

Entropic barriers as a reason for hardness in both classical and quantum algorithmsMatteo Bellitti *Department of Physics, Boston University, Boston, Massachusetts 02215, USA*

Federico Ricci-Tersenghi

*Dipartimento di Fisica, Sapienza Università di Roma, P.le A. Moro 2, I-00185 Roma, Italy;
CNR, Nanotec, Rome unit, P.le A. Moro 2, I-00185 Roma, Italy;
and INFN, Sezione di Roma I, P.le A. Moro 2, I-00185 Roma, Italy*

Antonello Scardicchio

*The Abdus Salam International Center for Theoretical Physics, Strada Costiera 11, 34151 Trieste, Italy
and INFN Sezione di Trieste, Via Valerio 2, 34127 Trieste, Italy*

(Received 14 March 2021; accepted 17 August 2021; published 7 October 2021)

We study both classical and quantum algorithms to solve a hard optimization problem, namely 3-XORSAT on 3-regular random graphs. By introducing a new quasi-greedy algorithm that is not allowed to jump over large energy barriers, we show that the problem hardness is mainly due to entropic barriers. We study, both analytically and numerically, several optimization algorithms, finding that entropic barriers affect in a similar way classical local algorithms and quantum annealing. For the adiabatic algorithm, the difficulty we identify is distinct from that of tunneling under large barriers, but does, nonetheless, give rise to exponential running (annealing) times.

DOI: [10.1103/PhysRevResearch.3.043015](https://doi.org/10.1103/PhysRevResearch.3.043015)**I. INTRODUCTION**

Hard discrete optimization problems are ubiquitous in scientific disciplines and practical applications. The problem of minimizing a complex cost function (or equivalently maximizing a reward function) naturally appears in many different contexts: e.g., in physics in the computation of ground state configurations, in statistics in the maximization of the likelihood, in machine learning in the training of artificial neural networks, and so on.

Although real-world problems have usually local structures that make their analysis difficult, it is commonly believed that the main source of computational hardness arises from the strong long-range correlations that exist among variables, and this effect can be studied also in more idealized and simple-to-solve models, where some analytic control is possible. In other words, in hard optimization problems, starting from an optimal or near-optimal configuration, the change of a single variable (or a small subset of variables) often requires the rearrangement of many more variables in order to remain close to optimality; often the variables to be rearranged are not even close to the modified variable. This property makes the search for the optimal configuration a challenging task

even for sophisticated algorithms (see for example the case of backtracking algorithms, like DPLL [1]).

An ideal setting for studying this kind of hard optimization problems is provided by constraint satisfaction problems defined on sparse random graphs. Such problems have a twofold benefit: They can be solved analytically using the cavity method, a tool from statistical physics of disordered systems, and they can be efficiently handled on a computer, as the finite mean degree of the graph makes the computational resources (CPU and memory) required to represent an instance of the problem on the machine grow only linearly with the problem size.

Random constraint satisfaction problems (rCSP) are optimization problems where N discrete variables need to be assigned in order to satisfy $M = \alpha N$ constraints, each one involving a small subset of variables. The most famous among rCSP is perhaps random K -SAT [2]. Recently these rCSP have been the subject of intense studies based on statistical physics ideas with the aim of understanding the origin of their computational hardness [3–16]. Indeed, a common feature of all the hard rCSP is the presence of a broad range of the constraints per variable ratio α such that solutions to the problem exists with high probability (in the large N limit), but all known solving algorithms are unable to find any solution in a time growing polynomially with the problem size N . In this hard region it is expected that any solving algorithm requires a time growing exponentially with the problem size, $t \sim \exp(\alpha N)$.

By defining an energy function that counts the number of violated constraints, one can visualize the rCSP as the problem of searching for a zero-energy configuration in a complex energy landscape. The hard phase in rCSP does

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

actually corresponds to an energy landscape with exponentially many (in the system size) local minima that can trap the searching dynamics. The energy barriers between these minima are usually considered the main source of computational complexity, as any local dynamics is required to jump over these barriers in order to proceed further in the search for the optimal configuration.

Based on the above picture, it is often believed that a quantum evolution—that allows for tunneling events—may escape local minima more efficiently than a classical stochastic dynamics. This physically reasonable expectation implies that quantum algorithms may be faster in the search for optimal configurations than their classical counterparts, and has fueled interest in quantum algorithms that could benefit from this phenomenon: quantum annealing [17–22], its more recent variant, the quantum approximate optimization algorithm [21], and population transfer [23–25] are some well-known examples. All these algorithms typically show a complexity growth comparable with the best classical algorithms but, despite their initial promise, it is entirely possible that the limitations they exhibit are insurmountable, and it is unlikely that they could solve NP-hard problems in polynomial time. The limitations might come from the exponentially small tunneling rate out of a local minimum (a phenomenon linked to many-body localization and the existence of an emergent integrable dynamical phase [20,26–28]) or might come from other dynamical phenomena [29]. In this paper we identify one such phenomenon.

We consider one of the hardest sparse rCSP, namely random 3-XORSAT, and show that the problem hardness can be interpreted as coming essentially from *entropic barriers*: A local algorithm searching for a solution is hindered not only by the energy barriers around a local minimum, but also by the possibility to explore a large number of directions and local minima with a small change in energy. This is an entropic barrier. We introduce a quasi-greedy algorithm, unable to jump over large energy barriers, and notice that it is able to solve the problem as efficiently as state-of-the-art algorithms designed to overcome large energy barriers. The latter algorithms efficiently solve problems where energy barriers are the only obstacle, but still suffer from entropic effects. While we expect entropic barriers to be a generic feature of rCSP, the observation that a quasi-greedy algorithm performs as well as the others suggests that this model is a perfect candidate to understand the effect of entropic barriers.

We investigate several algorithms, both classical and quantum, in order to better understand the effect of entropic barriers. For all the algorithms analyzed, we find that the time to reach a solution scales exponentially with the system size and quantum dynamics seem to suffer from the presence of entropic barriers as much as the classical algorithmic dynamics. As the effort to build a quantum computer are finally giving up some results [30], we believe it is important to identify all possible stumbling blocks for quantum architectures.

II. MODEL DEFINITION AND ITS KNOWN SOLUTION

The random 3-XORSAT problem is among the simplest rCSP [3]: it is made of N binary variables $x_i \in \{0, 1\}$ that have

to satisfy $M = \alpha N$ parity checks of the kind

$$x_i \oplus x_j \oplus x_k = b_{ijk}, \quad (1)$$

where the variables entering each constraint are randomly chosen and the parity check bit b_{ijk} is 0 or 1 with equal probability. In a K -XORSAT problem each constraint involves K variables, but for the sake of simplifying the presentation we restrict to the random 3-XORSAT problem, where each constraint involves exactly 3 randomly chosen variables. Increasing α the typical problem becomes more and more difficult to solve. Solutions exist with high probability in the large N limit until the sat-unsat threshold α_s [31]. However, the most interesting transition from the point of view of searching algorithms is the clustering (or dynamical) transition that takes place at α_d before α_s [6]. For $\alpha \in [\alpha_d, \alpha_s]$ the space of solutions is shattered in exponentially many cluster of solutions [6,32] and this is what makes the search for solutions much more difficult [3,6,15]. Actual thresholds for random 3-XORSAT are $\alpha_d = 0.818$ and $\alpha_s = 0.918$. This picture has been proven rigorously to a large extent [33]. The hardness of the problem of finding a solution for some classes of local algorithms in the region $\alpha \in [\alpha_d, \alpha_s]$ has been proven [34–36], and found to depend on the so-called “overlap gap property”, which in practice implies the clustering of solutions taking place after the dynamical transition α_d . In other words, for $\alpha > \alpha_d$ the geometry of solutions is such that the Hamming distance d between any pair of solutions is either very small $d < d_1$ (for pairs of solutions in the same cluster) or very large $d > d_2$ (for pairs of solutions in different clusters) [6]. It is exactly the existence of such a range of distances with no solutions that creates an algorithmic bottleneck, while for $\alpha < \alpha_d$ local algorithms can sample the space of solutions [33]. So we are going to focus our attention on a 3-XORSAT problem beyond the dynamical threshold α_d [37].

Given that we are interested in using this model as a benchmark for optimization, we need two more ingredients: (i) We need to define an energy function whose ground state configurations are the solutions to the problem; the simplest choice consists in just counting the number of violated parity checks via the following Hamiltonian:

$$\mathcal{H}_0[s] = \frac{1}{2} \left(M - \sum_{a=1}^M J_a \prod_{i \in \partial a} s_i \right), \quad (2)$$

where $s_i = (-1)^{x_i}$ are Ising spins and $J_a = (-1)^{b_a}$ the couplings, being a an index running over all interactions (triplets for 3-XORSAT) and ∂a the set of variables entering the a -th interaction. (ii) At least a solution must always exist, and this can be ensured by enforcing a specific configuration to satisfy all the constraints. For example, by setting all $b = 0$, the configuration $x_i = 0 \forall i$ is always a solution. One may think this way of building the model naturally favors the imposed or planted configuration, but this is not the case for the XORSAT problem. As noticed since Ref. [38], finding the imposed or planted solution is like finding the crystal in a model of a liquid that upon cooling spontaneously forms a glass: it is well known that crystallization requires an activated dynamical process (nucleation), which is exponentially rare in models with long range interactions, as the random XORSAT. Planted models, which are hard to solve are the most natural

candidates for optimization benchmarks and the planted XORSAT turns out to be the hardest among these [5]. Properties of planted models have been studied in a great detail [39,40] and reviewed in [41].

It is worth noticing that the planted model, being equivalent to the solution of a set of linear equations in $GF(2)$, can be seen as the noiseless limit of a fundamental problem in cryptography, namely “learning parity with noise” [42,43]. We stress that, restricting to local algorithms in the very sparse regime (as we do here), the noiseless problem is as hard as the problem with a very small amount of noise. So our results can shed light on the origin of the hardness in the problem of learning parity with noise.

The Hamiltonian for the planted model simplifies to the following:

$$\mathcal{H}[s] = \frac{1}{2} \left(M - \sum_{a=1}^M \prod_{i \in \partial a} s_i \right), \quad (3)$$

which is indeed minimized by the configuration $s_i = 1 \forall i$. This is the energy function we are going to minimize in order to test classical and quantum optimization algorithms.

The last relevant choice regards the interaction hypergraph, that is the set of M triplets. In the model where the M triplets are chosen randomly the degree of each variable is a Poisson random variable of mean 3α . A different choice is the one where the interaction hypergraph is chosen such that each variable has the same degree d : this is called a random regular hypergraph and can be generated via the configurational model where N variables are given d legs each and $M = Nd/3$ interactions are given 3 legs each, and then variables and interactions legs are coupled in a random way, just avoiding that the same variable enters more than once in the same interaction. We are going to use the random regular version for the numerical simulations, while the random Poisson version is used for some analytic computations. The two versions share the same physical behavior.

The statistical properties of the random regular XORSAT problem are well known [39,44–46]. Hereafter we are going to focus our studies on the $d = 3$ case. In this case the random and the planted models are equivalent in the large N limit for any positive temperature (at $T = 0$ the subextensive differences between the two models may lead to some discrepancy in the number of solutions discussed in Refs. [47,48]). We are going to use the planted model in order to be sure that a solution always exists even for finite (and small) values of N .

Having fixed the degree of the hypergraph such that the XORSAT problem is in its hard phase, we can consider now the Gibbs measure corresponding to Hamiltonian \mathcal{H} in Eq. 3 at any temperature T . The XORSAT problem corresponds to the problem of finding a zero-energy ground state, so it is somehow related to the $T = 0$ physics of the model, but the behavior of the model at $T > 0$ is interesting as well. Indeed it is well known [38] that when a model is beyond the dynamical transition point at $T = 0$ (e.g., for $\alpha > \alpha_d$) it undergoes a dynamical phase transition at a positive temperature T_d and for $T < T_d$ it has an exponentially large number of metastable

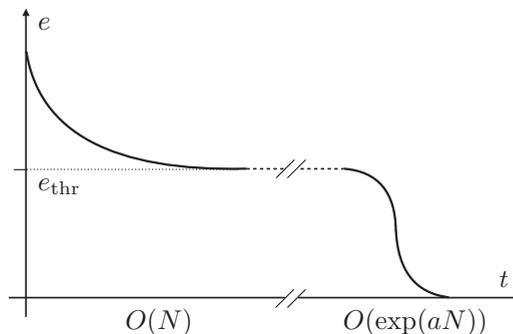


FIG. 1. Schematic picture for the energy relaxation in hard optimization problems. On short times the energy relaxes to a threshold value, while the ground state (solution) is reached only for times growing exponentially with the system size. The first regime can be described by ordinary differential equations taking the large N limit at t/N fixed, while the second requires the estimation of rare and large fluctuations.

states dominating the thermodynamics, $\mathcal{N} \sim \exp[N\Sigma]$, where Σ is the so-called complexity.

This is the case for the 3-regular 3-XORSAT model that shows a dynamical phase transition at $T_d = 0.255$ and a nonzero complexity of states at $T = 0$ that extends from $e \equiv \langle \mathcal{H} \rangle / N = 0$ to $e = e_d = 0.0206705$. Actually not all these states are expected to play a relevant role in the relaxation dynamics searching for ground states: from previous studies [45,46] we expect states above the marginal energy $e_{\text{marg}} = 0.018203$ to unlikely trap smart searching algorithms.

Unfortunately a precise connection between relaxation algorithms searching for low energy configurations (e.g., $T = 0$ Langevin dynamics) and the energy landscape that we can describe in a precise way via the computation of the complexity is still missing (and recent results have clarified that the situation is much more complicated than previously expected [49–51]). So we cannot make an analytical claim about the threshold energy, which is hard to go below by searching algorithms, but this threshold energy is certainly positive and close to e_{marg} . Reaching a solution, that is an $e = 0$ configuration, is a very hard problem and requires in general times scaling exponentially with the system size N .

III. THE OPTIMIZATION ALGORITHMS

In the study of the optimization algorithms, that is the out of equilibrium processes that try to minimize the energy, the order in which the large size limit ($N \rightarrow \infty$) and large times limit ($t \rightarrow \infty$) are taken is extremely important. We expect an algorithm-dependent threshold energy, $e_{\text{thr}} > 0$, to exist such that configurations with $e > e_{\text{thr}}$ can be reached in an “easy” way (e.g., in a time scaling linearly with N), while to reach a solution (i.e. a configuration with $e = 0$) a time growing exponentially in N is required in general for hard problems. (see the schematic picture in Fig. 1).

We will see that for some algorithms we are able to provide an approximate description of the dynamics in the regime where the $N \rightarrow \infty$ limit is taken before the $t \rightarrow \infty$ limit, thus estimating e_{thr} (that in the best cases is close to

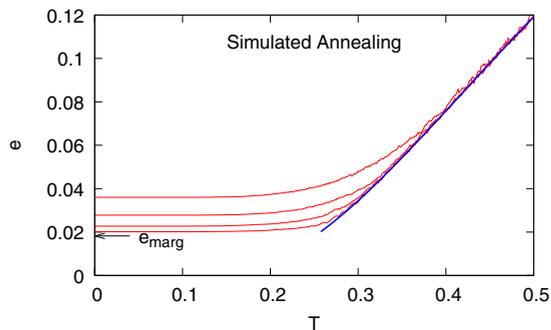


FIG. 2. Simulated annealing is not able solve the 3-regular 3-XORSAT problem, since it converges in the long time limit to a positive energy (close to the marginal energy e_{marg}). The four curves have annealing rates $\Delta T = 10^{-3}, 10^{-4}, 10^{-5}, 10^{-6}$ (from top to bottom). The blue (dark gray) curve represents the energy above the dynamical transition in the large N limit.

the marginal energy e_{marg} .) However the interesting question about the scaling of times to reach a solution requires a different analytic approach where fluctuations are taken into account. Most of our results are in the regime where times are made large while keeping N finite are based on numerical experiments.

The presentation of our results about optimization algorithms is somehow split in two parts. In Secs. III A and III B we discuss the regime of linear times where large sizes can be studied and analytical solutions in the large N limit can be obtained (thus estimating the threshold energy for various algorithms). In Secs. III C and III D we study the regime where times are made exponentially large in the system size N , and estimate the exponential growth rate of the time to reach a solution.

A. Simulated annealing: A warming up with the most widely used optimization algorithm

Simulated annealing (SA) is maybe the most widely used optimization algorithm. It consists in implementing a Monte Carlo Markov chain sampling from the Gibbs-Boltzmann distribution $P_{\text{GB}}(s) \propto \exp(-\mathcal{H}(s)/T)$ with a temperature T slowly decreasing towards zero. In Fig. 2 we report the results of Simulated Annealing run on samples of size $N = 10^5$ with a cooling schedule where the temperature is decreased by ΔT after each Monte Carlo sweep (MCS): the four curves corresponds to $\Delta T = 10^{-3}, 10^{-4}, 10^{-5}, 10^{-6}$ (from top to bottom).

From the simulated annealing results it is clear that an efficient, but still linear-time, algorithm is not able to solve the 3-regular 3-XORSAT problem, and seems to converge to configurations with a threshold energy close to e_{marg} . To achieve the zero energy configuration we need to use an algorithm that can go below the threshold energy, overcoming energetic and/or entropic barriers.

B. A broad class of stochastic search algorithms

Given a particular spin configuration, let us classify its variables according to the number of unsatisfied interactions

they belong to: we say a variable is of type k if it belongs to k unsatisfied interactions. In the present model $k \in [0, 3]$ since $d = 3$ for all variables. We call $f_k(t)$ the fraction of variables of type k at time t , that satisfy $0 \leq f_k(t) \leq 1$ and $\sum_{k=0}^3 f_k(t) = 1$ at any time.

The searching algorithm we propose is extremely simple and works as follows: at each time step it chooses one variable of type k with probability $p_k(t) \propto w_k f_k(t)$ and flips it. The time is then incremented by $1/N$ in order to have a well defined continuous process in the large N limit.

Starting from a random configuration the behavior of the algorithm is determined only by the vector of weights $\mathbf{w} = (w_0, w_1, w_2, w_3)$. The stopping condition depends on the weights: if all weights are non-null the algorithm never stops; while if $w_0 = 0$ the algorithm cannot flip variables participating only in satisfied interactions and thus any solution is a stopping configuration. We fix $w_0 = 0$ hereafter so as to make any solution a stopping configuration for the algorithm. We study several choices for the vector of weights (the weights need not be normalized, but the probabilities $p_k \propto w_k f_k$ are).

1. Analytic description

Before presenting the actual performance of this algorithm, we would like to stress that the evolution of the algorithm can be described analytically under some assumptions, which are similar to those already used in the literature to approximately describe the relaxation dynamics in model defined on a Bethe lattice [32,52,53]. In this section we derive a set of differential equations for $f_k(t)$, adapting the technique first introduced in [54].

At each time step the fractions $\{f_k(t)\}$ change depending on the variable chosen. For example, if a variable of type $k = 3$ is chosen and flipped, then that variable changes its type from 3 to 0 ($3 \rightarrow 0$), i.e., the fraction f_3 decreases by $1/N$ and the fraction f_0 increases by $1/N$; at the same time its 6 neighboring variables decrease by one the number of their types, and one needs to compute the number of changes $n_{3 \rightarrow 2}, n_{2 \rightarrow 1}$ and $n_{1 \rightarrow 0}$ in order to properly update the fractions $\{f_k\}$ (for example $\Delta f_2 = (n_{3 \rightarrow 2} - n_{2 \rightarrow 1})/N$). The three numbers $n_{3 \rightarrow 2}, n_{2 \rightarrow 1}$ and $n_{1 \rightarrow 0}$ are random variables distributed according to the multinomial distribution $\text{Mult}(\{p_{3 \rightarrow 2}, p_{2 \rightarrow 1}, p_{1 \rightarrow 0}\}, 6)$. Under the approximation that no correlation exists between the types of neighboring variables once the common interaction is removed (cavity approximation) we can compute the parameters of the multinomial distribution:

$$p_{3 \rightarrow 2} \propto 3f_3, \quad p_{2 \rightarrow 1} \propto 2f_2, \quad p_{1 \rightarrow 0} \propto f_1, \quad (4)$$

where the proportionality constant is fixed by $p_{3 \rightarrow 2} + p_{2 \rightarrow 1} + p_{1 \rightarrow 0} = 1$. One more example: if the variable chosen to be flipped is of type 2, then, apart from the change $2 \rightarrow 1$, we have that $n_{3 \rightarrow 2}, n_{2 \rightarrow 1}$ and $n_{1 \rightarrow 0}$ are distributed according to $\text{Mult}(\{p_{3 \rightarrow 2}, p_{2 \rightarrow 1}, p_{1 \rightarrow 0}\}, 4)$, while $n_{0 \rightarrow 1}, n_{1 \rightarrow 2}$ and $n_{2 \rightarrow 3}$ are distributed according to $\text{Mult}(\{p_{0 \rightarrow 1}, p_{1 \rightarrow 2}, p_{2 \rightarrow 3}\}, 2)$, where

$$p_{0 \rightarrow 1} \propto 3f_0, \quad p_{1 \rightarrow 2} \propto 2f_1, \quad p_{2 \rightarrow 3} \propto f_2, \quad (5)$$

with again the normalization condition $p_{0 \rightarrow 1} + p_{1 \rightarrow 2} + p_{2 \rightarrow 3} = 1$. So, the variations of the fractions due to a single

spin flip are random variables given by

$$\Delta f_k(t) = \frac{1}{N} [n_{k-1 \rightarrow k}^{(2(3-j))} - n_{k \rightarrow k+1}^{(2(3-j))} + n_{k+1 \rightarrow k}^{(2j)} - n_{k \rightarrow k-1}^{(2j)} + \delta_{j,3-k} - \delta_{j,k}] w / \text{prob. } p_j(t),$$

where the superscript (d) in $n^{(d)}$ refers to the total number of events in the multinomial distribution, and we fix $n_{-1 \rightarrow 0} = n_{0 \rightarrow -1} = n_{3 \rightarrow 4} = n_{4 \rightarrow 3} = 0$. The generalization of the above equation to other random graphs or interaction types is straightforward. Rescaling the time such that a spin flip happens every $\Delta t = 1/N$, and taking the average over a small, but finite, time interval, corresponding to $O(N)$ single spin flips, the stochastic equation above can be converted in the $N \rightarrow \infty$ limit to an ordinary first order differential equation in terms of mean values $\mathbb{E}[n_{k \rightarrow \ell}^{(d)}] = d p_{k \rightarrow \ell}$. For simplicity, we keep using the same notation, but now $\{f_i(t)\}$ are not single trajectories, but mean values over the trajectories:

$$f'_k(t) = \sum_{j=0}^3 p_j(t) [2(3-j)(p_{k-1 \rightarrow k} - p_{k \rightarrow k+1}) + 2j(p_{k+1 \rightarrow k} - p_{k \rightarrow k-1}) + \delta_{j,3-k} - \delta_{j,k}]. \quad (6)$$

An easy check is that $\sum_{k=0}^3 f'_k(t) = 0$, so the total probability is conserved. The solution to Eq. (6) can be easily achieved by any integration algorithm for ordinary differential equations. This analytic solution will be compared in the following with the actual evolution of the algorithm.

2. Actual performances

Let us start from a greedy version of the algorithm, that is an algorithm that can never increase the energy (it is the equivalent of gradient descent, but here the space of configurations is discrete, so gradients are not well defined). This greedy version of the algorithm requires $w_1 = 0$ because flipping a variable of type 1 would increase the energy (i.e., the number of unsatisfied interactions); only weights w_2 and w_3 can be non-null in the greedy case. For this greedy version the number of stopping configurations is very large (their entropy is computed in Appendix A): Any configuration where each variable participate to 0 or 1 unsatisfied interactions is a stopping configuration. These configurations are local minima of the energy function and we call them *blocked*.

In Fig. 3 we show the energy or fraction of unsatisfied constraints as a function of the running time of the greedy algorithm ($w_1 = 0$) for different choices of the ratio w_3/w_2 (the w 's are not normalized so only their ratio matters). Points are the numerical data and the lines are the solution to the ODE in Eq. (6). We notice the agreement between analytics and numerics is almost perfect, because at those energy values correlations are very weak. The same almost perfect agreement can be seen also at the level of fraction of variables f_k participating in k unsatisfied constraints (see Fig. 4 for the case $w_2 = w_3 = 1$).

Let us consider now non-greedy versions of this algorithm, that is cases with $w_1 > 0$ where the energy can sometimes increase during the evolution, although we expect on average the energy to relax to the threshold value.

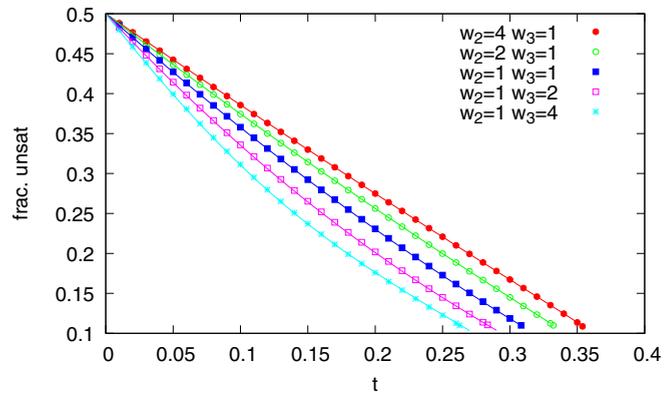


FIG. 3. Evolution of the energy for several greedy versions ($w_1 = 0$) of the algorithm. Points are numerical data and lines are the analytic description.

The first interesting case is the one that corresponds to WalkSAT [55]. A detailed description of this algorithm is given in Sec. III C; for the present purposes it is enough to say that in WalkSAT variables are flipped accordingly to the number of unsatisfied constraints they belong to, that is $w_k \propto k$. The numerical results and the comparison with the analytical description is provided in Fig. 5: the agreement is excellent. Although this version of the algorithm would run forever, after a finite time a stationary regime is reached and nothing interesting happens any more in the large N limit. Once the threshold energy has been reached, then the search for the solution takes place via rare fluctuations as discussed in the Sec. III C.

However, having achieved a good analytic description of the WalkSAT algorithm is not enough, because this algorithm works at very high energies: Its threshold energy is $1/4$ for 3-regular 3-XORSAT, more than twice the energy that the greedy version achieves and more than 10 times the marginal energy we expect to be the relevant threshold energy for smart searching algorithms.

So we are interested in studying more efficient versions of this algorithm and we explore the case where w_1 is much

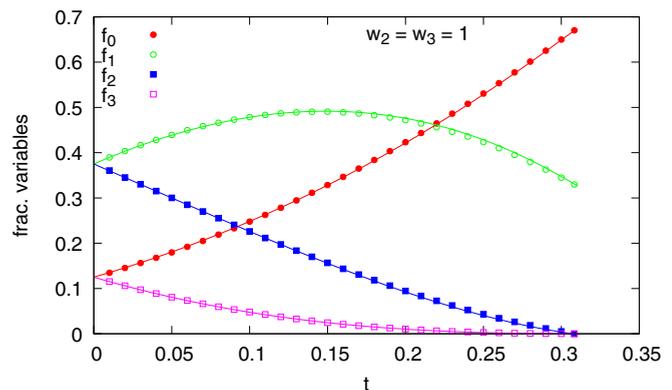


FIG. 4. The actual evolution of fraction of variables in a greedy version ($w_1 = 0$) of the algorithm (data points) are well described by the analytic solution (full curves). The algorithm stops when $f_2 = f_3 = 0$.

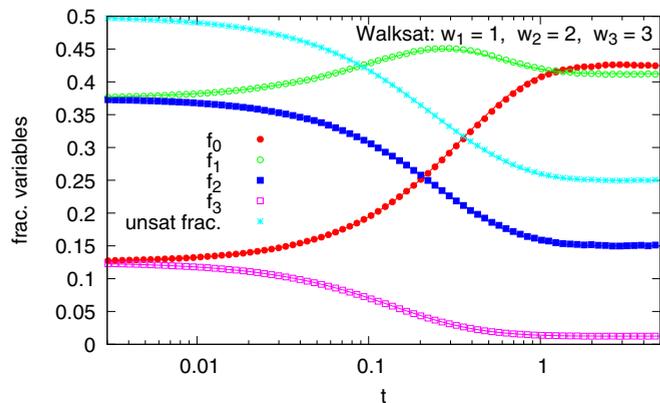


FIG. 5. Evolution of the energy and fraction of variables during the WalkSAT algorithm. The analytic description of the relaxation to the threshold energy is almost perfect.

smaller than w_2 and w_3 (in practice we fix $w_2 = w_3 = 1$ and explore the range $w_1 \ll 1$). Notice that if w_1 is made very small, the algorithm in practice is allowed to make an energy-increasing spin flip only when no other move is allowed, that is when the configuration is blocked. We have simulated also a version of the algorithm where this is made explicit and the results are equivalent.

In Fig. 6 we report the evolution of the energy for WalkSAT, the greedy version and several quasi-greedy versions of the algorithm. We observe that the threshold energy is $1/4$ for WalkSAT, close to 0.1 for the greedy version, but approaches very closely the marginal energy e_{marg} for $w_1 \ll 1$. So the quasi-greedy versions of this algorithm seems very effective in reaching the same energy values simulated annealing can achieve.

The analytic description of the algorithm changes a lot depending on the energy value. For high enough energies, correlations are very weak and the analytic solution matches perfectly the numerical data: this is true for WalkSAT and the greedy version as already shown in Figs. 3, 4, and 5. For

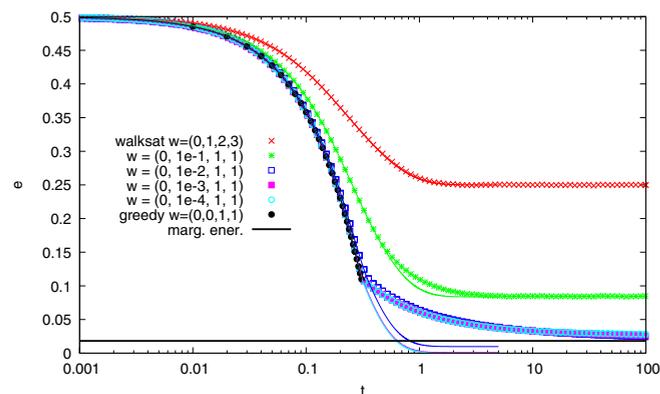


FIG. 6. Evolution of the energy for WalkSAT and several (quasi-)greedy versions of the algorithm. The vector of weights (w_0, w_1, w_2, w_3) appears in the legend. The analytic description under the assumption of lack of correlations is reported with a full curve and fails to describe the algorithms at the lowest energies, when correlations arise.

$w_1 = 0.1$ we start seeing some deviations at times $t \sim 1$, that luckily enough disappear for larger times.

The most interesting cases are those with $w_1 \leq 10^{-2}$. First of all we observe a very weak dependence on w_1 (the three values used $w_1 = 10^{-2}, 10^{-3}, 10^{-4}$ give practically the same results), so we are working in the limit where the algorithm does an energy-increasing spin flip only when no other move is available. So this algorithm is not able to jump over barriers of height larger than 1 (or few units at most), that correspond to $\Delta e = O(1/N)$.

Indeed this quasi-greedy algorithm follows closely the greedy version until the time, $t^* \sim 0.3$, when the latter reaches a blocked configuration that stops it. For $t > t^*$ the quasi-greedy algorithm keeps decreasing the energy at a much slower pace, eventually approaching an energy very close to the marginal one (that we expect to be a lower bound for the threshold energy in linear time algorithms).

In the regime $t > t^*$ the analytical approximation based on the lack of correlations fails dramatically in many important aspects: it keeps predicting a very fast energy decrease and it estimates a too small asymptotic energy. The reason for such a failure is clear: Low energy configurations are very peculiar and have strong correlations among variables, even among variables, which are far apart. In this energy range, our approximation fails and the analytic solution is meaningless.

It is worth stressing that we are still working on timescales growing linearly with N (the first regime depicted in Fig. 1). So the failure of the analytic approximation is not due to the fluctuations that become important in the second regime of exponentially large timescales. Here we are still working in a regime such that every spin variable has been flipped a finite number of times. Nonetheless the evolution of the quasi-greedy algorithm bring the system in configurations correlated up to some distance, such that hypothesis leading to Eq. (6) are no longer valid.

From a preliminary study we have understood that in this energy regime the algorithm proceeds by performing collective rearrangements of variables forming a local tree-like structure, and the size of these collective rearrangements seems to grow while approaching the threshold energy (this reminds us a lot what happens in similar problems approaching the dynamical transition [56,57]). The analytical description of the quasi-greedy algorithm in this low energy regime is deferred to a future work.

We move now to the problem of estimating the large deviation rate to reach a solution via a rare fluctuation from the threshold energy. Given the differences in the threshold energies clearly visible in Fig. 6 we expect quite different rates for WalkSAT and the several quasi-greedy versions of the algorithm.

C. WalkSAT: The large deviations rate to reach a solution from a simple argument

WalkSAT is a popular randomized (or stochastic) algorithm for solving constraint satisfaction problems [55]. In its simplest version works as follows: starts from a random configuration; at each time step, if all constraints are satisfied returns the solution, otherwise picks uniformly at random an unsatisfied constraint and flips a randomly chosen variable

participating that constraint. The flip certainly satisfies the chosen constraint (interaction) but may unsatisfy many other constraints: so the algorithm may increase the energy during the evolution.

An approximated analytic description of this algorithm exists [2,52,58] from which we learn that for the hard problem we are studying the algorithm relaxes to a positive energy in a linear time and then reaches a solution by a rare fluctuation taking an exponential time. However it is not clear which kind of barrier the algorithm crosses, given that the energy increase is in principle without bounds.

Actually the analytic description of the WalkSAT algorithm can be made even simpler than in previous approaches [2,52,58] by noticing that (i) the algorithm stays most of the time on configurations of very high energy where correlations are very weak and (ii) the fluctuation leading to a solution is so rapid that correlations do not arise.

These insights suggest that keeping track of the energy alone should capture the essential features of the algorithm. In practice, we think of the energy as a stochastic process, and write down a stochastic differential equation that describes its dynamics during the computation. The details of the calculation are presented in the following section.

1. Analytic description

In this section we focus on random p -XORSAT, with α constraints per spin: every spin participates in a Poisson distributed number of interactions (with average and variance $p\alpha$), while every interaction involves exactly p spins. Random p -XORSAT is known to be less correlated than the regular version [see the comment after Eq. (3)], so we expect it to be more amenable to a simple description.

Consider a uniformly random spin configuration s . Every spin is connected to a Poisson distributed number of broken constraints with average $p\alpha_u$, where $\alpha_u = \mathcal{H}[s]/N$. The same spin is also connected to a Poisson number of satisfied constraints with average $p(\alpha - \alpha_u)$. The fundamental idea of our approach is to assume that WalkSAT is incapable of building any correlations that violate this property, which is strictly true only on fully random configurations. This is reasonable, as most of the time is spent in high-energy states.

Given an initial random configuration s , we run WalkSAT for T steps and denote the number of broken constraints at this time $\mathcal{H}(T)$. When we take one more step, $\mathcal{H}(T)$ changes by

$$\Delta\mathcal{H}_T^{T+1} = -1 - u(T) + s(T), \quad (7)$$

where $u(T)$ is the number of *excess* broken constraints connected to the spin (i.e., excluding the one that was selected by WalkSAT) and $s(T)$ is the number of satisfied constraints connected to it. As explained earlier, we assume that at all times they are distributed as if the configuration was random. The simplicity of this expression is a peculiarity of XORSAT: Every time one variable is flipped, all the clauses connected to it change state (satisfied ones become broken and vice versa). The same approach could be adapted to other rCSP, with an appropriate choice of the energy increment.

After a number ΔT of steps, larger than one but small compared to the number of spins N , the total energy

change is

$$\Delta\mathcal{H}_T^{T+\Delta T} = -\Delta T - \sum_{k=0}^{\Delta T-1} u(T) + \sum_{k=0}^{\Delta T-1} s(T), \quad (8)$$

and we approximate the two sums using the central limit theorem:

$$\begin{aligned} \sum_{k=0}^{\Delta T-1} u(T) &\approx p\alpha_u\Delta T + R_1\sqrt{p\alpha_u\Delta T}, \\ \sum_{k=0}^{\Delta T-1} s(T) &\approx p(\alpha - \alpha_u)\Delta T + R_2\sqrt{p(\alpha - \alpha_u)\Delta T}, \end{aligned} \quad (9)$$

where R_1 and R_2 are two independent standard Gaussian random variables. In these expressions we assume ΔT is short enough that the energy density α_u does not appreciably change over the interval ΔT . The sum of the two terms involving R_1 and R_2 is again a Gaussian variable, so that the total energy change after ΔT steps is approximately

$$\Delta\mathcal{H}_T^{T+\Delta T} = -\Delta T + p(\alpha - 2\alpha_u)\Delta T + \sqrt{p\alpha}R\sqrt{\Delta T}. \quad (10)$$

We are interested in the large N limit, so it is convenient to change variable from T to $t = T/N$. Dividing the previous equation by N we obtain a stochastic differential equation describing an Ornstein-Uhlenbeck process:

$$d\alpha_u(t) = 2p\left(\frac{p\alpha - 1}{2p} - \alpha_u\right)dt + \sqrt{\frac{p\alpha}{N}}dW. \quad (11)$$

To estimate the scaling with N of the time necessary to reach a solution, it is best to study the Fokker-Planck equation associated to Eq. (11):

$$\frac{\partial P}{\partial t} = -\frac{\partial}{\partial \alpha_u}\left(2p\left(\frac{p\alpha - 1}{2p} - \alpha_u\right)P\right) + \frac{1}{2}\frac{p\alpha}{N}\frac{\partial^2 P}{\partial \alpha_u^2}. \quad (12)$$

Regardless of the boundary condition $P(\alpha_u, t = 0)$, after an initial transient, the energy is distributed according to the stationary solution of Eq. (12) (the normalization constant c_N is irrelevant here)

$$P_{\text{st}}(\alpha_u) = c_N \exp\left(-2N\frac{(\alpha_u - \alpha_0)^2}{\alpha}\right), \quad (13)$$

which is centered around the finite value (if $p\alpha > 1$)

$$\alpha_0 = \frac{p\alpha - 1}{2p}. \quad (14)$$

This indicates that the time necessary to reach the solution grows more rapidly than $O(N)$, but also gives us a concrete way to estimate it: the typical time over which a transition from energy α_0 to zero happens is the reciprocal of the rate

$$\begin{aligned} t_{\text{sol}} &\sim \frac{1}{\text{Prob}(\alpha_u = 0, t \rightarrow \infty | \alpha_u = \alpha_0, t = 0)} \\ &= \exp\left(N\frac{(p\alpha - 1)^2}{2\alpha p^2}\right). \end{aligned} \quad (15)$$

As expected, the solution time scales exponentially with the problem size. This analytical solution does show, in its simplicity, the mechanism by which a solution is found and why we call it *entropic barriers*. At equilibrium, reached after

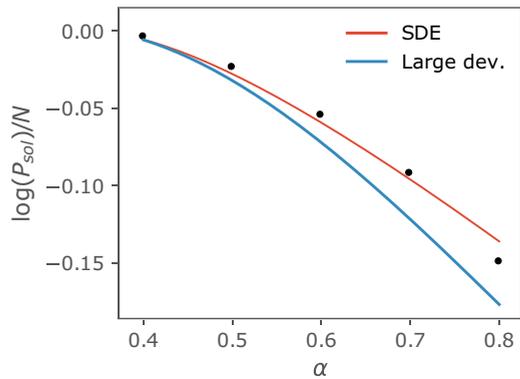


FIG. 7. Comparison to the numerics for 3-XORSAT. The vertical axis is the logarithm of the probability of solution in linear time divided by N . The black dots are the numerical data reported in [52] and the blue (lower) line is the analytic approximation derived in that reference, based on large deviation theory. The red (upper) line is our approximation [see Eq.(15)].

$O(N)$ steps, the system fluctuates around a very high energy threshold α_0 . This is well above any local minimum where the system could be stuck. With exponentially small probability however a rapid fluctuation of $O(N)$ spin flips can bring the system to the real solution, but this happens with exponentially small probability. So, only an exponentially small fraction of all the possible sequences of N spin flips will pick the right direction to the solution, hence the problem is of an entropic nature, not an energetic one. We will discuss this in more details after comparing with numerical results.

2. Comparison with numerics

A comparison with the actual rate measured in numerical experiments in Ref. [52] is provided in Fig. 7. The quality of the rate reported in Eq. (15) and obtained under the assumption of lack of any correlation is surprising and supports the idea that WalkSAT makes a random search without building any relevant correlation in the problem.

It is interesting to notice that the same simple argument can be made also in the 3-regular 3-XORSAT case and provides the following scaling for the time to a solution

$$t_{\text{sol}} \sim \exp(N/6). \quad (16)$$

The rate $\mu = 1/6 \simeq 0.167$ is again close to the numerics reported in Ref. [59], where $\mu \approx 0.124$ was measured.

Notice that the physical interpretation of this calculation is that the equilibrium distribution is centered around a *relatively large* nonzero value of the cost function, as seen in Fig. 8. The solution is found by one rare fluctuation, which brings the system out of equilibrium by a “lucky” series of $O(N)$ flips, which points in the exact direction to the ground state. During this series, no significant correlation is created between the values of the various random variables so the central limit theorem we used is valid. We expect this to be the behavior of a random algorithm, which works at high temperature, tackling a problem with entropic barriers.

In the analysis of more complex algorithms, the ones that work at lower temperatures, this assumption is most probably violated. Correlations between the random variables need to

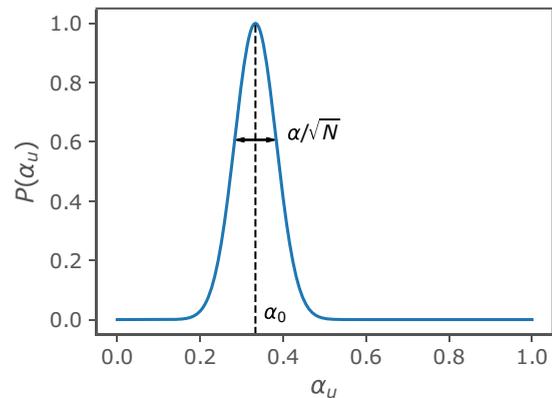


FIG. 8. Equilibrium distribution of the density of UNSAT clauses α_u for the Fokker-Planck equation (12). A solution is found when a rare fluctuation reaches the value $\alpha_u = 0$, which occurs with probability exponentially small in N .

be taken into account when computing the rate of the rare fluctuation that will lead to finding of the ground state. If one wants to make progress in the analysis of such algorithms one needs to go beyond what done in this paper.

D. Times to reach a solution via the quasi-greedy algorithm

Given that the quasi-greedy version of the algorithm has a much lower threshold energy than the WalkSAT algorithm we expect rare fluctuations leading to a solution to happen with a better exponential rate. Unfortunately in the quasi-greedy case the computation can not be done analytically because, as seen above, the approximation based on the lack of correlations provides very poor results and so it is not useful at all. We thus resort to a numerical computation.

We have observed above that in the limit $w_1 \ll 1$ the dependence on w_1 is extremely weak in the first regime of linear times. The same is true also for the second regime of exponential times. The exponential rate μ determining the mean time to reach a solution $t_{\text{sol}} \sim \exp(\mu N)$ does not show any visible dependence on w_1 . So, we have fixed $w_1 = 0.05$ to carry on our numerical experiments with the quasi-greedy version of the algorithm.

In Fig. 9 we report with blue points the average time to reach a solution running the quasi-greedy algorithm. The time we report is the mean number of sweeps, where a sweep corresponds to N spin flips. The average is taken over 10^6 different runs and the error is smaller than the symbol size. For large N the data is very well fitted by an exponential growth with rate $\mu_{\text{QG}} = 0.0835(5)$.

In the same figure (Fig. 9) we report also the mean time to reach a solution by the WalkSAT algorithm. It is clear that the rate is much larger, although for these sizes we are still observing the preasymptotic behavior and for larger sizes the rate will converge to $\mu_{\text{WS}} \approx 0.124$ [59].

A much more meaningful comparison is the one with the parallel tempering (PT) algorithm. PT is considered the state-of-the-art for thermalization (and optimization) in the field of disordered systems, as it can deal efficiently with very rough energy landscapes. For spin glass models, which are similar to the problem we are studying here, it has been recently shown

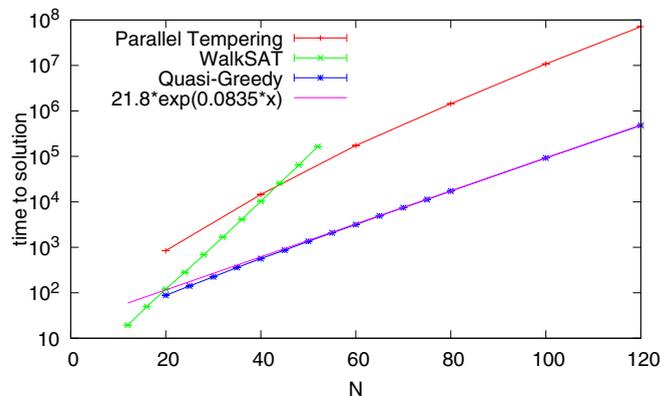


FIG. 9. Mean times to reach a solution measured as spin flips per variable. The exponential growth of the time to reach a solution by the quasi-greedy version of the algorithm has a rate $\mu_{\text{QG}} \approx 0.0835$, much smaller than the one of WalkSAT and comparable to the one of “parallel tempering” (which is however much slower because of the prefactor).

that PT performs comparably to population annealing (PA) [60] another standard algorithm to optimize complex energy functions. We prefer to study the performance of PT in finding a solution, rather than PA, because in PT it is straightforward to compute the time to solution once the temperature scheduling is fixed. While in PA one needs to increase the size of the population during the run and if the solution is not found a new run with larger sizes should be done; moreover comparing population sizes and running times requires some extra care.

We have run PT with an optimal temperatures scheduling derived in Appendix B. The results in terms of Monte Carlo sweeps per replica are shown in Fig. 9. The comparison with the quasi-greedy algorithm is very favorable for the latter: the growth rate for PT is slightly larger than μ_{QG} (but still decreasing and could eventually converge to the same value) and the prefactor for PT is larger by 2 orders of magnitude than the one for the quasi-greedy algorithm.

E. Quantum annealing

Originally proposed twenty years ago [17,61], quantum annealing is a quantum algorithm designed to solve *classical* optimization problems, exploiting the adiabatic theorem of quantum mechanics [62,63]. First, we encode the given classical problem in a *problem Hamiltonian* H_P , so that solutions correspond to its ground states. The problem Hamiltonian is typically in the form of a cost function with all terms commuting among themselves (in this sense it is a classical problem). A convenient and customary choice is to have the problem Hamiltonian be diagonal on the σ^z basis. Then, we choose a *fluctuation Hamiltonian* H_F , an arbitrary operator that should provide “quantum fluctuations” and must have a known and simple ground state. A popular choice is to use a uniform field in the σ^x direction to provide fluctuations:

$$H_F = \sum_i \sigma_i^x, \quad (17)$$

where the sum runs over all spin variables in the system. The quantum annealing algorithm consists then in time-evolving

the ground state of

$$H(t) \equiv \frac{t}{T} H_P + \left(1 - \frac{t}{T}\right) H_F \quad (18)$$

from $t = 0$ to $t = T$. For a long enough annealing time T , the adiabatic theorem guarantees that a system initially prepared in the ground state of $H(0) = H_F$, known by construction, will evolve into a state belonging to the ground state manifold of $H(T) = H_P$. Measuring this state on the σ^z basis we obtain a solution to the original problem. Notice that choosing the fluctuation Hamiltonian as in Eq. (17) guarantees the initial ground state has finite overlap with every state in the computational basis; the algorithm will not miss a solution only because the corresponding state had no initial amplitude. The adiabatic theorem also provides a lower bound on T : for the algorithm to succeed, the annealing time should be longer than

$$T \gg \frac{\max_t \langle \psi_1(t) | \partial_s H(s) | \psi_0(t) \rangle}{\min_t \Delta^2(t)} \quad s \equiv \frac{t}{T}, \quad (19)$$

where $|\psi_0(t)\rangle$ is the instantaneous ground state at time t , $|\psi_1(t)\rangle$ is the first excited state, and $\Delta(t)$ is the energy gap between them. We have already encountered a classical Hamiltonian encoding XORSAT in Eq. (3). The quantum mechanical version is simply

$$H_P = \frac{1}{2} \left(N - \sum_{a=1}^N \prod_{i \in \partial a} \sigma_i^z \right). \quad (20)$$

In our numerics, we use this as the problem Hamiltonian and a uniform transverse field as the fluctuation one. Once we fix the annealing time T and initial state $|\psi(0)\rangle$, QA will end up in some final state $|\psi(T)\rangle$. Let the probability of measuring an energy equal to the ground state energy, which in our model is $E_{\text{GS}} = 0$, be ε in this state. As T increases, ε will approach unity. The complexity of the algorithm is not expected to change with ε , as long as $\varepsilon = O(1)$. In fact, since ε is the probability to find the ground state after a time T , one can enhance this probability by repetition of the algorithm. A success probability of ε after two repetitions becomes $1 - (1 - \varepsilon)^2 = 2\varepsilon - \varepsilon^2$, and even a small ε , after $n = O(1/\varepsilon)$ repetitions can be made close to 1. This repetition, or restart technique, is commonly used in algorithms for CSP [2], like WalkSAT. One commonly adopted definition for the time complexity of QA is to invert the relation between ε and annealing time: one fixes ε and N , and asks what is the corresponding annealing time $T(N, \varepsilon)$. For large enough N it is reasonable to expect an exponential scaling of the form

$$T(N, \varepsilon) \sim A(N, \varepsilon) \exp(\mu N), \quad (21)$$

where the prefactor $A(N, \varepsilon)$ is allowed to have a polynomial dependence on N , while the rate μ does not depend on ε . In practice, $T(N, \varepsilon)$ is rarely estimated from real-time unitary evolution, since the required computational power quickly becomes unmanageable as the number of spins grows above $N = 20$. There are two common strategies adopted to sidestep this problem: estimate the energy gap between the ground and first excited state via exact diagonalization or quantum Monte Carlo [19,47,48] and invoke the adiabatic theorem to impose a bound on $T(N, \varepsilon)$, or perform the evolution in imaginary time assuming that the scaling of the imaginary analog of

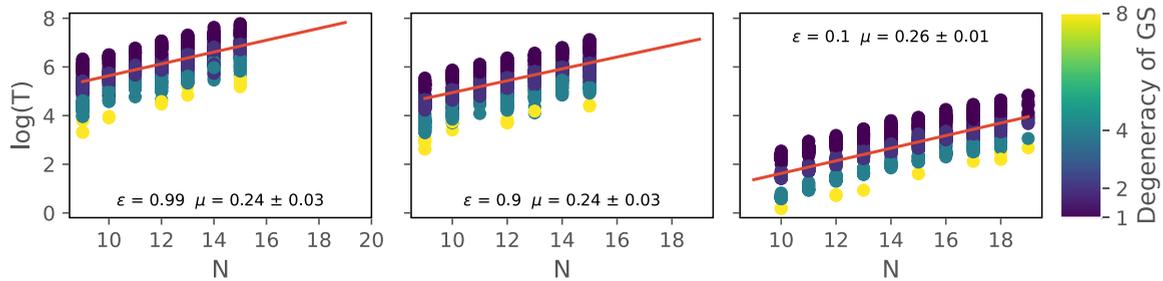


FIG. 10. Annealing time T as a function of instance size N and solution probability ε . Every dot corresponds to an instance, with darker color indicating a more degenerate ground state. Notice how the rate μ is insensitive to the target probability, as expected from the restart argument given in the main text. The red line is a linear fit performed on the median value of the annealing time, using the standard deviation of $\log(T)$ as error bar.

$T(N, \varepsilon)$ will be the same as the original quantity [18]. Both methods find exponential scaling of the time for hard classical problems, a fact that has been connected with the order of the thermodynamic transition [47] (however, there are models exhibiting a first order transition but only an algebraically small gap [64]). While both methods provide reasonable bounds on the scaling of the annealing time, it is unclear how the entropy of excited states affects those estimates. Since entropic effects are the primary focus of this work, we decided to perform real-time evolution to minimize confounding factors, even if this means limiting the simulations to moderate sizes. The role of entropic effects in similar problems have been discussed in the past, see for example [65] and [66]. We integrated the time-dependent Schrödinger equation via an explicit high-order adaptive Runge-Kutta method [67] implemented in the QuSpin library [68].

The results are presented in Fig. 10: it is clear that instances with a unique solution (dark-blue dots) are harder than instances with multiple solutions (lighter-blue dots), but not exponentially so. For any choice of the solution probability ε , the growth rate μ is compatible with

$$\mu = 0.25 \pm 0.07 \quad (22)$$

This value is higher than the one predicted in [19], that in our notation would read $\mu \simeq 0.167$. Although the difference is only one standard deviation and could be easily be due to statistics (or the more limited fitting region $N \in [12, 18]$), one should keep in mind that the simulations in that reference are run only on instances with a unique solution, and use a Quantum Monte Carlo method to estimate the gap, which is then connected to μ by the adiabatic theorem, upon making some assumption on the matrix elements of local operators. However one should point out the observed insensitivity of the rate μ to the degeneracy of the ground state: if one posits to find the solution of the problem with an $O(1)$ probability as N grows, the same rate of growth of the solution time is found for problems with one solution (for which one can apply the adiabatic theorem connecting the gap with the time, as in previous studies [69]) and those with many solutions, for which most of the action occurs at finite energy density, where entropy dominates. Again we find a trade-off between entropic effects at high energy density and energy barriers effects close to the ground state. We can also recognize here a striking parallel with the difficulties attributed to localization

phenomena [20], where many small avoided crossings occur between states, which are $O(N)$ spin flip apart.

IV. CONCLUSIONS

In Figs. 11 and 12 we compare the spectrum of a typical instance with unique solution to the spectrum of an instance with eightfold degenerate ground state, as a function of the transverse field. When the ground state is unique (Fig. 11), one of the states in the first excited manifold peels off and has an avoided crossing with the ground state. The performance of the quantum adiabatic algorithm is limited by the size of this gap. When the ground state is degenerate (Fig. 12) the picture is qualitatively different: as the transverse field is turned off, a few excited states go through a cascade of avoided crossings, eventually ending in the ground state manifold. In this case, knowledge of the minimum gap between the ground state and the first excited state is not enough to understand the scaling of the annealing time: we should know how many states end up in the ground state, and what avoided crossings those have. Typically, these states will start high in the spectrum, where there will be many crossings, and so where entropy is important. These instances are not naturally described by the picture of trapped states, which tunnel to the solution (like a single level crossing) and we cannot avoid noticing the similarity with the phenomenon of entropic barriers, in the sense that

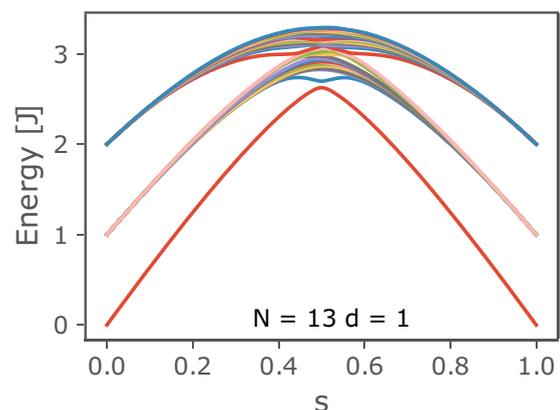


FIG. 11. Lowest 30 levels in the spectrum as a function of transverse field strength t/T for an instance with nondegenerate ground state. Compare to Fig. 20 in [48].

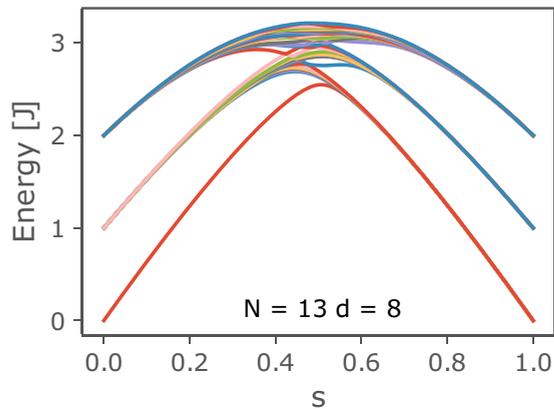


FIG. 12. Lowest 30 levels in the spectrum as a function of transverse field strength t/T for an instance with eightfold degenerate ground state.

only a “lucky” sequence of flips from one state to the other can lead to ground state manifold. And this, too, occurs with exponentially small probability.

Last, we must notice that the threshold energy for this problem is $e_{\text{marg}} \simeq 0.02$, so a more quantitative understanding of the difference between excited states below the threshold energy and above it in the context of quantum annealing requires studying instances with $N \simeq 100$ or larger. For the small sizes we can currently study via the exact integration of the Schrödinger equation the ground state degeneracy seems the best indicator for identifying hard instances.

For rCSP problems, some algorithms encounter entropic barriers rather than energetic barriers. This in particular happens when the algorithm works at energies higher than the threshold energies for the given problem and they find the ground state by a rare fluctuation, which picks the right sequence of $O(N)$ spin-flips among the exponentially many. We have shown this explicitly with a family of almost-greedy algorithms, which includes WalkSAT. These algorithms evolve in $t = O(N)$ to their equilibrium state (this part of the evolution can be followed by a system of coupled differential equations), and then must benefit from a rare stochastic fluctuation to find the ground state. In the case of WalkSAT we found the fluctuation rate—and hence the time to find a solution—with a simple Brownian motion (and a corresponding Fokker-Planck) analysis of the algorithm. This was possible exactly because the algorithm’s equilibrium state is at high energy, so the region where strong correlations exists between the variables is not explored.

We have also introduced a new class of quasi-greedy algorithms, which are not able to jump over large energy barrier. We have shown that algorithms belonging to this class are very effective in reaching the threshold energy in times $O(N)$ and also in reaching the ground state by rare fluctuations in times $O(\exp(\mu N))$, with a value for the μ exponent not larger than the one exhibited by state of the art algorithms. Given that these algorithms are not able to jump over large energy barriers, they provide a direct evidence that the most effective search for solutions in the hard XORSAT problems is limited mainly by entropic barriers. As it is likely to happen in many other hard optimization problems.

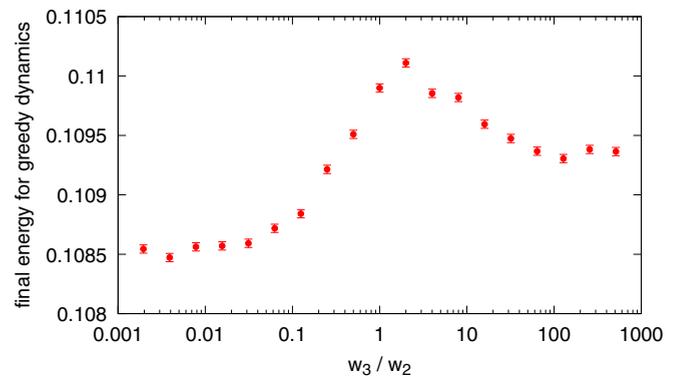


FIG. 13. Mean energy of the stopping configurations for the greedy dynamics as a function of the ratio w_3/w_2 .

Very recently, the role of the entropic barriers in the XORSAT problem has been thoroughly discussed in Ref. [70]. In the search for the planted solution, the authors compare the free-energy barrier that a thermal algorithm needs to overcome and the entropic barrier that limits the evolution of a microcanonical algorithm. They find that the free-energy barrier is just slightly smaller than the entropic one, thus supporting the idea that the entropic contribution is the leading one.

One more evidence that the entropic barriers are the leading one in the XORSAT problem comes from the effectiveness of the class of quasi-greedy algorithms we have introduced in this work. Indeed an algorithm from this class of quasi-greedy searches has been used in a challenge, whose participants were asked to find the ground state of a 3-regular 3-XORSAT problem [71]. A very optimized version of the quasi-greedy algorithm, running on Nvidia VT100 GPU, has rank first in the challenge [72], supporting our claim that the quasi-greedy algorithm introduced here are very effective, as the actual barriers are entropic in nature.

We have shown, by a time integration of the time-dependent Schrödinger equation, that a similar situation is encountered by quantum annealing algorithms.

Therefore we conclude that, in situations in which algorithms work in regions of the phase space in which trapping in local minima is not a real problem, the real problem is entropic barriers, which get for themselves the task of making the solution time exponential. The fact that two such phenomena, apparently so different from each other, can trade places and make sure $P \neq NP$ is a fascinating, and we believe not widely appreciated aspect of complexity theory. That quantum algorithms might suffer from a similar trade-off is, if possible, even more surprising.

ACKNOWLEDGMENTS

This work was supported by the European Research Council under the European Union’s Horizon 2020 research and innovation program (Grant No. 694925-Lotglassy). We warmly thank Guilhem Semerijan for the careful reading and suggestions.

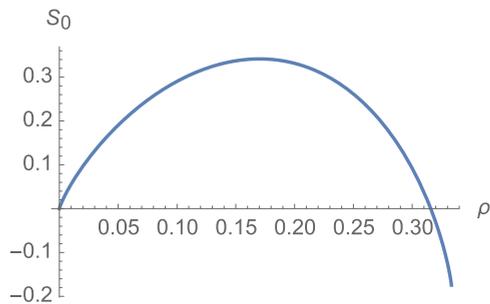


FIG. 14. Entropy S_0 of blocked configurations as a function of the energy ρ .

APPENDIX A: COUNTING BLOCKED CONFIGURATIONS AND THE SIZE OF THEIR BASINS OF ATTRACTION

The greedy version of the algorithm we have introduced has many possible stopping configurations. Indeed any configuration where each variable participate to at most one violated interaction is a local minimum of the energy function and thus can block the greedy dynamics. We call these configurations *blocked*, because by seeing the dynamics as the evolution of energy defects that can just annihilate in pairs, in a blocked configuration energy defects cannot evolve since they are isolated.

We start from the observation that all greedy dynamics tend to a threshold energy around 0.11; more precisely in Fig. 13 we plot the energy of the stopping configurations for several greedy dynamics as a function of the ratio w_3/w_2 (the only degree of freedom of the algorithm once we fix $w_0 = w_1 = 0$).

It is thus natural to ask whether it is possible to relate the stopping energy of the greedy dynamics to the entropy of blocked configurations. In order to compute the latter we use the replica symmetric cavity method, that is the Bethe approximation.

We consider the dual lattice to a 3-regular random 3-hypergraph, which is again a 3-regular random 3-hypergraph, where the variables are now the constraints and we assign variables $n_i \in \{0, 1\}$ indicating whether a constraint is satisfied ($n_i = 0$) or not ($n_i = 1$). The interaction among triplets of constraints (those shared by a variable in the original model) forbids any configuration with more than one violated constraint per triplet ($n_1 + n_2 + n_3 \leq 1$).

The energy of a spin configuration now corresponds to the density of variables $n_i = 1$ that we call ρ . Under the Bethe approximation the joint probability distribution can be factorized as follows:

$$P(\{n_i\}) \simeq \prod_{i=1}^N p_1(n_i) \prod_{(ijk)} \frac{p_3(n_i, n_j, n_k)}{p_1(n_i)p_1(n_j)p_1(n_k)} \quad (A1)$$

where the second product is over the N randomly chosen triplets forming the 3-regular 3-hypergraph. Given that the random hypergraph is regular we can assume the one-particle and 3-particles marginal probabilities, $p_1(n_i)$ and

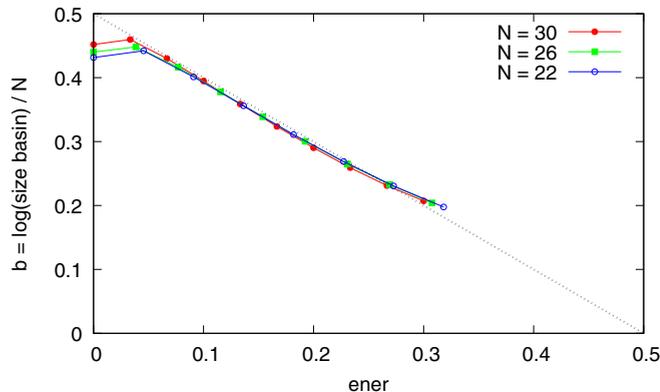


FIG. 15. The size of the basin of attraction grows exponentially with the size and the depth of the energy minimum. Finite size effects are evident only for the lowest energies.

$p_3(n_i, n_j, n_k)$, being site independent. These can be written explicitly in terms of the energy ρ as

$$p_1(n) = \begin{cases} 1 - \rho & \text{if } n = 0 \\ \rho & \text{if } n = 1 \end{cases} \quad (A2)$$

$$p_3(n_1, n_2, n_3) = \begin{cases} 1 - 3\rho & \text{if } n_1 + n_2 + n_3 = 0 \\ \rho & \text{if } n_1 + n_2 + n_3 = 1 \end{cases} \quad (A3)$$

From the Bethe approximation in Eq. (A1) the entropy of blocked configurations can be written as

$$\begin{aligned} S_0 &= - \sum_{n_1, n_2, n_3} p_3(n_1, n_2, n_3) \log p_3(n_1, n_2, n_3) \\ &\quad + 2 \sum_n p_1(n) \log p_1(n) \\ &= -(1 - 3\rho) \log(1 - 3\rho) - \rho \log \rho \\ &\quad + 2(1 - \rho) \log(1 - \rho) \end{aligned} \quad (A4)$$

which is plotted in Fig. 14.

S_0 is non-negative for $\rho \leq 0.315742$ and has a maximum in $\rho^* = 0.170209$. The threshold energy for the greedy algorithms we have studied is in the range $[0, \rho^*]$, but we cannot estimate it from S_0 alone, because each blocked configuration has a basin of attraction whose size matters as much as the entropy of blocked configurations.

Given the important role of the basins of attraction in predicting the large time limit of relaxation dynamics we have measure their sizes in problems of small size, $N \leq 30$, where an exact enumeration can be performed. For each blocked configuration we have measured the size of the basin of attraction as the number of initial conditions that a greedy dynamics brings to that blocked configuration. These sizes are in general exponentially large in N and thus we define $b = \log(\text{size of basin})/N$.

We report in Fig. 15 the results for b as a function of the energy of the blocked configuration for different sizes. The data have been averaged over different samples and different blocked configuration at fixed energy. It is remarkable that even for such small sizes the data show rather weak size dependence and are thus reliable.

The dotted line in the figure is just a guide for the eyes to convince the reader that the observed $b(e)$ is not far from

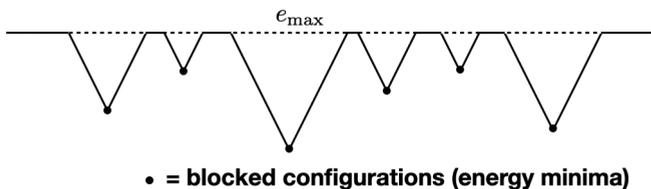


FIG. 16. A schematic picture of the energy landscape in the 3-regular 3-XORSAT problem. Most energy minima (including the solutions to the problem) have a basin of attraction, which is likely to be connected to the energy level e_{\max} where the dynamics starts. However the entropic barrier makes very unlikely to choose the pit leading to a solution.

following a linear behavior up to the energy of most numerous configurations $e_{\max} = 1/2$. According to this linear behavior the size of a basin of attraction depends linearly on the depth of the energy minimum (the blocked configuration). The simplest picture compatible with these data is the one schematically represented in Fig. 16 where each energy minimum corresponds to a pit whose edge is close to e_{\max} , such that the size of the pit grows exponentially with its depth, that is the distance from the edge.

According to the simplified picture in Fig. 16 every pit is in principle accessible from the initial configuration (which has typically an energy e_{\max}). However those leading to a solution correspond to a tiny minority of the configurations at e_{\max} and thus is extremely unlikely that the dynamics enter one of these (this is essentially the origin of the entropic barrier).

Combining the number of energy minima $\exp(NS_0)$ with their size $\exp(Nb)$ we can compute the large deviation rate to find one of the energy minima as a function of their energy. This is shown in Fig. 17, where indeed the maximum is achieved at $e \sim 0.11$ (marked by a vertical dotted line), which is the typical energy reached by the greedy algorithm in the large N limit. For finite values of N these data also provide the exponential rate to find a solution via the greedy

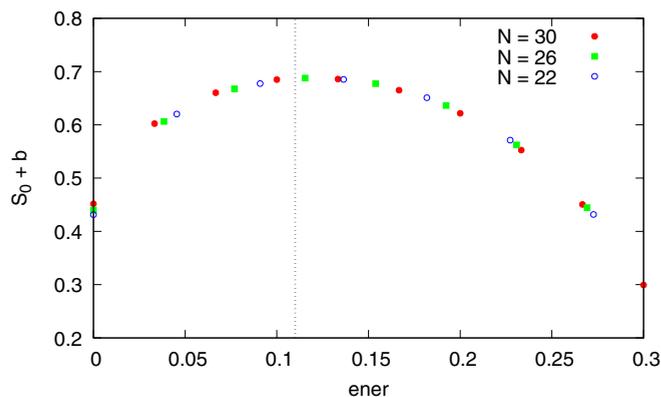


FIG. 17. Combining the entropy of blocked configurations S_0 and the log size of the basins of attractions b we can compute the rate of the large deviation probability of finding a block configuration of a given energy [shifted by $\log(2)$ because it is not normalized]. The maximum is achieved at an energy close to 0.11 (marked by a vertical-dotted line) where the greedy algorithm converges in the large N limit.

algorithm: This is just the probability of randomly choosing one of the pits having the bottom at $e = 0$ and corresponds to $\mu = \max_e (S_0 + b)|_{e=0} \approx 0.2$.

APPENDIX B: AN OPTIMAL TEMPERATURE SCHEDULING FOR PARALLEL TEMPERING

The main problem in setting up an optimized PT is the choice of the temperature set. However in the present model we are in a lucky situation because we know that the 3-regular random 3-XORSAT has no thermodynamical phase transition and so, after convergence, the energy sampled by the PT must be the one of the paramagnetic solution, $e(\beta) = [1 - \tanh(\beta/2)]/2$. In the large N limit we can assume that the extensive energy at inverse temperature β is a Gaussian variables with mean $E(\beta) = Ne(\beta)$ and variance $\sigma^2(\beta) = -Ne'(\beta)$. This Gaussianity assumption (which is rather well satisfied, but in the vicinity of the ground state) allows us to compute the probability of swapping two replicas at temperatures β_1 and β_2 , which has to be a function of the ratio of the energy difference to the energy fluctuation

$$p_{\text{swap}} = f\left(\frac{E(\beta_2) - E(\beta_1)}{\sqrt{\sigma^2(\beta_1) + \sigma^2(\beta_2)}}\right).$$

In the large N limit such a ratio can be written as

$$\frac{E(\beta_2) - E(\beta_1)}{\sqrt{\sigma^2(\beta_1) + \sigma^2(\beta_2)}} \simeq \frac{Ne'(\beta)\Delta\beta}{\sqrt{Ne'(\beta)}} \simeq \Delta\beta\sqrt{Ne'(\beta)}.$$

The explicit form for the function f is given by

$$f(z) = \int dx dy \frac{e^{-\frac{x^2}{2} - \frac{(y-z)^2}{2}}}{2\pi} \min(1, e^{z(x-y)}).$$

The best way to allow replicas to wander fast between temperatures is to fix a constant p_{swap} between any pair of successive temperatures, that is using temperature intervals set by

$$\Delta\beta = \frac{f^{-1}(p_{\text{swap}})}{\sqrt{Ne'(\beta)}},$$

which implies in the large N limit the following recursive equation for the temperatures to be used in the optimized parallel tempering

$$\beta_{n+1} = \beta_n + \frac{2r}{\sqrt{N(1 - \tanh(\beta_n/2)^2)}},$$

where $r = f^{-1}(p_{\text{swap}})$. The solution to the above recursive equation converges in the large N limit to the following set of temperatures:

$$\beta_n = 2 \operatorname{arcsinh}\left[\tan\left(\frac{rn}{\sqrt{N}}\right)\right], \quad (\text{B1})$$

with the index n running up to $\lfloor \sqrt{N}\pi/(2r) \rfloor$.

The optimal value for r is the one minimizing the mean traveling time between the extremal temperatures, which is proportional to $[r^2 f(r)]^{-1}$. The minimum is achieved at $r_{\text{opt}} \approx 1.68$, that corresponds to an optimal swapping rate $f(r_{\text{opt}}) \approx 0.23$ (this result is well known as the 0.23 rule). We run all out PT simulations with the temperature set defined in Eq. (B1) with $r = r_{\text{opt}}$.

- [1] S. Arora and B. Barak, *Computational Complexity: A Modern Approach* (Cambridge University Press, Cambridge, 2009).
- [2] A. K. Hartmann and M. Weigt, *Phase Transitions in Combinatorial Optimization Problems*, Vol. 67 (Wiley Online Library, Hoboken, NJ, 2005).
- [3] F. Ricci-Tersenghi, M. Weigt, and R. Zecchina, Simplest random k -satisfiability problem, *Phys. Rev. E* **63**, 026702 (2001).
- [4] M. Mézard, G. Parisi, and R. Zecchina, Analytic and algorithmic solution of random satisfiability problems, *Science* **297**, 812 (2002).
- [5] W. Barthel, A. K. Hartmann, M. Leone, F. Ricci-Tersenghi, M. Weigt, and R. Zecchina, Hiding Solutions in Random Satisfiability Problems: A Statistical Mechanics Approach, *Phys. Rev. Lett.* **88**, 188701 (2002).
- [6] M. Mézard, F. Ricci-Tersenghi, and R. Zecchina, Two solutions to diluted p -spin models and XORSAT problems, *J. Stat. Phys.* **111**, 505 (2003).
- [7] S. Cocco, O. Dubois, J. Mandler, and R. Monasson, Rigorous Decimation-Based Construction of Ground Pure States for Spin-Glass Models on Random Lattices, *Phys. Rev. Lett.* **90**, 047205 (2003).
- [8] A. Montanari, G. Parisi, and F. Ricci-Tersenghi, Instability of one-step replica-symmetry-broken phase in satisfiability problems, *J. Phys. A* **37**, 2073 (2004).
- [9] M. Mézard, T. Mora, and R. Zecchina, Clustering of Solutions in the Random Satisfiability Problem, *Phys. Rev. Lett.* **94**, 197205 (2005).
- [10] S. Mertens, M. Mézard, and R. Zecchina, Threshold values of random K -SAT from the cavity method, *Random Struct. Algor.* **28**, 340 (2006).
- [11] F. Krzakala, A. Montanari, F. Ricci-Tersenghi, G. Semerjian, and L. Zdeborová, Gibbs states and the set of solutions of random constraint satisfaction problems, *Proc. Natl. Acad. Sci. USA* **104**, 10318 (2007).
- [12] F. Krzakala and J. Kurchan, Landscape analysis of constraint satisfaction problems, *Phys. Rev. E* **76**, 021122 (2007).
- [13] A. Montanari, F. Ricci-Tersenghi, and G. Semerjian, Clusters of solutions and replica symmetry breaking in random k -satisfiability, *J. Stat. Mech.: Theory Exp.* (2008) P04004.
- [14] L. Zdeborová and F. Krzakala, Phase transitions in the coloring of random graphs, *Phys. Rev. E* **76**, 031131 (2007).
- [15] F. Altarelli, R. Monasson, and F. Zamponi, Relationship between clustering and algorithmic phase transitions in the random k -XORSAT model and its np-complete extensions, *J. Phys. Conf. Ser.*, **95**, 012013 (2008).
- [16] F. Krzakala, F. Ricci-Tersenghi, L. Zdeborova, R. Zecchina, E. W. Tramel, and L. F. Cugliandolo, *Statistical Physics, Optimization, Inference, and Message-Passing Algorithms: Lecture Notes of the Les Houches School of Physics-Special Issue, October 2013* (Oxford University Press, Oxford, 2016).
- [17] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, Quantum computation by adiabatic evolution, [arXiv-quant-ph/0001106](https://arxiv.org/abs/quant-ph/0001106).
- [18] G. E. Santoro and E. Tosatti, Optimization using quantum mechanics: Quantum annealing through adiabatic evolution, *J. Phys. A* **39**, R393 (2006).
- [19] E. Farhi, D. Gosset, I. Hen, A. W. Sandvik, P. Shor, A. P. Young, and F. Zamponi, Performance of the quantum adiabatic algorithm on random instances of two optimization problems on regular hypergraphs, *Phys. Rev. A* **86**, 052334 (2012).
- [20] B. Altshuler, H. Krovi, and J. Roland, Anderson localization makes adiabatic quantum optimization fail, *Proc. Natl. Acad. Sci. USA* **107**, 12446 (2010).
- [21] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, [arXiv:1411.4028](https://arxiv.org/abs/1411.4028).
- [22] C. R. Laumann, R. Moessner, A. Scardicchio, and S. L. Sondhi, Quantum annealing: The fastest route to quantum computation?, *Eur. Phys. J.: Spec. Top.* **224**, 75 (2015).
- [23] G. Mossi and A. Scardicchio, Ergodic and localized regions in quantum spin glasses on the bethe lattice, *Philos. Trans. R. Soc. A* **375**, 20160424 (2017).
- [24] G. Mossi *et al.*, The Transverse-Field Ising Spin Glass Model on the Bethe Lattice with an Application to Adiabatic Quantum Computing, Ph.D. thesis, SISSA (2017).
- [25] V. N. Smelyanskiy, K. Kechedzhi, S. Boixo, S. V. Isakov, H. Neven, and B. Altshuler, Nonergodic Delocalized States for Efficient Population Transfer within a Narrow Band of the Energy Landscape, *Phys. Rev. X* **10**, 011017 (2020).
- [26] D. M. Basko, I. L. Aleiner, and B. L. Altshuler, Metal-insulator transition in a weakly interacting many-electron system with localized single-particle states, *Ann. Phys. (NY)* **321**, 1126 (2006).
- [27] C. R. Laumann, A. Pal, and A. Scardicchio, Many-Body Mobility Edge in a Mean-Field Quantum Spin Glass, *Phys. Rev. Lett.* **113**, 200405 (2014).
- [28] J. Z. Imbrie, V. Ros, and A. Scardicchio, Local integrals of motion in many-body localized systems, *Ann. Phys. (NY)* **529**, 1600278 (2017).
- [29] E. Farhi, D. Gamarnik, and S. Gutmann, The quantum approximate optimization algorithm needs to see the whole graph: A typical case, [arXiv:2004.09002](https://arxiv.org/abs/2004.09002).
- [30] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature (London)* **574**, 505 (2019).
- [31] O. Dubois and J. Mandler, The 3-XORSAT threshold, *Comptes Rendus Math.* **335**, 963 (2002).
- [32] S. Cocco, R. Monasson, A. Montanari, and G. Semerjian, Approximate analysis of search algorithms with “physical” methods, [arXiv:cs/0302003](https://arxiv.org/abs/cs/0302003).
- [33] M. Ibrahim, Y. Kanoria, M. Kraning, and A. Montanari, The set of solutions of random XORSAT formulae, in *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms* (SIAM, Philadelphia, 2012), pp. 760–779.
- [34] W.-K. Chen, D. Gamarnik, D. Panchenko, M. Rahman, Suboptimality of local algorithms for a class of max-cut problems, *Ann. Probab.* **47**, 1587 (2019).
- [35] D. Gamarnik and M. Sudan, Limits of local algorithms over sparse random graphs, in *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science* (Association for Computing Machinery, New York, 2014), pp. 369–376.
- [36] D. Gamarnik and A. Jagannath, The overlap gap property and approximate message passing algorithms for p -spin models, *Ann. Probab.* **49**, 180 (2021).
- [37] Strictly speaking, XORSAT is in P, as the solution to a set of αN linear equations in N variables can be always found in time $O(N^3)$ by a simple Gaussian elimination algorithm (more complicated algorithms can do marginally better). However, such an algorithm is very special to XORSAT and cannot be generalized to other CSPs being inherently non-local (i.e. the

- operations performed involve in general many variables, also very distant on the interaction network). On the contrary, when *local* algorithms are run on XORSAT they are found to be very slow and affected strongly by the dynamical transition at α_d . Even the very same Gaussian elimination algorithm has an algorithmic phase transition at α_d from linear to cubic behavior [73] and it does not work as soon as the problem is slightly perturbed (e.g., by the addition of a tiny fraction of constraints with the OR operator).
- [38] S. Franz, M. Mézard, F. Ricci-Tersenghi, M. Weigt, and R. Zecchina, A ferromagnet with a glass transition, *Europhys. Lett.* **55**, 465 (2001).
- [39] F. Krzakala and L. Zdeborová, Following Gibbs states adiabatically, the energy landscape of mean-field glassy systems, *Europhys. Lett.* **90**, 66002 (2010).
- [40] L. Zdeborová and F. Krzakala, Generalization of the cavity method for adiabatic evolution of Gibbs states, *Phys. Rev. B* **81**, 224205 (2010).
- [41] L. Zdeborová and F. Krzakala, Statistical physics of inference: Thresholds and algorithms, *Adv. Phys.* **65**, 453 (2016).
- [42] O. Regev, On lattices, learning with errors, random linear codes, and cryptography, *J. ACM (JACM)* **56**, 1 (2009).
- [43] K. Pietrzak, Cryptography from learning parity with noise, in *International Conference on Current Trends in Theory and Practice of Computer Science* (Springer, New York, 2012), pp. 99–114.
- [44] S. Franz, M. Leone, F. Ricci-Tersenghi, and R. Zecchina, Exact Solutions for Diluted Spin Glasses and Optimization Problems, *Phys. Rev. Lett.* **87**, 127209 (2001).
- [45] A. Montanari and F. Ricci-Tersenghi, On the nature of the low-temperature phase in discontinuous mean-field spin glasses, *Eur. Phys. J. B* **33**, 339 (2003).
- [46] A. Montanari and F. Ricci-Tersenghi, Cooling-schedule dependence of the dynamics of mean-field glasses, *Phys. Rev. B* **70**, 134406 (2004).
- [47] T. Jörg, F. Krzakala, G. Semerjian, and F. Zamponi, First-Order Transitions and the Performance of Quantum Algorithms in Random Optimization Problems, *Phys. Rev. Lett.* **104**, 207206 (2010).
- [48] V. Bapst, L. Foini, F. Krzakala, G. Semerjian, and F. Zamponi, The quantum adiabatic algorithm applied to random optimization problems: The quantum spin glass perspective, *Phys. Rep.* **523**, 127 (2013).
- [49] G. Folena, S. Franz, and F. Ricci-Tersenghi, Rethinking Mean-Field Glassy Dynamics and its Relation with the Energy Landscape: The Surprising Case of the Spherical Mixed p -Spin Model, *Phys. Rev. X* **10**, 031045 (2020).
- [50] S. Sarao Mannelli, G. Biroli, C. Cammarota, F. Krzakala, P. Urbani, and L. Zdeborová, Marvels and Pitfalls of the Langevin Algorithm in Noisy High-Dimensional Inference, *Phys. Rev. X* **10**, 011057 (2020).
- [51] S. S. Mannelli and L. Zdeborová, Thresholds of descending algorithms in inference problems, *J. Stat. Mech.: Theory Exp.* (2020) 034004.
- [52] G. Semerjian and R. Monasson, Relaxation and metastability in a local search procedure for the random satisfiability problem, *Phys. Rev. E* **67**, 066103 (2003).
- [53] G. Semerjian and M. Weigt, Approximation schemes for the dynamics of diluted spin models: The ising ferromagnet on a bethe lattice, *J. Phys. A* **37**, 5525 (2004).
- [54] N. C. Wormald, Differential equations for random processes and random graphs, *Ann. Appl. Probab.* **5**, 1217 (1995).
- [55] B. Selman, H. A. Kautz, B. Cohen, Local search strategies for satisfiability testing., *Cliques, Coloring, and Satisfiability* **26**, 521 (1993).
- [56] A. Montanari and G. Semerjian, On the dynamics of the glass transition on Bethe lattices, *J. Stat. Phys.* **124**, 103 (2006).
- [57] G. Semerjian, On the freezing of variables in random constraint satisfaction problems, *J. Stat. Phys.* **130**, 251 (2008).
- [58] W. Barthel, A. K. Hartmann, and M. Weigt, Solving satisfiability problems by fluctuations: The dynamics of stochastic local search algorithms, *Phys. Rev. E* **67**, 066104 (2003).
- [59] M. Guidetti and A. P. Young, Complexity of several constraint-satisfaction problems using the heuristic classical algorithm walksat, *Phys. Rev. E* **84**, 011102 (2011).
- [60] W. Wang, J. Machta, and H. G. Katzgraber, Population annealing: Theory and application in spin glasses, *Phys. Rev. E* **92**, 063307 (2015).
- [61] T. Kadowaki and H. Nishimori, Quantum annealing in the transverse Ising model, *Phys. Rev. E* **58**, 5355 (1998).
- [62] A. Ambainis and O. Regev, An elementary proof of the quantum adiabatic theorem, [arXiv:quant-ph/0411152](https://arxiv.org/abs/quant-ph/0411152).
- [63] A. Messiah, *Quantum Mechanics* Vol. II (North-Holland Publishing, Amsterdam, 1962).
- [64] C. R. Laumann, R. Moessner, A. Scardicchio, and S. L. Sondhi, Quantum Adiabatic Algorithm and Scaling of Gaps at First-Order Quantum Phase Transitions, *Phys. Rev. Lett.* **109**, 030502 (2012).
- [65] L. Foini, G. Semerjian, and F. Zamponi, Solvable Model of Quantum Random Optimization Problems, *Phys. Rev. Lett.* **105**, 167204 (2010).
- [66] V. Bapst, G. Semerjian, and F. Zamponi, Effect of quantum fluctuations on the coloring of random graphs, *Phys. Rev. A* **87**, 042322 (2013).
- [67] J. R. Dormand and P. J. Prince, A family of embedded Runge-Kutta formulae, *J. Comput. Appl. Math.* **6**, 19 (1980).
- [68] P. Weinberg and M. Bukov, QuSpin: A Python package for dynamics and exact diagonalisation of quantum many body systems Part I: Spin chains, *Sci. Post Phys.* **2**, 003 (2017).
- [69] A. P. Young, S. Knysh, and V. N. Smelyanskiy, First-Order Phase Transition in the Quantum Adiabatic Algorithm, *Phys. Rev. Lett.* **104**, 020502 (2010).
- [70] H.-J. Zhou and Q. Liao, Circumventing spin-glass traps by microcanonical spontaneous symmetry breaking, *Phys. Rev. E* **103**, 042112 (2021).
- [71] M. Kowalsky, T. Albash, I. Hen, and D. A. Lidar, 3-regular 3-XORSAT planted solutions benchmark of classical and quantum heuristic optimizers, [arXiv:2103.08464](https://arxiv.org/abs/2103.08464).
- [72] M. Bernaschi, M. Bisson, M. Fatica, E. Marinari, V. Martin-Mayor, G. Parisi, and F. Ricci-Tersenghi, How we are leading a 3-XORSAT challenge: From the energy landscape to the algorithm and its efficient implementation on GPUs, *Europhys. Lett.* **133**, 60005 (2021).
- [73] A. Braunstein, M. Leone, F. Ricci-Tersenghi, and R. Zecchina, Complexity transitions in global algorithms for sparse linear systems over finite fields, *J. Phys. A* **35**, 7559 (2002).