Survey Propagation methods for efficient optimization and probing of glassy states

Thesis submitted for the degree of
Philosophiæ Doctor

Candidate:
Demian A. Battaglia

Supervisors:
Prof. Riccardo Zecchina
Prof. Giuseppe Santoro

Sector of
Condensed Matter Theoretical
and Computational Physics

To Enrica, the fulfillment
of my most important
research

Contents

Introduction

Very similar in spirit to a statistical physics model, a generic combinatorial optimization problem involves many discrete variables—like Boolean variables, finite sets of colors, etc.— which interact through constraints typically involving a limited number of them. Optimization consists actually in finding configurations which violates the smallest possible number of such pre-imposed constraints and this task might become extraordinarily difficult if the number of competing conditions is large. A compromise between the quality of a solution and the time required for finding it might have then to be found, leading only to sub-optimal results.

The number of violated constraints sum up to give a global cost-function which can be interpreted as the hamiltonian of a (generalized) spin system. One may decide to adhere completely to the physical analogy and to use statistical mechanics to devise new optimization strategies inspired by nature. The celebrated Simulated Annealing mimics, for instance, the thermal relaxation of a solid toward its optimal crystalline state. The more recently introduced Quantum Annealing is based on the simulation of quantum fluctuations which can trigger tunneling events into states with a better cost. These algorithmic methodologies based on physical, or even biological, *metaphors* are often very effective and may compete with *ad hoc* strategies developed for dealing with specific real-world optimization problem (like planning, mail delivery, timetable scheduling, chip design and verification, and many others). Nevertheless, despite the many important progresses achieved thanks to this interdisciplinary cross-fertilization [**65, 66**], vast classes of problems remain still today practically unsolvable, at least as far as the goal of finding true global optima is concerned.

The fields of combinatorial optimization [**130**] and complexity theory in computer science deal precisely with the design of efficient resolution strategies and with the classification of computational problems according to their "hardness". A first distinction might be introduced between tasks for which resolution can be done in polynomial time (**P** problems) and tasks for which such a resolution method is not known in general, but for which efficient verification of a guess is possible (**NP** problems). It is generally

believed that most **NP** problems cannot be solved in affordable time using ordinary hardware. Nevertheless the conjecture **P**=**NP** has not yet been ruled out by any formal proof, and the mirage still exists of finding an algorithm able to solve in polynomial time all the instances of one of the hardest problems in **NP** (the so-called **NP**-complete problems like boolean 3-satisfiability, or the Traveling Salesman Problem).

We would like at this point to stress that this classification in classes is based only on the complexity of the *worst-case*. For the sake of practical applications, in reality, the establishment of a complexity theory for the *typical cases* would be highly desirable. Just to give an example, most random instances of the 3-SAT problem can be solved without any effort, despite the fact that the hardest of all the 3-SAT instances have the maximum possible resolution complexity in **NP**. The efficient resolubility of many random 3-SAT formulas does not imply, of course, that **P**=**NP**, but, on the other hand, inferring from the **NP**-completeness of 3-SAT that the resolution of any instance is unaffordable would be too limiting.

However, when the problem instances are extracted at random from nontrivial ensembles (that is ensembles which contains many instances that are hard to solve), computer science meets statistical physics in another fundamental way: many models of crucial interest for Computer Science, like random 3-SAT, can be indeed identified with diluted spin glass models defined over random graphs with a small connectivity [**46, 70**]. The use of powerful statistical mechanics techniques, like the replica or the cavity method, developed for the study of disordered spin systems, allows to give a solid explanation of the onset of their hardness in certain specific range of parameters. The physical scenario predicts indeed that an exponentially large number of metastable states may trap for exponentially long times any stochastic local search process satisfying detailed balance, like Simulated Annealing. For small temperatures and for sub-exponential simulation times an extensive gap in energy will separate the true optimal assignments from the threshold states, hiding them to any heuristics. The physical analysis brings often to the derivation of *phase diagrams* giving a portrait of the complexity of the typical case. In several cases, the theoretical predictions are in remarkable agreement with numerical experimental investigations, and can be used to supplement the rigorous but incomplete informations provided by formal proofs.

New algorithmic results as well as many mathematical issues have been put forward by statistical physics studies, with examples ranging from the

identification of phase transitions [**117, 107**] and out-of-equilibrium analysis of randomized algorithms [**40**] to the introduction of entire new classes of algorithms [**107, 109, 26**].

An example is provided by the so-called Survey Propagation methods, based on a smart application of the cavity method at the level of one step of RSB. Survey Propagation (SP) is an efficient message-passing algorithm, which can be used to evaluate the local magnetizations of individual variables in arbitrary random instances of any constraint satisfaction problem. Assuming from the very beginning the existence of disconnected distant clusters of solutions, its applicability can be extended well inside the hard optimization region, characterized by a breakdown of the zero energy phase in an exponential number of components. The possibility of implementing the resolution of the cavity equations over individual samples, without the need of averaging over the quenched disorder (a step unavoidable, on the other hand, in applications based on the replica method) allows to extract microscopic detailed information of algorithmic relevance and optimal assignments can then be generated by resorting to a decimation process, in the spirit of the simplest real-space renormalization group.

In this work we will deal with new applications of SP to the optimization and constrained resolution of random constraint satisfaction problems (CSPs). We will focus our attention mainly on the K-SAT model because of its inherent simplicity of formulation and for comparison with the broad literature available on this subject. Nevertheless, our techniques could be generally applied to the optimization of sparse networks of arbitrary discrete constraints.

Really hard random samples can be generated in a narrow region centered on the transition point between generally satisfiable (SAT) and generally unsatisfiable (UNSAT) samples. For this reason (and independently from physics), during the last decade the problem of finding minimal energy configurations of random combinatorial problems similar to diluted spin-glasses has become a very popular algorithmic benchmark in computer science [**70**]. We will show that even the most performing optimized algorithm for the resolution of hard K-SAT instances, namely WalkSAT, undergoes an exponential slow-down in this hard region, despite the fact that it does not properly satisfy detailed balance like other more standard heuristics (in agreement with analysis presented in [**149, 11**]). SP, on the other hand, will be able to efficiently retrieve complete satisfying assignments up to almost the critical point itself. These results will be extended to the UNSAT region, where survey propagation techniques will be used to

produce sub-optimal assignments lying well below the lower bound to the glassy threshold states.

A further advantage of the SP approach, is the possibility of implementing modified equations that can be applied to the selection of an exponential number of solutions, according to a geometric criterion. It is indeed possible to assign different probabilistic weights to different clusters of assignments, in order to direct the SP convergence process towards clusters sufficiently close to a desired reference point in the configuration space. The resulting forcing mechanism is the perfect analogue of importance sampling in Montecarlo simulations, and can be added to the ordinary SP iterations without affecting their overall computational cost.

This cluster-addressing capabilities will be used to explore efficiently the geometry of the space of solutions of a given random combinatorial problem, by probing selectively the local neighborood of individual solutions. We will be able to provide then the first numerical evidence for the RSB geometric structure of the ground-state phase of large instances of NP-complete problems. In particular, we will observe the predicted landscape topology transition between a sharp clustered phase to a regime in which a more complex hierarchical organization of the states is present. Surprisingly enough, such geometrical insight could also become relevant for engineering applications in information theory, where the generation of graphs with specific properties is an integral part of the design process (*e.g.* LDPC error correcting codes [**140**]).

The efficient selection of states in hard phases containing an exponential number of well separated clusters open also the possibility, in perspective, of realizing new algorithms in which clustering is exploited for computational purposes. A solution cluster could be seen as an embedded attractor and the selection of a state as the retrieval of information contents. We will discuss here a preliminary implementation of "physical" lossy data compression, in which crucial will be the possibility of addressing reliably a single cluster among the exponential number available just by retaining a very limited amount of memory. Many other applications might be envisaged as well, ranging from pattern restoration and vector quantization to more general artificial intelligent associative systems in which the clustering of attractors could represent a defense against noise.

More importantly, these state-selecting capabilities can be achieved resorting only to completely local equations. Fully parallel versions can then be devised, particularly suitable for distributed computational environments like GRIDs. Implementations over massively parallel reconfigurable hardware (FPGA, etc.) are also possible, resulting in a computational cost which

scales only with the logarithm of the problem size. Cluster importance sampling implemented through simple trains of driving pulses, can be potentially used to produce state-selecting networks of basic units exchanging between them only digital elementary messages. Preliminary results are encouraging and seem to suggest that hard phases in random constraint satisfaction problems, longly considered as a curse to be defeated, may become, on the long run, a new powerful resource for technological and algorithmic applications, making intelligent computational use of the exponential number of selectable states.

* * *

This thesis is organized as follows:

Part 1 is about preliminaries on the characterization of complex problems. Chapter 1 will give an introduction to the Theoretical Computer Science view-point on complexity, and Chapter 2 will introduce the fundamental statistical mechanics method needed for the physical analysis of combinatorial optimization problems. The random K-SAT phase diagram will be introduced and explained.

In Chapter 3 we will discuss the behavior of some performing local search heuristics (WalkSAT and fRRT), highlighting the existence of the glassy threshold phenomenon. A comprehensive summary of the application of Quantum Annealing to the optimization of K-SAT will also be given, in order to offer an example of the most exotic approaches to optimization proposed in the last years of interdisciplinary research.

Part 2 will be entirely devoted to Survey Propagation and its generalizations. Chapter 4 will introduce the basic SP equations and will discuss their message-passing interpretation. We will present, moreover, the finite-pseudo-temperature version of SP (SP-Y) and its application to the optimization in the UNSAT region.

Chapter 5 will be about state selection. The SP-ext equations with external forcing will be studied, as well as their use for probing the cluster structure of the zero-energy phase. The distribution of solutions will be also investigated and described by mean of Principal Component Analysis visualizations. Lossy data compression of random unbiased and biased sources will be discussed as well.

Finally, Chapter 6 will be essentially about perspectives and will introduce the completely local version of SP (AdaSP). The first very preliminary results on simple warning propagation schemes with external forcing will

give us the opportunity of speculating about the possible relevance of state-selection mechanisms for neural processing of information.

The work is completed by some introductory appendices for the interested reader:

Appendix A is a quick overview of the principal replica calculations about the K-SAT phase diagram.

Appendix B is a concise introduction to worst-case Complexity Theory and Appendix C an elementary tutorial on information theory.

Appendix D discusses the SP equations in a generalized formalism, suitable for comparison with other message-passing algorithms like Belief Propagation (BP).

Part 1

Complexity and glassiness

CHAPTER 1

A computer scientist's look at Complexity

Real-life offer a plethora of examples in which combinatorial optimization problems have to be solved. Let us think about mail delivery. A given set of missives, each one labeled by a precise address, have to be dispatched to destination in the shortest possible time, taking in account the fact that the number of postmen is limited (as well as their working hours) and that fuel is expensive. Besides the omnipresent planning problem, we can also consider more academic tasks. One of the first mathematical investigations of the resolubility of a combinatorial problem was proposed in 1736 by Leonhard Euler [**49**]. The puzzle consisted in deciding if it was possible to walk through the city of Königsberg (today's Kaliningrad, in Russia, an important harbor over the Baltic Sea) crossing each one of the seven bridges over the river Pregel exactly once and returning home.

Euler had the great merit to identify the combinatorial structure underlying the problem in a seminal work that can be regarded as the founding act of graph theory. In a modern language, indeed, the Seven Bridge problem can just be considered an instance of the more general decision problem:

EULERIAN CIRCUIT: : Given a graph $G = (V, E)$, does G contain an eulerian circuit, connecting all the vertices in V, passing exactly once over each *edge* in E?

Of course, a possible solution method could be the exhaustive enumeration of all the possible paths over G, but already Euler himself remarked:

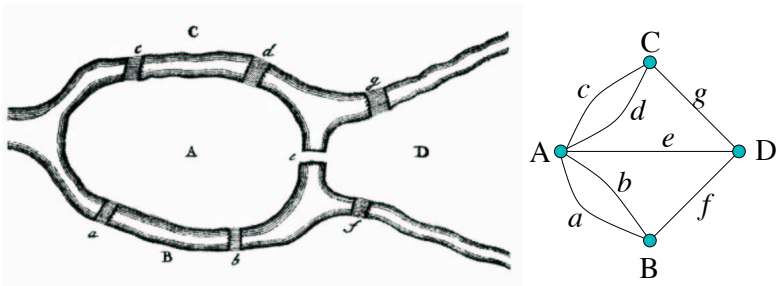

FIGURE 1. The Seven Bridges of Kaliningrad and the corresponding Eulerian circuit graph problem.

> *"The particular problem of the seven bridges of Königsberg could be solved by carefully tabulating all possible paths, thereby ascertaining by inspection which of them, if any, met the requirement. This method of solution, however, is too tedious and too difficult because of the large number of possible combinations, and in other problems where many more bridges are involved it could not be used at all"*

Euler solution was instead based on the basic observation that if there is an eulerian circuit, it must leave any vertex v via an edge e different from the edge e' used for reaching v. This simple consideration allows to state that having only even connectivities in G is a necessary condition for the existence of an eulerian circuit. This was enough to solve the Königsberg bridges problem (the answer was "no"), but the great mathematician was also able to prove the sufficient condition, devising then an efficient polynomial algorithm able to solve any random instance of the Eulerian Circuit problem, at least as far as we are only interested in *deciding* the existence or not of such a circuit.

In 1859, Sir William Rowan Hamilton sold to a London dealer for 25 pounds the idea for a new puzzle, which was commercialized under the mysterious name of "Icosian Game" [**64, 67**]. It consisted in a simple wooden pegboard with hemispheric holes in correspondence with the vertices of a dodecahedral graph. The game consisted in finding a route along the edges of the graph, passing now each *vertex* exactly once and returning to the starting edge. Despite the similarity with the Eulerian Circuit problem, no simple answers to this Hamiltonian Circuit problem are known up to date (see *eg.* [**158, 130**]). Hamiltonian Circuit Problem is generally intractable (no solutions accessible in polynomial time) and nobody knows clearly why.

1. Worst-case complexity

The preceding discussion, based on examples from years in which computers didn't exist yet, highlights the central question of an important branch of discrete mathematics, known under the name of Complexity Theory [**58, 129**]: given a specific problem, defined in precise mathematical (often combinatorial) terms, is it possible to solve it using a polynomial amount of time (and/or space) in a reasonable model of automated computation? The answer seem to be "yes" for the Eulerian Circuit problem that can then be put in the so-called class **P** of the polynomially solvable problems. A concise

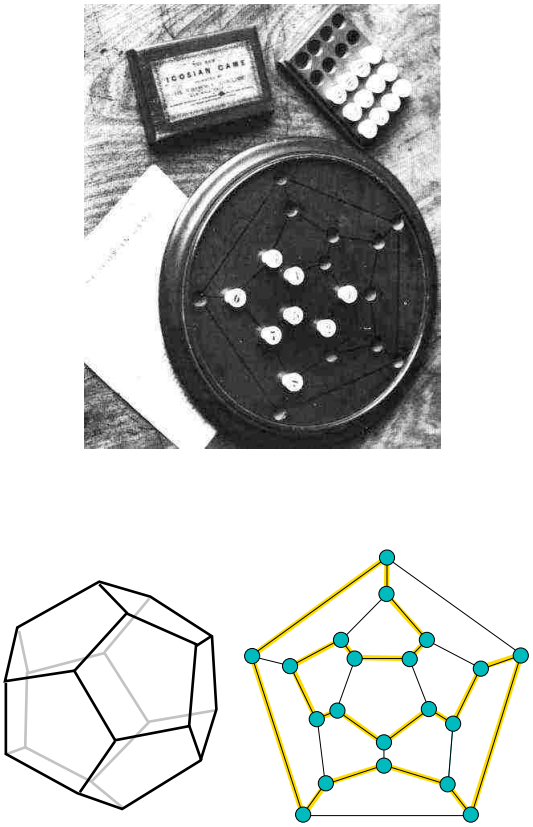

FIGURE 2. A box of the Hamilton's "Icosian Game" and its graph-theoretical solution.

but rigorous introduction to Complexity theory will be provided in the Appendix B, but we can easily understand that Hamiltonian Circuit problem will belong to a different *complexity class*. In particular it will belong to the class **NP** of the problems for which the verification of a tentative solution can be performed in polynomial time[1].

The mathematical analysis of complexity can be extended from *decision problems* (in which only the two answers "yes" and "no" can be given) to the most general case of *optimization problems*. A celebrated example is represented by the Traveling Salesman Problem (TSP). Here, the input data are a set of N cities and the matrix $\mathbb{D} = \|d_{ij}\|$ of their reciprocal distances (not necessarily corresponding to a distribution over an euclidean plain). The N cities can be seen as the vertices of a complete graph K_N, whose $N(N-1)/2$ edges are weighted by the entries of the symmetric matrix $\mathbb{D}$. The TSP consists then in finding the hamiltonian circuit over K_N with the

[1]"**NP**" does not stand for "non polynomial" as it could be suspected, but for "non-deterministic polynomial", see Appendix A.

smallest associated weight (given by the sum of all the individual edges' weights). Obviously the complexity of the optimization problem is at least equivalent to that of its decision version: finding a path of length ℓ allows also to decide if a path shorter than $\ell+\epsilon$ exists or not. Hence, TSP belongs to the **NP** class, since it can be very difficult to find a short path, but its length can be measured just by performing the sum of a polynomially bounded number of terms.

A special role in complexity theory is played by the so-called **NP**-complete problems, which are in a sense the most complex in the whole **NP** class. A problem A is **NP**-complete if all the problems in **NP** can be translated in polynomial time into the resolution of some instance of A. The existence of an efficient algorithm for the resolution of *all* the instances of A would then allow the efficient resolution of every instance of every possible **NP** problem. The 3-SAT problem which will be introduced in Section 2 and thoroughly discussed in the rest of this work constitutes the paradigm of all the **NP**-complete problems. It can be proved that also TSP is **NP**-complete (by showing that it is at least as complex as $3-SAT$) and more comprehensive lists of **NP**-complete problems can be found in [**58**].

We would like to stress here the words "all instances" appearing in the preceding paragraph about **NP**-completeness. Many instances of an arbitrary **NP**-complete problem, like TSP, can in reality be quite easy to solve. Nevertheless, it is possible to exhibit some extraordinarily difficult instances for which no resolution in polynomial time could be found until now. These maybe very few hard instances would be sufficient for classifying TSP as a **NP** problem. In particular, we would be able of show that the **NP** class is in reality included in the **P** class, if we could find a polynomial algorithm for the resolution of every possible instance of 3-SAT. Efficient resolution of just a subset of instances would not be enough for proving the discussed conjecture **P** = **NP**.

For these and other reasons, the complexity results briefly mentioned in the present section are said to belong to the field of *Worst-Case Complexity Theory* [**129**].

2. Typical-case complexity

In many practical applications, worst-case statements could be too limiting. One of the most widely used algorithms for the resolution of linear programming problems is the simplex method [**130**], which requires an exponential resolution time in the worst-case. Now, linear programming itself is known to be a polynomial problem, since the introduction of the ellipsoid algorithm in 1979. It turns out, however, that *in most cases* the simplex

method is very efficient and it is then preferred to the more awkward ellipsoid algorithm which is guaranteed to be always polynomial. Linear programming provides just an example in which the indications of worst-case analyses could be misleading for practical purposes.

Let us consider now the boolean K-satisfiability problem (K-SAT). We define a *K-clause* as the disjunction $A_1 \vee A_2 \vee \ldots A_K$ of K *literals* equal to boolean variables either directed or negated. A *formula* consists in the conjunction of M clauses. Given N Boolean variables and a formula, the K-SAT problem consists in deciding whether it exists or not an assignment of the variables that satisfies all M the clauses. Since a same variable might appear directed in some clauses, negated in others, the determination of a satisfying assignment could not be straightforward.

Many relevant problems can be mapped directly or indirectly to K-SAT. This is notably the case of planning problems, circuit fault analysis and many others [**76, 80, 81, 71**]. The logical nature of K-SAT allowed Thomas Cook to prove its **NP**-completeness for $K \geq 3$. Many hard instances of the 3-SAT problem can be downloaded from online benchmark-sets repositories [**71**]. In this work, we will be mainly interested in instances sampled from some defined random ensemble of formulas (random K-SAT). A first example, could be the ensemble of fixed connectivity random formulas, in which every variable participates to a fixed number γ of clauses, appearing each time negated or directed either with a given probability, or in an exactly balanced way. Another common ensemble is obtained when selecting in a completely random way the K variables involved in each one of the M clauses. It is possible to see that in this framework the number of clauses in which each variable is involved becomes a Poissonian random variable.

In the case of random Constraint satisfaction problems, the well-defined probabilistic set-up allows to make statements about properties of instances that are typical for certain choices of parameters. Many hard instances, however, are highly structured, far from being random, and the notion of typical-case analysis does not play any obvious role. The study of the connection (if any) between worst-case and typical-case complexity is indeed an open one and very few general results are known [**6**].

For random formulas, an extremely relevant parameter is the ratio $\alpha = M/N$. Extensive numerical experiments have been conducted in order to study the distribution of satisfiable (SAT) and unsatisfiable (UNSAT) instances. Figure 3 (adapted from [**83**]) shows the probability $P_{\mathrm{UNSAT}}(\alpha)$ of finding unsatisfiable random instances, obtained from the analysis of a large number of samples of different sizes and clauses-to-variables ratios α. A

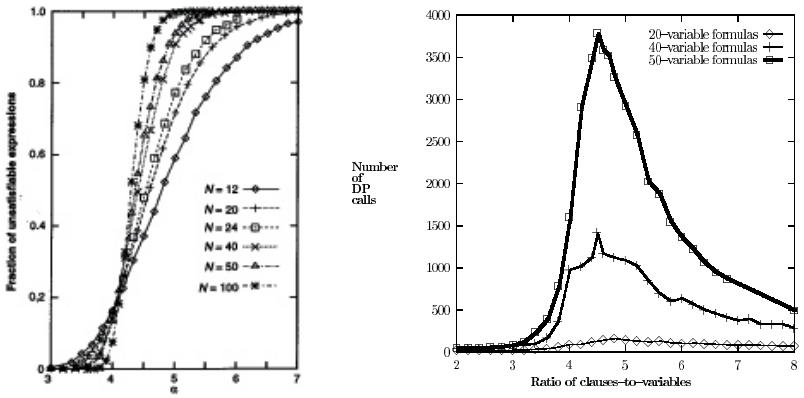

FIGURE 3. Phase transitions in random 3-SAT. (*Left*) The probability of randomly generating a satisfiable 3-SAT instance with random Poissonian connectivity is plotted for several sample sizes against the clause-to-variable ratio α. Two regions in which it is almost certain to produce SAT and UNSAT instances can be clearly seen. The complexity of proof resolution with a Davis-Putnam complete solver is also shown (*Right*) and seem to diverge for large instances across the phase-boundary region (figures adapted from [**111, 83**]).

rather sharp transition between two well defined regions can be clearly observed. At small α, it is almost certain that a randomly generated instance will be satisfiable, while for large α $P_{\mathrm{UNSAT}}(\alpha)$ seems to reach unity. The transition region becomes smaller with increasing sample size, and a finite-size scaling analysis allow to extrapolate from the data an approximate critical point location in $\alpha_c^{(\mathrm{num})} \simeq 4.3$.

Only relatively small sizes can be managed in this kind of experiments, since the use of complete solvers is needed to provide certificates of unsatisfiability. Indeed, satisfiability can be established by presenting just a single satisfying assignment, but to prove unsatisfiability the full configuration space must be explored to rule out the existence of exact solutions. In practice, this can be done by pruning considerable parts of the search tree when unavoidable contradictions are found. The Davis-Putnam algorithm is based on an efficient implementation of this selective branching strategy and constitutes the fastest known complete procedure for unsatisfiability proofs. It has been experimentally found in the case of random instances that the time cost of resolution is strongly affected by the value of α. The data presented in Figure 3 at right (adapted from [**111**])

indicates that the hardest instances are concentrated in a small region surrounding the SAT/UNSAT transition region. The observation of this EASY-HARD-EASY pattern (common in other combinatorial optimization problems) seems to suggest the existence of a second kind of transition besides the SAT/UNSAT critical point.

The study of the SAT/UNSAT phase transition is considered of crucial relevance for understanding the general mechanisms of the onset of computational complexity in typical instances. A lot of work aiming at a rigorous foundation to the emerging field of typical-case complexity has been focused on the study of both the decision problem and the optimization version of K-SAT (random MAX-K-SAT). The existence of a SAT/UNSAT critical point $\alpha_c(K)$ has been theoretically confirmed for 2-SAT, which is in **P** [**128**] and for which $\alpha_c(2) = 1$ [2]. Much less is known for the **NP**-complete case $K \geq 3$ [**60**]. A formal proof of the existence of the threshold phenomenon for large N has been published in [**55**], but the fact that the corresponding α_c has a limit when $N \to \infty$ has not yet been established rigorously (the threshold location could oscillate). Upper bounds $\alpha_{\rm UB}(K)$ on the position of α_c have been found using first moment methods [**45, 77**] and variational interpolation methods [**63, 54**], and lower bounds $\alpha_{\rm LB}(K)$ have been found using either explicit analysis of some algorithms [**5**], or second moment methods [**3**]. For random MAX-K-SAT theoretical bounds are also known [**4, 78**], as well as analytical results on the running times of random walk and approximation algorithms [**144, 136**].

Nevertheless, the experimental picture of the various phase transitions presented by the random K-SAT problem appears extremely more detailed than the very rough description based only on rigorously proved theorems. A similar state of advancement subsists for many other combinatorial optimization problems, like for instance random graph coloring [**55**]. It seems definitely that new concepts and tools are needed in order to produce further important theoretical progresses. In the next chapter, we will see how statistical physics might become a fundamental cornerstone in the construction of the new language of typical-case complexity theory.

[2] Note that MAX-2-SAT becomes again **NP**-complete for $\alpha > 1$.

CHAPTER 2

A physicist's look at Complexity

Statistical mechanical studies have often to deal with the exact or approximate solution of involved counting and optimization problems. This situation has already lead in the past to fruitful reciprocal fertilizations, transgressing the boundaries between physics and discrete mathematics. Two significant examples could be the Kasteleyn method for the counting of perfect matches over planar graphs (from the clever summing of the dimer partition function to an efficient algorithm [**79**]) and the study of percolation (from pioneering graph-theoretical works of Erdös and Rényi [**50**] to several sophisticated and interconnected models in statistical mechanics [**151**]).

Every optimization problem can be defined as the task of determining a variable assignment that makes a given lower-bounded *cost function* as small as possible. The cost of a configuration can be the length of the path in TSP, and the number of violated clauses in K-SAT. In any case it can be considered as a formal energy value, and each optimization problem can then be associated to the search for the ground state of a specific hamiltonian [**135, 123**]

To be more precise, we can define a generic constraint satisfaction problem (CSP) by introducing N discrete variables —*e.g.* Boolean variables, corresponding to Ising spins, or finite sets of colors, associated to Potts spins— and M constraints involving typically a small number of variables. The energy cost $C_a(\vec{\xi})$ of a single constraints a will be equal to 0 or not, depending if a is satisfied or not by a given assignment $\vec{\xi}$. Typically the individual costs are simple integer-valued functions of just a small subset of variables $V(a) = \{j_{a_1}, \ldots, j_{a_K}\}$ (every variable j is involved on the other hand in a subset $V(j)$ of constraints). For random CSPs, the sets $V(a)$ are generated selecting random groups of size K among the N variables. The CSP corresponds in the end to the determination of configurations minimizing the global energy function $\mathcal{E}[\vec{\xi}] = \sum_a C_a(\vec{\xi})$. We said that the CSP can be exactly solved if we are able to provide at least one assignment $\vec{\sigma}$ such that $\mathcal{E}[\vec{\sigma}] = 0$.

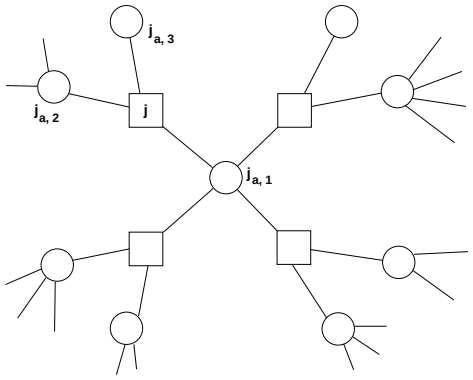

FIGURE 1. Factor graph representation of a random CSP. Variables are represented as circles connected to the constraints in which they are involved, depicted as squares.

The energy contribution associated to a single K-clause of a K-SAT formula can be for instance written as:

$$C_a^{\mathrm{SAT}}(\vec{\xi} = \{s_i\}, \{J_{a,i}\}) = 2 \prod_{l=1}^{K} \frac{1 + J_{a,l} s_{a,l}}{2} \tag{2.1}$$

where the Ising spins assume a value $s_{a,l} = \pm 1$ depending on the truth value of the l-th variable $j_{a,l}$ in the clause a and the coupling constants $J_{a,l} = \pm 1$ reflect if $j_{a,l}$ appears in a in the negated or the directed form. The energy contribution (2.1) vanishes if at least one literal is true (that is if there is at least one product $J_{a,l} s_{a,l} = -1$)[1]. The overall hamiltonian of a K-SAT instance reads then:

$$\mathcal{H}_{K\text{-SAT}}\left(\vec{\xi}, \{J_{a,i}\}\right) = 2 \sum_{a} C_a^{\mathrm{SAT}}(\vec{\xi}, \{J_{a,i}\}) \tag{2.2}$$

where the dependence from the random couplings has been explicited.

It is common to provide a graphical visualization of instances of CSPs, by drawing its factor graph representation [**86**]. Each variable is depicted as a circular node in a bipartite random graph and it is connected to square nodes representing the constraints in which they are involved (see Fig. 1). Random instances correspond to random graphs which are typically tree-like (that is the shortest loop-length is expected to be of order $\log N$ [**50, 22**]). Various kinds of connecting edges can be drawn, reflecting differences in the associated coupling constants. In a random K-SAT instance, in particular, each variable can appear directed and negated in different terms and, hence, give rise potentially to frustration.

[1]The multiplicative prefactor 2 is added for future convenience.

The cost function associated to any given random K-SAT formula can be seen then as the hamiltonian of a spin system with *quenched disorder* defined over a sparse random factor graph. Many powerful techniques have been developed in the last twenty years in order to deal with disordered statistical models and can be used in order to describe the different phases of random combinatorial optimization problems and achieve a better understanding of the mechanisms inducing resolution complexity [**135, 123**].

The way physicists and mathematicians proceed is quite different, since theoretical physicists generally do not prove theorems, but ordinarily prefer to maximize the number of approximate results based on reasonable hypothesis. If from the rigorous mathematical point of view it is only possible to state that some phase transition must take place in random 3-SAT between $\alpha = 3.14$ and $\alpha = 4.51$, statistical mechanics computations provide estimates which can be compared with the experimental observations and which result valid up to the second decimal digit at least. The extreme success of physical methods has triggered a large interest in the "orthodox" computer scientists and mathematicians community, and new proof techniques and definitions, which formalize on a solid ground the "quick and dirty" physicists' approaches, are being gradually introduced, giving rise to entire new original lines of research. This is the case of important achievements in Measure Theory [**155**] based on the cavity method, which will be reviewed in the next section 2.

Statistical mechanics method have been applied with success to the characterization of several CSPs, like graph partitioning [**56**], graph coloring [**27**], vertex cover [**160**], knapsack [**85, 73**], number partitioning [**98**], matching [**44, 104, 92, 93**], TSP [**103, 137**], and many other information theoretical applications [**123, 150**]. The dynamical properties of various local search or complete procedures have also been explored in detail, and the physical approach has provided the inspiration for new powerful algorithmic families, like the Survey-Propagation-based methods that will constitute the main topic of this work. We will focus hereafter on the K-SAT problem, as a case-study for different methodologies of analysis.

1. The replica approach

When considering disordered systems defined in terms of a set of random quenched interactions $\{J\}$, their partition function $\mathcal{Z}[\{J\}]$ itself becomes a random variable, which can be shown to be not *self-averaging*. The evaluation of free-energy and of other functionals of relevant physical interest, requires unfortunately the evaluation of the average over the quenched disorder of the logarithm of the partition function $\overline{\ln \mathcal{Z}[\{J\}]}$, and

the non self-averaging nature of $\mathcal{Z}[\{J\}]$ makes the so called annealed approximation $\overline{\ln \mathcal{Z}[\{J\}]} = \ln \overline{\mathcal{Z}[\{J\}]}$ inexact.

A solution to this problem was first introduced by Edwards and Anderson (for an excellent account on the development of spin glass theory see [**135**]). Since $f(x)^n = 1 + n\, f(x) + O\left(n^2\right)$, they could write the formal identity:

$$\overline{\ln \mathcal{Z}[\{J\}]} = \lim_{n \to 0} \frac{\overline{\mathcal{Z}[\{J\}]^n} - 1}{n} \tag{2.3}$$

When n is an integer larger than the unity, the replicated partition function $\mathcal{Z}[\{J\}]^n$ can be interpreted as the partition function of n independent replicas of the same system. Thanks to this fact, factorization becomes generally possible and, in the case of many mean-field models, an approximate sum of $\mathcal{Z}[\{J\}]^n$ can be found in the thermodynamical limit, by means of the saddle-point method. Only the dominant contribution $\overline{\mathcal{Z}[\{J\}]^n} \simeq \exp\left(-N\, f_{\text{opt}}/T\right)$ is retained, which is expressed in terms of the global extremum of an effective replicated free-energy functional. When the explicit dependence from the number of replicas n has been identified for all the relevant functionals, an analytical continuation to the complex plane is performed and the replica limit $n \to 0$ is taken[2].

Several aspects of the replica method seem rather cumbersome. In particular, the inversion of the $n \to 0$ and of the $N \to \infty$ limits and the paradoxical use of a number of replicas of the system smaller than unity itself can hardly be justified on a rigorous ground. The main argument in favor of the Replica method has historically been its remarkable success in many applications, both when considering statical and dynamical models. In a seminal paper, D. Sherrington and S. Kirkpatrick were able to find a solution of the full-connectivity version of the Edwards-Anderson model, based on a so called *Replica Symmetric* ansatz (RS). At a certain point of the derivation, an effective free energy functional is written which depends explicitly on the average overlaps between replicas $q_{ab} = \langle s^{(a)} s^{(b)} \rangle$ and on the average square magnetizations $m_a = \langle s^{(a)} s^{(a)} \rangle$, forming a square $n \times n$ symmetric matrix $\mathbb{Q}$.

Extremization with respect to this $n(n-1)/2$ overlaps has to be taken, and restrictive hypothesis over the properties of the minimizing set can be done in order to simplify the derivation. The RS ansatz consists in selecting a replica matrix $\mathbb{Q}$ in which all the off-diagonal entries are simply equal

[2]Note that performing the replica limit, the number of replicas involved becomes negative when $n < 1$. This extravagant fact forces us to maximize instead of minimize the free energy functional. After all, we are looking for the optimal value of the off-diagonal elements of a 0×0 matrix...

to q and all the diagonal entries to m, without further index-dependencies. Extremization consists then in finding the optimal values for the overlap scales q and m in the replica limit $n \to 0$.

The SK model zero-energy entropy computed under the RS ansatz was negative and the ground state energy slightly over-estimated. J.R.L. De Almeida and D.J. Thouless were able to show indeed that under certain conditions the RS optimum becomes unstable, since better optima can be found by looking for extremization in an extended functional space characterized by more complex replica matrices $\mathbb{Q}$. Better results were obtained by G. Parisi, who introduced the one step replica symmetry breaking (1-RSB) solution, containing two different inter-replica overlap scales $q_0 \leq q_1$ and one magnetization m. The construction can be iterated generating ansatzs with multiple steps of replica symmetry breaking. We can introduce $k+2$ integers $m_0, \ldots, m_{k+1}$ such that $m_0 = n$, $m_{k+1} = 1$ and m_i/m_{i+1} is an integer. The n replicas are then divided into n/m_1 groups of m_1 replicas, each one of them subdivided into m_1/m_2 groups of m_2 replicas, and so on. The inter-replica overlap scales are then set to k possible values, according to the rule:

$$q_{a,b} = q_i, \quad \text{if } \left\lfloor \frac{a}{m_i} \right\rfloor \neq \left\lfloor \frac{b}{m_i} \right\rfloor \text{ and } \left\lfloor \frac{a}{m_{i+1}} \right\rfloor = \left\lfloor \frac{b}{m_{i+1}} \right\rfloor \tag{2.4}$$

A graphical exemplification of the 2-step replica matrix for $n = 8$ is given in Figure 2. Free energy optimization must here be taken with respect to the k inter-replica overlaps and to the k parameters $m_1, \ldots, m_k$ which, after the analytical continuation needed for the replica limit $n \to 0$, are shown at the saddle point to form a sequence $0 \leq m_i \leq m_{i+1} \leq 1$. The saddle points q_i are shown to be on the other hand in a decreasing order, and the various order parameters can be collected into a single step function over the interval $0 < x < 1$:

$$q(x) = q_i \text{ if } m_i < x < m_i + 1 \tag{2.5}$$

The most precise results are obtained finally with the f-RSB ansatz, in which the free-energy is optimized with respect to the variations of a whole continuous function $q(x)$, constituting the generalization of (2.5) for an infinite number of replica symmetry breaking steps. It is now generally believed that the f-RSB solution is exact for the SK model.

1.1. Clustering of the states. RSB theory would have remained just a sophisticated technical tool its intimate connection with the complex geometrical organization of the glassy states had not been discovered.

In a ferromagnet, below the Curie point the free-energy function has a simple structure characterized by just two wells, which corresponds to two

$$\mathbb{Q}_{2-step} = \begin{pmatrix} q_2 & q_2 & q_1 & q_1 & q_0 & q_0 & q_0 & q_0 \\ q_2 & q_2 & q_1 & q_1 & q_0 & q_0 & q_0 & q_0 \\ q_1 & q_1 & q_2 & q_2 & q_0 & q_0 & q_0 & q_0 \\ q_1 & q_1 & q_2 & q_2 & q_0 & q_0 & q_0 & q_0 \\ q_0 & q_0 & q_0 & q_0 & q_2 & q_2 & q_1 & q_1 \\ q_0 & q_0 & q_0 & q_0 & q_2 & q_2 & q_1 & q_1 \\ q_0 & q_0 & q_0 & q_0 & q_1 & q_1 & q_2 & q_2 \\ q_0 & q_0 & q_0 & q_0 & q_1 & q_1 & q_2 & q_2 \end{pmatrix}$$

FIGURE 2. A 2-RSB replica matrix for 8 replicas, with $m_1 = 4$ and $m_2 = 2$.

states with opposite directions of magnetization. The free-energy of a spin-glass-like system is considered on the other hand to have an exponentially large number of minima, according to estimations based on the counting of TAP equations solutions, and the barriers between them are expected to grow indefinitely in the thermodynamic limit. Let us suppose now that the system is large but nor infinite and let us consider its evolution in time under thermal fluctuations. The system will be trapped around one of the free-energy minima for a really long time, but, ordinarily, statistical mechanics assumes that after extremely long times we will be able to observe behaviors reflecting the properties of all the other valleys. Usual averages are indeed typically taken with respect to a Gibbs measure, but the Gibbs equilibrium state can indeed be considered a mixture of several *pure states*. To be more precise, let consider the example of the Ising model. Here, the Gibbs average of a quantity, denoted by $\langle\cdot\rangle$ can be decomposed, as evident by symmetry arguments, in the convex sum of the averages over the pure states of opposite magnetization, corresponding to the free-energy minima:

$$\langle\cdot\rangle = \frac{1}{2}\langle\cdot\rangle_+ + \frac{1}{2}\langle\cdot\rangle_- \tag{2.6}$$

It can be proven [**141**] that such decomposition is always unique, if we formally define a pure state as a state in which all the connected correlation functions $\langle A(x)B(y)\rangle_{(a)} - \langle A(x)\rangle_{(a)}\langle B(y)\rangle_{(a)}$ go to zero when $|x-y| \to \infty$. The large-time behavior of the system will be then mainly dominated by pure states, since the equilibrium implies that the response to a perturbation is small and this implies by fluctuation-dissipation theorem the vanishing of connected correlation functions. This non rigorous argument by Giorgio Parisi [**131**], suggests that long-time dynamical behaviors and spin-ordering in independent valleys can be studied, ignoring the transitions to other valleys. In each pure state independently, we can define the Edwars-Anderson

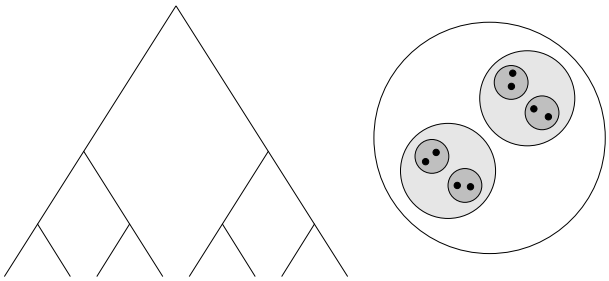

FIGURE 3. Ultrametric organization of the states, in clusters, clusters of clusters and so on. A larger number of steps of RSB corresponds to a more complex hierarchical organization.

order parameter:

$$q_{\mathrm{EA}} = \lim_{t\to\infty} \lim_{N\to\infty} \langle s_i(t_0) s_i(t_0 + t) \rangle_{(a)} \tag{2.7}$$

measuring how fast a system confined into a single state evolve in time, and corresponding with the squared magnetization inside the pure state a. The ordinary Gibbs average $\overline{q}$ will be generally different and the discrepancy between $\overline{q}$ and q_{EA} will be a sign of the existence of a multiplicity of pure states:

$$\overline{q} = \left(\sum_a P_{(a)} m_i^{(a)} \right)^2 \tag{2.8}$$

where the index (a) denotes quantities relatives to individual pure states a, each one contributing to the convex sum with a probabilistic weigth $P_{(a)}$.

A crucial result establishes a connection between the RSB order parameter $q(x)$ and the continuous spectrum between $\overline{q}$ and q_{EA}, corresponding to the variety of degrees of transitions between valleys. Let introduce indeed the probability distribution $P(q)$ of finding an overlap q between two pure states:

$$P(q) = \sum_{a,b} P_{(a)} P_{(b)} \delta_(q - q_{ab}) \tag{2.9}$$

It is then possible to see that the associated cumulative distribution:

$$x(q) = \int_0^q \mathrm{d}q'\, P(q'), \quad \frac{\mathrm{d}x}{\mathrm{d}q} = P(q) \tag{2.10}$$

is exactly the inverse of the RSB order parameter $q(x)$.

A remarkable feature of the Parisi RSB solution of the SK model (and of many others) is its *ultrametricity*. It is indeed possible to evaluate the

joint probability distribution of the overlaps between three states:

$$
\begin{aligned}
P(q_1, q_2, q_3) &= \sum_{a,b,c} P_{(a)} P_{(b)} P_{(c)} \delta(q_1 - q_{ab}) \delta(q_2 - q_{bc}) \delta(q_3 - q_{ca}) \\
&= \frac{1}{2} P(q_1)\, x(q_1)\, \delta(q_1 - q_2) \delta(q_1 - q_3) \\
&+ \frac{1}{2} P(q_1) P(q_2)\, \theta(q_1 - q_2)\, \delta(q_2 - q_3) \\
&+ \text{(other two terms with 1, 2, 3 permuted)}
\end{aligned}
\tag{2.11}
$$

This means that the distances between three states can form only equilateral or isosceles triangles. Composing then two jumps between states, it will be impossible to reach a state more distant from the initial one than the length of the longest individual jump. The states are then segregated in families of clusters and form a nested tree-like structure depicted in Figure 3. The different overlap scales individuated by the RSB theory are associated to the distance scales between states in the same cluster, in different clusters, in different clusters of clusters and so on. The 1-RSB solution introduces two overlap scales, q_0 corresponding to the typical distance between optimal configurations in different clusters and q_1 corresponding to the typical distance between optimal configurations in the same cluster. The larger the number of step of replica symmetry breaking, the larger will be the number of hierarchical levels of clustering taken in account in describing the organization of the pure states.

1.2. Applications to K-SAT. The random version of the K-SAT problem defined by the hamiltonian (2.2) has been object of careful analysis using the replica method [**113, 114, 115, 116, 117**] . The difficulty of the calculations involved has allowed to obtain only a description at the replica symmetric level. The application of replica theory to mean field models over diluted graphs presents indeed further considerable technical complications with respect to the fully connected case briefly outlined in the previous section.

In the RS approximation over a diluted model the order parameter is represented by a probability distribution $P(m)$ for the local magnetizations, identical over all the sites. It is possible to evaluate self-consistently $P(m)$ from the RS saddle point equations. The main passages of the derivation contained in [**113, 114**] are contained in the appendix A. Here, we are mainly interested in highlighting some results particularly relevant for typical case complexity theory. A variable will be relatively free to fluctuate if its magnetization is small and this will increase the probability that some fluctuations may cause the variable to satisfy many of the clauses to

which belongs. On the other hand, if a variable is strongly polarized, it will be *frozen* in a specific state, reducing then the probability of satisfying certain constraints. When performing an average over the quenched disorder — represented by the random couplings connecting the variables to the clauses in the energy terms (2.1) —, only the variables with a magnetization of order $m = \pm 1 - O\left(\mathrm{e}^{-|z|\beta}\right)$ will be relevant when computing the average number of violated clauses. It is in particular possible to express the RS ground state energy as a functional of the probability distribution $R(z)$ for the variable $z = (1/\beta)\mathrm{arctanh}\,(m)$. The resulting equations can be solved self-consistently and provide the following picture.

For $K = 2$, the RS theory predicts $E_{GS} = 0$ for $\alpha < 1$ and positive energy for $\alpha > 1$. For large α's, E_{GS} scales as $\alpha/4$ in agreement with numerical observations. The RS theory predicts then correctly the location of the SAT/UNSAT transition $\alpha_c(2) = 1$.

For $K \leq 3$, there is only a single solution corresponding to $E_{GS} = 0$ for $\alpha < \alpha_m(K)$. Then a second metastable solution leading initially to a negative E_{GS} appears and it becomes finally thermodynamically stable for $\alpha > \alpha_s(K)$, which corresponds then with the RS estimation of $\alpha_c(K)$. In reality, both $\alpha_m(3) \simeq 4.67$ and $\alpha_s(3) \simeq 5.18$ are larger than the experimentally observed $\alpha_c \simeq 4.2$, indicating that the RS theory ceases to be completely valid for $K \geq 3$.

It has been possible to evaluate perturbatively the entropy density of completely satisfying assignments in the $\beta \to \infty$ optimization limit, up to the seventh order. The results are shown in Figure 4. Let us first consider the $K = 2$ case for which RS theory is believed to be exact. at the phase transition, $S(\alpha = \alpha_c(2) = 1) = 0.38$ which is still a quite large value, when compared with $S(\alpha = 0) = \ln 2$. This means that the SAT/UNSAT transition is associated to the appearance of logical contradictory loops in all the solutions simultaneously and not to the progressive decrease of the number of these solutions down to zero at the threshold. The same happens in the $K = 3$ case ($S(\alpha = \alpha_c(3) = 4.2) = 0.1$), for which RS theory is just a rough approximation.

Analytical calculations based on variational approaches [**20**] shows that for $K \geq 3$ at a certain threshold still in the SAT region before $\alpha_c(K)$, RSB takes place[3]. Indeed, two typical distance scales appear around $\alpha \simeq 4$, corresponding to the distance between solutions within the same cluster or in different clusters. The existence of RSB will be confirmed also by analysis based on the cavity method and will be crucial for understanding the onset of computational complexity close to the SAT/UNSAT transition.

[3]In the same framework a better estimate $\alpha_v(3) = 4.48$ can be obtained for $\alpha_c(3)$.

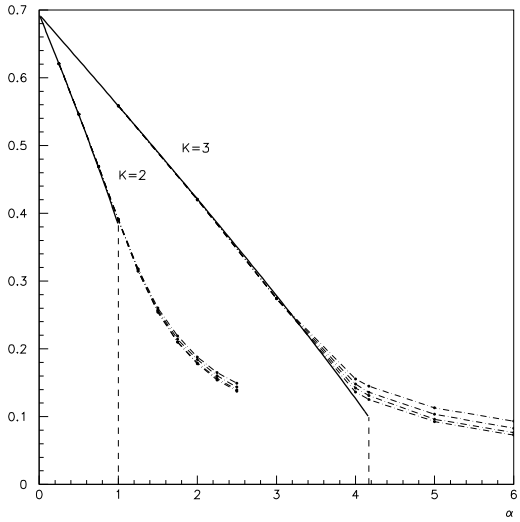

FIGURE 4. RS entropy density for random 2-SAT and 3-SAT. The RS estimates are in good agreement with exhaustive enumerations over small samples (N ranging from 20 to 30) before the experimental critical values $\alpha_c(2) = 1$ and $\alpha_c(3) = 4.3$. At the SAT/UNSAT transition an exponential number of completely satisfying assignments continues to exist (adapted from [**46**]).

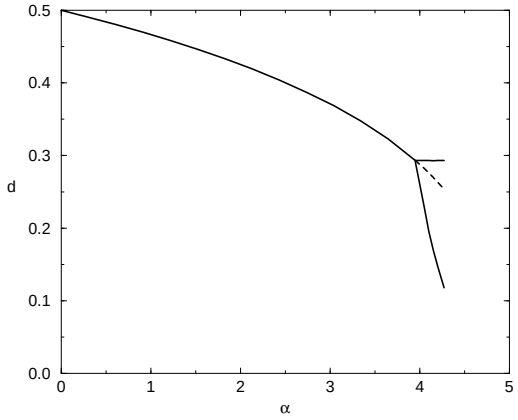

FIGURE 5. Variational RSB estimate for 3-SAT of the clustering of solutions. When RSB takes place two characteristic distance scales appear, one within each cluster and one within solutions belonging to different clusters (adapted from [**46**].

2. The "zero temperature" cavity approach

The cavity method is based on an original idea by Lars Onsager [**126**], introduced in 1936 for the study of charged particles in suspension in a dielectric medium, but it was developed in its modern form only rather recently [**105, 134, 135**], to provide a new method of resolution for mean-field spin glasses, not affected by the serious interpretation problems of the replica theory. The results provided by the cavity method are indeed equivalent to their analogues obtained in the replica framework (both at the RS and 1-RSB level), but they are based on pure (bayesian) statistical inference. Precise hypothesis about the correlations between variables in a system with N spins are indeed assumed, and they have to be self-consistently reproduced after the introduction of a $(n+1)$-th variable. The results become strictly speaking exact only in the thermodynamic limit (if it is possible to prove its existence for the problem under exam), but the introduction of precise models for describing the correlations in finite-size systems allows to use the cavity method as a tool for the definition and the rigorous study of new complex probability measures, describing ground state properties of frustrated spin systems in a language acceptable by orthodox mathematicians [**155**].

A second important advantage of the cavity method is that the site-dependence of the order parameter is explicitly taken in consideration; to be rigorous, the order parameter itself becomes the functional probability distribution of the magnetization probability distributions over the different sites. There is not then a real need to perform an average over the quenched disorder if one is interested in instance-specific features more than in typical properties, and this fact will be determinant for all the algorithmic applications described in chapters 4 and 5.

2.1. Local fields, cavity fields and cavity biases. Although the original cavity method is of more general applicability, we will focus hereafter on its version "directly at zero temperature" [**106, 109**], which is particularly suitable for the analysis of ground states of random combinatorial optimization problems. Let us consider the factor graph $\mathcal{G}^{(N)}$ of a random CSP instance with N variables and let us remove temporarily a single variable j_0 together with all its neighboring interactions $b \in V(j_0)$. The result is the so called cavity factor graph $\mathcal{G}_0^{(N-1)}$, depicted in Figure 6. Since the original instance had been randomly generated, and since N is large, we can consider that all the variables on the surface of the cavity (that is, every $j \in \cup_{b \in V(j_0)} V(b) \setminus j_0$) are approximately uncorrelated (*cavity ansatz*).

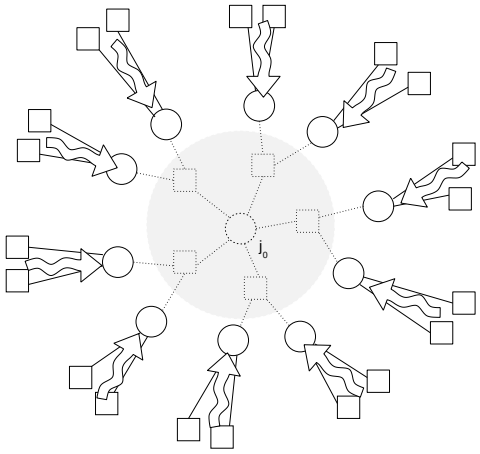

FIGURE 6. Cavity method. After the removal of the cavity variable j_0 the fields acting over the variables at the surface of the cavity can be considered independent.

The partial partition function of the cavity site σ_0 can now be written as:

$$\mathcal{Z}_{j_0} = \sum_{b \in V(j_0)} \sum_{s_i^{(b)} \in V(b) \setminus j_0} e^{-\beta C_b \left(s_1^{(b)}, \ldots, s_{K_b-1}^{(b)}, s_0 \right)} e^{\beta \sum_{j_i^{(b)} \in V(b) \setminus j_0} h_{j_i^{(b)} \to b} s_i^{(b)}} \tag{2.12}$$

where $s_i^{(b)}$ and s_0 are respectively the states of the variables $j_i^{(b)} in V(b) \setminus j_0$ and j_0 and C_b is the energy associated to the constraint b. The *cavity fields* $h_{j_i^{(b)} \to b} s_i^{(b)}$ summarize the action of the rest of the graph over the variables $j_i^{(b)}$, and can be considered independent random variables because of the cavity ansatz. We can then write out the partial resummation:

$$\mathcal{Z}_{j_0} = \sum_{b \in V(j_0)} e^{\beta \left(w_{b \to j_0} + s_0 u_{b \to j_0} \right)} \tag{2.13}$$

where the *cavity biases* $u_{b \to j_0}$ are the individual contributions of the interaction b to the *local field* H_{j_0} on the variable j_0:

$$H_{j_0} = \sum_{b \in V(j_0)} u_{b \to j_0} \tag{2.14}$$

and $w_{b \to j_0}$ is a free-energy shift depending only on the temperature and on the cavity fields $h_{j_i^{(b)} \to b}$, exactly as the cavity biases $u_{b \to j_0}$. Furthermore, the probability distribution of the state of the variable j_0 is given by:

$$P(s_0) = \frac{e^{\beta H_{j_0} s_0}}{2 \cosh \left(\beta H_{j_0} \right)} \tag{2.15}$$

The specific functional form of the free-energy shifts and of the cavity biases, can be obtained in the *optimization limit* by observing that when

$\beta \to \infty$ only the dominating contributions survive in the partial partition function, hence:

$$-\left(w_{b\to j_0} + s_0 u_{b\to j_0}\right) = \min_{s_1^{(b)},\dots,s_{K_b-1}^{(b)}} \left[C_b\left(s_1^{(b)},\dots,s_{K_b-1}^{(b)}, s_0\right) - \sum_{j_i^{(b)}\in V(b)\setminus j_0} h_{j_i^{(b)}\to b} s_i^{(b)} \right] \tag{2.16}$$

We remark that all the cavity biases, the cavity and the local fields are labeled by site and edge indices and are therefore dependent from the specific coupling constants of the instance under exam. Until now, no averages over the disorder have been taken, as already previously mentioned.

2.2. The RS cavity ansatz. The crucial statistical hypothesis behind the RS version of the cavity method in the optimization limit is that the state of the cavity graph $\mathcal{G}_0^{(N-1)}$ is correlated with the state of $\mathcal{G}^{(N)}$ *in a unique way* after the reintroduction of the cavity variable j_0. The optimal energy of $\mathcal{G}_0^{(N-1)}$ can be considered to be, under the cavity hypothesis, $A - \sum_{b\in V(j_0)} \sum_{j_i^{(b)}\in V(b)\setminus j_0} \left| h_{j_i^{(b)}\to b} \right|$ (all the variables at the surface of the cavity get aligned with the cavity fields). After the addition of the cavity site, the best energy becomes $A - \sum_{b\in V(j_0)} w_{b\to j_0} - \left|\sum_{b\in V(j_0)} u_{b\to j_0}\right|$. The introduction of the cavity site produces then finally an energy shift:

$$\Delta E_{j_0} = \sum_{b\in V(j_0)} \left(\sum_{j_i^{(b)}\in V(b)\setminus j_0} \left| h_{j_i^{(b)}\to b} - w_{b\to j_0} \right| \right) - \left| \sum_{b\in V(j_0)} u_{b\to j_0} \right| \tag{2.17}$$

We note now that the cavity biases and the cavity fields are random variables whose probability distributions can be linked by the following set of coupled integral equations:

$$P_{j\to a}(h_{j\to a}) = C_{j\to a} \int \mathcal{D}Q_{j,a}\ \delta\Big(h - \sum_{b\in V(j)\setminus a} u_{b\to j}\Big) \tag{2.18}$$

$$Q_{a\to i}(u) = \int \mathcal{D}P_{a,i}\ \delta\left(u - \hat{u}_{a\to i}\left(\left\{h_{j\to a}\right\}\right)\right) \tag{2.19}$$

where $\hat{u}^J_{a\to i}\left(\left\{h_{j\to a}\right\}\right)$ is the specific functional dependence of the cavity biases from the cavity fields and from the coupling constants and the $C_{i\to a}$'s are normalization constants. The integration measures are given explicitly by:

$$\mathcal{D}Q_{j,a} = \prod_{b\in V(j)\setminus a} Q_{b\to j}(u_{b\to j})\, du_{b\to j}, \tag{2.20a}$$

$$\mathcal{D}P_{a,i} = \prod_{j\in V(a)\setminus i} P_{j\to a}(h_{j\to a})\, dh_{j\to a} \tag{2.20b}$$

An analogous equation can be written for the probability distribution of the local fields:

$$P_i(H_i) = C_i \int \mathcal{D}Q_i \; \delta\Big(h - \sum_{b\in V(i)} u_{a\to i}\Big) \tag{2.21}$$

with the measure:

$$\mathcal{D}Q_i = \prod_{a\in V(i)} Q_{a\to i}(u_{a\to i})\, du_{a\to i} \tag{2.22}$$

Averaging over the distribution of the disordered couplings in (2.33) and over the distribution of connectivities in both (2.32) and (2.33) —the sets $\{V(a)\}$ and $\{V(j)\}$ are themselves random— have to be performed if we are interested in computing typical instance properties and phase diagrams. Many possible distributions are possible. Most commonly the couplings are selected to be $J = \pm 1$ with a flat probability, the constraint node connectivity K is fixed to a given value and the connectivities follow a poissonian distribution $f_{K\alpha}(\gamma) = (K\alpha)^\gamma \exp(-K\alpha)/\gamma!$ with average $K\alpha$ (this distribution corresponds to a large graph generation process in which a variable node is connected to a constraint node with uniform probability $2K\alpha/N^2$, see *eg.* [**22**]).

It is then possible to write an expression for the average energy contribution created by the addition of a node:

$$\Delta E_{\text{added}} = \overline{\int \prod_{a\in V(i)} \mathcal{D}P_{a,i} \Delta E_i} \tag{2.23}$$

We could naively think that ΔE_{site} itself offers an expression for the average energy density of a given CSP. In reality, when performing the averaging process we have supposed to connect the $(N+1)$-th cavity site to γ random constraint nodes according to the same probability distribution used for the generation of the topology of the cavity graph. There is then over-generation of random constraints. For instance, if the link generation probability is uniform, giving rise asymptotically to Erdös-Rényi random graphs with poissonian connectivity, we are generating constraints inside the cavity with a probability $2K\alpha/N^2$ instead of $2K\alpha/(N+1)^2$ and we have to delete a fraction $1 - N^2/(N+1)^2 \simeq 2/N$ of them at random. In the large N limit this correspond to an average number of deleted nodes equal to 2α. Each removed node a takes away the following energy term:

$$\Delta E_{\text{removed}} = \tag{2.24}$$

$$\overline{\int \mathcal{D}P_a \min_{s_1,\dots,s_{K_a}} \left[C_a(s_1,\dots,s_{K_a}) - \sum_{j\in V(a)} h_{j\to a} s_j \right] + \sum_{j\in V(a)} \left| h_{j\to a} \right|}$$

with a measure:

$$\mathcal{D}P_a = \prod_{j\in V(a)} P_{j\to a}(h_{j\to a})\, dh_{j\to a} \tag{2.25}$$

The final correct RS expression for the ground state energy density becomes then:

$$E_{GS} = \Delta E_{\text{added}} - 2\alpha \Delta E_{\text{removed}} \tag{2.26}$$

with the coefficient 2α to be appropriately modified in case of less usual random ensembles.

2.3. The 1-RSB cavity ansatz. A 1-RSB scenario is characterized by the presence of a multiplicity of different pure states, forming clusters of close configurations with the same energy. In the framework of the cavity approach, the existence of clusters have to be explicitly taken in account as a further hypothesis.

We suppose then that many different states exist. For each state separately, we can follow exactly the same passages of the RS case. Every state ℓ will be associated to its own distribution of cavity biases and cavity and local fields:

$$H_{j_0}^{(\ell)} = \sum_{a\in V(j_0)} u_{a\to j_0}\left(\{h_{j\to a}^{(\ell)}\}\right) \tag{2.27}$$

An important point is that now the addition of an additional cavity spin will produce *different energy shifts* in different pure states:

$$\Delta E_{j_0}^{(\ell)} = \sum_{b\in V(j_0)} \left(\sum_{j_i^{(b)}\in V(b)\setminus j_0} \left| h_{j_i^{(b)}\to b}^{(\ell)} - w_{b\to j_0}^{(\ell)} \right| \right) - \left| \sum_{b\in V(j_0)} u_{b\to j_0}^{(\ell)} \right| \tag{2.28}$$

We have then to take in account the possibility of energy level crossings. Let us introduce the so-called *complexity* functional:

$$N\Sigma(\alpha, E/N) = \log \mathcal{N}_{\text{clust}}(\alpha, E/N) \tag{2.29}$$

providing the logarithm of the number of pure states at a given energy density E/N. We want now to evaluate the probability distribution of the local field on an added site j_0 under the constraint that the energy density of the system is fixed to a precise value E. The local field (2.27) is correlated with the energy shift (2.28) in every state. We call $S_{j_0}(H, \delta E)$ the (unknown) joint probability considering *all* the states that the addition of j_0 produce an energy-shift δE and that a field H is observed in j_0. If the addition of the cavity spin induce a shift δE, the constraint of generating a state with energy density E can be obtained only if the spin is added to a state with

energy density $E - \delta E$. The probability distribution of the local field in j_0 at fixed energy density E can then be formally written:

$$\begin{aligned} P_{j_0}^{(E)}(H) &\propto \int \mathrm{d}(\delta E) S_{j_0}(H, \delta E) \exp\left[N\Sigma\left(\frac{E-\delta E}{N}\right)\right] \\ &\propto \int \mathrm{d}(\delta E) S_{j_0}(H, \delta E) \exp\left[-y\delta E\right] \end{aligned} \tag{2.30}$$

where we have defined the *pseudo-temperature*:

$$y = \frac{\partial \Sigma}{\partial E} \tag{2.31}$$

and we have expanded the complexity at the first order ($N \to \infty$), taking in account that the complexity is an increasing function of energy. We remark the crucial appearance of an exponential reweighing term $\chi_{j_0} = \exp\left(-y\Delta E_j\right)$. The y parameter can also directly put in correspondence with the single m_1 Parisi parameter appearing in the 1-RSB replica approach. Analogous considerations can be done for level crossings over cavity graphs. The coupled integral equations for the probability distribution functions of cavity biases and cavity fields become then:

$$\begin{aligned} &P_{j\to a}(h_{j\to a}) = \\ &\qquad C_{j\to a} \int \mathcal{D}Q_{j,a}\ \delta\Big(h - \sum_{b\in V(j)\setminus a} u_{b\to j}\Big)\ \chi_{j\to a}\Big(\{u_{b\to j}\}\Big) \end{aligned} \tag{2.32}$$

$$Q_{a\to i}(u) = \int \mathcal{D}P_{a,i}\ \delta\left(u - \hat{u}_{a\to i}\left(\left\{h_{j\to a}\right\}\right)\right), \tag{2.33}$$

with the cavity reweighing factors:

$$\chi_{j_\to a}\left(\{u_{b\to j}\}\right) = \mathrm{e}^{y\left(\left|\sum_{b\in V(j)\setminus a} u_{b\to j}\right| - \sum_{b\in V(j)\setminus a}\left|u_{b\to j}\right|\right)} \tag{2.34}$$

An analogous modification holds for the local fields probability distributions:

$$P_i(H_i) = C_i \int \mathcal{D}Q_i\ \delta\Big(h - \sum_{a\in V(i)} u_{a\to i}\Big)\ \chi_i\Big(\{u_{a\to i}\}_{a\in V(i)}\Big) \tag{2.35}$$

where the local reweighing factor is simply given by:

$$\chi_i\left(\{u_{a\to i}\}\right) = \mathrm{e}^{y\left(\left|\sum_{a\in V(i)} u_{a\to i}\right| - \sum_{a\in V(i)}|u_{a\to i}|\right)} \tag{2.36}$$

The internal energy functional could be computed in a way analogous to (2.26), but taking in account level crossing. In practice, we prefer to introduce a free-energy functional $\Phi(y)$ such that:

$$\Sigma(\alpha, E) - yE = -y\Phi(y) \tag{2.37}$$

From the knowledge of its dependence on the pseudo-temperature y, both $\Sigma(y) = \partial\Phi/\partial(1/y)$ and $E(y) = \partial(y\Phi(y))\partial y$ can be derived, and, hence, by

mean of a parametric plot, $\Sigma(E)$. We can relate the free-energy functional to the cavity and local fields probability distribution functions. The contributions associated to the variable nodes and to the constraint nodes can be isolated, in a way typical of the so-called Bethe approximation in variational approaches [**161**]:

$$N\Phi(y) = \sum_{a=1}^{M} \Phi_a^{(f)}(y) - \sum_{i=1}^{N} (\gamma_i - 1)\Phi_i^{(v)}(y) \tag{2.38}$$

Here, γ_i is the connectivity of the variable i. The contributions $\Phi_a^{(f)}(y)$ associated to every constraint a and $\Phi_i^{(v)}(y)$ associated to any variable node i can be written as:

$$\Phi_i^{v}(y) = -\frac{1}{y}\overline{\ln(C_i)} \tag{2.39}$$

$$\Phi_a^{(f)}(y) = -\frac{1}{y}\overline{\ln\left\{\int \prod_{i\in V(a)} \mathcal{D}Q_{i,a} \;\; \exp\left[\Delta E_a\right]\right\}} \tag{2.40}$$

where:

$$\Delta E_a = - y \min_{\{s_i, i\in V(a)\}} \left(C_a\left(\{s_i\}\right) - \sum_{i\in V(a)} \left[\sum_{b\in V(i)\setminus a} u_{b\to i} \right] \sigma_i + \sum_{b\in V(i)\setminus a} |u_{b\to i}| \right) \tag{2.41}$$

The equations (2.39) and (2.40) allow to compute all quantities of interest for any CSP, once known the probability distribution functions given by (2.32) and (2.35). Population dynamics algorithms can be used in order to obtain a self-consistent computation of all these probability distributions. In the next Chapter 4 we will study in detail a method allowing to derive them over single specific instances.

In practice, the collection of all the $P_i(H_i)$ defined over all the pure states and potentially different for every site constitute a sampling from a functional probability distribution of probability distributions $\mathcal{P}[P_i(H)]$. This "meta-probability" constitute the true order-parameter for the CSP. The 1-RSB approximation consists in doing a specific bayesian ansatz for the functional structure of the $P_i(H)$ (that is, a hypothesis on $\mathcal{P}[P]$). Under certain conditions, this ansatz can become unstable when embedded into a larger functional space and the 1-RSB description should have to be substituted by a more sophisticated one taking in account correlations between different pure states.

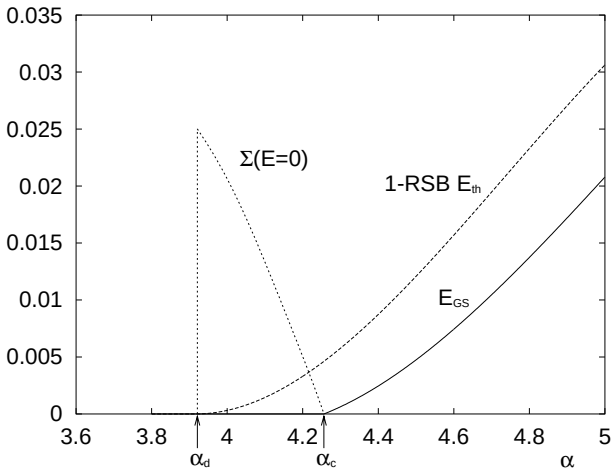

FIGURE 7. Complexity and ground state energy of random 3-SAT from the 1-RSB cavity method. The correct position of the SAT/UNSAT point can be identified. In addition, the hardness of resolution is explained by the proliferation of metastable states with an energy of scale E_{th} (adapted from [**109**]).

2.4. Applications to K-SAT. It is easy to verify starting from the hamiltonian (2.2) that the explicit form of the zero-temperature cavity bias for K-SAT becomes:

$$u_{a\to i} = \hat{u}_{a\to i}\left(\left\{h_{j\to a}\right\}\right) = -J_{a,i}\prod_{j\in V(a)\setminus i}\theta\left(J_{a,j}h_{j\to a}\right) \tag{2.42}$$

The use of RS cavity method at zero temperature for the K-SAT model allows to reproduce exactly the same result of the more standard replica approach discussed in Section 1.

More interesting is the 1-RSB case. Figure 7 shows the evaluation of the average ground state energy density and of the average zero-energy complexity for a standard random instances ensemble with $K = 3$ and fluctuating variable-nodes connectivity.

Starting from the inspection of E_{GS} we note that it becomes finite already at $\alpha_c = 4.267$, in excellent agreement with the experimental data. We remark that E_GS can easily be evaluated by finding the maximum value $E_{GS} = \max_y[\Phi(y)]$. Before the SAT/UNSAT transition, the Bethe free energy is always negative and reaches zero only asymptotically for $y \to \infty$. At the critical point, $\Phi(y)$ suddenly develops a positive maximum for a finite value y^* of the pseudo-temperature, and the value of $\Phi(y^*(\alpha))$ provide the plotted ground state energy density curve.

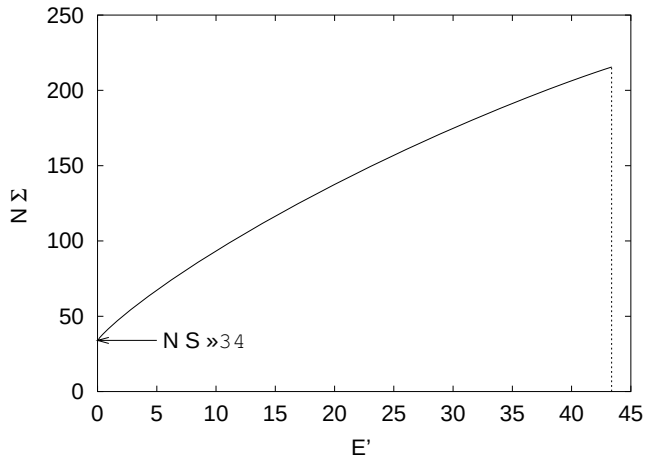

FIGURE 8. Complexity against energy density for random 3-SAT at $\alpha = 4.20$. Just before the SAT/UNSAT transition there is still an exponential number of clusters of solutions. However an exponentially larger number of states at excited energies exists, and it is supposed to be responsible of the slowing-down of search algorithms in the HARD region (adapted from [**109**]).

A second phase transition takes place at $\alpha_d = 3.921$, where suddenly the complexity of zero energy states jumps to a finite value. Then $\Sigma(E = 0)$ regularly decreases, vanishing exactly at α_c. Still, an exponential number of solutions exists exactly before the SAT/UNSAT transition, but they are all grouped in a single cluster. More remarkably, the number of solution clusters itself becomes exponential for α only slightly smaller than α_c.

Figure 8 shows that the complexity is an increasing function of the energy. If there is an exponential number of solution clusters, for $\alpha > \alpha_d$ there is also an exponentially larger number of clusters of metastable configurations at finite energy. This means that the probability of hitting a satisfying assignment is vanishingly small with respect to the probability of finding some local minima. When considering local search algorithms that explore the energy landscape of a CSP by performing only short-range moves if they are not too energetically unfavorable, it is quite likely that they will be stopped by glassy *threshold states* with an energy close to the threshold energy:

$$E_{th} = \operatorname{argmax} \Sigma(E) \tag{2.43}$$

There is indeed an exponentially larger number of clusters at energy density E_{th} than for any other lower value of energy. Such proliferation of

metastable states acting as entropic traps might provide a reasonable explanation for the observed EASY-HARD-EASY pattern in numerical experiments.

In Figure 7 the threshold energy evaluated from the 1-RSB complexity is plotted for comparison with E_{GS}. In reality, such E_{th} constitutes only an upper bound to the true energy scale of the threshold states, since the 1-RSB ansatz is known to become unstable [**118**] for any energy larger than the so-called *Gardner energy* E_G (plotted in Figure 9). Since, E_G is the largest energy for which the 1-RSB computations can be considered stable, the largest reliable estimated values of complexity lie on the line $\Sigma(E_G(\alpha))$. The Gardner energy can then be considered as a lower bound to the true threshold energy.

E_G crosses the 1-RSB ground-state energy density line in two points, $\alpha_G = 4.15$ in the SAT region and $\alpha_g = 4.39$ in the UNSAT region. Preliminary f-RSB corrections suggest that the true threshold states have energies very close to the lower bound. Consequently, only instances extracted from the interval $[4.15, 4.39]$ should be taken as really hard random benchmarks for algorithm testing. Furthermore, as displayed in the inset of Figure 9, the actual value of the energy gap $E_G - E_{GS}$ is very small close to α_G and α_g. In order to avoid systematic finite size errors, numerical simulations should be done close to the SAT/UNSAT point, *i.e.,* well at the center of the 1-RSB stable region. Consistently with the fact that finite size fluctuations are of order $(O(\sqrt{N}))$, even close to α_c problem sizes of the order at least of $N = 10^5$ might be needed in order to obtain a clear experimental evidence of the predicted entropic threshold phenomenon. Extensive experiments relative to the actual performance of local search heuristics in the hard region of the random K-SAT phase diagram will be the main topic of the next chapter.

3. A comprehensive summary of the K-SAT phase diagram

The identification of a SAT/UNSAT transition and the discovery of the ground-state clustering phenomenon are sufficient for explaining the main features of the observed typical-case complexity behavior of random 3-SAT. Nevertheless, a wide number of detailed theoretical investigations of the model has allowed to obtain a much richer description of its many phases, which can be summarized as follows:

For $\alpha < \alpha_D = 3.86$, the $T = 0$ phase is characterized by a zero ground-state energy. The entropy density is finite but the phase is Replica Symmetric (RS) and unfrozen. Roughly speaking, this means that there exists one

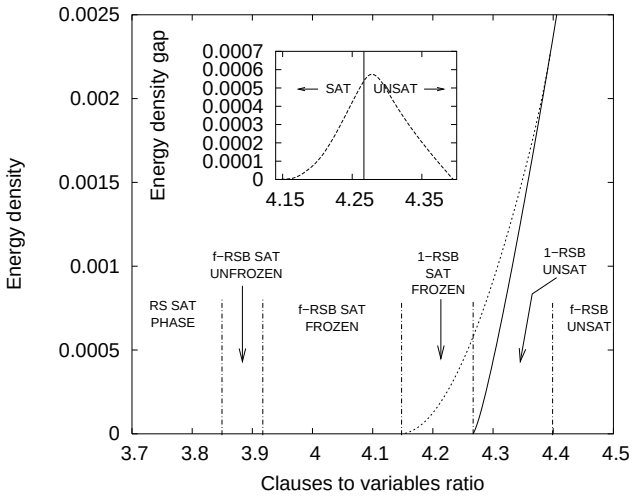

FIGURE 9. Summary of the random 3-SAT phase diagram. The solid line is an extimation for E_{GS}, while the dashed curve represents the Gardner energy E_G. (numerical data adapted from ref. [**118**]). In the inset we show that the gap $(E_G - E_{GS})$ is strictly positive only in the small 1-RSB stable region around the SAT/UNSAT transition critical point (indicated by the vertical line).

giant cluster of nearby solutions and that the effective fields vanish linearly with the temperature. This phase is SAT and EASY.

For $\alpha_D = 3.86 < \alpha < \alpha_d = 3.92$, there is a full RSB phase. The solution space breaks in clusters and the order parameter becomes a nested probability measure in the space of probability distribution describing cluster to cluster fluctuations. A peculiar property of this intermediate phase is the fact that the clusters have a vanishing backbone fraction. This is then a sort of "liquid" clustering scenario, where all the variables are unfrozen but present highly structured correlated patterns of fluctuations [**20, 133**]. This phase is SAT and EASY.

At $\alpha = \alpha_d \simeq 3.92$ there is a discontinuous transition toward a clustered and frozen phase [**107, 109**]. Up to $\alpha_g = 4.15$ the phase is f-RSB-like (a debate is open about the possible SUSY-breaking 1-RSB nature of this phase [**36**]) while for $\alpha_g = 4.15 < \alpha < \alpha_c = 4.2667$ the 1-RSB solution becomes stable [**118, 99**]. The 1-RSB complexity is finite in this region and its estimates are reliable, at least for $E < E_G(\alpha)$. The exponentially many 1-RSB metastable states become indeed unstable at the Gardner energy $E_G(\alpha)$, which constitutes a lower bound to the true dynamical *threshold energy*. This phase is SAT and HARD.

At $\alpha_c = 4.2667$ the ground state energy becomes positive and therefore the typical random 3-SAT problem becomes UNSAT. At the same point the complexity vanishes. The phase remains 1-RSB up to $\alpha_G = 4.39$ where an instability toward a zero complexity full RSB phase appears. This phase is UNSAT and HARD.

Finally, for $\alpha > \alpha_G = 4.39$, the phase is UNSAT and it is expected to be EASY, because of the absence of any sharp clustering.

A graphical sketch of this many phases is presented in Figure. 9.

CHAPTER 3

Local search fails

The general idea of local search is to explore the space of possible configurations initializing the search at some point and moving from there iteratively to a neighboring position, according only to information gathered looking at a local neighborood. An important advantage of the local search approach is its *universality*: a single algorithm can be applied to very different kinds of instances of the same problem, or even to different CSPs. Also the requirements in space resources are typically very small [**2, 130**].

A drawback is typically the *incompleteness* of the procedure, that is the impossibility to guarantee the finding of the best possible assignment even for satisfiable instances and, hence, the impossibility of providing UNSAT certificates. This limitation is essentially due to the non-systematic nature of the configuration-space exploration, which is exactly the most distinctive feature of local search itself. Another more fundamental flaw of the local search approach has already been pointed out in the introduction, and it's the possibility of being trapped by local minima and plateau regions of the search space [**14, 40**]. Nevertheless, until when brute-force exhaustive search will become a serious option, thanks to the introduction of some kind of superior hardware —but the present state-of-the-art in the implementation of both quantum and DNA-based computers seem to be rather deceiving [**122, 25**]— , local search will continue to be a valuable alternative to complete solvers, in particular if we are just interested in producing good suboptimal states in a short time and with reasonably small amount of memory.

In this chapter we will discuss the performance of some of the most performing local search heuristics known for the K-SAT problem, WalkSAT [**148, 88, 72**] and focused RRT [**145, 146**]. We will analyze briefly also the behavior of a recently introduced and quite "exotic" algorithm, based on techniques borrowed from quantum statistical mechanics, the so-called Quantum Annealing. The main conclusion will be, in both the classical and the quantum case, that the barrier created by the proliferation of the glassy threshold states typical of the hard phase cannot be overcome in reasonably short times by local search techniques.

Depending indeed on the models and on the details of the process —*e.g.,* the cooling schedule for SA— the long time dynamics will be dominated by different types of metastable states at different temperatures [**119**]. A common feature, confirmed by experiments, will be that at zero temperature and for simulation times which are sub-exponential in the size of the problem there exists an extensive gap in energy which separates the blocking states from the true ground states [**14**].

This substantial failure of local search will call for more elaborated optimization strategies, like the SP-based algorithms described in the chapters 4 and 5.

1. "Award-winning" heuristics

1.1. WalkSAT and its family. The first ancestor of the powerful WalkSAT algorithm is GSAT and was introduced in 1992 by Selman, Levesque and Mitchell [**147**]. GSAT was rather simple with respect to its further evolutions. We define the *score* of a variable as the variation in energy caused by flipping the variable itself. The main step consisted, after a random initialization of all the variables, in flipping exactly the variable with the highest *score*, choosing it among the ones involved in unsatisfied clauses. This greedy strategy suffered enormously the possibility of getting stuck in local trapping structures, and a large number of restarts was needed in order to generate good assignments. GWSAT (GSAT with *random walk*) introduced for the first time a *noise* parameter p [**148**]: at each step a random number r was extracted and, if $r < p$ a randomly selected variable appearing in unsatisfied clauses was flipped instead of the highest score one. The addition of the random walk feature gave to GWSAT the important property of *probabilistic approximate completeness* (PAC), shared with other heuristics based on metaphores derived from physics or biology (like simulated annealing or genetic algorithms): this means that the probability that GWSAT finds a solution converge to 1 when the number of allowed iterations before restart goes to infinity. Another mechanism proposed by the authors against trapping in local optima was *Tabu search*, that is preventing to flip again a variable before t_{tabu} steps are passed from its previous flipping. Note that while small values of t_{tabu} could not remove the need for restarts, too large values could hide the road to the true optima by forbidding too many moves. As a consequence, Tabu search may become fairly efficient is the tabu time is set carefully, but, as well as other history-based mechanisms, it cannot ensure alone the PAC property [**96, 59**].

WalkSAT [**148, 88**] is under many aspects similar to GWSAT, since it has a probability given by the noise setting p of performing a random

walk move, but it differs in the used scoring function (eligible variables continues to belong to unsatisfied clauses). The score of a variable i is given in WalkSat by the number of clauses that would be *broken* after an eventual flip, that is the number of satisfied constraints in which i is involved that become unsatisfied after inverting the value of i (when a greedy move is done, the variable with the lowest score is then selected). In addition to that, random walk moves are never performed if there are variables with score 0. This rather subtle differences are responsible for the largely superior performance of WalkSAT with respect to GWSAT. One may wonder which strategy is the most greedy, GWSAT or WalkSAT, and the answer is actually not evident. In a sense, WalkSAT is the greedier, since always flip a variable with null score if present. On the other hand, situations may happen where the move minimizing the number of broken clauses do not correspond to the choice that would have created the lowest total energy configuration: GWSAT could then be greedier choosing a flip that breaks many clauses but compensates this by satisfying a larger number of previously unsatisfied constraints.

Small refinements of the WalkSAT strategy have been proposed, like the introduction of Tabu search or novelty selection mechanisms [**88**], but no significant performance improvements have been induced. It seems, on the other hand, that an extremely careful optimization of the noise parameter p is needed in order to unleash the full power of the WalkSAT algorithm. The exact optimal value of p is strongly affected by the details of the instance under consideration or of the ensemble from which it is extracted, and it is then rather unpredictable without recurring to an exhaustive preliminary exploration of the possible settings. Optimal performance over the ensemble of random 3-SAT formulas is described in Fig. 1.

1.2. Focused local search. An important feature of WalkSAT was *focusing*, implemented by considering as eligible for flip just the variables belonging to violated clauses. The focusing property was already known to improve local search results, independently from the research related to WalkSAT [**128**]. Recent studies have addressed more specifically the efficiency of focused versions of Metropolis search (MS) and Record-to-record travel (RRT) [**145, 146**].

In the well-known Metropolis dynamics [**100**], a move is always accepted if produces a decrease in energy, and it is accepted with probability $\eta^{\Delta E}$ in the case of an increase. Typically the unit-increment acceptance probability is put in relation with a formal temperature $\eta = \exp[-1/T]$. If the temperature is slowly reduced according to some predefined schedule, Metropolis search is known under the name of *simulated annealing* [**82, 1**].

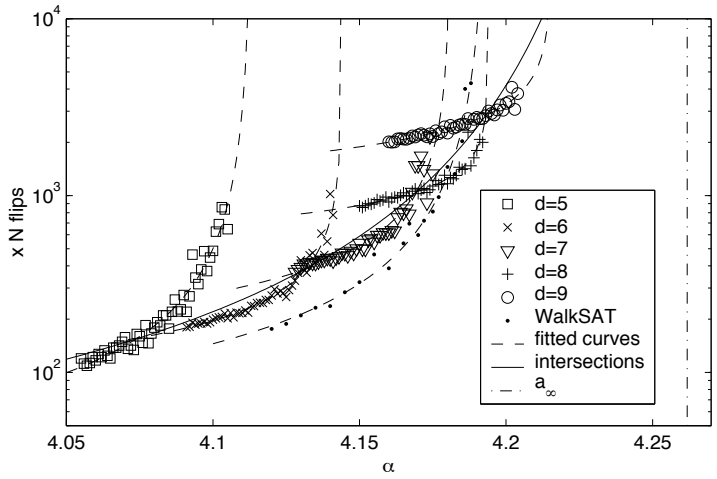

FIGURE 1. Performances of WalkSAT and fRRT. The number of flips per variable needed to find a complete solution is plotted against α for WalkSAT (optimal noise setting $p = 0.55$) and for fRRT with several acceptance spans d. For some settings solution can be found close to α_G, even slightly inside the hard region, but resolution times show a super-exponential divergence with α. From a power law fit for the location of the intersection of the curves for increasing d, an estimate a_∞ for the location of the critical point can be obtained in good agreement with previously established results (figure reproduced from [**145**]).

In Record-to-record travel [**47**], on the other hand, a move is accepted if and only if it produces a configuration with energy not larger than $E_{\rm record} + d$ where $E_{\rm record}$ is the best energy found so far.

The focused versions fMS and fRRT are obtained from the original ones just by allowing to select for spin-flip only the variables which belong to already violated clauses. The addition of focusing destroys the detailed-balance property of ordinary Metropolis search but boosts its efficiency in a remarkable way. While in the next section we will see that a simple implementation of simulated annealing is unable even to approach the Gardner lower bound to the glassy threshold states, fMS with fixed acceptance probability η is able sometimes to find complete solutions even inside the 1-RSB SAT region of random 3-SAT (although very close to the 1-RSB instability point $\alpha_G \simeq 4.15$). The performance of fRRT is only slightly better than that of fMS, and some results are shown in Fig.1, in comparison with WalkSAT data (obtained with optimal setting of the noise parameter p. Once again the setting of the acceptance span d strongly affect the observed behavior.

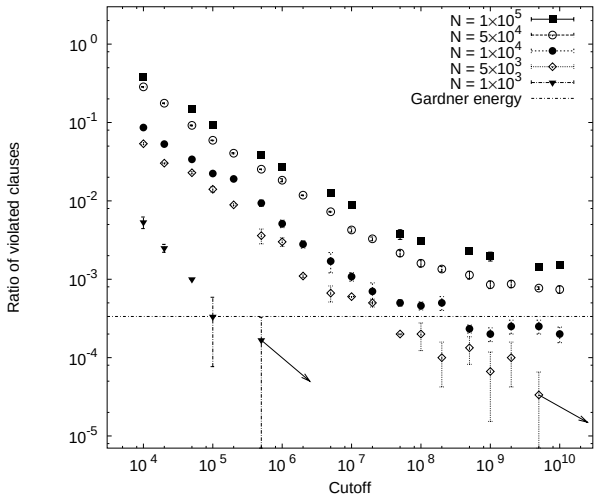

FIGURE 2. Threshold energy effect in SAT region. The WalkSat performance for various samples of different sizes and $\alpha = 4.24$ is presented. With increasing size, the curves appear to saturate above the Gardner energy. An arrow indicates that the next data point corresponds to a SAT assignment.

1.3. Evidence of the glassy threshold effect. The results presented in Fig. 1 and discussed in [**145**] clearly show divergence of resolution-times when entering the hard 1-RSB region. In particular, there are no published data about the resolution of large random instances with $N > 10^5$ and $\alpha \gtrsim 4.21$ using just local search techniques. The theory exposed in the preceding chapter predicts for random 3-SAT the existence of a SAT phase characterized by stable 1-RSB clustering between $\alpha_G \simeq 4.15$ and $\alpha_c \simeq 4.26$. In this region, the trapping effect into an exponential number of non-ergodic components of the phase space corresponding to clusters of clause-violating assignments should produce an exponential slowing-down of any local dynamics respecting detailed balance. Also WalkSAT and fRRT, although non respecting detailed balance in a proper sense, meet growing difficulties in finding solutions when N is increasing, and the average resolution time diverges in the thermodynamic limit.

In the case of WalkSAT precise theoretical results exist [**149, 11**], predicting that for α larger than some dynamical point new SAT clauses are produced at the same rate at which they are eliminated, creating a sort of energy freezing. Similar effects have been put in evidence by experiments on fRRT. Only exponentially rare fluctuations can allow to find a solution [**120, 14**]. Considering the small amplitude of the Gardner energy gap [**118**], the probability of such lucky fluctuations should be considered

safely vanishing in the hard 1-RSB region only for $N \simeq 10^5$ or larger [1]. An extensive set of experiments was carried out to check the actual existence of the threshold effect [**14**]. We ran WalkSat over different formulas in the hard-SAT region, with fixed $\alpha = 4.24$ (safely in the 1-RSB SAT region) and sizes varying between $N = 10^3$ and $N = 10^5$, reaching a maximum number of 10^{10} spin flips. The results are shown in Fig. 2 where the Gardner energy line is also reported for comparison. For small-size samples large fluctuations still allow the local search algorithm to find a SAT assignment. For larger formulas ($N \sim O(10^4)$) WalkSat invariantly fails in solving the proposed instances in the maximum allowed number of flips. The relaxation profile suffers of critical slowdown, and saturation at some well defined level is clearly visible. For the $N = 10^4$ formulas the saturation plateau is below the Gardner energy. The finite-size fluctuations are still of the same order of the energy gap between the ground and the threshold states and the experimental conditions are distant from the thermodynamical limit. When the size is increased up to 10^5 variables, the saturation level moves finally above the lower bound to e_{th}.

Since WalkSat does not respect detailed balance, we cannot conclude that its saturation level coincides with the precise threshold energy scale of the sample in exam. We can nevertheless claim that WalkSat is unable to explore the full energy landscape of the problem, without the help of unlikely strong fluctuations.

2. Quantum heuristics

It is generally claimed that an hypothetical Quantum Computer could exploit the peculiar properties of quantum states can in order to encode and manipulate information in a way quantitatively superior to classical computers [**122**]. A quantum computation is usually seen as a suitable sequence of unitary transformations acting over sets of *qubits* and maintaining the quantum state of a system under control for a sufficiently long time is an extremely demanding task from the point of view of practical implementation.

Most of the quantum algorithm proposed in the literature are based on an abstract model, the quantum Turing Machine, which have many points of contact with the models describing classical computing machines. In particular, the representation of information is digitalized, even resorting to complex error-correcting schemes in order to defeat the many sources of

[1] The Gardner energy becomes $O(1)$ only for $N \sim 10^4$ or larger, and for a smaller number of variables the threshold effect should be negligible when compared to finite size effects.

noise and decoherence that can destroy the fragile and essentially analogical nature qubits.

A different approach to quantum computation, more similar in spirit to local search heuristics, is taken by Quantum Ground State Search methods, like Quantum Adiabatic Computation and its variants [**51, 52, 53, 75**] and the semi-classical simulated Quantum Annealing algorithms [**143, 94, 95, 16, 17**]. These techniques work simulating (eventually with a quantum hardware, universal or dedicated) the externally driven evolution of an artificial quantum system, whose ground states are in a one-to-one mapping with the solution space of specific computational problems. Much of the enthusiasm about (simulated) quantum methods relies on the hypothesis that a a quantum system can dynamically explore the state space more efficiently than its classical counterpart. The positive effect of quantum phenomena has been verified in some cases, notably in relaxation experiments performed over amorphous anneals with magnetic impurities [**30**] and in numerical simulations of the Transverse field quantum Ising Spin Glass [**143, 94**] which can be considered a toy-model of the same class of materials. The success of the Path Integral Montecarlo (PIMC) scheme for Quantum Annealing pushed the authors to apply it to the optimization of hard instances of the Traveling Salesman Problem [**95**]. The results were encouraging, since PIMC-QA was able to outperform a simple implementation of classical Simulated Annealing: in particular, the convergence time was significantly shorter, even after taking in account the extra cost of simulating a quantum computation on a classical computer [**143, 95**].

Nevertheless, very few theoretical results are known to date [**52, 159, 110, 152, 153, 17**] and, In particular, it has not been proved that a good quantum optimization pathway must exist in general. Even in the cases in which the experiments show that the quantum strategy is more efficient than the classical analogue, no rigorous and definitive explanations for the superior quantum performance cannot be given.

In this section, we will discuss briefly the optimization of a hard instance of random 3-SAT with the PIMC-QA, in comparison with a basic implementation of the ordinary Simulated Annealing (SA) [**82**]. We do not expect that SA is particularly performing, since it is just a variant of Metropolis Search without the focusing feature, but one may hope that the introduction of a source of quantum fluctuations induce tunneling phenomena, able to overcome the energy barriers which isolate individual local minima. We expect, nevertheless, that this eventual feature is not particularly effective in a landscape punctuated by an exponential number of such traps: leaving one trap, would correspond to escape into another one and the

ground-state valleys would continue to be hidden because of mere entropic effects, exactly like in the classical case.

The hardness of the chosen 3-SAT instance deserves a further comment. It is common in the quantum computation literature to discuss the performance of newly introduced algorithms by testing them on simplified toy-problems of reduced size; this is essentially due to the demanding memory and time required by the classical simulation of a quantum computer. We applied instead PIMC-QA to the same hard samples used as benchmarks for WalkSAT and fRRT. The observed results will reflect then the typical behavior of the semiclassical algorithm applied to the optimization of a glassy 1-RSB landscape in the thermodynamic limit. As we shall see, the quantum pathways tend to visit basins of attraction with a considerably larger number of flat directions in order to minimize "kinetic" energy and this feature will turn out to be counterproductive for an efficient optimization of the 3-SAT landscape (unlikely the previously analyzed instances of Ising Glass and TSP).

2.1. PIMC-QA. In ordinary Simulated Annealing the configuration of the system is updated according to a Metropolis dynamics with unit-increment acceptance probability $\eta_t = \exp[-1/T_t]$. Here T_t is the temperature at time t. Temperature is gradually reduced according to a specific schedule that can be rather influent on the performance of the algorithm (see, for instance, Refs. [**1, 127**]). In order to make simpler the comparison with QA we will use just relatively inefficient linear schedules. We want indeed to stress our attention on the distinctive features of QA with respect to SA, more than achieving a good performance.

In order to introduce quantum fluctuations we substitute into the classical 3-SAT hamiltonian the Ising degrees of freedom with quantum SU(2) spin operators (oriented along the axis z) and we add a transverse field term (oriented alog the axis x):

$$\mathcal{H}_{\text{Quantum}} = \mathcal{H}_{\text{3-SAT}}\left(\left\{\hat{\sigma}_i^z\right\}, \{J_{a,i}\}\right) - \Gamma \sum_i \hat{\sigma}_i^x \tag{3.1}$$

The basic idea of Quantum Annealing is to drive the system (3.1) by varying adiabatically the intensity of the transverse field. A linear schedule is chosen for simplicity, varying between two extrema Γ_0 and Γ_f.

In the Path Integral Monte Carlo (PIMC) scheme, the statistical-mechanical behavior of the quantum model (3.1) is semi-classically approximated by mean of the Suzuki-Trotter transformation [**154, 37**]:

$$\mathcal{H}_{\text{ST}} = \frac{1}{P}\sum_{\rho=1}^{P} \mathcal{H}\left(\left\{S_{i,\rho}\right\}\right) - J_\Gamma \sum_{\rho=1}^{P}\sum_i S_{i,\rho} S_{i,\rho+1} \tag{3.2}$$

The Suzuki theorem ensures that when P goes to infinity the partition functions of the Hamiltonians (3.1) and (3.2) become identical, and the statistical-mechanical properties of the two systems become perfectly equivalent. The relaxation dynamics of the model (3.1) with N degrees of freedom can then be hopefully simulated by mean of a classical Metropolis dynamics over the model (3.2) with PN classical degrees of freedom, and P as large as possible. The system (3.2) is actually composed by P replicas (Trotter replicas) $\left\{S_{i,\rho}, \rho = 1 \cdots P\right\}$ of the original classical configuration $\{S_i\}$ at an effective quantum temperature $T_q = PT$, coupled among them by a nearest-neighbor transverse ferromagnetic coupling J_Γ. Periodic boundary conditions must be imposed along the transverse (Trotter) direction. The intensity of J_Γ is linked to the strength of the quantum perturbation term:

$$J_\Gamma = -\frac{T}{2} \ln \tanh\left(\frac{\Gamma}{PT}\right) > 0 \,. \tag{3.3}$$

The principal drawback of the PIMC approach is the impossibility of simulating the quantum system directly at zero temperature. For this purpose, other schemes should be developed, resorting for instance to the Green Function Monte Carlo method [**157**]. The simple PIMC-QA with linear schedule proved nevertheless to be very effective in the optimization of ising glass and TSP instances.

Let us describe now a first set of numerical experiments, performed over random 3-SAT instances with $N = 10^4$ and $\alpha = 4.24$, (*i.e.*. Empirical tuning of parameters was performed in order to determine the optimal setting for both the SA and QA linear schedules (for details, see [**16, 17**]). A comparison between the performance of the optimal SA and the optimal QA at $P = 50$, is shown in Fig. 3 at left. The graph shows the average residual energy (difference between the final energy reached by the algorithm and the known ground-state energy) as a function of the maximum number of total annealing iterations, ranging over several decades (from 10^2 to almost 10^6). Each point represents the average over 50 different realizations of the same experiment; in the case of QA, a second average was performed among the energies of the P replicas, which are in general different. It can be seen that the linear-schedule CA is everywhere more effective than the simple linear-schedule QA, simulated with $P = 50$ (white circles). Better performances would be produced by the use of a larger number of replicas (no further energy reductions are observed if $P \gtrsim 100$, see the inset of Fig. 3, the dashed line is the average result of SA of the same length), but a tradeoff is needed between accuracy of the simulation and need of extending as possible the total number of iterations. A slight improvement can be observed if *global moves* are performed in addition to ordinary local

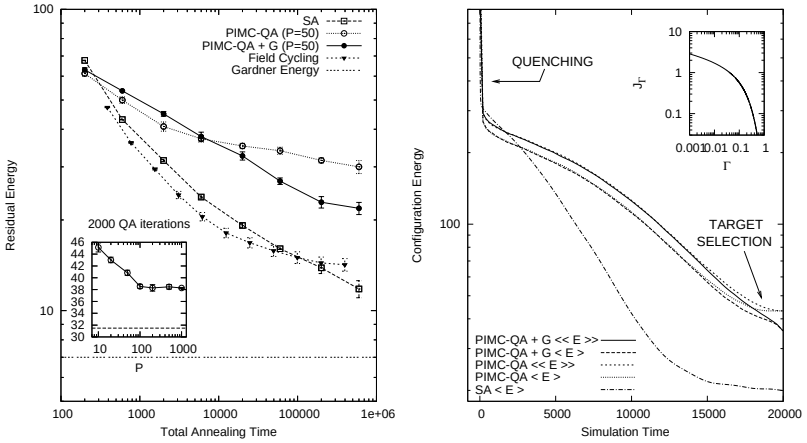

FIGURE 3. (*Left*) Comparison between optimal linear-schedule Classical (SA) and Quantum Annealing (QA). SA always performs better than QA simulated with $P = 50$ Trotter replicas. The solid triangles refer to data obtained by a field-cycling QA strategy. (*Right*) Evolution during QA of the averages $\langle E \rangle$ (average best replica) and $\langle\langle E \rangle\rangle$ (average of the average replica), compared with SA.

ones [**125**]. In a global move a same variable is flipped simultaneously in all the replicas, instead of independently. Global moves do not affect the quantum contribution to the energy and implement then a sort of "classical drift" of the system superposed to the quantum fluctuational dynamics. Despite their use, the asymptotic slope of linear-schedule QA cooling curves remains definitely less steep than that of CA, unlike the previously studied cases [**143, 95**].

In order to identify the causes of the bad quantum performance, we need a better characterization of the QA dynamical properties. We denote by $\langle\langle E \rangle\rangle$ the average over different experiments *and* over Trotter replicas of the configuration energy (this is the quantity plotted in Fig. 3); we denote, on the other hand, as $\langle E \rangle$ the average among different experiments of the *best* replica energy. In Fig. 3 on the right, the evolution profiles of $\langle\langle E \rangle\rangle$ and $\langle E \rangle$ are shown for a 2000-iterations-long linear-schedule QA. The strength of the quantum coupling (3.3), depending by the transverse field intensity (see inset of Fig. 3 on the right), determines the relative weigth of the classical and quantum terms in the Hamiltonians (3.1) and (3.2). In the very initial part of the simulated relaxation, the coupling J_Γ is small and each replica behaves almost independently from the others, according to P parallel Monte Carlo dynamic processes at a constant temperature. The replicas assume different configurations but at a similar energy scale (corresponding to the

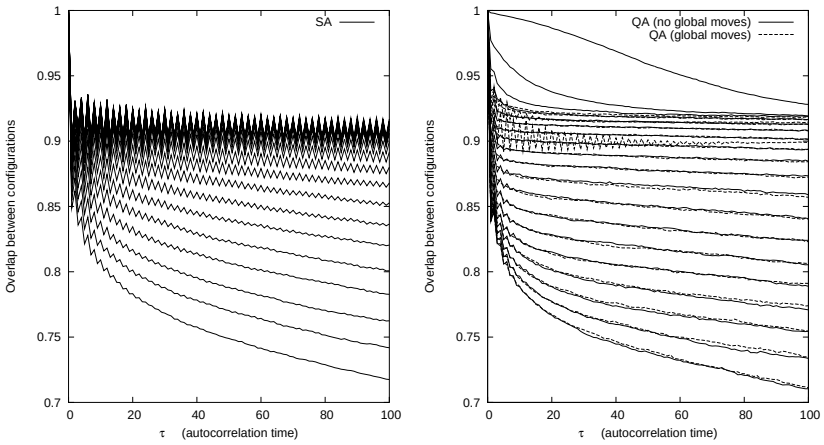

FIGURE 4. Autocorrelation functions $K(t^*, \tau)$. The different curves represent the decay with τ of several fixed-simulation-time snapshots of the autocorrelation function, for a SA experiment (*Left*) and for a QA experiment (*Right*).

states present in the incoherent mixture describing the quantum system at the beginning of the experiment), linked to the chosen fixed temperature T_q. When the coupling strength increases (see the inset of figure 3 on the right), the fluctuations of the different replicas become correlated. Some spin flips that would have been unlikely in absence of the replica-interaction term can be accepted, thanks to the kinetic term, and several replicas can visit configurations not typical for ordinary CA trajectories. When the transverse field vanishes, quantum fluctuations are gradually switched off, and the system seems to collapse completely into some low-energy suboptimal target state. The diverging value of the quantum coupling, essentially freezes the evolution of the configuration. The acceptance of local moves falls to zero and only small-range purely classical oscillations can take place if global moves are used.

Autocorrelation analysis can be carried out in order to visualize the described scenario. We define the autocorrelation function $K(t, \tau)$, for both SA and QA, as:

$$K(t, \tau) = \left\langle \frac{1}{N} \sum_{i=1}^{N} S_i(t) S_i(t - \tau) \right\rangle , \tag{3.4}$$

where an average over different realizations of the dynamics (and over replicas in the QA case) has to be understood. A fast decay of $K(t^*, \tau)$ with τ for fixed t^* indicates that at time t^* the spin configuration is still rapidly evolving, while a flat (or periodic) profile suggest that local stability has been reached and the system has entered into some attracting configuration (or

limit cycle). Autocorrelation plots are shown in fig. 4. The different curves represent the decay with τ of several fixed-simulation-time snapshots of the autocorrelation function, for SA experiments and for QA experiments with and without global moves; t^* is varying at fixed intervals between 200 and 2000, from bottom to top. In SA, when the temperature becomes too small the system is blocked in a restricted region of the phase space and only a limited subset of variables can undergo a cyclically repeated sequence. The target selection phase in QA is also clearly visible, characterized by damped classical oscillations around an equilibrium position or by a complete collapse into it, depending if global moves are allowed or not.

2.2. Field-cycling. A different way of describing the three-stage evolution (quenching, quantum search, target selection) of the spin configuration under QA borrows some terminology from the field of Genetic Programming [**87, 68**]. The P replicas can be seen as a population of individuals, carrying a genotype coded in the spin configuration. The classical Hamiltonian (2.2) provides a measure for the fitness of the individuals, and for the average fitness of the population. Contiguous replicas can "mate", exchanging sequences of their genotype thanks to the duplicating action of the transverse coupling. Good "schemata" (spin patterns common to many low-energy replicas) proliferates [**69**] and a global decrease of both the population and the best individual average fitness is observed. At this point the coupling has become rather strong. Even if some new exceptionally fit individual appeared by chance thanks to some mutation, the rare patterns accounting for its better energy would be overwritten with large probability by the corresponding sequences in the most common genotype. The population tends then at a certain point to collapse toward a group of "identical twins". Global moves can delay this fatal consanguineous interbreeding, but not avoid it.

In standard genetic algorithms (and in nature), crossover takes care of constantly renewing the population gene pool, but now we want to stitch faithfully to another metaphor, the quantum one. A trick allowed by the rules of quantum simulation consists in increasing the mutation rate by switching on and off repeatedly the quantum fluctuations, by mean of *field cycling*.

After the final target selection, the field Γ is raised again to the initial value, and a new descending ramp is started. In order to avoid a complete reinitialization of the configuration (analogous to the rejuvenation effects discussed in [**23, 24**]), before each new consecutive field ramp the temperature is slightly decreased in a stepwise manner (see the upper panels of Fig. 5). The results of the optimization with the best field-cycling piecewise-linear

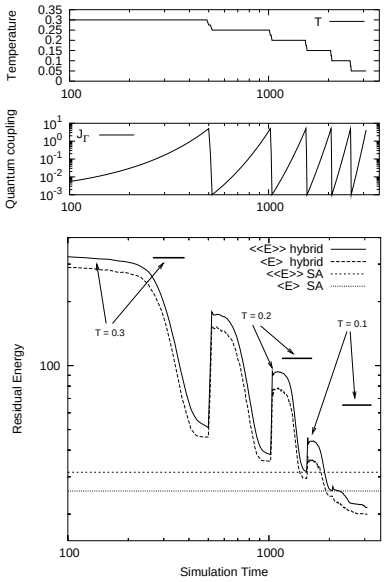

FIGURE 5. Energy evolution during a field-cycling hybrid strategy. Each ascending field-ramp unfreezes the system from a previously reached target state, and the target energies becomes increasingly smaller (and more distant from the quenching energy level associated to the temperature during the ramp)

.

QA schedule are shown in Fig. 3 on the left (black triangles, $\Gamma_0 \simeq 0.7$ to $\Gamma_f \simeq 0.001$, T_q reduced from 0.3 to 0.05). The same experiment has been repeated in absence of the transverse magnetic field, and with the same number of (now completely decoupled) Trotter replicas. The average over runs and replicas $\langle\langle E\rangle\rangle_{\rm SA}$, and the average over runs only of the best replica energy $\langle E\rangle_{\rm SA}$ have been computed. It is possible to see from Fig. 5 that $\langle\langle E\rangle\rangle_{\rm cycling}$ lies well below $\langle E\rangle_{\rm SA}$. The role played by the quantum coupling is then determinant and we can affirm that quantum restarts are more effective than classical ones [**120**], at least when short schedules are taken in account. The absence of temperature reinitialization is crucial for the survival of "seed patterns", peculiar of the target configurations obtained at the end of the i-th descending ramp, from which new and better low-lying configurations can be grown.

Over short time scales, the field-cycling strategy performs definitely better than linear-schedule SA, but asymptotically, the slope of its cooling curve reduces again to the one of simple linear-schedule QAs. Despite the efficiency of field-cycling in renewing the available gene pool without

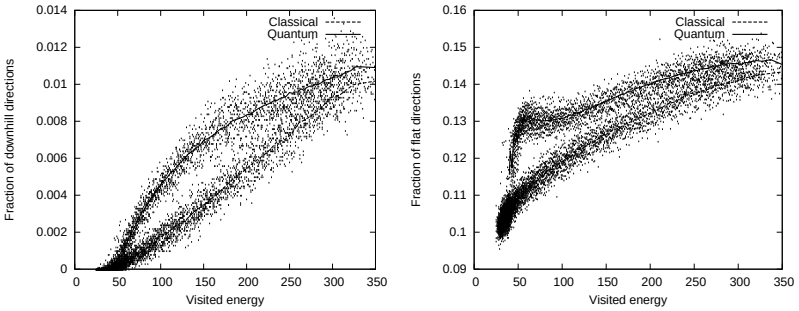

FIGURE 6. The local geometry of the visited regions of the phase space is probed by counting the fraction of directions in which the energy variation is negative, null or positive. Although both SA and QA get trapped in a local minimum, the quantum evolution tends to visit "valleys" that are more flat and with a larger number of downhill directions.

destroying completely the highly fit genes of the preceding generations, the deep reason of the QA failure must be researched elsewhere.

2.3. Landscape probing. Which kind of states are actually blocking the dynamics and are being visited during the simulated relaxation? Which is the geometry of the explored phase-space regions independently from the sampled energy values? To answer these and other questions, we looked at a local geometrical description during the dynamics.

When sitting in a precise given configuration of energy E, we counted the number of directions in which the energy variation was positive (uphill directions), null (flat) or negative (downhill). We show in Fig. 6 the data relative to the configurations sampled by a large number of linear-schedule SA and QA trials over the same hard $N = 10000$ sample (each run was 2000 iterations long). The fractions of downhill and flat directions are plotted against the energy of the visited configurations, for both SA and QA.

In the case of SA, the number of downhill directions falls to zero in correspondence of the lowest found energies, indicating that the SA dynamics ends in a local minimum. The number of remaining flat directions accounts exactly for the observed amplitude of the cycles shown in Fig. 4 for large t^*. When considering precise values of energy reached both by SA and QA, the quantum system configuration is characterized always by a significantly larger fraction of downhill and flat directions. If an abuse of language is tolerated, we could summarize by saying that SA optimization follows narrow

canyons, while QA prefers to take pathways exploring the edges of mid-altitude plateaus. This behavior, which is likely to be a genuinely quantum feature captured by the PIMC simulation, is strongly reminiscent of what happens in continuous space, where the choice of broader potential wells allow the system to reduce the kinetic contribution to the total energy (curvature-induced effects are also well known in the theory of instanton tunneling [**84**]).

The number of flat directions is comparable with the typical distance between the configurations of different replicas. All the states simultaneously visited by the quantum system belong then to a single valley explored in all its wideness. No quantum tunneling is then being simulated, but just a quantum mechanism responsible of the selection of a broad valley. When the field becomes too small, the QA dynamics enter into a classical regime, and starts looking for the the kind of local geometry preferred by SA. The dynamical collapse is indeed paralleled by an abrupt increase of the number of uphill directions. The existence of dead-end sinks, scattered in the broad basins of attraction selected in the quantum regime, which inexorably attract the quantum system at the transition to the classical regime like a "poisoned bait", could then be a possible explanation of the poor performance of quantum optimization in the random 3-SAT landscape[2].

[2]The presence of multiple and scattered wells is also known to strongly perturb genetic-like search [**43**].

Part 2

Survey Propagation and beyond

CHAPTER 4

Optimization with SP and SP-Y

The previous chapter described the failure of the more sophisticated local search approaches to the optimization of large random K-SAT instances in the hard phase. These negative results call for the development of radically different algorithms able to overcome the barrier created by the proliferation of clusters of metastable states.

1. Survey propagation

We have seen in Chap. 2 that the Cavity method allows to re-derive the phase diagram of many combinatorial optimization problems, including K-SAT. This particular technique, unlike the replica approach, allows to take the average over the quenched disorder only at the end of the computations. It becomes then possible to obtain information of algorithmic relevance about the statistical behavior of single variables in the stable and metastable states of any given energy density. The knowledge of local magnetizations will be helpful in the retrieval of configurations lying well beyond the glassy thresholds [**26, 14, 15**].

1.1. Message-passing interpretation of cavity equations. Let focus for simplicity on the case of the random K-SAT. We remind that every instance of the problem can be pictorially represented by a factor graph, in which circle nodes are associated to spin variables $i, j, k, \ldots$ and square nodes to K-body interactions, labeled by $a, b, c \ldots$, as shown in Figure 1 of Chapter 2. The locally tree-like structure of the factor graph allowed us to express the partition function and, therefore, the entropy, complexity and free-energy functionals in terms only of the probability distribution functions of the cavity biases $u_{i\to a}$, and of the cavity fields $h_{j\to a}$ and local fields H_i, depending on them. For the K-SAT model, we were able to write in the zero-temperature limit, of interest for optimization:

$$H_i = \sum_{b\in V(i)} u_{b\to i}, \quad h_{i\to a} = \sum_{b\in V(i)\setminus a} u_{b\to i} \tag{2.14}$$

$$u_{a\to i} = \hat{u}_{a\to i}\left(\left\{h_{j\to a}\right\}\right) = -J_{a,i} \prod_{j\in V(a)\setminus i} \theta\left(J_{a,j} h_{j\to a}\right), \tag{2.42}$$

After the removal of a variable i and of the interactions $b \in V(i)$ form the

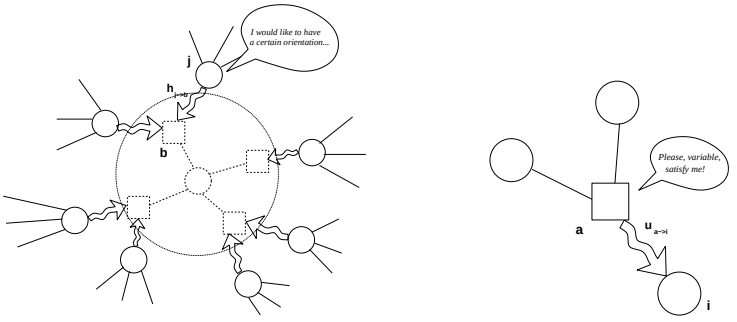

FIGURE 1. Message-passing interpretation. (*Left*) The variables on the surface of the cavity communicate their tendency to the clauses connecting them to the cavity variable; (*Right*) A clause send a warning signal to a variable if it cannot be satisfied by the other neighbors.

factor graph, the variables j on the surface of the cavity becomes uncorrelated and local energy minimization can be achieved simply by aligning them with the cavity fields $h_{j\to b}$:

$$s_j = \operatorname{sign} h_{j\to b} \tag{4.1}$$

The cavity fields $h_{j\to b}$ can then be interpreted as *messages* traveling from a variable j to a clause b, informing the clause b about its preferred orientation in absence of the constraint b itself (see Fig.1).

A cavity bias $u_{a\to i}$ becomes on the other hand different from zero only when for every $j \in V(a) \setminus i$ the cavity fields $h_{j\to a}$ have the same sign of the corresponding coupling $J_{a,j}$. Suppose that all the variables $s_j \in V(a) \setminus i$ are set according to the cavity fields $h_{j\to a}$. By substituting then (4.1) into the expression (2.1), one sees that the clause a can be satisfied only by setting $s_i = u_{a\to i}$. It would be then possible to interpret the cavity bias $u_{a\to i}$ as a message sent by the clause a pushing i in the direction required for solving an eventual contradiction (see Fig.1).

The sum of all the cavity biases converging into a single variable measures the local polarization induced by all the constraints actually relevant. It is indeed important to remark that the cavity biases (2.42) can be null. The magnetization of a variable i is given by:

$$m_i = P(H_i > 0) - P(H_i < 0) = W_i^+ - W_i^- \tag{4.2}$$

but a null magnetization can be achieved in two quite distinct ways. One possibility is that $W^+ = W^-$: in this case the variable would be forced by an equal number of constraints in opposite directions. Another possibility is that $P(H_i = 0) = 1$: this variable would be actually left completely free by all the clauses to which belongs, since they are already satisfied by the

assignments of other variables. Not all the variables are then completely *frozen* and not all the constraints equally effective. The existence of *unfrozen* variables is actually one of the most relevant manifestations of the clustering phenomenon. Our formalism will take explicitly into account the quantities $W_i^0 = P(H_i = 0)$ and this constitute a main difference with other message-passing algorithmic approaches like Belief Propagation (see Appendix D).

If we had not obtained the equations (2.14) and (2.42) from the assumptions of the cavity method in the optimization limit, we could have introduced them as defining dynamical laws for two coupled sets of fields, the "h-messages" and the "u-messages". We have indeed shown that, independently from their derivation, they carry extremely meaningful information along the edges of the factor graph. Let consider now a single solution $\vec{\sigma} = (\sigma_1, \sigma_2, \ldots, \sigma_N)$ of a given instance of K-SAT, and let artificially initialize the h-messages to $h_{i\to b} = \sigma_i$ for every i and every b. This initial set $\{h\}_0$ is used to compute the set of all u-messages $\{u\}_0$, that determines another set of h-messages $\{h\}_1$, and so on. This procedure is known under the name of *warning propagation* [**26**]. It is easy to see that at every instant t the variable configuration $\vec{s}_t = (\text{sign}\, H_{1,t}, \ldots, \text{sign}\, H_{N,t})$ would be identical to $\vec{\sigma}$ or would belong *to the same cluster* of it, that is, would differ from $\vec{\sigma}$ only in variables that were completely or almost unfrozen. Solutions in different clusters would be associated to different fixed points of the warning propagation equations (2.14) and (2.42). The same fixed point would correspond to all the solutions in a cluster, since unfrozen variables can freely be set.

Coming back to discussion of the cavity equations, the order parameter for a single specific instance of the K-SAT hamiltonian is characterized by the collection of all the probability distribution functions (pdf) of the cavity-biases (2.42). We remind that these pdfs are coupled between them by the following set of integral equations, in which no averages over the quenched disorder have to be performed:

$$P_{j\to a}(h_{j\to a}) = C_{j\to a} \int \mathcal{D}Q_{j,a}\; \delta\Big(h - \sum_{b\in V(j)\setminus a} u_{b\to j}\Big)\, \chi_{j\to a}\Big(\{u_{b\to j}\}_{b\in V(j)\setminus a}\Big) \tag{2.32}$$

$$Q_{a\to i}(u) = \int \mathcal{D}P_{a,i}\; \delta\left(u - \hat{u}_{a\to i}\left(\left\{h_{j\to a}\right\}\right)\right), \tag{2.33}$$

Analogue equations could be written for the local field pdfs, like in section 2 of the Chapter 2.

In terms of the message-passing interpretation, the pdf of a cavity bias $u_{a\to i}$ could be interpreted as an "opinion survey" among the different clusters about the value of the u-message sent along the edge $a \to i$. The reweighing term can be seen now as an energy filter:

$$\chi_{j\to a}\Big(\{u_{b\to j}\}_{b\in V(j)\setminus a}\Big) = \mathrm{e}^{y\left(|\sum_{b\in V(j)\setminus a} u_{b\to j}| - \sum_{b\in V(j)\setminus a}|u_{b\to j}|\right)} \tag{2.34}$$

It is indeed easy to see that sending $y \to -\infty$ only the configurations of messages associated to zero energy clusters can give a non null contribution to the integral (2.32)[1]. It becomes then possible to write formally:

$$Q_{a\to i}(u) = \frac{1}{\mathcal{N}_{\text{clusters}}} \sum_{\ell} \delta\left(u, u^{\ell}_{a\to i}\right) \tag{4.3}$$

where we suppose that there are $\mathcal{N}_\ell$ different solution clusters and that $\{u^{\ell}_{a\to i}\}$ is the fixed-point set of u-messages associated to the ℓ-th cluster. In the following, we will use the name of "u-survey" for the pdf's of u-messages.

1.2. Population dynamics. The relations (2.32) and (2.33) (asymptotically exact under the hypothesis of large tree-like instances) constitute a closed set of self-consistent equations for the pdf's that can be solved by iteration. In the case of the K-SAT problem, a u-survey can be parameterized by just one real number η, giving the probability of warning emission:

$$Q_{a\to i}(u) = \eta^{0}_{a\to i}\delta(u) + \eta_{a\to i}\delta(u + J_{a,i}) \tag{4.4}$$

where $\eta^{0}_{a\to i} = 1 - \eta_{a\to i}$. This compact functional form simplifies the algorithmic resolution of the equations. The full functional order parameter $Q[Q(u)]$ can then be described by a single function giving the distribution of the parameters η. A u-survey $Q_{a\to i}(u)$ is attached to every edge $a \to i$, from a function node a to a variable node i. From these u-surveys one can compute the cavity fields $h_{i\to b}$ for every $b \in V(i)$, which in turn determine new output u-surveys (see Fig. 2). Starting from a random initial population of u-surveys $Q_{a\to i}(u)$ (an initial population of η's), the function nodes are selected sequentially at random and the u-surveys updated according to (2.32) and (2.33). At each step of the iteration several convolution integrals have to be computed. This can always be done in a relatively efficient way, by a fast recursive procedure, described in Section 3.1. Further simplifications, described in Section 2 are possible when we are interested in the SAT region and the $y \to \infty$ limit can be taken. After several iterations of this "survey propagation" procedure, a fixed-point is typically reached. The self-consistent list of the u-surveys can then be used to compute the normalized pdfs of the local fields acting on each variable, and, hence, the local

[1]We will see in the following sections how the value of the Lagrange multiplier y can be set in order to slice group of states at different energies

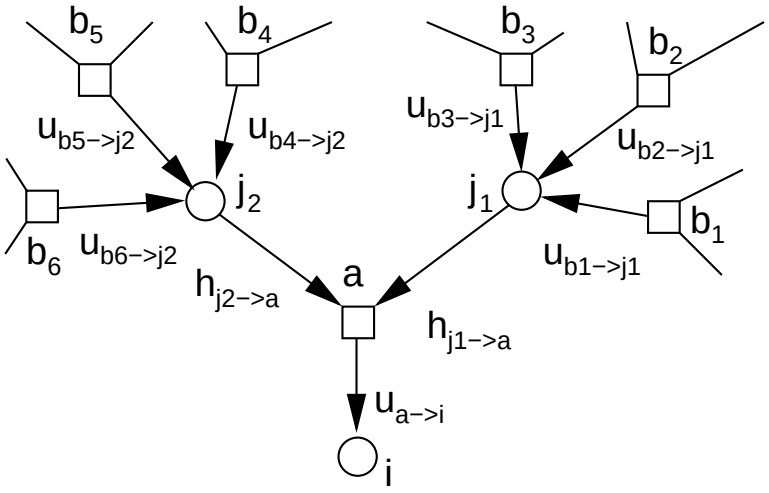

FIGURE 2. Cavity fields and Cavity biases. The u-survey for the u-message $u_{a\to i}$ depends on the pdf's of the cavity fields $h_{j_1\to a}$ and $h_{j_2\to a}$. These are on the other side dependent on the u-surveys for the u-messages incoming to the variables j_1 and j_2.

magnetizations and other functionals of interest, like Bethe free-energy and complexity.

2. The limit $y \to \infty$ of SP iterations

A first algorithmic application of the microscopic information obtained by the SP self-consistent computation is the resolution of satisfiable instances. In this section we will present results relative to the 3-SAT problem, but the strategy can be extended to arbitrary constraint satisfaction problems. A generalized formulation of the SP equations will be briefly presented in Appendix D.

Since we are interested in determining complete satisfying assignments in the SAT region, every configuration of messages associated to clusters of excited states have to be filtered out. At this purpose, we take the limit $y \to -\infty$ in the reweighing factor (2.34). Substituting then the parameterization (4.4) into the coupled convolution integrals (2.32) and (2.33), we obtain the following simplified relations:

$$\eta_{a\to i} = \prod_{j\in V(a)\setminus i} \left[\frac{\Pi^u_{j\to a}}{\Pi^u_{j\to a} + \Pi^s_{j\to a} + \Pi^0_{j\to a}} \right] \tag{4.5}$$

where:

$$\Pi^{u}_{j\to a} = \left[1 - \prod_{b\in V^{u}_{a}(j)} (1-\eta_{b\to j})\right] \prod_{b\in V^{s}_{a}(j)} (1-\eta_{b\to j}) \tag{4.6a}$$

$$\Pi^{s}_{j\to a} = \left[1 - \prod_{b\in V^{s}_{a}(j)} (1-\eta_{b\to j})\right] \prod_{b\in V^{u}_{a}(j)} (1-\eta_{b\to j}) \tag{4.6b}$$

$$\Pi^{0}_{j\to a} = \prod_{b\in V(j)\setminus a} (1-\eta_{b\to j}) \tag{4.6c}$$

$V^{u,s}_{a}(j)$ are shortcut notations for the sets:

$$\begin{aligned} &\text{if } J_{a,j} = +1 \quad : \quad V^{u}_{a}(j) = V^{+}(j)\setminus a \quad ; \quad V^{s}_{a}(j) = V^{-}(j) \\ &\text{if } J_{a,j} = -1 \quad : \quad V^{u}_{a}(j) = V^{-}(j)\setminus a \quad ; \quad V^{s}_{a}(j) = V^{+}(j) \end{aligned}$$

where $V^{\pm}(j)$ are the sets of the clauses b connected to j with a coupling $J_{b,j} = \pm 1$. We notice that $\left[1 - \prod_{b\in V^{u}_{a}(j)}(1-\eta_{b\to j})\right]$ is just the probability that the variable j is receiving at least one warning trying to align it to the direction $J_{a,j}$; multiplying it by the probability $\prod_{b\in V^{s}_{a}(j)}(1-\eta_{b\to j})$ that j is not receiving warning trying to align it on the direction $-J_{a,j}$, we verify that (4.5) corresponds exactly to the probability that a constrains actively the truth value of i.

Quite remarkably, the optimization limit have removed the need of solving complicated coupled convolution integrals, giving instead the much simpler set of coupled algebraic equations (4.5). At each iterative step, the new population of u-surveys can be computed just by the set of parameters describing the u-surveys of the previous generation. At the fixed point, from the self-consistent set of u-surveys $\{\eta^{*}\}$, we can compute the probabilities $W^{+,-,0}_{i}$, that picking up at random a cluster of solutions, the variable i is frozen in the state ± 1 or is under-constrained:

$$W^{\pm}_{i} = \frac{\widehat{\Pi}^{\pm}_{i}}{\widehat{\Pi}^{\pm}_{i} + \widehat{\Pi}^{\mp}_{i} + \widehat{\Pi}^{0}_{i}} \tag{4.7a}$$

$$W^{0}_{i} = 1 - W^{+}_{i} - W^{-}_{i} \tag{4.7b}$$

with:

$$\widehat{\Pi}^{\pm}_{i} = \left[1 - \prod_{b\in V^{\pm}(i)} (1-\eta^{*}_{b\to i})\right] \prod_{b\in V^{\mp}_{a}(i)} (1-\eta^{*}_{b\to j}) \tag{4.8a}$$

$$\widehat{\Pi}^{0}_{i} = \prod_{b\in V(i)} (1-\eta^{*}_{b\to j}) \tag{4.8b}$$

The SP equations can now be converted into a search algorithm by adopting a repeated *decimation* procedure, in the spirit of the simplest real space

renormalization group. The variables are sorted according to the absolute value of their magnetization:

$$|m_i| = \left|W_i^+ - W_i^-\right| \tag{4.9}$$

and a bunch of pre-ordered size f of the most biased of them is fixed according to the direction of their local field. After the fixing, we actually reduce the original N variable problem into a new one with $N - f$ variables. New u-surveys can then be computed over the simplified factor graph. The procedure of decimation continues until when one of the following possible outcomes is obtained:

(1) the algorithm is able to complete the construction of a full satisfying assignment;
(2) the algorithm do not converge yet after a predefined maximum number of iterations;
(3) the algorithm reach a paramagnetic state.

In the last case, the left subproblem is simplified is passed to a specialized heuristic, like WalkSat, in the attempt of completing the satisfying assignment. It is indeed expected that RS problems are easy to optimize for local search algorithms.

The algorithm described above is known in the literature with the name of Survey-Propagation-inspired decimation (SPID) algorithm and it has been used for the determination of solutions of hard random K-SAT sample, up to size of order $N \simeq 10^7$ [**26, 109**]. In the case of $K = 3$, solutions have been found up to $\alpha_c^{\text{num, 1}} = 4.253$, which is extremely close but slightly smaller than the expected value of $\alpha_c = 4.267$ [**133**]. In particular the hard 3-SAT instance attacked by WalkSat and Quantum Annealing in the previous chapter can be exactly solved by SPID.

An improvement up to $\alpha_c^{\text{num, 2}}$ has been obtained by Giorgio Parisi, introducing a backtracking strategy similar to the one discussed in Section 3.1 [**132**]. The SPID algorithm have been also succesfully used for the optimization of graph coloring and some of its variants [**27, 35**]. A version of the Survey Propagation equations in the $y \to \infty$ valid for every discrete CSP is outlined in Appendix D. Several substantial refinement to the SPID algorithm will be presented in the next Chapters 5 and 6.

3. SP for energy optimization

The $y \to \infty$ limit filters out completely any clause-violating truth value assignment during the computation of the convolution (2.32). This feature is desirable if we want to optimize instances extracted from the SAT region, but must be rejected when our sample is presumably unsatisfiable.

INPUT:: a Boolean formula $\mathcal{F}$ in conjunctive normal form
OUTPUT:: a simplified Boolean formula $\mathcal{F}'$ in conjunctive normal form (ideally empty) and a partial truth value assignment for the variables of $\mathcal{F}$ (ideally a complete one)

0. For each edge $a \to i$ of the factor graph, randomly initialize the $\eta_{a\to i} \in \{0, 1\}$
1. Compute all the fixed point u-surveys, using equations (4.5)
3. IF the population dynamics equations converge,
   3.1 FOR every unfixed variable i, compute the local magnetization m_i using (4.7)
   3.2 Sort the variables according to the $|m_i|$ and fix the highest ranked variable
   3.3 Simplify the graph accordingly
   3.4 IF all the variables are fixed, RETURN the full truth value assignment and an empty sub-formula, ELSE, go to 1.
4. ELSE IF the population dynamics equations do not converge, simplify the formula by imposing the already assigned truth values, RETURN the partial solution and pass the obtained sub-formula to WalkSat

FIGURE 3. The SP-inspired decimation algorithm (SPID).

In the UNSAT region the use of a finite value of the Lagrange multiplier y is required. The filtering action of the re-weighting factor (2.34) is then damped and the set of messages that vehicle information pointing to states with a non vanishing number of violated constraints can contribute to the computation of the integrals appearing in (2.32) and (2.33).

3.1. The SP-Y algorithm. We have seen in Section 2, that the optimization limit $y \to \infty$ transform the complicated convolutions (2.32) and (2.33) into a simple set of coupled algebraic equations. If the pseudo-temperature $1/y$ is kept finite, we have to deal with a non-trivial reweighing factor. Still, a relatively fast recursive procedure can be written [**14, 15**]. Let us take a variable j with Γ_j neighboring function nodes and let us compute the cavity field pdf $P_{j\to a}(h)$ where $a \in V(j)$. We start by randomly picking up a function node b_1 in $V(j) \setminus a$ and calculating the following "h-survey":

$$P^{(1)}_{j\to a}(h) = \eta^0_{b_1\to i}\,\delta(h) + \eta_{b_1\to i}\,\delta(h - J_{b_1,i}) \tag{4.10}$$

INPUT:: a Boolean formula $\mathcal{F}$ in conjunctive normal form; a backtracking ratio r; optionally, a fixed inverse pseudo-temperature y_{in}

OUTPUT:: a simplified Boolean formula $\mathcal{F}'$ in conjunctive normal form (ideally empty) and a partial truth value assignment for the variables of $\mathcal{F}$ (ideally a complete one)

0. For each edge $a \to i$ of the factor graph, randomly initialize the $\eta_{a\to i} \in \{0, 1\}$
1. IF there is a fixed y_{in} as input, put $y^* = y_{in}$, ELSE after a fixed number of steps, determine by bisection the position of the free-energy maximum y^*
2. Compute all the fixed point u-surveys, using equations (4.10), (4.11), (4.12) and putting $y = y^*$
3. IF the population dynamics equations converge,
   - 3.1 FOR every unfixed variable i, compute the local field pdf using (4.10), (4.11)
   - 3.2 Extract a random number q in [0,1]
   - 3.3 IF $q \geq r$, Sort the variables according to the index function (4.14), and fix the most biased variable
   - 3.4 ELSE IF $q < r$, Sort the variables according to the index function (4.17) and unfix the highest ranked variable
   - 3.5 IF all the variables are fixed, RETURN the full truth value assignment and an empty sub-formula, ELSE, go to 1.
4. ELSE IF the population dynamics equations do not converge, simplify the formula by imposing the already assigned truth values, RETURN the partial solution and pass the obtained sub-formula to WalkSat

FIGURE 4. The SP-Y simplification algorithm.

The function $P^{(1)}_{j\to a}(h)$ would correspond to the local field pdf of the variable j in the case in which b_1 was the only neighboring clause (as denoted by the upper index).

The recursive procedure continues then adding the contributions of all the other function nodes in $V(j) \setminus a$, clause by clause (Fig. 5):

$$\begin{aligned}\widetilde{P}^{(\gamma)}_{j\to a}(h) &= \eta^0_{b_\gamma\to j}\,\widetilde{P}^{(\gamma-1)}_{j\to a}(h) \\ &\quad + \eta_{b_\gamma\to j}\,\widetilde{P}^{(\gamma-1)}_{j\to a}(h - J_{b_\gamma,i})\,\exp\left[-2y\,\hat{\theta}(-J_{b_\gamma,i}h)\right]\end{aligned} \tag{4.11}$$

Here $\widetilde{P}^{(\gamma)}_{j\to a}(h)$ denotes an unnormalized pdf and $\hat{\theta}(h)$ is a step function equal to 1 for $h \geq 0$ and zero otherwise. The recursive computation ends after

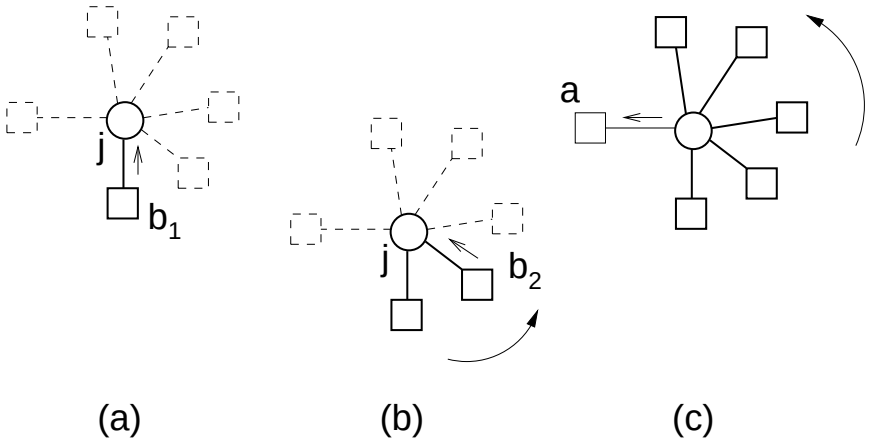

FIGURE 5. Recursive computation of a cavity field pdf. (*a*) To compute the cavity pdf $P_{j\to a}(h)$, a single clause b_1 in $V(j)\setminus a$ is picked up at random and the u-survey $Q_{b_1\to j}$ is used to compute the equation (4.10); (*b*) the contributions of the other function nodes in $V(j)\setminus a$ are added, clause by clause; (*c*) after $\Gamma_j - 1$ iterations we obtain a pdf which coincides with $P_{j\to a}(h)$.

$\gamma = \Gamma_j - 1$ steps, when the contribution of every clause in $V(j)\setminus a$ has been taken in account. By computing $\widetilde{P}_{j\to a}(h)$ for all the values of the field $-\Gamma_j+1 < h < \Gamma_j-1$ also the correct normalization constant can be obtained, finding finally $P_{j\to a}(h) = \widetilde{P}_{j\to a}(h)/C_{j\to a}$. Local fields pdfs $P_i(H)$ can be computed in exactly the same way, just by including the contributions of all the edges afferent to a single variable[2].

The knowledge of $K-1$ input cavity-field pdf's can now be used to obtain a single output u-survey. When taking the $y\to\infty$ limit, the extreme simplification has removed the need of this intermediate computation. Let us compute for instance the u-survey $Q_{a\to i}(u)$ (see always Fig. 2 for the notation). The cavity field pdf's $P_{j\to a}(h)$ for every $j\in V(a)\setminus i$ enter the computation of the parameters $\{\eta^0_{a\to i}$ and $\eta_{a\to i}$:

$$\eta_{a\to i} = \prod_{n=1}^{K-1} W^{J_{j_n,a}}_{j_n\to a}, \quad \eta^0_{a\to i} = 1 - \eta^{J_{a,i}}_{a\to i}, \tag{4.12}$$

[2]The preceding recursion brings to the computation of all the cavity fields pdfs in a time which is *cubic* in the average connectivity $\langle\Gamma\rangle$ ($\langle\Gamma\rangle = K\alpha$ for Erdös-Réni instances). It is also possible to compute per every single node the single local field pdf $P_i(H)$ and then to remove just the single contributions associated to the missing edges, in order to compute all the cavity pdfs. The overall procedure (requiring parallel update of the u-surveys instead of sequential) has an overall cost *quadratic* in the average connectivity.

where we introduced the cavity weight factors:

$$W^{+}_{j\to a} = \sum_{h=1}^{\Gamma_j - 1} P_{j\to a}(h), \quad W^{-}_{j\to a} = \sum_{h=-\Gamma_j+1}^{-1} P_{j\to a}(h). \tag{4.13}$$

A systematical application of the recursions (4.10), (4.11) and of the relation (4.12) brings once again to a full update of the population of u-surveys and the step is repeated until when self-consistency is reached.

The magnetization of the variable is computed as usual:

$$|m_j| = |W^{+}_j - W^{-}_j|. \tag{4.14}$$

with the cluster fractions given by:

$$W^{+}_j = \sum_{H=1}^{\Gamma_j} P_j(H), \quad W^{-}_j = \sum_{H=-\Gamma_j}^{-1} P_j(H), \tag{4.15}$$

Decimation once again can be performed by fixing the most biased variables. After the fixing we must take care of the fixed variables, returning always fully polarized cavity field pdfs, regardless of the recursions (4.10), (4.11):

$$P_{i\to a}(h) = \delta\left(h - J_{a,i}s_i\right); \tag{4.16}$$

The complete polarization reflects the knowledge of the truth value of the literals depending on the spin s_i.

An important difference with the SPID algorithm is that, when working at finite pseudo-temperature, the probability that some non optimal fixing is done in presence of "thermal noise" cannot be neglected. After several update sweeps, the biases of some fixed spins may become smaller than the value they had at the time when the corresponding spin was fixed. It is even possible that certain local fields revert completely their favored orientation. An index function like:

$$m_{\text{backtrack}}(j) = -s_j\left(W^{+}_j - W^{-}_j\right), \tag{4.17}$$

can be used to track the appearance of such prematurely fixed spins. We can then devise a simple "error-removal" strategy where both unfixing and fixing moves are performed at a fixed ratio $0 \le r < 0.5$ (another slightly more complex backtracking implementation can be found in [**132**]).

Another crucial free parameter for the optimization version of SP (SP-Y) is the inverse pseudo-temperature y. In the simplest version of the algorithm, y is kept fixed during the simplification, but it is also possible to update it dynamically in order to keep a value as close as possible to the maximum y^* of the free energy functional $\Phi(y)$ given by equation (2.38).

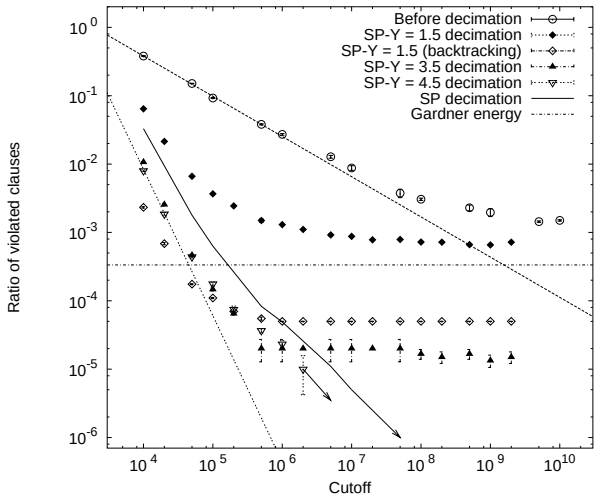

FIGURE 6. Efficiency of SP-Y in the SAT region (single sample with $N = 10^5$ variables and $\alpha = 4.24$). An arrow indicates that the next data point corresponds to a SAT assignment. After SP-Y simplification, WalkSat succeeds generally in finding assignments below e_G.

This choice ensure indeed[3] to select the pdfs of set of messages associated with states close to the instance ground-state. The free energy functional given by the equations (2.39), (2.40) can be easily evaluated algorithmically from the knowledge of the cavity and the local field pdfs:

$$\Phi_a^f(y) = -\frac{1}{y}\Big[\ \ln\ \Big(1 + (\mathrm{e}^{-y} - 1)\prod_{i\in V(a)} W_{i\to a}^{J_{a,i}}\Big) - \ \ln\Big(\prod_{i\in V(a)} C_{i\to a}\Big)\Big], \tag{4.18}$$

$$\Phi_i^y(y) = -\frac{1}{y}\ \ln\ (C_i)\,. \tag{4.19}$$

A summary of the SP-Y simplification procedure can be found in Fig. 4.

3.2. Overview of SP-Y results. As we have already discussed in Chapter 3, for sufficiently large instances, any local search algorithm gets trapped by the exponentially numerous threshold states with energy close to e_G. In this section, we will review results relative to the posiitve performance of the SP-Y algorithm when applied to the optimization of several hard samples taken both from the SAT and the UNSAT phase.

3.2.1. *SAT region.* Using the complicated SP-Y strategy in the SAT region where SPID is already so efficient, could seem cumbersome, but

[3]At least as far as the 1-RSB and the Bethe approximation have to be considered valid.

this will allow us to understand better the working of the new algorithm. In particular we will show experimentally how the choice of the pseudo-temperature $1/y$ affects the energy of the addressed layers of states.

The efficiency of the simplest SP-Y decimation strategy against the glass threshold is exemplified in Fig. 6. No backtracking has been used in this purely explorative experiment. We optimize a single randomly generated formula ($N = 10^5$, $\alpha = 4.24$, the same sample of Chap.2) at several fixed values of pseudo-temperature. For comparison we plot also the WalkSat results after a standard SPID decimation (solid line): as expected, the ground state, $E = 0$, is reached after only a rather small number of spin flip trials. The same happens after SP-Y decimation, if we choose a sufficiently large value of the inverse pseudo-temperature ($y > 4$); we remind indeed that in the SAT region the optimal value for y would be infinite, limit in which the SP-Y recursions reduce to the SP equations. If the decimation is performed with smaller values of y, SP-Y loses convergence before reaching a paramagnetic state. WalkSat is no more able to find a zero energy assignment of the left-subproblem. Nevertheless, if y is not too small, the cooling curves reach a saturation level which lies well below the Gardner energy. Although the exact ground state is not found, the trapping effect of the original threshold states can be successfully overcome. The failure is not surprising, since using a finite y we are sampling many assignments that should be completely filtered out in the SAT region. The procedure is intrinsically error prone, and it will allow in general to reach only "good states", but not the true optimal solutions. The use of backtracking can partially cure the accumulation of errors at finite y: the saturation level can in fact be significantly lowered by keeping the same pseudo-temperature and introducing a small fraction of backtrack moves during the simplification. In Fig. 6, the data for $y = 1.5$ saturate above Gardner energy, but the introduction of a tiny fraction of backtrack moves is enough to move the plateau well below the lower bound to the threshold.

3.2.2. *UNSAT region.* When entering the UNSAT region, the task of looking for the optimal state becomes harder (there is complexity, but the complexity of ground state is null). The optimal assignments contain unavoidably violated constraints and the use of a finite pseudo-temperature becomes compulsory. Unfortunately, after many spin fixings, the recursions (4.10), (4.11) stop to converge for y larger than some finite value smaller than the maximum of free energy y^*. This mean that simplification can be performed only in a suboptimal regime. This phenomenon takes most likely place when the sub-problem enters a full-RSB-like phase. At this point, a 2-RSB version of SP could be used to achieve convergence at larger y's, but

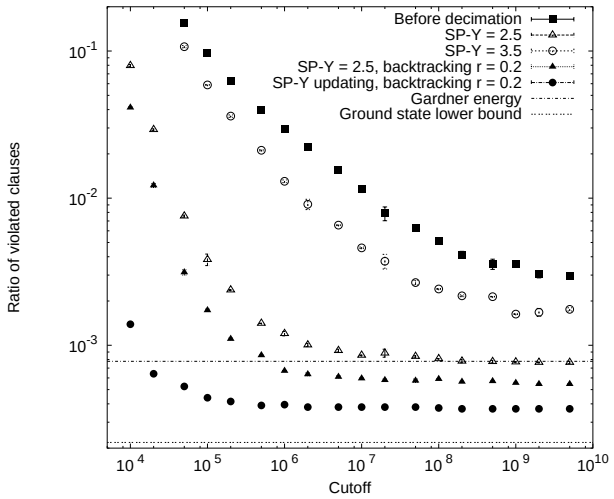

FIGURE 7. SP-Y performance in the UNSAT region (single sample with $N = 10^5$ variables and $\alpha = 4.29$). Several simplification strategies are compared; the combined use of backtracking and runtime y updating allows to reach energies closer to the ground state than to the Gardner lower bound.

the small improvement would not justify the computational effort required for achieving it. Alternatively, one could try to launch directly some local search heuristics (hoping that the full RSB sub-system is not exponentially hard to optimize). We decided more simply to continue the decimation with the largest possible value of y allowing convergence, until when either convergence is lost (for every value of y) or a paramagnetic state is reached.

In our experiments we studied several 3-SAT samples problems extracted from the 1-RSB stable UNSAT phase. WalkSat continued to be our heuristics of choice: although it is not optimized for unsatisfiable problems, in the 1-RSB stable UNSAT region it performs still much better than any basic implementation of SA. Nevertheless, even after the huge number of 10^{10} spin flips, the WalkSat best score was still quite distant from the Gardner energy.

In Fig. 7 we show the results relative to many different SP-Y simplifications with various values of y and r for a single sample with $N = 10^5$ and $\alpha = 4.29$. The simplification produced always an improvement with respect to the WalkSat performance. Nevertheless, we were unable to go below the Gardner lower bound, without backtracking, although we touched it in some cases: in Fig. 7 we show the data for a simplification at fixed $y = 2.5$; a simplification with runtime optimization of y but without backtracking reached the same level.

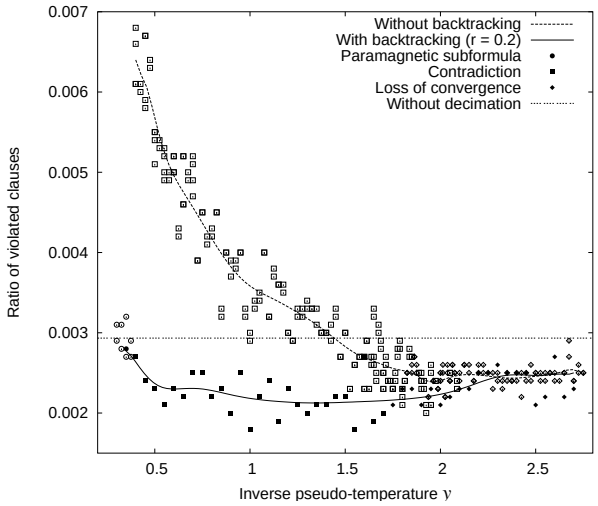

FIGURE 8. Backtracking efficiency. Many SP-y experiments at various fixed Y have been performed on a single sample with $N = 10^4$, $\alpha = 4.35$. The three possible outcomes of the simplification are coded by different shapes: a diamond indicates that the simplification process stopped because of loss of convergence, a circle because of finding a completely unbiased paramagnetic state, and the squares indicates that the loss of convergence happened at an advanced stage where some clause-violating assignments had already been introduced by SP-Y.

The exact choice of the value of y can affect enormously the overall performance. We conducted first an exhaustive simultaneous optimization of y and r, using smaller samples in order to produce more experimental points. After some trials, we identified the fraction $r = 0.2$ as optimal, at least for our implementation, and in the small region under investigation of the K-SAT phase diagram. The data in Fig. 8 refers to a formula with $N = 10^4$ variables and $\alpha = 4.35$. The dashed horizontal line shows the WalkSat best energy obtained on the original formula after 10^9 spin flips. The WalkSat performance was seriously degraded after simplifications at too small values of y, but the introduction of backtracking cured the problem, identifying and repairing most of the wrong assignments. The WalkSat efficiency became almost independent from the choice of pseudo-temperature, whereas in absence of error correction a time consuming parameter tuning was required for optimization.

Coming back to the analysis of the sample of Fig. 7, the backtracking simplifications allowed us to access states definitely below the Gardner

lower bound. The most effective strategy was a combination of runtime y-optimization and of error correction, which was able to produce a WalkSat saturation level strikingly closer to the ground state lower bound than to the Gardner energy. A further valuable effect of backtracking was the increased efficiency of the formula simplification itself: while without backtracking, the algorithm lose convergence on average after having fixed just the 40% of variables[4], in the backtracking experiments, only the 20% percent of truth value assignments had to be determined by the supporting heuristics.

All the previously analyzed samples were taken from the 1-RSB stable region of the 3-SAT problem, where the equations (2.32), (2.33) can be considered exact. For $\alpha > 4.39$, the phase is supposed to become full-RSB (the loss of convergence at $y < y^*$ takes place interestingly at the very first decimation step). In this region, SP-Y can still be used in a sub-optimal way, but is not useful, since Simulated Annealing alone is already able of finding close-to-optimum assignments efficiently (as expected for a full RSB scenario) and behaves definitely better than WalkSat.

[4]This proportion is comparable with the fixing percentage achieved by SPID.

CHAPTER 5

State selection with SP-ext

We have shown in the previous chapters how the cavity method can be used in the hard phase of the K-SAT problem to compute the local fields probability distribution functions, and how this information can be exploited in order to build complete satisfying assignments. The SPID algorithm is based on the use of a simple unbiased decimation strategy, which fixes at each step only the highest ranked variables according to some index function. In principle, the SP iterations have at least as many fixed points as clusters of solutions (as clarified also by the discussion in Appendix D). In practice, it turns out from experiments that, unless very specific initial conditions for the population dynamics are selected, a single fixed point is invariantly selected, which does not correspond on the other hand to any individual group of solutions. The convolutions (2.32) and (2.33) are computed using a flat measure over all the zero-energy clusters (as evident from (4.3)), and each cluster contribute then equally to the determination of the fixed-point pdfs for cavity and local fields. As a result of this, SPID retrieves invariantly always the same solution when the index-function for decimation and the problem instance have been fixed.

In the present chapter, we will show how the weight of the different clusters can be altered in a controlled way, without altering the overall computational cost of solution-finding, by the introduction of additional control messages. This new algorithmic feature will be used to investigate the geometry of the solution space on a local basis. The clustering phenomena will be then exploited as a computational resource for pattern restoration and lossy data compression [**12, 13**].

It is worth reminding that the formalism introduced is of general applicability to any random CSP characterized by clustered phases (see [**28**], or Appendix D). In the next sections we will continue for simplicity to address the case of random K-SAT.

1. The SP-ext forcing mechanism

Our goal is now to derive a generalization of the SPID algorithm allowing the retrieval not only of some solution, but, more specifically, of a

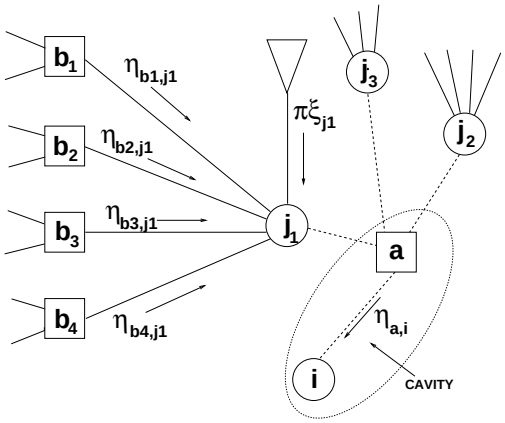

FIGURE 1. Factor graph with external messages. The triangle nodes represent external sources of messages, whose probability distribution functions are never updated but are preimposed in order to control efficiently the population dynamics convergence

solution close to any desired point (ideally, the closest solution to any given configuration $\vec{\xi}$).

A standard approach would be the resolution of the cavity equations in presence of a geometrical constraint, conjugated to an additional Lagrange Multiplier $\widetilde{y}$. One could hope then to slice states with specific overlaps with $\vec{\xi}$ just by carefully setting y, in a way analogous to the SP-Y strategy described in Section 3.1. It is quite likely, nevertheless, that this method involves the use of non integer fields, and the algorithmic implementation of the population dynamics could not be efficient.

An alternative approach has very recently been proposed in [**102**] and constitutes in applying the ordinary cavity method to a different combinatorial problem, defined over the space of the solutions of the original one. More specifically, if $\mathcal{S}$ is the space of satisfying assignments of a given K-SAT instance, a new problem with hamiltonian $C_d(\vec{s}) = d\left(\vec{s}, \vec{\xi}\right)$ is defined for every $\vec{s} \in \mathcal{S}$, where $d(\cdot, \vec{\xi})$ denotes the Hamming distance from the reference configuration $\vec{\xi}$ or another well defined metric. Geometric complexity functions, giving an estimation of the number of clusters of solutions at a given distance from the reference, can be obtained, and the framework is generalizable in order to obtain counts of clusters at a given energy, with a given internal entropy and at a given distance [**102**]. Nevertheless, no attempts of applying this promising formalism to the determination of assignments have been done up to date, and, in particular, it is not known yet

if a decimation strategy could be carried out without loosing convergence at a too early stage.

We shall now present, on the other hand, a quite natural extension of the SPID algorithm (SP-ext, [**12, 13**]), not supported by a complete theory, but extremely efficient in optimization applications. New external sources of messages are attached to each variable of the instance under consideration (denoted by triangles in the factor graph representation of Fig.1). The probabilities that these sources inject actually into the system a message are given *a priori* and encoded into additional surveys that are never updated according to the population dynamics equations (2.32) and (2.33), but just enter into the probability measures (2.20). Given an arbitrary reference configuration $\vec{\xi}$, we define as *forcing* in the direction $\vec{\xi}$ and of intensity $\vec{\pi}$ the collection of additional surveys:

$$Q_i^{\text{ext}}(u) = (1-\pi_i)\delta(u) + \pi_i\delta(u-\xi_i) \tag{5.1}$$

The parameter π_i can be interpreted as local strength of the perturbation to the population dynamics in direction ξ_i. Since we are mainly interested in addressing clusters of completely satisfying assignments, we will focus hereafter on the $y \to \infty$ limit.

After the application of the forcing, the cavity weights entering the SP equations (4.5) becomes:

$$\eta_{a\to i} = \prod_{j\in V(a)\setminus i}\left[\frac{\Pi^{J_{a,j}}_{j\to a,\pi_j}}{\Pi^{+}_{j\to a,\pi_j}+\Pi^{-}_{j\to a,\pi_j}+\Pi^{0}_{j\to a,\pi_j}}\right] \tag{5.2}$$

with:

$$\begin{aligned}\Pi^{\pm}_{j\to a,\pi_j} = &\left[1-(1-\pi_j\,\delta_{\xi_j,\pm})\prod_{b\in V^{\pm}\setminus a(j)}(1-\eta_{b\to j})\right]\cdot\\ &\cdot(1-\pi_j\,\delta_{\xi_j,\mp})\prod_{b\in V_a^{-\xi_j}(j)}(1-\eta_{b\to j})\\ \Pi^{0}_{j\to a,\pi_j} = &(1-\pi_j)\prod_{b\in V(j)\setminus a}(1-\eta_{b\to j})\end{aligned} \tag{5.3}$$

The relations (5.3) and the cavity analogues to equations (4.7) allow to evaluate the quantities $\mathcal{W}^{\pm,0}_{j\to a}(\pi,\vec{\xi})$. These so called cavity weights measure the probabilistic weight associated to the set of clusters of solutions of the cavity graph in which the variable j is respectively frozen in the ±1 state or is unfrozen, in absence of the interaction a and in presence of an external forcing along the direction $\vec{\xi}$ of intensity π. Considering the parameters $\eta_{a\to i}$ as dynamical variables updated by the population dynamics, it becomes possible to monitor the evolution of the cavity weights under the

SP-ext iterations. Let take an instantaneous snapshot of the values of all the parameters $\eta_{a\to i}$ and let compute the temporary values of the cavity weights $\mathcal{W}^{\pm,0}_{j\to a}(\pi,\vec{\xi})$. It is possible to separate explicitly a term $\mathcal{Y}^{\pm,0}_{j\to a}(\pi,\vec{\xi})$ depending on the forcing $\pi\vec{\xi}$ from a term $W^{\pm,0}_{j\to a}$ which would be unmodified in absence of the external messages:

$$\mathcal{W}^{\pm,0}_{j\to a}(\pi,\vec{\xi}) = W^{\pm,0}_{j\to a} + \mathcal{Y}^{\pm,0}_{j\to a}(\pi,\vec{\xi}) \tag{5.4}$$

Let denote as:

$$\varkappa_{j\to a}(\vec{\xi}) = W^{-\xi_j}_{j\to a} + W^{0}_{j\to a}, \tag{5.5}$$

the probabilistic weight of the locally misaligned clusters if π is temporarily set to zero, that is the unperturbed weight of the clusters in which the variable j is not frozen in the local direction ξ_j. Let also introduce the following notation:

$$\Delta_{j\to a}(\pi,\vec{\xi}) = \frac{\pi_j\left(1-\varkappa_{j\to_a,\xi_j}\right)}{1-\pi\,\varkappa_{j\to_a,\xi_j}} \tag{5.6}$$

It can then be easily shown starting from (5.3), the cavity version of (4.7) and (5.4), that:

$$\begin{aligned}\mathcal{Y}^{-\xi_j,0}_{j\to a}(\pi,\vec{\xi}) &= -\Delta_{j\to a}(\pi,\vec{\xi})W^{-\xi_j,0}_{j\to a}\\ \mathcal{Y}^{\xi_j}_{j\to a}(\pi,\vec{\xi}) &= +\varkappa_{j\to a}(\vec{\xi})\,\Delta_{j\to a}(\pi,\vec{\xi})\end{aligned} \tag{5.7}$$

Every misaligned cluster loses an amount $\Delta_{j\to a}(\pi,\vec{\xi})$ of probability weight. The total perturbations $\mathcal{Y}^{-\xi_j,0}_{j\to a}(\pi,\vec{\xi})$ to the fractions of misaligned clusters are indeed proportional to the unperturbed fractions $W^{-\xi_i,0}_{j\to a}$. The same amount of probability weight is transferred to the clusters with the correct local orientation.

Since $\pi_j < 1$, the transferred quantum of probability weight (5.6) is a decreasing function of the misaligned fraction $\varkappa_{j\to a}(\vec{\xi})$. The perturbation of the cavity weights is then *adaptive*: the perturbation operated by the forcing is indeed much more effective when the weight of the misaligned clusters is small. This specific form of the elementary probability transfer (5.6) ensures typically convergence to a self-consistent fixed point, unless the strength of the perturbation is too large. While intensities $\pi_j \simeq 1$ would produce a complete polarization of the messages if $\vec{\xi}$ was a solution, in the general case the use of a smaller forcing intensity allows the system to react to the contradictory driving and to converge to a set of surveys sufficiently biased in the desired direction. In particular, this reaction produces an average internal magnetization level [1] $\langle|W^+ - W^-|\rangle$ constantly larger than in

[1] Computed considering only the surveys if internal messages.

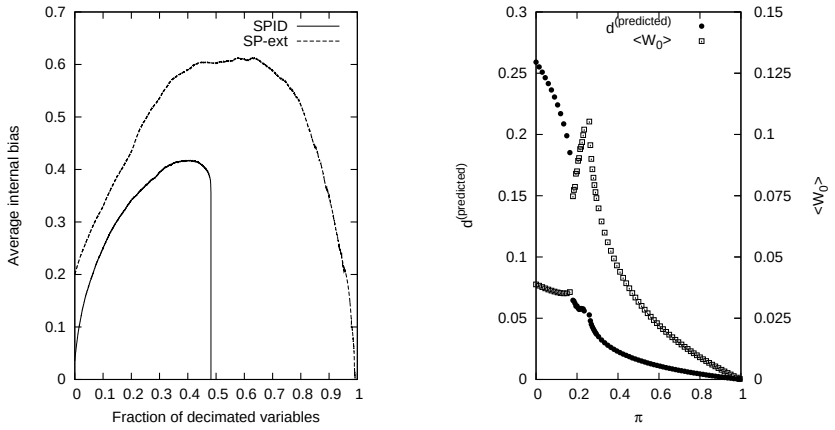

FIGURE 2. (*Left*) SP-ext decimation. While SPID undergoes a sudden paramagnetic collapse after the decimation of a certain fraction of the variables, SP-ext average magnetization is constantly larger and complete decimation becomes possible. (*Right*) The cluster-selection transition (forcing applied along a solution $\vec{\sigma}$). After a critical forcing strength π_c, strong polarization of the fields takes place and the predicted distance from $\vec{\sigma}$ vanishes. A cusp in the predicted fraction of unfrozen variables is observed at $\pi = \pi_c$ (note the different vertical scale).

the ordinary SP case. This allows a decimation strategy based on the SP-ext marginals to be used until the determination of complete assignments, without experiencing the sudden paramagnetic collapse typical of the SPID algorithm (see Fig. 2 at left).

An important observation is that the satisfying assignment $\vec{\sigma}_{\vec{\xi}}$ produced when forcing in direction $\vec{\xi}$ shows a significantly non random correlation with $\vec{\xi}$ (a typical distance scale for Erdös-Rény random samples extracted in the hard SAT region is $\simeq 0.3$). This basic experiment highlights that SP-ext can be practically and efficiently used to perform decimation under a lapse geometrical constraint. This important feature will be the base of the applications presented in Sections 2 and 3.

1.1. The cluster selection transition. Let us consider a given solution $\vec{\sigma}_0$ and let take a forcing of homogeneous intensity $\pi_i = \pi$ for all variables i parallel to $\vec{\sigma}$. The fixed-point of the (5.2) equations is computed for increasing values of π. After convergence, we can evaluate the two quantities $\langle W_i^0 \rangle$ and $d^{\text{predicted}} = \frac{1}{2}\left(1 - \sum_i W_i^{\sigma_i}\right)$ giving respectively the average weight of the

clusters in which a variable is free to fluctuate and the predicted average distance from the reference solution.

These two indicators are plotted against the forcing strength in Fig.2 (at right), for a given sample extracted from the hard-SAT region. A phase transition is clearly visible for a critical value π_c of the forcing strength. We call this phase transition *cluster-selection transition*, since for $\pi > \pi_c$ the average distance from $\vec{\sigma}_0$ rapidly rapidly vanishes, exactly as if only clusters with a significant overlap with $\vec{\sigma}_0$ were considered with a non negligible weight in the survey for the u-messages.

Unfortunately, since the SP-ext equations lack a rigorous derivation, no further conclusion can be extracted from the empirical study of these estimators. We do not provide any explanation, in particular, of the cusp in $\langle W^0 \rangle$. Hence, much progress has still to be done in order to use SP-ext-like methods for the derivation of geometric complexity curves.

2. Local probing of the solution space

SP-ext can be efficiently used to probe the local geometry of the zero energy phase. First, we select a random reference solution $\vec{\sigma}$ found by mean of ordinary SPID or also SP-ext decimation, with application of a random direction forcing. In the following of this section, we will refer to $\vec{\sigma}_0$ as to the *seed* of the grown cluster. SP-ext decimation performed with a forcing in direction $\vec{\sigma}$ obviously reproduce an assignment at least belonging to the same cluster of $\vec{\sigma}$. If π was large enough a true cluster-selection transition would take place, but now we want to keep π small enough in such a way that a decimation forced in a completely random direction can proceed until the end.

We flip then in $\vec{\sigma}$ a fraction δ of randomly chosen spins and we perform a new decimation, forcing along the obtained direction $\vec{\zeta}_\delta$ [2]. The result is a new satisfying assignment that we call $\vec{\sigma}_\delta$. This step is then repeated by generating new forcings $\vec{\zeta}_\delta$ from $\vec{\sigma}$ for increasing values of δ. The interesting quantity to monitor is the Hamming distance from the cluster seed $d_H(\vec{\sigma}, \vec{\sigma}_\delta)$ in dependence of the geometric perturbation δ.

In order to have a highly homogeneous distribution of the clusters, we have chosen for our experiments an ensemble of random K-SAT in which variables have a fixed even degree and are balanced (i.e. have an equal number of directed and negated occurrences in the clauses). For this specific ensemble and for $K = 5$, the critical points are given by $\alpha_d = 14.8$ and $\alpha_c = 19.53$. An estimation of the position of the Gardner instability point can be obtained by looking at the point in which the marginals computed

[2] $\vec{\zeta}_\delta$ in particular will no more be a solution, but a configuration with finite energy.

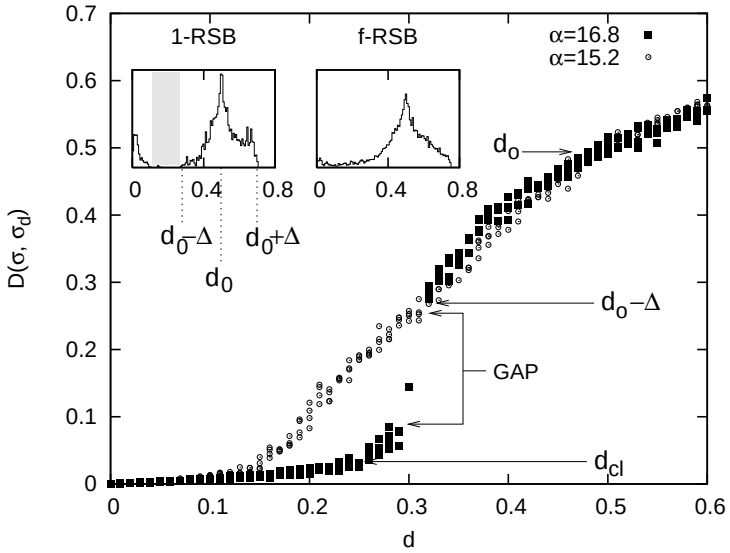

FIGURE 3. Stability diagram for fixed connectivity, balanced $K = 5$ formulas. A step-like transition is observed for the case $\alpha = 16.8$, substituted by a continuous convex curve for $\alpha = 15.2$. In the 1-RSB case, the sampled histogram of the reciprocal distances between found solutions is characterized by a small gap, disappearing below the estimated Gardner instability critical point.

by SP start to differ from the evaluation of an approximate 2-RSB version of SP (called SP-2 [**10**], see also Appendix D). According to our measures, the instability point should be positioned in proximity of $\alpha_G = 16.77$.

In the experiments we have taken instances of size $N = 10^4$ with an intensity $\pi = .35$ of the forcing (much close to the largest value of π for which the SP-ext equations always converged for this sample). The Hamming distance between σ and σ_δ is shown inf Fig. 3, plotted against δ. The first *stability diagram* (black data points) refer to a formula inside the 1-RSB stable region with $\alpha = 16.8$. Two zones can be clearly observed. First, for $\delta < \delta_c \simeq 0.3$, the distance $d_H(\vec{\sigma}, \vec{\sigma}_\delta)$ linearly increases with a very small slope. At $\delta \simeq 0.25$ we find still solutions at a distance $d_{\rm cl} \simeq 0.035$, which is in good agreement with the average fraction of unfrozen variables $\langle W^0 \rangle$, computed with SP (a first very rough estimation for the scale of the cluster diameter).

Conversely, for $\delta > \delta_c$, a sudden jump to a larger distance is obtained. Increasing δ up to 1, a new branch of the stability diagram is generated, extended over a span of distances of radius Δ centered in d_0. The symmetry of this set of distances can be better observed in the leftmost inset, containing

the histogram of all the reciprocal distances between the sampled solutions. In this histogram, the existence of a gap is even more visible. The measured value of d_0 (the distance obtained when using a completely random forcing at $\delta = 0.5$) is in good agreement with the analytic estimation coming from $\langle W^+ - W^- \rangle$. The average magnetization provides indeed the typical scale of the overlap between two randomly selected clusters $q_0 = \frac{1-d_0}{2}$ (under the hypothesis of homogeneous distribution of clusters, quite reasonable for the ensemble under consideration).

The gap between clusters is the main prediction of the 1-RSB cavity theory which finds in these experiments a nice confirm. In the intermediate region $0.25 < \delta < \delta_c$ a small deviation from the inter-cluster linearity of distances is observed[3]. This effect becomes less important for increasing sizes, but still persists for $N = 10^6$, and seem to be related to the peculiar irregular geometry of the surface of clusters. This point will further be discussed in the next Section 2.2.

A completely different behavior is observed when repeating the experiment for a sample in the expected f-RSB phase, with $\alpha = 15.2$. The gap in the histogram of the full set of reciprocal distances now disappears (rightmost inset of Fig. 3) and also the curve followed in the stability diagram (white data-points) has a characteristic continuous and convex shape. It must be remarked that various shapes of the stability plot, differing for their degree of convexity but often gapless can be obtained starting from different cluster seeds over the same sample. This hints to the existence of a mixed phase, in which higher order hierarchical clustering is present and in which many local cluster distribution topologies can coexist. Further experimental exploration of this region, eventually supported by the clustering techniques outlined in Section 2.2, would be needed in order to clarify the properties of the unstable region and the mechanisms bringing to 1-RSB phase instability.

2.1. Correlations with the forcing energy. The data obtained from the stability analysis of the previous section can be studied also in energetic terms. Plotting $d_H(\vec{\sigma}, \vec{\sigma}_\delta)$ against δ but also against the energy $E(\vec{\zeta}_\delta)$ of the associated forcing direction ζ_δ, three regimes become clearly distinguishable.

The inter-cluster regime is also the low forcing energy regime $E < E_c \simeq 0.46$ and it is characterized by an almost linear increase of $d_H(\vec{\sigma}, \vec{\sigma}_\delta)$ with the energy $E(\vec{\zeta}_\delta)$. Increasing the energy, approximate linearity is lost and an

[3]Interestingly the largest distance observed in this "bubble" of points coincides with the height of the cusp in $\langle W^0 \rangle$ at the cluster-selection phase transition.

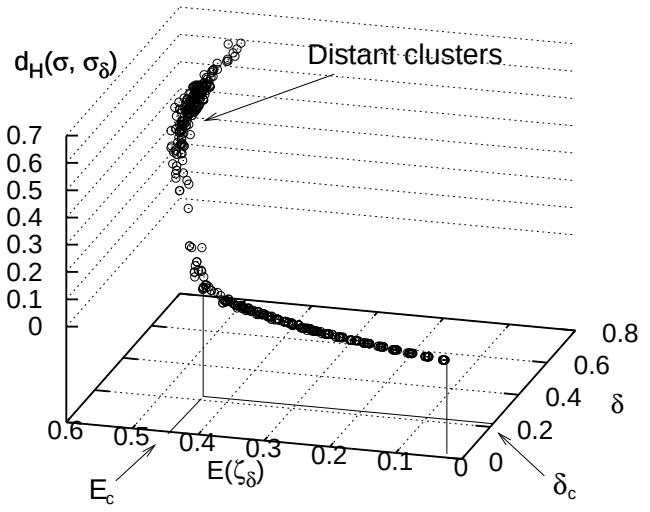

FIGURE 4. Stability diagram in energetic terms ($K = 5$, balanced, $\alpha = 16.8$). Instability is triggered by a forcing energy larger than E_c.

abrupt increase of $d_H(\vec{\sigma}, \vec{\sigma}_\delta)$ is noticed (transition regime, corresponding to the anomalous data-points in the gap region).

Then, after a small gap in both energy and distance, the full branch of Fig. 3 associated to distant clusters collapses into a small island. Its slightly curved shape, reminiscent of a boomerang, can be ascribed to the observed fact that configurations $\vec{\zeta}_\delta$ anticorrelated with $\vec{\sigma}$ ($\delta > 0.5$) have an energy smaller than that of purely random configurations.

This kind of analysis suggest that the instability of the SP-ext mechanism causing the arrest of the growth of the cluster from its seed could be triggered also by an exceedingly large energy of the applied forcing. When $E(\vec{\zeta}_d)$ becomes too large, several alternative ways of repairing the contradictions imposed by the forcing become possible, and different clusters tend to compete for the dominance over the population dynamics evolution, preventing then the onset of the cluster-selection transition, typical of the inter-cluster regime.

2.2. Principal component analysis of SP-ext data-sets. Well established clustering techniques can be employed in order to obtain visualizations of the distance relations existing between all the solutions generated

in extensive SP-ext experiments. Work is in progress [**18**], in order to apply *Principal Component Analysis* (PCA) to data-sets generated during the construction of stability diagrams. One of the goals is to achieve a better understanding of the geometrical distribution of solutions around a given point (*eg.* the seed of the grown cluster). No more information than the one already contained in the data-set itself can be accessed, but the new kind of statistical analysis could be helpful in determining a reduced number of relevant parameters in the description of the scenario.

PCA involves a mathematical procedure that transforms a number of possibly correlated variables into a small number of uncorrelated variables called principal components [**74, 124**]. The first principal component accounts for as much of the variability in the data as possible, and each succeeding component for as much as possible of the remaining variability. In one of the simplest approaches to PCA, a square symmetric covariance matrix is built from the data:

$$\mathbb{C}_{\vec{\sigma}} = \left\langle \left(\vec{\sigma} - \langle\vec{\sigma}\rangle\right)\left(\vec{\sigma} - \langle\vec{\sigma}\rangle\right)^T \right\rangle \tag{5.8}$$

where the row vectors $\vec{\sigma}$ represent solutions of an instance generated by SP-ext and averages are taken over the whole solution set. The first few eigenvalues and eigenvectors of $\mathbb{C}_{\vec{\sigma}}$ can easily be found by an efficient numerical procedure which starts with Householder reduction to tridiagonal form.

Once a partial orthogonal basis of eigenvectors $\vec{X}, \vec{Y}, \vec{Z}, \ldots$ with the desired size has been obtained, the principal components ($X_{\vec{\sigma}}$, $Y_{\vec{\sigma}}$, $Z_{\vec{\sigma}}, \ldots$ of a data-point $\vec{\sigma}$ are obtained just by projecting the solution onto the basis itself.

Significant data-sets can be built with SP-ext by mean of a multiple generation protocol. First a random solution $\vec{\sigma}^A$ is taken as seed for a cluster growing experiment. Using progressively more perturbed configurations $\vec{\zeta}^A_\delta$ as forcing, the full set of solutions $\vec{\sigma}^A_\delta$ needed for the construction of the stability diagram is generated by SP-ext decimation. The second generation is obtained by building the full stability diagrams using *each one* of the solutions $\vec{\sigma}^A_\delta$ as cluster seed: one obtains then the solutions $\vec{\sigma}^A_{\delta,\delta'}$ for as many as possible combinations $0 \le \delta, \delta' \le 1$. With the same mechanism, third generation data-points $\vec{\sigma}^A_{\delta,\delta',\delta''}$ can be produced and so on. Then the data-set can still be expanded by repeating the full procedure from new completely independent seeds $\vec{\sigma}^B, \vec{\sigma}^C, \ldots$. Figs. 5, 6 and 7 show preliminary results of a work in progress, relative to three $N = 10^4$ instances extracted from the ordinary Erdos-Reny random 3-SAT ensemble, the first one with $\alpha = 3.60$, the second one with $\alpha = 4.10$ and the third one with $\alpha = 4.21$. Only

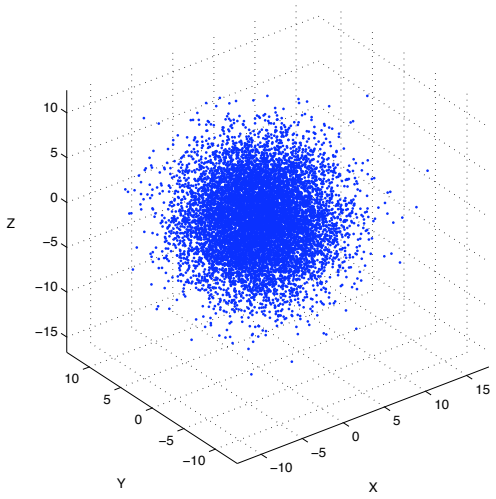

FIGURE 5. PCA clustering for an $\alpha = 3.60$ instance (RS).

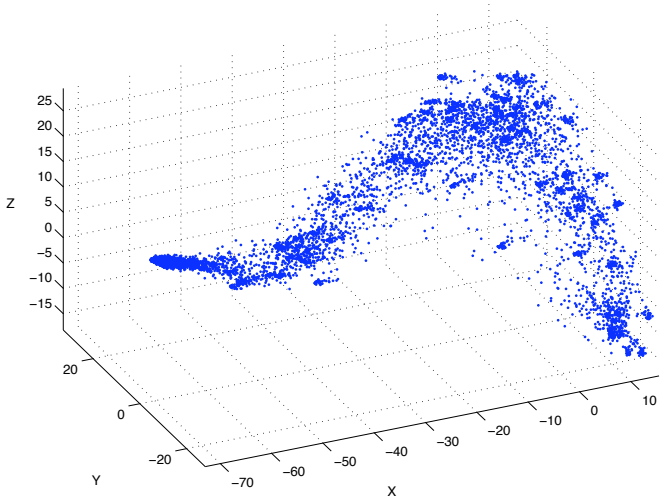

FIGURE 6. PCA clustering for an $\alpha = 4.10$ instance (f-RSB).

one "zeroth" generation cluster seed was used and the construction of the data-set stopped after producing a partial but homogeneous subset of the third-generation solutions.

In a space spanned by the first three principal components, the RS case with $\alpha = 3.60$ (Fig. 5) is characterized by the presence of a single giant globular amass of solutions, indicating that only a single scale of distance can be employed to model the main features of the zero-energy landscape.

When increasing α elongated structures begin to appear, connecting the "galaxy" of the solutions nearby the original cluster seed, to the region of

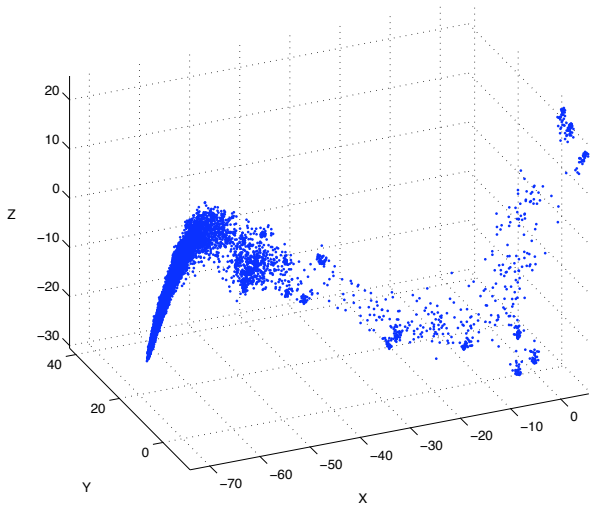

FIGURE 7. PCA clustering for an $\alpha = 4.21$ instance (1-RSB).

the external clusters. In the f-RSB-like case with $\alpha = 4.10$ (Fig. 6), several grains of solutions seem to be diffused in a powder of scattered points. Rarefactions in the tissue of data seem to mark the transition between different clusters of clusters, but more statistics and a better filtering would be needed in order to clarify the f-RSB clustering picture.

Better results are already available for the 1-RSB case (Fig. 7). Here only a significantly rarified powder is scattered in the gaps between grains representing the distant clusters, and between the distant clusters and the main grown cluster to which the seed belong. It must be remarked that the abnormally large size of the cluster containing the seed is a systematic artifact induced by the data-set building protocol. The result is a magnification of the details of the local neighborood surrounding the grown cluster seed. It is possible to identify a small group of satellite amasses of solutions, that seem to be closer to the main cluster than to the region of the distant ones. In particular the largest satellite group can be put in correspondence with the solutions in the small region of the stability diagram in Fig. 3 with $0.25 < \delta < \delta_c$, where the slope of the distance increase sensibly varies with respect to the proper intra-cluster region, but which lies definitely below the larger distance gap at $\delta_c \simeq 0.3$. These data-points constitute also the transition energy region discussed in Section 2.1. Whether these satellites are real disconnected subdominant clusters, really existing in the landscape of finite instances, or they are just a manifestation of the rough geometry of the cluster frontier region is still an open question, which will be addressed in ongoing investigations.

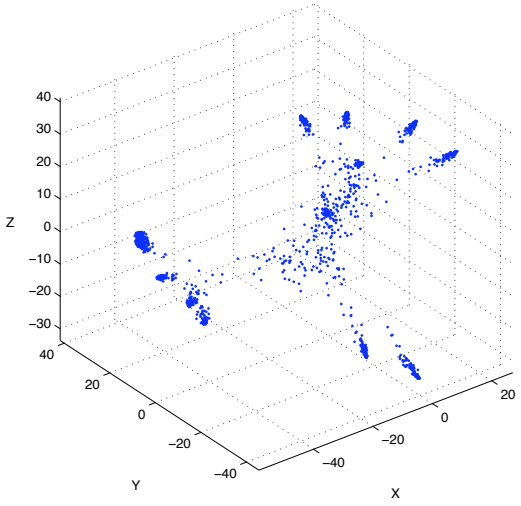

FIGURE 8. PCA clustering in the gap region. Despite the noise due to poor statistics, filamentous bridges connecting the main cluster region (not shown) to the sea of outer clusters seem to be identifiable.

Another open issue is relative to the existence itself of true gaps. The results available until now justify the claim that the probability of finding solutions at intermediate distances between the inter-cluster and the intra-cluster scale is rather small. Nevertheless, the observations do not rule out the hypothesis of the existence of a system of "filaments" of solutions, of subdominant measure when compared to the fraction of solutions belonging to groups. These filaments could actually create thin bridges between more or less close solution groups, destroying any true clustering, but creating a sort of spongy network of bubbles. No rigorous results about the existence of complete gaps exist indeed for K-SAT if $K < 9$ [**101, 121**]. Furthermore, arguments against an exact clustering can be taken by the analysis of the so-called whitening dynamics [**29, 91**]. PCA clustering analysis could be helpful in identifying the existence of such structures. In Fig. 8, we show very preliminary results relative to the clustering of solutions in the gap region. We built a normal three-generation data-set starting from a single seed $\vec{\sigma}$, but then we selected just the subset composed by the second and third generation offsprings of the first generation solutions $\vec{\sigma}_\delta$ with $0.25 < \delta < 0.4$. PCA clustering over this sample is rather noisy, but it is possible nevertheless to identify filamentous structures, which originates from clusters grains associated to "tentacles" at the surface of the main cluster and converge toward a central region containing distant clusters solutions, creating

a curious imploding firework look. Other clustering techniques, like superparamagnetic clustering [**21**] could be used to define more rigorously the boundaries of individual clusters.

3. Information-theoretical applications

The cluster-selection transition is quite sharp and produce an extremely quick polarization of the local fields over the variables belonging to the backbone of a cluster. This property can be exploited for practical purposes. The framework of Information-theory offer some example of nice applications. It is remarkable that the presence of clustering, usually considered an obstacle to an efficient optimization, can be used positively as a new computational resource.

3.1. Reconstruction of solutions. We have seen that the application of a random forcing allow to determine a close solution if the forcing intensity is not too large. On the other hand, if the forcing direction coincides with a solution, very large forcing strengths can be applied (eventually even $\pi = 1$): the result will always be a solution belonging to the same cluster of the forcing direction.

One may wonder if cluster selection can be obtained by applying a forcing only along a part of a complete solution. Since there exists at least one solution in which the selected variables are oriented according to the forcing direction, strong forcing intensities should be sustained, and if the information contained in the partial forced assignment is enough to fully identify a precise backbone, selection of an individual cluster should take place.

Let us consider a given instance and one solution $\vec{\sigma}$ and let us denote by $\vec{\sigma}^{(R)}$ the vector composed by the first NR entries of $\vec{\sigma}$. We call R *rate of compression*[4]. We apply now a really strong partial forcing in the direction of $\vec{\sigma}^{(R)}$ ($\pi \simeq 0.95$). If R is larger than a critical value R_c depending on the instance under consideration, all the non forced variable becomes immediately polarized. In particular, a complete solution $\vec{\sigma}'$ belonging to the same cluster of $\vec{\sigma}$ can be obtained *without* resorting to decimation, just by fixing all the variables according to their local field orientation after the first convergence. This behavior is the ultimate source of inspiration for the distributed implementation of SP-ext that will be presented in Section 1 of the next Chapter 6. If $R < R_c$, on the other hand, the configuration obtained by fixing all the variables after the first convergence has generally a finite energy.

[4]The quantity usually denoted by R in information theory (see appendix 2) is the inverse of our definition of rate.

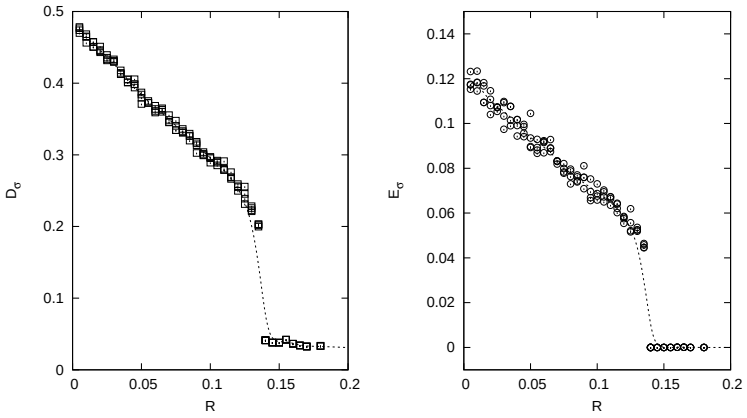

FIGURE 9. Solution reconstruction. The distance $D_{\vec{\sigma}}$ between the reconstructed configuration and the original solution, is plotted on the left against the compression rate. On the right, we plot the energy $E_{\vec{\sigma}}$ of the reconstructed configuration. If $R > R_c$ the cluster containing $\vec{\sigma}$ is selected. $E_{\vec{\sigma}}$ goes exactly to zero, differently from $D_{\vec{\sigma}}$ indicating that a cluster is selected but not a solution.

The results of an experiment performed over a $K = 5$, balanced, $\alpha = 16.8$ are shown in Fig. 9. Cluster selection can be observed for R as small as $R_c \simeq 0.15$, both looking at the distance plot (main plot) and at the energy plot (inset). Interestingly, the energy $E_{\vec{\sigma}}$ of the retrieved configuration falls immediately to zero, while $D_{\vec{\sigma}}$ jumps to a small value but only slowly decreases. This is compatible with the fact that SP-ext is able to select clusters more than individual solutions.

Further improvement of the critical rate will be obtained by resorting to the doping stage, as discussed in the next pages. We conclude this section remarking that SP-ext solution reconstruction could be used to implement a new kind of associative memories if it was possible to generate instances of CSPs with a predefined set of solutions. The resulting learning issue is still completely open.

3.2. Lossy data compression. A large part of information theory deal with the determination of alternative encodings for the same information content (a short introduction to Information Theory can be found in Appendix C, see also [**89**]). While more redundant encodings could result in increased robustness against noise during data transmission, more concise encodings would reduce the amount of space needed in order to store information. In many applications, where the space resource is critical or

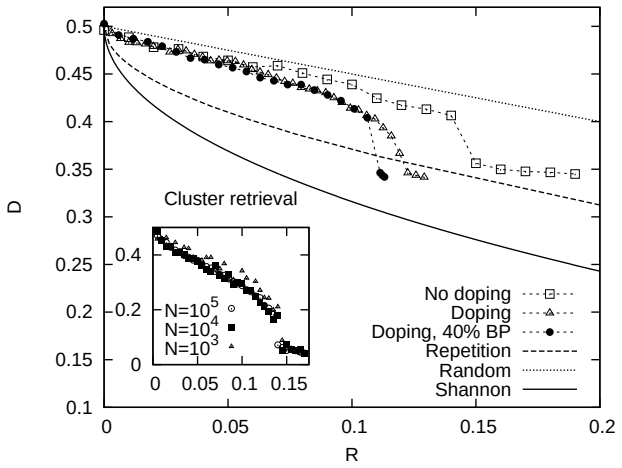

FIGURE 10. Rate-Distortion profile for the compression of a random unbiased source. Source reconstruction can be carried out thanks to the cluster-selection transition. The Shannon bound, and performance of random guessing and repetition decoding are shown too for comparison. The scaling of the cluster-selection transition associated to a successful decoding is also shown in the inset.

where the size of the data to compress is really large (notably multimedia contents), one can even accept a small pert of information in order to achieve larger compression [**142**]. This is the realm of lossy data compression (LDC) and of *rate-distortion theory*, establishing precise bounds about the better compression rate that can be achieved for any given maximum allowed distortion (the so-called Shannon bound).

A new kind of lossy data compressor (LD codec) can be devised, exploiting the cluster-selection transition. Let us suppose to have an input N-bit binary string $\vec{\xi}$ generated for simplicity from an unbiased and uncorrelated random source. Let us take then some balanced even connectivity K-SAT instance and let find by SP-ext decimation a solution $\vec{\sigma}_\xi$, by applying a forcing of the largest possible strength in direction $\vec{\xi}$. The distance $d_H(\vec{\sigma}_\xi, \vec{\xi})$ will be called the "a priori distortion" and it will coincide with the best possible distortion achievable after decoding by our LDC scheme. If the cluster distribution is homogeneous we can expect it to be of the order $d_0 - \Delta$. Furthermore, we take α only slightly larger than α_G in order to maximize the number of well separated addressable clusters, and minimize then the a priori distortion. At this point, we build a compressed string $\vec{\sigma}_\xi^R$ by retaining just the spins of the first NR variables of $\vec{\sigma}_\xi$. This concludes the basic encoding stage (compression).

Decoding is now performed resorting to the SP-ext solution reconstruction capabilities. A very intense forcing ($\pi = 0.95$) is applied along $\vec{\sigma}_\xi^R$ and after the first convergence all the variables are fixed according to their local field orientation. If $R > R_c$ a solution $\vec{\sigma}_\xi^{\text{retrieved}}$ belonging to the same cluster of $\vec{\sigma}_\xi$ is generated without creating contradictions. We consider such $\vec{\sigma}_\xi^{\text{retrieved}}$ as decoded string. The distortion accumulated with respect to $\vec{\xi}$ is of the order of $d_{\text{cl}} + d_0 - \Delta$.

The data relatives to the compression and decompression of an unbiased and uncorrelated source with a balanced instance of $N = 3 \cdot 10^4$ variables, with $K = 5$ and $\alpha = 16.8$ are shown in Fig. 10. The performance of a random guess decoder (deciding the missing bits at random) and of a basic repetition codec[5] [**89**] are also plotted for comparison. Although the behavior of the SP-ext-based codec is definitely not random, the achieved distortion is too bad.

In order to approach the Shannon Bound, we use the *iterative doping* [**34**] technique for chosing which bits to store. After the determination of $\vec{\sigma}_\xi$, SP-ext is run again *without* applying any forcing and a ranking of the most balanced variables is performed. More specifically, we look for the variable i which minimizes the index function $|W_i^+ - W_i^-| + W_i^0$. The highest ranked variables are then frozen in opposite directions in a similar number of clusters, are rarely unconstrained and consequently their fixing produces the largest possible information gain. We take as first bit of the compressed string $\vec{\sigma}_\xi^R$ the state s_{i_1} assumed in the solution $\vec{\sigma}_\xi$ by the highest ranked variable i_1. Then we fix i_1 to the state s_{i_1}. At this point a new doping step is done, the length of $\vec{\sigma}_\xi^R$ is increased of one bit, and the procedure is continued until when the desired compression rate is reached. An *identical* doping stage is then performed in decompression. During decompression, we read the first bit s_{i_1} in $\vec{\sigma}_\xi^R$ and we have to decide which variable has to be fixed to this state. The answer can be obtained, without need of transfering additional information, just by repeating also in decoding the iterative ranking based on the doping index function. The highest ranked variable (that will be again i_1) is fixed to the state s_{i_1}, then we repeat the step for the second variable, and so on, until when the compressed string $\vec{\sigma}_\xi^R$ has been completely read. At this point the SP-ext iterations are used to fix the remaining variables. Fixing a balanced variable "switches off" a

[5]Encoding works by subdividing the source in blocks of an even number of bits λ, and then substituting each block with a single bit decided by majority rule. Decoding expand a single bit to a word given by λ repetitions of the same bit. The rate of a repetition codec is set by the choice of the block-length λ. See also the Appendix C for more details about repetition codes.

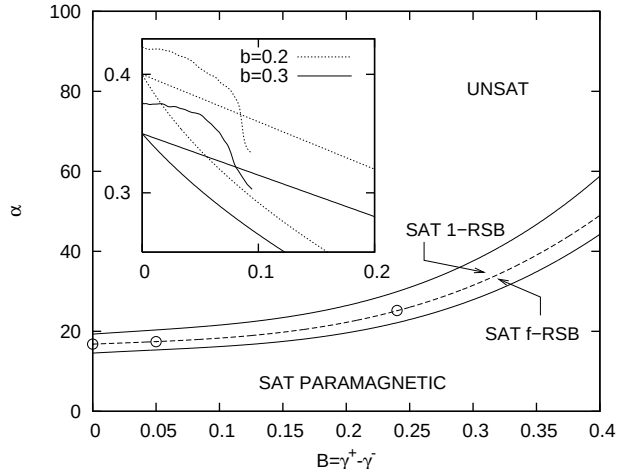

FIGURE 11. Phase diagram for fixed degree biased 5-SAT and compression performance for two biased sources.

larger number of clusters, and, hence, thanks to the less redundant coding, the critical rate is reduced to surprisingly small values (dash dotted curve in Fig. 10).

A further improvement can be obtained, by using in the doping stage modified iterations [**26, 91, 161**], that interpolates between SP and the well known Belief Propagation equations (solid curve in Fig. 10). It is supposed that this interpolation gives access also to microscopic information relative to the inter-cluster biases experimented by a variable. The use of the interpolated iterations allows to approach the Shannon Bound more than the simple repetition LD codec. It is nevertheless expected that much better performances can be achieved by a careful optimization of the graph ensemble as it was shown in the case of the iterative decoding with Belief Propagation of LDPC codes [**39, 90, 139**]. it is on the other hand surprising that the cluster-selection transition allows for a strongly non-trivial information manipulation over any typical random instance in a clustered phase.

SP-ext can also be applied to the encoding/decoding of *biased* sources, characterized by inhomogeneous occurrence probabilities $P_{\pm}$ for the (two) possible input symbols . The issue is not just academic but is also of practical interest since correlated sources can often be mapped to memoryless biased sources by mean of the Burrows-Wheeler Transform. [**33**]. Let us suppose that the bias of the source is $b = P_+ - P_- > 0$. We can engineer graphs with a cluster distribution concentrated in a hypercone around the ferromagnetic direction, just by taking for every variable i a fraction γ_+ of positive couplings larger than the fraction γ_- of negative couplings.

The phase diagram for biased fixed connectivity instances is shown in Fig. 11. The two control parameters are α and the balancing $B = \gamma_+ - \gamma_-$. A narrow SAT RSB stripe, characterized by a finite value of the 1-RSB complexity, is present for $B < 0.435$. When the balancing is too large, there is on the other hand a direct transition between an unfrozen SAT phase and the UNSAT region. The transition line is determined when possible by finding the point in which zero-energy phase complexity vanishes. For larger balancings, the line where WalkSAT starts to fail in determining a solution approximately continues the 1-RSB SAT/UNSAT transition line. The rate-distortion profile of the compression of two random uncorrelated sources with $b = 0.2$ and $b = 0.3$ is shown in the inset of the same figure. We have experimentally verified that the best decoding performances were obtained selecting graphs in proximity of the Gardner instability line (always determined by comparison with SP-2 marginals), while better critical rates but worse distortions are found for larger connectivities. The location of the employed instances is indicated on the phase diagram by empty circles.

The inset of Fig. 11 shows that cluster-selection is still possible and ensures again a non-random correlation between the input and the output sequence. The performance is nevertheless poorer than in the unbiased case. Once again it is expected that the codec results can be improved by mean of graph optimization.

CHAPTER 6

What next?

The algorithmic strategies based on the Survey Propagation formalism have completely fulfilled the goal of overcoming the barrier created by the proliferation of metastable states. Not only SP-Y has been successfully employed to generate configurations of the K-SAT model well beyond the lower bound to the threshold, but the introduction of external messages has allowed to develop the reliable SP-ext state-addressing protocol. The decimation process can be controlled without affecting the overall computational cost of the SP iterations, pushing it to retrieve assignments in a precise restricted portion of the phase-space. Lossy data compression has been proposed as a possible application, but the potentialities of the method go well beyond this first example and deserve new explorations in the immediate future.

It is worth to remind that the self-consistent computation of the pdfs for the cavity biases and for the local fields is a *fully local* procedure. Knowledge about all the variables is spread around the graph by mean of message-passing, but at each time-step only the information broadcasted by the neighboring nodes is needed for the update of all the outgoing surveys. The only step in both SPID and SP-ext requiring a global simultaneous awareness of the tendencies of all the variables is decimation. Nevertheless, we have already shown that, at least in cluster-reconstruction experiments, decimation becomes unnecessary thanks to the strong polarization induced by the SP-ext cluster-selection transition. In the next section 1 we will show how an adaptive forcing can be used to retrieve systematically solutions in the hard SAT region, after the computation only of a single population dynamics fixed point [**38**]. This new completely local algorithm will allow to find a solution in order $N \log N$ time step, against the quadratic cost of SPID and SP-ext (each one of the N computations of a fixed-point scales linearly with N). Furthermore, the new method seems to be particularly suitable for implementation over reconfigurable hardware devices. In this latter case, the cost for retrieving a solution would scale linearly with just $\log N$.

A further step toward the realization of efficient distributed algorithms for glassy states selection could be possibly the use of external forcing in

association with a pure warning propagation dynamics (WP). In section 2, we will show indeed that external trains of messages with fixed or variable frequency can be used to stabilize the evolution of warnings allowing them to converge toward solutions or toward nearby excited states well separated by the layers typically trapping WP.

Speculations about some possible implications of this behavior for the working of neural network associative memories will be briefly mentioned in the last section 3, and will finally close this incomplete but hopefully informative journey through the many and diverse applications of the Survey Propagation paradigm.

1. Distributed implementations

The equations (5.7) show how the effectiveness of an externally applied forcing is larger when there is a significant fraction of locally well-oriented clusters. The decompression strategy of the SP-ext codec provide an example in which the application of a "good" forcing remove the need for decimation. In the general case of a random forcing direction $\vec{\phi}$ not correlated with any solution cluster, the temporary configuration obtained by fixing all the variables according to their local field direction at the SP fixed-point contains a residual fraction of violated constraints. This is mainly due to the fact that the presence of many locally wrongly oriented clusters does not allow the use of strong forcing intensities. This fraction may be possibly reduced by adaptively building a particularly effective forcing pattern [**38**].

A first method (which we may call *extrinsic adaptation*) consists in applying initially a forcing along a random direction $\vec{\phi}$ and in observing which are the nodes that are reacting positively to its application. In the cases in which the system develops an internal local magnetization in agreement with the local direction ϕ_i, the forcing is kept active, but it is on the other hand *switched off* (with a certain probability) if the internal magnetization tends to contrast the action of the forcing. A second method (*intrinsic adaptation*), only slightly less effective, does not require the use of an external reference direction, but the forcing direction is always taken in agreement with the orientation of the internal local magnetizations.

In both cases, the applied adaptive forcing pattern is updated every predetermined fixed number of steps looking at the developed distribution of the internal fields. The general larger correspondence between the forcing direction and the spontaneous drift of the surveys, due to the topology itself of the network, allows in general the use of stronger forcing intensities and cluster-selection can finally be observed. Under application of adaptive forcings, full satisfying assignments are generally retrieved at the very first

INPUT:: a Boolean formula $\mathcal{F}$ in conjunctive normal form; a forcing strength π and an update frequency f; optionally an external reference direction $\vec{\phi}$

OUTPUT:: a full configuration of the variables appearing in $\mathcal{F}$ (ideally a satisfying assignment)

0. For each edge $a \to i$ of the factor graph, randomly initialize the $\eta_{a\to i} \in \{0, 1\}$
1. Compute all the fixed point u-surveys, using equations (5.2)
3. IF the step number is a multiple of f,
   3.1 FOR every unfixed variable i, compute the local magnetization m_i
   3.2 *Intrinsic method:* Set the forcing directions $\phi_i = \mathrm{sign}(m_i)$ and the forcing strengths $\pi_i = \pi$, otherwise $\pi_i = \pi$
   OR
   3.2 *Extrinsic method*: IF $\phi_i \neq \mathrm{sign}(m_i)$ set $\pi_i = 0$.
4. Fix all the variables to the states $s_i = \mathrm{sign}(m_i)$ and compute the configuration energy E; increment the step counter
5. GOTO (0.) UNTIL when $E = 0$ or the maximum allowed number of iterations has been reached.

FIGURE 1. The Adaptive SP online decimation algorithm (AdaSP).

convergence. In a sense, the adaptive forcing mechanism is equivalent to a runtime decimation with backtracking, where the variables are fixed just by imposing over them a strong field. For these reasons, we call the new algorithm *adaptive SP online decimation* (abbreviated as AdaSP).

In Fig. 2, we show the evolution of the instantaneous configuration $\vec{\xi}_t$ of the variables (according to the observed instantaneous magnetizations) during AdaSP online decimation with the extrinsic adaptive method of a $K = 5$, $\alpha = 16.8$ balanced sample. In a first stage, $\vec{\xi}_t$ rapidly converges toward the forcing direction $\vec{\phi}$, but, then, in order to obtain a further energy reduction, many forcing components must be switched off. This relaxation of the external constraints is responsible of the observed increase of the distance between $\vec{\xi}_t$ and ϕ, but remove any obstacle to selection of a solution cluster much close to $\vec{\phi}$. The observed a priori distortions are slightly better, but perfectly compatible, with what observed in SP-ext experiments over the same instances.

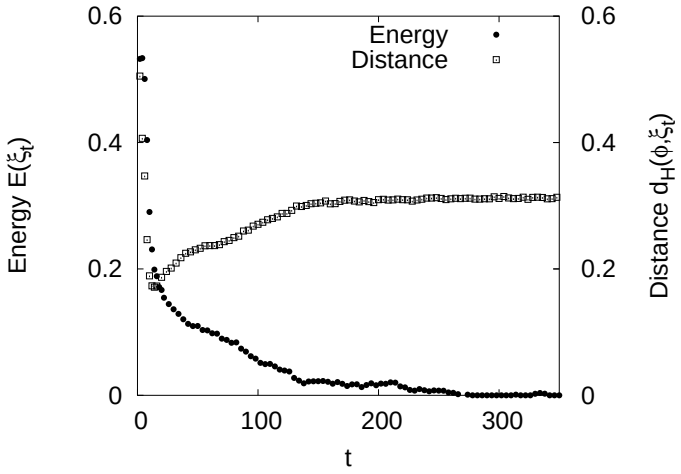

FIGURE 2. Adaptive SP online decimation (AdaSP). The distance of the instantaneous configuration $\vec{\xi}_t$ and its overlap with the applied forcing $\vec{\phi}$ are plotted against the number of sweeps.

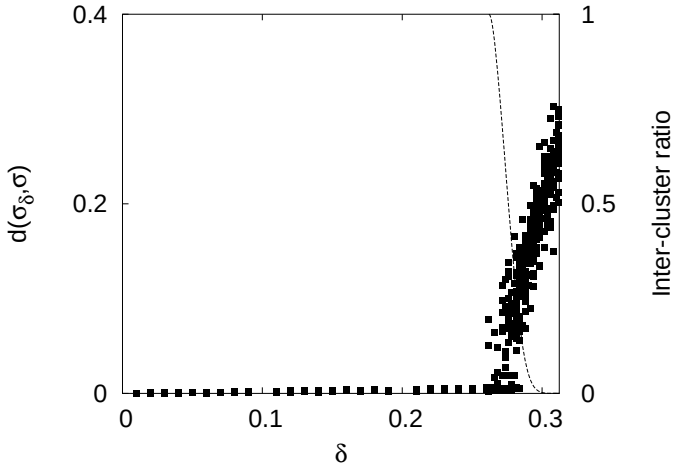

FIGURE 3. Stability diagram obtained with AdaSP. The ratio of found solutions falling inside the original cluster is also plotted for reference and falls to zero after a critical perturbation.

AdaSP cluster selection share several properties with the analogous mechanism observed in SP-ext. in particular stability experiments can be repeated, by adopting as reference forcing directions ϕ increasingly perturbed solutions. Fig. 3 shows the results obtained for the same $\alpha = 16.8$ balanced 5-SAT instance analyzed in Fig. 3. We see that the critical perturbation at which half of the generated assignments lies outside the cluster of $\vec{\sigma}_0$ has

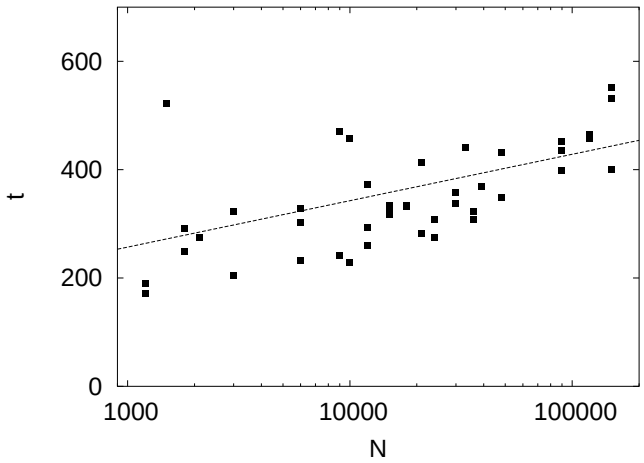

FIGURE 4. Logarithmic scaling of computational time. The data refer to the number of complete sweeps needed in order to find a solution with AdaSP. It would correspond with the true time cost of resolution using a FPGA electronic implementation of AdaSP.

now the slightly smaller value of $\delta_c \simeq 0.265$. Nevertheless, the sharpness of cluster-selection is definitely larger and the distant clusters branch begins with an almost vertical slope, after a very narrow gap. This feature may be used for implementing a very clean solution reconstruction algorithm.

We finally remark that AdaSP is a completely local approach to solution finding, and can therefore efficiently implemented over a distributed computing environment (GRID, etc.) or over massively parallel hardware. In particular, it is easy to write VHDL virtual hardware description for the implementation over reconfigurable logical-gates-arrays (like FPGAs) [**19**]. Simulations of a FPGA chip configured to implement the AdaSP algorithm show that over such a reconfigurable hardware the resolution cost would scale only with the logarithm of the instance size(see Fig. 4). For random CSPs, the information carried by local messages has to propagate through the whole graph before convergence, and this cannot be done faster than in a logarithmic time. Therefore, the solution retrieval time cost over distributed devices is, at least in terms of asymptotic scaling, the best possible.

Of course, hardware implementations suffer of important limitations, like the generally slow sustainable clock speeds and the large cost of the chips able to contain large instances[1]. Nevertheless, more powerful and economic devices are going to be developed in the next future at a fast pace

[1]Simulations for $N > 3000$ have been done with a C code, since they are beyond the technical possibilities of today available hardware.

and it is quite likely that they will play an increasingly important role in the field of high performance computation.

Moreover, the time cost of solution retrieval scales with $N \log N$ also when taking in account the cost of the simulation of the parallel hardware over an ordinary computer. This makes AdaSP one of the fastest existing incomplete solvers for hard random K-SAT instances, independently from the possibility of running it over a distributed computation platform.

2. Selecting states with Warning Propagation

We started our discussion of the Survey Propagation equations by observing that the SP marginals describe the probability of finding a certain warning traveling along a certain edge of the factor graph when considering all the Warning propagation (WP) fixed points associated to all the solution clusters.

No direct relation exists on the other hand in general between SP marginals and the time-series of the warnings generated during the evolution under WP. Many coexisting fixed points will compete for dominating the warning dynamics; the messages broadcasted by distant nodes will be associated to configurations pertinent to distant different solution clusters. As a result of this mismatch, the configurations associated to an instantaneous set of messages are typically characterized by an extremely high energy.

We have seen discussing the relations (5.7) that the main effect of the introduction of forcing is the induction of a non homogeneous measure over the different clusters. If only a limited number of clusters were allowed to broadcast messages driving toward them, global correlations could more easily be established and agreement among all the warnings may possibly be reached. We will see now external sources of warnings can be used to efficiently enhance the role of certain clusters with respect to the others. SP-ext forcing was implemented by mean of additional surveys of a predefined strengths π_i. In the case of WP, warnings have a digital nature (a message can just being or not being sent), but it is possible to tune accurately the frequencies ν_i of external sources of warnings, connected to the variables by a predefined set of couplings in direction $\vec{\xi}$. We will denote as WP-ext the resulting forced WP dynamics.

A Montecarlo simulation of the evolution of the warnings can easily be performed. At each time step the external warning sources inject or not in the system their pulse with a probability given by the preassigned frequencies ν_i. The messages are updated synchronously according to the

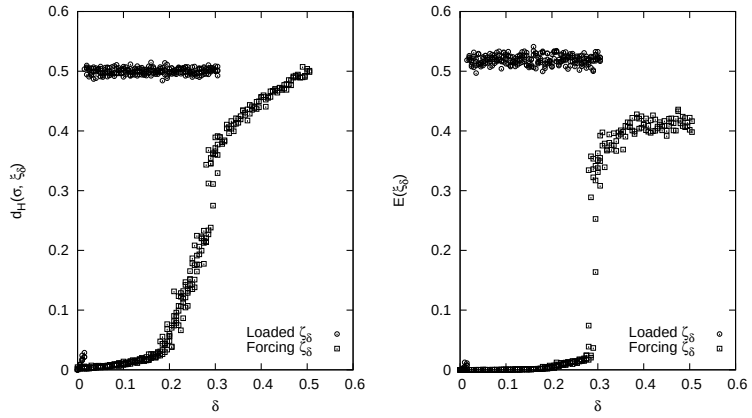

FIGURE 5. Warning propagation with extrinsic forcing ($K = 5, \alpha = 16.8$ balanced instance). The stability diagram construction experiment has been repeated with ordinary WP, loading perturbed solutions $\vec{\zeta}_\delta$ as initial configurations, and with WP-ext, starting from random initial conditions but applying a forcing along $\vec{\zeta}_\delta$ with a fixed frequency $\nu = 0.7$. Both when considering geometrical (*left*) and energetic (*right*) stability, the critical perturbation δ_c is comparable with the corresponding value obtained in SP-ext stability experiments.

dynamical equations (2.14) and (2.42). The configuration of the variables is determined according to the local instantaneous magnetizations[2].

The direction $\vec{\xi}$ of the forcing can be preassigned (extrinsic forcing) or adaptively determined by looking at the internal tendencies as in the case of intrinsic AdaSP. Let us first consider the extrinsic case. The application of a small frequency forcing along the direction of a solution $\vec{\sigma}$ is enough for filtering out completely after a certain number of updates all the contributions coming from different clusters. Stability experiments can also be performed, by forcing along a direction $\vec{\zeta}_\delta$ obtained by flipping a fraction δ of the spins in $\vec{\sigma}$. If δ is not too large, WP-ext reaches a fixed point after a rather small number of updates or enters into a limit cycle. An important difference with SP-ext is indeed the fact that it is not guaranteed that WP-ext finds always a fully satisfying assignment. Nevertheless, even in the case in which WP-ext enters a limit-cycle concatenating a small group of excited configurations, the energies are at least an order of magnitude smaller than in the case of the states trapping ordinary WP. Furthermore,

[2]It is eventually possible but definitely not needed to dampen fast fluctuations by averaging the magnetization over a small window of observation.

these excited configurations are geometrically close to the reference solution $\vec{\sigma}$. This means that the basins of attraction of WP-ext are not merely composed by the solutions clusters, but comprehend also surrounding regions of moderately excited configurations. Since they are geometrically pertinent to the solution cluster region, one can safely affirm that entering a limit cycle composed of such peripheric states is another way of selecting the state to which $\vec{\sigma}$ belongs.

In Fig. 5 at right it is indeed possible to observe two well distinguished regimes, depending on the perturbation entity δ. If $\delta < \delta_c \simeq 0.29$ only configurations with a very small energy are obtained at the equilibrium. When, on the other hand, δ exceeds the critical value, the energy explode and extremely excited configurations are visited by WP-ext. In Fig. 5 at left the same loss of stability phenomenon is monitored looking at the Hamming distance from the reference solution. All the obtained low energy configurations are also significantly correlated with $\vec{\sigma}$. Here, a transition regime between $\delta_{c,1} \simeq 0.18$ and $\delta_c \simeq 0.29$ where the distance increases more rapidly is also observed. Then, after the loss of stability, progressively uncorrelated configurations are generated.

The effectiveness of the external forcing can be underlined by conducting a second stability experiment in which no forcings are applied, but in which $\vec{\zeta}_\delta$ are just loaded as initial conditions of pure WP experiments. In this latter case, stability is lost after a noticeably smaller critical strength of the perturbation (circle points in Fig. 5). The external forcing plays then an important stabilizing role over the WP dynamics.

It must finally be remarked that, while SP-ext always find a solution —if the forcing direction is too distant from the reference cluster, then solutions inside different clusters are addressed—, in the case of WP-ext the alternative to state selection is simply absence of selection.

Let us now consider the case of intrinsic adaptive forcing. As in the case of AdaSP, the local direction of the forcing is determined by looking at the internal instantaneous magnetization. It is important to update the forcing direction less frequently than the warnings themselves in order to obtain actual cluster filtering (information has too spread in order to produce a matching between the warnings broadcasted by distant emitters). The way in which the energy is reduced is shown in Fig. 6. The inset shows the number of downhill directions surrounding the visited configurations, against their energy. The slowdown of energy relaxation is associated to a transition in the topology of the visited portion of landscape (visible from the strong change of slope in the inset). Intrinsic adaptive WP-ext seem to

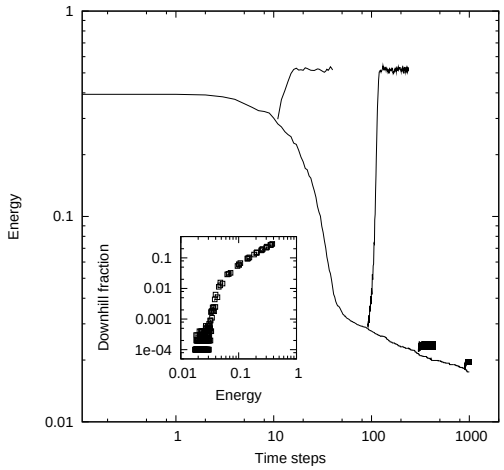

FIGURE 6. Energy evolution for intrinsic adaptive Warning Propagation ($K = 5, \alpha = 16.8$ balanced instance). Warning exchange allows the system to reach states with progressively smaller energies. As shown in the inset, the visited configurations with lower energies are characterized by a reduced number of downhill directions and are therefore close to the bottom of some local minima, attracting the warning propagation dynamics also when the external forcing is removed. On the other hand, is the energy is larger, in absence of a stabilizing forcing, the warnings dynamics is not persistent but escapes toward highly excited configurations.

be captured by local minima structures (with a large number of residual flat directions).

If the adaptive forcing is suddenly switched off, different behaviors are observed according to the already reached energy (and the associated topological properties of the neighborood). If the energy is too large, WP lose completely stability, the warnings cease to be associated coherently to a small number of clusters and push the system into the energy band that would have been reached performing pure WP from the beginning without the use of a forcing. On the other hand, if the energy is small enough, WP get trapped by the local minima even in absence of forcing, and *persistent dynamics* is observed (at least for a certain time in a transition region). This is a highly desirable feature in the context of associative short-term memory modeling and we will briefly return on this point in the next concluding section.

Further refinements of intrinsic adaptive WP-ext are needed in order to make it able to retrieve satisfying assignments, and not just low-lying energetic configurations. Nevertheless, the use of fixed frequency forcings has endowed WP with unexpected state selection capabilities. An important advantage of WP-ext strategies with respect to both SP-ext and AdaSP, is that it makes use of binary elementary messages instead of the real numbers parametrizing the surveys. This extreme simplicity would be considerably helpful in the eventual electronic implementation of WP-ext-like strategies, and could open the road to new applications in artificial pattern recognition. Possible implications for neural computing will be outlined in the next section.

3. A new kind of neural networks?

The fact that a simple dynamical scheme based on the exchange of elementary messages is able to exhibit sharp cluster-selection transitions (at least if weak external conditionings are also taken into account) is somehow surprising. These preliminary observations suggest the hypothesis that similar mechanisms are used by biologically realistic networks of spiking neurons acting as associative memories (retrieval would correspond to the selection of a specific cluster).

The experiments described in the preceding section could be interpreted as describing an elementary example of *pattern recognition*. More in general, state selection in clustered phases could be a relevant feature for the attractor neural network modeling of neural processes [**7, 48**]. The term *working memory* refers to active storage mechanisms which provide real-time information to be used for the execution of willful actions [**8, 31, 32**]. The sustained firing rate found in higher cortical areas during delayed response tasks is generally considered to be the physiological analogue of this mental process [**112, 61**]. In the WP-ext experiment described in Figure 6, persistent activity is observed also when the forcing stimulus is removed if a state is selected. The large gap existing between the two bands of the states visited, respectively, in presence or in absence of cluster selection creates indeed a sharp "digital-like" distinction between the two possible outcomes of a stimulus presentation, "*Pattern retrieved*" or "*Not retrieved*". Furthermore, since slightly perturbed stimuli produce an association with the same neighborood in the configuration space, the organization of states in clusters produces positively a considerable reliability against noise.

Once the dynamics has been trapped in a precise attracting region, the system can escape the region surrounding the ground-state cluster whenever a new stimulus acting as perturbation is applied. The network might then

enter into a new attracting state, non-trivially correlated *both* with the new stimulus and with the previously visited attractor. Interestingly enough, a similar *latching* dynamics has recently been proposed as supporting mechanism for evolved neurocomputing tasks like language generation [**62, 156**].

Of course, informations about neural processes cannot be extrapolated from experiments over factor graphs of K-SAT instances. A crucial goal becomes then the identification of clustered phases in biologically realistic models. A first and reasonably simple model for a neuron can be represented by a basic threshold node performing the weighted sum of γ_{in} inputs:

$$s_{out} = J_{out}\theta(\sum_{j\in In(i)} J_j s_j - \tau_j) \tag{6.1}$$

The synaptic weights are taken to be binary ($J = \pm 1$) and the single output connection (inactive or active inhibitory/excitatory depending on the node) is multiplied and sent to γ_{out} receiving nodes. Even if an energy function is not strictly defined for threshold nodes implementing a dynamical rule more than a local constraint in the usual sense of CSPs, an hamiltonian can be introduced resorting to a trick first introduced in [**41**]:

$$\mathcal{H}_j = \begin{cases} 0 & \text{if } s_j = J_j\theta(\sum_{j\in \text{In}(i)} J_j s_j - \tau_j) \\ 1 & \text{otherwise} \end{cases} \tag{6.2}$$

A node is not satisfied only if the local configuration of the variables violates the causal order imposed by the threshold dynamical rule. The ground states of a network of nodes (6.2) are in biunivocal correspondence with fixed points of the dynamics (6.1) over the same network. Message-passing optimization can then be used to characterize the statical properties of the model neural dynamics, and in particular to estimate the number of fixed points and of eventual clusters of fixed points.

It must be remarked that both ordinary Belief Propagation [**161**] and Survey Propagation are based on message exchange over a non-directed graph, while the dynamics (6.1) is inherently directed, as visible also by the non-symmetric nature of the hamiltonian (6.2). The signals actually exchanged between the neuronal units would not then correspond with the messages whose statistics is being sampled by BP or SP. Message-passing (online) decimation is then an artificial and not necessarily local procedure that can be used to identify the existence of phases with a large number of addressable attractors, but a more realistic description of state retrieval by the network will be obtained resorting to algorithms similar in spirit to a directed version of WP-ext.

Research is ongoing in order to determine network topologies suitable for cluster-selection. Relevant order parameters are the average connectivity γ, the fraction of leaves λ and the fraction of inhibitory connections ι. Non trivial distributions of the magnetizations are obtained for $\iota \simeq 0.5$, but preliminary investigations seem to suggest that more complex neuron models are needed in order to produce complex state organization. Networks of neurons in which the integrating properties of individual dendritic compartments are explicitly taken into account [**97**] may constitute an interesting and quite natural field of exploration.

Preliminary results are also available relative to the use of message-passing techniques for batch learning of patterns in a binary-synaptic-weights perceptron. An AdaSP-like online decimation protocol is used for retrieving assignments in the coupling space, endowing the perceptron with the ability of discriminating a linear number of patterns.

These first promising indications fully justify further experimental and theoretical investigations. Major efforts are necessary in order to clarify the relevance of the message-passing paradigm in neural information processing, and to support eventually the visionary description of the brain as a powerful distributed hardware running sophisticated message-passing algorithms as software.

Conclusions

Our exploration of complex optimization problems started from the consideration that phase transitions in the topology of the ground-state space are responsible of the bad performance of many complete solvers and heuristics. The proliferation of metastable states was shown indeed to constitute a practically insurmountable barrier, even when considering the most performing local search algorithms or exotic semi-classical approaches, like simulated quantum annealing.

The development of an iterative procedure deriving from the cavity method at the level of 1-RSB allowed us to compute the distribution of the local magnetizations over single problem instances. Such microscopic information could be firstly exploited in the SPID and in the SP-Y strategy that were successful in determining in polynomial time complete assignments in the hard-SAT region and configurations well beyond glassy threshold in the hard-UNSAT phase.

A further innovative step was represented by the SP-ext method, based on the use of external messages able to tune accurately the convergence behavior of standard Survey Propagation. The new technique could be used to generate a large (ideally exponential) number of solutions of really hard instances. The existence of a significant correlation between the activated pattern of forcing and the retrieved solution was employed to characterize the local geometry of the ground-state phase. More specifically, we were able to discriminate experimentally between samples extracted from the 1-RSB region, presenting sharp clustering, and samples from the unstable region, showing a continuous spectrum of possible overlap measures between the minimum and the maximum distance scale.

Quite surprisingly, the cluster selection phenomenon opened the road to new algorithmic strategies for which the presence of an exponential number of well separated clusters constituted a positive resource, instead of an insuperable obstacle. We examined briefly a preliminary application to the lossy data compression of random uncorrelated sources. The compression of biased sources achieved by the generation of suitably unbalanced samples was a first simple example of the possibility of engineering graphs with an *ad hoc* phase-space geometry, tailored to supply better results in predefined

tasks. Furthermore, it is quite likely that, in a near future, the geometry-sampling capabilities of SP-ext can be exploited to provide a better knowledge of the landscape features responsible for the superior performance of exceptionally good linear codes.

An important byproduct of the research on external forcings was the discovery that decimation is not actually needed for the determination of solutions. The iterative fixing of variables can indeed be substituted by the application of local stabilizing fields, that are adaptively updated in order to reduce the fraction of locally misaligned clusters. The onset of a cluster-selection transition induces strong polarization and the system collapses spontaneously in a single pure state, without resorting to global sorting and artficial fixings. The result is the AdaSP algorithm, which is able to retrieve complete satisfying assignments in the hard-SAT region in a time which scales just with $N \log N$ (or only with $\log N$, when implementing it in a massively parallel computing environment, like GRID or reconfigurable FPGA hardware).

The efficiency of the fully local AdaSP approach is remarkable and casts new light over the computational potentialities of networks of simple message-exchanging units. Quite surprisingly indeed, we were able to observe state-selection behaviors in simple warning propagation dynamics over random K-SAT factor graphs, subjected to external fixed-frequency trains of stimulations (WP-ext). These findings, among others, have pushed us to explore the possible relevance of state-selection phenomena for neural computation. Besides artificial intelligence applications, brain itself represents indeed the ultimate example of massively parallel hardware. Its component units, the neurons, realize complex computational and cognitive tasks through an intricate exchange of messages which follows a still largely unknown protocol.

Behaviors reminiscent of evolved neural features like working memory and latching have been observed in WP-ext experiments over random K-SAT graphs, which are far removed from ordinary models of neural networks. Yet, any dynamical rule for the local update of neuronal states can be mapped to an associated optimization problem over a non-directed factor graph. Techniques like forcing-driven BP or SP can then be used to characterize the ground states of the resulting CSP, which are in one-to-one correspondence with the fixed-points of the dynamics. It becomes then possible to investigate the phase diagrams of networks of arbitrarily complex and biologically realistic models of neurons, the only restriction being represented by the discrete nature of the admissible cell states.

It is worth to remind, however, that, in the recent years, many experimental indications have been found supporting the evidence of discrete jumps in the synaptic connection strengths of real neurons. The efficacy of a biological synapse is indeed bounded and quantized at the latest level of the release of single vesicles of neuro-transmitter [**57, 138**]. Remarkably, message-passing techniques —in which the signals exchanged between elementary units vehicle probabilistic information— prove to be remarkably performing when applied to discrete neural models. An example is provided by batch and online learning for a perceptron with $J = \pm 1$ synapses. While ordinary algorithms can only be used to store a logarithmic number of patterns [**9**], rather simple message-passing protocols allow the retrieval of a linear number of them.

Work is in progress in order to achieve a better understanding of the relation existing between statistical inference methods analogues to BP and SP and the neural processing of information (both at the network and at the intracellular level). Developing a new concise framework able to capture the many subtleties of the neural code elaboration is one of the most important and fascinating challenges for today's research in interdisciplinary science. Our modest hope is that our contributions and their developments may help to grasp in the future at least some details of the working of the most complex system known in nature, our brain.

Part 3

Appendices

APPENDIX A

Replica approach to K-SAT

The partition function of a given instance of K-SAT can be formally written as:

$$\mathcal{Z}[\{J\}] = \sum_{\{s\}} \exp\left(-\frac{\mathcal{H}_{K\text{-SAT}}\left(\vec{\xi},\{J\}\right)}{T}\right) \tag{A.1}$$

The ground state energy E_{GS} is a random variable which in the thermodinamical limit concentrates around the following approximate estimate:

$$E_{GS} = -T\overline{\ln \mathcal{Z}[\{J\}]} + O\left(T^2\right) \tag{A.2}$$

The direct computation of the average over the quenched disorder is awkward, since the logarithm of the Partition function is not a self-averaging quantity, but it can be performed resorting to the so-called *replica trick*. Observing that $f(x)^n = 1 + n\, f(x) + O\left(n^2\right)$, we can write the formal identity:

$$\overline{\ln \mathcal{Z}[\{J\}]} = \lim_{n\to 0} \frac{\overline{\mathcal{Z}[\{J\}]^n} - 1}{n} \tag{A.3}$$

The trick consists in evaluating directly $\overline{\mathcal{Z}[\{J\}]^n}$ as the partition function of n independent replicas of the original system, and only then in doing an analytical continuation of n into the real plane and taking the limit $n \to 0$. If we write:

$$\overline{\mathcal{Z}[\{J\}]^n} = \sum_{\{s^{(1)}\},\ldots,\{s^{(n)}\}} \prod_{a=1}^{n} z[\{s^{(a)}\}] = \sum_{\{s^{(1)}\},\ldots,\{s^{(n)}\}} \overline{\exp\left(-\sum_{a=1}^{n} \mathcal{H}_{K\text{-SAT}}\left(\vec{\xi},\{J\}\right)/T\right)} \tag{A.4}$$

we remark that each individual term factorizes over the sets of the M clauses because of the absence of any correlation in the formula generation process:

$$z[\{s^{(a)}\}] = \left(\zeta[\{s^{(a)}\}]\right)^M, \quad \zeta[\{s^{(a)}\}] = \overline{\exp\left(-\frac{1}{T}\sum_{a=1}^{n} \delta\left(\sum_{i=1}^{N} J_i s_i^{(a)} - K\right)\right)} \tag{A.5}$$

where now the coupling constants J_i are dummy random variables over which an average has to be performed, and can eventually assume also the value $J_i = 0$, to indicate that a variable is not participating to a given random interaction. After several straightforward but lengthy manipulations, it becomes possible to express (A.5) as a function only of the 2^N occupation fractions $x\left(\vec{\sigma}\right) = x\left(\sigma^{(1)},\ldots,\sigma^{(n)}\right)$, equal to the number of labels i for which $s_i^{(a)} = \sigma^{(a)}$ for every $a = 1,\ldots n$ (each possible n-components vector $\vec{\sigma}$ labels one of the occupation fraction):

$$\zeta[\{s^{(a)}\}] = \zeta[\{x\}] = \frac{1}{2^K} \sum_{J_1,\ldots,J_K=\pm 1} \sum_{\vec{\sigma}_1,\ldots,\vec{\sigma}_K} x\left(-C_1\vec{\sigma}_1\right)\ldots x\left(-C_K\vec{\sigma}_K\right)$$
$$\times \exp\left(-\frac{1}{T}\sum_{a=1}^{n}\prod_{\ell=1}^{K}\delta\left(\sigma_\ell^{(a)} - 1\right)\right)$$

In the thermodynamical limit only the leading order in N can be retained and we can resort to *saddle-point integration*. We can write then $\overline{\mathcal{Z}[\{J\}]^n} \simeq \exp\left(-N f_{\rm opt}/T\right)$, where $f_{\rm opt}$ correspond to the *minimum* of the functional:

$$f[\{x\}] = e[\{x\}] + \frac{1}{T}\sum_{\vec{\sigma}} x(\vec{\sigma}) \ln x(\vec{\sigma}) \tag{A.6}$$

with:

$$e[\{x\}] = \alpha \ln\left[\sum_{\vec{\sigma}_1,\ldots,\vec{\sigma}_K} x(\vec{\sigma}_1)\ldots x(\vec{\sigma}_K) \exp\left(-\frac{1}{T}\sum_{a=1}^{n}\prod_{\ell=1}^{K}\delta\left(\sigma_\ell^{(a)}-1\right)\right)\right] \tag{A.7}$$

and under the condition that $\sum_{\vec{\sigma}} x(\vec{\sigma}) = 1$.

We remark that the functional f is symmetric under exchange of the different replicas, and then it seems reasonable to adopt a Replica Symmetric (RS) ansatz for its optimization. We can then express all the occupation fractions in terms of a *single* local magnetization probability distribution $P(m)$ which will have to be determined self-consistently under the extremization condition of the formal free energy (A.6):

$$x(\vec{\sigma}) = \int_{-1}^{1} {\rm d}m\, P(m) \prod_{a=1}^{n} \frac{1+m\sigma^{(a)}}{2} \tag{A.8}$$

We remark that if the RS ansatz for the occupation fractions corresponded exactly to a global minimum of f, $P(m)$ would correspond to the histogram of all the local magnetizations sampled over all the optimal configurations (of course, after averaging over the quenched disorder).

At this stage, it is possible to perform the replica limit $n \to 0$, since all the functionals have been expressed coherently in term of the generic number of replicas n. This step brings to a self-consistent equation for the order parameter $P(m)$:

$$\begin{aligned} P(m) = \frac{1}{1-m^2}\int_{-\infty}^{\infty} {\rm d}u\, \cos\left[\frac{u}{2}\ln\left(\frac{1+m}{1-m}\right)\right] \\ \times \exp\left[-\alpha K + \alpha K \int_{-1}^{1}\prod_{\ell=1}^{K-1} {\rm d}m_\ell P(m_\ell)\cos\left(\frac{u}{2}\ln A_{K-1}\right)\right] \end{aligned}$$

with:

$$A_\lambda = 1 + \left(\mathrm{e}^{-\beta}-1\right)\prod -\ell = 1^\lambda\left(\frac{1+m_\ell}{2}\right) \tag{A.9}$$

Under these conditions, the corresponding RS optimal free-energy becomes:

$$\begin{aligned} -\beta f_{\rm opt}(\alpha,T) = \ln 2 + \alpha(1-K)\int_{-1}^{1}\prod -\ell = 1^K {\rm d}m_\ell P(m_\ell)\ln A_K \\ + \frac{\alpha K}{2}\int_{-1}^{1}\prod -\ell = 1^{K-1}{\rm d}m_\ell P(m_\ell)\ln A_{K-1} \\ - \frac{1}{m^2}\int_{-1}^{1}{\rm d}m P(m)\ln\left(1-m^2\right) \end{aligned}$$

The knowledge of the free-energy functional (A.10) allows the determination of relevant quantities like the ground-state entropy $S[E_{GS}]$. The limit $\beta \to \infty$ has to be taken in order to examine ground-state properties.

When considering the fraction of violated clauses at temperature $1/\beta$, the finite contributions to the ground-state energy will depend only upon the magnetizations of order at least $\pm 1 - O\left(\mathrm{e}^{-|z|\beta}\right)$ (corresponding to variables in the backbone). The saddle-point

equations can be used to determine self-consistently the probability distribution $R(z)$ of the variable $z = (1/\beta)\text{arctanh}\,(m)$:

$$R(z) = \sum_{\ell=-\infty}^{\ell=\infty} \mathrm{e}^{-\gamma} \mathrm{I}_\ell(\gamma)\delta(z-\ell) \tag{A.10}$$

where I_ℓ denote the ℓ-th modified Bessel function and γ satisfies the following implicit equation:

$$\gamma = \alpha K \left[\frac{1-\mathrm{e}^{-\gamma}\mathrm{I}_0(\gamma)}{2}\right]^{K-1} \tag{A.11}$$

The corresponding ground-state energy density can then be written as:

$$E_{GS}[\alpha] = \frac{\gamma}{2K}\left[1-\mathrm{e}^{-\gamma}\mathrm{I}_0(\gamma) - K\mathrm{e}^{-\gamma}\mathrm{I}_1(\gamma)\right] \tag{A.12}$$

The equations (A.11) and (A.12) can be solved numerically providing the RS description of the SAT/UNSAT transition.

APPENDIX B

Notions of Computability and Complexity theory

Theoretical computer science is a relatively young discipline, sprung from mathematical logic as a separated branch devoted to the study of potentialities limits of artificial computing machines. A common pedagogical approach to it (which follows also the historical evolution) starts with *computability theory*, whose main purpose is to identify problems which can or cannot be computed in a finite time (see *eg.* [**42**]). Undecidability is actually the highest form of complexity.

Next comes *complexity theory* in the proper sense (see *eg.* [**129**]). Its slightly more ambitious goal is not only to understand which computations can be performed in principle, but also which computations can be *efficiently* performed. The concepts of reduction and complexity classes are cornerstones of this field, which now comprises a major part of the active research in theoretical computer science.

1. Computability Theory

A central definition of computability theory is the one of *primitive recursive function*. This class includes several elementary functions over the integers, like constants, identities, successor function and it is closed with respect to the operations of: *composition* (if $f, g_1, \ldots, g_k$ are primitive recursive also $h(x_1, \ldots, x_m) = f(g_1(x_1, \ldots, x_m), \ldots, g_k(x_1, \ldots, x_m))$ is primitive recursive); and *primitive recursion* (given $h = h(x_0, x_1, x_2, \ldots, x_k)$ and $g = g(x_2, \ldots, x_k)$ primitive recursive, the uniquely defined f such that $f(0, x_2, \ldots, x_k) = g(x_2, \ldots, x_k)$ and $f(y+1, x_2, \ldots, x_k) = h(y, f(y, x_s, \ldots, x_k), x_2, \ldots, x_k)$ is also primitive recursive. The basic idea is that primitive recursive function are computable in finite time because they can be built out in finite time from elementary computable functions.

Nevertheless, it is easy to find computable functions which are not primitive recursive, like the *Ackermann function*, a sort of "hyper-power" operation:

$$A(0, x, y) = x + 1$$

$$A(1, x, y) = x + y$$

$$A(2, x, y) = x \cdot y \tag{B.1}$$

$$A(n+1, x, y) = \begin{cases} 1 & y = 0 \\ A(n, x, A(n+1, x, y-1)) & y \neq 0 \end{cases} \quad \text{for } n \geq 2$$

Its growth is incredibly fast (*eg.* $A(1,2,2) = 4$, $A(2,4,4) = 4 \cdot 4 = 16$, $A(3,4,4) = 4^4 = 256$, $A(4,4,3) = 4^{4^4} \simeq 1.34 \times 10^3 4$, $A(4,4,4) = 4^{4^{4^4}}$, etc.) and it is possible to prove that it cannot tracked by composition of successor operations. It is nevertheless computable in finite time, and there is then need for a more extended definition of recursive (computable) function.

1.1. Turing machines. A Turing Machine (TM) is composed by a control unit and (at least) one tape. It can be formally described by a quadruple (K, Σ, δ, s), where K is a finite set of states, Σ is a finite set of symbols, $s \in K$ is the initial state. Σ always contains the special symbols $\sqcup$ (blank space)and $\triangleright$ first symbol). Finally δ is a transition function which maps $K \times \Sigma$ to $(K \cup \{h, \text{“}yes\text{”}, \text{“}no\text{”}\}) \times \Sigma \times \{\rightarrow, \leftarrow, ---\}$. Strictly speaking h (the “halting” state), *“yes”* (the accepting state), *“no”* (the rejecting state and the cursor directions $\leftarrow$ for “left sliding” of the tape, $\rightarrow$ for “right sliding” and — for “stay”, do not belong to $K \cup \Sigma$, so some authors introduce in the definition of the TM an extended alphabet $T \supset \Sigma$ including all of them.

We will never use this definition directly, but we cited it to give the flavour of the language needed to give sound foundations to Computability theory. The control unit of the TM just scroll its tape until when it finds the starting character. At this point it starts reading toward the right the characters written on the tape until the first occurrence of the blank character. This marks the end of input. At this point, the transition function δ is used to determine the new state of control unit and to translate the input sequence into actions, that can be chosen in a very limited repertoire, including $\rightarrow$, $\leftarrow$, —, write/read character on current location, and very few else. Every TM executes then the peculiar *program* encoded in its transition function. The program execution continues until when one of the output special states h, *“yes”*, or *“no”* is reached.

TMs offer a way to define *recursive functions*: we say that $f : \Sigma \rightarrow \Sigma$ [1] is computable if there exists a TM such that reading an input $z \in \Sigma$ stops after a finite number of steps writing on the tape a symbol w such that $f(z) = w$ if f is defined in z, or never stops if f is undefined in z. It is easy to show that functions like the Ackermann function can be computed by a TM.

TMs may seem quite weak computational machines, especially when comparing them with the powerful personal computer surrounding us in everyday life. But TMs are a *universal* computation model. It is possible indeed to prove the existence of a TM U that receiving as input the description in some alphabet of a TM M operating over the input x (description denoted by the syntax $|M; x|$) produces as output exactly the computation of M over x:

$$\exists U, \forall M, \forall x / \, (|M; x|) = M(x) \tag{B.2}$$

It is then possible to build a TM that can simulate the computation operated by any other TM. This powerful property among others pushed Alonzo Church and others to do the following claim, generally accepted as a sort of defining axiom (and making the definition of computability in terms of TMs quite reasonable):

> **Church-Turing thesis**: A function computable in any reasonable model of computation can also be computed by a TM.

Particularly important are TMs able to *decide a language*: given a TM M with alphabet Σ_M and a *language* $L_M \subset \Sigma_M$, we say that M decides the language L_M if and only if the computation stops in *“yes”* state when the input $z \in L_M$, and stops in *“no”* state if the input $z \notin L_M$. A language deciding TM computes actually the characteristic function of a given subset of its input alphabet.

[1] Rigorously speaking, the input and output alphabet can be different and the output alphabet is included in the extended alphabet T appearing in the more complete definitions of TM.

1.2. Undecidability. Undecidability is in some sense the most lethal form of complexity. There is a trivial reason for the existence on undecidable languages, there are uncountably many more languages than TMs able to decide them. We have indeed already mentioned the fact that to every possible computation of the Universal Turing Machine one can associate a precise integer called its Gödel number. There is then only a countable infinity of recursive functions, but the number of languages in $\mathbb{N}$ is uncountable, corresponding to the cardinality of $2^{\mathbb{N}}$, the set of subsets of $\mathbb{N}$.

What appears striking is the fact that there are undecidable problems that at first glance would seem quite innocuous. It is for instance impossible to build a machine deciding the language H of the descriptions $|M; x|$ of the computations $M(x)$ that do not diverge. This means that the task of deciding if a given Turing Machine will generate a diverging computation starting from a given input is actually non computable in finite time (*halting problem*)

2. Complexity Theory

First of all, it is useful to remark that every *decision problem* can be put in correspondence with a specific language. It is indeed enough to find a reasonably succinct coding in some alphabet Σ for the instances of a given problem. A TM able to decide the language Σ_Π of the positive instances of a decision problem Π, will be of course able to decide if a given instance of Π is positive or not.

We can define a *complexity class* as a set of languages that can be decided by a given TM, using for every possible input x no more than $f(|x|)$ units of a given resource. Here $|x|$ is the length of the input and for "resources" we intend typically time steps required for the computation to halt, or space units occupied by symbols on the tape. We can then naturally define the complexity classes **TIME**$(f(n))$ and **SPACE**$(f(n))$. These classes (unless we use extremely pathological functions f) tend to form *hierarchies*. It can for instance proved rigorously that:

> **Time hierarchy theorem:** if $f(n) > n$ then **TIME**$(f(n)) \subset$ **TIME**$(f^3(n))$
> **Space hierarchy theorem:** if $f(n) > \log n$ then **SPACE**$(f(n)) \subset$ **SPACE**$(f(n) \log f(n))$

Another important theorem is the following:

> **Time Speed-up theorem:** if L is a language decidable in time $f(n) > n$ by a TM M, there exist for every $\epsilon > 0$ another TM M' that decides the same language L in a time $\epsilon f(n) + n + 2$

An analogous result holds for space. The proof relies on the definition of a new alphabet whose symbols encode for more than one symbol of the original alphabet. Speed-up theorems affirm than that optimization of the computation is always possible, but that the asymptotic behavior for large n cannot be modified, making then meaningful the definition of the classes **TIME**$(f(n))$.

We remark in the end that deciding if $x \in L$ can be more or less difficult depending on the element x for any assigned L. The definition of complexity classes refers to the maximum possible cost $f(|x|)$. What we are building is then a theory of *worst-case complexity*. Theories of typical-case complexity can only difficulty be founded on rigorous grounds, and only recently some progresses in this direction have been done thanks to the use of techniques borrowed from Statistical Mechanics (see Chap. 1 and Chap. 2).

2.1. Non-deterministic Turing Machines. Factoring an integer n, product of two prime factors, can be an extraordinarily difficult task if the factors are large. Nevertheless it is extremely easy to verify if a given prime number is or is not a factor of the input integer. Trying to divide n by a prime taken at random could be considered as a non deterministic algorithm for factoring. If it was possible to divide n by all the primes and to look *simultaneously* at all the outcomes, factoring would be easy.

A non-deterministic Turing Machine is a Turing Machine endowed with an additional component called hypothesis module. At the beginning of the computation the hypothesis module write on the tape an "hint string" that determines the further evolution of the computation together with the input. Different hints bring to different computations. We say that a non-deterministic Turing machine decides a language L if at least one computation path bring to a *"yes"* state for every $x \in L$.

The time and space costs of the non-deterministic computation are the smallest possible among the different possible computation paths. It is possible to define for non-deterministic TMs the classes **NTIME**$(f(n))$ and **NSPACE**$(f(n))$, in exactly the same way used for deterministic TMs.

2.2. The polynomial classes. At this point the definition of the well-known classes **P** and **NP** becomes trivial:

$$\mathbf{P} = \bigcup_{k=1}^{\infty} \mathbf{TIME}\left(n^k\right), \qquad \mathbf{NP} = \bigcup_{k=1}^{\infty} \mathbf{NTIME}\left(n^k\right)$$

Please remark that **NP** does not stand for "not polynomial". The most important open problem in complexity theory is actually deciding if **P** = **NP** or not. Many alternative proofs or counter-proofs of this equivalence can be defined. **coNP** is the class of the problems complementary to an **NP** problem. For instance, in the Traveling Salesman Problem, deciding if there is a path shortest than the length λ is a problem in **NP** (a certificate solving the decision problem is the length of a single path shorter than λ). An instance of the complementary problem consists in deciding if all the possible TSP paths are longer than another threshold-length λ', and appears to be much more complicated (no certificates more succinct than the list of all the path lengths are known up to date). While the *Pratt Theorem* states that the primality problem belongs to **NP** $\cap$ **coNP**, in the case of TSP there is an evident asymmetry between the complexity of the direct and complementary problem. What we know is that:

$$\mathbf{NP} \neq \mathbf{coNP} \Rightarrow \mathbf{P} \neq \mathbf{NP}$$

while the reverse is not necessarily true.

Considering now the space complexity classes, one can define:

$$\mathbf{PSPACE} = \bigcup_{k=1}^{\infty} \mathbf{SPACE}\left(n^k\right), \qquad \mathbf{NPSPACE} = \bigcup_{k=1}^{\infty} \mathbf{NSPACE}\left(n^k\right)$$

in addition to $\mathbf{L} = \mathbf{SPACE}(\log n)$ and $\mathbf{NL} = \mathbf{NSPACE}(\log n)$. A crucial property of the space resource is the possibility of being reused when some information is no more needed. Something more can then be proved: first of all, $\mathbf{PSPACE} \setminus \mathbf{L} \neq \emptyset$. And then:

Savitch's theorem[2]: $\mathbf{PSPACE} = \mathbf{NPSPACE}$

So this is the final tower of class inclusions:

$$\mathbf{L} \subseteq \mathbf{NL} \subseteq \mathbf{P} \subseteq \mathbf{NP} \subseteq \mathbf{PSPACE}$$

[2]To be more precise, just an important corollary to the Savitch's theorem!

2.3. Polynomial-time reductions and complete problems. Reductions play a crucial role in both computability and complexity theory. A language A is said to be Karp-reducible in polynomial time to another language B if a there is a function computable in polynomial-time by some TM such that $x \in A$ if and only if $f(x) \in B$.

A language L in a complexity class **C** is said to be **C**-complete if and only if every other language in **C** is Karp-reducible in polynomial time to L. Although the definition of complete problems is simple, it is not obvious that they really exist. A first positive answer comes from:

Cook's theorem: 3-SAT problem is **NP**-complete

Many other problems are proved to be **NP**-complete reducing 3-SAT to them and showing then that they have at least the same complexity. It can be shown that 3-SAT reduce to the search of an independent subset of vertices in a graph (IND), to the search of a clique (CLIQUE) and of an hamiltonian path (HAMILT). HAMILT reduce on the other hand to the decision version of TSP. A particular version of 3-SAT reduces to the coloring problem (COL), and so on. For a comprehensive list of **NP**-complete problems, the reader is asked to refer to [**58**].

Besides the undecidable languages, the most complex solvable problems belong to **PSPACE**-complete. Tasks like generalized version of the Go game, or the resolution of 3-SAT formulas with quantified variables ($\forall$, $\exists$) belong to this class.

2.4. Other complexity classes. Given an instance I of a problem Π the set $S_\Pi(I)$ of the solutions of the instance I and some specific subset $S \subset S_\Pi(I)$: the *decision versions* of Π consist in deciding if S is empty or not; the *optimization versions* consist in finding an element of S; the *enumeration versions* consist in counting the element of S.

The complexity of optimization and enumeration versions is at least equivalent to the one of the corresponding decision problem, since finding or counting elements shows also their existence. Optimization versions of **NP**-complete decision problems are often called **NP**-hard problems. Suppose to have a TM with an *oracle* for the decision version (that is a machine which is able to solve the decision problem without any effort; if consulting this machine, the optimization version becomes solvable in polynomial time, than the optimization problem is said to be **NP**-easy. If a problem is both **NP**-hard and **NP**-easy it is said to be **NP**-equivalent (this is the case of the optimization version of TSP).

For enumeration problems one can define the class $\sharp$**P** of the problems solvable by a non-deterministic TM such that all the computations ending with acceptance are polynomially bounded in time and are exactly as many as the solutions included in the set to enumerate. Some problems in $\sharp$**P** are so complicated that they would be eligible for intractability even if the identity **P** = **NP** was proved. Valiant's theorem states that *permanent* computation is a $\sharp$**P**-complete problem (completeness for enumeration problems require the use of *parsimonious reductions* conserving the number of solutions).

Other complexity classes can be defined when the computational model is modified, notably for *probabilistic* Turing machines (similar to non-deterministic TM in every respect but for the fact that different computation path are selected stochastically instead of simultaneously). If the probability of solving a problem becomes larger than $\frac{1}{2}$ after a polynomial number of steps, this problem is in **PP**. If the probability becomes larger than $\frac{1}{2} + e$ with $0 < e < \frac{1}{2}$ in polynomial time, the problem is in **BPP**. Another interesting class is **ZPP**, for which it is impossible to find a wrong answer but it is possible to end the computation in a "don't know" state.

A complexity theory for probabilistic TMs has been developed also for *quantum* Turing Machines, where the control unit states belong to a Hilbert Space: the classes **QP**, **BQP** and **ZQP** have been introduced in literature. Factoring belongs to **NP** but also to **BQP**.

APPENDIX C

Notions of Information theory

A concise description of what Information theory is about can be found in a sentence of its founder, Claude Shannon:

> *The fundamental problem of [Information theory] is that of reproducing a one point either exactly or approximately a message selected at another point.*

In the following appendix we will just give a concise list of definitions and statements, useful for a better understanding of the topics dealt with in Chapter 4. For a more complete introduction, the reader is asked to refer to other textbooks, like the excellent one by David MacKay [**89**], main inspiration for the following paragraphs.

1. Basic concepts and notions

1.1. Measures of information content. The Information content (or Shannon entropy) of an *outcome* x is defined as:

$$h(x) = -\log_2 P(x) \tag{C.1}$$

where $P(x)$ is the probability of finding x as outcome of a sampling for the random variable X. Information content is measured in *bits*. This definition can be extended to an *ensemble*, evaluating the average information content of a trial outcome ($\mathcal{A}_X$ is the set of possible outcomes of X):

$$H(X) = -\sum_{x\in\mathcal{A}_X} P(x)\log_2 P(x) \tag{C.2}$$

An example can be given in order to convince that Shannon entropy is a meaningful definition of information content. Suppose to play a simplified version of the battleship game over an 8×8 checkerboard. There are two players and each one put on the checkerboard just one single boat, occupying a single square. Each shot has two possible outcomes, either "miss" or "hit". Suppose that the first player hits the opponent's ship at the very first shot. According to the definition (C.1), the information gained after this outcome is $\log_2 64 = 6$ bits. Remark that exactly 6 bits are needed in order to identify the coordinate of the ship on the square 8×8 checkerboard. Suppose on the other end to initially miss the opponent's ship, hitting it only when there are n left positions. The information gained hitting the target at this point is: $-\log_2\frac{63}{64} - \log_2\frac{62}{63} - \ldots - \log_2\frac{n}{n+1} - \log_2\frac{1}{n} = \log_2 64 = 6$ bits independently of n.

Given two different ensembles $\mathcal{A}_X$ and $\mathcal{A}_Y$ of outcomes (statistically independent or not), the joint information content is analogously given by:

$$H(X,Y) = -\sum_{x\in\mathcal{A}_X, y\in\mathcal{A}_Y} P(x,y)\log_2 P(x,y) \tag{C.3}$$

The *joint* information about X and Y is $H(X,Y) = H(X) + H(Y)$ if X and Y are independent random variables.

The *conditional* information of a random variable X given $Y = b_y$ is:

$$H(X|y = b_k) = -\sum_{x \in \mathcal{A}_X} P(x|y = b_k) \log_2 P(x|y = b_k) \tag{C.4}$$

Once again this can be averaged over the ensemble, giving the conditional information content of X given Y:

$$H(X|Y) = -\sum_{x \in \mathcal{A}_X, y \in \mathcal{A}_Y} P(x,y) \log_2 P(x|y) \tag{C.5}$$

This measures how much uncertainty remains about x on average when y is known.

We can now define the *mutual* information between X and Y as:

$$I(X;Y) = I(Y;X) = H(X) - H(X|Y) \tag{C.6}$$

For proving the simmetry, note that $H(X,Y) = H(X) + H(Y|X) = H(Y) + H(X|Y)$, that is the uncertainty about x and y is the uncertainty about x plus the uncertainty about y given x or vice versa. The mutual information between X and Y measures then the average reduction in uncertainty about x that results from learning the value of y. It can also be intended as information that x conveys about y, or, equivalently, y about x.

1.1.1. *Data-processing theorem.* The data-processing theorem states the quite obvious but somehow disappointing fact that data-processing can only destroy information. The proof can be obtained considering the three ensembles W, D and R, where w is the state of the world, d is the data gathered and r is the processed data. The three random variables W, D and R form a Markov chain, such that:

$$P(w,d,r) = P(w)P(d|w)P(r|d) \tag{C.7}$$

At this point is straightforward to show that $I(W,R) \leq I(W,D)$, that is the information about the world conveyed by the processed data is smaller or equal than the one conveyed by the gathered data.

1.2. Codes. A *source* produce a signal $\mathbf{s}$. The signal is then *encoded* in a new message $\mathbf{t}$ which is then transmitted through a typically noisy *channel*. The output of the channel is the received message $\mathbf{r}$ which is then *decoded* into a message $\widehat{\mathbf{s}}$. The coding-decoding stages have been succesful if $\widehat{\mathbf{s}}$ conveys a sufficient amount of information about $\mathbf{s}$.

Error-correction codes can be used for a systemic approach against transmission errors induced by the channel. Using a binary alphabet a possible model noisy communication channel could be the *binary symmetric channel* (BSC), introducing bit-flips with probability f, such that:

$$P(\widehat{s} = 0,1|s = 0,1) = 1 - f, \;\; P(\widehat{s} = 1,0|s = 0,1) = f$$

Another noisy channel model could be the *binary erasure channel* (BEC), where the message bit is either exactly or confusedly transmitted:

$$P(\widehat{s} = 0,1|s = 0,1) = 1 - f, \;\; P(\widehat{s} = ?|s = 0,1) = f, \;\; P(\widehat{s} = 1,0|s = 0,1) = 0$$

An error-correcting code operates increasing *redundancy*, in order to compensate the loss of information occurring during transmission. The transmission *rate* of this codes is then smaller than 1 (less than one bit of information about the source is transmitted during each cycle).

Another possible purpose of source coding is to compress information in such a way that not too much information is lost after the encoding and decoding stages. Of course, lossless compression is possible only if the original code used by the source is already redundant.

1.2.1. *Repetition codes.* A trivial error-correction code (obviously increasing the coding redundance) consists in repeating an every bit of the source message an odd prearranged number of times. Decoding is then performed by a majority vote decoder. It is easy to show that, for the transmission of a single source bit, encoded with N repetitions, through a BSC with bit-flip probability f, the probability of wrong decoding is:

$$p_b = \sum_{n=(N+1)/2}^{N} \binom{N}{n} f^n (1-f)^{N-n} \tag{C.8}$$

Ordinary one gigabyte disk drives can be modeled by BSC with $f \simeq 0.1$. In order to achieve the industrial standard of reliability $p_b \simeq 10^{-15}$, one would need 60 one gigabyte drives instead of a single one. Obviously the redundancy introduced by repetition codes is too large...

1.2.2. *Linear codes.* A better strategy consists in adding redundancy to block of data instead of encoding one bit at a time. In a linear block code, each sequence of K bits is converted in a sequence of $N > K$ bits before transmission. The $N - K$ extra bits are called *parity-check* bits. The transmitted codeword $\mathbf{t}$ can be written form the source sequence $\mathbf{s}$ with the operation:

$$\mathbf{t} = \mathbb{G}\mathbf{s} \tag{C.9}$$

where $\mathbb{G}$ is a $N \times K$ matrix with suitable properties (well-known examples are the so-called Hamming Codes). All the encoding operations are performed with modulo-2 arithmetics.

Efficient decoding schemes have then to be devised. A common strategy is *syndrome decoding* (used for instance for the optimal decoding of Hamming codes), in which the inspection of the pattern of the alterations of the received block allows to decide with large accuracy which are the most likely corrections to apply in order to retrieve the original codeword.

Low-density parity check codes (LDPC) are a special family of linear codes representing the state of the art for communication over realistic channels. They involve just the use of *sparse* binary matrices and they can be practically and almost optimally decoded by message-passing algorithms, similar to BP described in Chap. 3. Moreover, they recently attracted the attention of physicists, because of their intimate relationship with the diluted p-spin model [**23, 108**].

2. Some fundamental theorems by Claude Shannon

2.1. Source coding theorem. The source coding theorem describes the properties of codes without redundancy. A possible verbal statement is the following: N independent identically distributed random variables each one with Shannon entropy $H(X)$ can be compressed into more than $NH(X)$ bits with negligible risk of information loss, as $N \to \infty$; conversely, if they are compressed into fewer than $NH(X)$ bits it is virtually certain that information will be lost.

In a sense the Source coding theorem gives the ultimate justification of the use of Shannon entropy as measure of information content. The shortest code for the position of the opponent's ship is given by the 6 bits required for giving its coordinates. A sequence

of n shots requires at least $6n$ bits in order to be described, but constitutes in the end just a redundant code for the 6-bit-encodable opponent's ship position.

An important part of Information Theory is devoted to the search for optimal loss-less coding strategies. Huffman codes, arithmetic codes, Lempel-Ziv-like methods are some of the most standard approaches to the problem.

2.2. Noisy-channel coding theorem. What performance can the best error-correcting codes achieve? We have seen that a reduction in the decoding error probability p_b can typically be obtained by increasing the transmission rate R. Shannon was able to prove that associated to any given channel there is a number C such that communication at an *arbitrarily small error probability* becomes theoretically possible if $R < C$. The number C is called the *capacity* of the channel, and it is given by:

$$C = \max_{\mathcal{A}_X} I(X, Y) \tag{C.10}$$

where X describes the input to the channel, Y its output and the maximum is taken over the possible input distributions. Optimal error-free communication (communication at capacity) is then obtained by maximizing the information conveyed by the channel output about the input.

The capacity for the BSC is given by:

$$C_{\mathrm{BSC}}(f) = 1 + f \log_2 f + (1 - f) \log_2(1 - f) \tag{C.11}$$

Coming back to the problem of the error correction for a one gigabyte disk drive, it is theoretically possible (and also practically in this case) to achieve the industrial standard $p_b = 10^{-15}$ using just two disks, against the 60 required by the repetition code...

2.3. Rate-distortion theory. It is possible to transform a noisy-channel into a noiseless one just by increasing the transmission rate. But what is now the maximum rate that we can achieve over a noiseless channel if we accept given amount of decoding error D?

Take a noiseless binary channel with a capacity of 1 bit, and try to transmit at a rate $R = 2$ (two bit per cycle). Half of the information will be thrown away. A trivial decoding is *random guessing*: the receiver just guesses the missing fraction $1 - 1/R$ at random, with an average probability of bit error given by:

$$p_b = \frac{1}{2}\left(1 - \frac{1}{R}\right) \tag{C.12}$$

It is in reality possible to do much better. Consider an optimal block code C with rate $R' = K/N$ able to achieve error error-free communication over a BSC with $f = D$ (K ¡ N, where N is the length of the codewords and K the source block-length). Let use it in the reverse way. We chop the source stream in blocks of length N and we compress them using the *decoder* of C. The results are compressed chunks of length K that can be expanded again to the original length by using the *encoder* of C. The expanded block will typically differ from the original one in DN bits. The capacity of the BSC with $f = D$ is given by eq. (C.11) and corresponds to the best possible rate for the error-correcting code C. The associated lossy-compressor with distortion D will have then a rate equal to $1/C_{\mathrm{BSC}}(D)$.

Attaching this lossy compressor to the noiseless channel with capacity C, we have then proved the achievability of communication with distortion D up to a rate of:

$$R = \frac{C}{1 + D \log_2 D + (1 - D) \log_2(1 - D)} \tag{C.13}$$

The results (C.12) and (C.13) can easily be generalized to the case of biased sources.

APPENDIX D

SP as a message-passing algorithm

The various algorithms based on the idea of Survey Propagation belong to the general class of *message-passing* algorithm, object of intense research in theoretical computer science and information theory during the last ten years. In particular, the SP equations (4.5), taken after the limit $y \to \infty$, share important similarities with one of the most important message-passing techniques, the so called *Belief Propagation* (BP) algorithm.

In this appendix we will introduce the BP equations for a generic combinatorial constraint satisfaction problem, and we will rewrite the SP equations (4.5) in a new generalized formalism, in which similarities and differences with BP will become evident. This comparison will introduce then to the statement of a general equivalence theorem, clarifying the role of unfrozen variables and opening the road to the realization of SP versions equivalent to ansatz with further step of RSB.

1. Belief Propagation for general CSPs

Given a CSP where the variables can assume q states (Potts spins) and a configuration $\vec{x}$, we can define three quantities associated with them:

Warning: : For any given edge $a \to i$ of the factor graph, we define as *warning* the q-component vector $\vec{u}_{a\to i}(\vec{x}) \in \{0,1\}^q$ with components:

$$\hat{u}^p_{a\to i}(\vec{x}) = C_a\left[(x_j)_{j\in V(a)\setminus i} |\, x_i \leftarrow p\right] \;,\; p = 1, ..., q,$$

given then by the value of constraint C_a in the configuration obtained from $\vec{x}$, by substituting p to the original value of x_i.

Cavity field: : For any given edge $a\not{o}i$ of the factor graph, we define as *cavity field* the q-component vector $\vec{h}_{i\to a}(\vec{x}) \in \{0,1\}^q$ with components:

$$h^p_{i\to a}(\vec{x}) = \max_{b\in V(i)\setminus a} u^p_{b\to i}(\vec{x})$$

Local field: : For any given node i, we define as **local field** the q-dimensional vector $\vec{h}_i(\vec{x}) \in \{0,1\}^q$ with components:

$$h^p_i(\vec{x}) = \max_{a\in V(i)} u^p_{a\to i}(\vec{x}) \;.$$

A warning component $\vec{u}^q_{a\to i}(\vec{x}) = 1$ have to be intended as a message sent from constraint a to variable i saying that x_i cannot be in the states q without violating constraint a. The local field on variable i summarizes on the other end all the warnings received by i. The only admissible states for i are then the one for which the components of its local field are null. The cavity fields are local fields on a cavity graph.

The messages above are defined in terms of a given configuration $\vec{x}$, which has not necessarily to be a solution. Interesting objects to evaluate are the histograms of the local

and the cavity fields, above all the solutions $\vec{s} \in S_C$ of the CSP. For instance:

$$H_i(\vec{h}) = \frac{1}{|S_C|} \sum_{\vec{s} \in S_C} \delta_{\vec{h},\vec{h}_i(\vec{s})} \,, \tag{D.1}$$

Local-field histograms constitute a statistical description of the set of solutions, which can be exploited algorithmically in order to recursively search an element of S_C (*eg.*. by decimation techniques). All the messages will have always *at least one component different from zero*, forbidding some state to each variable. In particular, if all but one components of a message will be non-zero, the variable will be *frozen* to one specific value in all solutions.

The histograms $H_i(\vec{h})$ can then be computed approximately using a message passing procedure. The histogram of $\vec{u}_{a\to i}(\vec{s})$ depends on the joint histogram of all the warnings $\vec{u}_{b\to j}(\vec{s})$ for every $j \in V(a) \setminus i$ and for every $b \in V(j) \setminus a$. Now the cavity ansatz enters into play. For every edge $a \to i$ of the factor graph, we define the *belief* $B_{a\to i}(\vec{u})$ as the histogram of the warning $\vec{u}_{a\to i}$ over the configurations $s \in S^{(i)}$ which are solutions of a problem with the same factor graph, but a cavity in i:

$$B_{a\to i}(\vec{u}) = \frac{1}{|S^{(i)}|} \sum_{\vec{s} \in S^{(i)}} \delta_{\vec{u},\vec{u}_{a\to i}(\vec{s})} \tag{D.2}$$

Let us first consider a set of incoming warnings $\{\vec{u}_{b\to j}\}$ for every $j \in V(a) \setminus i$ for every $b \in V(j) \setminus a$, and the associated set of cavity fields $\{\vec{h}_{j\to a}\}$. We denote by $\mathcal{T}(\{\vec{h}_{j\to a}\})$ the set of all the configurations of the s_j variables allowed by the incoming set of warnings:

$$T(\{\vec{h}_{j\to a}\}) = \left\{(s_j)_{j\in V(a)\setminus i} \mid h^{s_j}_{j\to a} = 0,\ \forall j \in I(a) \setminus i\right\} \tag{D.3}$$

For each (s_j) in $T(\{\vec{h}_{j\to a}\})$, it becomes then possible to evaluate the output warning $\vec{u}_{a\to i}$. Assuming independence of the incoming warnings, we can approximately write:

$$H_{j\to a}(\vec{h}) = \sum_{\{\vec{u}_{b\to j}\}_{b\in V(j)\setminus a}} \delta_{\vec{h},\vec{h}_{j\to a}} \prod_{b\in V(j)\setminus a} B_{b\to j}(\vec{u}_{b\to j}) \,. \tag{D.4}$$

and then:

$$B_{a\to i}(\vec{u}) = \frac{1}{\mathcal{Z}} \sum_{\{\vec{h}_{j\to a}\}_{j\in V(a)\setminus i}} \left[\sum_{\vec{s}\in\mathcal{T}(\{\vec{h}_{j\to a}\})} \delta_{\vec{u},\vec{u}_{a\to i}(\vec{s})} \right] \prod_{j\in V(a)\setminus i} H_{j\to a}(\vec{h}_{j\to a}) \tag{D.5}$$

where the prefactor $\mathcal{Z}^{-1}$ is a normalization constant.

These equations can tentatively be solved by iteration, starting form some randomly chosen population of beliefs. In some cases a fixed-point will be reached and the histograms of local fields (evaluated with a relation analogue to (D.4) including the contributions from all the incident edges) will be helpful for performing decimation.

2. Survey Propagation for general CSPs

A crucial difference between BP and SP is the fact that *unfrozen* variables are taken explicitly into account. If a given configuration of messages allows a variable to assume several possible states (this is the case in which the cardinality of $\mathcal{T}(\{\vec{h}_{j\to a}\})$ is larger than 1), in the SP formalism, we assign to it a special state called the *joker* state. A byproduct of the definition of this extra state is the collapse of several "BP-style" configurations into a single "SP-style" configuration. To be more rigorous, warnings, cavity and local fields in the context of the SP formalism will be functions of sets of allowed configurations, more than of allowed configurations themselves: For a given CSP, we define:

Generalized warning (or cavity bias): : For any given edge $a\text{ø}i$ of the factor graph, let S be a given set of possible values for the variables $(x_j)_{j\in V(a)\setminus i}$. We define as **warning** the q-component vector $\vec{\hat{u}}_{a\to i}(S) \in \{0,1\}^q$ with components:

$$\hat{u}^p_{a\to i}(S) = \min_{(x_j)\in S} C_a\left[(x_j)_{j\in V(a)\setminus i} |\, x_i \leftarrow p\right] \ , \ p = 1, ..., q,$$

Note that now the *null message* (0,0,0, ..., 0) enters the set of possible warnings. It is always sent when a variable is free to assume any state, that is when it is forced to assume a joker state. The definitions of cavity and local fields are exactly identical to the BP case, with generalized warnings depending on sets of configurations S instead of ordinary warnings depending on configurations.

The histograms of warnings and fields are now defined over *clusters of solutions* S_C instead of individual solutions $\vec{s}$. In the generalized formalism presented in this section, it is straightforward to see that solutions differing only in the assignments given to unfrozen variables generate exactly the same set of generalized warnings. The histogram of local fields is given by:

$$H_i(\vec{h}) = \frac{1}{\mathcal{N}_{cl}} \sum_{\ell=1}^{\mathcal{N}_{cl}} \delta_{\vec{h},\vec{h}_i(S_C^\ell)} \ . \tag{D.6}$$

We define on the other end the *surveys*, in term of clusters of solutions of the cavity graphs:

$$Q_{a\to i}(\vec{u}) = \frac{1}{\mathcal{N}^{(i)}_{cl}} \sum_{\ell=1}^{n^{(i)}_{cl}} \delta_{\vec{u},\vec{u}_{a\to i}(S_C^{\ell,(i)})} \ . \tag{D.7}$$

Now a fundamental difference enter into play. It is still possible to consider incoming warnings as independent but, since multiple allowed configurations collapse into a single extended-space configuration, contradictory messages have now to be *explicitly filtered*[1]. We have:

$$H_{j\to a}(\vec{h}) = \sum_{\{\vec{u}_{b\to j}\}_{b\in V(j)\setminus a}} \delta_{\vec{h},\vec{h}_{j\to a}} \prod_{b\in V(j)\setminus a} Q_{b\to j}(\vec{u}_{b\to j}) \ . \tag{D.8}$$

We introduce then the set of all non-contradictory cavity field configurations:

$$\mathcal{M}_{a\to i} = \left\{ \{\vec{h}_{j\to a}\}_{j\in V(a)\setminus i} \mid \forall j : \ \vec{h}_{j\to a} \in \{0,1\}^q, \ \vec{h}_{j\to a} \neq (1,\dots,1) \right\} \ . \tag{D.9}$$

Then, for an element of $\mathcal{M}_{a\to i}$ we define again:

$$\mathcal{T}\left(\{\vec{h}_{j\to a}\}\right) = \left\{ (s_j)_{j\in V(a)\setminus i} \, | h^{s_j}_{j\to a} = 0, \forall j \in V(a)\setminus i \right\} \tag{D.10}$$

All the elements of $\mathcal{T}(\{\vec{h}_{j\to a}\})$ give rise to a single outgoing warning, then, finally:

$$Q_{a\to i}(\vec{u}) = \frac{1}{\tilde{\mathcal{Z}}} \sum_{\{\vec{h}_{j\to a}\}\in\mathcal{M}_{a\to i}} \delta_{\vec{u},\vec{u}_{a\to i}\left(\mathcal{T}(\{\vec{h}_{j\to a}\})\right)} \prod_{j\in V(a)\setminus i} H_{j\to a}(\vec{h}_{j\to a}) \tag{D.11}$$

In BP, each configuration of input messages contributes with the full set of the possible outputs, introducing a bifurcation mechanism which may easily become unstable. On the contrary, in SP the clustering ansatz is done from the very beginning and a variable which receives a message having at least two zero components will be frozen in the unfrozen state!

[1]This is exactly the action of the reweighing term (2.34) in the limit $y \to \infty$.

3. Extended spaces and SP/BP equivalence

At least in the case of the K-SAT problem, rigorous results about the correspondence between BP and SP can be proved. It has indeed be shown in [**29**] that the marginals computed with SP equations over a given K-SAT instance coincides exactly with the marginals extracted by BP over a modified "local equilibrium" hamiltonian, defined in an extended configuration space, including besides the states ± 1 also the joker state, and over a suitably built dual graph.

The proof is simple but rather technical and the reader is asked to refer to the original paper for more details. An interesting fact is that the entropy computed by the "dual" BP coincides with the complexity estimated by SP. The link between the onset of clustering and the appearance of variables forced to be unfrozen can then be put on a rigorous ground, at least for the K-SAT model. It is indeed possible to build an exact mapping between a 1-RSB problem and a RS problem over an extended space and with a modified hamiltonian.

Another important feature of the dual construction is the possibility of iterating it more times. If SP coincides with BP over a dual problem, one could think to build SP over the dual problem itself. This would coincide with BP over a second dual problem, but also to a sort of 2-RSB version of SP. The second dual construction can at least approximately be completed, and a simplified version of its equations as been also implemented algorithmically. This is the SP-2 algorithm used tentatively in Chapter 5 for establishing the onset of 1-RSB instability: a difference between the marginals computed in the 1-RSB and the 2-RSB ansatz shows that the 1-RSB framework has become unstable with respect to the immersion in a larger functional space. Work is in progress for obtaining a rigorous and complete characterization of the behavior and meaning of the SP-2 equations ([**10**]).

Bibliography

[1] E. Aarts, J. Korst, *Simulated Annealing and Boltzmann Machines*, Wiley & Sons, Chichester (1989)

[2] E. Aarts, J.K. Lenstra, *Local Search in Combinatorial Optimization*, John Wiley and sons (1997)

[3] D. Achlioptas, C. Moore, Random k-SAT: Two Moments Suffice to Cross a Sharp Threshold, preprint `cond-mat/0310227` (2003)

[4] D. Achlioptas, U. Noar, Y. Peres, On the Maximum Satisfiability of Random Formulas, in *Proceedings of 44th IEEE Symposium on Foundations of Computer Science*, IEEE Computer Soc. Press, 362–370, Los Alamos (2003)

[5] D. Achlioptas, G. Sorkin, Optimal myopic algorithms for random 3-SAT, in *Proceedings of 41st IEEE Symposium on Foundations of Computer Science*, IEEE Computer Soc. Press, 590–600, Los Alamos (2000)

[6] N. Ajtai, Generating hard instances of lattice problems, Electronic Colloquium on Computational Complexity (ECCC) 3, 7 (1996)

[7] D. Amit, *Modeling Brain Function*, Cambridge University Press, Cambridge, UK (1989)

[8] D. Amit, N. Brunel, Model of global spontaneous activity and local structured delay activity during delay periods in the cerebral cortex, Cerebral Cortex 7, 237–252 (1997)

[9] D. Amit, S. Fusi, Learning in neural networks with material synapses, Neural Computation 6, 957 (1994)

[10] C. Baldassi, R. Zecchina, private communication

[11] W. Barthel, A. K. Hartmann, M. Weigt, Solving satisfiability problems by fluctuations: An approximate description of the dynamics of stochastic local search algorithms, lecture at Schloss Dagstuhl International Conference and Research center for Computer Science; download site: `cond-mat/0301271` (2003)

[12] D.A. Battaglia, A. Braunstein, J. Chavas, R. Zecchina, Exact probing of glassy states by Survey Propagation, Progr. Theor. Phys. Suppl. 157, 330–337 (2005)

[13] D.A. Battaglia, A. Braunstein, J. Chavas, R. Zecchina, Source coding by efficient selection of ground-state clusters, Phys. Rev. E 72, 015103 (2005)

[14] D.A. Battaglia, M. Kolář, R. Zecchina, Minimizing energy below the glass thresholds, Phys. Rev. E 70, 036107 (2004)

[15] D.A. Battaglia, M. Kolář, R. Zecchina, From statistical physics methods to algorithms, *proceedings of Quantum Many-Body theory XII-th workshop, Santa Fe*, to be published in Int. J. Mod. Phys. B (2005)

[16] D.A. Battaglia, G.E. Santoro, E. Tosatti, Optimization by quantum annealing: Lessons from hard satisfiability problems, Phys. Rev. E 71, 066707 (2005); reissued on Virtual Journal of Quantum Computation (july 2005)

[17] D.A. Battaglia, G.E. Santoro, L. Stella, E. Tosatti, O. Zagordi, contributed chapter in *Quantum Annealing and related optimization methods*, A. Das, B.K. Chakrabarti editors, Lecture notes in Physics series, vol. 679, Springer (2005)

[18] D.A. Battaglia, N. Shental, R. Zecchina, work in progress

[19] Bertin, D. Roncin, J. Vuillemin, Introduction to Programmable Active Memories, contributed chapter in *Systolic Array Processors*, J. McCanny, J. McWhirter, E. Swartzlander editors, Prentice Hall (1989)

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

[21] M. Blatt, S. Wiseman, E. Domany, Super-paramagnetic clustering of Data, Physical Review Letters 76, 3251 (1996)

[22] B. Bollobas, *Random Graphs*, 2nd edition, Cambridge University Press, UK (2001)

[23] J.- Bouchaud, L. F. Cugliandolo, J. Kurchan, M. Mézard, Out of equilibrium dynamics in spin-glasses and other glassy systems, in *Spin Glasses and Random Fields*, edited by A. Young (World Scientific 1997)

[24] J. Bouchaud, V. Dupuis, V. Hammann, E. Vincent, Separation of time and length scales in spin glasses: Temperature as a microscope, Phys. Rev. B 65, 24439 (2002)

[25] R.S. Braich, N. Chelyapov, C. Johnson, W.K. Rothemund, L. Adleman, Solution of a 20-variable 3-SAT Problem on a DNA Computer, Science 296, 499–502 (2002)

[26] A. Braunstein, M. Mézard, R. Zecchina, Survey propagation: an algorithm for satisfiability, Random Structures and Algorithms 27, 201–226 (2005)

[27] A. Braunstein, R. Mulet, A. Pagnani, M. Weigt, R. Zecchina, Polynomial iterative algorithms for coloring and analyzing random graphs, Phys. Rev E 68, 036702 (2003)

[28] A. Braunstein, M. Mézard, M. Weigt, R. Zecchina, Constraint satisfaction by Survey Propagation, preprint `cond-mat/0212451`

[29] A. Braunstein, R. Zecchina, Survey Propagation as local equilibrium equations, JSTAT P06007 (2004)

[30] J. Brooke, D. Bitko, T.F. Rosenbaum, G. Aeppli, Quantum annealing of a disordered magnet, Science 284, 779 (1999)

[31] N. Brunel, Persistent activity and the single cell f-I curve in a cortical network model, Network 11, 261–280 (2000)

[32] N. Brunel, X.J. Wang, Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition, J. Com Neuroscience 11, 63–85 (2001)

[33] M. Burrows, A block-sorting lossless data compression algorithm, Tech. Re 124, Digital SRC, Palo Alto(1994)

[34] G.Caire, S. Shamai, S.Verdù, A new data compression algorithm for sources with memory based on error correcting codes, in *proc. IEEE Information Theory Workshop*, 291–295, Paris, (2003)

[35] T. Castellani, V. Napolano, F. Ricci-Tersenghi, R. Zecchina, Bicoloring random hypergraphs, J. Phys. A 36, 11037 (2003)

[36] A. Cavagna, I. Giardina, G. Parisi, Cavity method for Supersymmetry breaking spin glasses, Phys. Rev. B 71, 024422 (2005)

[37] B.K. Chakrabarti, A. Dutta, Sen, *Quantum Ising Phases and Transitions in Transverse Ising Models*, Springer, Berlin–Heidelberg (1996)

[38] J. Chavas, D. Battaglia, A. Cicuttin, R. Zecchina, Construction and VHDL implementation of a fully local network with good reconstruction properties of the inputs, in *Artificial Intelligence and Knowledge Engineering Applications: a Bioinspired approach*, J. Mira, J.R. Alvarez editors, Part II, 385–394, Lecture Notes in Computer Science series, vol. 3562, Springer (2005)

[39] S.Y. Chung, G.D. Forney, T.J. Richardson, R. Urbanke, On the design of low-density-parity-check codes within 0.0045 dB of the Shannon limit, IEEE Communications Letters, vol. 5, no. 2, 58 - 60 (2001)

[40] S. Cocco, R. Monasson, A. Montanari, G. Semerjian, Approximate analysis of search algorithms with "physical" methods, in *Computational Complexity and Statistical Physics*, edited by G. Istrate, C. Moore, A. Percus, Oxford University Press, Oxford (2004)

[41] L. Correale, M. Leone, A. Pagnani, M. Weigt, R. Zecchina, Complex regulatory control in Boolean networks, preprint `cond-mat/0412443`

[42] N. Cutland, *Computability: an introduction to recursive function theory*, Cambridge University Press (1980)

[43] K.A. De Jong, Adaptive system design: a genetic approach, IEEE Transactions on Systems, man and Cybernetics, vol. SMC-10, 566–574 (1980)

[44] V. Dotsenko, Exact solution of the random bipartite matching model, J. Phys A 33, 2015 (2000)

[45] O. Dubois, Y. Boufkhad, J. Mandler, Typical random 3-SAT formulae and the satisfiability threshold, in *Proc. 11th ACM-SIAM Sym on Discrete Algorithms*, 124, San Francisco (2000)

[46] Special Issue on *NP-hardness and Phase transitions*, edited by O. Dubois, R. Monasson, B. Selman and R. Zecchina, Theor. Com Sci. 265, Issue: 1-2 (2001)

[47] G. Dueck, New optimization heuristics: the great deluge algorithm and the record-to-record travel, J. Comput. Phys. 104, 86–92 (1993)

[48] A. Engel, C. L Van den Broeck, *Statistical Mechanics of Learning*, Cambridge University Press, Cambridge, UK (2001)

[49] L. Euler, Solutio problematis ad geometriam situs pertinentis, Commentarii academiae scientiarum Petropolitanae 8, 128–140 (1741), reprinted in N. Biggs, E.K. Lloyd, R.J. Wilson, *Graph Theory, 1736–1936*, Clarendon Press, New York (1986)

[50] Erdos, A. Rényi, On the evolution of random graphs, Publ. Math, Inst. Hungar. Acad. Sci. 5, 17–61 (1960)

[51] E. Fahri, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, D. Preda, A quantum adiabatic evolution algorithm applied to random istances of an NP-complete problem, Science 292, 472 (2001)

[52] E. Fahri, S. Gutmann, Analog analogue of a quantum digital computation, Phys. Rev. A 57 (1996), 2403

[53] A.B. Finnila, M.A. Gomez, C. Sebenik, C. Stenson, J.D. Doll, Quantum annealing: a new method for minimizing multidimensional functions, Chem. Phys. Lett. 219, 343 (1994)

[54] S. Franz, M. Leone, Replica bounds for optimization problems and diluted spin systems, J. Stat. Phys. 111, 535 (2003)

[55] E. Friedgut, Necessary and sufficient conditions for sharp thresholds of graph properties, Journal of the A.M.S. 12, 1017 (1999)

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

[57] S. Fusi, J. Drew, L.F. Abbott, Cascade models of synaptically stored memories, Neuron 45, 599-611 (2005)

[58] M. R. Garey, D. S. Johnson, *Computers and intractability*, Freeman, New York (1979)

[59] I. Gent, T. Walsh, Towards an understanding of hill-climbing procedures for SAT, in *proc. of AAI '93*, 28–33, MIT-AAAI Press (1993)

[60] I. Gent, T. Walsh, The SAT Phase Transition, in *Proceedings of ECAI-94*, A.G. Cohn editor, 105, John Wiley & Sons (1994)

[61] M. Griniasty, M.V. Tsodyks, D.J. Amit, Conversion of temporal correlations between stimuli to spatial correlations between attractors, Neural Computation 5, 1–9 (1993)

[62] A. Grüning, Stack-and-queue-like dynamics in recurrent neural networks, Connection Science 17, to appear (2005)

[63] F. Guerra, Broken Replica Symmetry Bounds in the Mean Field Spin Glass Model, Comm. Math. Phys. 233, 1 (2003)

[64] W.R. Hamilton, Quart. J. Pure Applied Math. 5, 305 (1862)

[65] A. K. Hartmann, H. Rieger, *Optimization Algorithms in Physics* (Wiley-VCH, Berlin, 2001)

[66] A. K. Hartmann, H. Rieger, *New and Advanced Optimization Algorithms in Physics and Computational Science* (Wiley-VCH, Berlin, 2004)

[67] A.S. Herschel, Sir Wm. Hamilton's Icosian Game, Quart. J. Pure Applied Math. 5, 305 (1862)

[68] J.H. Holland, Outline for a logical theory of adaptive systems, Journ. ACM 9(3), 297–314 (1962)

[69] J.H. Holland, Processing and Processors for schemata, contributed chapter in *Associative Information Processing*, E.L. jacks editor, 127–146, Elsevier, New York (1971)

[70] T. Hogg, B.A. Huberman, C. Williams, Phase transitions and the search problem, Artificial Intelligence 81 I & II, 1–15 (1996)

[71] H.H. Hoos, T. St utzle, SATLIB: An online resource for research on SAT, in *SAT 2000*, I. Gent, H. Van Maren, T. Walsh editors, 283–292, IOS Press (2000). Online repository: `http://www.satlib.org`

[72] H.H. Hoos, T. St utzle, Local Search algorithms for SAT: an empirical evaluation, J. Automated Reasoning 24(4), 421–481 (2000)

[73] J. Inoue, Statistical mechanics of the multi-constraint continuous knapsack problem, J. Phys. A 30, 1047–1058

[74] M. Inoue, K. Hukushima, M. Okada, A PCA approach to Sourlas code analysis, Progr. Theor. Phys. Suppl. 157, 246–249 (2005)

[75] T. Kadowaki, H. Nishimori, Quantum annealing in the transverse Ising model, Phys. Rev. E 58, 5355 (1998)

[76] A. Kamath, N.K. Karmarkar, K.G. Ramakrishnan, M.G.C. Resende, A continuous approach to inductive inference, Mathematical Programming 57, 215–238 (1992)

[77] A. Kaporis, L. Kirousis, E. Lalas, The probabilistic analysis of a greedy satisfiability algorithm, in *Proceedings of the 4th European Symposium on Algorithms*, Lecture Notes in Computer Science 973, 574–585 (2002)

[78] H. KARLOFF, U. ZWICK, A 7/8-approximation algorithm for MAX-3-SAT?, in *Proceeding of 38th IEEE Symposium on Foundations of Computer Science*, IEEE Computer Soc. Press, 406–415, Los Alamos (1997)

[79] W. KASTELEYN, The statistics of dimers on a lattice I: The number of dimer arrangements on a square lattice, Physica 27, 1209–1225 (1961)

[80] H. KAUTZ, D. MCALLESTER, B. SELMAN, Encoding Plans in Propositional Logic, in *Proceedings of the 5th International Conference on the Principle of Knowledge Representation and Reasoning* (KR-96), 374–384, AAAI-Press (1996)

[81] H. KAUTZ, B. SELMAN, Pushing the Envelope: Planning, Propositional Logic, and Stochastic Search, in *Proc. of the Thirteenth National Conference on Artificial Intelligence*, 1194, AAAI-Press (1996)

[82] S. KIRKPATRICK, C. D. GELATT, M. VECCHI, Optimization by Simulated Annealing, Science 220, 671–680 (1983)

[83] S. KIRKPATRICK, B. SELMAN, Critical Behavior in the Satisfiability of Random Boolean Expressions, Science 264, 1297 (1994)

[84] H. KLEINERT, *Path Integrals in Quantum Mechanics, Statistics and Polymer Physics*, 2nd ed., World Scientific, Singapore (1995)

[85] E. KORUTCHEVA, M. OPPER, B. LOPEZ, Statistical mechanics of the knapsack problem, Jour. Phys A 27(18), L645 (1994)

[86] F. R. KSCHISCHANG, B. J. FREY, H.-A. LOELIGER, Factor graphs and the sum-product algorithm, IEEE Trans. Inf. Theory 47, 498 (2002)

[87] J.R. KOZA, *Genetic Programming: On the Programming of Computers by Means of Natural Selection*, MIT Press, Cambridge, MA (1992)

[88] D. MCALLESTER, B. SELMAN, H. KAUTZ, Evidence for invariants in local search, in *Proc. of AAAI '97*, 459–465, MIT-AAAI Press (1997)

[89] D.J.C. MACKAY, Information Theory, Inference, and Learning Algorithms, Cambridge University Press, Cambridge, MA (2003)

[90] D.J.C. MACKAY, R.M. NEAL, Near Shannon-limit performance of low density parity check codes, Electronics Letters vol. 32, 1645–1655; reprinted with corrected errors in vol. 33, 457–458 (1996)

[91] E. MANEVA, E. MOSSEL, M.J. MAINWRIGHT, A new look at Survey Propagation and its generalizations, in *Proc. SODA 2005*, ACM Press (2005)

[92] O. MARTIN, M. MÉZARD, O. RIVOIRE, Frozen glass phase in the Multi-Index Matching problem, Phys. Rev. Lett. 93, 217205 (2004)

[93] O. MARTIN, M. MÉZARD, O. RIVOIRE, Random Multi-Index Matching problems, preprint `cond-mat/0507180` (2005)

[94] R. MARTOŇÁK, G.E. SANTORO, E. TOSATTI, Quantum annealing by the path-integral Monte Carlo method: the two-dimensional random Ising model, Phys. Rev. B 66, 94203 (2002)

[95] R. MARTOŇÁK, G.E. SANTORO, E. TOSATTI, Quantum annealing of the Traveling Salesman Problem, Phys. Rev. E 70, 057701 (2004)

[96] B. MAZURE, L. SAIS, É. GRÉGOIRE, Tabu Search for SAT, in *Proc. of AAAI '97*, 281–285 (1997)

[97] B.W. MEL, Information processing in dendritic trees, Neural Computation 6(6), 1031–1085 (1994)

[98] S. MERTENS, Phase transition in the Number Partitioning Problem, Phys. Rev. Lett. 81(20), 4821–4284 (1998)

[99] S. Mertens, M. Mézard, R. Zecchina, Threshold values of Random K-SAT from the cavity method, preprint, cs.CC/0309020 (2003); to appear on Random Structures and Algorithms (2005)

[100] N. Metropolis, S. Ulam, The Monte Carlo method, J. Amer. Stat. Assoc. 44, 335–341 (1949)

[101] M. Mézard, T. Mora, R. Zecchina, Clustering of solutions in the random satisfiability problem, Phys. Rev. Lett. 94, 197205 (2005)

[102] M. Mézard, M. Palassini, O. Rivoire, Landscape of solutions in constraint satisfaction problems, preprint cond-mat/0507451

[103] M. Mézard, G. Parisi, A replica analysis of the traveling salesman problem, J. Physique 47. 1285–1296 (1986)

[104] M. Mézard, G. Parisi, On the solution of the random link matching problem, J. Phys. 48, 1451–1459 (1987)

[105] M. Mézard, G. Parisi, The Bethe lattice spin glass revisited, Euro. Phys. J. B 20(2), 217–233 (2001)

[106] M. Mézard, G. Parisi, The cavity method at zero temperature, J. Stat. Phys 111, 1 (2003)

[107] M. Mézard, G. Parisi, R. Zecchina, Analytic and algorithmic solutions of random satisfiability problems, Science 297, 812 (2002)

[108] M. Mézard, F. Ricci-Tersenghi, R. Zecchina, Alternative solutions to diluted p-spin models and XORSAT problems, J. Stat. Phys. 111, 505 (2003)

[109] M. Mézard, R. Zecchina, Random K-Satisfiability: from an analytic solution to an efficient algorithm, Phys. Rev. E 66, 056126 (2002)

[110] D.R. Mitchell, C. Adami, W. Lue, C. Williams, A random matrix model of adiabatic quantum computing, Phys. Rev. A 71, 052324 (2005)

[111] D. Mitchell, B. Selman, H. Levesque, Hard and Easy Distribution of SAT Problems, in *Proc. of the Tenth National Conference on Artificial Intelligence*, 459, AAAI-Press (1992)

[112] Y. Miyashita, H.S. Chang, Neuronal correlate of pictorial short term memory in the primate temporal cortex, Nature 331, 61 (1988)

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

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

[115] R. Monasson, R. Zecchina, Tricritical points in random combinatorics: the $(2 + p)$-SAT case, J. Phys. A 31, 9209–9217 (1997)

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

[117] R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, L. Troyanski, Determining computational complexity from characteristic phase transitions, Nature 400, 133 (1999)

[118] A. Montanari, G. Parisi, F. Ricci-Tersenghi, Instability of one-step replica-symmetry-broken phase in satisfiability problems, J. Phys. A 37, 2073 (2004)

[119] A. Montanari, F. Ricci-Tersenghi, On the cooling-schedule dependence of the dynamics of mean-field glasses, Phys. rev. B 70, 134406 (2004)

[120] A. Montanari, R. Zecchina, Optimizing searches via rare events, Phys. Rev. Lett. 88, 178701 (2002)

[121] T. Mora, M. Mézard, R. Zecchina, Pairs of SAT assignments and clustering in Random Boolean Formulae, preprint cond-mat/0506053

[122] M. Nielsen, I.L. Chuang, *Quantum Computation and Quantum Information*, Cambridge University press, Cambridge (2000)

[123] H. Nishimori, *Statistical Physics of Spin Glasses and Information Processing* (Oxford University Press, 2001)

[124] E. Oja, Neural networks, principal components and subspaces, Int. J. of Neural Systems 1(1), 61–68 (1989)

[125] Y. Okabe, M. Kikuchi, Vectorized coding for Monte Carlo simulation of the one-dimensional quantum spin-system, Phys. Rev. B 34(11), 7896-7900 (1986)

[126] L. Onsager, Electric moments of molecules in liquids, J. Am. Chem. Soc. 58, 1486–1493 (1936)

[127] R.H. Otten, L. van Ginneken, *The annealing algorithm*, Kluwer Academic Pub., Boston (1989)

[128] C.H. Papadimitriou, On selecting a satisfying truth assignment, in Proc. 32nd IEEE Symposium on the foundations of computer science, 163–169 (1991)

[129] C.H. Papadimitriou, *Computational Complexity*, Addison-Wesley, New York (1994).

[130] C. H. Papadimitriou, K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity* (Prentice-Hall, Englewood Cliffs, NJ, 1982)

[131] G. Parisi, *Statistical Field Theory*, Addison-Wesley, New York (1988)

[132] G. Parisi, A backtracking survey propagation algorithm for K-satisfiability, preprint cond-mat/0308510 (2003)

[133] G. Parisi, Some remarks on the survey decimation algorithm for K-satisfiability, intervention at the conference SAT 2004, Vancouver; preprint cs.CC/0301015 (2003)

[134] G. Parisi, M. Mézard, M.Virasoro, SK Model: The Replica solution without replicas, Europhys. Lett., 1(2), 77–82 (1986)

[135] G. Parisi, M. Mézard, M.Virasoro, *Spin Glass Theory and Beyond*, World Scientific, Singapore (1987)

[136] A.J. Parkes, Scaling properties of Pure Random walk on Random 3-SAT, in *Proceedings of the 8th International Conference on Principles and Practice of Constraint Programming*, Cornell University, Ithaca, NY; Lecture Notes in Computer Science 2470, 708 (2002)

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

[138] C.C.H. Petersen, R.C. Malenka, R.A. Nicoll, J.J. Hopfield, All-or-none potentiation at CA3-CA1 synapses, Proc. Nat. Acad. Sci. 95, 4723–4737 (1998)

[139] T. Richardson, A. Shokrollahi, R. Urbanke, Design of capacity-approaching irregular low-density-parity-check codes, IEEE Trans. Inform. Theory 47, 619–637 (2001)

[140] T. Richardson, R. Urbanke, An introduction to the analysis of iterative coding systems, in *Codes, Systems, and Graphical Models*, edited by B. Marcus and J. Rosenthal, Springer, New York, (2001)

[141] D. Ruelle, *Statistical Mechanics*, Benjamin, New York (1969)

[142] D. Salomon, *Data compression: the complete reference*, Springer-Verlag, New York (1998)

[143] G.E. Santoro, R. Martoňák, E. Tosatti, R. Car, Theory of quantum annealing of a an Ising spin glass, Science 295, 2427 (2002)

[144] U. Schöning, A Probabilistic Algorithm for K-SAT based on Limited Local Search and Restart, Algorithmica 32(4), 615-623 (2002)

[145] S. Seitz, Orponen, An efficient local search method for random 3-satisfiability, in *Proceedings LICS'03 Workshop on Typical Case Complexity and Phase Transitions*, Ottawa; Electronic Notes in Discrete Mathematics 16, Elsevier, Amsterdam (2003)

[146] S. Seitz, M. Alava, Orponen, Focused local search for random 3-satisfiability, J. Stat. Mech. P06006 (2005)

[147] B. Selman, H. Levesque, D. Mitchell, A new method for solving hard satisfiability problems, in *Proc. of AAAI '92*, 440–446, MIT-AAAI press (1992)

[148] B. Selman, H. Kautz and B. Cohen, Noise strategies for improving local search, in *Proceedings of AAAI*, 337–343, MIT-AAAI Press, Seattle (1994)

[149] G. Semerjian, R. Monasson, Relaxation and Metastability in a local search procedure for the random satisfiability problem, Phys. Rev. E 67, 066103 (2003)

[150] N. Sourlas, Spin-glass models as error-correcting codes, Nature 339, 693–695 (1989)

[151] D. Stauffer, A. Aharony, *Introduction to Percolation theory*, 2nd edition, Taylor and Francis, London (1992)

[152] L. Stella, G.E. Santoro, E. Tosatti, Optimization by quantum annealing: lessons from simple cases, Phys. Rev. B 72, 014303 (2005)

[153] L. Stella, G.E. Santoro, E. Tosatti, to appear

[154] M. Suzuki, Relationship between d-dimensional quantal spin systems and $(d+1)$-dimensional Ising systems: Equivalence, critical exponents and systematic approximants of the partition function and spin correlations, Prog. Theor. Phys. 56(5), 1454–1469 (1976)

[155] M. Talagrand, *Spin glasses: A challenge for mathematicians*, Series of Modern Surveys in Mathematics, vol. 46, Springer (2003)

[156] A. Treves, Frontal latching networks: a possible neural basis for infinite recursion, Cognitive Neurophysiology 21, 276–291 (2005)

[157] N. Trivedi, D.M. Ceperley, Green-function Monte Carlo study of the quantum antiferromagnets, Phys. Rev. B 40, 2737 (1989)

[158] W.T. Tutte, On Hamiltonian circuits, J. London Math. Soc. 21, 98-101 (1946)

[159] W. Van Dam, M. Mosca, U. Vazirani, How powerful is adiabatic Quantum Computation?, in *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science*, 279, IEEE press, Los Alamos (2001)

[160] M. Weigt, A.K. Hartmann, Number of guards needed by a museum: a phase transition in vertex covering of random graphs, Phys. Rev. Lett. 84(26), 6118-6121 (2000)

[161] J.S. Yedidia, W.T. Freeman and Y. Weiss, Generalized Belief Propagation, in *Advances in Neural Information Processing Systems 13* , 689–695, MIT Press (2001)

Debts

The three years of my PhD in Trieste have constituted for me a period of intense work but also of radical change in my life. I will never forget this strange and fascinating city, suspended at the frontier of "nowhere" [2].

I wish, first of all to thanks all my collaborators, Michal Kolář for his rigorousness and patience, Joël Chavas for his enthusiasm, Alfredo Braunstein for his continuous support and his changeless kindness.

I have not profited as much as I would of the wisdom of Prof. Erio Tosatti, but I am grateful to him for some deep suggestions and for his miraculous craft which literally transfigured one of my worst manuscripts into an unexpected small success.

I am greatly indebted to Giuseppe Santoro, for his unlimited assistance and his friendship. And, last but not least, to my supervisor Riccardo Zecchina, for the opportunity he gave me of conducting innovative and exciting research, for his bursting creativity, and for his extreme tolerance towards my disorder and my arrogance.

A warm gratitude goes finally to my first mentor Mario Rasetti for his solicitude; to my friends in SISSA and Via del Ronco (naming all of them would be too long), for the many memorable and pleasant days; and to my parents, who always sustained me and accepted my eccentric choice of becoming a wanderer-researcher.

Trieste, September 20th, 2005

[2] Cfr. Jan Morris, *Trieste and the meaning of nowhere*, New York, Simon and Schuster (2001).