems and absent in the more intensively studied K -SAT and Coloring, is responsible also for the existence of the second turning point, and the second unphysical branch at high values of y , which is indeed a new feature of this system.

APPENDIX D: STABILITY OF 1RSB

In this appendix we describe how to check the self-consistency (stability) of the 1RSB solution, with a treatment similar to the one in section III E, for the replica-symmetric solution. We do it only for the solution at zero energy, $y \rightarrow \infty$: as we will see, this is sufficient to determine the SAT/UNSAT

transition line in an interval of ϵ near $\epsilon = 0$, thus complementing the informations we already have for the neighbourhood of $\epsilon = 1/2$.

The stability analysis of the replica-symmetric solution investigates if the replica-symmetric state tends to split into exponentially many states. In the case of 1RSB we have two stability conditions to test this. The type I stability condition determines the tendency of 1RSB states to aggregate, and the type II determines the tendency of the states to split. The names type I and type II comes from [11, 29]. In the case that the 1RSB solution is not stable, i.e. the states tend to split or aggregate, we would deduce instability towards the 2-step of Replica-Symmetry-Breaking (2RSB). However, this is not expected to be the correct picture, while what more probably occurs is what is called full-RSB [31], or something else still unknown.

The method for computing these stabilities for models on random graphs was introduced in [29] and applied to K -SAT [11, 32] and later to many other problems. For both types of 1RSB stability there exist several equivalent analyses.

For the stability of type II we choose the *bug proliferation*. This is developed concisely in section D 1 for the 1RSB solution of zero energy ($y \rightarrow \infty$). For more theoretical background for this method see papers [29, 30, 32].


For type I instability it is possible to consider the convergence of survey propagation equations (42, 43) on a single graph. However, we choose to consider the *noise propagation* method, which uses a population-dynamics technique similar to that used for stability of the replica-symmetric solution, eq. (31). Again we give just the formulas and results in section D 2, for general explanation see [29, 30, 32].

1. Stability of the second kind, bug proliferation

Suppose that, in a neighbourhood of edge (a, i) , with j and k being the other two variables incident on a , an incoming warning $u_{b \rightarrow j}$ is changed from value u to another value u' . Assume also that there are $n-1$ other clauses c_2, \dots, c_n (besides a and b) incoming to node j , the warnings having values respectively u_2, \dots, u_n , and that there are m other clauses c'_1, \dots, c'_m (besides a) incoming to node k , the warnings having values respectively $v_1 \dots v_m$. Conditioned to the existence of the path $((bj), (ja), (ai))$, the other coordinations m and $n-1$ are decorrelated and Poissonian distributed, with rate γ . Denote with the letter \mathcal{J} the set of parameters describing the characteristics of the graph in this neighbourhood

$$\mathcal{J} = (n-1, m; J_{ai}, J_{aj}, J_{ak}); \quad (\text{D1})$$

Labels w and w' will describe the value of the output warning, $u_{a \rightarrow i}$, under $u_{b \rightarrow j} = u$ and u' respectively:



Define the six-dimensional transition matrix, in the index pairs (u, u') and (w, w') , pairs of distinct elements in $\{0, \pm 1\}^2$

$$\mathcal{P}_{\mathcal{J}}(w \rightarrow w' | u \rightarrow u') = \frac{1}{\mathcal{N}} \sum_{\substack{u_2 \dots u_n \in \{0, \pm 1\} \\ v_1 \dots v_m \in \{0, \pm 1\}}} q_{u_2}^{c_2 \rightarrow j} \dots q_{u_n}^{c_n \rightarrow j} q_{v_1}^{c'_1 \rightarrow k} \dots q_{v_m}^{c'_m \rightarrow k} \cdot \delta_{\mathcal{F}^{a \rightarrow i}(u, u_2, \dots, v_m), w} \delta_{\mathcal{F}^{a \rightarrow i}(u', u_2, \dots, v_m), w'} \delta_{\Delta E^{a \rightarrow i}(u', u_2, \dots, v_m), 0}. \quad (\text{D2})$$

The quantity $\mathcal{P}(w \rightarrow w' | u \rightarrow u')$ is proportional to the probability that the change $u \rightarrow u'$ in warning $u_{b \rightarrow j}$ has induced a change $w \rightarrow w'$ in the warning $u_{a \rightarrow i}$. Here $\mathcal{F}^{a \rightarrow i}(u^{b \rightarrow j}, u^{c_2 \rightarrow j}, \dots, u^{c_n \rightarrow j}, u^{c'_1 \rightarrow k}, \dots, u^{c'_m \rightarrow k})$ implements the cavity equations at zero energy (12), with the appropriate disorder parameters (J_{ai}, J_{aj}, J_{ak}) , so the delta's force this value to be equal to $u^{a \rightarrow i}$, in the two cases ($u^{b \rightarrow j} = u, u^{a \rightarrow i} = w$)

and ($u^{b \rightarrow j} = u', u^{a \rightarrow i} = w'$). The delta in the energy shift is the residual of the reweighting factor $\exp(-y\Delta E^{a \rightarrow i})$ in the limit $y \rightarrow \infty$ (note that it only appears on the pair (u', w') of biases along the chain). Normalizations $\mathcal{N} = \mathcal{N}_{j \rightarrow a} \mathcal{N}_{k \rightarrow a} \mathcal{N}_{a \rightarrow i}$ are those from eqs. (42, 43).

The transition matrix defined above determines if a small fluctuation in the equilibrium distribution of the fields is reinforced through the cavity iterations, thus leading to an instability. After d iterations, the modulus of the fluctuation changes on average by some factor $(2\gamma)^d \text{Tr} \langle \mathcal{P}_{\mathcal{J}_1} \dots \mathcal{P}_{\mathcal{J}_d} \rangle$, because there are on average $(2\gamma)^d$ chains of length d ending on a given edge, and the trace of a ‘‘chain’’ of transition matrices estimates the influence of changing a bias at an edge at distance d upstream. So we define the (finite- d and $d \rightarrow \infty$) type-II stability parameters

$$\mu_{\text{II}}(d) = 2\gamma \text{Tr} \langle \mathcal{P}_{\mathcal{J}_1} \dots \mathcal{P}_{\mathcal{J}_d} \rangle^{\frac{1}{d}}; \quad \mu_{\text{II}} = \lim_{d \rightarrow \infty} \mu_{\text{II}}(d); \quad (\text{D3})$$

where the average is over the connectivity distribution and the disorder in negations, i.e. the parameters globally identified above with the letter \mathcal{J} . The various \mathcal{J} 's refers to the different segments of the chain, and thus are independent. Again, the stability condition reads

$$\mu_{\text{II}} < 1. \quad (\text{D4})$$

The matrix \mathcal{P} for the ϵ -1-in-3 SAT problem is six-dimensional, and we computed it for general realizations of negation and connectivities. It has a block-triangular form (a change $0 \rightarrow \pm 1$ never induces changes other than $0 \rightarrow \pm 1$ and a change $\pm 1 \rightarrow 0$ never induces $\pm 1 \rightarrow \mp 1$). Moreover two of the three 2×2 blocks on the diagonal, B , are equal and have elements always larger than those of the third block B'

$$\mathcal{P}(w \rightarrow w' | u \rightarrow u') : \quad \begin{array}{l} 0 \rightarrow \pm 1 \\ \pm 1 \rightarrow 0 \\ \pm 1 \rightarrow \mp 1 \end{array} \begin{array}{|c|c|c|} \hline & \begin{array}{c} 0 \rightarrow \pm 1 \\ \pm 1 \rightarrow 0 \end{array} & \begin{array}{c} \pm 1 \rightarrow 0 \\ \pm 1 \rightarrow \mp 1 \end{array} \\ \hline \begin{array}{c} 0 \rightarrow \pm 1 \\ \pm 1 \rightarrow 0 \end{array} & \begin{array}{|c|c|} \hline B & 0 \\ \hline * & B \end{array} & \begin{array}{|c|} \hline 0 \\ \hline \end{array} \\ \hline \begin{array}{c} \pm 1 \rightarrow 0 \\ \pm 1 \rightarrow \mp 1 \end{array} & \begin{array}{|c|c|} \hline * & * \\ \hline \end{array} & \begin{array}{|c|} \hline B' \\ \hline \end{array} \\ \hline \end{array}$$

So in the large- d limit we need to analyze only the 2×2 block $B_{\mathcal{J}}$ of elements, propagating the ‘bug’ $0 \leftrightarrow \pm 1$. For $\mathcal{J} = (n-1, m; J_{ai}, J_{aj}, J_{ak})$, the block elements take the form

$$B = \left[\delta_{J_{aj}, +1} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} + \delta_{J_{aj}, -1} \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \right] \begin{pmatrix} 0 & \alpha \\ \beta & 0 \end{pmatrix} \left[\delta_{J_{ai}, +1} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} + \delta_{J_{ai}, -1} \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \right]; \quad (\text{D5})$$

$$\alpha = \prod_{\ell=1}^m q_0^\ell \prod_{j=2}^n q_0^j; \quad \beta = \prod_{\ell=1}^m (q_0^\ell + q_{-J_{ak}}^\ell) \prod_{j=2}^n q_0^j; \quad (\text{D6})$$

and triplets (q_0^e, q_-^e, q_+^e) for different edges are independently sampled from the stationary distribution.

Results for the type-II stability are presented in figure 5.

2. Stability of the first kind, the noise propagation

In a similar way to bug proliferation, we write a sort of transfer matrix T

$$(T_{b \rightarrow j}^{a \rightarrow i})_{\sigma, \tau} = \frac{\partial q_\tau^{a \rightarrow i}}{\partial q_\sigma^{b \rightarrow j}}, \quad \sigma, \tau \in \{+, -\}. \quad (\text{D7})$$

The dependence of $q^{a \rightarrow i}$ on $q^{b \rightarrow j}$ is given by the survey propagation equations (42,43).

We perform a population-dynamics analysis, where to every edge in the population is associated a triple for the surveys, (q_-, q_0, q_+) , updated with the cavity equations (42,43), and a pair of *noise* parameters, $\vec{v} = (v_+, v_-)$, which are updated according to

$$\vec{v}^{a \rightarrow i} = \sum_{b \in \partial j \setminus a} T_{b \rightarrow j}^{a \rightarrow i} \vec{v}^{b \rightarrow j} + \sum_{c \in \partial k \setminus a} T_{c \rightarrow k}^{a \rightarrow i} \vec{v}^{c \rightarrow k}. \quad (\text{D8})$$

The motivation is to compute whether a small change in the equilibrated incoming survey $q^{b \rightarrow j}$ is dumped under cavity iterations.

The analysis goes in complete analogy with the one in section III E. We initialize the noise parameters with an arbitrary random procedure, and wait for equilibration of the distribution, up to a scaling overall, $\|v\|_t := \sum_e (|v_+^e| + |v_-^e|)$ where t denotes the iteration time, and the sum is over the population. The stability parameter is now, for some time t larger than equilibration,

$$\mu_I = \left\langle \frac{\|v\|_{t+1}}{\|v\|_t} \right\rangle, \quad (\text{D9})$$

and the stability condition is $\mu_I < 1$.

-
- [1] Cook S, *The Complexity of Theorem Proving Procedures*, Proc. of the third annual ACM symposium on Theory of computing, (1971), 151-158
- [2] Garey M R and Johnson D S *Computers and Intractability: A Guide to the Theory of NP-Completeness*, WH Freeman & Co. New York, NY, USA, 1979.
- [3] Mézard M and Parisi G, *The Bethe lattice spin glass revisited*, 2001 *Eur. Phys. J. B* **20** 217
- [4] Mézard M and Parisi G, *The cavity method at zero temperature*, 2003 *J. Stat. Phys.* **111** N. 1-34
- [5] Biroli G, Monasson R and Weigt M, *A variational description of the ground state structure in random satisfiability problems*, 2000 *Eur. Phys. J. B* **14** 551
- [6] Theoretical Computer Science 265 (2001), Special Issue on NP Hardness and Phase Transitions.
- [7] Levin L A, *Average case complete problems*, 1986 *SIAM Journ. on Comp.* **15** 285-286
- [8] Mézard M, Parisi G and Zecchina R, *Analytic and Algorithmic Solution of Random Satisfiability Problems*, 2002 *Science* **297** 812.
- [9] Mézard M and Zecchina R, *The random K-satisfiability problem: from an analytic solution to an efficient algorithm*, 2002 *Phys. Rev. E* **66** 056126
- [10] Mézard M, Parisi G and Virasoro M A, *Spin-glass Theory and Beyond*, vol. 9 of Lecture notes in Physics. World Scientific, Singapore, 1987
- [11] Montanari A, Parisi G and Ricci-Tersenghi F, *Instability of one-step replica symmetry broken phase in satisfiability problems*, 2004 *J. Phys. A* **37**, 2073
- [12] Krzakala F, Pagnani A and Weigt M, *Threshold values, stability analysis and high-q asymptotics for the coloring problem*, 2004 *Phys. Rev. E* **70**, 046705
- [13] Mulet R, Pagnani A, Weigt M and Zecchina R, *Coloring random graphs*, 2002 *Phys. Rev. Lett.* **89** 268701
- [14] Rosenthal J W, Speckenmeyer R and Kemp R, *Exact satisfiability, a natural extension of set partition, and its average case behavior*, Annals of Mathematics and Artificial Intelligence, Volume 6, N. 1-3/March, 1992
- [15] Achlioptas D, Chtcherba A D, Istrate G and Moore C, *The phase transition in 1-in-k SAT and NAE 3-SAT*, 2001 *SODA*, 721-722,
- [16] Monasson R, Zecchina R, *Tricritical Points in Random Combinatorics: the (2+p)-SAT case*, 1998 *J. Phys. A* **31**, 9209
- [17] Monasson R, Zecchina R, Kirkpatrick S, Selman B and Troyansky L, *Determining computational complexity from characteristic ‘phase transitions’*, 1999 *Nature* **400**, 133
- [18] Achlioptas D, Kirousis L M, Kranakis E and Krinzanc D, *Rigorous Results for (2 + p) – SAT*, Theoretical Computer Science, 265 (1-2), (2001), 109-129.
- [19] Walsh T, *From P to NP: COL, XOR, NAE, 1-in-k, and Horn SAT*, Proceedings of AAAI-2002, 2002.
- [20] Walsh T, *2+p-COL* Proceedings of the Computational Symposium on Graph Coloring and its Generalizations, held alongside CP-2002, Cornell, 2002.
- [21] Deroulers C and Monasson R, *Criticality and Universality in the Unit-Propagation Search Rule*, 2006 *Eur. Phys. J. B* **49** 339
- [22] Kalapala V and Moore C, *The Phase Transition in Exact Cover*, [cs.CC/0508037](https://arxiv.org/abs/cs.CC/0508037)
- [23] Knysh S, Smelyanskiy V N and Morris E R, *Approximating the satisfiability transition by suppressing fluctuations*, [cond-mat/0403416](https://arxiv.org/abs/cond-mat/0403416)
- [24] Ricci-Tersenghi F, Weigt M and Zecchina R, *Simplest random K-satisfiability problem*, 2001 *Phys. Rev. E* **63**, 026702
- [25] Mézard M, Ricci-Tersenghi F and Zecchina R, *Alternative solutions to diluted p-spin models and XORSAT problems*, 2003 *J. Stat. Phys.* **111** 505
- [26] Chvátal V and Szemerédi E, *Many hard examples for resolution*, 1988 *Journal of the ACM* **35** 759
- [27] Pearl J, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference* (2nd edition). Morgan Kaufmann Publishers, San Francisco, 1988.
- [28] Yedidia J S, Freeman W T and Weiss Y, *Understanding Belief Propagation and Its Generalizations*, Exploring Artificial Intelligence in the New Millennium, Chap. 8, pp. 239-269, 2003.
- [29] Montanari A and Ricci-Tersenghi F, *On the nature of the low-temperature phase in discontinuous mean-field spin glasses*, 2003 *Eur. Phys. J. B* **33** 339

- [30] Rivoire O, Biroli G, Martin O C and Mézard M, *Glass models on Bethe lattices*, 2004 *Eur. Phys. J.* **B 37** 55-78
- [31] Parisi G, *A sequence of Approximated Solution to the S-K Model for Spin Glasses*, 1980 *J. Phys. A* **13** L115-L121
- [32] Mertens S, Mézard M and Zecchina R, *Threshold Values of Random K-SAT from the Cavity Method*, 2006 *Rand. Struct. Alg.* **28**, 340-373
- [33] Mézard M, Palassini M and Rivoire O, *Landscape of solutions in constraint satisfaction problems*, 2005 *Phys. Rev. Lett.* **95**, 200202
- [34] Krzakała F, Montanari A, Ricci-Tersenghi F, Semerjian G and Zdeborová L, *Gibbs states and the set of solutions of random constraint satisfaction problems*, to be published in PNAS (2007).
- [35] Franz S and Leone M, *Replica bounds for optimization problems and diluted spin systems*, 2003 *J. Stat. Phys.* **111** 535
- [36] Franz S, Leone M and Toninelli F L, *Replica bounds for diluted non-Poissonian spin systems*, 2003 *J. Phys. A* **36** 10967
- [37] Selman B, Kautz H A and Cohen B, *Noise strategies for improving local search*, Proc. of the 12th nat. conf. on Artificial Intelligence (vol. 1), p.337-343, Seattle, United States, 1994
- [38] Seitz S and Orponen P, *An efficient local search method for random 3-satisfiability*, Proc. LICS03 Workshop on Typical Case Complexity and Phase Transitions. *Electr. Notes in Discr. Math.*, **16** (2003).
- [39] Seitz S, Alava M and Orponen P, *Focused local search on random 3-Satisfiability*, *J. Stat. Mech.* P06006 (2005)
- [40] Ardelius J and Aurell E, *Behavior of heuristics and state space structure near SAT/UNSAT transition*, *Phys. Rev. E* **74**, 037702 (2006).
- [41] Zdeborová L and Krzakała F, *Phase transitions in the Coloring of Random Graphs*, arXiv:0704.1269v1 [cond-mat.dis-nn]