Statistical mechanics methods and phase transitions in optimization problems.

Olivier C. Martin,[a,1] Remi Monasson,[b,2] and Riccardo Zecchina[c,3]

[a]LPTMS, Universite Paris-Sud, Orsay, France

[b]The James Franck Institute, The University of Chicago, Chicago, IL.

[c]International Centre for Theoretical Physics, Trieste, Italy

Abstract

Recently, it has been recognized that phase transitions play an important role in the probabilistic analysis of combinatorial optimization problems. However, there are in fact many other relations that lead to close ties between computer science and statistical physics. This review aims at presenting the tools and concepts designed by physicists to deal with optimization or decision problems in an accessible language for computer scientists and mathematicians, with no prerequisites in physics. We first introduce some elementary methods of statistical mechanics and then progressively cover the tools appropriate for disordered systems. In each case, we apply these methods to study the phase transitions or the statistical properties of the optimal solutions in various combinatorial problems. We cover in detail the Random Graph, the Satisfiability, and the Traveling Salesman problems. References to the physics literature on optimization are provided. We also give our perspective regarding the interdisciplinary contribution of physics to computer science.

Key words: statistical physics, phase transitions, optimization
PACS: 64.60.Cn, 75.10.Nr, 02.60.Pn

[1] E-mail: martino@ipno.in2p3.fr

[2] Permanent address: CNRS-Laboratoire de Physique Theorique de l'ENS, Paris, France; E-mail: monasson@lpt.ens.fr

[3] E-mail: zecchina@ictp.trieste.it

1 Introduction

At the heart of statistical physics, discrete mathematics, and theoretical computer science, lie mathematically similar counting and optimization problems. This situation leads to a transgression of boundaries so that progress in one discipline can benefit the others. An old example of this is the work of Kasteleyn (a physicist) who introduced a method for counting perfect matchings over planar graphs (a discrete mathematics problem). Our belief is that a similar cross-fertilization of methods and models should arise in the study of combinatorial problems over random structures. Such problems have attracted the attention of a large community of researcher in the last decade, but a transgression of boundaries has only just begun. One of the many potential spin-offs of this kind of cross-fertilization would be the use of computer science and graph theoretical methods to tackle unsolved problems in the statistical physics of "complex" (disordered) systems. But we also hope that the benefits can go the other way, i.e., that the recent developments in statistical physics may be of use to the other two communities; such is our motivation for this article.

This review does not assume any knowledge in physics, and thus we expect it to be accessible to mathematicians and computer scientists eager to learn the main ideas and tools of statistical physics when applied to random combinatorics. We have chosen to illustrate these "physical" approaches on three problems: the Random Graph, the Satisfiability, and the Traveling Salesman problems. This particular focus should help the interested reader explore the statistical physics literature on decision and optimization problems. Furthermore, we hope to make the case that these methods, developed during the last twenty years in the context of the so called spin glass theory [1,2], may provide new concepts and results in the study of phase transitions, and average case computational complexity, in computer science problems. Some examples of this kind of methodological transfer can also be found in three other papers of this TCS special issue, dealing with statistical mechanics analyses of vertex covering on random graphs [3], of number partitioning [4] and of learning theory in artificial neural networks [5].

Random combinatorics became a central part of graph theory following the pioneering work by Erdos and Renyi. Their study of clusters in random graphs (percolation for physicists) showed the existence of zero-one laws (phase transitions in the terminology of physics). More recently, such phenomena have played a fundamental role when tackling average{case complexity. Indeed, numerical evidence suggests that the onset of intractability in random NP-complete problems can be put in relation with the appearance of phase transitions analogous to the percolation transition. Interestingly, the concept of random structures is present in most natural sciences, including biology, chem-

istry, or physics. But in the last two decades, the theoretical framework developed in physics has lead to new analytical and numerical tools that can be shared with the more mathematical disciplines. The potential connections between discrete mathematics, theoretical computer science and statistical physics become particularly obvious when one considers the typical properties of random systems. In such cases, percolation, zero-one laws, or phase transitions are simply different names describing the same phenomena within the different disciplines. It seems to us that much can be gained by exploring the complementary nature of the different paradigms in mathematics and physics. In what follows, we shall try to make this happen by giving a thorough statistical mechanics analysis of three prototype problems, namely percolation in random graphs, satisfiability in random K-Satisfiability, and optimization via the Traveling Salesman Problem. The review is preceded by a general discussion of some basic concepts and tools of statistical mechanics. We have also included simple exercises to help the interested reader become familiar with the methodology; hopefully he (she) will be able to adapt it to the study of many other problems, e.g., matching, number partitioning [4], etc... When appropriate, we compare the results of statistical physics to those of discrete mathematics and computer science.

From a statistical mechanics perspective, a phase transition is nothing but the onset of non-trivial macroscopic (collective) behavior in a system composed of a large number of "elements" that follow simple microscopic laws. The analogy with random graphs is straightforward. There the elements are the edges of the graph which are added at random at each time step and the macroscopic phenomenon is the appearance of a connected component of the graph containing a finite fraction of all the vertices, in the limit of a very large number of vertices. If a system has a phase transition, it can be in one of several "phases", depending on the values of some control parameters. Each phase is characterized by a different microscopic organization. Central to this characterization is the identification of an order parameter (usually the expectation value of a microscopic quantity) which discriminates between the different phases. Once again the analogy with random graphs is appropriate. An order parameter of the percolation transition is the fraction of vertices belonging to the giant connected component. Such a fraction is zero below the percolation transition, that is, when the connectivity of the random graph is too small, and becomes strictly positive beyond the percolation threshold.

While in percolation it is proven that the order parameter is indeed the fraction of vertices belonging to the infinite giant component, in more complicated systems the determination of an order parameter is generally an open problem. Though not rigourous, statistical mechanics provides numerous specific methods for identifying and studying order parameters, and we shall illustrate this on the K-Satisfiability problem. This step is useful of course for providing a good intuitive view of the system's behavior, but more importantly it also

gives information on the microscopic structure of the phases, information that can be used both in deriving analytical results and in interpreting numerical simulations.

The way physicists and mathematicians proceed is quite different. Theoretical physicists generally do not prove theorems, rather they attempt to understand problems by obtaining exact and approximate results based on reasonable hypotheses. In practice, these hypotheses are "validated" a posteriori through comparison with experiments or numerical simulations, and through consistency with the overall body of knowledge in physics. In this sense, theoretical physics must be distinguished from mathematical physics whose scope is to make rigorous statements. Of course, exact solutions play an important role in statistical physics in that they represent limiting cases where analytical or numerical techniques can be checked, but they are not the main focus of this discipline.

For the sake of brevity we left out from this review some very relevant and closely connected topics such as exact enumeration methods [6] or applications of computer science algorithms to the study of two dimensional complex physical systems [7,8]. Furthermore we do not claim to present a complete picture of what has been done by physicists on decision and optimization problems. Rather, we hope that what we do present will enable readers from the more mathematical disciplines to understand in detail the majority of what has been done by physicists using the methods of statistical mechanics.

2 Elements of Statistical Physics

In this section, the reader will be introduced to the basic notions of statistical mechanics. We start by illustrating on various examples the existence of phases and phase transitions, ubiquitous in physics and more surprisingly in other fields of science too. The concepts of microscopic and macroscopic levels of description naturally appear and allow for a rapid presentation of the foundations of statistical mechanics. We then expose in greater detail the combinatorial interpretation of statistical mechanics and introduce some key vocabulary and definitions. An accurate investigation of the properties of the so-called Ising model on the complete graph K_N exemplifies the above concepts and calculation techniques. In order to bridge the gap with optimization problems, we then turn to the crucial issue of randomness and present appropriate analytical techniques to deal with random structures, e.g., the celebrated replica method.

This section has been elaborated for a non physicist readers and we stress that no a priori knowledge of statistical mechanics is required. Exercises have

been included to illustrate key notions and should help the reader to acquire a deeper understanding of concepts and techniques. Solutions are sketched in Appendix A. Excellent presentations of statistical mechanics can be found in textbooks e.g. [9–11] for readers wanting further details.

2.1 Phases and transitions

Many physical compounds can exist in nature as distinct "states", called phases, depending on the values of control parameters, such as temperature, pressure, ... The change of phase happens very abruptly at some precise values of the parameters and is called transition. We list below a few well-known examples from condensed matter physics as well as two cases coming from biology and computer science.

2.1.1 Liquid-gas transition.

At atmospheric pressure water boils at a "critical" temperature $T_c = 100°C$. When the temperature T is lower than T_c, water is a liquid while above T_c it is a gas. At the critical temperature T_c, a coexistence between the liquid and gas phases is possible: the fraction of liquid water depends only on the total volume occupied by both phases. The coexistence of the two phases at criticality is an essential feature of the liquid-gas transition. Transitions sharing this property are called first order phase transitions for mathematical reasons exposed later.

2.1.2 Ferromagnetic-paramagnetic transition.

It is well-known that magnets attract nails made out of iron. The magnetic field produced by the magnet induces some strong internal magnetization in the nail resulting in an attractive force. Materials behaving as iron are referred to as ferromagnetic. However, the attractive force disappears when the temperature of the nail is raised above $T_c = 770°C$. The nail then enters the paramagnetic phase where the net magnetization vanishes. There is no phase coexistence at the critical temperature; the transition is said to be of second order.

The ferromagnetic-paramagnetic transition temperature T_c varies considerably with the material under consideration. For instance, $T_c = 1115°C$ for cobalt, $T_c = 454°C$ for nickel and $T_c = 585°C$ for magnetite (Fe_3O_4). However, remarkably, it turns out that some other quantities – the critical exponents related to the (drastic) changes of physical properties at or close to the transition – are equal for a large class of materials! The discovery of such universality

was a breakthrough and led to very deep theoretical developments in modern physics. Universality is characteristic of second order phase transitions.

2.1.3 Conductor-superconductor transition.

Good conductors such as copper are used to make electric wires because of their weak resistance to electric currents at room temperature. As the temperature is lowered, electrical resistance generally decreases smoothly as collisions between electrons and vibrations of the metallic crystal become weaker and weaker. In 1911, Kammerling Onnes observed that the electrical resistance of a sample of mercury fell abruptly down to zero as temperature passed through $T_c \simeq 4.2^\circ K$ ($0^\circ K$ being the absolute zero of the Kelvin scale.) This change of state, between a normal conductor (finite resistance) and a superconductor (zero resistance) is a true phase transition: a very small variation of temperature at T_c is enough to change resistance by four or five orders of magnitude!

2.1.4 DNA denaturation transition.

In physiological conditions, DNA has the double helix structure discovered by Watson and Crick in 1953. The two strands carry complementary sequences of A, T, G or C bases and are intertwined, forming either A-T or G-C pairs. Bases in a pair are attached together by hydrogen bonds. As the temperature is raised or ionic conditions are appropriately modified, bonds weaken and break up. The strands may then separate so that the double helix structure is lost: the DNA is denatured. This transition is abrupt on repeated homogeneous DNA sequences [12].

Recent micromanipulation experiments on individual DNA molecules have shown that denaturation can also be obtained through a mechanical action on DNA. When imposing a sufficient torque to the molecule to unwind the double helix, the latter opens up and DNA denaturates. At a fixed critical torque, denaturated and double helix regions may coexist along the same molecule [13] so this transition is like a liquid-gas one.

2.1.5 Transition in the random K-Satisfiability problem.

Computer scientists discovered some years ago that the random K-Satisfiability problem exhibits a threshold phenomenon as the ratio α of the number of clauses (M) over the number of Boolean variables (N) crosses a critical value $\alpha_c(K)$ depending on the number of literals per clause K. When α is smaller than the threshold $\alpha_c(K)$, a randomly drawn formula is almost surely satisfiable while, above threshold, it is unsatisfiable with probability reaching one in the $N \to \infty$ limit.

For $K = 2$, the threshold is known exactly: $_c(2) = 1$. For $K \quad 3$, there is no rigorous proof of the existence of a phase transition so far but many theoretical and numerical results strongly support it, see articles by Achlioptas & Franco and Dubois & Kirousis in the present issue. Current best estimates indicate that the threshold of random 3-SAT is located at $_c(3) \prime 4.25$. Statistical physics studies show that the order of the phase transition depends on K, the transition being continuous for 2-SAT and of first order for 3-SAT (and higher values of K).

2.1.6 Macroscopic vs. microscopic descriptions.

What can be inferred from the above examples? First, a (physical) system may be found in totally different phases with very different macroscopic properties although its intrinsic composition at a microscopic level (molecules, magnetic spins, base pairs, clauses, ...) is the same. However, from a physical, mechanical, electrical, biological, computational, ... point of view, essential properties of this system change completely from a phase to another. Second, the abrupt change of phase follows from very slight modifications of a control parameter e.g. temperature, torque, ratio of clauses per variable ... about a critical value. Thirdly, critical exponents, that characterize quantitatively second order phase transitions, are universal, that is, insensitive to many details of the systems under study. Last of all, transitions appear for large systems only.

The above points raise some fundamental questions: how can the main features of a system at a macroscopic level, defining a phase, change abruptly and how are these features related to the microscopic structure of the system? Statistical physics focuses on these questions.

2.2 Foundations of statistical mechanics and relationship with combinatorics.

2.2.1 Needs for a statistical description.

Statistical physics aims at predicting quantitatively the macroscopic behaviour of a system (and in particular its phases) from the knowledge of its microscopic components and their interactions. What do we mean by interaction? Consider for instance a liquid made of N small particles (idealized representations of atoms or molecules) occupying positions of coordinates r_i in Euclidean space where label i runs from 1 to N. Particle number i is subject to a force f_i (interaction) due to the presence of neighboring particles; this force generally depends of the relative positions of these particles. To determine the positions of the particles at any later time t, we must integrate the equations of motion

given by Newton's fundamental law of mechanics,

$$m_i \frac{d^2 \mathbf{r}_i}{dt^2} = \mathbf{f}_i(\{\mathbf{r}_j\}) \; , \qquad (i = 1, \ldots, N) \; , \tag{1}$$

where m_i is the mass of particle i. Solving these equations cannot be done in practice. The forces $\mathbf{f}_i$ are indeed highly non linear functions of the particle positions $\mathbf{r}_j$. We therefore wind up with a set of complicated coupled differential equations whose number N, of order 10^{23}, is gigantic and not amenable to analytical treatment.

This impossibility, added to the intuitive feeling that understanding macroscopic properties cannot require the exact knowledge of all microscopic trajectories of particles has been circumvented by a totally different approach. The basic idea is to describe the system of particles in a probabilistic way in order to deduce macroscopic features as emergent statistical properties.

2.2.2 Probability distribution over the set of configurations.

The implementation of this idea has required the introduction of revolutionary concepts at the end of the ninteenth century by Boltzmann and followers, and in particular, the ideas of ergodicity and thermodynamical equilibrium. We shall not attempt here to provide an exposition of these concepts. The interested reader can consult textbooks e.g. [9–11]. As far as combinatorial aspects of statistical mechanics are concerned, it is sufficient to start from the following postulate.

A configuration C of the system, that is, the specification of the N particle positions $\{\mathbf{r}_j\}$, has a probability $p(C)$ to be realized at any time when the system is in equilibrium. In other words, the system will be in configuration C with probability $p(C)$. The latter depends on temperature T and equals

$$p(C) = \frac{1}{Z} \exp\left(-\frac{1}{T} E(C)\right) \; : \tag{2}$$

In the above expression, E is the energy and is a real-valued function, over the set of configurations. The partition function Z ensures the correct normalization of the probability distribution p,

$$Z = \sum_C \exp\left(-\frac{1}{T} E(C)\right) \; : \tag{3}$$

Note that we have used a discrete sum over configurations C in (3) instead of an integral over particle positions $\mathbf{r}_j$. This notation has been chosen since all

the partition functions we shall meet in the course of studying optimization problems are related to nite (i.e. discrete) sets of con gurations.

Consider two limiting cases of (2):

in nite temperature $T = 1$: the probability $p(C)$ becomes independent of C. All con gurations are thus equiprobable. The system is in a fully "disordered" phase, like a gas or a paramagnet.

zero temperature $T = 0$: the probability $p(C)$ is concentrated on the minimum of the energy function E, called the ground state. This minimum corresponds to a con guration where all particles are at mechanically stable positions, that is, occupy positions r_i carefully optimized so that all forces f_i vanish. Often, these strong constraints de ne regular packings of particles and the system achieves a perfect crystalline and "ordered" state.

When varying the temperature, intermediate situations can be reached. We now examine some simple examples.

2.2.3 Cases of one and two spins.

We now consider the case of a single abstract particle that can sit at two di erent positions only. This simple system can be recast as follows. Let us imagine an arrow capable of pointing in the up or down directions only. This arrow is usually called a spin and the direction is denoted by a binary variable , equal to $+1$ if the spin is up, to -1 if the spin is down.

In this single particle system, there are only two possible con gurations $C = \{+1\}$ and $C = \{-1\}$ and we choose for the energy function $E( ) = \quad$. Note that additive constants in E have no e ect on (2) and multiplicative constants can be absorbed in the temperature T. The partition function can be easily computed from (3) and reads $Z = 2\cosh \quad$ where $= 1/T$ denotes the inverse temperature. The probabilities that the spin points up or down are respectively $p_+ = \exp( )/Z$ and $p = \exp( )/Z$. At in nite temperature ($= 0$), the spin is indi erently up or down: $p(+1) = p(-1) = 1/2$. Conversely, at zero temperature, it only points upwards: $p(+1) = 1; p(-1) = 0$. $C = \{+1\}$ is the con guration of minimum energy.

The average value of the spin, called magnetization is given by

$$m = \langle \rangle_T = \sum_{= \pm 1} p( ) = \tanh( ) \quad : \tag{4}$$

The symbol $\langle \rangle_T$ denotes the average over the probability distribution p. Notice that, when the temperature is lowered from $T = 1$ down to $T = 0$, the magnetization increases smoothly from $m = 0$ up to $m = 1$. There is no

abrupt change (singularity or non analyticity) in m as a function of $$ and therefore no phase transition.

Exercise 1: Consider two spins $_1$ and $_2$ with energy function

$$E(_1; _2) = \quad _1 \, _2 \quad : \tag{5}$$

Calculate the partition function, the magnetization of each spin as well as the average value of the energy. Repeat these calculations for

$$E(_1; _2) = \quad _1 \quad _2 \quad : \tag{6}$$

How is the latter choice related to the single spin case?

2.2.4 Combinatorial meaning of the partition function.

We have so far introduced statistical mechanics in probabilistic terms. There exists also a close relationship with combinatorics through the enumeration of configurations at a given energy; we now show this relationship.

The average value of the energy may be computed directly from the definition

$$\langle E \rangle_T = \sum_C p(C) E(C) \quad ; \tag{7}$$

or from the partition function Z via the following identity

$$\langle E \rangle_T = \quad \frac{d}{d} \ln Z \quad ; \tag{8}$$

that can easily derived from (3). The identity (8) can be extended to higher moments of the energy. For instance, the variance of E can be computed from the second derivative of the partition function

$$\langle E^2 \rangle_T \quad \langle E \rangle_T^2 = \frac{d^2}{d^2} \ln Z \quad : \tag{9}$$

Such equalities suggest that Z is the generating function of the configuration energies. To prove this statement, let us rewrite (3) as

$$Z = \sum_C \exp(\quad E(C))$$

$$= \sum_E N(E) \exp(-\beta E) \quad ; \tag{10}$$

where $N(E)$ is the number of configurations C having energies $E(C)$ precisely equal to E. If $x = \exp(-\beta)$, $Z(x)$ is simply the generating function of the coefficients $N(E)$ as usually defined in combinatorics.

The quantity $\hat{S}(E) = \ln N(E)$ is called the entropy associated with the energy E. In general, calculating $\hat{S}(E)$ is a very hard task. Usually, it is much more convenient to define the average entropy $\langle S \rangle_T$ at temperature T as the contribution to the partition function which is not directly due to energy,

$$\langle S \rangle_T = \frac{1}{T} \left(F(T) - \langle E \rangle_T \right) \quad ; \tag{11}$$

where

$$F(T) = -T \ln Z(T) \tag{12}$$

is called the free-energy of the system.

In general, the above definitions for the energy and temperature dependent entropies do not coincide. However, as explained in next Section, in the large size limit $\langle S \rangle_T$ equals $\hat{S}(E)$ provided that the energy E is set to its thermal average $E = \langle E \rangle_T$.

The entropy is an increasing function of temperature. At zero temperature, it corresponds to the logarithm of the number of absolute minima of the energy function $E(C)$.

Exercise 2: Prove this last statement.

2.2.5 Large size limit and onset of singularity.

We have not encountered any phase transition in the above examples of systems with one or two spins. A necessary condition for the existence of a transition in a system is indeed that the size of the latter goes to infinity. The mathematical reason is simple: if the number of terms in the sum (3) is finite, the partition function Z, the free-energy F, the average energy, ... are analytic functions of the inverse temperature β and so do not have singularities at finite temperature.

Most analytical studies are therefore devoted to the understanding of the emergence of singularities in the free-energy when the size of the system goes

to in nity, the so-called thermodynamic limit.

An important feature of the thermodynamic limit is the concentration of measure for observables e.g. energy or entropy. Such quantities do not uctuate much around their mean values. More precisely, if we call N the size, i.e. the number of spins, of the system, the moments of the energy usually scale as

$$\begin{aligned} \langle E \rangle_T &= O(N) \\ \langle E^2 \rangle_T - \langle E \rangle_T^2 &= O(N) \end{aligned} \quad ; \tag{13}$$

and, thus the energy of a con guration is with high probability equal to the average value up to $O(\sqrt{N})$ uctuations. Such a result also applies to the entropy, and $\langle S \rangle_T = \hat{S}(\langle E \rangle_T)$ up to $O(\sqrt{N})$ terms. Measure concentration in the thermodynamic limit is a very important and useful property, see [14].

2.3 Spin model on the complete graph.

We shall now study a system of N spins, called the Ising model, exhibiting a phase transition in the limit $N \to \infty$. We consider the complete graph K_N; each vertex is labelled by an integer number $i = 1, \ldots, N$ and carries a binary spin σ_i. The energy function of a con guration $C = \{\sigma_1, \ldots, \sigma_N\}$ is given by

$$E(\sigma_1, \ldots, \sigma_N) = -\frac{1}{N} \sum_{i<j} \sigma_i \sigma_j - h \sum_i \sigma_i \quad . \tag{14}$$

2.3.1 Remarks on the energy function.

The rst term in (14) is called the interaction term. The sum runs over all pairs of spins, that is over all edges of K_N. The minus sign ensures that the minimum of energy is reached when all spins point in the same direction. This direction depends on the second term of (14) and, more precisely, upon the sign of the "magnetic eld" h. If the latter is positive (respectively negative), the ground state is obtained when all spins are up (resp. down).

In the absence of eld ($h = 0$), we know the two ground states. The energy and entropy at zero temperature can be computed from (14) and (11),

$$\langle E \rangle_{T=0} = -\frac{1}{2}(N - 1) \quad ; \tag{15}$$

$$\langle S \rangle_{T=0} = \ln 2 \quad . \tag{16}$$

Notice that the ground state energy is $O(N)$ due to the presence of the factor $1/N$ in (14) whereas the entropy is $O(1)$.

At infinite temperature, all configurations are equiprobable. The partition function is simply equal to the total number of configurations: $Z_{T=\infty} = 2^N$, leading to

$$\langle E \rangle_{T=\infty} = 0 \quad ; \tag{17}$$

$$\langle S \rangle_{T=\infty} = N \ln 2 \quad : \tag{18}$$

When the temperature is finite, a compromise is realized in (10) between energy and entropy: the configurations with low energies E have the largest probabilities but the most probable energy also depends on the entropy, i.e. on the size of the coefficients $N(E)$. Temperature tunes the relative importance of these two opposite effects. The phase transition studied in this section separates two regimes:

- a high temperature phase where entropy effects are dominant: spins configurations are disordered and spins do not point in any priviledged direction (for $h = 0$). The average magnetization m vanishes.
- a low temperature phase where energy effects dominate: spins have a tendency to align with each other, resulting in ordered configurations with a non zero magnetization $m = \langle \sigma_i \rangle_T \neq 0$.

Let us stress that the energy and the entropy must have the same orders of magnitude ($= O(N)$) to allow for such a compromise and thus for the existence of a phase transition at finite strictly positive temperature.

2.3.2 The magnetization is the order parameter.

We start by defining the magnetization of a configuration $C = \{\sigma_1, \ldots, \sigma_N\}$ as

$$m(C) = \frac{1}{N} \sum_{i=1}^{N} \sigma_i \quad : \tag{19}$$

The calculation of the partition function relies on the following remark. The energy function (14) depends on the configuration C through its magnetization $m(C)$ only. More precisely,

$$E(C) = N \left(\frac{1}{2} m(C)^2 + h\, m(C) + \frac{1}{2} \right) \quad : \tag{20}$$

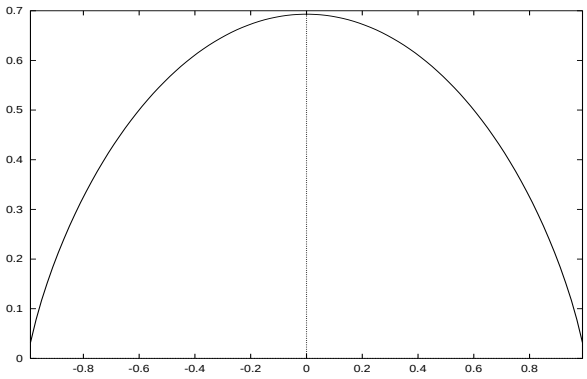

Fig. 1. Entropy $s(m)$ of the Ising model on the complete graph as a function of magnetization m.

In the following, we shall also need the entropy at fixed magnetization $S(m)$. Configurations with a fixed magnetization m have N_+ spins up and N_- spins down with

$$N_+ = N \frac{1+m}{2} \quad ; \qquad N_- = N \frac{1-m}{2} \quad : \tag{21}$$

The number of such configurations is therefore given by the binomial coefficient

$$e^{S(m)} = \frac{N!}{N_+! N_-!} \quad : \tag{22}$$

In the large N limit, Stirling's formula gives access to the asymptotic expression of the entropy density, $s(m) = S(m)/N$, at fixed magnetization,

$$s(m) = -\frac{1-m}{2} \ln \frac{1-m}{2} - \frac{1+m}{2} \ln \frac{1+m}{2} \quad ; \tag{23}$$

Figure 1 displays $s(m)$ as a function of m. The maximum is reached at zero magnetization ($s(0) = \ln 2$) and the entropy vanishes on the boundaries $m = \pm 1$.

Let us stress that $S(m)$ defined in (23) is the entropy at given magnetization and differs a priori from the energy and temperature dependent entropies, $\hat{S}(E)$ and $\langle S \rangle_T$, defined above. However, in the thermodynamic limit, all quantities are equal provided that m and E coincide with their thermal averages,

$\langle m \rangle_T$ and $\langle E \rangle_T$.

The average value $\langle m \rangle_T$ of the magnetization will be shown to vanish in the high temperature phase and to be different from zero in the low temperature phase. The magnetization is an order parameter: its value (zero or non-zero) indicates in which phase the system is.

2.3.3 Calculation of the free-energy.

The partition function Z reads

$$Z = \sum_{\sigma_1,\ldots,\sigma_N = \pm 1} \exp[-\beta E(\sigma_1,\ldots,\sigma_N)] = \sum_{m = -1, -1+\frac{2}{N}, \ldots, 1-\frac{2}{N}, 1} \exp\left[N \hat{f}(m) \right] \quad ; \tag{24}$$

where

$$\hat{f}(m) = -\frac{1}{2} m^2 - h m - T s(m) \quad ; \tag{25}$$

up to $O(1/N)$ terms. For the moment, we shall take $h = 0$.

In the limit of an infinite number N of spins, the free-energy may be computed by means of the saddle-point (Laplace) method. We look for the saddle-point magnetization m (that depends upon temperature T) minimizing $\hat{f}(m)$ (25). The latter is plotted in Figure 2 for three different temperatures.

It can be seen graphically that the minimum of $\hat{f}$ is located at $m = 0$ when the temperature is larger than $T_c = 1$ while there exist two opposite minima, $m = -m^*(T) < 0$, $m = m^*(T) > 0$ below this critical temperature. The optimum magnetization is solution of the saddle-point equation,

$$m^* = \tanh(\beta m^*) \quad ; \tag{26}$$

while the free-energy is given by

$$f(T) = \lim_{N \to \infty} -\frac{T}{N} \ln Z = \hat{f}(m^*) \quad : \tag{27}$$

The average energy and entropy per spin (divided by N) can be computed from (27, 8, 11),

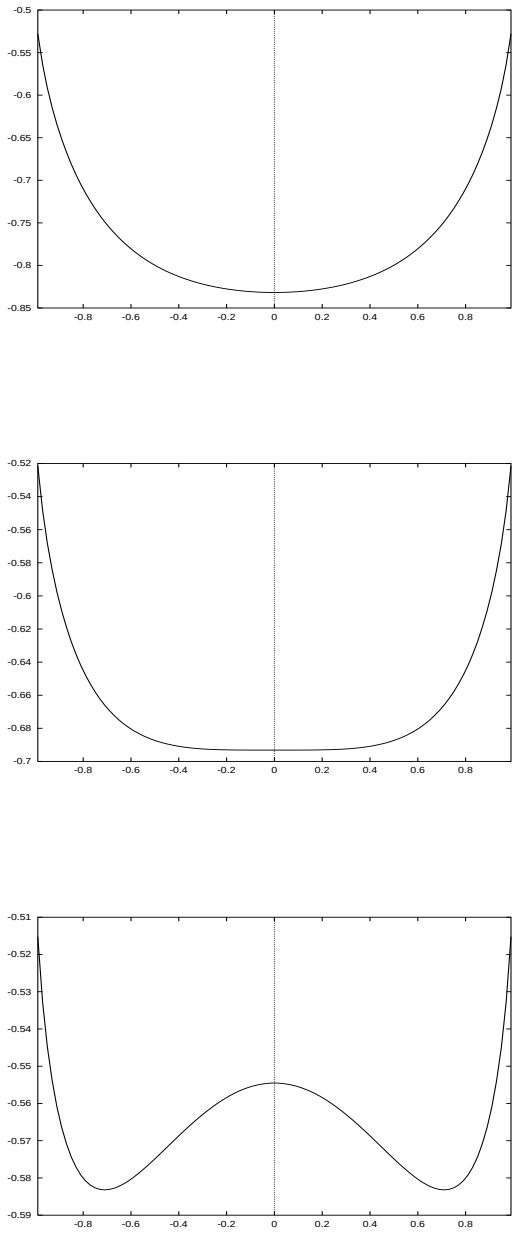

Fig. 2. Free-energy function $\hat{f}(m)$ of the Ising model on the complete graph as a function of the magnetization m in zero magnetic field h and for three different temperatures. a: high temperature $T = 1.2$, b: critical temperature $T = 1$, c: low temperature $T = 0.8$.

$$\langle e \rangle_T = -\frac{1}{2}(m^*)^2 \quad ; \tag{28}$$

$$\langle s \rangle_T = s(m^*) \quad : \tag{29}$$

2.3.4 Phase transition and symmetry breaking.

In the absence of a magnetic field, the energy (14) is an even function of the spins: the probability of two opposite configurations $\{\sigma_1, \ldots, \sigma_N\}$ and $\{-\sigma_1, \ldots, -\sigma_N\}$ are equal. As a consequence, the thermal average $\langle \sigma_i \rangle_T$ of any spin vanishes. This result is true for any N and so, in the large N limit,

$$\lim_{h \to 0} \lim_{N \to \infty} \langle \sigma_i \rangle_T = 0 \quad : \tag{30}$$

It is thus necessary to unveil the meaning of the saddle-point magnetization m arising in the computation of the partition function.

To do so, we repeat the previous calculation of the free-energy in presence of a magnetic field $h > 0$. The magnetization is now different from zero. At high temperature $T > T_c$, this magnetization decreases as the magnetic field h is lowered and vanishes when $h = 0$,

$$\lim_{h \to 0^+} \lim_{N \to \infty} \langle \sigma_i \rangle_T = 0 \qquad (T > T_c) \quad : \tag{31}$$

Therefore, at high temperature, the inversion of limits between (30) and (31) has no effect on the final result.

The situation drastically changes at low temperature. When $T < T_c$, the degeneracy between the two minima of f is lifted by the magnetic field. Due to the field, a contribution $-hm$ must be added to the free-energy (25) and favours the minimum in m over that in $-m$. The contribution to the partition function (24) coming from the second minimum is exponentially smaller than the contribution due to the global minimum in m by a factor $\exp(-2Nhm)$. The probability measure on spins configurations is therefore fully concentrated around the global minimum with positive magnetization and

$$\lim_{h \to 0^+} \lim_{N \to \infty} \langle \sigma_i \rangle_T = m \qquad (T < T_c) \quad : \tag{32}$$

From (30) and (32), the meaning of the phase transition is now clear. Above the critical temperature, a small perturbation of the system (e.g. a term in the energy function pushing spins up), is irrelevant: as the perturbation disappears ($h \to 0$), so do its effects ($m \to 0$), see (31). Conversely, below the critical temperature, a small perturbation is enough to trigger strong effects: spins point up (with a spontaneous magnetization $m > 0$) even after the perturbation has disappeared ($h = 0$), see (32). At low temperature, two phases with opposite magnetizations m and $-m$ coexist. Adding an infinitesimal field h favours and selects one of them. In more mathematical terms, the magnetization m is a non-analytic and discontinuous function of h at $h = 0$.

So, the phase transition here appears to be intimately related to the notion of symmetry breaking. In the case of the Ising model, the probability distribution over configurations is symmetrical, that is, left unchanged under the reversal of spins ! . A high temperature, this symmetry also holds for average quantities: h i$_T$ = 0. At low temperature, the reversal symmetry is broken since, in presence of an infinitesimal perturbation, h i$_T$ = m 6 0. The initial symmetry of the system implies only that the two possible phases of the system have opposite magnetizations m and m .

In the present case, the symmetry of the system was easy to identify, and to break! We shall see that more abstract and complex symmetries may arise in other problems, e.g. the random graph and K-Satisfiability. The understanding of phase transitions very often will rely on the breaking of associated symmetries.

Exercise 3: How does equation (26) become modified when there is a non-zero magnetic field? Calculate explicitely the free-energy in presence of a magnetic field and check the correctness of the above statements.

2.3.5 Vicinity of the transition and critical exponents.

To complete the present analysis, we now investigate the properties of the Ising model close to the critical temperature $T_c = 1$ and define $T = 1 + $ with $| | \ll 1$. The spontaneous magnetization reads from (26),

$$m( ) = \begin{cases} 0 & \text{if } \quad 0 \ , \\ \sqrt{\frac{3}{\quad}} & \text{if } \quad 0 \ . \end{cases} \tag{33}$$

Thus the magnetization grows as a power of the shifted temperature : $m( ) \sim ( )^{\ }$ with $= 1/2$. , not to be confused with the inverse temperature, is called a critical exponent since it characterizes the power law behaviour of a physical quantity, here the magnetization, close to criticality. Such exponents are universal in that they are largely independent of the "details" of the definition of the model. We shall come back to this point in the sections devoted to the random graph and the K-Satisfiability models.

Another exponent of interest is related to the finite size effect at the transition. So far, we have calculated the average values of various quantities in the infinite size limit N ! 1 . We have in particular shown the existence of a critical temperature separating a phase where the sum of the spins is on average zero (> 0) from a phase where the sum of the spins acquires an O(N) mean (< 0). At the transition point (= 0), we know that the sum of spins cannot be of order N ; instead we have a scaling in N with < 1.

What is the value of ? From expression (24), let us expand the free-energy function $\hat{f}(m)$ (25) in powers of the magnetization $m = O(N^{-1})$,

$$f(m) - f(0) = \frac{}{2} m^2 + \frac{1}{12} m^4 + O(m^6; \ m^4) \quad ; \tag{34}$$

with $f(0) = T \ln 2$. Above the critical temperature, > 0, the average magnetization is expected to vanish. Due to the presence of the quadratic leading term in (34), the fluctuations of m are of the order of $N^{-1=2}$. The sum of the spins, $N m$, has a distribution whose width grows as $N^{1=2}$, giving $= 1=2$.

At the critical temperature, the partition function reads from (24),

$$Z \sim 2^N \int dm \, e^{-N m^4=12} \quad : \tag{35}$$

The average magnetization thus vanishes as expected and fluctuations are of the order of $N^{-1=4}$. The sum of the spins, $N m$, thus has a distribution whose width grows as $N^{3=4}$, giving $= 3=4$.

The size of the critical region (in temperature) is defined as the largest value $_{max}$ of the shifted temperature leaving unchanged the order of magnitude of the fluctuations of the magnetization m. A new critical exponent that monitors this shift is introduced: $_{max} \quad N^{-1=}$. Demanding that terms on the r.h.s. of (34) be of the same order in N, we find $= 2$.

2.4 Randomness and the replica method.

The above analysis of the Ising model has been useful to illustrate some classic analytical techniques and to clarify the concept of phase transitions. However, most optimization or decision problems encountered in computer science contain another essential ingredient we have not discussed so far, namely randomness. To avoid any confusion, let us stress that randomness in this case, e.g. a Boolean formula randomly drawn from a well-defined distribution, and called quenched disorder in physics, must be clearly distinguished from the probabilistic formulation of statistical mechanics related to the existence of thermal disorder, see (2). As already stressed, as far as combinatorial aspects of statistical mechanics are concerned, we can start from the definition (10) of the partition function and interpret it as a generating function, forgetting the probabilistic origin. On the contrary, quenched disorder cannot be omitted. We are then left with combinatorial problems defined on random structures, that is, with partition functions where the weights themselves are random variables.

2.4.1 Distribution of "quenched" disorder.

We start with a simple case:

Exercise 4: Consider two spins σ_1 and σ_2 with energy function

$$E(\sigma_1, \sigma_2) = -J\,\sigma_1\,\sigma_2\,; \tag{36}$$

where J is a real variable called coupling. Calculate the partition function, the magnetization of each spin as well as the average value of the energy at given (quenched) J. Assume now that the coupling J is a random variable with measure $\rho(J)$ on a finite support $[J_-, J_+]$. Write down the expressions of the mean over J of the magnetization and energy. What is the value of the average ground state energy?

The meaning of the word "quenched" is clear from the above example. Spins are always distributed according to (2) but the energy function E now depends on randomly drawn variables e.g. the coupling J. Average quantities (over the probability distribution p) must be computed keeping these random variables fixed (or quenched) and thus are random variables themselves that will be averaged over J later on. To distinguish both kinds of averages we hereafter use an overbar to denote the average over the quenched random variables while brackets still indicate a thermal average using p.

Models with quenched randomness are often very difficult to solve. One of the reasons is that their physical behaviour is more complex due to the presence of frustration.

2.4.2 Notion of frustration.

Frustration is best introduced through the following simple example.

Exercise 5: Consider three spins σ_1, σ_2 and σ_3 with energy function

$$E(\sigma_1, \sigma_2, \sigma_3) = -\sigma_1\,\sigma_2 - \sigma_1\,\sigma_3 - \sigma_2\,\sigma_3\,; \tag{37}$$

Calculate the partition function, the magnetization of each spin as well as the average value of the energy. What are the ground state energy and entropy?

Repeat the calculation and answer the same questions for

$$E(\sigma_1, \sigma_2, \sigma_3) = -\sigma_1\,\sigma_2 - \sigma_1\,\sigma_3 + \sigma_2\,\sigma_3\,; \tag{38}$$

Note the change of the last sign on the r.h.s. of (38).

The presence of quenched disorder with both negative and positive couplings generates frustration, that is conicting terms in the energy function. A famous example is the Sherrington-Kirkpatrick (SK) model, a random version of the Ising model on the complete graph whose energy function reads

$$E_{SK}(\sigma_1,\ldots,\sigma_N) = -\frac{1}{\sqrt{N}} \sum_{i<j} J_{ij}\,\sigma_i\,\sigma_j \quad , \tag{39}$$

where the quenched couplings J_{ij} are independent random normal variables. In the SK model, contrarily to the Ising model, the product of the couplings J_{ij} along the loops of the complete graph K_N may be negative. The ground state is no longer given by the "all spins up" configuration, nor by any simple prescription and must be sought for among the set of 2^N possible configurations. Finding the ground state energy for an arbitrary set of couplings J_{ij} is a hard combinatorial optimization task which in this case belongs to the class of NP-hard problems [15,16].

2.4.3 Thermodynamic limit and self-averaging quantities.

Though physical quantities depend a priori on quenched couplings, some simplications may take place in the large size limit $N \to \infty$. Many quantities of interest may exhibit less and less uctuations around their mean values and become self-averaging. In other words, the distributions of some random variables become highly concentrated as N grows. Typical examples of highly concentrated quantities are the (free-)energy, the entropy, the magnetization, ... whereas the partition function is generally not self-averaging.

Self-averaging properties are particularly relevant when analyzing a problem. Indeed, for these quantities, we only have to compute their average values, not their full probability distributions. We shall encounter numerous examples of concentrated random variables later in this article.

Exercise 6: Show that the partition function of the SK model is not self-averaging by calculating its rst two moments.

2.4.4 Replica method.

We consider a generic model with N spins σ_i and an energy function $E(C;J)$ depending on a set of random couplings J. Furthermore we assume that the free-energy $F(J)$ of this model is self-averaging and would like to compute its

quenched averaged value $\overline{F(J)}$ or, equivalently from (12), the averaged logarithm of the partition function $\overline{\ln Z(J)}$. Though well posed, this computation is generally a very hard task from the analytical point of view. An original but non rigorous method, the replica approach, was invented by Kac in the sixties to perform such calculations. The starting point of the replica approach is the following expansion

$$Z(J)^n = 1 + n \ln Z(J) + O(n^2) \quad ; \tag{40}$$

valid for any set of couplings J and small real n. The identity (40) may be averaged over couplings and gives the mean free-energy from the averaged n^{th} power of the partition function

$$\overline{F(J)} = -T \lim_{n \to 0} \frac{\overline{Z(J)^n} - 1}{n} \quad . \tag{41}$$

If we restrict to integer n, the n^{th} moment of the partition function Z can be rewritten as

$$\begin{aligned} \overline{Z(J)^n} &= \overline{\left[\sum_C \exp -\frac{1}{T} E(C;J) \right]^n} \\ &= \sum_{C^1;\ldots;C^n} \overline{\exp -\frac{1}{T} \sum_{a=1}^{n} E(C^a;J)} \quad . \end{aligned} \tag{42}$$

This last expression makes transparent the principle of the replica method. We have n copies, or replicas, of the initial problem. The random couplings disappear once the average over the quenched couplings has been carried out. Finally, we must compute the partition function of an abstract system of N vectorial spins $\tilde{\sigma}_i = (\sigma_i^1;\ldots;\sigma_i^n)$ with the non random energy function

$$E_{eff}(\{\tilde{\sigma}_i\}) = -T \ln \left[\exp -\frac{1}{T} \sum_{a=1}^{n} E(C^a;J) \right]^3 \quad . \tag{43}$$

This new partition function can be estimated analytically in some cases by means of the saddle-point method just as we did for the Ising model. The result may be written formally as

$$\overline{Z(J)^n} = \exp -N\tilde{f}(n) \quad ; \tag{44}$$

to leading order in N. On general grounds, there is no reason to expect the partition function to be highly concentrated. Thus, $\tilde{f}(n)$ is a non linear function of its integer argument n satisfying $\tilde{f}(0) = 0$. The core idea of the

replica approach is to continue analytically f^n to the set of real n and obtain $\overline{F}(J) = TN\, df^n/dn$ evaluated at $n = 0$. The existence and uniqueness of the analytic continuation is generally ensured for finite sizes N due to the moment theorem. In most problems indeed one succeeds in bounding $|Z(J)|$ from above by a (J independent) constant C. The moments of Z grow only exponentially with n and their knowledge allows for a complete reconstruction of the probability distribution of $Z(J)$. However this argument breaks down when the saddle-point method is employed and the upper bound $C = \exp(O(N))$ becomes infinite.

Though there is generally no rigorous scheme for the analytic continuation when $N \to \infty$, physicists have developped in the past twenty years many empirical rules to use the replica method and obtain precise and sometimes exact results for the averaged free-energy. We shall see in the case of the K-Satisfiability problem how the replica approach can be applied and how very peculiar phase transitions, related to the abstract "replica" symmetry breaking, are present.

The mathematician or computer scientist reader of this brief presentation may feel uneasy and distrustful of the replica method because of the uncontrolled analytic continuation. To help him/her loose some inhibitions, he/she is asked to consider the following warming up exercise:

Exercise 7: Consider Newton's binomial expression for $(1+x)^n$ with integer n and perform an analytic continuation to real n. Take the $n \to 0$ limit and show that this leads to the series expansion in x of $\ln(1+x)$.

3 Random Graphs

In this section, we show how the statistical mechanics concepts and techniques exposed in the previous section allow to reproduce some famous results of Erdos and Renyi on random graphs [17].

3.1 Generalities

First let us define the random graphs used. Consider the complete graph K_N over N vertices. We define G_{N,N_L} as the set of graphs obtained by taking only $N_L = \alpha N/2$ among the $\binom{N}{2}$ edges of K_N in all possible different ways. A random graph is a randomly chosen element of G_{N,N_L} with the flat measure. Other random graphs can be generated from the complete graph K_N through a random deletion process of the edges with probability $1 - \alpha/N$. In the large

N limit, both families of random graphs share common properties and we shall mention explicitely the precise family we use only when necessary.

3.1.1 Connected components.

We call "clusters" the connected components of a given graph G; the "size" of a cluster is the number of vertices it contains. An isolated vertex is a cluster of size unity. The number of connected components of G is denoted by $C(G)$ and we shall indicate its normalized fraction by $c(G) = \frac{C}{N}$. If c is small, the random graph G has few big clusters whereas for c approaching unity there are many clusters of small size. Percolation theory is concerned with the study of the relationship between the probability p of two vertices being connected with the typical value of c in the $N \to \infty$ limit. The scope of this section is to show how such a relationship can be exploited by the study of a statistical mechanics model, the so called Potts model, after a suitable analytic continuation. As a historical note, let us mention that analytic continuations have played an enormous role in physics this last century, leading often to unexpected deep results, impossible or very difficult to obtain by other means.

3.1.2 Generating function for clusters.

Let $P(G)$ be the probability of drawing a random graph G through the deletion process from the complete graph K_N. Since the edge deletions are statistically independent, this probability depends on the number of edges N_L only, and factorizes as

$$P(G) = p^{N_L(G)}(1-p)^{\frac{N(N-1)}{2} - N_L(G)} \quad ; \tag{45}$$

where

$$1 - p = 1 - \frac{}{N} \tag{46}$$

is the probability of edge deletion. We want to study the probability density (c) of generating a random graph with c clusters,

$$(c) = \sum_G P(G) \quad (c - c(G)) \quad ; \tag{47}$$

where $$ indicates the Dirac distribution.

We can introduce a generating function of the cluster probability by

$$
\begin{aligned}
Y(q) &= \int_0^{+\infty} dc\ \Gamma(c) q^{N c} \\
&= \int_0^{+\infty} dc q^{N c} \sum_{G \in K_N} P(G)\ \delta(c - c(G)) \\
&= \sum_{G \in K_N} P(G) q^{C(G)} = \sum_{G \in K_N} p^{L(G)} (1-p)^{\frac{N(N-1)}{2} - L(G)} q^{C(G)} \ ; \qquad (48)
\end{aligned}
$$

with q being a formal (eventually real) parameter.

3.1.3 Large size limit.

In the large size limit, $\Gamma(c)$ is expected to be highly concentrated around some value $c(\alpha)$ equal to the typical fraction of clusters per vertex and depending only the average degree of valency α. Random graphs whose $c(G)$ differs enough from $c(\alpha)$ will be exponentially rare in N. Therefore, the quantity

$$
\omega(c) = \lim_{N \to \infty} \frac{1}{N} \log \Gamma(c) \qquad (49)
$$

should vanish for $c = c(\alpha)$ and be strictly negative otherwise. In the following, we shall compute $\omega(c)$ and thus obtain information not only on the typical number of clusters but also on the large deviations (rare events).

Defining the logarithm $f(q)$ of the cluster generating function as

$$
f(q) = \lim_{N \to \infty} \frac{1}{N} \log Y(q) \ ; \qquad (50)
$$

we obtain from a saddle-point calculation on c, see (48,49),

$$
f(q) = \max_{0 \le c \le 1} c \ln q + \omega(c) \ . \qquad (51)
$$

In other words, f and ω are simply conjugated Legendre transforms. It turns out that a direct computation of f is easier and thus prefered.

3.2 Statistical mechanics of the random graph.

Hereafter, we proceed to compute the properties of random graphs by using a mapping to the so-called Potts model. Some know results can be rederived by the statistical mechanics approach, and additional predictions are made.

3.2.1 Presentation of the Potts model.

The Potts model[18] is defined in terms of an energy function which depends on N spin variables σ_i, one for each vertex of the complete graph K_N, which take q distinct values $\sigma_i = 0, 1, \ldots, q-1$. The energy function reads

$$E[\{\sigma_i\}] = \sum_{i<j} \delta(\sigma_i, \sigma_j) \; , \tag{52}$$

where $\delta(a,b)$ is the Kronecker delta function: $\delta(a,b) = 1$ if $a = b$ and $\delta(a,b) = 0$ if $a \neq b$. The partition function of the Potts model is

$$Z_{Potts} = \sum_{\{\sigma_i = 0, \ldots, q-1\}} \exp\Big[\beta \sum_{i<j} \delta(\sigma_i, \sigma_j)\Big] \tag{53}$$

where β is the inverse temperature and the summation runs over all q^N spin configurations.

In order to identify the mapping between the statistical mechanics features of the Potts model and the percolation problem in random graphs we compare the expansion of Z_{Potts} to the definition of the cluster generating function of the random graphs.

3.2.2 Expansion of the Potts partition function.

Following Kasteleyn and Fortuin [19], we start by rewriting Z_{Potts} as a dichromatic polynomial. Upon posing

$$v = e^{\beta} - 1 \; , \tag{54}$$

one can easily check that (53) can be recast in the form

$$Z_{Potts} = \sum_{\{\sigma_i\}} \prod_{i<j} \big[1 + v\, \delta(\sigma_i, \sigma_j)\big] \; : \tag{55}$$

When σ_i and σ_j take the same value there appears a factor $(1+v)$ in the product (corresponding to a term e^{β} in (53)); on the contrary, whenever σ_i and σ_j are different the product remains unaltered. The expansion of the above product reads

$$Z_{Potts} = \sum_{\{\sigma_i\}} \Big[1 + v \sum_{i<j} \delta(\sigma_i, \sigma_j) + v^2 \sum_{i<j;\, k<l = (i,j) \neq (k,l)} \delta(\sigma_i, \sigma_j)\, \delta(\sigma_k, \sigma_l) + \ldots\Big] : \tag{56}$$

We obtain $2^{\frac{N(N-1)}{2}}$ terms each of which composed by two factors, the first one given by v raised to a power equal to the number of σ's composing the second factor. It follows that each term corresponds to a possible subset of edges on K_N, each edge weighted by a factor v. There is a one–to–one correspondence between each term of the sum and the sub–graphs G of K_N. The edge structure of each sub–graph is encoded in the product of the σ's. This fact allows us to rewrite the partition function as a sum over sub–graphs

$$Z_{Potts} = \sum_{\{i\}} \sum_{G \subset K_N} [v^{L(G)} \prod_{k=0}^{L(G)} \delta(\sigma_{i_k}, \sigma_{j_k})] \tag{57}$$

where $L(G)$ is the number of edges in the sub–graph G and i_k, j_k are the vertices connected by the k-th edge of the sub–graph. We may now exchange the order of the summations and perform the sum over the spin configurations. Given a sub–graph G with L links and C clusters (isolated vertices included), the sum over spins configurations will give zero unless all the σ's belonging to a cluster of G have the same value (cf. the δ functions). In such a cluster, one can set the σ's to any of the q different values and hence the final form of the partition function reads

$$Z_{Potts} = \sum_{G \subset K_N} v^{L(G)} q^{C(G)} \ . \tag{58}$$

3.2.3 Connection with the cluster generating function

If we now make the following identification

$$p = 1 - e^{-\beta} = v/(1+v) \ , \tag{59}$$

we can rewrite the partition function as

$$\begin{aligned} Z_{Potts} &= \sum_{G \subset K_N} \left(\frac{p}{1-p}\right)^{L(G)} q^{C(G)} \\ &= (1-p)^{\frac{N(N-1)}{2}} \sum_{G \subset K_N} p^{L(G)} (1-p)^{-\frac{N(N-1)}{2} + L(G)} q^{C(G)} \ . \end{aligned} \tag{60}$$

Computing the prefactor on the r.h.s. of (60), we have

$$Z_{Potts} = e^{-\frac{N}{2}\beta} Y(q) \ , \tag{61}$$

for terms exponential in N. Y is the cluster generating function of the graph (48). The large N behaviour of the cluster probability $\Gamma(c)$ is therefore related

to the Potts free{energy,

$$f_{Potts}(q) = \lim_{N \to \infty} \frac{1}{N} \ln Z_{Potts} \; ; \tag{62}$$

through

$$\frac{}{2} \quad f_{Potts}(q) = \max_{0 \; c \; 1} (c \ln q + \; ! \; (c)) \; : \tag{63}$$

We are interested in nding the value $c(q)$ which maximizes the r.h.s. in (63); since

$$\left. \frac{d \; ! \; (c)}{dc} \right|_{c \; (q)} = \quad \ln q \tag{64}$$

it follows that ! takes its maximum value for $q = 1$. Dierentiating eq. (63) with respect to q, we have

$$\frac{df_{Potts}}{dq} = \frac{d}{dq}(c \ln q + \; ! \; (c)) = \frac{@}{@c}(c \ln q + \; ! \; (c)) \frac{@c}{@q} + \frac{c}{q} \; ; \tag{65}$$

which, in virtue of eq. (64) becomes:

$$c(q) = \quad q \frac{df_{Potts}}{dq}(q) \quad : \tag{66}$$

It is now clear that the typical fraction of clusters per site, $c(q = 1)$, can be obtained, at a given connectivity , by computing the Potts free-energy in the vicinity of $q = 1$. Since the Potts model is originally dened for integer values of q only, an analytic continuation to real values of q is necessary. We now explain how to perform this continuation.

3.2.4 Free-energy calculation.

As in the case of the Ising model of section II, a careful examination of the energy function (52) shows that the latter depends on the spin conguration only through the fractions $x(; f_{i}g)$ of variables $_{i}$ in the -th state (= $0; 1; \quad ; q \quad 1$) [20],

$$x(; f_{i}g) = \frac{1}{N} \sum_{i=1}^{N} \quad (_{i}; \;); \quad (= 0; 1; \ldots; q \quad 1) \quad : \tag{67}$$

Of course, $\sum_{\alpha} x(\alpha;\{\sigma_i\}) = 1$. Note that in the Ising case ($q = 2$) the two fractions $x(0)$ and $x(1)$ can be parametrized by a unique parameter e.g. the magnetization $m = (x(1) - x(0))/2$.

Using these fractions, the energy (52) may be rewritten as

$$E[\{\sigma_i\}] = -\frac{N}{2}\sum_{\alpha=0}^{q-1}[x(\alpha;\{\sigma_i\})]^2 + \frac{N}{2} \quad . \tag{68}$$

Note that the last term on the r.h.s. of (68) can be neglected with respect to the first term whose order of magnitude is $O(N^2)$.

The partition function (53) at inverse temperature $\beta = \beta/N$ now becomes

$$\begin{aligned} Z_{Potts} &= \sum_{\{\sigma_i=0,1,\ldots,q-1\}} \exp\left(-\frac{\beta}{2}N\sum_{\alpha=0}^{q-1}[x(\alpha;\{\sigma_i\})]^2\right) \\ &= \sum_{\{x(\alpha)=0,1/N,\ldots,1\}}^{(R)} \exp\left(-\frac{\beta}{2}N\sum_{\alpha=0}^{q-1}[x(\alpha)]^2\right) \frac{N!}{\prod_{\alpha=0}^{q-1}[Nx(\alpha)]!} \\ &= \int_0^{1(R)} \prod_{\alpha=1}^{q-1} dx(\alpha)\, \exp(-\beta N f[\{x(\alpha)\}]) \end{aligned} \tag{69}$$

to the leading order in N. The subscript (R) indicates that the sum or the integral must be restricted to the normalized subspace $\sum_{\alpha=0}^{q-1} x(\alpha) = 1$. The "free-energy" density functional f appearing in (69) is

$$f[\{x(\alpha)\}] = \sum_{\alpha=0}^{q-1}\left[\frac{\beta}{2}[x(\alpha)]^2 + x(\alpha)\ln x(\alpha)\right] \quad . \tag{70}$$

In the limit of large N, the integral in (69) may be evaluated by the saddle-point method. The Potts free-energy (62) then reads

$$f_{Potts}(q) = \min_{\{x(\alpha)\}} f[\{x_\alpha\}] \tag{71}$$

and the problem becomes that of analyzing the minima of f. Given the initial formulation of the problem, each possible value of α among $0,\ldots,q-1$ plays the same role; indeed f is invariant under the permutation symmetry of the different q values. However, we should keep in mind that such a symmetry could be broken by the minimum (see section 2). We shall see that depending on the value of the connectivity c, the permutation symmetry may or may not be broken, leading to a phase transition in the problem which coincides with the birth a giant component in the associated random graph.

3.2.5 Symmetric saddle-point.

Consider first the symmetric extremum of f,

$$x^{sym}(\alpha) = \frac{1}{q}, \qquad \forall \alpha = 0,\dots,q-1. \tag{72}$$

We have

$$f^{sym}_{Potts}(q) = -\ln q - \frac{1}{2q}. \tag{73}$$

Taking the Legendre transform of this free-energy, see (63,66), we get for the logarithm of the cluster distribution density

$$\omega^{sym}(c) = -\frac{1}{2}(1-c)(1+\ln 2 - \ln[2(1-c)]). \tag{74}$$

$\omega^{sym}(c)$ is maximal and null at $c^{sym}(q) = 1 - \frac{1}{2}$, a result that cannot be true for connectivities larger than two and must break down somewhere below. Comparison with the rigorous derivation in random graph theory indicates that the symmetric result is exact as long as $c_c = 1$ and is false above the percolation threshold c_c. The failure of the symmetric extremum in the presence of a giant component proves the onset of symmetry breaking.

To understand the mechanism responsible for the symmetry breaking, we look for the local stability of the symmetric saddle-point (72) and compute the eigenvalues of the Hessian matrix

$$M_{\alpha,\beta} = \frac{\partial^2}{\partial x(\alpha)\partial x(\beta)} f[\{x(\gamma)\}]_{sym;(R)}, \tag{75}$$

restricted to the normalized subspace. The simple algebraic structure of M allows an exact computation of its $q-1$ eigenvalues for a generic integer q. We find a non degenerate eigenvalue $\lambda_0 = q(q-2)$ and another eigenvalue $\lambda_1 = -q$ with multiplicty $q-2$. The analytic continuation of the eigenvalues to real $q \to 1$ lead to the single value $\lambda = -1$ which changes sign at the percolation threshold c_c. Therefore, the symmetric saddle-point is not a local minimum of f above c_c, showing that a more complicated saddle-point has to be found.

3.2.6 Symmetry broken saddle-point.

The simplest way to break the symmetry of the problem is to look for solutions in which one among the q values appears more frequently than the others.

Therefore we look for a saddle-point of the form

$$x(0) = \frac{1}{q}[1 + (1 - q)s]$$
$$x(\ell) = \frac{1}{q}[1 - s]; \qquad (\ell = 1, \ldots, q-1). \tag{76}$$

The symmetric case can be recovered in this enlarged subspace of solutions by setting $s = 0$. The free-energy of the Potts model is obtained by plugging the fractions (76) into (70). In the limit $q \to 1$ of interest,

$$f[\{x_\ell\}] = -\frac{\epsilon}{2} + (q-1) f_{Potts}(s, \epsilon) + O((q-1)^2) \tag{77}$$

with

$$f_{Potts}(s, \epsilon) = \frac{\epsilon}{2}(1 - \frac{1}{2}s^2) - 1 + s + (1-s)\ln(1-s) \tag{78}$$

Minimization of $f_{Potts}(s, \epsilon)$ with respect to the order parameter s shows that for $\epsilon \leq 1$ the symmetric solution $s = 0$ is recovered, whereas for $\epsilon > 1$ there exists a non vanishing optimal value $s(\epsilon)$ of s that is solution of the implicit equation

$$1 - s = \exp(-\epsilon s) \quad . \tag{79}$$

The stability analysis (which we will not give here) shows that the solution is stable for any value of ϵ. The interpretation of $s(\epsilon)$ is straightforward: s is the fraction of vertices belonging to the giant cluster. The average fraction of connected components $c(\epsilon)$ equals $-f_{Potts}(s(\epsilon), \epsilon)$, see (66), in perfect agreement with exact results by Erdos and Renyi.

3.3 Discussion.

Further results on the properties of random graphs can be extracted from the previous type of calculation. We shall examine two of them.

3.3.1 Scaling at the percolation point.

Given the interpretation of $s(\ )$ for any large but finite value of N, we may define the probability of existence of a cluster containing N_s sites as follows

$$P(s;N) \sim \frac{\exp(N f(s;\ ))}{\exp(N f(s_{\ };\ ))} \tag{80}$$

In the infinite size limit this leads to the expected result

$$\lim_{N \to \infty} P(s;N) = \ (s - s(\ )) \tag{81}$$

In order to describe in detail how sharp (in N) the transition is at $= 1$, we need to consider corrections to the saddle point solutions by making an expansion of the free-energy $f_{Potts}(s;\ = 1)$ in the order parameter s. At threshold, we have $s(1) = 0$ and $f_{Potts}(s;1) = -s^3/6 + O(s^4)$ and therefore

$$P(s;N) \sim \exp(-N s^3/6) \tag{82}$$

In order to keep the probability finite at the critical point the only possible scaling for s is $s = O(N^{-1/3})$ which leads to a size of the giant component at criticality $N - N^{1/3} = N^{2/3}$, in agreement with the Erdos-Renyi results.

3.3.2 Large deviations.

The knowledge of the Potts free-energy for any value of q allows one to compute its Legendre transform, $-(c)$. The computation does not show any difficulty and we do not reproduce the results here [21]. Phase transitions are also found to take place for rare events (graphs that do not dominate the cluster probability distribution). Notice that we consider here random graphs obtained by deleting edges from K_N with a fixed probability. Large deviations results indeed depend strongly on the process of generating graphs.

As a typical example of what can be found using statistical mechanics, let us mention this simple result

$$-(c = 1) = \ -\frac{}{2} \ ; \tag{83}$$

for all connectivities . The above identity means that the probability that a random graph has $N - o(N)$ connected components decreases as $\exp(-N/2)$ when N gets large. This result may be easily understood. Consider for instance graphs with N edges made of a complete graph on $\sqrt{2N}$ vertices plus N

$\sqrt{2N}$ isolated vertices. The fraction of connected components in this graph is $c = 1 - O(1/\sqrt{N}) \to 1$. The number of such graphs is simply the number of choices of $\sqrt{2N}$ vertices among N ones. Taking into account the edge deletion probability $1 - p = 1 - e^{-c/N}$, one easily recovers (83).

3.3.3 Conclusion.

The random graph problem is a nice starting point to test ideas and techniques from statistical mechanics. First, rigorous results are known and can be confronted to the outputs of the calculation. Secondly, analytical calculations are not too difficult and can be exploited easily.

As its main focus, this section aimed at exemplifying the strategy used in more complicated, e.g. K-Satisfiability, problems. The procedure of analytic continuation, which is at the root of the replica approach, appears nicely in the computation of the Potts free-energy and is shown to give exact results (though in a non rigorous way). The power of the approach is impressive. Many quantities can be computed and rather subtle effects such as large deviations are easily obtained in a unique framework.

At the same time, the main weakness of the statistical mechanics approach is also visible. Most interesting effects are obtained when an underlying symmetry is broken. But the structure of the broken saddle-point subspace is far from obvious, in contrast to the Ising case of the previous section. There is at first sight some kind of arbitrariness in the search of a saddle-point of the form of (76). In the absence of a well-established and rigorous procedure, the symmetry breaking schemes to be used must satisfy at least basic self-consistency checks (plausibility of results, local stability, ...). In addition, theoretical physicists have developed various schemes that are known to be efficient for various classes of problems but (fail in other cases). A kind of standard lore, of precious help to solve new problems, exists and is still waiting for firm mathematical foundations.

4 Random K-satisfiability problem

In what follows we shall describe the main steps of the replica approach to the statistical mechanics analysis of the Satisfiability problem. The interested reader may find additional details concerning the calculations in several published papers [22–28] and in the references therein.

The satisfaction of constrained Boolean formulae is a key issue in complexity theory. Many computational problems are known to be NP-complete [15,29]

through a polynomial mapping onto the K-Satisfiability (SAT) problem, which in turn was the first problem shown to be NP-complete by Cook in 1971 [30].

Recently [31], there has been much interest in a random version of the K-SAT problem defined as follows. Consider N Boolean variables x_i, $i = 1;:::;N$. Call a clause C the logical OR of K randomly chosen variables, each of them being negated or left unchanged with equal probabilities. Then repeat this process by drawing independently M random clauses C_ℓ, $\ell = 1;:::;M$. The logical AND of all these clauses is a "formula", referred to as F. It is said to be satisfiable if there exists a logical assignment of the xs evaluating F to true, and unsatisfiable otherwise.

Numerical experiments have concentrated on the study of the probability $P_N(\ ;K)$ that a randomly chosen F having $M = \ N$ clauses be satisfiable. For large sizes, a remarkable behaviour arises: P_N seems to reach unity for $< \ _c(K)$ and vanishes for $> \ _c(K)$ when $N \ ! \ 1$ [32,31]. Such an abrupt threshold behaviour, separating a SAT phase from an UNSAT one, has indeed been rigorously confirmed for 2-SAT, which is in P, with $\ _c(2) = 1$ [33,34]. For larger $K \quad 3$, K-SAT is NP-complete and much less is known. The existence of a sharp transition has not been rigorously proved but estimates of the thresholds have been found : $\ _c(3) \ ' \ 4:3$ [35]. Moreover, some rigorous lower and upper bounds to $\ _c(3)$ (if it exists), $\ _{lb:} = 3:14$ and $\ _{ub:} = 4:51$ respectively have been established (see the review articles dedicated to upper and lower bounds contained in this TCS special issue).

The interest in random K-SAT arises partly from the following fact: it has been observed numerically that hard random instances are generated when the problems are critically constrained, i.e., close to the SAT/UNSAT phase boundary [32,31]. The study of such hard instances represents a theoretical challenge towards an understanding of complexity and the analysis of exact algorithms. Moreover, hard random instances are also a test-bed for the optimization of heuristic (incomplete) search procedures, which are widely used in practice.

Statistical mechanics provides new intuition on the nature of the solutions of random K-SAT (or MAX-K-SAT) through the introduction of an order parameter which describes the geometrical structure of the space of solutions. In addition, it gives also a global picture of the dynamical operation of search procedures and the computational complexity of K-SAT solving.

4.1 K-SAT energy and the partition function.

To apply the statistical physics approach exemplified on the random graph problem, one has to identify the energy function corresponding to the K-SAT

problem.

The logical values of an x_i can be represented by a binary variable S_i, called a spin, through the one-to-one mapping $S_i = -1$ (respectively $+1$) if x_i is false (resp. true). The random clauses can then be encoded into an $M \times N$ matrix $C_{\alpha i}$ in the following way : $C_{\alpha i} = -1$ (respectively $+1$) if the clause C_α includes $\overline{x_i}$ (resp. x_i), $C_{\alpha i} = 0$ otherwise. It can be checked easily that $\sum_{i=1}^{N} C_{\alpha i} S_i$ equals the number of wrong literals in clause α. Consider now the cost-function $E[C;S]$ defined as the number of clauses that are not satisfied by the logical assignment corresponding to configuration S.

$$E[C;S] = \sum_{\alpha=1}^{M} \Theta\left(K - \sum_{i=1}^{N} C_{\alpha i} S_i \right) \quad ; \tag{84}$$

where $\Theta(j) = 1$ if $j = 0$, zero otherwise, denotes the Kronecker function. The minimum (or ground state -GS) $E[C]$ of $E[C;S]$, is the lowest number of violated clauses that can be achieved by the best possible logical assignment [23]. $E[C]$ is a random variable that becomes highly concentrated around its average value $E_{GS} \equiv \overline{E[C]}$ in the large size limit [36]. The latter is accessible through the knowledge of the averaged logarithm of the generating function

$$Z[C] = \sum_{S} \exp(-E[C;S]/T) \tag{85}$$

since

$$E_{GS} = -T \lim \overline{\log Z[C]} + O(T^2) \quad ; \tag{86}$$

when the auxiliary parameter T is sent to zero. Being the minimal number of violated clauses, E_{GS} equals zero in the sat region and is strictly positive in the unsat phase. The knowledge of E_{GS} as a function of α therefore determines the threshold ratio $\alpha_c(K)$.

4.2 The average over the disorder.

The calculation of the average value of the logarithm of the partition function in (86) is an awkward one. To circumvent this difficulty, we compute the n^{th} moment of Z for integer-valued n and perform an analytic continuation to real n to exploit the identity $\overline{Z[C]^n} = 1 + n \overline{\log Z[C]} + O(n^2)$. The n^{th} moment of Z is obtained by replicating n times the sum over the spin configurations S

and averaging over the clause distribution [23]

$$\overline{Z[C]^n} = \sum_{S^1,S^2,\ldots,S^n} \overline{\exp\left(\sum_{a=1}^{n} E[C;S^a]/T\right)} \quad ; \tag{87}$$

which in turn may be viewed as a generating function in the variable $e^{-1/T}$.

In order to compute the expectation values that appear in eq.(87), one notices that each individual term

$$z[\{S^a\}] = \overline{\exp\left(-\frac{1}{T}\sum_{a=1}^{n} E[C;S^a]\right)} \tag{88}$$

factorises over the sets of different clauses due to the absence of any correlation in their probability distribution. It follows

$$z[\{S^a\}] = \left(\psi_K[\{S^a\}]\right)^M \quad ; \tag{89}$$

where each factor is defined by

$$\psi_K[\{S^a\}] = \overline{\exp\left[-\frac{1}{T}\sum_{a=1}^{n}\sum_{i=1}^{N} C_i S_i^a + K\right]} \quad ; \tag{90}$$

with the bar denoting the uniform average over the set of $2^K \binom{N}{K}$ vectors of N components $C_i = 0,\ 1$ and of squared norm equal to K.

Resorting to the identity,

$$\sum_{i=1}^{N} C_i S_i^a + K = \prod_{i=C_i\neq 0} (S_i^a + C_i) \quad ; \tag{91}$$

one may carry out the average over in disorder in eq.(90) to obtain

$$\psi_K[\{S^a\}] = \frac{1}{2^K} \sum_{C_1,\ldots,C_K = \pm 1} \frac{1}{N^K} \sum_{i_1,\ldots,i_K=1}^{N} \exp\left(-\frac{1}{T}\sum_{a=1}^{n}\prod_{\ell=1}^{K} h S_{i_\ell}^a + C_\ell i\right) \tag{92}$$

up to negligible $O(1/N)$ contributions.

The averaged term in the r.h.s. of (87) depends on the $n \cdot N$ spin values only through the 2^n occupation fractions $x(\tilde{\sigma})$ labeled by the vectors $\tilde{\sigma}$ with

n binary components; $x(\underline{\sigma})$ equals the number (divided by N) of labels i such that $S_i^a = \sigma^a, \forall a = 1,\ldots,n$. It follows that $\Gamma_K[\{S^a\}] = \Gamma_K[x]$ where

$$\Gamma_K[x] = \frac{1}{2^K} \sum_{C_1,\ldots,C_K} \sum_{\underline{\tau}_1,\ldots,\underline{\tau}_K} x(C_1\,\underline{\tau}_1)\ldots x(C_K\,\underline{\tau}_K) \qquad (93)$$
$$\exp\left(\frac{1}{T}\sum_{a=1}^{n}\prod_{\ell=1}^{K}[\tau_\ell^a - 1]\right) \;.$$

To leading order in N (e.g., by resorting to a saddle point integration), the final expression of the n^{th} moment of Z can be written as $\overline{Z[C]^n} \simeq \exp(-N f_{opt}/T)$ where f_{opt} is the optimum (in fact the minimum for integer n) over all possible xs of the functional [23]

$$f[x] = e[x] + \frac{1}{T}\sum_{\underline{\tau}} x(\underline{\tau})\log x(\underline{\tau}) \quad ; \qquad (94)$$

with

$$e[x] = -\ln\left[\sum_{\underline{\tau}_1,\ldots,\underline{\tau}_K} x(\underline{\tau}_1)\ldots x(\underline{\tau}_K)\exp\left(\frac{1}{T}\sum_{a=1}^{n}\prod_{\ell=1}^{K}[\tau_\ell^a - 1]\right)\right] \;. \qquad (95)$$

Note the similarities between equations (94) and (70). While in the random graph or Potts model case τ took on q values, the K-SAT model requires the introduction of 2^n vectors $\underline{\tau}$. In both cases, an analytic continuation of the free-energy to non integer values of q or n has to be performed. Finally, note that the optimum of f fulfills $x(\underline{\tau}) = x(-\underline{\tau})$ due to the uniform distribution of the disorder C.

4.3 Order parameter and replica-symmetric saddle-point equations.

The optimization conditions over $f[x]$ provide 2^n coupled equations for the xs. Notice that f is a symmetric functional, invariant under any permutation of the replicas a, as is evident from equation (87). An extremum may thus be sought in the so-called replica symmetric (RS) subspace of dimension $n+1$ where $x(\underline{\tau})$ is left unchanged under the action of the symmetric group. In the limit of interest, $T \to 0$, and within the RS subspace, the occupation fractions may be conveniently expressed as the moments of a probability density $P(m)$

over the range $-1 \le m \le 1$ [23],

$$x(\alpha_1, \alpha_2, \ldots, \alpha_n) = \int_{-1}^{+1} dm\, P(m) \prod_{a=1}^{n} \frac{1+m}{2}^{a} \quad . \tag{96}$$

$P(m)$ is not uniquely defined by (96) for integer values of n but acquires some precise meaning in the $n \to 0$ limit. It is the probability density of the expectation values of the spin variables over the set of ground states. Consider a formula F and all the spin configurations $S^{(j)}$, $j = 1, \ldots, Q$ realizing the minimum $E[C]$ of the cost-function $E[C, S]$, that is the solutions of the MAX-SAT problem defined by F. Then define the average magnetizations of the spins

$$m_i = \frac{1}{Q} \sum_{j=1}^{Q} S_i^{(j)} \quad , \tag{97}$$

over the set of optimal configurations. Call $H(C, m)$ the histogram of the m_is and $H(m)$ its quenched average, i.e., the average of $H(C, m)$ over the random choices of the formulae F. $H(m)$ is a probability density over the interval $-1 \le m \le 1$ giving information on the distribution of the variables induced by the constraint of satisfying all the clauses. In the absence of clauses, all assignments are solutions and all magnetizations vanish: $H(m) = \delta(m)$ and variables are not constrained. Oppositely, variables that always take the same value in all solutions, if any, have magnetizations equal to $+1$ (or -1): such variables are totally constrained by the clauses.

As discussed in ref. [23], if the RS solution is the global optimum of (94) then $H(m)$ equals the above mentioned $P(m)$ in the limit of large sizes $N \to \infty$. Therefore, the order parameter arising in the replica calculation reflects the "microscopic" structure of the solutions of the K-SAT problem.

At this stage of the analysis it is possible to perform the analytic continuation $n \to 0$ since all the functionals have been expressed in term of the generic number of replicas n. Such a process leads to a self-consistent functional equation for the order parameter $P(m)$, which reads

$$P(m) = \frac{1}{\sqrt{1-m^2}} \int_{-\infty}^{+\infty} du\, \cos\left(\frac{u}{2} \ln \frac{1+m}{1-m}\right) \exp\left[-K + K \sum_{\ell=1}^{K-1} \binom{K-1}{\ell} \int_{-1}^{+1} dm' P(m') \cos\left(\frac{u}{2} \ln A_{(K-1)}\right) \right] \tag{98}$$

with

$$A_{(K\ 1)}\quad A_{(K\ 1)}(fm_{\cdot},g;\ )=1+\ (e\qquad 1)\prod_{\cdot=1}^{K\ 1}\frac{1+m_{\cdot}}{2}\qquad ; \tag{99}$$

and $1=T$. The corresponding replica symmetric free{energy density reads

$$\begin{aligned} f_{opt}(\ ;T)= & \ln 2+\quad (1\quad K)\int_{1}^{1}\prod_{\cdot=1}^{K}dm_{\cdot}P(m_{\cdot})\ln A_{(K)} \\ & +\frac{K}{2}\int_{1}^{1}\prod_{\cdot=1}^{K\ 1}dm_{\cdot}P(m_{\cdot})\ln A_{(K\ 1)} \\ & \frac{1}{2}\int_{1}^{1}dm\,P(m)\ln(1\quad m^{2})\qquad : \end{aligned} \tag{100}$$

It can be checked that equation (98) is recovered when optimizing the free-energy functional (100) over all (even) probability densities $P(m)$ on the interval [-1,1].

4.4 The simple case of $K=1$.

Before entering in the analysis of the saddle-point equations for general K, it is worth considering the simple $K = 1$ case which can be solved either by a direct combinatorial method or within the statistical mechanics approach. Though random 1-SAT does not present any critical behaviour (for nite), its study allows an intuitive understanding of the meaning and correctness of the statistical mechanics approach.

For $K = 1$, a sample of M clauses can be de ned completely by giving directly the numbers t_i and f_i of clauses imposing that a certain Boolean variable S_i must be true or false respectively. The partition function corresponding to a given sample reads

$$Z[ft;fg]=\prod_{i=1}^{N}(e^{\ t_i}+e^{\ f_i})\ ; \tag{101}$$

and the average over the disorder gives

$$\frac{1}{N}\overline{\ln Z[ft;fg]}=\frac{1}{N}\sum_{ft_i;f_ig}\frac{M\ !}{\Theta_{i=1}^{N}(t_i!f_i!)}\ln Z[ft;fg]$$

$$= \ln 2 \quad \frac{}{2} + \sum_{l=1}^{\infty} e \quad I_l( ) \ln \quad \cosh \quad \frac{l}{2} \quad ; \qquad (102)$$

where I_l denotes the l^{th} modified Bessel function. The zero temperature limit gives the ground state energy density

$$e_{GS}( ) = \frac{}{2}[1 \quad e \quad I_0( ) \quad e \quad I_1( )] \qquad (103)$$

and the ground state entropy density

$$s_{GS}( ) = e \quad I_0( ) \ln 2 \quad : \qquad (104)$$

For any > 0, the ground{state energy density is positive and therefore the overall Boolean formula is false with probability one. Also, the entropy density is finite, i.e., the number of minima of the energy for any is exponentially large. Such a result can be understood by noticing that there exist a fraction of unconstrained variables $e \quad I_0( )$ which are subject to equal but opposite constraints $t_i = f_i$.

The above results are recovered in the statistical mechanics framework, thereby showing that the RS Ansatz is exact for all and when $K = 1$.

The solution of the saddle-point equation (98) can be found for any temperature T leading to the expression

$$P(m) = \sum_{`=1}^{\infty} e \quad I_`( ) \quad m \quad \tanh \quad \frac{`}{2} \quad : \qquad (105)$$

In the limit of interest $! \ 1$, this formula reads

$$P(m) = e \quad I_0( ) \quad (m) + \frac{1}{2}(1 \quad e \quad I_0( ))((m \quad 1) + \quad (m + 1)) \quad : (106)$$

As shown in figure 3, the fraction of unconstrained variables is simply associated with the unfrozen spins and thus gives the weight of the {function at $m = 0$. On the contrary, the non-zero value of the fraction of violated clauses, proportional to the ground-state energy density, is due to the presence of completely frozen (over constrained) spins of magnetizations $m = 1$. Such a feature remains valid for any K.

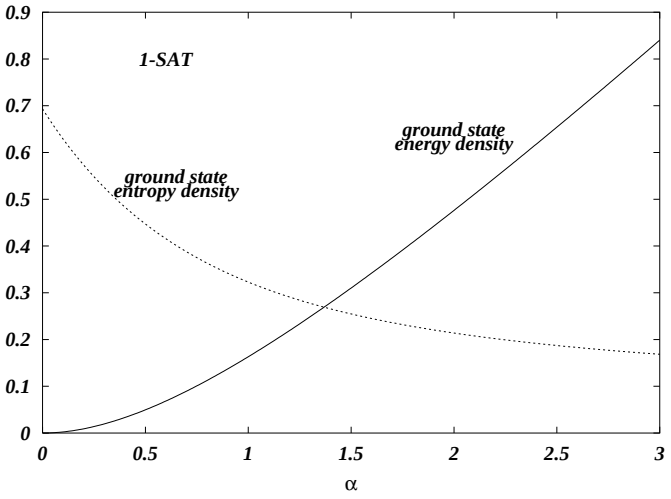

Fig. 3. Energy density (bold line) and entropy density (thin line) versus α in a random 1-SAT formula, in the limit $N \to \infty$.

4.5 Sat phase: structure of the space of solutions.

We start by considering the sat phase. An interesting quantity to look at is the typical number of solutions of the random K-SAT problem; this quantity can be obtained from the ground state entropy density $s_{GS}(\alpha)$ given by eq.(100) in the $N \to \infty$ limit.

In the absence of any clauses, all assignments are solutions: $s_{GS}(\alpha = 0) = \ln 2$. We have computed the Taylor expansion of $s_{GS}(\alpha)$ in the vicinity of $\alpha = 0$, up to the seventh order in α. Results are shown in Figure 4. It is found that $s_{GS}(\alpha_c = 1) = .38$ and $s_{GS}(\alpha = 4.2) = .1$ for 2-SAT and 3-SAT respectively: just below threshold, solutions are exponentially numerous. This result is confirmed by rigorous work [37].

More involved calculations, including replica symmetry breaking (RSB) effects [28], have shown that the value of the entropy is insensitive to RSB in the sat phase. Therefore the RS calculation provides a quite precise estimate of the entropy (believed to be exact at low α ratios, see Talagrand's paper in this volume for a discussion).

Recent analytical calculations for 3-SAT [28] (also confirmed by numerical investigations) indicate that the RS theory breaks down at a definite ratio α_{RSB} below α_c, where the solutions start to be organized into distinct clusters. The meaning of this statement is as follows. Think of the space of spins configurations as the N-dimensional hypercube. Optimal assignments are a subset of

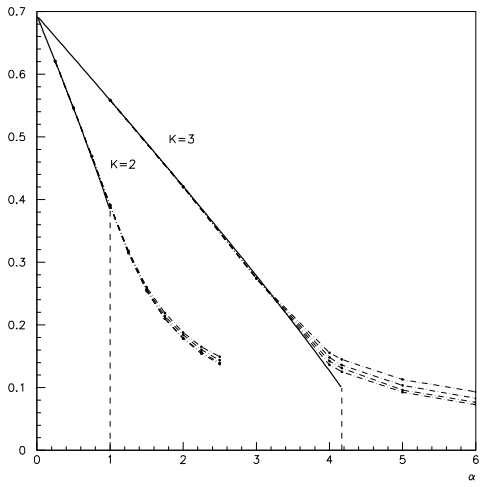

Fig. 4. RS estimate for the entropy density in random 2-SAT and 3-SAT below their thresholds. RSB corrections due to clustering are absent in 2-SAT and very small (within few a percent) in 3-SAT. The dots represent the results of exact enumerations in small systems (N ranging from 20 to 30, see ref. [22])

the set of 2^N vertices on the hypercube. Replica symmetry amounts to assuming that any pair of vertices are a.s. separated by the same Hamming distance d, defined as the fraction of distinct spins in the corresponding configurations. In other words, solutions are gathered in a single cluster, of diameter dN. RSB variational calculations [28] show that this simplifying assumption is not generally true in the whole sat phase and that another scenario may take place close to threshold:

> Below $_{RSB}$ the space of solutions is replica symmetric. There exist one cluster of solutions characterized by a single probability distribution of local magnetizations. The Hamming distance d is a decreasing function of , starting at $d(0) = 1/2$.
>
> At $_{RSB} \simeq 4.0$, the space of solutions breaks into a large number (polynomial in N) of different clusters. Each cluster contains an exponential number of solutions. The typical Hamming distance d_0 between solutions belonging to different clusters is close to 0.3 and remains nearly constant (it is slightly decreasing) up to $_c$, indicating that the centers of these clusters do not move on the hypercube when more and more clauses are added. Within each cluster, solutions tend to become more and more similar, with a rapidly decreasing intra-cluster Hamming distance d_1.

Figure 5 provides a qualitative representation of the clustering process. The fact that the Hamming distance can take two values at most is a direct consequence of the RSB Ansatz. In reality, the distance distribution could be more complicated. The key point is that statistical mechanics calculations strongly support the idea that the space of solutions has a highly organized structure, even in the sat phase.

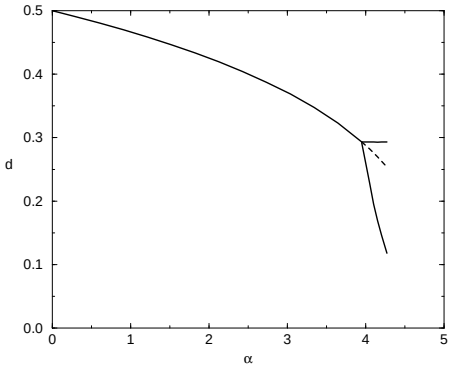

Fig. 5. Variational RSB estimate for the clustering of solutions below α_c for 3-SAT. d is the typical Hamming distance between solutions. The splitting of the curves at $\alpha \simeq 4$ corresponds to clustering. There appear two characteristic distances, one within each cluster and one between solutions belonging to different clusters.

Recently, the exact solution of the balanced version of random K-SAT [38] has provided a concrete example in which the appearance of clustering before the sat/unsat transition can be studied both analytically and numerically. Note that this phenomenon is strongly reminiscent of what happens in some formal multi-layer neural networks models [5].

4.6 Unsat phase: the backbone and the order of the phase transition.

In the unsat phase, it is expected that $O(N)$ variables become totally constrained, i.e. take on the same value in all the ground states. Such a hypothesis, which of course needs to be verified a posteriori, corresponds to a structural change in the probability distribution $P(m)$ which develops Dirac peaks at $m = \pm 1$.

In the limit of interest ($T \to 0$), to describe the accumulation of the magnetization on the borders of its domain ($m \in [-1;1]$), we introduce the rescaled variable z, implicitly defined by the relation $m = \tanh(z/T)$, see equation (106). Calling $R(z)$ the probability density of the zs, the saddle-point equations read

$$R(z) = \int_{-\infty}^{+\infty} \frac{du}{2} \cos(uz) \exp\left[-\frac{K}{2^{K-1}} + \alpha K \int_0^{+\infty} \prod_{\ell=1}^{K-1} dz_\ell R(z_\ell) \cos(u \min(1, z_1, \ldots, z_{K-1})) \right]^3 . \tag{107}$$

The corresponding ground state energy density reads, see (100),

$$e_{GS}(\alpha) = (1-K)\int_0^1 dz\, R(z)\, \min(1, z_1, \ldots, z_K)\prod_{\ell=1}^{K}$$
$$+ \frac{K}{2}\int_0^1 dz\, R(z)\, \min(1, z_1, \ldots, z_{K-1})\prod_{\ell=1}^{K-1} \int_0^1 dz R(z) z \quad : \qquad (108)$$

It is easy to see that the saddle–point equation (107) is in fact a self–consistent identity for $R(z)$ in the range $z \in [0,1]$ only. Outside this interval, equation (107) is merely a definition of the functional order parameter R.

As discussed in detail in ref. [23], equations (107) admit an infinite sequence of more and more structured exact solutions of the form

$$R(z) = \sum_{l=-\infty}^{+\infty} r_{\ell q + \ell}\, z^{\ell q + \ell}, \qquad (109)$$

having exactly q peaks in the interval $[0,1[$, whose centers are $z_{\ell q+\ell} = \frac{\ell}{q}$, $\ell = 0, \ldots, q-1$. The corresponding energy density reads, from (109) and (108),

$$e_{GS} = \frac{(1-K)}{q}\left[\frac{1-r_0}{2}\right]^K + \sum_{j=1}^{q-1}\left(\frac{1-r_0}{2}\sum_{l=1}^{j} r_l\right)^K$$
$$+ \frac{K}{2q}\left[\frac{1-r_0}{2}\right]^{K-1} + \sum_{j=1}^{q-1}\left(\frac{1-r_0}{2}\sum_{l=1}^{j} r_l\right)^{K-1}$$
$$\sum_{j=1}^{q} \frac{j}{q}\left(\frac{r_0}{2} + \frac{r_j}{2} + \sum_{l=1}^{j-1} r_l\right) \quad : \qquad (110)$$

Though there might be continuous solutions to (107), it is hoped that the energy of ground state can be arbitrarily well approximated by the above large q solutions.

The location of the sat/unsat threshold can be obtained for any K by looking at the value of α beyond which the ground state energy becomes positive. For 2–SAT the exact result $\alpha_c(2) = 1$ is recovered whereas for $K > 2$ the RS energy becomes positive at a value of α (e.g., $\alpha_c(3) \simeq 4.6$ as shown in figure 6) which is slightly higher than the value estimated by numerical simulations.

4.6.1 A hint at replica symmetry breaking.

The RS theory provides an upper bound for the thresholds for any $K > 2$, whereas the exact values can be obtained only by adopting a more general

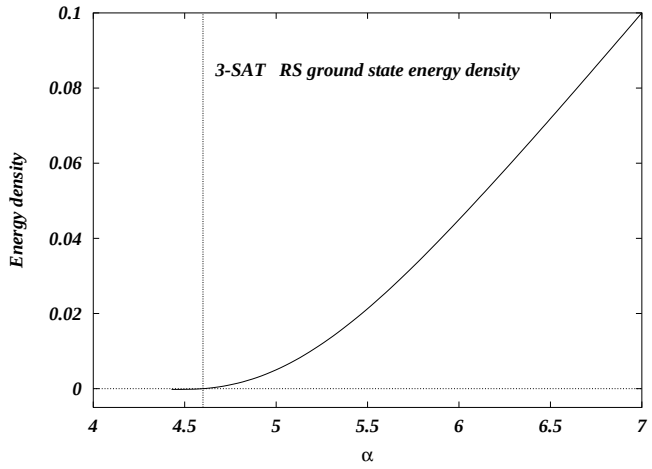

Fig. 6. RS estimate for the ground state energy density, *i.e.*, the number of violated clauses divided by N in random 3-SAT. The prediction is given as a function of α, for $q \to 1$ and in the limit $N \to \infty$. See ref. [23] for details.

functional form for the solution of the saddle-point equations which explicitly breaks the symmetry between replicas (see ref. [27] for a precise discussion). Such an issue is indeed a relevant, and largely open, problem in the statistical physics of random systems [39–46].

The general structure of the functional order parameter which describes solutions that break the permutational symmetry among replicas consists of a distribution of probability densities: each Boolean variable fluctuates from one cluster of solutions to another, leading to a site dependent probability density of local Boolean magnetizations. The distribution over all different variables then provides a probability distribution of probability distributions. The above scheme can in principle be iterated, leading to more and more refined levels of clustering of solutions. Such a scenario would correspond to the so-called continuous RSB scheme [1]. However the first step solution could suffice to capture the exact solution of random K-SAT, as happens in other similar random systems [1].

4.6.2 Abrupt vs. smooth phase transition.

Of particular interest are the fully constrained variables – the so called backbone component –, that is the x_is such that $m_i = \pm 1$. Within the RS Ansatz, the fraction of fully constrained variables $\Sigma(\alpha, K)$ can be directly computed

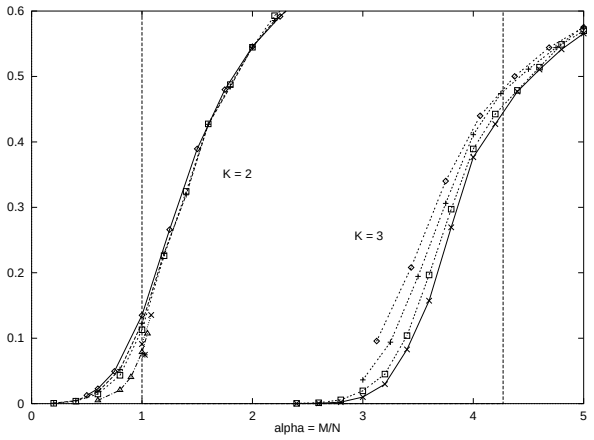

Fig. 7. Numerical estimates of the value of the backbone order parameter in 2-SAT and 3-SAT. The curves [25] are obtained by complete enumerations in small systems (up to $N = 500$ variables for 2-SAT and $N = 30$ for 3-SAT) averaged over many samples.

from the saddle-point equations. Clearly, $(;K)$ vanishes in the SAT region otherwise the addition of N new clauses to F would lead to a contradiction with a nite probability for any > 0. Two kinds of scenarii have been found when entering the unsat phase. For 2-SAT, $(;2)$ smoothly increases above the threshold $_c(2) = 1$. For 3-SAT (and more generally $K \quad 3$), $(;3)$ exhibits a discontinuous jump to a nite value $_c$ slightly above the threshold. A nite fraction of variables become suddenly over constrained when crossing the threshold! Numerical results on the growth of the backbone order parameter are given in gure 7.

4.6.3 The random 2+p-SAT model.

The sat/unsat transition is accompanied by a smooth (respectively abrupt) change in the backbone component and therefore in the structure of the solutions of the 2-SAT (resp. 3-SAT) problem. A better way to understand how such a change takes place is to consider a mixed model, which continuously interpolates between 2-SAT and 3-SAT. The so-called 2 + p-SAT model [25] includes a fraction p (resp. $1 \quad p$) of clauses of length two (resp. three). 2-SAT is recovered for $p = 0$ and 3-SAT when $p = 1$. The RS theory predicts that, at the sat/unsat transition, the appearance of the backbone component becomes abrupt when $p > p_0 ' 0:4$ (see gure 8). On the contrary, when $p < p_0$, the transition is smooth as in the 2-SAT case. Such a scenario is consistent with both rigorous results (see the paper by Achlioptas et al. in this volume) based on the probabilistic analysis of simple algorithm and with variational calculations [28] which include RSB e ects.

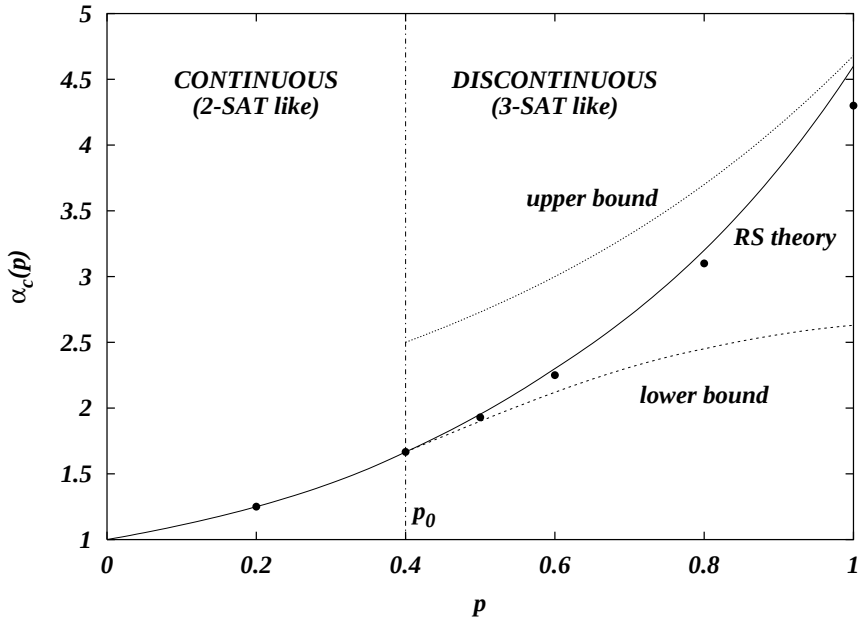

Fig. 8. $_c(p)$ versus p in random 2+p-SAT. Up to $p_0 \prime$:4 $_c(p) = 1=(1 \quad p)$, in agreement with rigorous results. For $p > p_0$ the transition becomes discontinuous in the backbone order parameter and the RS theory provides an upper bound for $_c(p)$ which is within a few percent of the results of numerical simulations (dots) [25,26].

An additional argument in favor of the above picture is given by the analysis of the nite-size e ects on $P_N(;K)$ and the emergence of some universality for $p < p_0$. (The de nition of P_N was given when we began discussing the properties of K-SAT.) A detailed account of these ndings may be found in [25,26]. For $p < p_0$ the size of the critical window where the transition takes place is observed to remain constant and close to the value expected for 2-SAT. The critical behaviour is the same as for the percolation transition in random graphs (see also ref. [47]). For $p > p_0$ the size of the window shrinks following some non-universal exponents toward its statistical lower bound [48] but numerical data do not allow for any precise estimate. The balanced version of 2+p-SAT can be studied exactly and both the phase diagram and the critical exponents turn out to behave very similarly to the ones of 2+p-SAT [49].

As we shall conclude in the next section, the knowledge of the phase diagram of the 2+p-SAT model is very precious to understand the computational complexity of 3-SAT solving.

4.7 Computational complexity and dynamics.

Numerical experiments have shown that the typical solving time of search algorithms displays an easy-hard-easy pattern as a function of with a peak of complexity close to the threshold. Since computational complexity is strongly affected by the presence of a phase transition, it is appropriate to ask whether the nature of this phase transition plays an important role too. The peak in the search cost seems indeed to scale polynomially with N (even using Davis-Putnam-like procedures) for the 2-SAT problem, where the transition is continuous, and exponentially with N in the 3-SAT case, for which the birth of the backbone is known to be discontinuous.

Precise numerical simulations [25,26] on the computational complexity of solving critical 2+p-SAT instances support the view that the crossover between polynomial and exponential scalings takes place at p_0, the very value of p separating continuous from discontinuous transitions. Though investigated 2+p-SAT instances are all critical and the problem itself is NP-complete for any $p > 0$, it is only when the phase transition is abrupt that hardness shows up (including the fastest known randomized search algorithms such as walksat [50]).

To understand why search algorithms require polynomial or exponential computational efforts, statistical studies of the solutions cannot be sufficient. A full dynamical study of how search procedures operate has to be carried out. Such studies had already been initiated by mathematicians in the easy region, where search tree are particularly simple and almost no backtracking occurs. Franco and Chao [51] have in particular analyzed the operation of DP algorithms with different kinds of heuristics and have shown that at small values of the typical complexity is linear in N.

Recently, the whole range of values of , including the hard phase, has been investigated, using dynamical statistical mechanics tools [52]. During the search process, the search tree built by DP grows with time and this growth process can be analyzed quantitatively. The key idea is that, under the action of DP, 3-SAT instances are turned into mixed 2+p-SAT instances (some clauses are simplified into clauses of length two, other are satisfied and eliminated). The parameters p and of the instance under consideration dynamically evolve under the action of DP. Their evolution can be traced back as a trajectory in the phase diagram of the 2+p-SAT model of figure 8. Depending on whether trajectories cross or not the sat/unsat boundary, easy or hard resolutions take place, and the location of crossings can be used to quantitatively predict the scaling of the resolution times [52].

5 The traveling salesman problem and the cavity method

In Section 3, we derived partition functions using statistical physics representations based on analytic continuations. Furthermore, we used the saddle point method on these partition functions and that allowed us to reproduce a number of exact results. Then we moved on in Section 4 and applied these methods to models with quenched disorder. However, because of the greater complexity of such models, we resorted to an additional tool of statistical physics: the replica method. Though this kind of approach is non-rigourous, it is believed that it provides new exact results for a number of different problems, in particular in optimization.

The replica method is not the only technical tool that physicists have developed in the past years. Another approach, called the cavity method, will be exposed in the present Section. The cavity approach gives, at the end of the computation, the same results as the replica approach. Yet the assumptions it relies upon turn out to be much more intuitive and its formalism is closer to a probabilistic theory formulation. Because of this, it can be used to prove some of the results derived from statistical mechanics; see [53,54] for recent progress in this direction. In the rest of this section, we show how this cavity method can be used to "solve" a case of the Traveling Salesman Problem (TSP).

The TSP is probably the world's most studied optimization problem. As usually formulated for a weighted graph, one considers all Hamiltonian cycles or "tours" (closed circuits visiting each vertex once and only once) and asks for the shortest one. The total length is given by the sum of the weights or "lengths" of the edges making up the tour. Since the Hamiltonian cycle problem is NP-complete, certainly the TSP is very difficult. However, in most cases considered, the graph is complete (there is an edge for each pair of vertices), so the difficulty lies in determining the shortest tour. Without further restrictions on the nature of the graph, the TSP is NP-hard [15]. One speaks of the asymmetric TSP when the edges on the graph are oriented, and of the symmetric TSP for the usual (unoriented) case. Both types are frequently used models in scheduling and routing problems, though the industrial applications tend to move away from the simple formulations considered in academia. The symmetric TSPs are further divided into "metric" and non-metric according to whether or not the triangle inequality for the edge lengths is satisfied. The so-called Euclidean TSP is probably the best known TSP and it is metric; the vertices are points (cities, or sites) in the plane, and the length of the edge connecting cities i and j is given by the Euclidean distance between i and j. Even within this restricted class of weighted graphs, the problem of finding the optimum tour remains NP-hard [15].

The TSP has been at the forefront of many past and recent developments

in complexity. For instance, pretty much all general purpose algorithmic approaches have been first presented and tested for the TSP. This tradition begins back in 1959 when Beardwood et al. [55] published tour lengths obtained from hand-drawn solutions! Later, the idea of optimization by local search was introduced in the context of the TSP by Lin [56], and simulated annealing [57,58] was first tested on TSPs also. The list continues with branch and bound [59], until today's state of the art algorithms based on cutting planes (branch and cut) [60], allowing one to solve problems with several thousand cities [61]. Many physicists have worked on these kinds of algorithmic questions from a practical point of view; in most cases their algorithms incorporate concepts such as temperature, mean field, and renormalization, that are standard in statistical physics, leading to some of the most effective methods of heuristic resolution [62]. It might be argued that these approaches can also be used to improve the heuristic decision rules at the heart of exact methods (for instance in branching strategies), but more work has to be done to determine whether this is indeed the case.

The widespread academic use of the TSP also extends to other issues in complexity. For instance, there has been much recent progress in approximability of the TSP [63]. However statistical physics has nothing to say about worst case behavior; instead it is relevant for describing the typical behavior arising in a statistical framework and tends to focus on self-averaging properties. Thus we are lead to consider TSPs where the edge lengths between vertices are chosen randomly according to a given probability distribution; the corresponding problem is called the stochastic TSP.

5.1 The stochastic TSP.

Statistical physicists as well as probabilists are not interested per-se in any particular instance of the TSP, rather they seek "generic" properties. This might be the typical computational complexity or the typical length of TSPs with N cities. It is then necessary to consider the stochastic TSP where each instance (the specification of the weighted graph) is taken at random from an ensemble of instances; this defines our "quenched disorder". Although one may be interested in many different ensembles, only a few have been the subject of thorough investigation. Perhaps the most studied stochastic TSP is the Euclidean one where the cities are randomly distributed in a given region of the plane [55]. This is a "random point" ensemble. Another ensemble that has been much considered consists in having the edge lengths all be independent random variables, corresponding to a "random distance" ensemble. (This terminology is misleading: the problem is not metric as the triangle inequality is generally not satisfied.) Random distance ensembles have been considered for both the symmetric [64] and the asymmetric [65] TSP.

For any of these ensembles, one can ask for the behavior of the optimum tour length, or consider properties of the tour itself. Most work by probabilists has focused on the first aspect (see [14] for a review), starting with the seminal work of Beardwood, Halton, and Hammersley [55] (hereafter referred to as BHH). Those authors considered the Euclidean ensemble where points are randomly (and independently) distributed in a bounded region of d-dimensional Euclidean space according to the probability density $\rho(X)$. Given a not too singular ρ, BHH proved that the optimum tour length, L_E, becomes peaked at large N, and that with probability one as $N \to \infty$

$$\frac{L_E}{N^{1-1/d}} \to \beta(d) \int \rho^{1-1/d}(X)\,dX \tag{111}$$

Here β is a constant, independent of ρ, depending only on the dimension of space. Some comments are in order. The first is that the relative fluctuations of the tour length about its mean tend to zero as $N \to \infty$, allowing one to meaningfully define a "typical" or generic tour length at large N. This fundamental property was initially proven using sub-additivity properties of the tour length, but from a more modern perspective, it follows from considering the passage from N to $N+1$ cities, corresponding to a martingale process (see [66]). The second point is that the N dependence of this typical length is such that the rescaled length $L_E/N^{1-1/d}$ converges in probability at large N. In the language of statistical physics, this quantity is just the ground state energy density of the system where one increases the volume linearly with N so that the mean density of points is N-independent. In general such an energy density is expected to be self-averaging, i.e., have a well defined large N limit, independent of the sequence of randomly generated samples (with probability one) as in Eq.(111). In some problems, the self-averaging property can be derived, while it will simply be assumed to hold when using the cavity approach.

Another comment is that given Eq.(111), the essence of the problem is the same for any $\rho(X)$; it is thus common practice to formulate the Euclidean TSP using N points laid down independently in a unit square (or hypercube if $d > 2$), the distribution being uniform.

There has been much work [14] on obtaining bounds and various estimates of the constants $\beta(d)$, but no exact results are known for $d > 1$. However, Rhee [67] has proved that

$$\frac{\beta(d)}{\sqrt{d}} \to \frac{1}{\sqrt{2e}} \quad \text{as} \quad d \to \infty \tag{112}$$

From the point of view of a statistical physics analysis, the difficulty in com-

puting (d) arises from the correlations among the point to point distances. Indeed, in the Euclidean ensemble, there are dN random variables associated with the random positions of the points, and N (N 1)=2 distances; these distances are thus highly redundant (and a fortiori correlated). When these distances are instead taken to be random and independent, the "cavity" method of statistical physics allows one to perform the calculation of the corresponding . Because of this, we will focus on that quenched disorder ensemble.

In the "independent edge-lengths ensemble" (as opposed to the independent points ensemble), it is the distances or edge lengths between points that are independent random variables. Let d_{ij} be the "distance" between points i and j (the problem is not metric, but we nevertheless follow the standard nomenclature and refer to d_{ij} as a distance). In the most studied case, d_{ij} is taken from a uniform distribution in $[0;1]$. From a physicist's perspective, it is natural to stay "close" to the Euclidean random point ensemble [64] by taking the distribution of d_{ij} to be that of two points randomly distributed in the unit square (hypercube when $d > 2$). The independent points and independent edge-lengths ensembles then have the same distribution for individual distances, and in the short distance and large N limit they also have the same distribution for pairs of distances. The main difference between the ensembles thus arises when considering three or more distances; in the Euclidean case, these have correlations as shown for instance by the triangle inequality.

The minimum tour length in these random edge-lengths models is expected to be self-averaging; the methods of Rhee and Talagrand [66] show that the distribution of TSP tour lengths becomes peaked at large N in this case, but currently there is no proof of the existence of a limit as in the Euclidean case. Nevertheless, this seems to be just a technical difficulty, and it is expected that the rescaled tour length indeed has a limit at large N ; we thus define (d) in analogy to the expression in Eq.(111) with the understanding that the s are different in the independent points and independent edge-lengths ensembles.

5.2 A statistical physics representation.

Following the notation of Section 2, we introduce the generating or partition function

$$Z(T) = \sum_{} \exp\left(-\frac{L( )}{T} \right) \tag{113}$$

where is a permutation of the vertices and determines uniquely a tour. In effect we have identified configurations with tours, that is with permutations; furthermore, the energy of a configuration is simply the length of its tour.

This construction amounts to introducing a probability $e^{-L(\sigma)/T}/Z$ for each tour. When $T = 1$, all tours are equally probable, while when $T \to 0$ only the shortest tour(s) survive. As before, T is the temperature, and the averages $\langle\cdot\rangle_T$ using this probability distribution are the thermal averages. From them one can extract most quantities of interest. For instance

$$< L >_T = \frac{1}{Z}\frac{dZ}{d(1/T)} \tag{114}$$

gives the mean tour length at temperature T. We then have for the TSP tour length: $L_{min} = \lim_{T\to 0} < L >_T$.

The generating function Z requires performing a sum over all permutations and is a difficult object to treat. To circumvent this difficulty, a different representation is used. We first introduce what is called a "spin" S, having now m-components, S_α, $\alpha = 1,\ldots,m$. These components are real and satisfy the constraint $\sum_\alpha (S_\alpha)^2 = m$. Such a spin can be identified with a point on a sphere in m-dimensional Euclidean space. Note that when $m = 1$, we recover the kinds of spins considered in the previous sections. Now for our statistical physics representation of the TSP, a spin S_i is associated to each vertex V_i of the graph, $i = 1,\ldots,N$. Define $R_{ij} = e^{-d_{ij}/T}$ and introduce a new generating function

$$G(T,m,\omega) = \int dS_1 dS_2 \ldots dS_N \exp\Big(\omega \sum_{i<j} R_{ij} S_i \cdot S_j\Big) \tag{115}$$

In this expression, $\cdot$ is the usual scalar product, and dS is associated with the uniform measure on the sphere in dimension m. We have normalized it so that $\int dS = 1$; then $\int dS S_\alpha S_\beta = \delta_{\alpha,\beta}$. The claim is now that the initial generating function Z is equivalent to using an analytic continuation of G in m:

$$\lim_{\substack{m\to 0\\ \omega\to\infty}} \frac{G}{\omega^N} = \sum_{\sigma} \exp\Big(-\frac{L(\sigma)}{T}\Big) \tag{116}$$

Comparing to the Potts model of Section 3, we see that m is analogous to the Potts parameter q: the partition function is defined for integer values of the parameter, and then has to be analytically continued to real values.

The derivation of equation (116) is based on showing the equality of both sides when performing a power series in $1/T$. First expand the exponential in the integral:

$$G = \int dS_1 dS_2 \ldots dS_N \Big[1 + \omega \sum_{i<j} R_{ij}(S_i \cdot S_j) + \frac{\omega^2}{2!} \Big]^3 \tag{117}$$

Now integrate term by term; each resulting contribution can be associated with a subgraph (but where edges can appear multiple times) whose weight is given in terms of its edges and its cycles. (Note that each vertex must be covered an even number of times because the integrand is even under $S_i \to -S_i$.) Each edge E_{ij} appearing in the subgraph contributes a multiplicative factor R_{ij} to its total weight. A further factor comes from the loops (cycles) of the subgraph. It is not difficult to see that each such loop leads to a factor m in the total weight because of the integration over the m-dimensional spins. Thus as $m \to 0$ only subgraphs having a single loop survive in G and then vertices cannot belong to more than two edges. Finally, when $\omega \to \infty$, the loops with the most vertices dominate, leading to tours. Thus if we first take $m \to 0$ and then $\omega \to \infty$, the expansion of G^{-1} reduces to a sum over all the tours of the graph. Furthermore, the weight of each tour is proportional to the product of the R_{ij} belonging to the tour, so that one recovers the total weight $m^N \omega^N \exp(-\omega L/T)$ where L is the tour length. In conclusion, Eq. (116) is justified to all orders in $1/T$, and thus for any finite N it holds as an identity.

Whether one uses Z or G^{-1} does not matter as they differ only by an irrelevant multiplicative factor (we assume m and $1/\omega$ infinitesimal). From G^{-1}, one can compute the optimum tour and not just the optimum tour length; indeed, at finite temperature, the probability that a tour contains the edge E_{ij} is given by the mean occupation of that edge. Defining $n_{ij} = 1$ if the edge is used by the tour and $n_{ij} = 0$ otherwise, the probability of occupation is

$$\langle n_{ij} \rangle_T = \omega R_{ij} \langle S_i S_j \rangle_T \tag{118}$$

where from now on $\langle \cdot \rangle_T$ means thermal average using either Z or G^{-1}; the one that is used should be clear from the observable considered. Now if we take in Eq. 118 the limit $T \to 0$, we find those edges that are occupied and thus the optimal tour (assuming it is unique). Note also that Eq. (118) has a simple justification: $\langle S_i S_j \rangle_T$ has a numerator whose expansion gives m^{N-1}/R_{ij} times the weighted sum over all tours containing the edge ij, while the denominator is m^N times the weighted sum over all tours. The identity Eq. (118) then follows immediately.

5.3 The cavity equations.

The partition function G^{-1} gives the "statistical physics" of the TSP for any given graph. Using this formalism to determine analytically the optimum tour in a general case seems an impossible task. Nevertheless, G is a good starting point for following the passage from N to $N+1$ vertices as in a martingale process, and the derivation of a recursion in N is the heart of

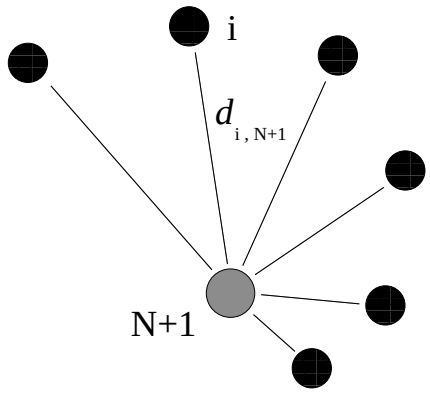

Fig. 9. $(N+1)$th spin and its ordered neighbors.

the cavity method. The term cavity comes from the fact that the system at $N+1$ is compared to the one at N by removing the $(N+1)$th spin, thereby creating a cavity. In figure 9, we have represented in counter-clockwise order the nearest, next-nearest, etc... neighbors of site $N+1$ which is at the center of the cavity. Because the total number of spins will be sometimes N and sometimes $N+1$, we indicate the number via a subscript on G. Thus for instance G_N 1 is to be used when considering quantities for the system with N spins. Now for every quantity associated with the system having $N+1$ spins, if we integrate explicitly over spin $N+1$, we are left with quantities defined in the system having only N spins. Consider for instance G_{N+1} 1 itself. When expanding the exponentials depending on S_{N+1}, we obtain: (i) terms linear in S_{N+1} that integrate to zero; (ii) terms quadratic in S_{N+1} that upon integration give products S_i S_j; (iii) higher powers in S_{N+1} that do not contribute as $m \to 0$. A simple calculation leads to the identity

$$\frac{G_{N+1} \quad 1}{G_N \quad 1} = \beta^2 \sum_{1 \le j<k \le N} R_{j,N+1} R_{k,N+1} \langle S_j \quad S_k \rangle_T^0 = \frac{Z_{N+1}}{Z_N} \tag{119}$$

where $\langle \cdot \rangle_T^0$ is a "cavity average", to be taken in the system having only the first N spins, spin $N+1$ being absent. Note that Z_N and Z_{N+1} are the partition functions of Eq. 113 when there are N and $N+1$ vertices; also, it is easy to see that one need not restrict the sum to $j \neq k$ because the term $j = k$ vanishes as $m \to 0$.

Straightforward calculations in this same spirit lead to relations between thermal expectation values using $N+1$ spins and those using N spins. For instance

$$\langle S_{N+1} \rangle_T \quad (G_{N+1} \quad 1) = \sum_{j=1}^{N} \beta R_{j,N+1} \langle S_j \rangle_T^0 \quad (G_N \quad 1) \tag{120}$$

Similarly, one has for the two-spin average:

$$\langle S_{N+1} S_i \rangle_T (G_{N+1} - 1) = \sum_{j \neq i} R_{j,N+1} \langle S_i S_j \rangle_T^0 (G_N - 1) \qquad (121)$$

More generally, the numerator in any observable depending on spin $N+1$ has a simple expression in terms of the numerators of observables in the absence of that spin. Furthermore, one can use Eq. 119 to eliminate all reference to G_N and G_{N+1} in these relations. The conclusion is that if we know how to compute the properties of systems with N spins, we can then deduce those of systems with $N+1$ spins; the cavity method is thus a recursion on N for all the properties of such a system.

5.4 The factorization approximation.

Unfortunately, these recursion equations cannot be solved, but let us approximate them by neglecting certain correlations. Clearly, S_{N+1} is strongly correlated with its nearest neighbors because the corresponding Rs are important. More generally, two spins whose joining edge length is short (are near neighbors) will be strongly correlated because short tours will often occupy that edge. Thus we must and will take into account the correlations between S_{N+1} and its near neighbors. However, we will neglect here the correlations among these neighbors themselves, so that in the absence of S_{N+1}, their joint probability distribution factorizes, so that in particular

$$\langle S_i S_j \rangle_T^0 = \langle S_i \rangle_T^0 \langle S_j \rangle_T^0 \qquad (122)$$

This property implies that replica symmetry is not broken, and this is indeed believed to be the case for the TSP. Factorization makes the cavity approach particularly tractable, as we shall soon see. (In systems where replica symmetry is broken, it is necessary to find ways to parametrize these correlations; this is quite complex and not well resolved, even within the statistical physics approach.)

A second point concerns the meaning of $\langle S_{N+1} \rangle_T$. G_{N+1} is rotationally symmetric; there is no preferred direction, so the thermal average of any spin vanishes. Note however that we have seen a similar situation before in the context of the Ising model (c.f. Section 2). Here as before, the interactions tend to align the spins. Thus, when the temperature is low enough, we expect to have a spontaneous magnetization when $N \to \infty$. To make this more explicit, we can introduce a small magnetic field, i.e., an interaction term of the type $-h \sum_i S_i$ for each spin; we then take the limit $N \to \infty$ and only after take $h \to 0$. This magnetic field breaks the rotational symmetry, and so the

system has a preferred direction, even after the field has been removed. By convention, we shall take this direction to be along the first axis.

Given these two remarks, we can use the exact equations (120) and (121) to obtain the cavity equations assuming factorization. Denoting by S^1 the component along the first axis of S, one has

$$\langle S^1_{N+1}\rangle_T = \frac{\sum_{j=1}^{N} R_{j,N+1}\langle S^1_j\rangle^0_T}{1 + \sum_{1\le j<k\le N} R_{j,N+1}R_{k,N+1}\langle S^1_j\rangle^0_T\langle S^1_k\rangle^0_T} \tag{123}$$

Similarly, one has for the two-spin average (see Eq. 118):

$$\langle n_{i,N+1}\rangle_T = R_{i,N+1}\langle S^1_i\rangle^0_T \frac{\sum_{j\neq i} R_{j,N+1}\langle S^1_j\rangle^0_T}{1 + \sum_{1\le j<k\le N} R_{j,N+1}R_{k,N+1}\langle S^1_j\rangle^0_T\langle S^1_k\rangle^0_T} \tag{124}$$

These are the standard cavity recurrence equations, first derived by Mezard and Parisi [68]. We also note that in this factorization approximation, one has $\langle S_{N+1}\, S_i\rangle_T = \langle S^1_{N+1}\rangle_T \langle S^1_i\rangle^0_T$

5.5 The $N \to \infty$ and $T \to 0$ limits.

The last step of the cavity method is to assume that the recurrence equations, when considered in the disorder ensemble, give rise to a stationary stochastic process when $N \to \infty$. Consider for instance the individual magnetizations $\langle S_i\rangle_T$; they are random variables because the d_{ij} themselves are. If we want them to have a limiting distribution at large N, (i.e., in physical terms, to have a thermodynamic limit), we have to rescale the d_{ij} by $N^{1/d}$ or equivalently set $T = \tilde{T}N^{-1/d}$ with $\tilde{T}$ fixed. (Note that in the case of the Euclidean TSP, the rescaling of lengths can be interpreted as taking the limit $N \to \infty$ while keeping the density of points fixed, that is by increasing the size of the volume linearly with N.) The important point is that the "environment" seen by the spins must have limiting statistical properties as $N \to \infty$, and this translates to having N-independent statistics for the distances of a spin to its near neighbors. Then it is assumed that the probability density of the $\langle S^1_i\rangle_T$ converges to a limiting distribution P_1 when $N \to \infty$. The cavity method is thus a kind of bootstrap approach where P_1 is assumed to exist and it is determined by its stationarity property under the cavity recurrence.

That such a stationary limit exists can be motivated by the large N behavior of the tour length in the stochastic TSP. In fact, it is expected that all quantities associated with any fixed number of edges will converge in the thermodynamic limit, so it should be possible to look at 2, 3, or k edge constructs. At present

though, because of the technical difficulty, only the single edge computations have been carried out. Fortunately, that is enough for getting the value of , and allows one to obtain the so called link-length distribution, i.e., the distribution of the edge lengths appearing in the optimal tours.

Equation (123) with the condition of stationarity of the stochastic process leads to a complicated implicit equation for P_1 . Fortunately, in the zero temperature limit (which is where we recover the usual stochastic TSP), the recurrence relations are much simpler. Following Krauth and Mezard [69], one defines $_i$ for any vertex i (i = 1; :::;N) via:

$$hS_i^{1}i^{0} = \frac{\exp(_i=T)}{!^{1=2}} \tag{125}$$

One also defines $_{N+1}$ analogously using $hS_{N+1}^{1}i$. Now re-order the indices of the first N vertices so that

$$N^{1=d}d_{1;N+1} \quad _1 \quad N^{1=d}d_{2;N+1} \quad _2 \quad ::: \quad N^{1=d}d_{N;N+1} \quad _N \tag{126}$$

Then the zero-temperature limit of Eq. 123 leads to

$$_{N+1} = d_{2;N+1}N^{1=d} \quad _2 \tag{127}$$

while Eq. 124 shows that the optimum tour uses the edges connecting N + 1 to vertices 1 and 2, i.e., $n_{1;N+1} = n_{2;N+1} = 1$, all others are equal to zero.

If we have a stationary stochastic process, Eq. (127) leads to a self-consistent equation for the probability density P of the s. We also see that the random variables $_i = N^{1=d}d_{i;N+1} \quad _i$ (i = 1; :::;N) play a fundamental role. By hypothesis, they are uncorrelated: the $d_{i;N+1}$ because we are dealing with the independent edge-lengths ensemble, and the $_i$ because we have explicitly neglected the correlations between the spins in the absence of S_{N+1}. Denote by ( ) the probability density of these random variables; ( ) is uniquely determined in terms of P , assuming the distribution of $d_{i;N+1}$ given. From here on, take for simplicity these edge lengths to be uniformily distributed in [0;1]. (This corresponds to the 1-dimensional case d = 1; we refer the reader to [69] for more general distributions.) The relation between and P then becomes

$$(x) = \frac{1}{N} \int_0^{N} P (l \quad)dl \tag{128}$$

Now a self-consistent equation for P is obtained by using the fact that $_{N+1}$

is the second smallest of the N different s:

$$P(\ ) = N(N\ \ 1)\ (\ )\Big(\int_{\ }^{1}(u)du\Big)\Big(\int_{\ }^{1}(u)du\Big)^{N\ \ 2} \tag{129}$$

In the large N limit, this integral non-linear implicit equation simplifies to

$$P(\ ) = \frac{dG(\ )}{d}\,G(\ )e^{\ G(\ )} \quad \text{where} \quad G(\ ) = \int_{0}^{1} uP(u\ \ )du \tag{130}$$

Plugging the expression for P into this last equation leads to

$$G(\ ) = \int_{\ }^{1}[1 + G(t)]\,e^{\ G(t)}\,dt \tag{131}$$

This cannot be solved analytically, but can easily be treated numerically, and one can obtain machine precision results for G and thus P without too much effort.

Assuming G and P have been computed, one can find in a similar way the distribution of $d_{1;N+1}$ and $d_{2;N+1}$. For instance, the distribution of the rescaled distance $N\,d_{1;N+1} = \tilde{1}_1$ is given by

$$P_1(\tilde{1}_1) = \int_{1}^{1} P(\tilde{1}_1\ \ )e^{\ G(\ )}d \tag{132}$$

This, along with the analogous distribution for $d_{2;N+1}$, gives the distribution of edge lengths in the optimum tour, and thus also the mean tour length, i.e., when $d = 1$, the value of . Krauth and Mezard [69] showed that this constant could be written in terms of G alone,

$$= \frac{1}{2}\int_{1}^{1} G(t)\,[1 + G(t)]\,e^{\ G(t)}\,dt \tag{133}$$

and they found $= 2{:}041\ldots$ (Note that when $d = 1$, as suggested by Eq. (111), the tour length becomes independent of N. This can be understood qualitatively by observing that each vertex can connect to one of its near neighbors that is at a distance $O(1{=}N)$.)

5.6 "Exact" solution in the independent edge-lengths ensemble.

As described, the cavity method involves an uncontrollable approximation associated with ignoring certain correlations. It is natural to ask whether those correlations might in fact be absent in certain ensembles. A simple case is when the graph considered is a Cayley tree with the root (corresponding to vertex $N+1$) removed. Then the different neighbors of S_{N+1} are uncoupled and have no correlations at all. Unfortunately, this type of graph will not do for the TSP as it has no Hamiltonian cycles, but it can do for other problems close to the TSP such as the minimum matching problem.

So let us consider instead the structure of independent edge-lengths graphs. Locally their properties ressemble those of Cayley trees, so that with some luck the previous reasoning can hold for these types of graphs as $N \to \infty$. Although the correlations that were neglected in the cavity calculation will always be present at finite N in the independent edge-lengths model, they have every reason to go to zero as $N \to \infty$. The justification is that the close neighbors of vertex $N+1$ are "infinitely" far from one-another when $N \to \infty$. In the language of tours (rather than spins), this means that the probability for the tour to have an edge connecting two of the finite order neighbors of vertex $N+1$ should go to zero at large N. Clearly this is not the case in the Euclidean stochastic TSP because of the triangle inequality: the neighbor of a neighbor is itself a neighbor. But in the independent edge-length model, the neighbors represented in figure 9 are "far away" from one-another with a probability tending towards 1 as $N \to \infty$. This kind of random "geometry" is then expected to lead to uncorrelated spins among the finite order neighbors of S_{N+1} and so the cavity calculation may become exact as $N \to \infty$.

Although it is not clear yet that the correlations go away as $N \to \infty$ in the independent edge-lengths ensemble, the reasoning above is supported by extensive simulational results. In these kinds of tests, one generates weighted graphs in the ensemble of interest, determines the optimum tour for different sizes N, and then estimates the statistical properties in the large N limit. All such simulational studies to date have confirmed the validity of the cavity method. Both the assumptions of no replica symmetry breaking [70] and the predictions for $\quad$ and $P(d_{1,N+1})$ have been validated [69,71,70] in that way. Although these tests have limited precision in the context of the TSP, more stringent tests [72,73] have been performed on matching problems. For instance, using the cavity and replica methods, Mezard and Parisi predicted [68] that the length of a minimum matching of N points would have the large N limit $\quad^2/12$ when the d_{ij} are uniformily distributed in $[0;1]$. The numerical simulations confirm this value at the level of 0.05%.

The consensus is thus that the cavity method gives exact results at large N

for all independent edge-lengths disorder ensembles. But for the physicist, this is not the only interest of the cavity method: even as an approximation, it is useful for understanding the effects of quenched disorder. For instance, one can ask [69] how bad is the factorization approximation when applied to the Euclidean TSP in $d = 2$. For that, we compare Krauth and Mezard's cavity prediction $(2) = 0.7251\ldots$ to the best estimate from numerical simulations [74,71] $0.7120 \quad 0.0004$. We see that in fact the prediction is quantitatively good, and it turns out that this approximation becomes even better as the dimension of space d is increased.

5.7 Remarks on the cavity approach and replica symmetry breaking.

In some respects, the cavity method is complementary to the replica method, but both become unwieldy when replica symmetry is broken. In the case of the TSP, it turns out that only the cavity method has allowed a complete solution, but that model has no replica symmetry breaking. When replica symmetry breaking does arise, the situation is far more complex, and to date only models defined on graphs with infinite connectivity have been solved exactly (though not rigorously). Nevertheless, recent progress [75] in using the cavity method may soon lead to "exact" solutions of other models such as K-SAT in spite of the presence of replica symmetry breaking.

6 Related topics and conclusion.

6.1 Other optimization problems investigated in physics.

This article has focused on presenting statistical physics tools in the context of a few well-known problems. But many other random combinatorial problems have been considered by physicists, often using nearly identical techniques to the ones we have presented. For the reader interested in having a more complete view of such work, we give here a partial list of problems and pointers to the litterature.

Graph bipartitioning.

Given a graph G, partition its N vertices into two sets of equal size. The cost of the partition is the number of edges connecting vertices in different sets. The graph bipartitioning (or graph bisection) problem consists in finding the minimum cost partition.

This problem is readily reformulated in the physics language of spins: to each vertex i attach a spin S_i and set it to $+1$ if the vertex is assigned to the first set and -1 if it is assigned to the second set. Calling G_{ij} the adjacency matrix of the graph G, the number of edges "crossing" the partition can be identified with an energy:

$$E = \frac{1}{2} \sum_{i<j} G_{ij} (1 - S_i S_j) \quad . \tag{134}$$

Since the partition is assumed balanced, the global magnetization $M = \sum_i S_i$ is constrained to be zero. In physics studies, researchers enforce this constraint in a soft way by adding $M^2/2$ to the energy E, where α is a positive parameter. As a result, spins interact through effective couplings $J_{ij} = (G_{ij} - \alpha)/2$ that can be positive or negative. The corresponding energy function is then seen to be a spin glass Hamiltonian, similar to the Sherrington–Kirkpatrick model exposed in Section 2.4.2. The first authors to notice this identification were Fu and Anderson [76,77]. They then applied the Parisi solution of the Sherrington–Kirkpatrick model to give the large N value of the minimum cost partition when G has connectivities growing linearly with N. These results generalize to weighted graphs straightforwardly.

Weighted minimum bipartite matching.

Let I and J be two sets containing N points each. We assume given an $N \times N$ matrix of "distances" d_{ij} defined for each pair $i \in I, j \in J$. For any complete matching (a one-to-one map or a pairing between I and J, more commonly known as a bipartite matching), its cost is defined as the sum of the distances between paired points. In the minimum weighted bipartite matching problem one is to find the complete matching of lowest cost. Naturally, one can consider a stochastic version where the entries of the distance matrix are independent random variables, drawn from a probability distribution $p(d)$. This problem is close in its technical aspects to the stochastic TSP, and like the non-bipartite case it has been "solved" both via the replica and the cavity methods [68,78]. In the special case where $p(d)$ is the uniform distribution in $[0;1]$, Mezard and Parisi have computed the large N limit of the typical cost to be $\pi^2/6$. In fact, in a real tour de force, they also obtained the form of the $1/N$ correction to this limit. More recently, Parisi considered the special case $p(d) = \exp(-d)$ and conjectured [79] that for any N the mean minimum cost is given by $\sum_{k=1,\ldots,N} 1/k^2$. All current evidence, both numerical and analytical for small N values [80], indicates that this formula at finite N could be exact.

Number partitioning.

This problem can be motivated by the need to divide an estate between two inheritors in a fair way. It is usually formulated as follows. Let $\{x_1, x_2, \ldots, x_N\}$ be N real numbers in $[0,1]$ and consider a partition of the x_i into two (unbalanced) sets. The "unfairness" of a partition is the sum of the x's in the first set minus the sum of the x's in the second. The number partitioning problem consists in determining the partition that minimizes the absolute value of the unfairness. When the x_i are independent random numbers, it is possible to derive some statistical properties of the minimum. We refer the reader to Mertens' detailed review in the present issue [4] of his recent work.

Vertex cover

Very recently, A. Hartmann and M. Weigt studied the minimum size of vertex coverings of random graphs. Phase transitions take place, accompanied by drastic changes of the computational complexity of finding optimal vertex coverings using branch{and{bound algorithms. See the article in the present volume [3].

Neural Networks.

To a large extent, learning and generalization properties of formal neural networks are optimization problems. These properties have been the subject of intense studies by statistical physicists in the last fifteen years. A quite complete review of these works and results are exposed in the article by A. Engel in this volume [5].

6.2 Further statistical properties.

Statistical physics concepts and techniques are powerful tools to investigate the properties of ground states, that is the solutions of combinatorial optimization problems. So far, we have concentrated on the large size (large "N") limit of these problems, but one can also consider finite N. In addition, it may be of interest to know the properties of the near-optimum solutions.

Finite-size corrections and scaling.

Mean-field models can be solved through saddle-point calculations in the infinite size limit only. Clearly, optimization problems usually deal with a finite number of variables. It is therefore crucial to achieve a quantitative under-

standing of the finite size corrections to be expected, e.g., on the ground state energy.

Far from phase transitions, corrections to the saddle-point value can usually be computed in a systematic way using perturbation theory. An example of such a calculation to determine finite-size corrections has been mentioned previously (see the bipartite matching problem discussed in Section 6.1). For any quantity or "observable" associated with the optimum solution of a problem, one can ask how its disorder-average depends on the system size. Similarly, fluctuations, which disappear in the infinite volume limit, generally matter for finite sizes. Both effects are well-known in the statistical physics of systems without disordered interactions and have been the subject of many theoretical studies [81,82].

Close to transition points, the handling of finite-size corrections is much more involved. Few results are avalaible for disordered systems [83]. Generally speaking, the transition region is characterized by a window, the width of which scales as some negative power of the system size, shrinking to zero in the infinite size limit. We have already discussed the critical scaling properties of some systems in Sections 2.3.5 and 3.3.1. No similar theoretical study of critical exponents has been performed so far for complex optimization problems, e.g. K-SAT; only numerical data or bounds on the exponents are currently available.

Finite-dimensional energy landscapes and robustness.

Realistic physical systems and certain optimization problems such as the Traveling Salesman Problem live in a finite-dimensional world. Thus, although we considered in Section 3 a percolation model on a random graph, the physics of the problem is usually modeled using a lattice in two or three-dimensional space, edges joining vertices only if they are close in Euclidean space. Models based on random graphs are considered to describe physical systems only when the dimension goes to $+\infty$.

Finite-dimensionality may have dramatic consequences on some properties of the models; for instance it is known that the critical exponents depend on the dimension of the embedding space. More crucially, in low dimensions, the correct order parameter could be quite different from what it is in infinite dimension. This issue is particularly acute in the physics community in the case of spin glasses: so far, no consensus has been reached concerning the correct description of these systems in dimension 3. Two main theories exist:

Parisi's hierarchical picture. This sophisticated theory comes from extending mean-field theory to finite dimensional spaces. It states that low lying configurations, i.e. having an energy slightly larger than the ground state,

may be very far away, in the configuration space, from the ground state. These excited configurations are organized in a complex hierarchical fashion, in fact an ultrametric structure.

The droplet picture. Conversely, the droplet picture is based on simple scaling arguments inspired from ferromagnetic systems and claims that low-lying configurations stand close to the ground state. Higher and higher energy excitations will be obtained when flipping more and more spins from the ground state.

A detailed presentation of the theories can be found in [2]. Knowing which picture is actually correct could have deep consequences for dynamical issues (see the next paragraph), and also for the robustness of the ground state. For instance, it can be important from a practical point of view to know how much a perturbation or modification of the energy function affects the ground state properties. Consider in particular the problem of image reconstruction. Can a small change in the data modify macroscopically the reconstructed image? Within the droplet picture, the answer would be generally no, while Parisi's theory would support the view that disordered systems often have non-robust ground states.

6.3 Perspectives.

The study of the statistical properties of disordered systems has witnessed major advances in the last two decades, but the most recent trend has been towards trans-disciplinary applications. Although it is difficult to guess what new directions will emerge, there has been a clear and growing interest in using statistical physics tools for investigating problems at the heart of computer science. In this review, we illustrated this for decision and optimization problems, but many other problems should follow. Looking at the most recent work, we see emerging efforts to extend these methods to understand the statistical properties of the corresponding algorithms, be-they exact or heuristic. Let us first sketch these issues and then mention some further possible directions.

Typical case computational complexity.

The notion of typical case computational complexity is appealing, and statistical physics tools may help one understand how that kind of classification of decision problems may be reached. But clearly the methods needed to do so go much beyond what we have presented: partition functions and analogous tools describe the solutions of a problem, not how long it can take to find them. Nevertheless, as we mentioned in Section 4.7 in the context of the Davis-Putnam

tree search, physical arguments can shed new light on how algorithms such as branch and bound behave near a phase transition. Thus these methods may tell us what is the typical computational complexity of an instance chosen at random in an ensemble, given a particular tree search algorithm. Extending this classification to obtain an algorithm-independent definition of typical case computational complexity may follow, but so far it remains largely open.

Long time (stationary) limit of stochastic search algorithms.

Consider heuristic algorithms that are based on stochastic search. Examples are simulated annealing, G-Walk, or deterministic limits of these such as local search. These kinds of algorithms define random walks, i.e., stochastic dynamics on a discrete space of solutions (boolean assignments for K-SAT, tours for the TSP, etc...) and these dynamics are "local": just a few variables are changed at each time step. Assume for simplicity that the initial position of the walk is chosen at random. At long times, the search settles in a steady state where the distribution of energies becomes stationary, that is time-independent. (The energy at any given time is a random variable, depending on the starting point of the search and also on all the steps of the walk up to that time. The energy thus has a distribution when considering all initial positions and all possible walks.) An obvious question is whether this distribution becomes peaked in the large size limit. Indeed, in most cases, one can show that the energy of a random solution is self-averaging; note that this corresponds simply to the self-averaging property of the thermodynamic energy at infinite temperature. In fact, for the problems we have focused upon, the energy is expected to be self-averaging at all temperatures. By a not so bold extrapolation, one may conjecture that any local stochastic search algorithm leads to self-averaging energies in the long time limit. (Naturally, we also have to assume that the algorithms do not have too much memory; using a simulated annealing with temperatures changing periodically in time will not do!) There is numerical evidence [84] in favor of this conjecture, and it may be possible to use statistical physics methods to prove it in some limiting cases. One can also ask what is the limiting shape of the distribution of energies. This is a difficult question, but it may be easier in this context than when considering the optimum.

Dynamics of stochastic search algorithms.

Is the self-averaging behavior just mentioned restricted to long times? Since the initial energies are those of random solutions and are thus self-averaging, it is quite natural to generalize the conjecture to all times: "the energy at any given time of a local stochastic search algorithm is self-averaging".

Quite a bit of intuition about this issue can be obtained by considering what happens by analogy with a physical system relaxing towards equilibrium. The main characteristic of the dynamics in a physical system is the property called detailed balance; this condition puts very stringent restrictions on the transition probabilities. But within this specific framework, there has been much progress recently in describing the time dependence of the dynamical process. In particular, the conjecture introduced above is confirmed in the context of mean field p-spin glass models. The exact solution of these models has led to new results on entropy production while the phenomenon of "ageing" has been explained theoretically. Clearly an important goal is to extend these results to arbitrary stochastic dynamics without the hypothesis of detailed balance. But perhaps one of the most remarquable results coming from these studies (see for instance the contribution of Bouchaud et al. in [2]) is a relation between the relaxation during these dynamics and the effects of a perturbation: the prediction, called the generalized fluctuation-dissipation relation, seems numerically to be quite general and it would be of major interest to test it in the context of more general stochastic dynamics.

Further directions.

We will be brief and just give a list of what we consider to be promising topics. First, just as the notion of computational complexity has to be generalized to a typical case description, the analogous generalization of approximability is of interest. In its stochastic or typical extension, an algorithm provides an typical case approximation to a problem if with probability tending towards 1 in the large size limit, its output is within ϵ of the actual solution. Naturally results that hold in the worse case also hold stochastically, but one may expect new properties to hold in this generalized framework. Second, there has been an upsurge of interest in physics for combinatoric problems, using techniques from field theory and quantum gravity. The problems range from coloring graphs to enumerating meanders. Although the initial problem has no disorder, the approaches use identities relating systems with disorder to systems without disorder that are as yet still in the conjectoral stage. Third, is there a relation between replica symmetry breaking and typical case complexity? Forth, will the statistical physics approaches in artificial neural networks and learning lead to new developments in artificial intelligence? Fifth, an active subject of study in decision science concerns "belief propagation" algorithms which are extensions of the cavity method. Can these extensions lead to better understanding of physical systems, and inversely, will the use of physics concepts such as temperature, mean field, scaling, and universality continue to lead to improved algorithms in practice?

Acknowledgements | O.C.M. acknowledges support from the Institut Universitaire de France. We thank C. Kenyon and P. Flajolet for comments and encouragements.

A Answers to Exercises

A.1 Exercise 1: System with two spins and statistical independence.

The partition function (3) at temperature $T = 1/\beta$ reads

$$\begin{aligned} Z(T) &= \sum_{\sigma_1,\sigma_2=\pm 1} \exp\left[-\frac{1}{T} E(\sigma_1,\sigma_2)\right] \\ &= \sum_{\sigma_1,\sigma_2=\pm 1} \exp(\beta J \sigma_1 \sigma_2) \\ &= 4\cosh(\beta J) \quad . \end{aligned} \tag{A.1}$$

The magnetization $m(T)$ and the average value of the energy $\langle E \rangle_T$ can be computed from the knowledge of Z, see (8). One obtains

$$m(T) = \langle \sigma_1 \rangle_T = 0 \quad ; \tag{A.2}$$

and

$$\langle E \rangle_T = -J \tanh(\beta J) \quad . \tag{A.3}$$

The magnetization vanishes since any configuration $\{\sigma_1, \sigma_2\}$ has the same statistical weight as its opposite, $\{-\sigma_1, -\sigma_2\}$.

These calculations can be repeated for the second choice of the energy function, $E(\sigma_1,\sigma_2) = -J\sigma_1\sigma_2$, with the following results:

$$\begin{aligned} Z(T) &= (2\cosh\beta J)^2 \\ m(T) &= \tanh \\ \langle E \rangle_T &= -2J\tanh(\beta J) \quad . \end{aligned} \tag{A.4}$$

We see that the partition function is the square or the single spin partition function. The magnetization and the energy (once divided by the number of spins) are equal to the ones of a single spin, see expression (4).

This coincidence is a direct consequence of the additivity property of the energy. More precisely, whenever the energy of a system can be written as the sum of two (or more) energies of disjoint subsystems, i.e., involving disjoint configuration variables, the partition function is simply the product of the subsystem's partition functions. Such disjoint subsystems do not interact and are statistically independent.

A.2 Exercise 2: Zero temperature energy and entropy.

Let us suppose that the configurations C form a discrete set. Let us call E_0 the smallest energy and N_0 the number of configurations having this energy. Similarly we call E_1 the immediately higher value of energy, with degeneracy N_1. This process can be repeated for more and more excited energies. At the end, configurations are sorted according to their energies with $E_0 < E_1 < E_2 < \ldots$.

From the definition (3) of the partition function, we write

$$\begin{aligned} Z &= \sum_{j \geq 0} N_j e^{-\beta E_j} \\ &= e^{-\beta E_0} \left(N_0 + N_1 e^{-\beta G_1} + N_2 e^{-\beta G_2} + \ldots \right) ; \end{aligned} \tag{A.5}$$

where $G_j = E_j - E_0$ is the gap between the j^{th} excited energy and the minimal one. By construction, all gaps G_j are strictly positive ($j \geq 1$). Thus, in the small temperature (large β) limit, we obtain

$$Z(T) = N_0 e^{-\beta E_0} \left(1 + O\left(e^{-\beta G_1} \right) \right) ; \tag{A.6}$$

from which we deduce the free-energy,

$$F(T) = -T \ln Z(T) = E_0 - T \ln N_0 + O\left(\frac{1}{\beta} e^{-\beta G_1} \right) . \tag{A.7}$$

From the definition of entropy (11), it appears that the zero temperature entropy $\langle S \rangle_{T=0}$ is simply the logarithm of the number of absolute minima of the energy function $E(C)$.

A.3 Exercise 3: Spins on the complete graph in the presence of a field.

The calculations are immediate from (25). The only difference is that, in the presence of a small but non zero field h, the two minima of the free-energy shown on figure 2 are now at two different heights. One of the two minima (with the opposite sign of h) is exponentially suppressed with respect to the other.

A.4 Exercise 4: Quenched average.

Using the results of Exercise 2, we write the partition function, magnetization and the average value of the energy,

$$\begin{aligned} Z(T;J) &= 4\cosh(\beta J) \\ m(T;J) &= 0 \\ \langle E\rangle_T(J) &= -J\tanh(\beta J) \quad : \end{aligned} \tag{A.8}$$

All these statistical quantities depend on the quenched coupling J.

We now average over the coupling J, with distribution $\rho(J)$ on the support $[J_-;J_+]$. We obtain for the quenched average magnetization and energy,

$$\begin{aligned} \overline{m(T)} &= 0 \\ \overline{\langle E\rangle_T} &= \int_{J_-}^{J_+} dJ\ \rho(J)\,(-J\tanh(\beta J)) \quad : \end{aligned} \tag{A.9}$$

In the zero temperature limit, the spins align (respectively anti-align) onto each other if the coupling J is positive (resp. negative). The resulting ground state energy equals $-|J|$. Averaging over the quenched coupling, we obtain

$$\overline{\langle E\rangle_{T=0}} = \int_{J_-}^{J_+} dJ\ \rho(J)(-|J|) \quad : \tag{A.10}$$

A.5 Exercise 5: Frustrated triangle of spins.

Both energies are even functions of the spins; the magnetization is thus always equal to zero.

We first consider the energy function

$$E(\sigma_1,\sigma_2,\sigma_3) = -\sigma_1\sigma_2 - \sigma_1\sigma_3 - \sigma_2\sigma_3 \quad : \tag{A.11}$$

The partition function and the average value of the energy read respectively,

$$\begin{aligned} Z(T) &= 2e^{3\beta} + 6e^{-\beta} \\ \langle E\rangle_T &= \frac{-3+3e^{-4\beta}}{1+3e^{-4\beta}} \quad : \end{aligned} \tag{A.12}$$

In the zero temperature limit, the ground state energy and entropy are given by

$$\begin{aligned} \langle E \rangle_{T=0} &= -3 \\ \langle S \rangle_{T=0} &= \ln 2 \quad . \end{aligned} \tag{A.13}$$

There are indeed two configurations with minimal energy; all their spins are aligned in the same direction.

We now consider the energy function

$$E(\sigma_1, \sigma_2, \sigma_3) = -\sigma_1\sigma_2 - \sigma_1\sigma_3 + \sigma_2\sigma_3 \quad . \tag{A.14}$$

The partition function and the average value of the energy now read respectively,

$$\begin{aligned} Z(T) &= 6e^{\beta} + 2e^{-3\beta} \\ \langle E \rangle_T &= \frac{-3 + 3e^{4\beta}}{3 + e^{4\beta}} \quad . \end{aligned} \tag{A.15}$$

In the zero temperature limit, the ground state energy and entropy are given by

$$\begin{aligned} \langle E \rangle_{T=0} &= -1 \\ \langle E \rangle_{T=0} &= \ln 6 \quad . \end{aligned} \tag{A.16}$$

As a result of frustration, the ground state energy is higher than in the previous case, as well as the number of ground states. Note also that the gap between the lowest and second lowest energy levels has become smaller.

A.6 Exercise 6: Partition function of the Sherrington-Kirkpatrick model.

The partition function of the Sherrington-Kirkpatrick (SK) model reads

$$Z(J) = \sum_{\{\sigma_i = \pm 1\}} \exp\left(-\beta \sqrt{\frac{1}{N}} \sum_{i<j} J_{ij}\sigma_i\sigma_j\right) \quad , \tag{A.17}$$

where the quenched couplings $J = \{J_{ij}; 1 \leq i < j \leq N\}$ are randomly drawn from the Gaussian distribution

$$P(J) = \prod_{1 \leq i<j \leq N} \sqrt{\frac{1}{2\pi}} \exp\left(-\frac{1}{2}J_{ij}^2\right) \quad . \tag{A.18}$$

To compute the average value of the partition function, we first average the couplings out and only then calculate the sum over the spins

$$\begin{aligned}\overline{Z(J)} &= \int dJ\, P(J)\, Z(J) \\ &= \sum_{i=1} \exp\left(\frac{2}{2N} \sum_{i<j} (\sigma_i \sigma_j)^2 A\right) \\ &= 2^N \exp\left(\frac{2}{4}(N-1)\right) ; \end{aligned} \tag{A.19}$$

We now calculate the second moment of the partition function by rewriting the squared sum as the product of two independent sums, see Exercise 1,

$$\begin{aligned}\overline{Z(J)^2} &= \int dJ\, P(J)\, Z(J)^2 \\ &= \int dJ\, P(J) \sum_{i=1} \sum_{i=1} \exp\left(\frac{1}{\sqrt{N}} \sum_{i<j} J_{ij}(\sigma_i \sigma_j + \tau_i \tau_j) A\right) \\ &= \sum_{i=1} \sum_{i=1} \exp\left(\frac{2}{2N} \sum_{i<j} (\sigma_i \sigma_j + \tau_i \tau_j)^2 A\right) \\ &= \overline{Z(J)}^2 Y ; \end{aligned} \tag{A.20}$$

where Y equals

$$\begin{aligned} Y &= \frac{1}{4^N} \sum_{i=1} \sum_{i=1} \exp\left(\frac{2}{N} \sum_{i<j} \sigma_i \sigma_j \tau_i \tau_j A\right) \\ &= \frac{1}{4^N} \exp\left(\frac{2}{2}\right) \sum_{i=1} \sum_{i=1} \exp\left(\frac{2}{2N} \left[\sum_i \sigma_i \tau_i\right]^2 A\right) : \end{aligned} \tag{A.21}$$

The calculation proceeds as in the case of the spin model on the complete graph, see section 2.3. We define for each configuration $C = \{\sigma_i, \tau_i\}$ of the $2N$ spins, the overlap function

$$q(C) = \frac{1}{N} \sum_{i=1}^{N} \sigma_i \tau_i : \tag{A.22}$$

The effective energy function appearing in the last term of the pseudo partition function Y (A.21) depends on the configuration through $q(C)$ only. Following the steps of section 2.3, a saddle-point calculation leads to the asymptotic

behaviour of Y,

$$Y = \exp\left(N^{2} + o(N) \right) \quad ; \tag{A.23}$$

where is the minimum over q of the "free-energy" functional $\hat{f}(q)$ defined in (25) with T^2 instead of T. The results of section 2.3 teach us that there is a "critical" temperature $T_c = 1$ such that $= 0$ for temperatures above T_c and < 0 when $T < T_c$.

Above T_c, the partition function does not fluctuate too much around the average value $Z(J)$; the partition function is itself self-averaging and the free-energy per spin simply equals $f(T) = T \ln 2 \quad 1=(4T)$, see the paper by M. Talagrand in the same volume. At low temperatures, below T_c, the second moment of $Z(J)$ is exponentially larger than the squared average; there are huge fluctuations and the partition function is not self-averaging. It is therefore much more complicated to calculate the value of the free-energy.

A.7 Exercise 7: A toy replica calculation.

We want to compute the series expansion of $\ln(1+x)$ starting from the identity (for small real n)

$$(1+x)^{n} = 1 + n \ln(1+x) + O(n^{2}) \quad ; \tag{A.24}$$

and the series expansion of $(1+x)^n$ for integer n. To do so, we use Newton's binomial formula

$$(1+x)^{n} = \sum_{k=0}^{n} \frac{n!}{k!(n \quad k)!} x^{k} \quad ; \tag{A.25}$$

valid for positive integers n. n play two roles in formula (A.25). First, it is the upper bound of the sum over k. Secondly, n appears in the combinatorial factor in the sum. Factorials may be continued analytically to real values of n using Euler's Gamma function. As (z) has poles at negative integer values of the argument z, we may extend the sum in expression (A.25) to integer values of k larger than n without changing the final result,

$$(1+x)^{n} = \sum_{k=0}^{1} \frac{n!}{k!(n \quad k)!} x^{k} \quad : \tag{A.26}$$

Let us focus now on the combinatorial factor

$$C(n;k) = \frac{n!}{k!(n-k)!} = \frac{n(n-1)(n-2)\cdots(n-k+1)}{k!} \quad : \quad (A.27)$$

For $k = 0$, we have $C(n;0) = 1$ for all n. When $k \geq 1$, the r.h.s. of (A.27) is a polynomial of n and can be immediately continued to real n. In the small n limit, we obtain

$$C(n;k) = n\,\frac{(-1)(-2)\cdots(-k)}{k!} + o(n) = n\,\frac{(-1)^{k-1}}{k} + o(n) \qquad (k \geq 1) : (A.28)$$

Finally, we write the small n continuation of equation (A.25) as

$$(1+x)^n = 1 + n\sum_{k=1}^{\infty}\frac{(-1)^{k-1}}{k}x^k + o(n) \quad : \quad (A.29)$$

Comparing equation (A.24) and (A.29), we obtain the correct result

$$\ln(1+x) = \sum_{k=1}^{\infty}\frac{(-1)^{k-1}}{k}x^k \quad : \quad (A.30)$$

The above calculation is a simple application of the replica trick. Obviously, the calculation of the free-energy of disordered models, e.g. the K-Satisfiability or the TSP models, are much more involved from a technical point of view.

References

[1] M. Mezard, G. Parisi, M. A. Virasoro (Eds.), Spin Glass Theory and Beyond, World Scientific, Singapore, 1987.

[2] A. P. Young (Ed.), Spin Glasses and Random Fields, World Scientific, Singapore, 1998.

[3] A. Hartmann, M. Weigt, Statistical mechanics perspective on the phase transition in vertex covering of finite-connectivity random graphs, in the present issue of Theoretical Computer Science.

[4] S. Mertens, A physicist's approach to number partitioning, in the present issue of Theoretical Computer Science.

[5] A. Engel, Complexity of learning in artificial neural networks, in the present issue of Theoretical Computer Science.

[6] R. Baxter, Exactly solved models in Statistical Mechanics, Academic Press, San Diego, 1982.

[7] H. Rieger, Frustrated systems: Ground state properties via combinatorial optimization, in: J. Kertesz, I. Kondor (Eds.), Advances in Computer Simulation, Vol. 501 of Lecture Notes in Physics, Springer-Verlag, Heidelberg, 1998.

[8] P. Young, Informatics – 10 years back, 10 years ahead, celebration of the 10th aniversary of Schloss Dagstuhl.

[9] R. Reif, Fundamentals of Statistical and Thermal Physics, McGraw-Hill, New-York, 1965.

[10] S. Ma, Statistical Mechanics, World Scientific, Singapore, 1985.

[11] K. Huang, Statistical Mechanics, Wiley, New-York, 1967.

[12] W. Saenger, Principle of Nucleic Acid Structure, Springer-Verlag, New-York, 1984.

[13] T. Strick, J. Allemand, D. Bensimon, R. Lavery, V. Croquette, Phase coexistence in a single DNA molecule, Physica A 263 (1998) 392.

[14] J. M. Steele, Probability Theory and Combinatorial Optimization, SIAM, Philadelphia, 1997.

[15] M. R. Garey, D. S. Johnson, Computers and Intractability: A Guide to the Theory of NP-Completeness, Freeman, New York, 1979.

[16] C. H. Papadimitriou, K. Steiglitz, Combinatorial Optimization: Algorithms and Complexity, Prentice Hall, Englewood Cliffs, NJ, 1982.

[17] B. Bollobas, Random Graphs, Academic Press, New-York, 1985.

[18] R. Potts, Proc. Camb. Phil. Soc. 48 (1952) 106.

[19] P. Kasteleyn, C. Fortuin, J. Phys. Soc. Japan Suppl. 26 (1969) 1114.

[20] F. Wu, The potts model, Rev. Mod. Phys. 54 (1982) 235.

[21] A. Morgante, Large deviations in random graphs, Tech. rep., Laboratoire de Physique Theorique de l'ENS, rapport de stage (1998).

[22] R. Monasson, R. Zecchina, Entropy of the K-satisfiability problem, Phys. Rev. Lett. 76 (1996) 3881–3885.

[23] R. Monasson, R. Zecchina, Statistical mechanics of the random K-Sat problem, Phys. Rev. E 56 (1997) 1357–1361.

[24] R. Monasson, R. Zecchina, Tricritical points in random combinatorics: the (2+p)-SAT case, J. Phys. A 31 (1998) 9209–9217.

[25] R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, L. Troyansky, Computational complexity from 'characteristic' phase transitions, Nature 400 (1999) 133–137.

[26] R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, L. Troyansky, 2+p-sat: Relation of typical-case complexity to the nature of the phase transition, Random Structures and Algorithms 3 (1999) 414.

[27] R. Monasson, Optimization problems and replica symmetry breaking in finite connectivity spin glasses, J. Phys. A 31 (1998) 513.

[28] G. Biroli, R. Monasson, M. Weigt, A variational description of the ground state structure in random satisfiability problems, Euro. Phys. J. B 14 (2000) 551.

[29] C. H. Papadimitriou, Computational Complexity, Addison–Wesley, 1994.

[30] S. Cook, The complexity of theorem–proving procedures, in: Proc. 3rd Ann. ACM Symp. on Theory of Computing, Assoc. Comput. Mach., New York, 1971, p. 151.

[31] T. H. Eds., B. A. Huberman, C. Williams, Issue 1-2, special issue on phase transitions, Artificial Intelligence 81.

[32] D. Mitchell, B. Selman, H. Levesque, Hard and easy distributions of sat problems, in: Proc. of Am. Assoc. for Artif. Intell. AAAI-92, 1992, pp. 456–465.

[33] A. Goerdt, A threshold for unsatisfiability, J. Comput. System Sci. 53 (1996) 469.

[34] V. Chvatal, B. Reed, Mick gets some (the odds are on his side), in: Proc. 33rd IEEE Symp. on Foundations of Computer Science, 1992, p. 620.

[35] S. Kirkpatrick, B. Selman, Critical behaviour in the satisfiability of random boolean expressions, Science 264 (1994) 1297.

[36] A. Broder, A. Frieze, E. Upfal, On the satisfiability and maximum satisfiability of random 3-cnf formulas, Proc. 4th Annual ACM-SIAM Symp. on Discrete Algorithms (1993) 322.

[37] O. Dubois Private communication.

[38] F. Ricci-Tersenghi, M. Weigt, R. Zecchina, The simplest k-satisfiability model, Phys. Rev. E 63 (2001) 026702, arXiv:cond-mat/0011181.

[39] J.R. Banavar, D. Sherrington, N. Sourlas, Graph bipartitioning and statistical mechanics, J. Phys. A Lett. 20 (1987) L1{L8.

[40] M. Mezard, G. Parisi, Mean-field theory of randomly frustrated systems with finite connectivity, Europhys. Lett. 3 (10) (1987) 1067.

[41] L. Viana, A. J. Bray, Phase diagrams for dilute spin-glasses, J. Phys. C 18 (1985) 3037{3051.

[42] I. Kanter, H. Sompolinsky, Mean-field theory of spin-glasses with finite coordination number, Phys. Rev. Lett. 58 (1987) 164.

[43] C. D. Dominicis, Y. Goldschmidt, Replica symmetry breaking in finite connectivity systems: a large connectivity expansion at finite and zero temperature, J. Phys. A 22 (1989) L775.

[44] P. Mottishaw, C. D. Dominicis, On the stability of randomly frustrated systems with finite connectivity, J. Phys. A 20 (1987) L375.

[45] C. D. Dominicis, P. Mottishaw, Replica symmetry breaking in weak connectivity systems, J. Phys. A 20 (1987) L1267.

[46] Y. Goldschmidt, P. Lai, The finite connectivity spin glass: investigation of replica symmetry breaking of the ground state, J. Phys. A 23 (1990) L775.

[47] B. Bollobas, C. Borgs, J. Chayes, J.H. Kim, D. Wilson, The scaling window of the 2-sat transition, Random Structures and Algorithms 2000, in press.

[48] D. Wilson, The empirical values of the critical k-sat exponents are wrong 2000 preprint arXiv:math/0005136.

[49] M. Leone, F. Ricci-Tersenghi, R. Zecchina, Phase coexistence and finite size scaling in random combinatorial problems, J. Phys. A, in press (2001).

[50] B. Selman, H. Kautz, B. Cohen, Local search strategies for satisfiability testing, in: Proceedings of DIMACS, 1993, p. 661.

[51] J. F. M-T. Chao, Probabilistic analysis of a generalization of the unit-clause literal selection heuristics for the k-satisfiability, Information Science 51 (1990) 289{314.

[52] S. Cocco, R. Monasson, Trajectories in phase diagrams, growth processes and computational complexity: how search algorithms solve the 3-satisfiability problem ,, Phys. Rev. Lett. 86 (2001) 1654, arXiv:cond-mat/0009410.

[53] M. Talagrand, Rigorous low temperature results for the p-spin mean field spin glass model, Probability Theory and Related Fields 117 (2000) 303–360.

[54] D. J. Aldous, The zeta(2) limit in the random assignment problem math.PR/0010063.

[55] J. Beardwood, J. H. Halton, J. M. Hammersley, The shortest path through many points, Proc. Camb. Phil. Soc. 55 (1959) 299–327.

[56] S. Lin, Computer solutions of the traveling salesman problem, Bell System Technical Journal 44 (1965) 2245–2269.

[57] S. Kirkpatrick, C. D. Gelatt Jr., M. P. Vecchi, Optimization by simulated annealing, Science 220 (1983) 671–680.

[58] V. Cerny, Thermodynamical approach to the traveling salesman problem: An efficient simulation algorithm, J. Optimization Theory Appl. 45 (1985) 41.

[59] E. L. Lawler, D. E. Wood, Branch-and-bound methods: A survey, Operations Res. 14, No. 4 (1966) 699.

[60] M. W. Padberg, G. Rinaldi, A branch and cut algorithm for the resolution of large-scale symmetric traveling salesman problems, SIAM Review 33 (1991) 60.

[61] D. Applegate, R. Bixby, V. Chvtal, W. Cook, On the solution of traveling salesman problems, Documenta Mathematica, J.D.M. ICM III (1998) 645–656.

[62] D. S. Johnson, L. A. McGeoch, The traveling salesman problem: A case study in local optimization, in: E. H. L. Aarts, J. K. Lenstra (Eds.), Local Search in Combinatorial Optimization, J. Wiley & and Sons, New York, 1997, pp. 215–310.

[63] S. Arora, Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems, Journal of the ACM 45 (1998) 753–782.

[64] J. Vannimenus, M. Mezard, On the statistical mechanics of optimization problems of the travelling salesman type, J. Physique Lett. 45 (1984) L1145–L1153.

[65] R. M. Karp, A patching algorithm for the nonsymmetric travelling salesman problem, SIAM J. Comput. 8 (1979) 561–573.

[66] W. T. Rhee, M. Talagrand, Martingale inequalities and NP-complete problems, Mathematics of Operation Research 12 (1) (1987) 177–181.

[67] W. T. Rhee, On the travelling salesperson problem in many dimensions, Random Structures and Algorithms 3 (3) (1992) 227–233.

[68] M. Mezard, G. Parisi, Mean-field equations for the matching and the travelling salesman problem, Europhys. Lett. 2 (1986) 913–918.

[69] W. Krauth, M. Mezard, The cavity method and the travelling-salesman problem, Europhys. Lett. 8 (1989) 213–218.

[70] A. G. Percus, O. C. Martin, The stochastic traveling salesman problem: Finite size scaling and the cavity prediction, J. Stat. Phys. 94 (5/6) (1999) 739–758.

[71] D. S. Johnson, L. A. McGeoch, E. E. Rothberg, Asymptotic experimental analysis for the Held-Karp traveling salesman bound, in: 7th Annual ACM–SIAM Symposium on Discrete Algorithms, Atlanta, GA, 1996, pp. 341–350.

[72] R. Brunetti, W. Krauth, M. Mezard, G. Parisi, Extensive numerical simulations of weighted matchings: Total length and distribution of links in the optimal solution, Europhys. Lett. 14 (1991) 295–301.

[73] J. Houdayer, J. H. Boutet de Monvel, O. C. Martin, Comparing mean field and Euclidean matching problems, Eur. Phys. Jour. B 6 (1998) 383–393.

[74] A. G. Percus, O. C. Martin, Finite size and dimensional dependence in the Euclidean traveling salesman problem, Phys. Rev. Lett. 76 (1996) 1188–1191.

[75] M. Mezard, G. Parisi, The Bethe lattice spin glass revisited cond-mat/0009418.

[76] Y. Fu, P. W. Anderson, Application of statistical mechanics to NP-complete problems in combinatorial optimization, J. Phys. A 19 (1986) 1605–1620.

[77] K. Fischer, J. Hertz, Spin glasses, Cambridge University Press, Cambridge, 1991.

[78] M. Mezard, G. Parisi, On the solution of the random link matching problems, J. Physique 48 (1987) 1451–1459.

[79] G. Parisi, A conjecture on random bipartite matching, cond-mat/9801176 (1998).

[80] V. Dotsenko, Exact solution of the random bipartite matching model, J. Phys. A 33 (2000) 2015, cond-mat/9911477.

[81] J. Cardy, Finite size scaling, vol.2 of Current Physics Sources and Comments, North-Holland, Amsterdam, 1988.

[82] V. Privman, Finite size scaling and numerical simulations of statistical systems, World Scientific, Singapore, 1988.

[83] G. Parisi, F. Ritort, F. Slanina, Several results on the finite-size corrections in the Sherrington-Kirkpatrick spin glass model., J. Phys. A 26 (1993) 3775.

[84] G. R. Schreiber, O. C. Martin, Cut size statistics of graph bisection heuristics, SIAM Journal on Optimization 10(1) (1999) 231–251.