

Thermal, quantum and simulated quantum annealing: analytical comparisons for simple models

V Bapst and G Semerjian

LPTENS, Unité Mixte de Recherche (UMR 8549) du CNRS et de l'ENS, associée à l'UPMC
Univ Paris 06, 24 Rue Lhomond, 75231 Paris Cedex 05, France.

E-mail: victor.bapst@polytechnique.edu, guilhem@lpt.ens.fr

Abstract. We study various annealing dynamics, both classical and quantum, for simple mean-field models and explain how to describe their behavior in the thermodynamic limit in terms of differential equations. In particular we emphasize the differences between quantum annealing (i.e. evolution with Schrödinger equation) and simulated quantum annealing (i.e. annealing of a Quantum Monte Carlo simulation).

1. Introduction

Combinatorial optimization is a branch of computer science concerned with the task of finding the minimum of a cost function $C(\underline{\sigma})$ of discrete variables $\underline{\sigma} = (\sigma_1, \dots, \sigma_N)$. The computational complexity theory [1] classifies the difficulty of such a task, depending on the problem type (the “shape” of the function), in terms of the time and memory requirements of algorithms that find optimal configurations of the variables (or approximations with a controlled accuracy). If the cost function is viewed as an energy, optimal configurations correspond to groundstates and this analogy leads to physics-inspired annealing procedures aiming at discovering the groundstate of a given system. Thermal (simulated) annealing [2] in particular exploits thermal fluctuations to achieve this goal: in a Monte Carlo numerical experiment a fictitious temperature can be slowly decreased down to very low temperatures. If this decrease is slow enough the system remains close to equilibrium at all times, and at the end of this annealing, where the temperature is very low, the typical configurations encountered will be the groundstates of the system. Instead of thermal fluctuations, or in addition to them, one can envision the use of quantum fluctuations to accelerate the exploration of the configuration space towards the discovery of the groundstate. This idea appeared with several variants and names [3, 4, 5, 6, 7], in particular quantum annealing and quantum adiabatic algorithm (see [8, 9, 10, 11] for reviews). In one of these variants one considers the evolution of an isolated system at zero temperature, thus described by the Schrödinger equation, with a quantum Hamiltonian interpolating during the course of the evolution between a simple Hamiltonian with an easy to prepare groundstate (typically a transverse field) and the Hamiltonian encoding the cost function one wants to minimize. If this interpolation is slow enough (i.e. if the adiabaticity condition is fulfilled), a system initially prepared in the groundstate of the simple Hamiltonian will end up in the sought-for groundstate of the cost function. This procedure, that we shall call quantum annealing in the following, could be implemented in an ideal quantum computer (with a perfect coherence), but its simulation on a classical computer is of course very costly (the dimension of the quantum

Hilbert space growing exponentially with the number of variables). Another variant, that we shall call simulated quantum annealing in the following (also known as Path-Integral Monte Carlo Quantum Annealing in [8, 10] or Quantum Monte Carlo Annealing in [9]), amounts to perform a Quantum Monte Carlo numerical experiment, based on an imaginary-time path integral representation of the quantum partition function, at positive temperature, in which some parameters, in particular the transverse field, are slowly modified during the simulation. Recent claims that such a procedure can be implemented experimentally with a rather large number of variables can be found in [12].

A satisfactory understanding of these various annealing procedures would require being able to answer at least the following two questions: i) how long should be the annealing to reach the groundstate? ii) if the annealing is too fast, what is the residual energy (i.e. the difference between the energy of the configuration reached at the end of the annealing and the one of the groundstate) as a function of the annealing time? These questions, and in particular the behavior of these quantities in the thermodynamic limit, have notably been studied in the context of mean-field spin-glasses [13], that are relevant as typical instances of random combinatorial optimization problems [14]. However, even in this mean-field setting, only partial answers to questions i) and ii) are known. Consider for instance the thermal annealing case; mean-field spin-glasses do exhibit a so-called dynamic transition temperature T_d , below which metastable pure states proliferate and drive the system out-of-equilibrium if the annealing is not slow enough to cross the free energy barriers between the pure states [15]. In many of these models these barriers are extensive, hence the time required for an annealing to bring the system in its groundstate is exponentially large in the size of the system. Even if the computation of T_d is now well-understood, the rate of extensive growth of the barriers, hence the answer to question i), is generically unknown. Question ii) can be more easily addressed in a particular limit, i.e. for annealing durations which go to infinity after the thermodynamic limit has been taken, see in particular [16, 17, 18] for recent works on this question with the so-called state following method.

The current knowledge for the analysis of the quantum annealing procedure is similar and slightly less advanced: from the equilibrium analysis of quantum mean-field spin glasses, via the replica method for fully-connected models or the cavity one for diluted models (see [11] for a recent review of these works), the location of the quantum phase transition where the gap between the groundstate and the first excited one closes in the thermodynamic limit can be computed (similarly to T_d for the classical case), yet the quantitative behavior of the gap is known analytically only in some simple cases [19, 20, 21]. The equivalent of the state following method in the quantum case has not yet been developed. For what concerns the simulated quantum annealing, the dynamic transition of the path-integral model representing the quantum problem at finite temperature has been computed; as demonstrated in [11] this is the point at which a QMC annealing falls out-of-equilibrium, but again no prediction is available for its long (but subexponential) time behaviour. Let us also mention the existence of worst-case bounds on the annealing schedules of both quantum and simulated quantum annealing that ensure the convergence to the groundstate in an infinite time (see [10] for a review of these results) and of several numerical experiments comparing the efficiency of the various annealing procedures (for instance [22, 23], and [8, 9] and references therein).

The strategy followed in this paper consists in studying simple models on which these questions can be investigated analytically, namely non-disordered fully-connected models. The thermal annealing of these models is rather simple to analyze, see [24] for early results. We previously considered in great details the quantum annealing procedure (i.e. the real-time Schrödinger evolution) [21], here we show that the simulated quantum annealing case is also tractable (which was previously discovered in [25]) and we emphasize the difference between the last two dynamics.

The rest of this paper is organized as follows. In Sec. 2 we give explicit definitions of the

various annealing dynamics and of the models to be studied. Then in Sec. 3 we present the evolution equations that govern the dynamics in the thermodynamic limit, and study their long-time behavior; this unveils a crucial difference in the behavior of the quantum annealing and the simulated quantum annealing (even at very low temperature). We finally come to our conclusions in Sec. 4, sketching some extensions of these results to more complicated models and other time regimes. For reason of space limitations we refrain from giving details of the computations that shall be presented elsewhere [26].

2. Definitions

2.1. Various annealing dynamics

We consider models whose classical configurations are given by the values of N Ising spins, denoted $\underline{\sigma} = (\sigma_1, \dots, \sigma_N)$. In the quantum case we consider the 2^N -dimensional Hilbert space spanned by the vectors denoted $\{|\underline{\sigma}\rangle\}$, that form the so-called computational basis. We write $\hat{\sigma}_i^z$ and $\hat{\sigma}_i^x$ for the Pauli matrices acting on the i -th spin according to $\hat{\sigma}_i^z|\underline{\sigma}\rangle = \sigma_i|\underline{\sigma}\rangle$ and $\hat{\sigma}_i^x|\underline{\sigma}\rangle = |\underline{\sigma}^{(i)}\rangle$, where $\underline{\sigma}^{(i)}$ is the configuration obtained from $\underline{\sigma}$ by flipping the i -th spin.

In the various annealing procedures we denote \mathcal{T} the total physical time of the evolution, which is parametrized by a reduced time parameter $s \in [0, 1]$. All parameters are allowed to evolve with s , i.e. $\beta(s)$ is the time varying inverse temperature, and in the most general case we consider a time-dependent quantum Hamiltonian written as

$$\hat{H}(s) = \sum_{\underline{\sigma}} E(\underline{\sigma}; s) |\underline{\sigma}\rangle \langle \underline{\sigma}| - \Gamma(s) \sum_{i=1}^N \hat{\sigma}_i^x. \quad (1)$$

The ‘‘classical’’ (i.e. diagonal in the computational basis) energy $E(\underline{\sigma}; s)$ can a priori vary in time, as well as the intensity $\Gamma(s)$ of the transverse field.

Let us now define the dynamics to be studied.

2.1.1. Thermal annealing This first case is purely classical ($\Gamma(s) = 0$), with a dynamics made of a Markov stochastic process in the configuration space, described by a continuous time master equation for the probability $P(\underline{\sigma}; s)$ to observe the system in the configuration $\underline{\sigma}$ at (reduced) time s :

$$\frac{1}{\mathcal{T}} \frac{d}{ds} P(\underline{\sigma}; s) = \sum_{\underline{\sigma}'} W(\underline{\sigma}, \underline{\sigma}'; s) P(\underline{\sigma}'; s). \quad (2)$$

The jump rates from configuration $\underline{\sigma}'$ to $\underline{\sigma}$ (with diagonal terms chosen to ensure the conservation of probabilities) are assumed to verify the detailed balance condition with respect to the instantaneous Gibbs-Boltzmann distribution, i.e.

$$W(\underline{\sigma}, \underline{\sigma}'; s) e^{-\beta(s)E(\underline{\sigma}'; s)} = W(\underline{\sigma}', \underline{\sigma}; s) e^{-\beta(s)E(\underline{\sigma}; s)}. \quad (3)$$

In the usual form of simulated annealing $E(\underline{\sigma}; s)$ is independent of s and equal to the cost function $C(\underline{\sigma})$ to be minimized, and $\beta(s)$ increases from 0 to ∞ as s goes from 0 to 1. Note that β and E always appear through the combination βE ; to simplify the discussion below we shall assume that β is fixed to some arbitrary value and that all the time-dependence of the annealing schedule is included in $E(\underline{\sigma}; s)$. We also assume that the initial condition at $s = 0$ is at equilibrium.

We shall further assume that the dynamics flips one spin at a time, W is hence non-zero only if $\underline{\sigma}$ and $\underline{\sigma}'$ are at Hamming distance at most 1. For definiteness we further specialize this choice to

$$W(\underline{\sigma}^{(i)}, \underline{\sigma}; s) = \frac{1}{2} \left(1 - \tanh \left(\beta \frac{E(\underline{\sigma}^{(i)}; s) - E(\underline{\sigma}; s)}{2} \right) \right). \quad (4)$$

This form of the transition rate corresponds to the Glauber (or heat bath) rule, and can be expressed in terms of the effective magnetic field h_i acting on spin i , defined as

$$h_i(\underline{\sigma}; s) = \sigma_i \frac{E(\underline{\sigma}^{(i)}; s) - E(\underline{\sigma}; s)}{2} . \quad (5)$$

This expression is actually independent of σ_i .

2.1.2. Quantum annealing This case corresponds to the evolution of an isolated system. Its wave function $|\psi(s)\rangle$ thus follows Schrödinger's equation,

$$\frac{i}{\mathcal{T}} \frac{d}{ds} |\psi(s)\rangle = \hat{H}(s) |\psi(s)\rangle . \quad (6)$$

We shall assume that the interpolation starts from a purely transverse field, $E(\underline{\sigma}; s = 0) = 0$, with $|\psi(s = 0)\rangle$ the groundstate of the transverse field (i.e. the uniform superposition of the 2^N states of the computational basis), and that at the end of the annealing the transverse field is turned off and the classical energy corresponds to the cost function to be optimized: $\Gamma(s = 1) = 0$ and $E(\underline{\sigma}; s = 1) = C(\underline{\sigma})$.

2.1.3. Simulated quantum annealing Quantum models can be represented as classical models with an additional imaginary time dimension (using for instance the Suzuki-Trotter identity [27]), of length corresponding to the inverse temperature. In the case of Ising spins in a transverse field one is thus led to consider a configuration space of trajectories, that we shall denote $\underline{\sigma} = (\sigma_1, \dots, \sigma_N)$, with the convention that bold font symbols stand for imaginary-time dependent quantities, here $\sigma_i = \{\sigma_i(t) : t \in [0, \beta]\}$ with $\sigma_i(t) = \pm 1$ is the imaginary time trajectory of an Ising spin. The quantum partition function and the thermodynamic average can then be computed by introducing the following equilibrium measure on the space of trajectory configurations (for a classical energy $E(\underline{\sigma})$, a transverse field Γ and an inverse temperature β):

$$P_{\text{eq}}(\underline{\sigma}) = \frac{1}{Z} \prod_{i=1}^N \Gamma^{|\sigma_i|} e^{-\int_0^\beta E(\underline{\sigma}(t)) dt} , \quad (7)$$

where $|\sigma_i|$ denotes the number of spin flips in the trajectory of the i -th spin, and we assume periodic boundary conditions in the imaginary time direction, $\underline{\sigma}(t = 0) = \underline{\sigma}(t = \beta)$. A Quantum Monte Carlo simulation proceeds by a stochastic exploration of the space of configurations $\underline{\sigma}$. One can use for instance a quantum version of the heat-bath algorithm, where elementary moves correspond to the update of one spin trajectory σ_i (for the whole imaginary time axis), redrawn from the equilibrium measure conditioned on the current state of all trajectories on sites distinct from i . What we call simulated quantum annealing is precisely this type of stochastic process for $\underline{\sigma}$, described by the master equation,

$$\frac{1}{\mathcal{T}} \frac{d}{ds} P(\underline{\sigma}; s) = \sum_{\underline{\sigma}'} W(\underline{\sigma}, \underline{\sigma}'; s) P(\underline{\sigma}'; s) , \quad (8)$$

where the transition rates are those of heat-bath updates, respecting detailed balance with respect to the instantaneous equilibrium distribution (7) with the current values of $E(\cdot; s)$ and $\Gamma(s)$ (for simplicity of notation we assume the inverse temperature $\beta(s)$ to be held constant). The initial condition will be considered to be equilibrated at $s = 0$, i.e. in a purely transverse field.

2.2. Mean-field fully-connected models

The models we shall consider are generalizations of the (quantum) Curie-Weiss model (also related to the Lipkin-Meshkov-Glick model [28]), with their energy depending on the configurations only through the total magnetization:

$$E(\underline{\sigma}; s) = Ng \left(\frac{1}{N} \sum_{i=1}^N \sigma_i; s \right), \quad (9)$$

$g(m; s)$ being the energy density for a configuration of magnetization density m at reduced time s . The effective field (see the definition in Eq. (5)) created by such a configuration then reads

$$h(m; s) = -\frac{\partial}{\partial m} g(m; s), \quad (10)$$

at the leading order in the large N limit.

As the transverse field $\Gamma(s)$ is also conjugated to the total (transverse) magnetization, such quantum models are invariant under any permutation of the spins (or in more physical terms their Hamiltonians commute with the total spin), which greatly simplifies their analysis.

3. Evolution equations

3.1. Differential equations

As we shall see now the high level of symmetry in the definition of these models allows one to reduce the study of their dynamics, in the thermodynamic limit $N \rightarrow \infty$ with the annealing time \mathcal{T} kept finite, to simple differential equations, for the three type of dynamics introduced above and arbitrary time dependencies of their parameters.

3.1.1. Thermal annealing If the initial condition $P(\underline{\sigma}; s = 0)$ of the master equation (2) is invariant under the permutation of the sites (i.e. only depends on the magnetization of the initial condition), then this property is conserved along the dynamical evolution (even for finite N), allowing one to deduce a master equation for $P(m; s)$. In addition in the thermodynamic limit this probability distribution concentrates around its most likely value $m(s)$ (if the initial condition is not too wide), the latter being solution of a simple, deterministic, differential equation:

$$\frac{1}{\mathcal{T}} \frac{dm}{ds} = \tanh(\beta h(m(s); s)) - m(s). \quad (11)$$

This expression, that can be easily obtained after a short computation, has a natural interpretation: when a spin is updated its average magnetization corresponds to that of a single spin at equilibrium at inverse temperature β , in an effective magnetic field created by the rest of the configuration and thus equal to $h(m(s); s)$.

3.1.2. Quantum annealing The Schrödinger evolution of such mean-field spin models was studied in details in [21] (see also [29, 30] for similar treatments of particle models), where it was shown that the thermodynamic limit leads to semi-classical Hamiltonian equation of motions for the typical magnetization $m(s)$ and its conjugated momentum $\tilde{m}(s)$:

$$\begin{cases} \frac{1}{\mathcal{T}} \frac{d}{ds} m(s) = \frac{\partial}{\partial \tilde{m}} \mathcal{H}(m(s), \tilde{m}(s); s) \\ \frac{1}{\mathcal{T}} \frac{d}{ds} \tilde{m}(s) = -\frac{\partial}{\partial m} \mathcal{H}(m(s), \tilde{m}(s); s) \end{cases}, \quad (12)$$

where the effective Hamiltonian \mathcal{H} is given in terms of the parameters of the model as

$$\mathcal{H}(m, \tilde{m}; s) = g(m; s) - \Gamma(s) \sqrt{1 - m^2} \cos(2\tilde{m}). \quad (13)$$

The initial condition is $(m(s=0), \tilde{m}(s=0)) = (0, 0)$.

3.1.3. Simulated quantum annealing Let us finally consider the case of the annealing of a Quantum Monte Carlo simulation. Its analysis is at first sight much more complicated than the thermal annealing one, as the configurations have an additional imaginary time dependence. It turns out however that in the large N limit the magnetization density concentrates around a typical value $m(s)$, independently of the imaginary time (i.e. the static approximation of [25] becomes exact), and that the correlations along this axis are irrelevant (see [26] for a detailed derivation of this statement). This magnetization density is found to obey a simple differential equation,

$$\frac{1}{\mathcal{T}} \frac{dm}{ds} = \frac{h(m(s); s)}{\sqrt{h(m(s); s)^2 + \Gamma(s)^2}} \tanh(\beta \sqrt{h(m(s); s)^2 + \Gamma(s)^2}) - m(s), \quad (14)$$

as was already obtained in [25] for a particular choice of $h(m, s)$. As it should this equation reduces to (11) when the transverse field $\Gamma(s)$ vanishes at all times. In fact the intuitive explanation given after (11) remains valid here, as the first term in (14) is precisely the average value of the z -magnetization of a quantum spin 1/2 in a longitudinal field h and transverse field Γ .

3.2. Analysis in the large annealing time limit

For any finite value of \mathcal{T} the evolution equations (11), (12) and (14) can be easily integrated numerically, for any annealing schedule. One can however make further analytical progresses by considering the limit $\mathcal{T} \rightarrow \infty$ of a very slow annealing (we recall that we have already taken the thermodynamic limit $N \rightarrow \infty$). As the thermal annealing is a particular case of the simulated quantum annealing we shall not discuss explicitly the thermal case, and contrast the behavior of the quantum and simulated quantum annealing.

Consider the evolution equation (14) of the latter. In the large \mathcal{T} limit one can check that $m(s)$ will be close to $m_*(s)$, where $m_*(s)$ is such that the right-hand side of (14) vanishes. This condition is easily seen to be equivalent to $m_*(s)$ being an extremum of the free energy $f(m; s)$, a standard computation (see for example [21]) leading indeed to

$$f(m; s) = \text{ext}_{\lambda} \left[g(m; s) + m\lambda - \frac{1}{\beta} \ln 2 \cosh(\beta \sqrt{(1-s)^2 + \lambda^2}) \right]. \quad (15)$$

Of course the free energy can have several extrema, in that case one has to be more precise to define $m_*(s)$. The rules to remove this ambiguity are as follows. $m_*(0)$ has to be chosen as the local minimum reached by a steepest-descent from the initial condition. Then $m_*(s)$ remains stuck in the same local minimum as long as the latter exists. At all bifurcations where $m_*(s)$ loses its stability, it flows to a new local minimum following a steepest-descent search of the free energy. For future use let us write down a more explicit form of the above free energy in the zero-temperature limit:

$$f(m; s) = g(m; s) - (1-s) \sqrt{1-m^2}. \quad (16)$$

Let us now turn to the analysis of the slow annealing limit of the quantum annealing equations (12). These have the form of Hamiltonian equations of motion, and their large \mathcal{T} limit can thus be analyzed using the theory of adiabatic invariants of classical mechanics. The initial condition $(m(s=0), \tilde{m}(s=0)) = (0, 0)$ is a local minimum of the Hamiltonian $\mathcal{H}(m, \tilde{m}; s=0)$, hence for large annealing times \mathcal{T} the solution of (12) remains close to $(m(s), \tilde{m}(s)) = (m_*(s), 0)$, as long as this point corresponds to a local minimum of \mathcal{H} . This is actually also the minimum of the zero-temperature free energy (16), hence the quantum annealing and the zero-temperature simulated quantum annealing coincide in this limit as long as no bifurcation is encountered. However a drastic difference between the two dynamics appears as soon as the local minimum




Figure 1. Sketch of the difference of behavior between the zero temperature simulated quantum annealing (orange circle) and the quantum annealing (red circle and wavelet) when crossing a spinodal point along their annealing.

loses its stability: the semi-classical dynamics describing the quantum annealing is indeed conservative, hence a non-zero value of the momentum \tilde{m} is induced at the bifurcation, and an oscillatory behavior for $m(s)$ in its potential energy well takes place. On the contrary the zero-temperature simulated quantum annealing flows to the newly accessible local minimum. This difference of behavior is illustrated on figure 1.

4. Conclusions

In this paper we have studied various types of annealing dynamics of quantum spin models, simple enough to be reducible to unidimensional problems, their magnetization playing the role of a particle coordinate. The main message to be emphasized is the difference between a quantum annealing (an evolution with Schrödinger equation) and a low temperature simulated quantum annealing (a time-varying Quantum Monte Carlo simulation) that appears when a bifurcation occurs in the structure of the effective energy potential along the annealing (see also [31] for a discussion of this difference from another point of view). This difference persists even in the limit of zero-temperature, which can be understood physically as a system coupled to a “zero-temperature heat bath” can release part of its energy, at variance with an isolated system. Microscopic descriptions of the coupling to the environment might thus be necessary to interpret the results of experiments on quantum annealers [12].

When the groundstate of such a model encounters a first-order quantum phase transition both quantum and simulated quantum annealing remains stuck in the metastable continuation of the groundstate between the transition and the spinodal point. Right after the spinodal the energy is strictly smaller in the simulated quantum annealing case, as the coupling to the heat bath provides a mechanism for energy relaxation. This could seem to suggest that including both thermal and quantum fluctuations is the best strategy to attain the lowest possible residual energy. This conclusion is however wrong in general, as the subsequent relative evolution of the depth and width of the multiple energy wells between the spinodal and the end of the annealing can in some cases reverse this ordering [26].

The long-term goal sketched in the introduction would be to understand these annealing dynamics in disordered mean-field systems that correspond to random optimization problems. The discussion presented here in terms of effective free energy is reminiscent of the quantum and dynamical TAP equations approach of [32, 33] for disordered fully-connected models. Such models, even if obviously much more intricate than the simple ones discussed here, share with the latter the possibility of first-order phase transitions accompanied by metastable states and spinodals (thanks to their mean-field character), one can thus hope for some level of universality for the mechanism encountered at spinodals. As an intermediate step in this program we shall present in [26] further results on the quantum random subcubes model [34, 35]; in the regime where the number of cubes is finite one can exploit the permutation symmetry between sites

inside a few blocks of spins, which leads to a similar analysis of the annealing dynamics with the magnetization replaced by a finite dimensional vector of block magnetizations.

Let us finally emphasize that we concentrated here on a particular time regime, namely with the limit $\mathcal{T} \rightarrow \infty$ of long annealing times taken after the thermodynamic limit $N \rightarrow \infty$. In presence of metastability this time scale is however too short to reach the groundstate energy, the adiabatic time (for quantum or simulated annealing) growing exponentially fast with the system size. It is thus also important to consider the regime where \mathcal{T} and N are taken to infinity simultaneously, with $\mathcal{T} = \exp[N\tau]$. We have shown in [21] how to compute the residual energy of the quantum annealing as a function of τ via the determination of the exponentially small avoided level crossings along the metastable branches of the eigenvalues of the Hamiltonian $\widehat{H}(s)$. In the case of simulated quantum annealing one can also investigate this exponential time regime by computing the height of the free energy barriers, the adiabatic time being controlled by these barriers through an Arrhenius law [26].

- [1] Papadimitriou C H 1994 *Computational Complexity* (Addison-Wesley)
- [2] Kirkpatrick S, Gelatt Jr C D and Vecchi M 1983 *Science* **220** 671–80
- [3] Apolloni B, Carvalho C and de Falco D 1989 *Stoc. Proc. Appl.* **33** 233–44
- [4] Finnila A B, Gomez M A, Sebenik C, Stenson C and Doll J D 1994 *Chem. Phys. Lett.* **219** 343–8
- [5] Kadowaki T and Nishimori H 1998 *Phys. Rev. E* **58** 5355–63
- [6] Brooke J, Bitko D, Rosenbaum T F and Aeppli G 1999 *Science* **284** 779–81
- [7] Farhi E, Goldstone J, Gutmann S, Lapan J, Lundgren A and Preda D 2001 *Science* **292** 472–5
- [8] Santoro G E and Tosatti E 2006 *J. Phys. A: Math. Gen.* **39** R393
- [9] Das A and Chakrabarti B K 2008 *Rev. Mod. Phys.* **80** 1061–81
- [10] Morita S and Nishimori H 2008 *J. Math. Phys.* **49** 125210 .
- [11] Bapst V, Foini L, Krzakala F, Semerjian G and Zamponi F 2013 *Physics Reports* **523** 127–205
- [12] Boixo S *et al.* 2013 Quantum annealing with more than one hundred qubits (*Preprint* arXiv:1304.4595 [quant-ph])
- [13] Mézard M and Montanari A 2009 *Information, Physics and Computation* (New York: Oxford University Press)
- [14] Cheeseman P, Kanefsky B and Taylor W M 1991 Where the really hard problems are *Proc. 12th IJCAI (San Mateo CA USA)* (Morgan Kaufmann) pp 331–7
- [15] Cugliandolo L F and Kurchan J 1993 *Phys. Rev. Lett.* **71** 173–6
- [16] Zdeborová L and Krzakala F 2010 *Phys. Rev. B* **81** 224205
- [17] Franz S and Parisi G 2013 *Journal of Statistical Mechanics: Theory and Experiment* **2013** P02003
- [18] Krzakala F and Zdeborová L 2013 *this volume*.
- [19] Jörg T, Krzakala F, Kurchan J and Maggs A C 2008 *Phys. Rev. Lett.* **101** 147204
- [20] Farhi E, Goldstone J, Gosset D, Gutmann S and Shor P 2011 *Quantum Information & Computation* **11** 840–54
- [21] Bapst V and Semerjian G 2012 *Journal of Statistical Mechanics* **2012** P06007
- [22] Santoro G E, Martonak R, Tosatti E and Car R 2002 *Science* **295** 2427
- [23] Stella L, Santoro G E and Tosatti E 2006 *Phys. Rev. B* **73** 144302
- [24] Griffiths R B, Weng C-Y and Langer J S 1966 *Phys. Rev.* **149** 301–5
- [25] Inoue J I 2010 *Journal of Physics: Conference Series* **233** 012010
- [26] Bapst V and Semerjian G 2013 in preparation
- [27] Suzuki M 1976 *Progress of Theoretical Physics* **56** 1454–69
- [28] Lipkin H J, Meshkov N and Glick A J 1965 *Nucl. Phys.* **62** 188–98
- [29] Sciolla B and Biroli G 2011 *J. Stat. Mech.* **2011** P11003
- [30] Schiró M and Fabrizio M 2010 *Phys. Rev. Lett.* **105** 076401
- [31] Hastings M B and Freedman M H 2013 Obstructions to classically simulating the quantum adiabatic algorithm (*Preprint* arXiv:1302.5733 [quant-ph])
- [32] Biroli G 1999 *Journal of Physics A: Mathematical and General* **32** 8365–88
- [33] Biroli G and Cugliandolo L F 2001 *Phys. Rev. B* **64** 014206
- [34] Mora T and Zdeborová L 2008 *J. Stat. Phys.* **131** 1121–38
- [35] Foini L, Semerjian G and Zamponi F 2010 *Phys. Rev. Lett.* **105** 167204