Quantum glasses, and quantum optimization by quantum annealing

This text was written for my CNRS evaluation in February 2011, and is based on a submitted (and rejected) ERC Starting Grant project.

This part of my research activity was devoted to the development of analytical and numerical methods to describe a wide class of *strongly interacting quantum disordered systems*, in which the interplay between quantum fluctuations and disorder/frustration plays a key role. This project was motivated by the potentially numerous applications of the techniques that we developed, that range from quantum computing to the physics of localization in condensed matter systems. The basic methodology is that of developing a mean field treatment of these systems that allows for a full analytical solution, at the same time keeping track of local fluctuations which are essential in the study of disordered phases and localization effects. This project was realized in close collaboration with F.Krzakala (ESPCI), A.Rosso (CNRS), G.Semerjian (ENS), and M.Tarzia (Univ. Paris 6). It was originally conceived as an ANR "Jeunes Chercheurs" grant request for the 2008 and 2009 calls, and was rejected in both cases. I then improved the project and submitted it to the ERC Starting Grant calls in 2010 and 2011; the project was classified as eligible for funding but not financed in 2010, and it was rejected in 2011. Despite failure in raising funds, we were able to carry on the project and reach most of the original goals (yes, theoretical physicists don't need money to work!). Part of the project was the subject of the PhD thesis of Laura Foini and of Giuseppe Carleo, that worked on these topics under my supervision.

In the following I will:
A) begin by a short description of these problems, explaining what are their common ingredients and why they can be treated by similar methods;
B) explain what are the challenges that one meets when studying these problems; then explain the methodology that we used to solve them, and the key physical ideas beyond it;
C) for each application, make a short list of concrete results that we obtained using this method.

A. Overview of the physical problems
The method we developed allowed to obtain interesting results in the following areas:
(i) *Adiabatic quantum computing:* The theoretical research on quantum computing is motivated by the exciting perspective of computers that take intrinsically advantage of the laws of quantum mechanics. Besides the great effort of research towards the physical realization of these devices, a lot of activity has been devoted to the development of "softwares", that is algorithms that could use the specific properties of quantum computers to achieve a faster velocity in performing computational tasks with respect to classical devices, for instance minimizing irregularly shaped cost functions. In this context the frustration comes from the large number of contradictory constraints that have to be simultaneously satisfied. Thanks to the analogy between these problems and spin-glasses, the Hamiltonian of the quantum computer is akin to a quantum spin glass, where the quantum character arises typically from the inclusion of a transverse magnetic field. One of the aims of this project was precisely to study quantitatively the phase diagrams of quantum spin glasses in a transverse field. Solving these models resulted in a better understanding of the performances of a class of quantum algorithms for solving difficult optimization problems.
(ii) *Superfluidity and superconductivity in disordered systems:* In this context there is a wide number of open problems where the interplay between quantum fluctuations and disorder is important: for instance in cases where localization drives the formation of a quantum glassy phase, such as the Bose glass, or for disordered solid phases, such as the recently proposed superglass phase. These exotic phases are observed in Helium 4 when disorder is introduced by absorbing it in porous media or by producing "dirty" crystals by fast quenches from the liquid phase. Moreover, they have very recently been observed in cold atoms assemblies subjected to disordered potentials. These problems are also relevant for the superconductor-insulator transition in disordered electron systems, where it has been proposed that (Bosonic) Cooper pairs form a Bose glass. A paradigmatic example of a quantum strongly-interacting particle system displaying these phenomena is the (Bosonic or Fermionic) Hubbard model, in which particles can hop between neighboring sites of a lattice and interact via short-range potentials.
(iii) *Stochastic dynamics of disordered systems:* It is well known that the master equation describing the stochastic dynamics of a classical system can be mapped onto a quantum Hamiltonian. Hence, similar methods can be used to tackle both problems. Dynamical problems involving a large number of interacting entities in presence of disorder are encountered in a wide class of applications: examples are the analysis of

stochastic algorithms in computer science and of gene regulatory networks in biology. It is difficult to exhaust here all the potential applications of the method is this context. At the beginning the project was more focused on applications to quantum problems; hence in the following I will discuss first the quantum case. The study of classical dynamics will be discussed at the end.

All the examples above belong to a wider class of problems where the interplay between *quantum fluctuations*, *strong interactions* and *disorder/frustration* plays a key role. Indeed, the models that have been used to investigate these problems are strikingly similar. Typically, the quantum Hamiltonian is the sum of a "classical" part that contains strong interactions and disorder, and a term inducing the quantum fluctuations (hopping for particles, or a transverse field for spins). The "classical" part of the Hamiltonian can typically be reduced to a classically frustrated system, such as a Random Field Ising Model or a spin glass. Therefore, one can expect that these problems can be tackled by properly translating to the quantum world our experience in the study of classically frustrated systems, as will be detailed in the next section.

B. Overview of our research methodology

The common way to attack these problem is to use a new method, the quantum version of the so-called "cavity method", that has very recently been developed, first by a Princeton group, then by ourselves. Our main aim in this project was that of computing quantitatively, within reasonable approximation, the phase diagram of a given quantum strongly interacting disordered Hamiltonian. This is extremely difficult, mainly because due to the strong interactions, perturbation theory (that is the standard way of tackling difficult quantum problems) breaks down. Note that we are not interested here in the critical regime close to a phase transition, hence sophisticated techniques such as the Renormalization Group are of little help in this case. Apart from special cases (mainly one dimensional) where one can obtain exact solutions, two main strategies to tackle these problems have become standard nowadays: on the analytical side, one often resorts to mean-field like theories; on the numerical side, one makes use of Quantum Monte Carlo simulation schemes. However, both methods, despite many successes, suffer from severe drawbacks when applied to disordered strongly interacting systems:

- Numerical simulations are difficult mostly because in strongly disordered systems there are ergodicity problems, and additionally one has to consider very large samples in order to properly take into account rare realizations of the disorder, that are often relevant in determining the physical behavior of the system in the thermodynamic limit. The behavior of small systems is sometimes misleading, as the scaling with system size can change dramatically for larger systems. Striking examples of this are the Bose glass phase and Griffiths phases in quantum disordered magnets. Hence, although numerical simulations can give very precious indications, analytical methods that are capable of taking directly the thermodynamic limit and the average over the disorder are very much needed.
- In the mean field approach, one assumes that the behavior of the system can be captured by a "representative" degree of freedom subject to a "mean external field" which is averaged over the whole system. The latter is determined self-consistently by assuming that the representative degree of freedom describes the environment of its neighbors. However, disordered systems are strongly heterogeneous, and each degree of freedom feels a very different local field coming from its local environment. The interplay of heterogeneity and quantum fluctuations is at the basis of the interesting phenomena observed in such systems. Another important ingredient that has to be taken into account by the theory is that frustration often induces the existence of a very large number of *metastable states* (i.e. local minima of a suitable free energy functional), each described by a different set of local fields.

Hence, solving these problems requires the development of a method capable to take into account the strong fluctuations of the local fields, while at the same time being analytically solvable in the thermodynamic limit, and being capable to perform the correct average over the disorder. In the classical case, this has been achieved by the so-called *cavity method* [Mézard, Parisi, EPJB 2001]**.** This method is an extension of older ideas (known as the "Bethe approximation" in condensed matter or "Belief Propagation" in information theory) and is designed to tackle exactly this physical situation, where local fields display strong fluctuations from site to site, and from one metastable state to the other. Its more elaborate versions allow to deal with glassy phases that arise in frustrated spin models and which manifest themselves by the proliferation of the number of metastable states. This method has been very successfully applied in the context of random combinatorial optimization, thanks to the analogy between these problems and finite connectivity mean-field spin-glasses.

In summary, the cavity method can be seen as a "refined" mean-field like approximation (the so-called *Bethe approximation*) of models defined, for instance, on a square or cubic lattice. Yet, it has several important advantages:

1) It takes into account local spatial fluctuations of the environment, and the existence of many different states; it allows to take the thermodynamic limit and in this limit it allows to compute probability distributions of local observables with respect to disorder;
2) It is exact for models defined on random lattices (or Bethe lattices); this means that there is a class of concrete models, defined by a local Hamiltonian, of which the method gives the exact solution (for instance allowing for a direct comparison with Monte-Carlo simulations);
3) It can be formulated via a variational principle: in other words there is a suitable free energy functional whose minimization yields the cavity equations; this allows a direct access to the free energy of the system, and makes easy to devise variational approximations to the true solution.
4) It allows to define a distance between two spins as the number of interactions on a shortest path between them, leading to a consistent definition of a correlation length.
More specifically, in the quantum case
5) It allows to study phenomena that are intrinsically related to the notion of distance, the most striking of them being Anderson localization [Abou-Chacra, Thouless, Anderson, J.Phys.C 1973], that can be studied on a Bethe lattice while it disappears in the standard mean field limit of infinite connectivity.

From a technical point of view the goal of our project was to develop a full generalization of the cavity method to quantum models, suitable to be applied to the diverse problems outlined above. This program was started recently and is based on a discrete imaginary time (Suzuki-Trotter) path integral formulation [Laumann, Scardicchio, Sondhi, PRB 2008]. We improved over this preliminary investigation in several ways, mainly by showing how to perform the continuum imaginary-time limit and by extending the method to Bosons, as will be detailed below.

It is important to stress at this point the relation between our formulation of the quantum cavity method and other attempts to go beyond the simplest mean-field theory for disordered quantum systems. Indeed, our formulation of the quantum cavity method is *exact* on sparse random graphs, and contains, as special limits, many different methods that have been very recently developed to investigate these problems. These all correspond to large connectivity limits, in particular:
1) The stochastic mean-field theory of [Bissbort, Hofstetter, EPL 2009] and the closely related method of [Ioffe, Mézard, PRL 2010] both correspond to the leading order in a large connectivity expansion of the cavity method, for bosons and spins respectively.
2) A certain class of extensions of the Fermionic DMFT [Dobrosavljevic, Kotliar, PRL 1997] correspond to the leading order in large connectivity of the cavity method for fermions.
3) The recently formulated Bosonic DMFT [Byczuk, Vollhardt, PRB 2008] and the original formulation of [Laumann, Scardicchio, Sondhi, PRB 2008] take into account the next-to-leading order in the same expansion for bosons and spins, respectively.
The fact that so many cavity-like approximations have been recently derived testifies the vitality of the method in the present moment.

C. Overview of the results

The cavity method is exact on random or Bethe lattices in the sense that it allows to reduce the solution of a problem involving a large number of interacting degrees of freedom to that of a single functional equation. However, the cavity equation is a complicated recurrence equation for the local effective action, and one has to find its fixed point. At variance with DMFT equations, in which only the leading Gaussian term is kept, in the cavity method one explicitly takes into account all many-points local correlations in time, therefore having access to the full local effective action. However, the equations are difficult to handle. We obtained the following technical results that constitute the basis of all our results:

1. We introduced a new strategy to solve the quantum cavity equations, that consists in constructing a sample of spin trajectories from the local effective action and finding statistically a fixed point by iterating the cavity equation on this sample. We showed how to perform the continuum imaginary-time limit and we applied the method to solve exactly the quantum Ising ferromagnet in a transverse field on a random regular lattice [Krzakala, Rosso, Semerjian, Zamponi, PRB 2008].
2. In a subsequent paper we showed how to generalize the method to Bosonic systems [Semerjian, Tarzia, Zamponi, PRB 2009].
3. More recently, we included in the quantum cavity method some developments of the classical case (known as replica symmetry breaking effects) that are needed to treat situations where many metastable states are present [Jorg, Krzakala, Semerjian, Zamponi, PRL 2010].

Therefore, we showed that our method works well for spin and Bosonic models where the matrix elements of the Hamiltonian in the Suzuki-Trotter representation are all positive and the effective action is then a

probability distribution. However, this method is limited because it cannot access directly the zero-temperature limit, and it does not work for cases (such as fermions or the real-time dynamics of bosons) where a sign problem arises. It would be desirable to develop alternative strategies that could address, possibly in an approximate way, these more complicated cases.

Having developed the necessary technical tools, we applied the method to some interesting physical problems obtaining the following results:

(i) *Failure of adiabatic quantum computing for some difficult problems* - Some of the quantum algorithms proposed up to now have been written with a specific task in mind, for instance factoring a large integer [Shor, SIAM J.Comput. 1997]; the quantum adiabatic algorithm (QAA) [Farhi et al., Science 2001; Kadowaki, Nishimori, PRE 1998] is, on the contrary, able to tackle a variety of optimization problems in a universal way. Its functioning is based on a slow interpolation between a simple Hamiltonian, such that the quantum computer can be easily initialized in its ground state, and the Hamiltonian (cost function) to be minimized. If the interpolation is slow enough the adiabatic theorem ensures that the system will remain at all times in the ground state of the interpolating Hamiltonian. At the end of the evolution, it will be in the ground state of the problem Hamiltonian which will be solved. However, physical intuition suggests that the time needed to ensure the adiabaticity condition will diverge in the thermodynamic limit whenever a quantum phase transition is encountered during the interpolation. Since the instantaneous Hamiltonian can typically be written as a spin glass model in a transverse field, the problem is turned into the problem of determining the phase diagram of a broad class of quantum spin glasses.

At first, we studied the simplest optimization problem, namely the random XORSAT problem in a transverse field. The structure of the solution space of this problem in the classical case has been rigorously established, and we used the cavity method to understand how this structure is modified when adding quantum fluctuations to the problem. We showed that this problem displays a first order quantum phase transition as a function of the quantum fluctuations intensity. This transition is accompanied by an exponentially small (in system size) gap in the thermodynamic limit, which makes the QAA ineffective for this problem [Jorg, Krzakala, Semerjian, Zamponi, PRL 2010].

An important ingredient that is missing in XORSAT is that classical combinatorial problems usually have an exponentially large (in the number of variables) degeneracy of their ground and excited states. In such a situation the relevant gap probably involves the excited states that are not continuously transformed to the degenerate classical ground states. Preliminary results pointing in this direction have been recently obtained [Altshuler et al., Farhi et al., preprint 2009]. We formulated and investigated a toy model that is exactly solvable, yet which presents such a degeneracy of levels [Foini, Semerjian, Zamponi, PRL 2010]. We showed that entropy causes unexpected effects: in some situations, quantum fluctuations might promote the existence of a glass phase!

(ii) *Superfluidity and superconductivity in disordered systems* - We focused on Bosonic systems, since the method does not work for Fermions. Still, the investigation of interacting Bosons is very timely since cold-atoms experiments recently provided access to exotic phases of disordered Bosons [Roati et al., Billy et al., Nature 2008]. We confirmed, by means of the cavity method, an earlier suggestion [Boninsegni, Prokof'ev, Svistunov, PRL 2006] that geometrical frustration alone can induce a "superglass" phase of bosons in which the system supports at the same time glassy ordering and superfluidity [Carleo, Tarzia, Zamponi, PRL 2009]. We also investigated quantum extensions of lattice glass models [Biroli, Mézard, PRL 2001] and established their phase diagram by the cavity method [Foini, Semerjian, Zamponi, PRB 2011]: the main outcome of this study was the existence of a first order superfluid-glass phase, therefore accompanied by the coexistence of the two phases, which could induce an heterogeneous “superglass” phase in which some region would support the glassiness while other regions would support the superfluidity. Moreover, this study confirmed that quantum fluctuations promote glassiness at low enough temperature, a result that was also recently obtained by mean of quantum Mode-Coupling Theory [Markland et al., Nature Physics 2011].

(iii) *Perspectives: Stochastic dynamics of disordered systems* - As already discussed, there are a lot of potential applications of the method in this context. We are now starting to investigate the potentialities of the method in this domain. In the following we discuss two examples of concrete problems that can be tackled with the cavity method, but we hope that many other interesting application will emerge during the development of the project.

1. Classical optimization problems are often solved by mean of stochastic algorithms that perform a random walk in the space of configurations according to some local update rule. It can be proven in many cases that for random instances, a given algorithm will be effective, in the thermodynamic limit, up to a given

algorithmic threshold in the ratio of clauses to variables (which is a simple measure of the "difficulty" of a problem). Above the threshold the algorithm fails with very high probability. The analytical computation of algorithmic thresholds has been performed only in very simple cases; by mean of the dynamic cavity method we will able to compute them for a much wider class of algorithms, and hopefully we will obtain insight into the structural reasons that make these algorithms fail.
2. In many applications to chemistry and biology one is interested in the computation of *transition rates* between different stable states of a network. For large systems, it is tempting to identify these states as metastable states. Hence, one would like to compute transition rates between different metastable states of the network. We expect this to be possible for sparse random networks by mean of the dynamic cavity method, combined with standard sampling methods like the one proposed in [Dellago, Bolhuis, Csajka, Chandler, JCP 1998].

Summary
This project provided some new insight into some difficult open problems, characterized by a common ingredient: the interplay of quantum fluctuations, disorder and frustration. Nice perspectives of future development of these ideas are offered by the application to the stochastic dynamic of classical systems, with potential applications in chemistry, biology and information theory. I have been invited to write a News & Views paper for Nature Physics, that reviews our results and those of Markland et al. mentioned above [Zamponi, Nature Physics, 2011].