


LABORATOIRE DE PHYSIQUE THEORIQUE
DE L'ECOLE NORMALE SUPERIEURE



THESE D'HABILITATION A DIRIGER DES RECHERCHES

présentée par

Guilhem Semerjian

**Mean-field disordered systems :
glasses and optimization problems,
classical and quantum**

Soutenue le 31 janvier 2013

Composition du jury :

M.	Amin COJA-OGHLAN	(Rapporteur)
Mme	Leticia CUGLIANDOLO	
M.	Bernard DERRIDA	
M.	Benoît DOUÇOT	
M.	Satya MAJUMDAR	(Rapporteur)
M.	Marc MÉZARD	
M.	Rémi MONASSON	
M.	Giorgio PARISI	(Rapporteur)

Contents

1	Introduction	3
1.1	Statistical mechanics of disordered systems	3
1.2	Random optimization problems	4
1.3	Quantum computing and statistical mechanics	6
2	Phase transitions in random constraint satisfaction problems	7
2.1	Definitions and basic properties	7
2.1.1	Constraint satisfaction problems	7
2.1.2	Random constraint satisfaction problems	9
2.2	Structural phase transitions	11
2.2.1	The clustering transition of random XORSAT formulas	11
2.2.2	The condensation transition	14
2.2.3	The freezing transition	16
2.3	Algorithmic/dynamic consequences	16
2.3.1	Dynamical signature of the clustering transition	18
2.3.2	Belief inspired decimation	20
2.4	Methods	22
3	The effect of quantum fluctuations	29
3.1	Introduction to (adiabatic) quantum computing	29
3.2	Quantum cavity method	32
3.3	Quantum annealing of random CSPs	33
4	Perspectives	39
	Acknowledgments	41
	Bibliography	43

Chapter 1

Introduction

This Habilitation Thesis manuscript summarizes my main research activities since the end of my Ph.D., to let the committee assess my ability to supervise the doctoral works of future students. A large part of this activity has taken place at the border between physics and other disciplines, namely discrete mathematics and theoretical computer science, the first goal of this introduction is thus to make a short history and give a justification for this interdisciplinary approach.

1.1 Statistical mechanics of disordered systems

On the physics side, the line of research that eventually led to these interdisciplinary applications was triggered, in the 70's, by the discovery of materials with peculiar magnetic properties [1]. These so-called spin-glasses are very imbalanced mixtures of two elements, the majority one being metallic and non-magnetic, while the minority one has magnetic properties. The compound is prepared in a liquid phase at high temperature, then cooled down to a solid phase. The location of the magnetic impurities is thus frozen (“quenched”); it turns out moreover that the interaction between the spins of the impurities is mediated by the delocalized electrons of the metal, in such a way that the coupling strength between two spins has an oscillating sign as a function of the distance between them [2]. The existence of both ferromagnetic and antiferromagnetic interactions between the spins induces frustration in the system, i.e. these contradictory interactions forbid a simple groundstate (a configuration of minimal energy) as in the ferromagnetic case where all spins are aligned in the same direction. Yet, the competing effects of energy and entropy lead to a phase transition as a function of temperature between a paramagnetic phase where interactions are irrelevant and spins fluctuate under the effect of thermal noise, and a low-temperature phase where the spins get polarized in a preferred orientation, even though this orientation is different for each spin as a consequence of the frustrating interactions.

Structural glasses, e.g. window glasses, are liquids that, upon cooling, become so viscous that they solidify. Yet they are amorphous solids, their microscopic structure shows no trace of long-range crystalline order, at variance with many “ordered solids”, most pure metals for instance. The “local freezing” of the positions of the molecules in a structural glass, with no crystalline order, is the analog of the freezing of the spin orientations, with no ferromagnetic order, of a spin-glass, hence the name of the latter. Glasses and spin-glasses thus share several properties (slow dynamics, aging and memory effects), even if they also differ on crucial aspects (continuity or not of their order parameter at the transition, presence

of quenched disorder). Structural glasses are still the object of intense experimental and theoretical works (see the book [3] for recent reviews). A central puzzle of this field is the difficulty to find noticeable signatures of the growth of spatial correlations in the vicinity of the glass transition, whereas on the contrary the temporal correlations of the dynamics exhibit a dramatic increase. Some of the works discussed in the following will be directly linked to this point.

The built-in (“quenched”) disorder of spin-glasses makes their theoretical study much harder than their pure counterparts, i.e. ferromagnets. In the latter case one has a very detailed picture of the finite-dimensional Ising models, for instance rigorous proofs for the existence of a phase transition in the thermodynamic limit (in dimension ≥ 2), exact solution in dimension 2, knowledge of the upper critical dimension above which the critical exponents of the mean-field approximation are correct, and field theoretical approaches for systematic computations of the critical behavior below the upper critical dimension. By contrast, and despite decades of research, the status of finite-dimensional spin-glasses is much more uncertain [4, 5, 6], and is still the object of (controversial) numerical studies [7, 8]. Fortunately, for the interdisciplinary applications to be studied here, mean-field models, which from a physical point of view only represent a gross approximation to be improved upon, are the most relevant ones and shall be the focus of the following.

The simplest mean-field spin-glass model was introduced by Sherrington and Kirkpatrick (SK) in 1975 [9], shortly after the finite-dimensional modelization and mean-field approximate treatment by Edwards and Anderson [10]. It is the disordered analog of the Curie-Weiss model of a ferromagnet: N Ising spins interact all with each other, but the $N(N - 1)/2$ coupling constants are drawn as independent random variables. The computation of the free-energy per spin was completed by Parisi in 1980 [11], via the development of the replica method with replica symmetry breaking [12], that encodes the complex structure of the configuration space of the SK model. A rigorous proof of this formula was obtained much more recently by Talagrand [13], building on the interpolation method developed by Guerra and Toninelli [14]. In the SK model every spin interacts (weakly) with all others, in other words the graph of interaction is the complete graph on N vertices, where the degree of each spin is $N - 1$, that diverges in the thermodynamic limit. There exists however mean-field models which keep a finite degree (connectivity) in the thermodynamic limit, the first of which was introduced by Viana and Bray [15]. The mean-field nature of these models arises from the randomness in the choice of the interaction graph, they do not have an Euclidean finite-dimensional geometric structure underlying the notion of “neighboring” vertices. These models defined on sparse random graphs are tightly connected with random optimization problems, as will be expanded upon below.

1.2 Random optimization problems

Combinatorial optimization problems [16] of theoretical computer science have many points in common with statistical mechanics. In both cases one considers a discrete set of configurations (aka microstates, or assignments) built from a large number of elementary degrees of freedom (spins, boolean variables...), endows it with some real function, to be called cost function or energy depending on the context, and one is mainly interested in the minima of this function (i.e. groundstates or optimal assignments). The positive temperature partition function can be naturally interpreted as a generating function for the number of configurations at a given energy, and allows to smoothly interpolate towards the counting of the optimal assignments recovered in the zero-temperature limit. If the objects studied are thus very similar, one however encounters at this point a deep divergence of methods

and objectives between statistical mechanics and one branch of theoretical computer science, namely the computational complexity theory [17, 18]. The latter classifies indeed families of problems, in other words “shapes” of cost functions, according to the algorithmic complexity of the exploration of the low energy configurations of all instances of a family of problems (dividing, very roughly, families where the number of required operations is a polynomial in the number of variables from those where it is exponential). This worst case point of view has been very fruitful for this classification perspective and has led to a rich and beautiful theory of complexity classes, with still famous pending conjectures, most notably the P vs NP question. Physics tend on the contrary to disregard worst case events, and while acknowledging that it is possible for all molecules of air inside a room to spontaneously leave half of the room empty, the (hopefully) extremely small probability of this event makes it considered irrelevant, because atypical. It turned out that a similar concern for “typicality” emerged from inside computer science, and eventually made the two fields closer than could be initially expected. This concern raised in the 80’s, because of the following disturbing situation. Some combinatorial problems, notably k -satisfiability and q -coloring, to be defined precisely later, were known to be NP-complete, and thus widely believed to be “hard”, in the sense that some of their instances should require an exponentially large time to be solved. At the same time most instances that could be explicitly constructed were easily solved by known algorithms [19]. This observation raised two kind of questions: could it be that some families of problems are difficult only because of a few atypical instances? Is it possible to generate such hard instances in an efficient way? To answer these questions, and to give a precise definition of typicality, random ensembles of instances of combinatorial optimization problems have been introduced [20, 21, 22]. “Typical” is thus meant in this context for “with a probability which goes to one in the large size (thermodynamic) limit”. This construction actually leads back to the first subject of this introduction, and explains the relevance of statistical mechanics to this approach to combinatorial optimization problems [12, 23]. The randomness in the ensembles of instances is the exact parallel of the quenched disorder of spin glasses. In addition the constraints in hard optimization problems have competing requirements, similarly to the physical frustration induced by coexisting ferromagnetic and antiferromagnetic interactions. Finally, there is a priori no reason to embed optimization problems in a finite-dimensional Euclidean space (though there exist structured instances relevant for applications that could have some remnants of finite-dimensionality), hence the random ensembles of optimization problems devised by computer scientists turned out to be (finitely-connected) mean-field spin-glasses. Not surprisingly with this analogy in mind, (many distinct) phase transitions do occur in random ensembles of optimization problems. During the last twenty years many works used the methods originally devised for the study of physical spin-glasses in order to better understand the properties of random optimization problems. Several insights, based in particular on the “glassy” aspects of such models [24, 25, 26], have been fruitfully translated into this different context, and for some of them have been turned into rigorous statements [27, 28, 29, 30]. In fact, if computer scientists could early on identify one of the phase transitions occurring in these problems, it seems fair to say that the intuition about glasses brought by physicists was very helpful to understand that several other sharp thresholds were to be looked after in these models. Note that the interdisciplinary character of this field of research has been also beneficial for physics, as some concepts first elaborated in the mathematical literature turned out to be useful for glassy physics [31]. The main features of the current picture that emerged from this line of research will be presented in chapter 2, along with some of the works I participated to.

A similar change of perspective from worst-case to typical analysis was also performed in

the 60's in another domain of discrete mathematics, namely graph theory [32], even if with slightly different motivations. Random graph theory [33] has since then become a branch of probability on its own, with again a wealth of phase transitions occurring in the limit of infinitely large random graphs. Tools of statistical mechanics have also been successfully applied to these problems, I shall briefly mention some examples later on.

1.3 Quantum computing and statistical mechanics

Chapter 3 will be devoted to studies of quantum generalizations of the optimization problems, or more generically disordered models, introduced above. To give a brief definition for a non-physicist reader, one should view the configurations of the classical model as indices for an orthonormal basis of vectors in a Hilbert space, and define a Hamiltonian operator on this space with diagonal matrix elements equal to the classical cost function, and some non-zero off-diagonal elements that allow the system to “jump” from one classical configuration to another. These quantum fluctuations are such that the Hamiltonian is Hermitian, the operator has thus as many real eigenvalues as there were configurations in the classical problem; the quantum partition function takes the same generating function form as in the classical case, with these eigenvalues replacing the values of the cost function on the original configuration space.

The motivation for these studies is two-fold. From a physical point of view, one can model in this way for instance systems of interacting particles constrained to live on the vertices of a lattice [34] (this is either a simplification of a system of particles in the continuum or a very good description of some cold atoms experiments [35]), the quantum fluctuations corresponding to the hopping events of one particle from one vertex to a neighboring one. Understanding the interplay between interactions, in particular disordered and frustrating ones, and quantum fluctuations is an important issue in condensed matter. One can notably wonder about the existence of phases of matter that display simultaneously both quantum and glassy features [36]. Quantum mean-field disordered systems are a natural (simplified) playground to investigate this question.

The perspectives offered by quantum computing [37] form the second side of the motivation. Very roughly, a quantum computer would be a device whose internal state is described by an arbitrary vector of the Hilbert space constructed above, in other words it can exist in a coherent superposition of all the classical configurations. Presented in this way one could think that such a device offers for free an exponential speedup with respect to classical computers, as it seems to operate in a parallel way all possible inputs of a problem. This is of course too naive because of the laws of projective measurement in quantum mechanics. Nevertheless, specific quantum algorithms have been devised that solve computational tasks by exploiting some of the specificities of the quantum mechanical behavior of a quantum computer, the most famous one being the integer factoring algorithm by Shor [38]. For what concerns optimization problems, a versatile algorithm has been proposed under several names, i.e. quantum annealing or quantum adiabatic algorithm [39, 40, 41, 42], that proceeds by a progressive vanishing of the quantum fluctuations (the off-diagonal elements of the Hamiltonian), similarly to simulated annealing where the thermal fluctuations are progressively reduced. This procedure can be applied to any optimization problem, however its efficiency, i.e. the time required for it to be successful, will depend very much on the problem considered. In a series of works summarized in chapter 3 and more extensively in the recent review [43] we have attempted to understand this efficiency in the typical case, by adding quantum fluctuations to random optimization problems.

Chapter 2

Phase transitions in random constraint satisfaction problems

2.1 Definitions and basic properties

2.1.1 Constraint satisfaction problems

The optimization problems considered here are Constraint Satisfaction Problems (CSPs): their energy function can be interpreted as the number of unsatisfied constraints, or clauses, where each of these constraints involves a small number of variables and limits the authorized values for them. To be specific, we shall consider a discrete configuration space of N variables denoted $\sigma_1, \dots, \sigma_N$, each of them taking values in a finite set χ , and denote a global configuration $\underline{\sigma} = (\sigma_1, \dots, \sigma_N) \in \chi^N$. In most computer science applications the variables considered are boolean, and will be encoded here as Ising spins, $\chi = \{+1, -1\}$. Let us give some examples of the energy, or cost function, $E(\underline{\sigma})$, corresponding to some well-known CSPs:

- In the graph q -coloring problem, in short q -COL, one is given a graph $G = (V, L)$ with V a set of N vertices and L a set of M edges between pairs of vertices. The variables take values in $\chi = \{1, \dots, q\}$, with $q \geq 2$ an arbitrary integer, so that each configuration $\underline{\sigma}$ corresponds to the coloring where vertex i is given the color σ_i . The cost function is

$$E(\underline{\sigma}) = \sum_{a=1}^M \delta_{\sigma_{i_a^1}, \sigma_{i_a^2}} \quad , \quad (2.1)$$

where the sum runs over the M edges (i_a^1, i_a^2) of the graph, and δ denotes the Kronecker symbol. The cost function thus counts the number of monochromatic edges in the configuration $\underline{\sigma}$.

The following examples involve binary variables, i.e. $\chi = \{+1, -1\}$.

- The k -XORSAT problem is defined on a k -hypergraph $G = (V, L)$: each of the M hyper-edges in L involves a k -uplet of variables, with $k \geq 2$, thus generalizing the notion of usual graphs that corresponds to $k = 2$. We label the hyper-edges with an index $a = 1, \dots, M$, and denote i_a^1, \dots, i_a^k the indices of the vertices linked by the a -th

hyper-edge. In addition to the hyper-graph the problem is defined by M constants $J_a \in \{+1, -1\}$, and the cost function reads

$$E(\underline{\sigma}) = \sum_{a=1}^M \frac{1 - J_a \prod_{j=1}^k \sigma_{i_a^j}}{2}. \quad (2.2)$$

It is easily seen that this sum equals the number of hyper-edges a for which the condition $\sigma_{i_a^1} \dots \sigma_{i_a^k} = J_a$ is violated. There are various equivalent interpretations of this condition; in the language of coding theory [44] this is a parity check rule. By associating Ising spins to $\{0, 1\}$ variables according to $\sigma_i = (-1)^{x_i}$ it is also equivalent to a linear equation of the form $x_{i_a^1} + \dots + x_{i_a^k} = y_a$, where $J_a = (-1)^{y_a}$ and the additions are interpreted modulo 2. Finally it can be seen as a condition on the eXclusive OR of k boolean variables $\{\text{True}, \text{False}\}$, hence the name of the problem.

- In the k -SAT problem one is given an hypergraph and, for each hyper-edge, k constants $J_a^1, \dots, J_a^k \in \{+1, -1\}$; the cost of a configuration is then defined as

$$E(\underline{\sigma}) = \sum_{a=1}^M \prod_{j=1}^k \frac{1 - J_a^j \sigma_{i_a^j}}{2}. \quad (2.3)$$

Each term of the sum is equal to 1 if, for all the k vertices involved in that hyper-edge, one has $\sigma_{i_a^j} = J_a^j$; on the contrary it vanishes as soon as one of the k vertices fulfill $\sigma_{i_a^j} \neq J_a^j$. In terms of boolean variables this is the disjunction (logical OR) of k literals, that are equal to a variable or its logical negation depending on the sign of J_a^j . The formula encoded by this set of constraints is thus a SAT formula in Conjunctive Normal Form.

From the perspective of the computational complexity theory, the decision problem associated to these CSPs is: “given an instance (or formula), i.e. a (hyper)graph and the coupling constants (if necessary), decide whether there exists an assignment $\underline{\sigma}$ satisfying simultaneously the M constraints or not”. Such an assignment is usually called a solution of the problem; in physical terms this corresponds to a zero-energy groundstate. The three examples given above belong to different complexity classes: k -XORSAT is in the (easy) class P for any value of k ; there exists indeed an algorithm running in Polynomial time (with respect to N, M) that provides the answer yes or no for the existence of a zero-energy configuration. Using the interpretation in terms of linear equations given above, this can indeed be solved by Gaussian elimination. On the contrary k -SAT and q -COL, for $k, q \geq 3$ are in the NP-complete class (2-SAT and 2-COL turn out to be in P). This means that no algorithm is known to solve all instances of these decision problems in polynomial time, and if one were discovered it could be adapted to solve all problems of the class NP; the current belief is that no such algorithm exists, but this remains a conjecture (the famous P \neq NP one).

In the following we shall study random instances, for which the worst-case distinction of computational complexity theory is much less relevant, and for instance the behavior of random XORSAT instances is in many respects similar to the one of SAT instances. Before that we shall make a few remarks: there are many other computational tasks related to these CSP besides the decision on the existence of a solution. For instance, one can ask to decide the existence of a configuration of energy smaller or equal than some threshold E_0 . This problem contains the one introduced above (it is recovered when $E_0 = 0$), hence it is

at least as hard. But in fact it can be strictly harder, and this is the case for XORSAT; if the Gaussian elimination algorithm concludes that no assignment can satisfy all the linear equations, it gives no indication on the optimal configuration that minimizes the number of unsatisfied constraints. This observation is made here to emphasize that there are many different tasks associated to the same CSPs (finding the minimal energy reachable by a configuration, constructing such an optimal configuration, counting their number, ...) and that even when the decision problem on the existence of solutions is easy all these tasks are generally (much) harder.

2.1.2 Random constraint satisfaction problems

The motivations for the study of optimization problems from a typical instead of worst-case point of view have been discussed in the introduction. For the three examples given above the simplest random ensembles, that were first (as far as I know) introduced in [20] for k -SAT, are defined in a very natural way. For each of the M constraints one chooses the k -uplet i_a^1, \dots, i_a^k of variables involved uniformly at random among the $\binom{N}{k}$ possible ones ($k = 2$ in the coloring case), and the coupling constants (J_a for XORSAT, J_a^1, \dots, J_a^k for k -SAT) are taken to be ± 1 with equal probability $1/2$. These choices are done independently, in the same way, for each constraint. The large size (thermodynamic) limit corresponds to $N, M \rightarrow \infty$, with $\alpha = M/N$ a fixed parameter (the other scaling regimes of M with respect to N being trivial). In the coloring case one thus studies Erdős-Rényi $G(N, M)$ random graphs, the cases $k > 2$ corresponding to the natural hypergraph generalization of these ensembles.

The authors of [21, 22] performed extensive numerical experiments on such randomly generated instances. Using complete algorithms they determined the probability $P(\alpha, N)$ that an instance with N variables and $M = \alpha N$ constraints has a ground state of vanishing energy (i.e. is q -colorable, or satisfiable depending on the case). This probability is obviously a decreasing function of α : it can only become harder to satisfy all the constraints as their number is increased. What came as a surprise at that time is the fact that for larger and larger values of N the probability of satisfiability decreased in a steeper and steeper way, which suggested the following *satisfiability conjecture*:

$$\lim_{N \rightarrow \infty} P(\alpha, N) = \begin{cases} 1 & \text{if } \alpha < \alpha_s \\ 0 & \text{if } \alpha > \alpha_s \end{cases}, \quad (2.4)$$

where α_s is some fixed threshold value, that depends on the problem considered (coloring or satisfiability), and on the parameters k, q . In more physical terms this threshold phenomenon corresponds to a phase transition between a SAT (or COL) phase where almost all instances are satisfiable (colorable) and their ground state energy is zero to an UNSAT (UNCOL) phase in which almost none of them is, and the average ground state energy is positive. A rigorous proof of a weaker version of the satisfiability conjecture for k -SAT and q -COL, where α_s is allowed to depend on N , can be found in [45]; lower and upper bounds on $\alpha_s(N)$ are known, see for instance [46, 47, 48] and references therein, that become tight in the large k limit [49], what remains to be proven to settle the (widely held to be true) conjecture is the convergence of $\alpha_s(N)$ as $N \rightarrow \infty$. The satisfiability transition was demonstrated for k -XORSAT in [50, 51, 52].

Statistical mechanics methods applied to these random CSPs have provided a framework to actually compute the satisfiability threshold α_s , and maybe more importantly suggested the existence of other phase transitions in the satisfiable regime $\alpha < \alpha_s$, concerning the different possible organizations of the set of solutions inside the space of all configurations.




Figure 2.1: An example of a factor graph.

Before entering this discussion, let us introduce here a convenient representation of the interactions present in Eqs. (2.1,2.2,2.3). All of them can be written generally

$$E(\underline{\sigma}) = \sum_{a=1}^M \varepsilon_a(\underline{\sigma}_{\partial a}) , \quad (2.5)$$

where $\partial a = (i_a^1, \dots, i_a^k)$ denotes the set of variables involved in the a 'th constraint, for a subset S of the variables $\underline{\sigma}_S$ means $\{\sigma_i | i \in S\}$, and ε_a is the energy associated to the a 'th term of the cost function. A factor graph [53], see Fig. 2.1 for an example, is a bipartite graph where each of the N variables σ_i is represented by a circle vertex, while the M interactions are associated to square vertices. An edge is drawn between a variable i and an interaction a if and only if ε_a actually depends on σ_i , i.e. $i \in \partial a$. In a similar way we shall denote ∂i the set of interactions in which σ_i appears, i.e. the graphical neighborhood of i in the factor graph. One has a natural notion of graph distance between two variable nodes i and j , defined as the minimal number of interaction nodes on a path linking i and j .

To describe the set of solutions of a CSP it will be convenient to define the uniform measure over these configurations,

$$\mu(\underline{\sigma}) = \frac{1}{Z} \prod_{a=1}^M w_a(\underline{\sigma}_{\partial a}) , \quad Z = \sum_{\underline{\sigma} \in \mathcal{X}^N} \prod_{a=1}^M w_a(\underline{\sigma}_{\partial a}) , \quad (2.6)$$

where the partition function Z ensures the normalization of the probability law. The uniform measure over solutions of a CSP corresponds to the choice $w_a(\underline{\sigma}_{\partial a}) = (1 - \varepsilon_a(\underline{\sigma}_{\partial a}))$, then Z counts the number of solutions. For physical applications (and to regularize the above definition if there are no solution) one introduces the Gibbs-Boltzmann measure at inverse temperature β by setting $w_a(\underline{\sigma}_{\partial a}) = e^{-\beta \varepsilon_a(\underline{\sigma}_{\partial a})}$. The former case is thus the zero temperature limit of the latter. The associated intensive thermodynamic potential is $\frac{1}{N} \ln Z$; at zero temperature this is the entropy, at finite temperature we shall call it free-entropy (equal to the usual free-energy multiplied by $-\beta$).

The random choices of the indices of variables appearing in an interaction obviously turn the factor graphs into random objects. In the following we shall use some elementary properties of these random factor graphs, in the thermodynamic limit:

- the probability that a randomly chosen variable i has degree $|\partial i| = d$ is $q_d = e^{-\alpha k} (\alpha k)^d / d!$, the Poisson law of mean αk .

- if one chooses randomly an interaction a , then a variable $i \in \partial a$, the probability that i appears in d interactions *besides* a , i.e. that $|\partial i \setminus a| = d$, is $\tilde{q}_d = e^{-\alpha k} (\alpha k)^d / d!$.
- the random factor graphs are locally tree-like: choosing at random a vertex i , the subgraph made of all nodes at graph distance from i smaller than some threshold t is, with a probability going to 1 in the thermodynamic limit with t fixed, a tree.

More general ensembles of random factor graphs can be constructed, by fixing a degree distribution q_d and drawing at random from the set of all graphs of size N with Nq_0 isolated vertices, Nq_1 vertices of degree 1, and so on and so forth. Then the two distributions q_d and \tilde{q}_d are different in general, and related through $\tilde{q}_d = (d+1)q_{d+1} / \sum_{d'} d' q_{d'}$. An important example in this class corresponds to regular random graphs, where q_d is supported by a single integer.

2.2 Structural phase transitions

The satisfiability phase transition described in Eq. (2.4) was identified already in the first numerical studies of the random CSP ensembles. In the recent years several other phase transitions in the satisfiable phase $\alpha < \alpha_s$ have been unveiled, that concern the typical structure of the set of solutions of satisfiable problems. Their non-trivial organization is not only interesting per se, but has also important consequences for the performance of some algorithms to solve these problems, and understanding the organization of the solutions for $\alpha < \alpha_s$ is actually crucial to understand the satisfiability phase transition α_s itself. The mechanism for the disappearance of the solutions at this transition is indeed strongly dependent on their structure just before their limit of existence.

The presentation in the following pages does not follow the chronological order of the wealth of works that applied statistical mechanics tools to random CSPs since several years [54], but is rather an attempt to present the current understanding of their behavior with (unfortunately non-exhaustive) references to the works that shaped this view.

The qualitative picture of these phase transitions is to some extent independent of the details of the random CSP under study (for instance k -SAT and q -COL, for all large enough values of k, q behave in the same way), and corresponds to the mean-field phenomenology of (spin-)glasses. There are however some exceptions, and we shall begin the discussion with the XORSAT case which turns out to be simpler than generic random CSPs.

2.2.1 The clustering transition of random XORSAT formulas

The first studies of the random XORSAT problem that unveiled the rich structure of its satisfiable phase are [50, 51], with a construction that can be described as follows. Consider an arbitrary XORSAT formula, and suppose that one of the variables σ_i appears in a single interaction, call it a . A moment of thought reveals that the formula is satisfiable if and only if the formula obtained from it by removing a is satisfiable. One can iterate this process and reduce further the formula, removing at each step the interactions in which appears a variable of degree 1. At the end of this “leaf-removal” process one ends up with a reduced formula called the 2-core of the original one, that corresponds to the largest subgraph in which all variables have degree at least 2; we shall denote C the set of these variables. By analyzing the dynamics of the leaf removal process one can show that a phase transition occurs, in the thermodynamic limit, for some value α_d of the parameter α controlling the density of the initial graph. For $\alpha < \alpha_d$, the 2-core is typically empty, i.e. the leaf removal process is able to prune all interactions. On the contrary for $\alpha > \alpha_d$, with high probability




Figure 2.2: A sketch of the configuration space of random XORSAT problems. For low values of α the solutions, represented as black dots, are evenly spread on the N -dimensional hypercube. In the intermediate regime they are grouped in clusters, symbolized by the circles. For $\alpha \geq \alpha_s$ there are no more solutions.

the 2-core is non-trivial. Moreover for $k \geq 3$, which will be understood from now on in this discussion, the transition is discontinuous: the typical fraction of vertices in the 2-core jumps from 0 to a strictly positive value as α increases across α_d .

Let us now discuss the consequences of this transition in the factor graph in terms of the solutions of the corresponding formula. Consider first the case where the 2-core is empty; then the original formula is obviously satisfiable. One can indeed assign satisfying values for the variables in the last removed interaction, and then reintroduce the interactions in the reverse order of their removal, using the fact that at each step at least one variable (the leaf) can be freely chosen to satisfy the re-introduced interaction. Indeed, if one calls l_a the number of yet non-assigned variables among the k variables of interaction a at the moment it is reintroduced, by construction $l_a \geq 1$. When l_a is strictly greater than 1 there are 2^{l_a-1} possible choices for the free variables in a , so that the total number of solutions is $\exp[(\ln 2)(\sum_a l_a - M)]$. A detailed analysis [50, 51] reveals that typically in the thermodynamic limit, $\frac{1}{N} \sum_a l_a \rightarrow 1$, hence the total number of solutions for $\alpha < \alpha_d$ is $e^{N s_{\text{tot}}(\alpha)}$, with an entropy density $s_{\text{tot}}(\alpha) = (\ln 2)(1 - \alpha)$. One can moreover show [55, 56] that all these solutions are in some precise sense close one to each other.

On the other hand, when $\alpha > \alpha_d$ the 2-core contains typically an extensive number of variables and interactions. This reduced formula, in which all variables are involved in at least two interactions, goes from satisfiable to unsatisfiable at an higher threshold $\alpha_s > \alpha_d$ [50, 51]. As a formula is satisfiable if and only if its 2-core is satisfiable, this is also the satisfiability threshold for the original ensemble (before leaf removal).

Let us consider the intermediate regime $\alpha \in [\alpha_d, \alpha_s]$, where the reduced formula on the 2-core is non-trivial but still has some solutions; we shall denote $e^{N \Sigma(\alpha)}$ the typical number of solutions of the reduced formula on the variables C of the core, and call its exponential rate $\Sigma(\alpha)$ the complexity (or configurational entropy). From every solution of the 2-core reduced formula one can construct different solutions of the full formula, by reintroducing the interactions in reverse order, as explained above. All the solutions that emerge from the same seed, i.e. from the same solution $\underline{\sigma}_C$ of the 2-core, will be said to belong to the same cluster (or pure state). The number of free choices while constructing a cluster from one seed is $\prod'_a 2^{l_a-1}$, where the product is restricted to the interactions outside the 2-core. This geometrical construction makes it clear that the number of solutions of the full formula emerging from one solution $\underline{\sigma}_C$ of the 2-core is independent of $\underline{\sigma}_C$ (this can be traced back to the linear structure of the constraints), and we shall denote $e^{N s_{\text{int}}(\alpha)}$ this number. In this intermediate regime the total entropy density $s_{\text{tot}}(\alpha)$ is decomposed as $\Sigma(\alpha) + s_{\text{int}}(\alpha)$, with a contribution from the cluster degeneracy and one from the internal




Figure 2.3: The total entropy of solutions $s_{\text{tot}}(\alpha)$ and the complexity (or configurational entropy) $\Sigma(\alpha)$ for the 3-XORSAT problem [50, 51]. The inset is a zoom of the region close to α_d and α_s ; the internal entropy s_{int} is given by the difference $s_{\text{tot}} - \Sigma$.

entropy of the clusters. A very important point is that two distinct solutions $\underline{\sigma}_C$ and $\underline{\sigma}'_C$ of the 2-core formula are far away from each other, in the Hamming distance sense (the number of variables taking opposite values in $\underline{\sigma}_C$ and $\underline{\sigma}'_C$). Indeed, because of the form of the constraints of the XORSAT problem, each interaction must contain an even number of spins i with $\sigma_i \neq \sigma'_i$. In other words, as the 2-core does not contain leaves, at least one loop of disagreeing spins between $\underline{\sigma}_C$ and $\underline{\sigma}'_C$ has to be closed. As the random graphs are locally tree-like such a loop has necessarily a length diverging with N in the thermodynamic limit (this is not strictly true: in random graphs one typically find a finite number of finite cycles, hence clusters whose core solution only differs on a small “flippable cycle” should be grouped together [55, 56]).

The fact that solutions inside one cluster are close to each other, while solutions belonging to distinct clusters are necessarily separated by a large Hamming distance, is represented pictorially in Fig. 2.2; the outer square contains all the 2^N configurations (obviously the representation of this N -dimensional hypercube on a two-dimensional drawing is only a cartoon), while the black area stands for the $e^{N s_{\text{tot}}(\alpha)}$ solutions. For $\alpha < \alpha_d$ all solutions are somehow close to each other, while for the intermediate regime $\alpha \in [\alpha_d, \alpha_s]$ they are broken in $e^{N \Sigma(\alpha)}$ clusters, each containing $e^{N s_{\text{int}}(\alpha)}$ solutions. The functions $s_{\text{tot}}(\alpha)$, $s_{\text{int}}(\alpha)$, $\Sigma(\alpha)$, as well as the thresholds α_d, α_s can be computed explicitly [50, 51], they are plotted in Fig. 2.3 for $k = 3$. Further results on the organization of the solution space of random XORSAT formulas can be found in [57, 58, 55, 56].

Let us summarize the main messages on the properties of random CSPs that should be drawn from this particular case. For low values of the control parameter α the exponentially many solutions are spread in the whole configuration space, and close-by one to the other. Increasing α there appears a clustering transition at α_d , after which there are still exponentially many solutions, yet they are grouped in clusters of close-by configurations, the clusters being separated one from the other; in this regime the complexity or configurational entropy Σ counts the exponential number of clusters. The total entropy density of solutions, s_{tot} , is the sum of the complexity Σ and the internal entropy density s_{int} of each cluster, which is here the same for all clusters: $s_{\text{tot}}(\alpha) = \Sigma(\alpha) + s_{\text{int}}(\alpha)$. For even larger values of α the satisfiability transition α_s is due to the vanishing of Σ , i.e. the disappearance of the




Figure 2.4: Illustration of Eq. (2.7). The solid curve is the function $\Sigma(s)$, dashed lines have slope -1 . On the left panel, $\alpha < \alpha_c$, the supremum is reached in an interior point (s_*, Σ_*) . On the right the system is in a condensed phase, $\alpha > \alpha_c$, the supremum is reached at the upper edge s_{\max} , the typical complexity Σ_* vanishes, $s_{\text{tot}} = s_{\max}$. On both panels Σ_{SP} is the maximum of $\Sigma(s)$.

clusters of solutions. It is thus crucial to identify the relevant decomposition into clusters of the set of solutions: what vanishes at α_s is the number of clusters, not the total number of solutions, because the last clusters to disappear still contain an exponential number of solutions.

2.2.2 The condensation transition

For a generic random CSP one can expect, from what we learnt by studying the XOR-SAT problem, that by increasing the density α of constraints a clustering transition will be encountered at some critical value α_d , above which the set of solutions breaks into well-separated clusters of close-by solutions (even though an explicit definition of the clusters is more complicated, as discussed at the end of this section). There is however a priori no reason to assume that all clusters will be of the same size; in XORSAT this feature was a side effect of the peculiar linear structure of the constraints. We will now discuss the consequences of a non-trivial distribution of cluster sizes, that were first envisioned in this context in [24, 59]. Assume that, for a typical random instance of density α , the number of clusters containing $e^{N s_{\text{int}}}$ solutions is, at the leading exponential order, $e^{N \Sigma(\alpha, s_{\text{int}})}$. Suppose in addition that the range of allowed internal entropies is some interval $[s_{\min}, s_{\max}]$, and that in this range Σ is some positive concave function of s_{int} , vanishing at the upper limit s_{\max} (see Fig. 2.4). Then the computation of the total entropy reads

$$e^{N s_{\text{tot}}(\alpha)} = \int_{s_{\min}}^{s_{\max}} ds e^{N(\Sigma(\alpha, s) + s)}, \quad s_{\text{tot}}(\alpha) = \sup_{s \in [s_{\min}, s_{\max}]} \Sigma(\alpha, s) + s, \quad (2.7)$$

where the second expression is obtained from the first via an evaluation of the integral by the Laplace method, at the leading exponential order.

Two qualitatively distinct situations can then occur (and are illustrated on the two panels of Fig. 2.4). If the supremum in Eq. (2.7) is reached at a point s_* inside the interval $[s_{\min}, s_{\max}]$, in other words if there exists a point where the tangent of Σ has slope -1 , then the total entropy is dominated by the $e^{N \Sigma_*}$ clusters with internal entropy s_* . This means that a randomly chosen solution of such a formula would belong to such a cluster, i.e. the vast majority of solutions are spread in an exponential number of clusters. On the contrary




Figure 2.5: Schematic representation of the phase transitions in a random CSP ensemble.

in the situation depicted on the right panel of Fig. 2.4 the supremum is reached at the upper limit s_{\max} , the associated complexity Σ_* vanishes, hence a typical configuration belongs to one of the sub-exponentially many largest clusters of internal entropy s_{\max} .

Explicit computations for k -SAT and q -COL (for $k, q \geq 4$, the cases $k, q = 3$ being peculiar) [26, 60, 61] have shown that the first situation is encountered in a regime of parameter $\alpha \in [\alpha_d, \alpha_c]$, while the second one occurs for $\alpha \in [\alpha_c, \alpha_s]$ (see Fig. 2.5 for an illustration). The so-called “condensation” transition that occurs at α_c is the analog of the thermodynamic, also called Kauzmann, transition of mean-field spin-glasses, that occurs already in the simplest of such models, the Random Energy one [62].

The condensation transition can thus be identified as the threshold for the vanishing of the complexity Σ_* of the point of $\Sigma(s)$ that has a tangent slope of -1 . Another point of this curve also plays a special role, namely the maximum of $\Sigma(s)$ (with tangent slope 0), denoted Σ_{SP} on the figure. This is indeed the exponential rate for the total number of clusters, irrespectively of their size. As a consequence the satisfiability transition α_s corresponds to the vanishing of Σ_{SP} , beyond which there are no cluster anymore, hence no solution. The index SP stands for Survey Propagation, Σ_{SP} is indeed the quantity computed by the Survey Propagation algorithm [25, 63] and the criterion of its vanishing was used for the first computation of the currently held values of α_s in [25, 64].

Note that the case of random 3-SAT and random 3-COL is peculiar: for this value of k, q one has $\alpha_d = \alpha_c$, the clustering transition is continuous and there is no phase $\alpha \in [\alpha_d, \alpha_c]$ with an exponential number of clusters in the relevant part of μ (i.e. in the support of typical solutions). Conversely in the case of k -XORSAT the two thresholds α_c and α_s coincide, as all clusters have the same size the disappearance of the relevant ones (at α_c) marks simultaneously the disappearance of all of them (at the satisfiability threshold α_s).

We come back now on the issue of the definition of the clusters in a generic CSP. The intuitive notion that solutions inside the same cluster should be close (in Hamming distance), while two distinct clusters should be far away, is slightly delicate to formalize (see [26] for one such proposal). For a given instance of finite size, the set of clusters should provide a partition of the set of solutions; this partition should be such that the “frontiers” between clusters contain much less solutions than their inside, i.e. that “entropic barriers” separate them, mimicking the more usual notion of pure states decomposition for translation invariant statistical mechanics models. This definition should imply some correlation decay condition for the Gibbs-Boltzmann probability measure restricted to one cluster (pure state). Note however that the condition to be fulfilled for a partition to be an acceptable cluster decomposition becomes sharp only in the thermodynamic limit, some arbitrariness should remain for a finite given instance, non ambiguous statements on the number of clusters are restricted to the exponential rate of growth of this number in the large N limit. In the rigorous works [28, 29] “clusters” are defined as connected components on the hypercube (two configurations being considered adjacent if the Hamming distance between them is 1), and are grouped in “regions” that fulfill conditions similar to the intuition explained above.

2.2.3 The freezing transition

By their very definition the clusters of XORSAT do contain a finite fraction of frozen variables, in the sense that some variables σ_i take the same value in all solutions of a given cluster (these are the variables of the 2-core C , and actually a few more, as during the re-introduction of the constraints some variables are directly implied by the choice of $\underline{\sigma}_C$). This is however a particular case: for generic random CSPs the clustering threshold α_d does not coincide with the freezing threshold α_f above which typical solutions contain an extensive number of frozen variables, in general $\alpha_d < \alpha_f$. From a technical point of view this statement can be deduced by inspection of the solution of the equations describing the clustered phase in the statistical mechanics formalism, that will be presented in Sec. 2.4 (for the expert reader, there can be non-trivial solutions of the 1RSB equations without any hard field).

A more intuitive and richer discussion of this phenomenon can be obtained in terms of the optimal rearrangements of the solutions. Consider indeed an instance of a CSP, and one of its solution $\underline{\sigma}$; a rearrangement of $\underline{\sigma}$ at vertex i will be another solution $\underline{\sigma}'$ of the instance, such that $\sigma'_i \neq \sigma_i$. In pictorial terms one perturbs the solution $\underline{\sigma}$ by flipping its i -th variable, and one rearranges the new configuration (eliminates the excess energy introduced by the defect at i) in order to fall back to a different groundstate of zero energy. Among all possible rearrangements one can study for instance the ones with minimal “size” (number of flipped variables during the rearrangement), i.e. the solutions $\underline{\sigma}'$ at minimal Hamming distance n from $\underline{\sigma}$, under the constraint that $\sigma'_i \neq \sigma_i$. This minimal distance n is a random variable because of the choice of the instance and of the reference solution $\underline{\sigma}$. Its probability distribution has been studied in [57] for XORSAT and in [65] for arbitrary CSPs. These studies showed the existence of a threshold α_f that separates a phase where n remains bounded in the thermodynamic limit from another one, at $\alpha > \alpha_f$, where n diverges in the thermodynamic limit with finite probability. In this framework infinite minimal size rearrangements are associated to frozen variables; one can actually be more precise and study the precursors of this divergence, i.e. the critical behavior of the distribution of n when $\alpha \rightarrow \alpha_f^-$. On the plots of Fig. 2.6 one sees a plateau growing in this limit in the integrated distribution of n , i.e. its probability law is bimodal with a second peak that moves towards larger values of n when α increases towards α_f . This divergence is controlled by a non-trivial exponent, as well as the scaling behavior around the plateau, that was computed in [57, 65] (for a technical coincidence this critical behavior is the one of the schematic Mode Coupling Theory [66] of liquids).

As said above for XORSAT one has the coincidence $\alpha_d = \alpha_f$; for this model further properties of the rearrangements were studied in [57], in particular the scaling of the minimal energy barrier to be crossed during a rearrangement. The low temperature dynamics of this model, for $\alpha \leq \alpha_d$, could thus be understood in terms of the Arrhenius law, with a barrier for relaxation fluctuating from spin to spin, and a crossover to a super-Arrhenius law when these barrier diverge at the freezing transition. For k -SAT and q -COL one finds instead $\alpha_f > \alpha_d$; the relative position of α_f and α_c do change with k and q [65, 61], for large values of these parameters one finds back the simpler situation of XORSAT with $\alpha_f \sim \alpha_d$. Rigorous proofs for the existence of frozen variables in this limit $k, q \rightarrow \infty$ can be found in [28, 29, 30].

2.3 Algorithmic/dynamic consequences

We shall now briefly discuss some studies of time-dependent processes related to disordered systems. In the computer science point of view they correspond to algorithms devised to




Figure 2.6: The integrated distribution Q_n of minimal size rearrangements, defined as the probability that a rearrangement has minimal size $\geq n$, for various values of α growing towards α_f , here in the case of 3-XORSAT.

solve a computational task, for instance decide whether there exists a solution or not. In a physical context these processes will mimic the real time dynamics of a system, for instance in contact with a thermostat.

There exists many different types of algorithms to solve the decision problem for a CSP, that can be roughly divided in two broad categories: local search and sequential assignment procedures. In the first one a random walk is performed in the configuration space, with transition rules tuned to bias the walk towards the solutions. This kind of algorithm is called incomplete: it cannot prove the absence of solution if it fails to find one. The second category proceeds differently: at some step of the algorithm only one part of the variables has a definite value, the others still being free. Each step thus corresponds to the choice of one free variable and of the value it will be assigned to, the CSP on the remaining variables being consequently simplified. The heuristics guiding these choices can be more or less elaborate. In the simplest cases one only takes into account simple properties of the free variables, such as the number of their occurrences in the remaining CSP. A rigorous analysis of these simple (“myopic”) approaches is possible and is at the basis of most of the lower bounds on the satisfiability transition [46, 47]. Such algorithms can be made complete if backtracking is allowed, i.e. choices which have led to a contradiction can be corrected in a systematic way. This complete version of the procedure is called the DPLL algorithm.

It is only natural, and very important for applications, to investigate the effect of the structural, static, phase transitions exposed above on the behavior of such algorithms. Intuitively the complex configuration space fractured in clusters for values of α relatively close to α_s should be more difficult to explore by simple algorithms than the well-connected phase at low α . It is however not easy to turn this intuition into precise and generic statements, of the type “all algorithms of the family X will always fail to solve CSPs at density larger than α_Y ”, where X could denote local-search vs sequential assignment procedures (or more precise subdivisions of these algorithms), and Y one of the phase transitions (clustering, condensation, freezing,...). We shall content ourselves in the next two subsections with two specific studies dealing with one representative of each family of algorithms, and briefly discuss at the end of each subsection the difficulty in the generalization of such statements.




Figure 2.7: The clustering and condensation transitions of a random CSP extend in the (α, T) phase diagram, at positive temperature $T = 1/\beta$.

2.3.1 Dynamical signature of the clustering transition

The clustering threshold at α_d is actually the zero-temperature remnant of a dynamical transition (hence the subscript d) for a whole class of stochastic processes verifying the detailed balance condition. This statement can be more easily understood by considering the finite temperature generalization of the measure μ discussed after Eq. (2.6). The zero temperature thresholds α_d and α_c extend as transition lines in the phase diagram function of (α, T) , as shown in Fig. 2.7. For all values of $\alpha > \alpha_c$ one encounters, upon lowering the temperature, the two transitions ubiquitous in mean-field spin-glasses with interactions between $p \geq 3$ spins, first studied in the p -spin fully-connected spherical model [67, 68, 69], and on which has been based the RFOT theory of glasses [70, 71]. When the temperature $T_d(\alpha)$ is crossed the free-energy of the system is analytic, even though the Gibbs-Boltzmann measure is fractured in an exponentially large number of pure states (the finite-temperature equivalent of clusters), and it is only at the condensation, or Kauzmann, temperature $T_c(\alpha)$ that a thermodynamic second order phase transition can be detected from the free-energy. However the structural transition at T_d has a drastic consequence for any dynamics that respect the detailed balance condition (in mathematical terms, that is reversible) with respect to the finite temperature Gibbs-Boltzmann measure. Indeed the equilibrium autocorrelation time of such a dynamics diverges when $T \rightarrow T_d^+$, and becomes exponentially large in the system for $T < T_d$, with a very rich out-of-equilibrium dynamics exhibiting an aging behavior [69].

This dynamic characterization of the transition at T_d was first understood via explicit resolutions of fully-connected models (that correspond here to the limit $\alpha \rightarrow \infty$). This can be put in a more general perspective in the context of finite connectivity disordered models by combining the following two observations:

- One technical definition for the transition at T_d is the appearance of a non-trivial 1RSB solution with Parisi parameter $m = 1$. It was first shown in [72], then expanded upon in [57, 26], that an equivalent definition can be given in terms of a specific form of spatial correlations, namely a point-to-set correlation function. To define this concept let $B(i, \ell)$ be the set of variables at graph distance ℓ from the vertex i in a factor graph, and define the correlation between i (a point) and $B(i, \ell)$ (a set) as

$$C(i, \ell) = \sum_{\underline{\sigma}_{B(i, \ell)}} \mu(\underline{\sigma}_{B(i, \ell)}) \sum_{\sigma_i} |\mu(\sigma_i | \underline{\sigma}_{B(i, \ell)}) - \mu(\sigma_i)|. \quad (2.8)$$

The inner sum quantifies the effect of a boundary condition $\underline{\sigma}_{B(i, \ell)}$ on the central variable σ_i (and indeed vanishes if they are independent one from each other). The

correlation function is an average measure of this dependency, when the boundary condition itself is taken with the Gibbs-Boltzmann probability law (a “thermalized boundary condition”, related to the thought experiment discussed in [73]). What happens at T_d is the birth of long-range point-to-set correlation length, i.e. for $T < T_d$ $C(i, \ell)$ remains strictly positive in the limit $\ell \rightarrow \infty$ (the thermodynamic limit and the average over the quenched disorder being taken first).

- It has been shown in [31] that the temporal correlations of a (local) dynamics that respect detailed balance can be related to the spatial correlation length defined in terms of point-to-set correlation. More precisely, one can define ℓ_i as the minimal distance such that $C(i, \ell)$ drops below a given small reference value, and similarly τ_i as the minimal time t such that the equilibrium auto-correlation function at site i , $\langle \sigma_i(0)\sigma_i(t) \rangle - \langle \sigma_i \rangle^2$ is smaller than some arbitrary threshold. Then bounds of the form

$$\ell_i \leq \tau_i \leq \exp[V(\ell_i)] \quad (2.9)$$

were proven in [31], where several constants have been hidden for simplicity, and where $V(\ell_i)$ is the number of sites contained in the ball of radius ℓ_i around site i .

Now the “dynamical” character of the clustering transition is demonstrated: on this line the point-to-set correlation length diverges, and from (2.9) this implies a divergence of the equilibrium correlation time of a reversible dynamics.

Note that the bounds between length and time scales obtained in [31] are valid for any topology of interactions, be it mean-field or finite-dimensional. This has some relevance in the context of glasses, as it sheds light on the apparent paradox of the glass transition phenomenology, where a dramatic increase of the relaxation time of a liquid occurs in a range of temperature where the structure factor (i.e. a two-point correlation function) barely changes. The upper bound in (2.9) shows that some correlation length has to increase (even in a modest way) when a relaxation time grows, but that more complicated correlation functions than a two-point function might be necessary to detect it. The point-to-set correlation functions were later on computed with numerical simulations of model liquids, see for instance [74, 75]. It is fair to say that connections between spatial and temporal correlations is a theme extensively studied in the mathematical physics literature on Markov Chains [76]. One originality of the treatment in [31] relied in the consideration of local measures of relaxation, i.e. in times τ_i a priori distinct for each site, instead of global measures as the gap of the Markov Chain generator. In this way the bounds in (2.9) are valid site by site, which is useful in presence of quenched disorder that can induce very heterogeneous distributions of correlation times.

Coming back to the issue of the efficiency of local search algorithms for solving random CSPs, one can deduce from the above considerations that for any $\alpha > \alpha_d$, a simulated annealing procedure slowly decreasing the temperature will necessarily cross the dynamical transition line at $T_d(\alpha)$ and then fall out-of-equilibrium (assuming the annealing time does not scale exponentially with the system size). One should however not conclude that such an algorithm is unable to reach a solution: even if the probability distribution at the end of the annealing cannot be the uniform measure over solutions, it can well be supported on a subset of the zero energy configurations. For further studies of such thermal annealing the reader is referred to [77, 78], where it is argued that the freezing transition is actually the relevant one up to which simulated annealing might be able to reach a (possibly non-uniform) solution (see also [79] for a more general discussion of the effect of this transition on local search algorithms). A further difficulty for reaching generic conclusions on the efficiency of local search algorithms is the fact that most of them do not respect the detailed

balance condition. Hence the dynamical transition has no reason to be an obstacle for them, and there are indeed evidence for some of them to be successful for clause densities $\alpha > \alpha_d$ [80, 81, 82, 83]. The algorithmic threshold of success is very sensitive to the details of the stochastic procedure for the exploration of the configuration space. The simplest of them, called (Pure Random) WalkSAT, chooses uniformly at random first an unsatisfied constraint a , then one of the variables $i \in \partial a$ of this constraint, and changes the value of σ_i to satisfy a . Numerical and approximate analytical studies of this procedure can be found in [84, 85], while a rigorous lowerbound on its threshold for k -SAT in the large k limit was given in [86].

2.3.2 Belief inspired decimation

We turn now to an analysis of a sequential assignment procedure to solve CSPs, that has been published in [87, 88]. Part of these results were turned into rigorous statements in [89, 90]. This procedure mimics the following idealized one, which assigns the values of the N variables in N steps $T = 1, \dots, N$ (not to be confused with temperature). Let us call $i(1), \dots, i(N)$ the order in which the variables are assigned, $D_T = \{i(1), \dots, i(T)\}$ the set of variables assigned at time T , and $\underline{\tau}_{D_T}$ the values chosen for these variables. One can define the uniform probability measure over the solutions of the CSP which are still compatible with the previous choices, let us call it $\mu_T(\underline{\sigma}) = \mu(\underline{\sigma} | \underline{\sigma}_{D_T} = \underline{\tau}_{D_T})$. If one were able to compute the marginal probabilities of this (conditional) probability measure and use them to draw the value of the newly assigned variable at each step, i.e. draw $\sigma_{i(T+1)}$ with probability $\mu_T(\sigma_{i(T+1)})$ and set $\tau_{i(T+1)} = \sigma_{i(T+1)}$, the Bayes rule would easily show that the output configuration after the N steps is an uniformly drawn random solution, and this would in particular lead to an algorithm for finding one solution of the CSP. The computation of these marginal probabilities $\mu_T(\sigma_i)$ is a computationally intractable task, in practice one has to resort to approximate algorithms, and in this context a natural choice is Belief Propagation (BP), which is often able to compute good approximations of these marginal probabilities. We shall come back on this algorithm in Sec. 2.4, at this point we may only say that it is expected to yield exact predictions for the free-entropy density, in the thermodynamic limit, for constraint densities $\alpha < \alpha_c$. Analyzing the behavior of the Belief inspired decimation procedure thus amounts to control the small (in the large N limit) error which accumulates at each of the N steps by using the BP approximate estimates of the marginal probabilities instead of the exact ones. A theoretical understanding and quantitative description of the deviations between exact and BP-computed marginal probabilities for graphical models is a formidable problem. In the works [87, 88] we followed a slightly simpler route, and studied the following question: assuming that all steps up to some time T have been performed with an hypothetical algorithm able to compute exactly the marginal probabilities required, what is the structure of the support of μ_T , i.e. the shape of the set of solutions compatible with the previous choices ?

The analytical description of the dynamics followed by this ideal process seems very difficult: at each time step the probability of the evolution $\underline{\tau}_{D_{T-1}} \rightarrow \underline{\tau}_{D_T}$ depends in a non-trivial way on all the choices made in the previous steps. However one can study this dynamical process in terms of a series of static problems, thanks to the following elementary observation. At any step T , the Bayes rule can be used to show that $\underline{\tau}_{D_T}$ is distributed according to the marginal of $\mu(\cdot)$. One can state this in a slightly different way: $\underline{\tau}_{D_T}$ can be obtained by drawing uniformly a solution $\underline{\tau}$ from $\mu(\cdot)$, retaining the configuration of the variables in D_T , and discarding the rest of the configuration. We shall further assume that the permutation $i(1), \dots, i(N)$ is drawn uniformly at random, such that D_T is a random set of T variables among N . In the thermodynamic limit we shall define $\theta = T/N$, the




Figure 2.8: Sketch of the phase diagram in the (α, θ) plane for the analysis of the decimation process of random CSP formulas.

fraction of assigned variables, and consider for simplicity that $D_{\theta N}$ is built by retaining independently each variable with probability θ (we only make an error of order $1/\sqrt{N}$ on the fraction of variables thus included in D).

These considerations led us to the definition of an ensemble of CSP instances parameterized by α and θ , generalizing the original one which corresponds to $\theta = 0$. Explicitly this ensemble of formulas corresponds to the following generation process:

1. draw a satisfiable CSP with parameter α
2. draw a uniform solution $\underline{\tau}$ of this CSP
3. choose a set D by retaining each variable independently with probability θ
4. consider the residual formula on the variables outside D obtained by imposing the allowed configurations to coincide with $\underline{\tau}$ on D .

Let us emphasize that, apart from simple cases like the XORSAT model, these ensembles do not coincide in general with randomly uniform formulas conditioned on their degree distributions. The fact that the generation of the configuration $\underline{\tau}$ depends on the initial CSP induces non-trivial correlations in the structure of the final formula.

We could however devise a generalization of the methods used to treat the original CSP ensembles (at $\theta = 0$), and compute phase diagrams for these new ensembles in the (α, θ) plane of parameters. As shown in a schematic way in Fig. 2.8 the phase transitions $\alpha_{d,c}$ extend as transition lines in some region of the phase diagram. More precisely, both lines exhibit a reentrance behavior (they go towards lower values of α when θ increases), down to some values α_* where they meet. In the crescent-shaped region of the phase diagram enclosed by these two lines the measure $\mu_{\theta N}$ is typically split on an exponentially large number of clusters, outside of it is dominated by a single or a sub-exponential number of clusters. Consider for instance some value $\alpha \in [\alpha_d, \alpha_c]$; at $\theta = 0$ there exists exponentially many clusters. However when θ increases more and more variables are revealed and the set of solutions compatible with the assignments in $D_{\theta N}$ shrinks; when the line $\theta_c(\alpha)$ is crossed the single cluster containing the reference solution $\underline{\tau}$ dominates the measure.

One can then try to interpret the behavior of the realistic implementation of this procedure, i.e. the decimation algorithm where the estimation of the probability $\mu(\sigma_{i(T+1)} | \underline{\sigma}_{D_T} = \underline{\tau}_{D_T})$ is made approximately by using Belief Propagation. In the case of XORSAT an exact

analysis of this process can be performed and confronted with the computation of the phase diagram in the (α, θ) plane (see also [91] for a similar study). The comparison shows explicitly that for $\alpha < \alpha_*$ the probability of success of the BP guided decimation is strictly positive in the thermodynamic limit, while it vanishes for $\alpha > \alpha_*$. For k -SAT one has to resort to numerical simulations to compute this probability of success. The results presented in [88] are in reasonable agreement with the conjecture that α_* marks the limit above which BP guided decimation typically fails. Indeed during the execution of the algorithm θ increases from 0 to 1, as symbolized with arrows in Fig. 2.8. For $\alpha > \alpha_*$ the condensation transition is thus crossed at $\theta_c(\alpha)$ after which the marginals computed by BP probably acquire a systematic error. Note that this interpretation, even if consistent with the numerical results, is rather speculative. In addition the analysis crucially relied on the randomness in the order of the assignment of the variables. In practice a modified algorithm where one uses the marginal probabilities computed by BP to set in priority the most biased variables (assuming that the error made by BP is likely to be smaller for these variables) remains successful for larger values of α , and it does not seem obvious to adapt the theoretical analysis to this case. Similarly the decimation based on the informations provided by Survey [25, 63] instead of Belief Propagation (i.e. concentrating on the most numerous clusters and not on the typical solutions) works for larger values of α , close to the satisfiability threshold α_s . As anticipated at the beginning of Sec. 2.3 the performances of algorithms are often very much dependent on fine details in their definitions, and thus difficult to analyze theoretically in complete generality.

From a technical point of view the computations in the (α, θ) plane are similar to those of the Franz-Parisi quenched potential [92]. In both cases one has to deal with two replica configurations of the same system, the first one (denoted $\underline{\sigma}$ here) being thermalized, while the second one ($\underline{\tau}$) is constrained by the first. Here the constraint is that the two configurations coincide on a subset (D) of variables, i.e. they are infinitely coupled on a finite fraction of variables, while in the usual potential the two replica have a finite coupling uniform on all variables (fixing in average the overlap between the two replicas). The technique introduced in [88], that can be easily adapted to the usual Franz-Parisi potential, was, to the best of my knowledge, the first computation of a two replica potential in a finite connectivity mean-field model. These potentials have also been studied in [77, 78] to interpret the annealing dynamics mentioned in Sec. 2.3.1. The decimated ensemble defined above has also been recently re-investigated in the perspective of glasses, see [93, 94] for a discussion of the effect of “random pinning” of particles in a liquid.

2.4 Methods

This section contains an overview of the methods used to derive in a quantitative way the picture of the various phase transitions discussed above. Generalizing the definition of Eq. (2.6), we consider a factor graph model defining a probability measure over the set of configurations χ^N as

$$\mu(\underline{\sigma}) = \frac{1}{Z} \prod_{a=1}^M w_a(\underline{\sigma}_{\partial a}) \prod_{i=1}^N v_i(\sigma_i), \quad Z = \sum_{\underline{\sigma} \in \chi^N} \prod_{a=1}^M w_a(\underline{\sigma}_{\partial a}) \prod_{i=1}^N v_i(\sigma_i), \quad (2.10)$$

where for future convenience we allowed some weights v_i (external fields) to act on the variables. As explained after Eq. (2.6) we consider random ensembles of such models, the indices of the variables involved in an interaction being chosen randomly, with the constraint

that the degree distribution of the variable nodes has to be q_d . We also let the weight functions w_a be independent identically distributed functions (and the same for the v_i 's).

The measure μ and the partition function Z are thus random variables. Thanks to self-averaging (concentration) properties in the thermodynamic limit, the typical features of such a random model are encoded in the average free-entropy

$$\phi = \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}[\ln Z], \quad (2.11)$$

where $\mathbb{E}[\bullet]$ denotes the average with respect to the quenched randomness. For the zero-temperature CSPs this potential is the total entropy density s_{tot} .

Averaging the logarithm of a random variable is in general a difficult task. The replica method bypasses this difficulty by computing moments of order n of the partition function, which is relatively easy for integer n , and by using the identity

$$\mathbb{E}[\ln Z] = \lim_{n \rightarrow 0} \frac{1}{n} \ln \mathbb{E}[Z^n]. \quad (2.12)$$

This method has been originally developed for mean-field fully connected models defined on the complete graph, most notably for the Sherrington-Kirkpatrick model [9]. The continuation from integer values of n (the number of “replicas” of the system) towards $n = 0$ involves a very subtle mechanism of replica symmetry breaking [11], that encodes the complicated landscape of mean-field spin-glasses. The replica method has been extended to deal with finite connectivity models, see in particular [95, 24] and references therein. However the formally equivalent cavity method turned out to be more adapted to finite connectivity models [96], and has a more transparent interpretation that will be described now.

The cavity method relies crucially on the local convergence of random factor graph models to random trees explained at the end of Sec. 2.1.2. Let us assume momentarily that the factor graph representing the model under study is a finite tree. Then the problem of characterizing the measure (2.10) and computing the associated partition function Z can be solved exactly in a simple, recursive way. One introduces on each directed edge $i \rightarrow a$ from one variable i to an adjacent function node a a “message” $\mu_{i \rightarrow a}$, which is a probability measure on the alphabet \mathcal{X} , that would be the marginal probability of σ_i if the interaction a were removed from the graph. These messages are easily seen to obey the recursive (so-called Belief Propagation) equations depicted in Fig. 2.9,

$$\mu_{i \rightarrow a}(\sigma_i) = \frac{1}{z_{i \rightarrow a}} v_i(\sigma_i) \prod_{b \in \partial i \setminus a} \sum_{\underline{\sigma}_{\partial b \setminus i}} w_b(\underline{\sigma}_{\partial b}) \prod_{j \in \partial b \setminus i} \mu_{j \rightarrow b}(\sigma_j), \quad (2.13)$$

with $z_{i \rightarrow a}$ ensuring the normalization of the law $\mu_{i \rightarrow a}$. On a tree factor graph there exists a single solution of these equations, which is easily determined starting from the leaves of the graph (for which the empty product above is conventionally equal to 1) and sweeping towards the inside of the graph. Once the messages $\mu_{i \rightarrow a}$ have been determined all local averages with respect to μ can be computed, as well as the partition function. The Belief Propagation algorithm, that was mentioned in Sec. 2.3.2, consists in looking for a fixed-point solution of (2.13), iteratively, even if the factor graph is not a tree.

Of course random graphs are only locally tree-like, they do possess loops, even if their lengths typically diverge in the thermodynamic limit. The cavity method amounts thus to a series of prescription to handle these long loops and to describe the boundary condition they impose on the local tree neighborhoods inside a large random graph. The simplest prescription, that goes under the name of replica symmetric (RS) and that is valid in the low




Figure 2.9: Graphical representation of Eq. (2.13).

α regime for random CSPs (more precisely $\alpha < \alpha_d$), assumes some spatial correlation decay properties of the probability measure μ . Consider the neighborhood of variable i represented in Fig. 2.9, viewed as a small portion of a large random factor graph. If the interaction c were removed from the graph, the variables j_1 and j_2 would typically be far away (in terms of the graph distance), and because of the correlation decay the probability $\mu_{\setminus c}(\sigma_{j_1}, \sigma_{j_2})$ in the amputated factor graph could be factorized as $\mu_{\setminus c}(\sigma_{j_1})\mu_{\setminus c}(\sigma_{j_2})$. The recursive computation of Eq. (2.13), exact on a tree, would thus be asymptotically exact in a random factor graph with fast enough spatial correlation decay. To compute the average thermodynamic potential (2.11) it is enough in this case to study the statistics with respect to the quenched disorder of the messages $\mu_{i \rightarrow a}$ on the edges of the random factor graph. In other words the order parameter of the RS cavity method is $\mathcal{P}(\eta)$, the probability (over the disorder) that the messages $\mu_{i \rightarrow a}$ in Eq. (2.13) (which are themselves probability distributions on \mathcal{X}) are equal to η . \mathcal{P} obeys a self-consistent functional equation, which is more simply written in a distributional form as

$$\eta \stackrel{d}{=} g(\eta_{1,1}, \dots, \eta_{1,k-1}, \dots, \eta_{d,1}, \dots, \eta_{d,k-1}, v, w_1, \dots, w_d) . \quad (2.14)$$

In this equation all the η 's are drawn independently from \mathcal{P} , and $\stackrel{d}{=}$ denotes the equality in distribution between random variables. Moreover d is drawn according to the law \tilde{q}_d , the v and w_a 's are independent copies of the variable and interaction random weights, and the function g in the r.h.s. is defined by

$$\eta(\sigma) = \frac{1}{z(\{\eta_{a,i}\}, v, \{w_a\})} v(\sigma) \times \sum_{\{\sigma_{a,i}\}_{a \in [1,d]}^{i \in [1,k-1]}} \left(\prod_{a,i} \eta_{a,i}(\sigma_{a,i}) \right) \prod_{a=1}^d w_a(\sigma, \sigma_{a,1}, \dots, \sigma_{a,k-1}) , \quad (2.15)$$

$z(\{\eta_{a,i}\}, v, \{w_a\})$ ensuring the normalization of η . The RS prediction for ϕ can then be expressed as a functional of \mathcal{P} , more precisely as the average over copies η_i drawn from \mathcal{P} of some simple functions obtained by analogy from the exact computation of the partition function of a finite tree. Note that the equation (2.14), if it has in general no analytic solution,

lends itself to a very natural numerical resolution where \mathcal{P} is approximately represented as an empirical distribution over a set of representatives η (a population representation [97, 96]).

The assumption of fast correlation decay that underlies the RS cavity method can fail in presence of frustration, for instance in the case of random CSPs with $\alpha > \alpha_d$, the clustering transition (we announced in Sec. 2.3.1 the interpretation of this transition as the birth of long-range point-to-set correlations). Indeed the configuration space of these models gets fractured in a large number of pure states (or clusters), and the correlation decay hypothesis only holds for the Gibbs measure restricted to one pure state, not for the complete Gibbs measure. This complication can be handled by the cavity method with “replica symmetry breaking” (RSB). It amounts to make further self-consistent hypotheses on the organization of these pure states, and on the correlated boundary conditions they induce on the tree-like portions of the factor graph. Inside each pure state the RS computation holds true, the RSB computation is then a study of the statistics of the pure states. Let us explain how this is done in practice at the first level of RSB (1RSB cavity method). The partition function is written as a sum over the pure states γ , that form a partition of the configuration space, $Z = \sum_{\gamma} Z_{\gamma}$, where Z_{γ} is the partition function restricted to the pure state γ . It can be written in the thermodynamic limit as $Z_{\gamma} = e^{N\phi_{\gamma}}$, with ϕ_{γ} the internal free-entropy. As in Sec. 2.2.2 one further assumes that the number of pure states with a given value of ϕ is, at the leading exponential order, $e^{N\Sigma(\phi)}$, with the complexity Σ a concave function of ϕ , positive on the interval $[\phi_{\min}, \phi_{\max}]$. In order to compute Σ one introduces a parameter m (called Parisi breaking parameter) conjugated to the internal thermodynamic potential, and the generating function of the Z_{γ} as $Z(m) = \sum_{\gamma} Z_{\gamma}^m$. In the thermodynamic limit its dominant behavior is captured by the 1RSB potential $\Phi(m)$,

$$\Phi(m) = \lim_{N \rightarrow \infty} \frac{1}{N} \log Z(m) = \sup_{\phi} [\Sigma(\phi) + m\phi] , \quad (2.16)$$

where the last expression is obtained by a saddle-point evaluation of the sum over γ . The complexity function is then accessible via the inverse Legendre transform of $\Phi(m)$ [98], or in a parametric form

$$\phi(m) = \Phi'(m) , \quad \Sigma(\phi(m)) = \Phi(m) - m\Phi'(m) , \quad (2.17)$$

where $\phi(m)$ denotes the point where the infimum is reached in Eq. (2.16). One has $\Sigma'(\phi(m)) = -m$, i.e. the introduction of the parameter m allows to explore the complexity curve as shown in Fig. 2.4 by tuning the tangent slope of the selected point.

The actual computation of $\Phi(m)$ is done as follows. One introduces on each edge of the factor graph a distribution $P_{i \rightarrow a}(\eta)$ of messages, which is the probability over the different pure states γ , weighted proportionally to Z_{γ}^m , that $\mu_{i \rightarrow a}^{\gamma} = \eta$, where $\mu_{i \rightarrow a}^{\gamma}$ are the messages that appear in Eq. (2.13), for the measure restricted to the pure state γ . Because $P_{i \rightarrow a}(\eta)$ fluctuates from instance to instance, the order parameter now becomes the distribution of $P_{i \rightarrow a}(\eta)$ with respect to the disorder, which we call $\mathcal{P}^{(1)}(P)$ and is solution of a self-consistent functional equation written as

$$P \stackrel{d}{=} G(P_{1,1}, \dots, P_{1,k-1}, \dots, P_{d,1}, \dots, P_{d,k-1}, v, w_1, \dots, w_d) . \quad (2.18)$$

Similarly to the RS case the P 's are independent copies drawn from $\mathcal{P}^{(1)}$, and the random

weights v and w_a are also independently generated. The r.h.s. of this equation stands for:

$$P(\eta) = \frac{1}{Z(\{P_{a,i}\}, v, \{w_a\}, m)} \times \int \prod_{\substack{a \in [1,d] \\ i \in [1,k-1]}} dP_{a,i}(\eta_{a,i}) \delta(\eta - g(\{\eta_{a,i}\})) z(\{\eta_{a,i}\}, v, \{w_a\})^m, \quad (2.19)$$

with g and z defined above in Eq. (2.15), and m is the Parisi parameter. From the solution of this equation (that again can be found numerically with the population dynamics method [96]) one computes the 1RSB potential $\Phi(m)$ via an expression similar to the one giving the expression of ϕ at the RS level, with now averages over $\mathcal{P}^{(1)}$.

There are different justifications for the appearance of the ‘‘reweighting factor’’ $z(\{\eta_{a,i}\}, v, \{w_a\})^m$ in Eq. (2.19). The argument in [96] is based on the exponential distribution of the free-entropy $N\phi_\gamma$ of the pure states with respect to some reference value, and on consistency requirements on the evolution of the pure states when the cavity factor graph is modified. One can also study the statistics of the many fixed point solutions of the BP equations (2.13) and devise a dual factor graph for the counting of these fixed points [23], the reweighting factor allowing to select the fixed points associated to some internal free-entropy. Another interpretation was proposed in [26], associating the pure states of a large but finite factor graph model to boundary conditions on trees.

The thresholds α_d and α_c (or their finite temperature counterparts) separate regimes where the 1RSB equations have different kind of solutions.

- It can happen that only trivial solutions of (2.18) exist, i.e. the P 's are supported on a single value of η . This is the translation of the existence of a single pure state, in which case the whole 1RSB machinery reduces to the RS case. This case is realized at high temperatures/low connectivity parameter α , on the left of the dynamical transition line in Fig. 2.7.
- If on the contrary non-trivial solutions of the 1RSB equations appear, one has to investigate them more carefully in order to obtain the 1RSB prediction for the free entropy density. From the definition of $Z(m)$ one would naturally take $\Phi(1)$ for it. This is indeed the case, provided the corresponding complexity $\Sigma(\phi(m=1))$ is positive. In this case, corresponding to the regime of (α, T) between the two lines of Fig. 2.7, an exponential number $e^{N\Sigma(\phi(1))}$ of pure states contribute to the Gibbs-Boltzmann measure, each with an internal free-entropy $\phi(m=1)$. It turns out that in this case the prediction $\Phi(m=1)$ coincides with the RS one, hence the analyticity of the thermodynamic properties upon crossing the line $T_d(\alpha)$, that has however drastic consequences on the dynamics as discussed in Sec. 2.3.1. Such a phase is usually called for this reason a dynamic 1RSB (d1RSB) phase.
- If there are non-trivial solutions of the 1RSB equations, but with $\Sigma(\phi(1)) < 0$, the system is in a condensed phase, corresponding to the rightmost part of the phase diagram in Fig. 2.7. By definition the number of clusters is expressed in terms of the complexity Σ as $e^{N\Sigma}$: if the complexity at $m=1$ is negative there are typically no clusters with internal free-entropy $\phi(1)$, its value corresponds to large deviation events with rare disorder realizations [99]. As depicted in the right panel of Fig. 2.4 the thermodynamic properties are dominated by the few largest clusters, one has thus to find the value m_s of the Parisi parameter such that $\Sigma(\phi(m_s)) = 0$, m_s is thus equal to minus the tangent slope of $\Sigma(\phi)$ in ϕ_{\max} .

We can now clarify the technical remark made in Sec. 2.2.3 on the criterion of existence of frozen variables in the clustered phase. A “hard field” is a probability distribution on χ that allows a single value of the variable, i.e. of the form $\eta(\sigma) = \delta_{\sigma, \sigma_0}$. Typical solutions of a CSP will thus have an extensive number of frozen variables in their respective clusters whenever the solution of the 1RSB equation (2.18), for the relevant value m_s of the Parisi parameter ($m_s = 1$ in a d1RSB phase), gives a non-zero probability to hard fields. However there can be in general non-trivial solutions of (2.18) without any weight on the hard fields, as happens for k -SAT and q -COL in the regime $\alpha \in [\alpha_d, \alpha_f]$.

Even though the main focus chosen in this manuscript was the application of statistical mechanics methods to random CSPs of computer science, it should be clear that the method sketched here applies more generally to any factor graph model that converges locally to a tree. It can thus be applied to other problems, in particular some emerging from random graph theory. One can for instance count the typical number of copies of some large subgraph inside a random graph. This strategy was applied to the counting of matchings [100], i.e. non-overlapping dimers, cycles [101] and k -regular subgraphs [102].

Many steps in the derivation of quantitative results on sparse random graphs by the cavity method are heuristic and far from mathematical rigor. Some of its predictions have however received rigorous confirmations. One can distinguish two main paths that have been followed to reach them. The first one is an extension to the finite-connectivity setting of the interpolation method of Guerra and Toninelli [14], first developed for fully connected models, that led Talagrand to the proof of the Parisi formula for the free energy of the Sherrington-Kirkpatrick model [13]. The finite-connectivity version of the interpolation method was first treated in [103, 104], where bounds on the free-energy of sparse mean-field models could be given in terms of the replica/cavity method predictions, including replica symmetry breaking ones. The second approach is closer in spirit to the heuristic version of the cavity method, as it exploits directly the local convergence of random graphs to trees, and amounts to prove rigorously the spatial correlation decay assumed in the heuristic version. Two examples of this approach are [105], where ferromagnetic Ising models are shown to obey the predictions of the cavity method, and [106], where the results on the number of matchings first obtained in [100] are recovered. Note that in both cases the model is replica-symmetric; obtaining the local convergence of μ towards a 1RSB measure remains a challenging open problem.

Chapter 3

The effect of quantum fluctuations

In this chapter I will present my recent works on quantum disordered models, which were motivated by applications both to computing issues and to more physical problems. In a first part (Sec. 3.1) I will briefly introduce some notions on quantum computing, then in Sec. 3.2 I will discuss one method which was developed to study how random optimization problems would be solved by a quantum computer, and finally show some results obtained in this way (in Sec. 3.3). A more extensive discussion can be found in the review [43].

3.1 Introduction to (adiabatic) quantum computing

The classical computational complexity theory classifies the difficulty of computational tasks in terms of the time (number of steps) and space (memory usage) requirements of algorithms that perform these tasks. One could think that these requirements are very much dependent on the actual device on which they are implemented. This is fortunately not the case: up to polynomial reductions these are universal quantities, and one can use a simplified model of a computer, known as a Turing machine, to formalize the notion of computation. This universality result is the content of the Church-Turing hypothesis.

Quantum computers are however a class of devices that fall out of this paradigm, because their basic functioning is ruled by the laws of quantum mechanics, instead of the classical vision that underlies the Turing machine modelization. In consequence quantum computers are potentially more efficient than classical ones, and a parallel classification in quantum complexity classes has been developed [37, 107, 108, 109, 110] (to name only two main classes, P and NP are respectively replaced by BQP and QMA in the quantum context).

The paradigmatic shift from classical computer science is the assumption that the elementary variables at the core of the computer behave quantumly: instead of bits which can take either the value 0 or the value 1, one deals with qubits which can be in a coherent superposition of the two values. More generically the classical configuration space χ^N is replaced by an Hilbert space spanned by the orthonormal basis $\{|\underline{\sigma}\rangle : \underline{\sigma} \in \chi^N\}$, the so-called computational basis. According to the laws of quantum mechanics the state of the system is described by a vector $|\psi\rangle$ of this Hilbert space, i.e. a (complex) linear combination of the vectors of the computational basis, which has norm 1. The state of the computer evolves during the execution of an algorithm; quantum mechanics teaches that this evolution is represented by the action of a linear operator on the Hilbert space, $|\psi\rangle \rightarrow \widehat{U}|\psi\rangle$, where the

linear operator \widehat{U} must be unitary in order to conserve the norm of $|\psi\rangle$. Every quantum algorithm thus corresponds to an unitary operator; in principle this operator acts on all the qubits of the system, making a practical implementation of non-trivial algorithms a seemingly impossible task. Fortunately it has been shown [111, 112, 113, 114] that any unitary operator can be factorized (with arbitrary precision) as a product of simple operators, called gates in this context, that act only on one or two qubits (this is to some extent similar, in the classical case, to the reducibility of any Boolean function to a combination of NotAND gates). Moreover there exist universal sets of gates that contain only a finite number of operators. The complexity of a quantum algorithm can thus be quantified as the number of elementary gates it involves.

Even if a quantum computer can be in a superposition of all the classical states, this does not mean that it can treat its exponentially many inputs in parallel (as a non deterministic Turing machine), because the output must be read out according to the quantum law of measurement. To state clearly this important point let us consider a binary function $f(\underline{\sigma})$ from $\{0, 1\}^N$ to $\{0, 1\}^M$. In the quantum setting this function is implemented as an unitary linear operator \widehat{U}_f ; note that an unitary transformation is invertible, hence \widehat{U}_f must somehow keep trace both of the input $\underline{\sigma}$ and of the output $f(\underline{\sigma})$ of the function f . A convenient way to fulfill this request is to let \widehat{U}_f act on the Hilbert space of $N + M$ qubits, its action being defined on the computational basis as

$$\widehat{U}_f|\underline{\sigma}, \underline{\sigma}'\rangle = |\underline{\sigma}, \underline{\sigma}' \oplus f(\underline{\sigma})\rangle , \quad (3.1)$$

where \oplus is here the bitwise addition modulo 2. Suppose that the quantum computer is initially prepared in the state

$$\frac{1}{2^{N/2}} \sum_{\underline{\sigma} \in \mathcal{X}^N} |\underline{\sigma}, \underline{0}\rangle , \quad (3.2)$$

where $\underline{0} = (0, \dots, 0)$. Then \widehat{U}_f transforms this state into

$$|\psi\rangle = \frac{1}{2^{N/2}} \sum_{\underline{\sigma} \in \mathcal{X}^N} |\underline{\sigma}, f(\underline{\sigma})\rangle , \quad (3.3)$$

which seems indeed to contain all the information about the behavior of f on its 2^N possible inputs. However this information is not reachable by an observer, because of the measurement axiom of quantum mechanics. For the non-physicist reader, a quantum mechanical observable is an Hermitian operator \widehat{O} of the Hilbert space, with distinct eigenvalues O_1, \dots, O_n , and associated eigenprojectors $\widehat{P}_1, \dots, \widehat{P}_n$. If the system is, right before the measurement, in the state $|\psi\rangle$ (a normalized vector of the Hilbert space), the measured value of \widehat{O} will be one of the eigenvalues O_i , with probability $|\langle \psi | \widehat{P}_i | \psi \rangle|^2$. In that case the state of the system right after the measurement is proportional to $\widehat{P}_i |\psi\rangle$, i.e. the state is projected to the eigenspace corresponding to the outcome of the measurement. A measurement of the qubits in the state $|\psi\rangle$ written above thus leads to nothing but a random choice of a single configuration $\underline{\sigma}$ among the 2^N possible ones, and the associated value $f(\underline{\sigma})$, with absolutely no gain with respect to a classical computation. This trivial observation explains why some thought has to be put in devising quantum algorithms that outperform classical ones; the availability of linear superpositions is not enough for that, one has to use in a more clever way the possibility of interference between states.

Several quantum algorithms have been devised in this framework, most of them dealing with arithmetic problems. This line of research culminated with the discovery by Shor [38] of

a quantum algorithm that factors integers in a time growing polynomially with the number of bits of the integer in input. This is a breakthrough with respect to classical algorithms that all require a superpolynomial time for this task. It would have deep practical consequences if a large quantum computer was built, as it destroys the security of some widely used cryptographic protocols, based precisely on the difficulty of integer factoring. However this problem is not known to be NP-complete, hence the existence of Shor’s algorithm does not imply a “quantum collapse” of the whole NP class.

We shall now consider a generic quantum algorithm to solve any optimization problem, namely the quantum annealing or quantum adiabatic algorithm [39, 40, 41, 42], that exploits quantum fluctuations to find low-energy configurations of an Hamiltonian, in a similar way as simulated annealing [115] uses thermal fluctuations for the same purpose. The state of a quantum system evolves in time according to the Schrödinger equation,

$$i \frac{d}{dt} |\psi(t)\rangle = \widehat{H}(t) |\psi(t)\rangle, \quad (3.4)$$

in units where $\hbar = 1$, $\widehat{H}(t)$ being the time-varying Hamiltonian of the system. The quantum annealing algorithm corresponds to a specific choice of the time dependence of the Hamiltonian, that is assumed to be at the disposal of the algorithm designer. To be more precise, to a cost function $E(\underline{\sigma})$ defined on the space of classical configurations we associate an operator \widehat{H}_f , diagonal in the computational basis, with $\widehat{H}_f|\underline{\sigma}\rangle = E(\underline{\sigma})|\underline{\sigma}\rangle$. We also consider another Hermitian operator \widehat{H}_i , that has non-zero off-diagonal elements in the computational basis, and such that one can easily prepare the system at time $t = 0$ in the groundstate of \widehat{H}_i . In the case of Ising spins the natural choice for \widehat{H}_i is a uniform transverse field $-\sum_i \widehat{\sigma}_i^x$, such that $\langle \underline{\sigma} | \widehat{H}_i | \underline{\sigma}' \rangle$ is equal to -1 if and only if $\underline{\sigma}$ and $\underline{\sigma}'$ are at unit Hamming distance one from the other. Now the quantum adiabatic algorithm, for an execution time \mathcal{T} , corresponds to the Schrödinger evolution (3.4) with an Hamiltonian interpolating linearly between these two operators, $\widehat{H}(t) = (1 - \frac{t}{\mathcal{T}}) \widehat{H}_i + \frac{t}{\mathcal{T}} \widehat{H}_f$ with $t \in [0, \mathcal{T}]$. This evolution defines an unitary operator $\widehat{U}_{\mathcal{T}}$ such that $|\psi(\mathcal{T})\rangle = \widehat{U}_{\mathcal{T}} |\psi(0)\rangle$. In the limit $\mathcal{T} \rightarrow \infty$, i.e. for infinitely slow interpolations, this operator becomes an “intertwiner” that associates to each eigenvector of \widehat{H}_i an eigenvector of \widehat{H}_f , respecting the order of the eigenvalues. In particular if $|\psi(0)\rangle$ is the groundstate of \widehat{H}_i , then the final state of the algorithm is the groundstate of \widehat{H}_f , and the classical optimization problem is solved, as a projective measurement of $|\psi(\mathcal{T})\rangle$ in the computational basis outputs one of the configurations minimizing $E(\underline{\sigma})$. Of course an algorithm taking an infinite time is of little practical use. Fortunately for large but finite times $\widehat{U}_{\mathcal{T}}$ remains close to its limit. A more quantitative condition is provided by the adiabatic theorem [116]; roughly speaking, the evolution time should be larger than the inverse square of the minimal gap between the two lowest eigenvalues of $\widehat{H}(t)$ along the interpolation for the groundstate of \widehat{H}_i to be transformed in a state close to the groundstate of \widehat{H}_f .

The computational complexity of solving optimization problems via quantum annealing is thus reduced to the determination of the minimal gap between the groundstate and the first excited state along the interpolation $\widehat{U}_{\mathcal{T}}$. Generically a quantum phase transition in the thermodynamic limit of a model is associated to a vanishing of this gap. For what concerns typical random optimization problems, one can expect that the classical phase transitions discussed above translate into quantum phase transitions as a function of the interpolating parameter. Indeed quantum phase transitions have been known to occur for fully-connected mean-field quantum spin glasses for a while [117, 118, 119, 120, 121]. This motivated the development of a quantum extension of the cavity method, able to produce quantitative predictions for the typical behavior of random CSPs in a transverse field.

3.2 Quantum cavity method

For the reasons exposed above we want to determine the spectrum of an Hermitian operator \hat{H} on the Hilbert space spanned by $\{|\underline{\sigma}\rangle, \underline{\sigma} \in \chi^N\}$, with diagonal elements given by a cost function $E(\underline{\sigma})$ and off-diagonal elements encoding some quantum fluctuations. Equivalently the eigenvalues can be obtained from their generating function, the partition function $Z = \text{Tr} e^{-\beta\hat{H}}$. For simplicity we shall discuss only the case of Ising spins, with the off-diagonal part given by a transverse field of intensity Γ , i.e. the off-diagonal part of \hat{H} is $-\Gamma \sum_{i=1}^N \hat{\sigma}_i^x$. This quantum computation can in fact be turned into a classical one, by modifying the nature of the basic degrees of freedom. We shall first state the result and explain later on its derivation. One finds that

$$Z = \int_{\underline{\sigma}(0)=\underline{\sigma}(\beta)} \left[\prod_{i=1}^N D\sigma_i v_i(\sigma_i) \right] e^{-\int_0^\beta E(\underline{\sigma}(t)) dt} . \quad (3.5)$$

Several new notations have been introduced: bold symbols denote functions of an ‘‘imaginary’’ time $t \in [0, \beta]$; in particular here σ_i is a piecewise constant function $\sigma_i(t) : [0, \beta] \rightarrow \{-1, +1\}$. The integration measure $D\sigma_i$ is decomposed as a sum over the number n of discontinuities of the function $\sigma_i(t)$, the times $t_1 \leq \dots \leq t_n$ at which they occur, and the value σ_i^0 at initial time (then $\sigma_i(t)$ flips to its opposite value at each discontinuity):

$$\int D\sigma_i \equiv \sum_{n=0}^{\infty} \sum_{\sigma_i^0} \int_0^\beta dt_1 \int_{t_1}^\beta dt_2 \dots \int_{t_{n-1}}^\beta dt_n . \quad (3.6)$$

For such a function the weight $v_i(\sigma_i)$ reads $v_i(\sigma_i) = \Gamma^n$. Such a path integral representation can be devised for any matrix element of $e^{-\beta\hat{H}}$, by simply fixing the initial and final configurations $\underline{\sigma}(0)$ and $\underline{\sigma}(\beta)$. Here we let them free, under the condition $\underline{\sigma}(0) = \underline{\sigma}(\beta)$, to compute the trace of $e^{-\beta\hat{H}}$. As $\sigma_i(0) = \sigma_i(\beta)$ only even values of the number n of discontinuities in the trajectory $\sigma_i(t)$ do contribute.

There are several ways to obtain such a path integral representation. The most pedestrian one is to use the Suzuki-Trotter formula, decomposing the imaginary time interval $[0, \beta]$ in N_s slices,

$$e^{\hat{X}_1 + \hat{X}_2} = \lim_{N_s \rightarrow \infty} \left(e^{\frac{1}{N_s} \hat{X}_1} e^{\frac{1}{N_s} \hat{X}_2} \right)^{N_s} , \quad (3.7)$$

with the two non-commuting operators $\hat{X}_1 = -\beta \sum_{\underline{\sigma}} E(\underline{\sigma}) |\underline{\sigma}\rangle \langle \underline{\sigma}|$ and $\hat{X}_2 = \beta \Gamma \sum_i \hat{\sigma}_i^x$. One can then introduce N_s representations of the identity between the terms of the product, as sums over the configurations $\underline{\sigma}(t = \alpha\beta/N_s)$ with α a discrete time index. The continuous time limit $N_s \rightarrow \infty$ then yields the path integral (3.5). A maybe more elegant way to obtain this result is to use the following operator identity,

$$e^{\hat{X}_1 + \hat{X}_2} = \sum_{p=0}^{\infty} \int_0^1 dt_1 \int_{t_1}^1 dt_2 \dots \int_{t_{p-1}}^1 dt_p e^{t_1 \hat{X}_1} \hat{X}_2 e^{(t_2 - t_1) \hat{X}_1} \hat{X}_2 \dots \hat{X}_2 e^{(1 - t_p) \hat{X}_1} ,$$

with the same values of $\hat{X}_{1,2}$ as above; one then inserts p representations of the identity in the p -th term of the sum, rescales the time integrals and reorders the summation according to the number of times each flipping operator $\hat{\sigma}_i^x$ is picked in the expansion of \hat{X}_2 . Finally these path integral representations can also be handled in a mathematically rigorous way,

see for instance [122, 123, 124, 125, 126, 127, 128], by starting from Poisson point processes for the candidate times of discontinuities of the variable trajectories, a distribution which is then properly biased by the classical energy term.

Suppose now that the classical energy function is decomposed along a factor graph as $E(\underline{\sigma}) = \sum_{a=1}^M \varepsilon_a(\underline{\sigma}_{\partial a})$. The quantum partition function can then be rewritten as

$$Z = \sum_{\underline{\sigma} \in \widehat{\mathcal{X}}^N} \delta_{\underline{\sigma}(0), \underline{\sigma}(\beta)} \prod_{a=1}^M w_a(\underline{\sigma}_{\partial a}) \prod_{i=1}^N v_i(\sigma_i), \quad w_a(\underline{\sigma}_{\partial a}) = e^{-\int_0^\beta dt \varepsilon_a(\underline{\sigma}_{\partial a}(t))}, \quad (3.8)$$

where we introduced $\widehat{\mathcal{X}}$ the space of periodic piecewise constant functions from $[0, \beta]$ to $\{-1, +1\}$, and used the notation $\sum_{\sigma_i \in \widehat{\mathcal{X}}}$ as a synonym of the integration $\int D\sigma_i$ defined in Eq. (3.6). This notation was chosen to emphasize the similarity with the classical partition function (2.10): the quantum computation is reduced to a classical one, the cost to be paid being the replacement from a discrete variable σ_i in \mathcal{X} to a function (trajectory) $\sigma_i = \{\sigma_i(t) | t \in [0, \beta]\}$ in $\widehat{\mathcal{X}}$ as basic degrees of freedom. Note that apart from this replacement, the ‘‘spatial’’ structure of the interactions w_a encoded in the factor graph is the same in the classical and in the quantum case. In particular as soon as the classical energy part of the quantum Hamiltonian falls into the category of models that can be solved by the cavity method (i.e. sparse random graphs), then this is also true for the quantum problem. This observation was first exploited in [129] to study the quantum spin glass on a random regular graph, at the replica symmetric level and within a finite number of Suzuki-Trotter slices. The formulation of the quantum cavity method in continuous imaginary time was then presented in [130] at the RS level, for a ferromagnetic model, and in [131] for the Bose-Hubbard model of bosonic particles. Results at the 1RSB level were given in [132] for quantum XORSAT problems on random regular graphs (see next subsection for some of these outcomes), and in [133] for interacting particle models. A more complete presentation of this series of works, treating the case of generic χ , as well as further original results on a quantum version of the coloring problem, can be found in the review [43]. Of course technical complications do arise when replacing a discrete spin σ_i by a function $\sigma_i(t)$, but conceptually the quantum cavity method remains very close to the classical one; in fact the number of discontinuities in the trajectory $\sigma_i(t)$ is of order $\beta\Gamma$, and as long as this parameter remains bounded the description of a trajectory in terms of the times of its discontinuities is relatively compact and manageable numerically. This estimation shows on the other hand the main drawback of this approach: the zero temperature ($\beta \rightarrow \infty$) limit cannot be taken explicitly here. Alternative, approximate, versions of the quantum cavity method that bypass this limitation can be found in [134, 135, 136, 137, 138, 139, 140, 141].

3.3 Quantum annealing of random CSPs

To assess the time requirements of the quantum adiabatic algorithm for the resolution of typical CSPs one has to evaluate the scaling of the minimal gap along its interpolation path, i.e. when the transverse field Γ controlling the intensity of the quantum fluctuations is reduced from a very large value down to zero. Quantum phase transitions, where this gap closes in the thermodynamic limit, are thus a major issue in this perspective; numerous studies of mean-field fully connected quantum spin-glasses [117, 118, 119, 120, 121], which share a large part of their phenomenology with random CSPs, have unveiled quantum phase transitions as a function of the transverse field. With this application in mind we have used the quantum cavity method to study several models in [132, 133, 43, 142], and also studied




Figure 3.1: Left panel: phase diagram for the 3-XORSAT model on 3-regular random graphs with a transverse field Γ . The vertical line at $\Gamma = 1$ marks the location of a first-order transition between a (dynamic) classical and a quantum paramagnet phase, the two lines around it are the spinodal limit of existence of these two phases. The line starting from $\Gamma = 0$ is the clustering (dynamic) 1RSB transition. Right panel: numerical results for the minimum gap Δ_{\min} of this model, showing its exponential scaling as a function of the size N of the instances. Full black points represent the average of Δ_{\min} obtained by exact diagonalization up to $N = 24$ (error bars are of the order of the symbol size except when explicitly shown for $N = 24$). Dashed bars represent the standard deviation of a single realization of the random variable Δ_{\min} . Open blue diamonds are Quantum Monte Carlo data for the median minimal gap from [145]. Here again error bars are of the order of symbol size. Dashed line is a fit to $\Delta_{\min}(N) = 0.911 \exp(-0.081N)$.

the quantum version of a toy model [143] which captures classically [144] the clustering and condensation phenomenon. To keep this manuscript within the required length we shall only discuss briefly the results for one of these models and mention the other cases afterwards.

For technical reasons models defined on random regular (hyper)graphs, in which all variables are connected to the same number of interactions, are simpler to analyze, while exhibiting the same phenomenology as models on Erdős-Rényi random graphs. For this reason we concentrated in [132] on random regular XORSAT formulas, of different connectivities. The results presented in Fig. 3.1 are for the simplest non-trivial case we encountered, with the same degree for the interactions and for the variables (all interactions involve $k = 3$ variables, all variables appear in 3 interactions). Classically this model exhibits a dynamical (clustering) transition at some temperature T_d , with a complexity at $m = 1$ that remains positive for all positive temperature, and vanishes at $T = 0$. In other words the Kauzmann temperature vanishes in this case, the whole low-temperature phase $T \leq T_d$ is a d1RSB phase, with thermodynamic properties described by the RS cavity method. The effect of the transverse field Γ in this model is summarized in the phase diagram on the left panel of Fig. 3.1. The most important phenomenon is the appearance of a thermodynamic first order transition, at low enough temperature, at the critical value of the field $\Gamma = 1$ (this value was determined, within numerical accuracy, in [132], and later shown to be exactly 1 with a duality argument between high and low values of Γ [146, 145]). At this point the first derivative of the free-energy with respect to Γ (proportional to the transverse magnetization) has a discontinuity. This is the trace of the crossing of the free-energies of two distinct phases, the (dynamic) classical paramagnet which emerges continuously from the classical model at $\Gamma = 0$, and the quantum paramagnet which is the continuation of the

phase at $\Gamma = \infty$ with all spins aligned along the transverse field. As usual in mean-field these two phases coexist in a region of the phase diagram, delimited by the spinodals also drawn on the phase diagram. The quantum cavity method does not give access directly to the zero-temperature limit, however by extrapolating the free-energies obtained at small temperatures one obtains a fairly good estimate of the groundstate energy per spin, and of the location of the zero-temperature quantum phase transition. It is generically expected that first-order transition are accompanied by an exponentially small gap with respect to the size of the system. This is indeed the case in this model, as demonstrated in the right panel of Fig. 3.1 with data obtained numerically for small system sizes. The phase diagram on the left of the figure also displays the dynamical (clustering) transition line $T_d(\Gamma)$ that we computed with the 1RSB quantum cavity method. On this line the structure of the measure on imaginary-time trajectories change, with no effect on the thermodynamics of the model.

The other models considered in [132, 143, 133, 43, 142] display some variety in the type of quantum phase transitions encountered upon lowering the transverse field Γ . In some of them it remains thermodynamically of first order, but towards a spin-glass phase emerging from a condensed situation in the classical limit $\Gamma = 0$. It can also occur that the quantum phase transition is of second order, but towards a complex spin-glass phase with a continuum of avoided crossings because of a competition between clusters that have classically a higher energy but which are bigger, and hence have an energy that decrease faster with the transverse field. This scenario can be explicitly checked in the toy model studied in [143], but this mechanism is probably at play in generic random CSPs in presence of quantum fluctuations [43, 142], and implies an exponentially small gap between the groundstate and the first excited state in the whole low Γ phase, below the phase transition. A seemingly similar phenomenon of “perturbative avoided crossings” was studied in [147, 148, 149, 150].

These statements sound rather negative for the interest of the quantum adiabatic algorithm, as it appears that for most types of CSPs an exponentially small gap will be encountered during the annealing of typical instances, hence the scaling of its execution time \mathcal{T} should be exponential in the size of the system to reach the groundstate of the classical problem. This contrasts the hope, triggered by numerical simulations on small size systems [42], that the quantum adiabatic algorithm could solve in polynomial time problems for which all classical algorithms take an exponential time. This negative conclusion was probably anticipated for a long-time by most practitioners of mean-field disordered systems, from the intuition gained on fully-connected quantum spin-glasses [117, 118, 119, 120, 121], demonstrating it on finite-connectivity models more closely related to optimization problems might help to export this opinion to other communities.

Despite this caveat there might be a strong potential for the quantum annealing to solve approximation problems in a more efficient way than classical algorithms. To explain this statement one should recall that it was proven in classical computational complexity theory that for some models, including k -SAT or k -XORSAT, it is hard to find an approximate estimate of the groundstate energy asymptotically better than the one obtained by picking at random a configuration [151]. If one runs the annealing algorithm for a time \mathcal{T} smaller than the adiabatic time (i.e. the inverse square of the minimal gap), the classical energy reached at the end, call it $E_{\text{fin}}(\mathcal{T})$, will certainly be greater than the groundstate one E_{gs} . However one could hope that for some models the scaling of the residual energy $E_{\text{fin}}(\mathcal{T}) - E_{\text{gs}}$, for large but sub-adiabatic times \mathcal{T} , would lead to better estimates of E_{gs} in a “reasonable” execution time of the quantum annealing.

Unfortunately there does not appear to be generic tools for the quantitative computation of $E_{\text{fin}}(\mathcal{T})$ in the non-adiabatic regime. To investigate this question we considered in [152] a very simple model, fully-connected and ferromagnetic, known as the p -spin quantum Curie-




Figure 3.2: Spectrum of the model (3.9) for $p = 3$, in the completely symmetric (maximal spin) sector. The right panel is a zoom around the first-order transition, that shows the metastable continuations of the groundstates from both phases, and the spinodal of the large s phase. All crossings are avoided but the exponentially small gaps are not distinguishable on the figure.

Weiss model (and related to the Lipkin-Meshkov-Glick Model [153]). Its Hamiltonian is given by

$$\hat{H}(s) = -Ns(\hat{m}^z)^p - N(1-s)\hat{m}^x, \quad \hat{m}^x = \frac{1}{N} \sum_{i=1}^N \hat{\sigma}_i^x, \quad \hat{m}^z = \frac{1}{N} \sum_{i=1}^N \hat{\sigma}_i^z, \quad (3.9)$$

where the parameter s is related to the transverse field Γ via $\Gamma = (1-s)/s$; $\hat{\sigma}_i^z$ is defined by its action in the computational basis, $\hat{\sigma}_i^z |\underline{\sigma}\rangle = \sigma_i |\underline{\sigma}\rangle$. The operator $\hat{H}(s)$ is invariant under all permutation of site indices, which allows to diagonalize it in subspaces classified by representations of the symmetry group (in more physical terms, by the total spin operator that commutes with \hat{H}). The lowest energy states are found in the maximally symmetric subspace, of dimension $N+1$, the eigenvalues in this sector are plotted in Fig. 3.2 for $p=3$, as a function of s . For this and all larger values of p there is a first-order phase transition, more easily seen on the blow up of the right panel. What appear as level crossings on this picture are actually avoided crossings with exponentially small gaps. This plot also exhibits in a striking way the metastability phenomenon, as an analytic continuation of one branch of the eigenvalue across the avoided crossing; one also sees the spinodal for the branches coming for large values of s , in particular the metastable continuation of the large s ground-state disappears around $s \approx 0.4$. Thanks to the permutation symmetry the diagonalization of this Hamiltonian, inside one symmetry sector, amounts to an unidimensional quantum problem, the magnetization playing the role of the coordinate of a particle. Moreover in this interpretation the role of \hbar is played by the inverse size of the system, $1/N$, hence in the thermodynamic limit all properties of the system, including the exponentially small gaps, can be computed with semi-classical methods, either instantonic [154, 155] or WKB-like [152].

One can then relate in an intuitive way these features of the spectrum to the annealing dynamics of a mean-field model exhibiting a first-order transition (see the sketch on the left panel of Fig. 3.3). Suppose first than the annealing time \mathcal{T} is very large, but sub-exponential in N . Then during the annealing all the exponentially small gaps encountered after the first-order transition cannot be resolved, and the dynamics follows the metastable continuation of the initial groundstate. Once the spinodal is reached there are no more

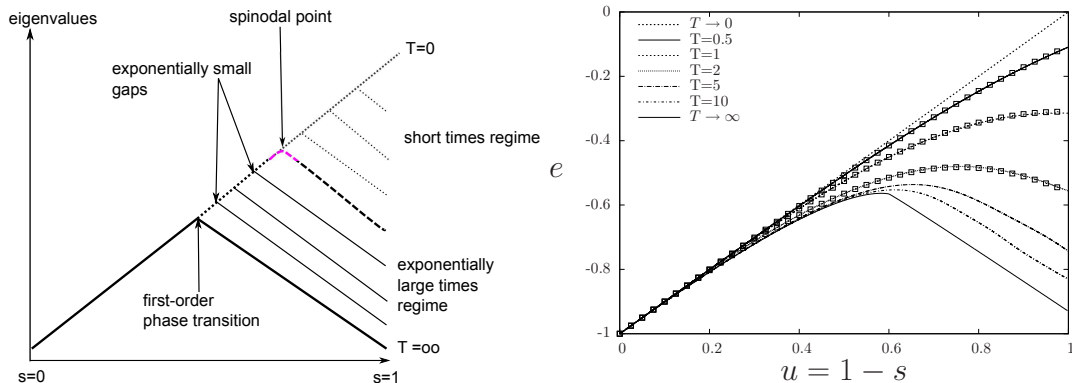


Figure 3.3: Left panel: a schematic view of annealing through a first-order phase transition. Right panel: explicit analytical and numerical results for the model (3.9), in the regime where \mathcal{T} (denoted T in this plot) is finite in the thermodynamic limit. Lines are the results of the numerical integration of the Hamilton equations resulting from the semi-classical equations in the thermodynamic limit, symbols correspond to the integration of Schrödinger equation at finite N .

avoided crossings, and as \mathcal{T} is large from this point the dynamics is adiabatic, the final energy reached at the end of the annealing is the one of the level connected to the spinodal point, let us call \hat{e}_{fin} the corresponding energy density. This trajectory actually separates two distinct regimes: for times $\mathcal{T} = e^{N\tau}$, with $\tau > 0$ but smaller than the exponential rate for adiabaticity, the metastable continuation of the initial groundstate is only followed up to a turning point before the spinodal, when a gap large enough to be resolved on the time-scale $e^{N\tau}$ is encountered. After this point the dynamics is again adiabatic, and leads to the energy density $e_{\text{fin}}(\tau) < \hat{e}_{\text{fin}}$ continuously connected to the turning point. This energy density decreases when τ increases, and reaches the groundstate one when τ crosses the exponential rate of the adiabatic time. The second regime corresponds on the contrary to times \mathcal{T} finite in the thermodynamic limit, and leads to energy densities $\tilde{e}_{\text{fin}}(\mathcal{T}) > \hat{e}_{\text{fin}}$, as illustrated on the right panel of Fig. 3.3.

In the simple model defined in Eq. (3.9) one can perform explicit calculations for the finally reached energy densities in these two regimes, $e_{\text{fin}}(\tau)$ and $\tilde{e}_{\text{fin}}(\mathcal{T})$, in the first one thanks to the WKB computation of the exponentially small gaps encountered along a metastable branch, and in the second one via a semi-classical reduction of the quantum dynamics to a Hamiltonian one. Maybe the most important outcome of these computations are the scaling laws found in the neighborhood of the separating trajectory, i.e. for “exponential but small” times ($\tau \rightarrow 0$) or “finite but large” times ($\mathcal{T} \rightarrow \infty$ after $N \rightarrow \infty$):

$$\hat{e}_{\text{fin}} - e_{\text{fin}}(\tau) \propto \tau^{\frac{4}{5}}, \quad \tilde{e}_{\text{fin}}(\mathcal{T}) - \hat{e}_{\text{fin}} \propto \mathcal{T}^{-\frac{4}{5}}. \quad (3.10)$$

Indeed one can conjecture that this exponent $4/5$ is universal for all mean-field models with a first-order transition, as this scaling behavior is controlled by the neighborhood of the spinodal point. A first step to check this idea would be to compute \hat{e}_{fin} in a less trivial model.

Chapter 4

Perspectives

I have tried in this manuscript to summarize the main results recently obtained in this interdisciplinary field at the interface between statistical mechanics, theoretical computer science and discrete mathematics. Let me now give some perspectives on possible directions for future research.

For what concerns the applications to computer science, two issues could be considered. The first one concerns the type of instances of CSPs studied up to now. Random ensembles have nice properties, as they are defined in terms of a single parameter α , and possess a natural thermodynamic limit in which sharp thresholds appear, with a very rich phenomenology as we have seen. However the satisfiability instances that are relevant for industrial applications are very different from random ones. They do exhibit some local structure, either because they derive from problems embedded in a finite-dimensional space, or because they are produced by encoding a problem from another language, this encoding repeating small “gadgets” that translate the basic unit of the more elaborated language. It would thus be very important to adapt, if possible, the statistical mechanics tools to these industrial formulas, without sacrificing the notions of typicality and thermodynamic limit that underlies the statistical mechanics approach (even if single sample algorithms as Survey Propagation [25, 63] deviate somehow from this paradigm). An intermediate step in this project would consist in devising a “random ensemble of structured formulas”, i.e. a probability law over instances biased to exhibit some of the features held relevant for industrial instances. The second issue would be the study of more general problems than CSPs, higher in the “polynomial hierarchy” of computational complexity theory, in particular Quantified Boolean Formulas. In these problems the existential quantifier “is there a configuration of the variables \underline{x} that satisfy all constraints” is complemented by an universal one, and leads to questions of the type “for all assignments of the variables \underline{x}_U , is there a configuration of the variables \underline{x}_E that satisfy all constraints”, with a partition of the variables between Universal and Existential ones. Statistical mechanics tools have recently been applied to similar problems in [156, 157].

On the physics side, besides technical issues as a better understanding of the Hilbert space interpretation of the replica symmetry breaking in quantum models, and a further development of the cavity method for stochastic dynamics, there is probably hope for new results on (many-body) localization problems [158] in the mean-field framework. This issue of transport properties of (interacting or not) quantum particles in a disordered environment is a long-standing subject of research in condensed matter [159] and in mathematical physics [160], with some controversial views on the effect of interactions on the localization phenomenon. Models of interacting quantum particles defined on random graphs should

enlighten these issues and provide a mean-field picture of the many-body localization; however the combination of disorder, quantum fluctuations, finite-connectivity and the necessity of exploring low-temperature regimes make the resolution of such models, even if they are mean-field, a challenging task, and up to now only approximate versions of the quantum cavity method could be used [139, 140, 141] for such problems. Note also that the mean-field non-interacting case [97], even if studied for a very long time, still exhibits not completely understood features. For instance it was rigorously proven very recently [161] that, at variance with all intuitive expectations, the mobility edge in the (energy-disorder) phase diagram does not extend down to zero disorder, i.e. that the spectrum of the Anderson model on the Bethe lattice remains entirely absolutely continuous in presence of a small but finite disorder in the local potentials. This phenomenon could not be detected via numerical methods, see for instance the (incorrect) phase diagram in [162], because of the Lifshitz tail phenomenon that brings the density of states to very small values close to the band edge [163].

In the context of quantum adiabatic computation it would be desirable to obtain a quantitative understanding of the energy reached by real-time Schrödinger and Quantum Monte Carlo annealing of non-trivial mean-field models, as attempted for classical simulated annealing in [77, 78], in particular in view of the approximability issues in computational complexity discussed at the end of Sec. 3.3.

Finally a large part of the results presented in this manuscript await a rigorous proof. From a physicist point of view the most satisfying scheme of proof would be the closer possible to the heuristic picture proposed by the cavity method, and would thus be based on the local convergence of random graphs to random trees supplemented by a precise control of the correlation decay. Some models have been treated rigorously along these lines, in particular the ferromagnetic Ising model [105] and the matchings problem [106], both being replica-symmetric. A local convergence result at the replica-symmetry breaking level would be a breakthrough for this line of research; more modest objectives would be generic proofs valid for a larger class of RS models, and generic validation of the exactness of the RS prediction for the free-energy in a dynamic 1RSB regime.

Acknowledgments

I would like first to express my gratitude towards David Sherrington, Leticia Cugliandolo and Rémi Monasson who guided my first steps in research before and during my Ph.D.

I also warmly thank all the other colleagues with whom I had the pleasure to co-author some works, Fabrizio Altarelli, Giulio Biroli, Simona Cocco, Silvio Franz, Enzo Marinari, Valery Van Kerrebroeck, Martin Weigt, and in particular for the main topics presented in this manuscript Victor Bapst, Laura Foini, Thomas Jörg, Florent Krzakala, Andrea Montanari, Federico Ricci-Tersenghi, Alberto Rosso, Marco Tarzia, Francesco Zamponi and Lenka Zdeborova. A special greeting goes to Francesco Zamponi who has made our common office a lively and stimulating place to work in.

I am extremely grateful to the members of the committee who accepted to fulfill this demanding task, and in particular the referees Amin Coja-Oghlan, Satya Majumdar and Giorgio Parisi.

Bibliography

- [1] J. Mydosh. *Spin glasses: an experimental introduction*. Taylor and Francis, 1993.
- [2] K. Fischer and J. Hertz. *Spin Glasses*. Cambridge University Press, 1991.
- [3] L. Berthier, G. Biroli, J. Bouchaud, L. Cipelletti, and W. van Saarloos. *Dynamical heterogeneities in glasses, colloids, and granular media*. OUP Oxford, 2011.
- [4] D. S. Fisher and D. A. Huse. Equilibrium behavior of the spin-glass ordered phase. *Phys. Rev. B*, **38**(1), 386–411 (1988).
- [5] C. Newman and D. Stein. Ground-state structure in a highly disordered spin-glass model. *J. Stat. Phys.*, **82**(3), 1113–1132 (1996).
- [6] E. Marinari, G. Parisi, F. Ricci-Tersenghi, J. Ruiz-Lorenzo, and F. Zuliani. Replica Symmetry Breaking in Short-Range Spin Glasses: Theoretical Foundations and Numerical Evidences. *J. Stat. Phys.*, **98**(5), 973–1074 (2000).
- [7] M. Baity-Jesi, R. Banos, A. Cruz, L. Fernandez, J. Gil-Narvion, A. Gordillo-Guerrero, M. Guidetti, D. Iniguez, A. Maiorano, F. Mantovani, E. Marinari, V. Martin-Mayor, J. Monforte-Garcia, A. Munoz Sudupe, D. Navarro, G. Parisi, M. Pivanti, S. Perez-Gaviro, F. Ricci-Tersenghi, J. Ruiz-Lorenzo, S. Schifano, B. Seoane, A. Tarancon, P. Tellez, R. Tripiccion, and D. Yllanes. Reconfigurable computing for Monte Carlo simulations: Results and prospects of the Janus project. *The European Physical Journal Special Topics*, **210**, 33 (2012).
- [8] B. Yucesoy, H. G. Katzgraber, and J. Machta. Evidence of Non-Mean-Field-Like Low-Temperature Behavior in the Edwards-Anderson Spin-Glass Model. *Phys. Rev. Lett.*, **109**, 177204 (2012).
- [9] D. Sherrington and S. Kirkpatrick. Solvable Model of a Spin-Glass. *Phys. Rev. Lett.*, **35**(26), 1792–1796 (1975).
- [10] S. F. Edwards and P. W. Anderson. Theory of Spin-Glasses. *J. Phys. F*, **5**, 965–974 (1975).
- [11] G. Parisi. A sequence of approximated solutions to the S-K model for spin glasses. *J. of Phys. A*, **13**(4), L115 (1980).
- [12] M. Mézard, G. Parisi, and M. Virasoro. *Spin glass theory and beyond*. World Scientific, Singapore, 1987.
- [13] M. Talagrand. The Parisi formula. *Annals of Mathematics*, **163**, 221 (2006).

- [14] F. Guerra and F. L. Toninelli. The Thermodynamic Limit in Mean Field Spin Glass Models. *Comm. Math. Phys.*, **230**, 71–79 (2002).
- [15] L. Viana and A. Bray. Phase diagrams for dilute spin glasses. *J. Phys. C*, **18**(15), 3037–3051 (1985).
- [16] C. Papadimitriou and K. Steiglitz. *Combinatorial Optimization: Algorithms and Complexity*. Dover, New York, 1998.
- [17] M. R. Garey and D. S. Johnson. *Computers and Intractability, A Guide to the Theory of NP-Completeness*. W.H. Freeman and Company, New York, 1979.
- [18] C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
- [19] A. Goldberg. Average case complexity of the satisfiability problem. In *Proc. 4th Workshop on Automated Deduction*, pages 1–6, Austin, TX, USA, 1979.
- [20] J. Franco and M. Paull. Probabilistic analysis of the Davis Putnam procedure for solving the satisfiability problem. *Discrete Applied Mathematics*, **5**(1), 77 – 87 (1983).
- [21] P. Cheeseman, B. Kanefsky, and W. M. Taylor. Where the really hard problems are. In *Proc. 12th IJCAI*, pages 331–337, San Mateo, CA, USA, 1991. Morgan Kaufmann.
- [22] D. G. Mitchell, B. Selman, and H. J. Levesque. Hard and Easy Distributions for SAT Problems. In *Proc. 10th AAAI*, pages 459–465, Menlo Park, California, 1992. AAAI Press.
- [23] M. Mézard and A. Montanari. *Information, Physics and Computation*. Oxford University Press, 2009.
- [24] G. Biroli, R. Monasson, and M. Weigt. A variational description of the ground state structure in random satisfiability problems. *Eur. Phys. J. B*, **14**, 551 (2000).
- [25] M. Mézard, G. Parisi, and R. Zecchina. Analytic and Algorithmic Solution of Random Satisfiability Problems. *Science*, **297**, 812–815 (2002).
- [26] 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. National Academy of Sciences*, **104**(25), 10318–10323 (2007).
- [27] H. Daudé, T. Mora, M. Mézard, and R. Zecchina. Pairs of SAT Assignments and Clustering in Random Boolean Formulae. *Th. Comp. Science*, **393**, 260–279 (2008).
- [28] D. Achlioptas and F. Ricci-Tersenghi. On the Solution-Space Geometry of Random Constraint Satisfaction Problems. *Proc. of the 38th annual ACM symposium on Theory of computing*, (2006).
- [29] D. Achlioptas and A. Coja-Oghlan. Algorithmic barriers from phase transitions. In *Foundations of Computer Science, 2008. FOCS'08. IEEE 49th Annual IEEE Symposium on*, pages 793–802. IEEE, 2008.
- [30] M. Molloy. The freezing threshold for k-colourings of a random graph. In *Proceedings of the 44th symposium on Theory of Computing*, page 921. ACM, 2012.
- [31] A. Montanari and G. Semerjian. Rigorous inequalities between length and time scales in glassy systems. *J. Stat. Phys.*, **125**(1), 23–54 (2006).

- [32] P. Erdős and A. Rényi. On the evolution of random graphs. *Magyar Tud. Akad. Mat. Kutató Int. Közl*, **5**, 17–61 (1960).
- [33] S. Janson, T. Luczak, and A. Rucinski. *Random graphs*. John Wiley and Sons, New York, 2000.
- [34] M. P. A. Fisher, P. B. Weichman, G. Grinstein, and D. S. Fisher. Boson localization and the superfluid-insulator transition. *Phys. Rev. B*, **40**(1), 546 (1989).
- [35] I. Bloch, J. Dalibard, and W. Zwerger. Many-body physics with ultracold gases. *Rev. Mod. Phys.*, **80**, 885 (2008).
- [36] G. Biroli, C. Chamon, and F. Zamponi. Theory of the superglass phase. *Phys. Rev. B*, **78**(22), 224306 (2008).
- [37] M.A. Nielsen and I.L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge, 2000.
- [38] P. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *Foundations of Computer Science, 1994 Proceedings., 35th Annual Symposium*, page 124, 1994.
- [39] B. Apolloni, C. Carvalho, and D. de Falco. Quantum stochastic optimization. *Stoc. Proc. Appl.*, **33**, 233 (1989).
- [40] A. Finnila, M. Gomez, C. Sebenik, C. Stenson, and J. Doll. Quantum annealing: A new method for minimizing multidimensional functions. *Chem. Phys. Lett.*, **219**, 343 (1994).
- [41] T. Kadowaki and H. Nishimori. Quantum annealing in the transverse Ising model. *Phys. Rev. E*, **58**(5), 5355 (1998).
- [42] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda. A Quantum Adiabatic Evolution Algorithm Applied to Random Instances of an NP-complete Problem. *Science*, **292**(5516), 472–475 (2001).
- [43] 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. *Physics Reports*, (0), – (2012).
- [44] T. Richardson and R. Urbanke. *Modern Coding Theory*. Cambridge University Press, Cambridge, 2008.
- [45] E. Friedgut. Sharp Thresholds of Graph Properties, and the k-Sat Problem. *J. of the American Mathematical Society*, **12**, 1017 (1999).
- [46] J. Franco. Results related to threshold phenomena research in satisfiability: lower bounds. *Theor. Comput. Sci.*, **265**, 147 (2001).
- [47] D. Achlioptas. Lower bounds for random 3-SAT via differential equations. *Theor. Comput. Sci.*, **265**(1-2), 159–185 (2001).
- [48] O. Dubois. Upper bounds on the satisfiability threshold. *Theor. Comput. Sci.*, **265**, 187 (2001).

- [49] D. Achlioptas and Y. Peres. The threshold for random k -SAT is $2^k \log 2 - O(k)$. *J. American Math. Soc.*, **17**, 947 (2004).
- [50] M. Mézard, F. Ricci-Tersenghi, and R. Zecchina. Two solutions to diluted p -spin models and XORSAT problems. *J. Stat. Phys.*, **111**(3-4), 505–533 (2003).
- [51] 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**(4), 047205 (2003).
- [52] O. Dubois and J. Mandler. The 3-XORSAT Threshold. In *Proceedings of the 43rd Symposium on Foundations of Computer Science, FOCS '02*, pages 769–778, 2002.
- [53] F. Kschischang, B. Frey, and H. Loeliger. Factor graphs and the sum-product algorithm. *IEEE Transactions on Information Theory*, **47**(2), 498 (2001).
- [54] R. Monasson and R. Zecchina. Entropy of the K -Satisfiability Problem. *Phys. Rev. Lett.*, **76**, 3881–3885 (1996).
- [55] M. Ibrahimi, Y. Kanoria, M. Kranning, and A. Montanari. The set of solutions of random XORSAT formulae. In *Proc. of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '12*, pages 760–779. SIAM, 2012.
- [56] D. Achlioptas and M. Molloy. The solution space geometry of random linear equations. [arXiv:1107.5550](https://arxiv.org/abs/1107.5550), (2011).
- [57] A. Montanari and G. Semerjian. On the dynamics of the glass transition on Bethe lattices. *J. Stat. Phys.*, **124**, 103 (2006).
- [58] T. Mora and M. Mézard. Geometrical organization of solutions to random linear Boolean equations. *J. Stat. Mech.*, page P10007 (2006).
- [59] M. Mézard, M. Palassini, and O. Rivoire. Landscape of Solutions in Constraint Satisfaction Problems. *Phys. Rev. Lett.*, **95**(20), 200202 (2005).
- [60] A. Montanari, F. Ricci-Tersenghi, and G. Semerjian. Clusters of solutions and replica symmetry breaking in random K -satisfiability. *J. Stat. Mech.*, **2008**(04), P04004 (2008).
- [61] L. Zdeborová and F. Krzakala. Phase transitions in the coloring of random graphs. *Phys. Rev. E*, **76**(3), 031131 (2007).
- [62] B. Derrida. Random-energy model: An exactly solvable model of disordered systems. *Phys. Rev. B*, **24**(5), 2613–2626 (1981).
- [63] A. Braunstein, M. Mézard, and R. Zecchina. Survey propagation: an algorithm for satisfiability. *Random Struct. Algorithms*, **27**(2), 201–226 (2005).
- [64] M. Mézard and R. Zecchina. Random K -satisfiability problem: From an analytic solution to an efficient algorithm. *Phys. Rev. E*, **66**(5), 056126 (2002).
- [65] G. Semerjian. On the freezing of variables in random constraint satisfaction problems. *J. Stat. Phys.*, **130**, 251 (2008).
- [66] W. Gotze and L. Sjogren. Relaxation processes in supercooled liquids. *Reports on Progress in Physics*, **55**(3), 241–376 (1992).

- [67] A. Crisanti and H.-J. Sommers. The spherical p -spin interaction spin glass model: the statics. *Zeitschrift für Physik B Condensed Matter*, **87**, 341–354 (1992).
- [68] A. Crisanti, H. Horner, and H.-J. Sommers. The spherical p -spin interaction spin-glass model: the dynamics. *Zeitschrift für Physik B Condensed Matter*, **92**, 257–271 (1993).
- [69] L. Cugliandolo and J. Kurchan. Analytical solution of the off-equilibrium dynamics of a long-range spin-glass model. *Phys. Rev. Lett.*, **71**(1), 173 (1993).
- [70] T. R. Kirkpatrick and P. G. Wolynes. Stable and metastable states in mean-field Potts and structural glasses. *Phys. Rev. B*, **36**, 8552 (1987).
- [71] T. R. Kirkpatrick and D. Thirumalai. Mean-field soft-spin Potts glass model: Statics and dynamics. *Phys. Rev. B*, **37**, 5342 (1988).
- [72] M. Mézard and A. Montanari. Reconstruction on trees and spin glass transition. *J. Stat. Phys.*, **124**(6), 1317–1350 (2006).
- [73] J.-P. Bouchaud and G. Biroli. On the Adam-Gibbs-Kirkpatrick-Thirumalai-Wolynes scenario for the viscosity increase in glasses. *The Journal of Chemical Physics*, **121**(15), 7347–7354 (2004).
- [74] G. Biroli, J.-P. Bouchaud, A. Cavagna, T. S. Grigera, and P. Verrocchio. Thermodynamic signature of growing amorphous order in glass-forming liquids. *Nature Physics*, **4**, 771 (2008).
- [75] L. Berthier and W. Kob. Static point-to-set correlations in glass-forming liquids. *Phys. Rev. E*, **85**, 011102 (2012).
- [76] F. Martinelli. Lectures on Glauber dynamics for discrete spin models. *Lectures on probability theory and statistics*, pages 93–191 (2004).
- [77] F. Krzakala and L. Zdeborová. Following Gibbs states adiabatically. The energy landscape of mean-field glassy systems. *Europhysics Lett.*, **90**(6), 66002 (2010).
- [78] L. Zdeborová and F. Krzakala. Generalization of the cavity method for adiabatic evolution of Gibbs states. *Phys. Rev. B*, **81**(22), 224205 (2010).
- [79] F. Krzakala and L. Zdeborová. Phase transitions and computational difficulty in random constraint satisfaction problems. *Journal of Physics: Conference Series*, **95**(1), 012012 (2008).
- [80] B. Selman, H. A. Kautz, and B. Cohen. Noise strategies for improving local search. In *Proc. of the Twelfth National Conference on Artificial Intelligence (AAAI'94)*, pages 337–343, Seattle, 1994.
- [81] D. McAllester, B. Selman, and H. Kautz. Evidence for Invariants in Local Search. In *Proc. of the Fourteenth National Conference on Artificial Intelligence (AAAI'97)*, pages 321–326, Providence, Rhode Island, 1997.
- [82] J. Ardelius and E. Aurell. Behavior of heuristics on large and hard satisfiability problems. *Phys. Rev. E*, **74**(3), 037702 (2006).
- [83] M. Alava, J. Ardelius, E. Aurell, P. Kaski, S. Krishnamurthy, P. Orponen, and S. Seitz. Circumspect descent prevails in solving random constraint satisfaction problems. *Proceedings of the National Academy of Sciences*, **105**, 15253 (2008).

- [84] G. Semerjian and R. Monasson. Relaxation and metastability in a local search procedure for the random satisfiability problem. *Phys. Rev. E*, **67**(6), 066103 (2003).
- [85] W. Barthel, A. K. Hartmann, and M. Weigt. Solving satisfiability problems by fluctuations: The dynamics of stochastic local search algorithms. *Phys. Rev. E*, **67**(6), 066104 (2003).
- [86] A. Coja-Oghlan and A. Frieze. Analyzing Walksat on random formulas. *Proc. 9th ANALCO*, pages 48–55 (2012).
- [87] A. Montanari, F. Ricci-Tersenghi, and G. Semerjian. Solving Constraint Satisfaction Problems through Belief Propagation-guided decimation. (2007). [arXiv:0709.1667](https://arxiv.org/abs/0709.1667), Proceedings of the 45th Allerton Conference.
- [88] F. Ricci-Tersenghi and G. Semerjian. On the cavity method for decimated random constraint satisfaction problems and the analysis of belief propagation guided decimation algorithms. *Journal of Statistical Mechanics: Theory and Experiment*, page P09001 (2009).
- [89] A. Coja-Oghlan. On belief propagation guided decimation for random k -SAT. *Proc. 22nd SODA*, page 957 (2011).
- [90] A. Coja-Oghlan and A. Y. Pachon-Pinzon. The decimation process in random k -SAT. *SIAM Journal on Discrete Mathematics*, **26**, 1471–1509 (2012).
- [91] 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. *Journal of Physics: Conference Series*, **95**(1), 012013 (2008).
- [92] Silvio Franz and Giorgio Parisi. Recipes for Metastable States in Spin Glasses. *J. Phys. I France*, **5**(11), 1401–1415 (1995).
- [93] C. Cammarota and G. Biroli. Ideal glass transitions by random pinning. *Proceedings of the National Academy of Sciences*, **109**, 8850 (2012).
- [94] C. Cammarota and G. Biroli. Random Pinning Glass Transition: Hallmarks, Mean-Field Theory and Renormalization Group Analysis. [arXiv:1210.8399](https://arxiv.org/abs/1210.8399), (2012).
- [95] R. Monasson. Optimization problems and replica symmetry breaking in finite connectivity spin glasses. *J. Phys. A*, **31**(2), 513–529 (1998).
- [96] M. Mézard and G. Parisi. The Bethe lattice spin glass revisited. *Eur. Phys. J. B*, **20**, 217 (2001).
- [97] R. Abou-Chacra, D. Thouless, and P. Anderson. A selfconsistent theory of localization. *J. Phys. C*, **6**, 1734 (1973).
- [98] R. Monasson. Structural Glass Transition and the Entropy of the Metastable States. *Phys. Rev. Lett.*, **75**, 2847–2850 (1995).
- [99] O. Rivoire. Properties of Atypical Graphs from Negative Complexities. *Journal of Statistical Physics*, **117**, 453–476 (2004).
- [100] L. Zdeborová and M. Mézard. The number of matchings in random graphs. *Journal of Statistical Mechanics: Theory and Experiment*, **2006**(05), P05003 (2006).

- [101] E. Marinari and G. Semerjian. On the number of circuits in random graphs. *Journal of Statistical Mechanics: Theory and Experiment*, **2006**(06), P06019 (2006).
- [102] M. Pretti and M. Weigt. Sudden emergence of q -regular subgraphs in random graphs. *Europhys. Lett.*, **75**(1), 8 (2006).
- [103] S. Franz and M. Leone. Replica bounds for optimization problems and diluted spin systems. *J. Stat. Phys.*, **111**(3-4), 535–564 (2003).
- [104] D. Panchenko and M. Talagrand. Bounds for diluted mean-fields spin glass models. *Probab. Theory Relat. Fields*, **130**(3), 319–336 (2004).
- [105] A. Dembo and A. Montanari. Ising models on locally tree-like graphs. *Ann. Appl. Probab.*, **20**, 565–592 (2010).
- [106] C. Bordenave, M. Lelarge, and J. Salez. Matchings on infinite graphs. *Probability Theory and Related Fields*, pages 1–26 (2012).
- [107] N.D. Mermin. *Quantum Computer Science*. Cambridge University Press, Cambridge, 2007.
- [108] D. Bruß and G. Leuchs, editors. *Lectures on Quantum Information*. Wiley-VCH, Weinheim, 2007.
- [109] A.Y. Kitaev, A.H. Shen, and M.N. Vyalyi. *Classical and quantum computation*. AMS, Providence, 2002.
- [110] C.R. Laumann, R. Moessner, A. Scardicchio, and S.L. Sondhi. Statistical mechanics of classical and quantum computational complexity. *arXiv:1009.1635*, (2010).
- [111] D. P. DiVincenzo. Two-bit gates are universal for quantum computation. *Phys. Rev. A*, **51**, 1015–1022 (1995).
- [112] D. Deutsch, A. Barenco, and A. Ekert. Universality in Quantum Computation. *Proceedings: Math. and Phys. Sciences*, **449**(1937), 669 (1995).
- [113] S. Lloyd. Almost Any Quantum Logic Gate is Universal. *Phys. Rev. Lett.*, **75**, 346–349 (1995).
- [114] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter. Elementary gates for quantum computation. *Phys. Rev. A*, **52**, 3457–3467 (1995).
- [115] S. Kirkpatrick, C. D. Gelatt Jr., and M. P. Vecchi. Optimization by Simulated Annealing. *Science*, **220**, 671–680 (1983).
- [116] A. Messiah. *Quantum mechanics vol. 2*. North-Holland, Amsterdam, 1962.
- [117] Y. Y. Goldschmidt. Solvable model of the quantum spin glass in a transverse field. *Phys. Rev. B*, **41**(7), 4858–4861 (1990).
- [118] T. Nieuwenhuizen and F. Ritort. Quantum phase transition in spin glasses with multi-spin interactions. *Physica A*, **250**(1), 8–45 (1998).
- [119] G. Biroli and L. F. Cugliandolo. Quantum Thouless-Anderson-Palmer equations for glassy systems. *Phys. Rev. B*, **64**(1), 014206 (2001).

- [120] L. F. Cugliandolo, D. R. Grempel, and C. A. da Silva Santos. Imaginary-time replica formalism study of a quantum spherical p -spin-glass model. *Phys. Rev. B*, **64**(1), 014403 (2001).
- [121] T. Jörg, F. Krzakala, J. Kurchan, and A. C. Maggs. Simple Glass Models and Their Quantum Annealing. *Phys. Rev. Lett.*, **101**(14), 147204 (2008).
- [122] J. Ginibre. Reduced density matrices of the anisotropic Heisenberg model. *Commun. Math. Phys.*, **10**, 140–154 (1968).
- [123] G. Gallavotti, S. Miracle-Sole, and D. Robinson. Analyticity properties of the anisotropic Heisenberg model. *Commun. Math. Phys.*, **10**, 311–324 (1968).
- [124] E. Farhi and S. Gutmann. The functional integral constructed directly from the Hamiltonian. *Annals of Physics*, **213**, 182–203 (1992).
- [125] M. Aizenman and B. Nachtergaele. Geometric aspects of quantum spin states. *Commun. Math. Phys.*, **164**, 17 (1994).
- [126] L. Chayes, N. Crawford, D. Ioffe, and A. Levit. The Phase Diagram of the Quantum Curie-Weiss Model. *J. Stat. Phys.*, **133**, 131–149 (2008).
- [127] D. Ioffe. Stochastic Geometry of Classical and Quantum Ising Models. In *Methods of Contemporary Mathematical Statistical Physics*, volume 1970 of *Lecture Notes in Mathematics*, page 87. Springer Berlin / Heidelberg, 2009.
- [128] F. Martinelli and M. Wouts. Glauber Dynamics for the Quantum Ising Model in a Transverse Field on a Regular Tree. *J. Stat. Phys.*, **146**, 1059–1088 (2012).
- [129] C. Laumann, A. Scardicchio, and S. L. Sondhi. Cavity method for quantum spin glasses on the Bethe lattice. *Phys. Rev. B*, **78**(13), 134424 (2008).
- [130] F. Krzakala, A. Rosso, G. Semerjian, and F. Zamponi. Path-integral representation for quantum spin models: Application to the quantum cavity method and Monte Carlo simulations. *Phys. Rev. B*, **78**(13), 134428 (2008).
- [131] G. Semerjian, M. Tarzia, and F. Zamponi. Exact solution of the Bose-Hubbard model on the Bethe lattice. *Phys. Rev. B*, **80**(1), 014524 (2009).
- [132] 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**(20), 207206 (2010).
- [133] L. Foini, G. Semerjian, and F. Zamponi. Quantum Biroli-Mézard model: Glass transition and superfluidity in a quantum lattice glass model. *Phys. Rev. B*, **83**, 094513 (2011).
- [134] M. Hastings. Quantum belief propagation: An algorithm for thermal quantum systems. *Phys. Rev. B*, **76**(20), 201102 (2007).
- [135] M. S. Leifer and D. Poulin. Quantum graphical models and belief propagation. *Annals of Physics*, **323**(8), 1899 (2008).
- [136] D. Poulin and E. Bilgin. Belief propagation algorithm for computing correlation functions in finite-temperature quantum many-body systems on loopy graphs. *Phys. Rev. A*, **77**(5), 052318 (2008).

- [137] E. Bilgin and D. Poulin. Coarse-grained belief propagation for simulation of interacting quantum systems at all temperatures. *Phys. Rev. B*, **81**(5), 054106 (2010).
- [138] D. Poulin and M. Hastings. Markov entropy decomposition: a variational dual for quantum belief propagation. *Phys. Rev. Lett.*, **106**(8), 80403 (2011).
- [139] L. B. Ioffe and M. Mézard. Disorder-Driven Quantum Phase Transitions in Superconductors and Magnets. *Phys. Rev. Lett.*, **105**(3), 037001 (2010).
- [140] M. V. Feigel'man, L. B. Ioffe, and M. Mézard. Superconductor-insulator transition and energy localization. *Phys. Rev. B*, **82**(18), 184534 (2010).
- [141] O. Dimitrova and M. Mézard. The cavity method for quantum disordered systems: from transverse random field ferromagnets to directed polymers in random media. *J. Stat. Mech.*, **2011**, P01020 (2011).
- [142] V. Bapst, G. Semerjian, and F. Zamponi. The effect of quantum fluctuations on the coloring of random graphs. *in preparation*.
- [143] L. Foini, G. Semerjian, and F. Zamponi. Solvable Model of Quantum Random Optimization Problems. *Phys. Rev. Lett.*, **105**(16), 167204 (2010).
- [144] T. Mora and L. Zdeborová. Random Subcubes as a Toy Model for Constraint Satisfaction Problems. *J. Stat. Phys.*, **131**, 1121–1138 (2008).
- [145] E. Farhi, D. Gosset, I. Hen, A. Sandvik, P. Shor, A. Young, and F. Zamponi. The performance of the quantum adiabatic algorithm on random instances of two optimization problems on regular hypergraphs. *arXiv:1208.3757*, (2012).
- [146] D. Gosset, 2011. PhD Thesis, Case Studies in Quantum Adiabatic Optimization.
- [147] M. H. S. Amin and V. Choi. First-order quantum phase transition in adiabatic quantum computation. *Phys. Rev. A*, **80**(6), 062326 (2009).
- [148] B. Altshuler, H. Krovi, and J. Roland. Anderson localization makes adiabatic quantum optimization fail. *Proceedings of the National Academy of Sciences*, **107**(28), 12446 (2010).
- [149] E. Farhi, J. Goldstone, D. Gosset, S. Gutmann, H. B. Meyer, and P. W. Shor. Quantum adiabatic algorithms, small gaps, and different paths. *Quantum Information & Computation*, **11**(3&4), 181–214 (2011).
- [150] N. G. Dickson. Elimination of perturbative crossings in adiabatic quantum optimization. *New J. of Physics*, **13**(7), 073011 (2011).
- [151] J. Hastad. Some optimal inapproximability results. *J. of the ACM*, **48**, 798 (2001).
- [152] V. Bapst and G. Semerjian. On quantum mean-field models and their quantum annealing. *J. Stat. Mech.*, **2012**(06), P06007 (2012).
- [153] H. Lipkin, N. Meshkov, and A. Glick. Validity of many-body approximation methods for a solvable model: (I). Exact solutions and perturbation theory. *Nucl. Phys.*, **62**(2), 188 – 198 (1965).
- [154] T. Jörg, F. Krzakala, J. Kurchan, and A. C. Maggs. Quantum Annealing of Hard Problems. *Progress of Theoretical Physics Supplement*, **184**, 290–303 (2010).

- [155] T. Jörg, F. Krzakala, J. Kurchan, A.C. Maggs, and J. Pujos. Energy gaps in quantum first-order mean-field-like transitions: The problems that quantum annealing cannot solve. *EPL*, **89**(4), 40004 (2010).
- [156] F. Altarelli, A. Braunstein, A. Ramezanpour, and R. Zecchina. Stochastic Matching Problem. *Phys. Rev. Lett.*, **106**, 190601 (2011).
- [157] M. Castellana and L. Zdeborova. Adversarial satisfiability problem. *Journal of Statistical Mechanics: Theory and Experiment*, **2011**(03), P03023 (2011).
- [158] P. W. Anderson. Absence of Diffusion in Certain Random Lattices. *Phys. Rev.*, **109**(5), 1492 (1958).
- [159] E. Abrahams, editor. *50 years of Anderson localization*. World Scientific, 2010.
- [160] R. Carmona and J. Lacroix. *Spectral Theory of Random Schrödinger Operators*. Birkhäuser, 1990.
- [161] M. Aizenman and S. Warzel. Resonant delocalization for random Schrödinger operators on tree graphs. *arXiv:1104.0969*, (2011).
- [162] G. Biroli, G. Semerjian, and M. Tarzia. Anderson Model on Bethe Lattices: Density of States, Localization Properties and Isolated Eigenvalue. *Progress of Theoretical Physics Supplement*, **184**, 187–199 (2010).
- [163] V. Bapst and G. Semerjian. Lifshitz Tails on the Bethe Lattice: A Combinatorial Approach. *Journal of Statistical Physics*, **145**, 51–92 (2011).

Résumé

Ce mémoire présente mes activités de recherche dans le domaine de la mécanique statistique des systèmes désordonnés, en particulier sur les modèles de champ moyen à connectivité finie. Ces modèles présentent de nombreuses transitions de phase dans la limite thermodynamique, avec des applications tant pour la physique des verres que pour leurs liens avec des problèmes d'optimisation de l'informatique théorique. Leur comportement sous l'effet de fluctuations quantiques est aussi discuté, en lien avec des perspectives de calcul quantique.

Abstract

This manuscript presents my research activities in the field of statistical mechanics of disordered systems, in particular mean-field models with finite connectivity. These models exhibit several phase transitions in the thermodynamic limit, with applications both for the physics of glasses and for their links with optimization problems from theoretical computer science. Their behavior under the effect of quantum fluctuations is also discussed, within the perspective of quantum computation.