



SAPIENZA
UNIVERSITÀ DI ROMA

Sapienza University of Rome

Scuola di Dottorato Vito Volterra - XXXIII ciclo
PhD in Physics

DOCTORAL THESIS

Exploiting atypicality in the Continuous Coloring

what makes *smart* algorithms actually smart

Thesis Advisor
Prof. Federico Ricci-Tersenghi

Candidate
Angelo Giorgio Cavaliere

April 2021

Contents

Introduction	1
I Theory and phenomenology of glassy systems	3
Preamble	4
1 Glass transition in supercooled liquids	6
1.1 A dramatic arrest	8
1.2 A precursor: two step relaxation	9
1.3 On the thermodynamic relevance of T_k	11
1.3.1 Adam-Gibbs-DiMarzio theory	11
1.3.2 Mosaic theory	12
1.3.3 Point-to-set correlation length	13
2 Glass transition in the mean-field limit	15
2.1 d 1RSB + s 1RSB framework	16
2.1.1 Equilibrium dynamics above T_d	17
2.1.2 Thermodynamics of the non-ideal glassy phase $T_k < T < T_d$	20
3 Finite connectivity mean-field: sparse graphs	23
3.1 Belief Propagation	24
3.1.1 Pairwise models	27
3.1.2 Replica Symmetric Population Dynamics	28
3.1.3 Long range correlations: need for 1RSB	29
3.2 Constraint satisfaction problems	31
3.2.1 Phase transitions in random CSPs	32
II Continuous coloring problem	35
Preamble	36
4 The model	39
4.1 Definitions	39
4.1.1 CSP definition	39
4.1.2 Planted model	41
4.1.3 Community detection definition	42
4.1.4 Mixed model	43
4.2 Belief propagation for CCP	45
4.2.1 RS population dynamics	47
4.2.2 Reconstruction equations for $m = 1$	48
4.2.3 Stability analysis	50
4.3 Continuous vs q -coloring	52

5	Simulations	56
5.1	Numerical determination of α_d	56
5.1.1	Population dynamics	56
5.1.2	Zero Energy Monte Carlo	60
5.2	Numerical determination of α_c	64
5.3	Further details	65
5.3.1	Extrapolation of the continuous limit	65
5.3.2	Global rotational symmetry	66
6	Postponing the dynamical transition	68
6.1	Biased potential	68
6.1.1	Optimized interaction for the Continuous Coloring	70
6.1.2	Discussion: the role of competing interactions	72
6.2	Complexity maximization	74
6.2.1	Computation of the gradient of the complexity	76
7	Algorithmic behaviour	80
7.1	<i>Smart</i> algorithms beating the threshold	80
7.2	A novel insight into a CSP phase diagram	82
7.2.1	$T = 0$ Monte Carlo algorithm	85
7.2.2	Numerical results	86
7.3	Quench and simulated annealing	91
7.3.1	Increasingly slow annealing	96
	Conclusions and Perspectives	99
	Bibliography	103
A	1RSB equations for $m = 1$	110

Introduction

Random constraint satisfaction problems are a fascinating subject. On the one hand, they are “realistic” enough to be numerically handled on a computer, thanks to the finite connectivity of the interaction network. On the other hand, they are truly mean-field in nature, hence something diametrically opposite to real world situation, but at least analytically solvable from the statistical physics perspective. A lot of effort has been devoted to the study of the average properties of typical random instances drawn from a given ensemble. Among the others, the most spectacular phenomenon affecting the structure of the space of solutions is perhaps the *clustering* (or dynamic) transition, happening at a ratio of constraints per variable dubbed α_d .

The clustering transition deeply affects the behaviour of local dynamics searching for solutions to a constraint satisfaction problem. The connection between static properties and dynamics is however not that simple and still far from being completely understood. Simple local search algorithms are often *smart* enough to construct solutions to real instances of the problem in linear or polynomial time in the system size up to the so called *easy-hard* transition threshold, which is algorithmic dependent and usually strictly *bigger* than α_d . Over the years, this intriguing phenomenology has been traced back to the fact that these procedures are intrinsically out-of-equilibrium, and are hence attracted towards atypical configurations with respect to the ones dominating the partition function of the model.

A promising scheme in this field consists in trying to connect non-trivial dynamical behaviour with the thermodynamic properties as obtained from a suitably tilted statistical measure, in the attempt to move some probability weight to solutions underrepresented in the original measure. We will pursue this route with what we believe to be the most simple strategy to count on [18]: adding a soft bias to the Hamiltonian of the model, in order to directly break uniformity in the sampling of solutions. To this end, we choose to work with a constraint satisfaction problem, the continuous coloring, that is defined on continuous variables and is characterized by an excluded volume interaction analogous to hard-spheres systems.

The thesis is organized as follows. In the first part a brief introductory overview is given on some topics that are functional to the development of the thesis: the discontinuous glass transition at the mean-field level, its finite connectivity realization on sparse random geometries and some of its connections to the more “traditional” physics of real supercooled liquids. The second part is entirely devoted to the analysis of the clustering phenomenon affecting the continuous coloring. In Chapter 4 a definition of the model is given along with the relevant belief propagation or cavity equations describing its phases. The phase diagram is presented and compared to the one of discrete q -coloring. In Chapter 5 we give the details of the numerical solution of the BP equations and of the estimation of the dynamic and condensation thresholds, along with the results of direct Monte

Carlo simulations supporting our analytical predictions. Chapter 6 is devoted to the study of the optimum bias that allows for the postponement of the dynamic transition threshold to the greatest connectivities. We will discuss of a methodology, which we have found most successful, that consists in iteratively maximizing the complexity of the typical clusters forming right above the dynamic threshold. Finally, in Chapter 7 our previous thermodynamic results are applied to the question of algorithmic behaviour, in order to gain some knowledge on the atypicality of the configurations visited by simple out-of-equilibrium Monte Carlo based procedures.

Part I

Theory and phenomenology of glassy systems

Preamble

Glasses are familiar, reassuring materials. They have accompanied the growth of mankind since the neolithic age, in the form of natural glasses, such as obsidian, before actual glass manufacturing was developed (cautious estimates date back to 3500 BC, in the Middle East region) [66]. For some sort of irony that will become clearer in a moment, even after such a huge amount of time, our theoretical comprehension of the glass state is far from being settled.

The first surprising fact about real glasses is their microscopic *amorphous* nature, meaning that they lack of the structural long-range order typical of solid crystals, despite being mechanically rigid. The radial distribution function, defined as the probability of finding another particle at a given distance from a reference one, is practically indistinguishable from the liquid case [50]. Essentially, glasses are liquids flowing so slowly upon decreasing temperature, that for all the practical purposes they can be considered as arrested. They are not, however, ever in equilibrium (even metastable) on experimental timescales.

Amorphism, on a more general ground, is not exclusive to particle glass-former liquids. These are also commonly referred to as *structural glasses*, in contrast to another very popular category of systems known as *spin-glasses*. Spin-glasses are complex models for magnetic systems that introduce a component of randomness in the Hamiltonian, also called quenched *disorder*. As a consequence, energy minima of the Hamiltonian, and in general the stable states of the system, do not usually possess any predictable notion of order, as it happens on the contrary for the crystalline (magnetized) states of matter.

The analogy with the spin-glass situation suggests that also in the case of structural glasses some sort of effective disorder is silently at work. This usually goes under the name of *self-induced disorder*, to underline its nature as a spontaneous emergent phenomenon in the huge space of configurations, rather than being directly encoded in random parameters entering the Hamiltonian. This similarity is even more compelling, as soon as one realises that the presence of *frustration* (which is the capability of the quenched randomness to destroy the ordered solutions) is not a necessary condition for a spin system to develop a spin-glass phase [14]. Emblematic in this respect is the case of the ferromagnetic p -spin on sparse graphs, which shows for low temperatures a ferromagnetic (crystalline) equilibrium phase along with a metastable glassy phase [40].

The important common denominator between this two types of “glasses”, is the presence of a highly non-trivial (energy) landscape in a high dimensional space, the number of translational degrees of freedom being $3N$ for a 3-dimensional particle system constituted of N particles. With highly non-trivial, we expect this landscape to be extremely rugged into a lot of (exponentially in N) stationary point, *i.e.* local minima and saddles. Being so numerous, most of these minima can only be amorphous, hardly distinguishable and hardly predictable. This situation is completely general,

and can be regarded as a key ingredient underlying a (possible) broader definition of glassiness.

These first introductory chapters will provide the reader with some background about the main topics that play a central role in the development of this thesis. First of all, some arbitrarily chosen results concerning the nature of the glass transition in supercooled liquids is briefly reviewed in Chapter 1. This introduction is utterly non exhaustive and awfully biased, given the vastness of the subject, and the fact that an unanimous theoretical framework is still lacking. Throughout this thesis, we will indeed adopt a point of view explicitly aimed to underline the deep analogies between the world of structural glasses and a certain class of spin-glasses undergoing the so called discontinuous glass transition, and which become manifest in the *mean-field* limit addressed in Chapter 2. That the mean-field description has actually something to say about real, finite dimensional world is still a matter of vigorous debate. The converse, however, is not necessarily so a delicate matter: the reader will take advantage of the beautiful analogy with real glassy systems in order to digest even the most exotic infinite dimensional models. As an incentive to follow this program, it comes the realization that the mean-field scenario is not limited to fully connected geometries or infinite spatial dimensions. Chapter 3 briefly presents the subject of models defined on sparse random geometries, which are both analytically tractable through belief propagation or cavity formalism and numerically manageable on a computer through direct simulation.

Chapter 1

Glass transition in supercooled liquids

When decreasing temperature down to the melting point T_m , a liquid generally undergoes a first-order phase transition towards a solid crystal state. Nevertheless, it is sometimes possible (but it is not always easy [23]) to avoid crystallization, for example by following a “fast enough” cooling schedule. Having succeeded in this program, one is left with a supercooled liquid¹.

Cooling further, at some point quite a striking phenomenon takes place: a dramatic increase (around 14 ~ 15 orders of magnitude) of the shear viscosity η on a relatively small interval of temperatures. This is also equivalent to consider a huge increase of the relaxation time τ , which is usually linked through direct proportionality to the viscosity. When τ is bigger than the available experimental time t_{exp} , the system is stuck out of equilibrium (it cannot be described as in a supercooled liquid metastable phase anymore). This is exactly the *empirical* glass transition T_g . Conventionally, one has $\eta(T_g) = 10^{12}\text{Pa} \cdot \text{s} = 10^{13}\text{Poise}$, which in turn corresponds to a relaxation time of $\tau \sim 10^2 - 10^3\text{s}$.

Below T_g , the system has not enough time to ergodically explore the phase space, and it suddenly loses a relevant portion of its degrees of freedom: a glass is formed. This is signaled by an abrupt drop of the constant pressure specific heat c_p . But what would happen if one was able to “equilibrate” the supercooled liquid even below T_g ? This impenitent speculation was firstly proposed by Kauzmann [53]. Let us imagine for a moment we succeeded in restoring ergodicity. The number of accessible microscopic configurations is in this case greatly enhanced with respect to the “arrested” glass phase. The usual way of quantifying this difference is by arguing that the total phase space of the putative supercooled liquid (SL) below T_g is divided into a number \mathcal{N}_c (where the subscript stands for *configurational*) of glassy states of typical internal size $\mathcal{N}_{\text{glass}}$, so that

$$\mathcal{N}_c = \frac{\mathcal{N}_{\text{SL}}}{\mathcal{N}_{\text{glass}}} = e^{N S_{\text{SL}} - N S_{\text{glass}}}, \quad (1.1)$$

where we have introduced the intensive entropies S_{SL} and S_{glass} .

A standard working practice is to approximate the S_{glass} contribution with the vibrational entropy typical of the solid crystal state, which is experimentally accessible. Calling excess entropy

¹Whenever a stable crystal phase is present, the supercooled liquid phase is metastable. This implies that it can be experimentally defined only on a time scale smaller than the nucleation time, and bigger than the relaxation time of the liquid, by maintaining an appropriate cooling rate. Such a time scale is not guaranteed to exist for every temperature. These problems can be disregarded if one studies suitable models where crystallization is suppressed, such as polydisperse sphere systems or frustrated spin-glass models.




Figure 1.1: Relevant temperatures and time scales in the though life of supercooled liquids. Below the liquid-crystal transition point T_m , a relentless non-trivial cooling schedule should be maintained in order to keep the system safe from unwanted crystallization. T_x is the Goldstein’s temperature below which activated dynamics kicks in. Activated dynamics is believed to prevent finite dimensional systems from experiencing a purely dynamic transition associated with the actual divergence of relaxation times at T_c (we will call it T_d in the rest of the thesis), as predicted instead from mode coupling theory and other mean-field models. At the glass transition T_g the relaxation time exceeds the experimental convention $\tau \sim 10^2 - 10^3$ s. At the Kauzmann temperature T_k the extrapolated excess entropy of the liquid with respect to the crystal appears to vanish. Finally, T_0 is the temperature where the Vogel-Fulcher-Tamman fit predicts a super-Arrhenius divergence of viscosity and relaxation time. Reprinted from [23].

S_{exc} the gap between the entropy of the liquid and that of the crystal, one finally obtains

$$S_c \sim S_{\text{exc}} = S_{\text{SL}} - S_{\text{crystal}}, \quad (1.2)$$

where $S_c = N^{-1} \log \mathcal{N}_c$ is known as *configurational entropy*.

This quantity, if extrapolated to low temperature from measurements in the metastable supercooled liquid phase, in some cases, as pointed out by Kauzmann, is observed to vanish at a nonzero temperature T_k . The great popularity of T_k is essentially due to two circumstances. On the one hand, as it will be discussed in the next section, there is an intriguing coincidence: the experimental viscosity data from several materials seems to diverge at a finite temperature $T_0 \sim T_k$. On the other hand, there is an analogy with the mean-field description, where T_k (*i.e.* the vanishing of a rigorously defined configurational entropy) can be proved to represent a true thermodynamic phase transition. Nevertheless, its occurrence in finite dimensional systems is still a point of debate, as it will be addressed in the rest of the chapter.

At least another key temperature can be identified in the quite hard life of real supercooled liquids, fig. 1.1. It corresponds to the Goldstein’s temperature $T_x > T_g$, that signals the point where activated dynamics between “metastable states”² kicks in, before ergodicity is finally lost at

²When talking about metastability in finite dimensions, one should keep in mind Goldstein’s scenario [44], for which the system is never *globally* found in a potential energy local minimum. Metastable states can be however reasonably defined *locally* both in space and in time, as addressed later on in this chapter.

T_g . Above T_x , experimental results are found in accordance with mean-field descriptions such as the mode coupling theory (MCT). These predict a sharp dynamic transition for $T_d \sim T_x$, with an actual divergence of shear viscosity and relaxation time, connected to the appearance of an exponential in N number of metastable states of infinite lifetime. However, in finite dimensional systems, this sharp transition happens to be avoided, and only a crossover takes place at T_x towards Goldstein's activated dynamics [44]. The underlying reason is that in real systems free-energy barriers are finite, and metastable states always possess a finite lifetime.

1.1 A dramatic arrest

A very popular manner to represent the dramatic slowing down of the dynamics of different glass-forming liquids is by plotting $\log \eta$ (or equivalently $\log \tau$) with respect to the inverse of temperature normalized to the specific value of T_g . This is usually called the Angell's plot [?], see fig. 1.2. This choice is particularly convenient, as it directly relates with the general Arrhenius law for




Figure 1.2: The Angell's plot is the traditional way to represent the striking increase of shear viscosity and relaxation time close to T_g . Glass-formers exhibiting a linear behaviour follow the Arrhenius law and are called “strong” liquids; “fragile” liquids exhibit instead a super-Arrhenius behaviour, characterized by an effective activation energy which increases when decreasing temperature. The whole range of intermediate behaviours between the truly strong and fragile liquids that are shown in the figure can be found in nature for different materials. Reprinted from [36].

activated events $\tau \propto \exp(\Delta/k_B T)$, where Δ is a constant (in temperature) activation energy. Liquids displaying an Arrhenius behaviour align on a straight line on the Angell's plot. These are called “strong” glass-formers. Remarkably, the vast majority of liquids does not follow an Arrhenius behaviour. A liquid is called as the more “fragile” the bigger the discrepancy from the Arrhenius law.

This classification hints to a very different behaviour of activation energies when lowering tem-

perature. Close to T_g , fragile glass-formers show a progressively steeper, super-Arrhenius increase of viscosity, that can be related to the growing of barriers. An empirical fit that reasonably suits the data is the Vogel-Fulcher-Tamman (VFT) law

$$\tau(T) = \tau_0 \exp\left(\frac{A}{T - T_0}\right). \quad (1.3)$$

The popularity of VFT is connected to the prediction, for fragile materials, of a finite temperature divergence of relaxation time at T_0 , compatible to the Kauzmann temperature T_k . In fact, thermodynamic theories for the glass transition have been proposed, linking the growth of barriers with the vanishing of S_c at T_k , and through the introduction of a diverging static correlation length. In this way, one can build a theoretical framework justifying the use of VFT.

At this point, some words of caution are required. The increase of viscosity is so sharp, and the range of low temperatures experimentally accessible is so limited, consistently far from the expected divergences, respectively at $T = 0$ or T_0 for strong or fragile glass-formers, that a definite validation of the different functional laws used to fit the data cannot be assured. In particular, fragile liquids obeying VFT, are found to be equally well described by fits with the same number of parameters, and only displaying a divergence for $T = 0$ [47]. On the other hand, even Arrhenius type glass-formers can be adequately fitted by crucially different functional forms, such as the Bässler law $\tau \propto \exp(A/T^2)$ [7]. In conclusion, these fits are not sufficient to validate, nor to reject, the existence of a finite temperature divergence $T_0 > 0$, to be associated to a hypothetical thermodynamic transition T_k .

1.2 A precursor: two step relaxation

A clear signature of the glass transition T_g can be obtained from the equilibrium relaxation in the supercooled liquid phase. This in turn allows us to put the physical relevance of T_g on a firmer basis, and underlies the exquisitely dynamical nature of the glass transition T_g . A dynamic correlation function is a quantity of the form

$$C(t_1, t_2) = \frac{1}{N} \sum_{i=1}^N \langle \phi_i(t_1) \phi_i(t_2) \rangle, \quad (1.4)$$

where $\phi_i(t)$ is a single particle observable at time t , and the average $\langle \cdot \rangle$ is carried over any stochasticity entering the dynamics, as the choice of the initial condition and the random “noise” component (*e.g.* thermal noise in Langevin dynamics, random sequences in Monte Carlo simulations).

At equilibrium time translational invariance (TTI) holds, hence two times quantities depend only on the time difference and $C(t_1, t_2) = C(t_2 - t_1) = C(t)$, where it was assumed $t_2 > t_1$. An experimentally accessible dynamic correlation, tracking in Fourier space the evolution of density fluctuations in supercooled liquids, is the self-intermediate scattering function $F_s(q, t) = \frac{1}{N} \sum_j \langle e^{i\mathbf{q}\cdot\mathbf{r}_j(t)} e^{-i\mathbf{q}\cdot\mathbf{r}_j(0)} \rangle$, fig. 1.3a. The decay of correlations with time noticeably depends on temperature. For high enough temperature, the decay is found to obey a simple exponential relaxation $\exp(-t/\tau)$. On the other hand, when lowering T , correlations start to exhibit a distinctive two-step relaxation, with the emergence of a plateau of increasing duration. The long time relaxation after the plateau is generally well fitted by a stretched exponential of the form $\exp[-(t/\tau)^\beta]$, with $\beta < 1$.




Figure 1.3: Two step equilibrium relaxation approaching T_g from above, as based on numerical simulations of Lennard-Jones glass-forming liquids. Left: self-intermediate scattering function, evaluated at the value of q corresponding to the first peak in the static structure factor. Reprinted from [36]. Right: for low temperatures, the crossover in the mean square displacement from ballistic to diffusive motion is divided by a plateau, similar in spirit to what happens with the scattering function and other dynamic correlations. Reprinted from [3].

A very peculiar property worth to stress, is that the plateau is *discontinuous* in nature. This means that for the first temperatures for which the decay is no longer exponential the dynamic correlation already shows the emergence of a plateau of nonzero height.

The two-step decay of the correlation hints at the presence of (at least) two distinct relaxation processes, whose timescales are progressively decoupled when lowering temperature. A suggestive physical interpretation for this decoupling can be found rooted in real space. To this end, one should consider the behaviour, as depicted in fig. 1.3b, of the mean square displacement (MSD)

$$\Delta(t) = \frac{1}{N} \sum_{i=1}^N \left\langle |\mathbf{r}_i(t) - \mathbf{r}_i(0)|^2 \right\rangle. \quad (1.5)$$

An early regime $\Delta \sim t^2$ is associated to the ballistic motion for times shorter than the typical collisional timescale. It is followed, for very long times, by a diffusive regime $\Delta \sim t$, dominated by collisions. In between, we assist to the appearance of a plateau, for sufficiently low T , which poses a stiff bound on the average distance a particle can explore even extending the available timescale t . The physical intuition behind the plateau in the MSD is that particles are trapped inside *cages* formed by the neighbours. The two-step relaxation of correlation functions can thus be interpreted as first, fast relaxation of each particle inside a cage, then followed by a second slow relaxation (the proper structural relaxation restoring ergodicity) only on timescales bigger than the confinement or caging time, *i.e.* the duration of the plateau.

This simple pictorial representation, albeit very appealing and often useful, has its own limitations. For instance, it does not clearly explain how cages are actually composed, which particle is the “jailer”, which one is the “prisoner”, or how they can be both at the same time. In fact, we expect the dramatic growth of the structural relaxation time to be connected to the need of cooperative rearrangements of an increasing number of particles, a situation very distant from a single particle caging-decaging description. It is worth to remember, however, that it can be misleading to regard

a unique mechanism as responsible for relaxation at *every* temperature $T > T_g$. If caging description may reasonably account for the appearance of the plateau at high enough temperatures, below the Goldstein crossover $T_x > T_g$ instead, where activated dynamics kicks in, different processes are expected to be at work, and even coexist in the same temperature range (this is the case of the so called slow α processes, responsible of the proper structural relaxation below T_x , and the Johari-Goldstein β -processes, or secondary relaxation, which occur on a timescale slower than the rattling-in-a-cage but faster than the α processes, and do not contribute to the structural relaxation of the liquid [52, 66]).

1.3 On the thermodynamic relevance of T_k

In this section I will elaborate a bit more on the presumed thermodynamic transition T_k , in order to clarify how a vanishing S_c should be connected with the divergence of a true static correlation length ξ . This perspective, albeit hypothetical, has the merit of having produced the most thriving and promising schemes that foreshadow a reconciliation between mean-field description (which predicts an analogous T_k scenario) and finite-dimensional world.

1.3.1 Adam-Gibbs-DiMarzio theory

One of the first organic theories proposing the existence of a true thermodynamic glass transition is due to Adam, Gibbs and DiMarzio [1, 42, 43]. As already argued when discussing caging at the end of the previous section, the increasingly slow relaxation close to T_g should be interpreted in terms of the cooperative rearrangement of a growing number of particles. These rearrangements have to be localised in space, allowing one to define a Cooperative Rearranging Region (CRR) as the smallest group of particles that can be rearranged independently from the rest of the system [1]. Let $n(T)$ be the typical number of particles constituting a CRR at a given temperature, and Ω be the number of “local states” a CRR can be found in. The definition of such local states (in place of a unique global metastable state) here essentially follows Goldstein’s idea of activated dynamics, for which, since activated events only involve a subset of the particles, many independent rearrangements can simultaneously occur in the global system, that is therefore always in the middle of a transition between energy landscape minima, rather than *in* a minimum [44].

We work under the assumption that Ω is independent from temperature and from the size of the CRR³. Treating different CRRs as uncorrelated, the number \mathcal{N} of global states the system as a whole can visit is simply

$$\mathcal{N} = \Omega^{N/n}, \quad (1.6)$$

from which the configurational entropy

$$S_c(T) = \frac{1}{N} \log \mathcal{N} = \frac{1}{n(T)} \log \Omega. \quad (1.7)$$

This implies that the size $n(T)$ of a typical rearranging region diverges at T_k , where the configurational entropy S_c goes to zero. From $n(T)$ we can extract a correlation length $\xi \sim n^{1/d}$, and relate

³This hypothesis may appear hard to digest. A more satisfying description is provided in this section in terms of mosaic theory and point-to-set correlation length.

it to the scaling of the barrier $\Delta \sim \xi^\psi$ for some $\psi > 0$. Adam-Gibbs' choice was simply $\Delta \sim n(T)$, thus leading to

$$\tau \propto \exp\left(\frac{A}{TS_c(T)}\right). \quad (1.8)$$

Remarkably, this expression reduces to the VFT law close to T_k , since

$$S_c(T) \sim S_{\text{exc}}(T) = \int_{T_k}^T \frac{\Delta c_p}{T} \sim \Delta c_p \log(T/T_k) \sim \Delta c_p \frac{T - T_k}{T_k}, \quad (1.9)$$

where we have used $S_{\text{exc}}(T_k) = 0$, $\Delta c_p = c_p^{\text{SL}} - c_p^{\text{crystal}} \sim \text{const}$, and expanded the logarithm for $T \sim T_k$. The extra T factor multiplying S_c in the activation time τ can be finally treated as constant T_k , if T_k is sufficiently far from $T = 0$, thus obtaining back eq. (1.3).

1.3.2 Mosaic theory

The *mosaic* theory shares many of the key ideas of Adam-Gibbs' scenario. It proposes a nucleation mechanism between a *very large* number of metastable states [56], a picture directly borrowed from mean-field phenomenology. In this framework, contrary to Adam-Gibbs' formulation, the number of potentially accessible states to each cooperative unit does scale (exponentially) with its size. A variation on classic nucleation theory is assumed to be

$$\Delta F(R) \sim Y(T)R^\theta - TS_c(T)R^d, \quad (1.10)$$

where R is the size of a spontaneous nucleus of a given state. $Y > 0$ is a generalized surface tension that accounts for phase mismatch at the boundary of the nucleated region ($\theta \leq d - 1$), and represents a cost in terms of free-energy balance. The second term (entropic *driving force*), replaces the usual bulk free-energy gain, that is here zero since all the states have the same free-energy (they are truly indistinguishable indeed).

According to the nucleation mechanism, small R fluctuations are then hindered by surface tension, while become advantageous only beyond a typical scale ξ for which the barrier $\Delta F(\xi)$ reaches its maximum and starts decreasing⁴,

$$\xi(T) \sim \left(\frac{Y(T)}{TS_c(T)}\right)^{\frac{1}{d-\theta}}, \quad (1.11)$$

$$\Delta(T) \sim \Delta F(\xi) \sim Y(T)\xi^\theta. \quad (1.12)$$

This result straightforwardly redirects to a generalized VFT law for the relaxation time (an extra free parameter is given by the exponent θ).

The suggestive physical picture arising from the nucleation scenario is the one interpreting the structure of real supercooled liquids as a patchwork (mosaic⁵) of local cooperative regions of size ξ . The most convincing intuition about the nature of the entropic driving term proceeds as follows [13, 23]. Regions smaller than ξ cannot decorrelate autonomously, due to surface tension

⁴Here we keep track only of the dependence on temperature, disregarding constant prefactors. Notice that for $R = \xi$ the two contributions to ΔF scale with T in the same manner.

⁵A very unexciting one, actually, as long as one is concerned about "artistic" merit. All the local metastable states would indeed appear identically amorphous to our eye.

free-energy price that pushes back fluctuations attempting to leave the current state. Above ξ , however, fluctuations leaving the state are secured since the group of particles cannot find its way back to the original state, due to the exponential proliferation of accessible basins with the cluster size. Because of the local and cooperative nature of rearrangements, these typically occur, independently throughout the system, *only* on length scales close to ξ , being the rate of more extended rearrangements exponentially suppressed in the size of the region.

1.3.3 Point-to-set correlation length

The above construction, in particular the physical interpretation of the entropic driving term, and how it affects the instability of a system towards mosaic fragmentation, may appear a bit exotic, if not delusive. A recent formulation [13] provides a perhaps more convincing argument supporting the mosaic picture, while highlighting at once the importance of locality, both in space and in time, in the definition of finite-dimensional metastability.

Consider the following experiment, in which one “freezes” the positions of all the particles outside a spherical cavity of radius R from a reference equilibrium liquid configuration. Particles inside the sphere are left free to move, but are subjected to the “pinning” field due to the particular choice of fixed boundary conditions. One can show that the notion of metastable state is thermodynamically consistent only on length scales smaller than ξ , which correspond to the typical size of a patch in the mosaic state. To this end, suppose the initial configuration was prepared in a true global state α . Particles inside the spherical cavity, on the other hand, are now free to choose one of the many possible $\exp[S_c(T)R^d]$ different states. The thermodynamic partition function of the particles inside the sphere accounts for two class of contributions, depending of whether state α is selected (Z_{in}) or not (Z_{out}) in the summation over all the possible states,

$$Z_{\text{sphere}} = Z_{\text{in}} + Z_{\text{out}} = \exp(-\beta f R^d + \beta Y R^\theta) + \exp(-\beta f R^d + S_c R^d), \quad (1.13)$$

where f is the internal free energy of every dominating metastable state, and Y is the surface tension (this term can be viewed as a positive contribution to Z_{in} coming from the favourable pinning field at the boundary of the region, or can be equivalently moved inside Z_{out} , with a change of sign, representing the surface tension cost payed by all the states but α).

The probabilities of finding (not finding) the spherical region in the original state α are then

$$p_{\text{in}}(R) = \frac{Z_{\text{in}}}{Z_{\text{in}} + Z_{\text{out}}} = \frac{\exp(\beta Y R^\theta)}{\exp(\beta Y R^\theta) + \exp(S_c R^d)}, \quad (1.14)$$

$$p_{\text{out}}(R) = \frac{Z_{\text{out}}}{Z_{\text{in}} + Z_{\text{out}}} = \frac{\exp(S_c R^d)}{\exp(\beta Y R^\theta) + \exp(S_c R^d)}, \quad (1.15)$$

which exemplify the roles of $Y(T)$ and $S_c(T)$ respectively as an obstacle and a driving force to rearrangements.

In this case, R represents an externally tunable parameter (at least theoretically and in numerical simulations), that can be used at each temperature to probe different length scales of the system in a rigorous way. Two regimes are obtained. For $R < \xi(T)$, p_{in} dominates ($\theta < d$) thanks to the exponential form of the probability, while for $R > \xi(T)$ one has the opposite. The length ξ , for which the exponents are balanced, is given exactly by eq. (1.11). We conclude that regions larger

than ξ are eventually able to decorrelate, *even* in the presence of a “symmetry breaking” pinning field. The assumption of a “global” metastable state is hence disproved, and metastability can be defined only up to length scale ξ . Indeed, for $R < \xi$, we can safely affirm that the sphere is stuck inside state α , as long as the constraint on the boundary conditions is not relaxed.

This rigorous construction, albeit fictitious, provides the most convincing, closest realisation of mean-field like scenario in finite-dimensional systems. As a byproduct, we are left with an operative definition of a static *point-to-set* correlation length ξ , that can be directly obtained in numerical simulations [9–11, 21]. This is very convenient, since an experimental validation of any of the theories exposed in this section can be rather difficult [51], and far from being conclusive⁶.

⁶Recall, for instance, that the experimental agreement of shear viscosity and relaxation time with VFT-like fits cannot be considered as a proof of the theories.

Chapter 2

Glass transition in the mean-field limit

So far our attention was mostly focused, for obvious anthropical reasons, on the finite-dimensional nature of real systems. The quite not insignificant detail that real world physics is generally bounded to three spatial dimensions, has not prevented anyway physicists from building unrealistic theoretical models which are infinite dimensional by construction. Spheres systems in $d \rightarrow \infty$ spatial dimensions have been exploited rather recently [68, 90–92]. Another very popular example are models defined on a fully connected geometry, as the Curie-Weiss (Ising) model or its spin-glass counterpart, the Sherrington-Kirkpatrick model. In all these cases, each variable interacts with $O(N)$ other variables, and fluctuation from the average behaviour is not tolerated. The many body problem of N interacting particles can then be generally reduced to a single particle description in an “average” external field.

Fully connected models are particularly suitable to analytical treatment, but may appear rather abstract and unsatisfactory from the physical point of view. In particular, real world degrees of freedom usually share a finite number of interactions. Furthermore, a major limitation to this approach is that numerical simulations scale prohibitively with the system size. Another class of mean-field models trying to answer to these issues are those defined on finite connectivity, locally tree-like structures, also called *sparse graphs*. Among them we have diluted spin-glass models, constraint satisfaction problems from computer science (that will be discussed in the next chapter), but also particle systems such as lattice glass models on Bethe lattices [12, 27, 62, 94] or the Mari-Kurchan-Krzakala (MKK) model [70]. To the peculiarity of tree-like geometries is entirely devoted the next chapter. In a few words, it strongly derives from the absence of short *loops* in the network of interactions, this washing out any Euclidean space structure. This is also what happens, for instance, in another mean-field realization of particle models due to Mari-Kurchan [71], which although living in finite spatial dimensions, lacks of the triangular property typical of short-range interacting particles (if particle A is close to particle B, that in turn interacts with particle C, then A and C must be relatively close together and likely interact) thanks to the introduction of arbitrarily large random shifts between each pair of particles.

The fundamental characteristic of mean-field systems, which radically distinguishes them from finite-dimensional ones, is that free energy barriers scale with the system size, and hence are truly infinite in the thermodynamic limit. This in particular allows for a rigorous definition of global metastables states, which in this context have infinite lifetime, in turn making it possible to *count* them. A suggestive idea according to which mean-field physics could be used to build a glass

transition theory for real systems was steadily proposed by Kirkpatrick, Thirumalai and Wolynes. The inspiration, and hope, comes from the field theory for standard second order critical phenomena, where an expansion out of mean-field allows one to recover three dimensional world (in a very recent formulation [2], such an expansion is performed starting from a theory defined on a tree-like structure called Bethe lattice, thus testifying the importance and versatility of the study of such geometries). In the case of the glass transition, however, the situation is crucially complicated by the presence of activated events in finite dimensions, which drastically alters the mean-field description.

It is not in the purpose of this thesis to pursue this direction. However, it may be of general interest to briefly consider some of the key outcomes of this approach. The same authors started from the realisation [54] that the equations describing the dynamics of a fully connected spin-glass model, the p -spin for $p \geq 3$, are formally equivalent to those predicted by Mode Coupling Theory (MCT), an approximate scheme for closing the complicated dynamic equations in real liquids [8, 45, 46]. Remarkably, MCT predictions agree fairly well with experimental data in a range of temperatures bigger than but not too close to T_g . This opens to two important lines of thought. Firstly, it proposes MCT as a sort of mean-field-like theory for the dynamics of liquids. This has been essentially confirmed by the study of spheres systems in infinite dimensions [68], meaning that MCT in the $d \rightarrow \infty$ limit is found to be qualitatively very close to the exact mean-field solution, even though important quantitative differences are generally to be taken into account [48]. Secondly, it poses a parallelism between structural glasses and a certain class of spin-glasses undergoing the so called discontinuous glass transition, thus envisaging a careful study of the latter type of systems as a relevant tool in order to also gain comprehension on the structural glass transition itself [57, 58]. This whole framework goes under the name of Random First Order Transition theory (RFOT). It is strongly oriented towards finite dimensional systems, as testified by the very fabric of the mosaic theory developed by the same authors in this context [56], where the “first order” nature of the transition is encoded in all those metastability and nucleation arguments that have been briefly addressed in the previous chapter. Despite this fact, the term RFOT is commonly used also to refer, in general, to the transition in any kind of glassy models, even spin systems, that undergo the so called one-step Replica Symmetry Breaking (1RSB) transition at the mean-field level. It is this second meaning that will be referred to in the rest of the thesis.

2.1 $d1\text{RSB} + s1\text{RSB}$ framework

The 1RSB transition, or *discontinuous* glass transition, or RFOT, predicts *two* important transition temperatures:

- T_d , or dynamic transition temperature, is connected to the appearance of an exponential in N number of metastable states for the system. Despite not being signaled by thermodynamics, it has strong consequences on the dynamics, with a true divergence of the relaxation time τ and the development of a persistent plateau in the dynamic correlation functions for $T \rightarrow T_d^+$. The divergence of τ is only power-law, rather than showing an exponential behaviour as in real liquids close to T_g . The dynamic transition is actually believed to be an *avoided* transition in finite dimensions, and it is associated to the crossover between non activated and activated dynamics $T_d \sim T_x > T_g$.

- T_s , or static transition temperature, is the point where the number of dominating metastable states becomes sub-exponential in the system size. This is equivalent to say that the *complexity* Σ , *i.e.* the logarithm of their number divided by N , vanishes at T_s . The definition of the complexity closely reminds us the concept of configurational entropy in the context of supercooled liquids. In fact, it can be viewed as its “rigorous” counterpart for mean-field systems, where global metastability is properly defined. This implies that the Kauzmann scenario, or ideal glass transition, can be exactly realized within mean-field ($T_s = T_k$ is a true thermodynamic transition, since states dominating the partition function have now associated an individual, non-vanishing statistical weight).

2.1.1 Equilibrium dynamics above T_d

The equilibrium behaviour close to T_d is conveniently studied from the point of view of dynamics. For definiteness, let us consider the case of the fully-connected spherical p -spin model, which for $p \geq 3$ is known to undergo a discontinuous glass transition [30, 55], and that is defined by the following Hamiltonian

$$\mathcal{H} = - \sum_{i_1 < \dots < i_p} J_{i_1 \dots i_p} \sigma_{i_1} \cdots \sigma_{i_p}, \quad \sigma_i \in \mathbb{R}, \quad \sum_{i=1}^N \sigma_i^2 = N \quad \forall i, \quad (2.1)$$

where the $J_{i_1 \dots i_p}$'s are quenched Gaussian variables of zero mean and variance $p!/(2N^{p-1})$. The relaxational dynamics for the $\sigma_i(t)$'s is assumed to be given by the set of N coupled Langevin equations

$$\partial_t \sigma_i(t) + \mu(t) \sigma_i(t) + \frac{\partial \mathcal{H}_J}{\partial \sigma_i} = \xi_i(t), \quad (2.2)$$

$$\langle \xi_i(t) \xi_j(t') \rangle = 2T \delta_{ij} \delta(t - t'), \quad (2.3)$$

where $\mu(t)$ is a Lagrange multiplier enforcing the spherical constraint at each time and to be consistently derived from the equations, and ξ_i is the usual delta correlated noise (white noise) of zero mean.

The subscript in \mathcal{H}_J signals the fact that the problem still depends on the quenched disorder. Following De Dominicis [35], one can use the dynamic functional integral method of Martin, Siggia and Rose [73] to average out the quenched disorder without the use of the replica method. This in turn allows one to rewrite the original set of Langevin equations (heavily coupled through the term $\partial \mathcal{H}_J / \partial \sigma_i$) as a single-particle Langevin process with colored noise

$$\partial_t \sigma(t) + \mu(t) \sigma(t) - h(t) - \frac{p(p-1)}{2} \int dt' R(t, t') C^{p-2}(t, t') \sigma(t') = \xi(t), \quad (2.4)$$

$$\langle \xi(t) \xi(t') \rangle = 2T \delta(t - t') + \frac{p}{2} C^{p-1}(t, t'), \quad (2.5)$$

having introduced the average correlation and response functions

$$C(t, t') = \langle \sigma(t) \sigma(t') \rangle, \quad R(t, t') = \frac{\partial \langle \sigma(t) \rangle}{\partial h(t')}, \quad (2.6)$$

where $\langle \cdot \rangle$ means average over the probability measure of ξ , and $h(t)$ is an external magnetic field.

From relations (2.4)-(2.5), one can derive self-consistent equations for the correlation and the response function [29, 55], which for $h = 0$ read

$$\frac{\partial R(t_1, t_2)}{\partial t_1} = -\mu(t_1)R(t_1, t_2) + \frac{p(p-1)}{2} \int_{t_2}^{t_1} dt R(t_1, t) C^{p-2}(t_1, t) R(t, t_2) + \delta(t_1 - t_2), \quad (2.7)$$

$$\begin{aligned} \frac{\partial C(t_1, t_2)}{\partial t_1} = & -\mu(t_1)C(t_1, t_2) + \frac{p(p-1)}{2} \int_{-\infty}^{t_1} dt R(t_1, t) C^{p-2}(t_1, t) C(t, t_2) + \\ & + 2TR(t_2, t_1) + \frac{p}{2} \int_{-\infty}^{t_2} dt C^{p-1}(t, t_1) R(t_2, t), \end{aligned} \quad (2.8)$$

$$\mu(t_1) = T + \frac{p^2}{2} \int_{-\infty}^{t_1} dt R(t_1, t) C^{p-1}(t_1, t). \quad (2.9)$$

In the equilibrium regime $T > T_d$, they drastically simplify thanks to TTI and to the fluctuation dissipation theorem (FDT), which allows us to express the response in terms of the sole correlation

$$R(t) = -\frac{\Theta(t)}{T} \frac{\partial C(t)}{\partial t}, \quad t = t_1 - t_2, \quad (2.10)$$

where $\Theta(t < 0) = 0$ and $\Theta(t > 0) = 1$. Moreover, ergodicity in the paramagnetic phase requires the correlation to satisfy $\lim_{t \rightarrow \infty} C(t) = 0$. The self-consistency equations then reduce to

$$\partial_t C(t) = -TC(t) - \frac{p}{2T} \int_0^t du C^{p-1}(t-u) \partial_u C(u). \quad (2.11)$$

Very interesting physics is encoded in this equation. To elucidate this point, it is convenient to recast it in the following form

$$\begin{aligned} \partial_t C(t) = & -TC(t) + \frac{p}{2T} C^{p-1}(t) [1 - C(t)] - \frac{p}{2T} \int_0^t du [C^{p-1}(t-u) - C^{p-1}(t)] \partial_u C(u) \\ = & -TC(t) + \frac{p}{2T} C^{p-1}(t) [1 - C(t)] + A_C(t) \leq 0, \end{aligned} \quad (2.12)$$

where we have imposed the condition $\partial_t C(t) \leq 0$ since $C(t)$ is a decreasing function of time. From $0 \leq C(t) \leq 1$ it also follows that the first term in the rhs of (2.12) is negative, while the other two are positive. In particular, $A_C(t) \geq 0$, so that we can write

$$-TC(t) + \frac{p}{2T} C^{p-1}(t) [1 - C(t)] \leq 0. \quad (2.13)$$

This leads us to define [22]

$$g(C) \equiv C^{p-2}(t) [1 - C(t)] \leq \frac{2T^2}{p}. \quad (2.14)$$

The function $g(C)$ is equal to 0 for both $C = 1$ and $C = 0$ and has a maximum in $C = q_d$, where

$$q_d = \frac{p-2}{p-1}. \quad (2.15)$$

Inequality (2.14) can be studied graphically. For high temperatures the inequality is satisfied for all times from $t = 0$ ($C(0) = 1$, $g(C = 1) = 0$) to $t = \infty$ ($C(\infty) = 0$, $g(C = 0) = 0$). Decreasing T , the difference between the rhs and $g(q_d)$ becomes smaller. From equation (2.12) we recognize

the difference $\Delta(T, t) = 2T^2/p - g(C(t))$ to be “related” to the time derivative of the correlation. If we neglect for the moment $A_C(t)$, we obtain a direct proportionality between this difference and $\partial_t C(t)/C(t)$: for small times, and C close to 1, $\Delta(T, t)$ is maximum as a function of t and the correlation decays fast; for “intermediate times”, when C is close to q_d , $\Delta(T, t)$ approaches zero and hence dynamics slows down; once we get past q_d , $\Delta(T, t)$ starts increasing again and the correlation can decay down to $C(\infty) = 0$. We put the second time-window name “intermediate times” under quotes, since we expect that the time to reach q_d becomes bigger and bigger, and actually diverges, when decreasing temperature. We can lower T down to the point for which (2.14) is saturated, $g(q_d) = 2T^2/p$. This defines the dynamic transition temperature

$$T_d = \sqrt{\frac{p(p-2)^{p-2}}{2(p-1)^{p-1}}} \quad (2.16)$$

as the temperature for which the system finally gets stuck in the long time limit to a plateau $\lim_{t \rightarrow \infty} C(t) = q_d$, $\lim_{t \rightarrow \infty} \partial_t C(t) = 0$. Notice that at this point the ergodicity assumption breaks down and we should go back to the original equations if we want to study the dynamics for $T < T_d$. A final remark is required in order to convince the reader that the presence of $A_C(t)$ is actually harmless to our conclusions. The delicate point here is the fact that the quantity $C(t)\Delta(T, t)$ is not in general proportional to $\partial_t C(t)$ for arbitrary shapes of $A_C(t)$. However, the whole picture would be safe¹ if $A_C(t)$ resulted to vanish for $t \rightarrow \infty$. The fact that this is verified can be readily seen from the definition of $A_C(t)$, equation (2.12), once one realizes that the two terms under integration alternatively vanish on complementary time windows. To make it more explicit, it is convenient to divide the range of integration $(0, t)$ in two intervals $(0, X)$ and (X, t) , where X is a diverging timescale satisfying $X \ll t$ for $t \rightarrow \infty$, e.g. $X = t^{1/2}$.

Mode coupling theory provides very important relations for the asymptotic behaviour of the correlation function. In particular, the way it decays at the dynamical transition T_d towards its long time plateau (also known as critical decay) or how it departs from it close to T_d (also known as the Von Schweidler law), is governed by two power-laws with different exponents that are commonly called a and b respectively

$$C(t) \simeq q_d + (t/t_0)^{-a} + O\left[(t/t_0)^{-2a}\right], \quad t_0 \ll t \ll t_\sigma, \quad T \rightarrow T_d^+, \quad (2.17)$$

$$C(t) \simeq q_d - (t/\tau)^b + O\left[(t/\tau)^{2b}\right], \quad t_\sigma \ll t \ll \tau, \quad T \rightarrow T_d^+, \quad (2.18)$$

where t_0 is a characteristic microscopic timescale, t_σ is a diverging timescale associated with the approach to the plateau (β -process), while the relaxation time $\tau \gg t_\sigma$ is a second diverging timescale associated with the proper structural relaxation of the liquid (α -process). The previous relations are to be intended valid up to terms of order $\epsilon = (T - T_d)/T_d$. The timescales are defined in terms of the same exponents a and b as

$$t_\sigma \sim (T - T_d)^{-\frac{1}{2a}}, \quad (2.19)$$

¹Suppose that $\lim_{t \rightarrow \infty} C(t) = q'$ and $\lim_{t \rightarrow \infty} \partial_t C(t) = 0$ for $T = T'$, with $q' \geq q_d$ and $T' \geq T_d$. From (2.12) we have then $\lim_{t \rightarrow \infty} \partial_t C(t) = -\frac{p}{2T} q' \Delta(q', T') + \lim_{t \rightarrow \infty} A_C(t) = 0$. In general, if $\lim_{t \rightarrow \infty} A_C(t) > 0$, then $\Delta(q', T') > 0$, which implies $T' > T_d$ or $T' = T_d$, $q' \neq q_d$. We may be worried in this case for the emergence of a plateau with $q' > q_d$ and/or $T' > T_d$. But since we show in the text that $\lim_{t \rightarrow \infty} A_C(t) = 0$ for *any* temperature T' , then it follows that the only condition for having a long time plateau with vanishing $\lim_{t \rightarrow \infty} \partial_t C(t) = 0$ is that $\Delta(q', T') = 0$, which gives us back $q' = q_d$ and $T' = T_d$.

$$\tau \sim (T - T_d)^{-\gamma}, \quad \gamma = \frac{1}{2a} + \frac{1}{2b}. \quad (2.20)$$

Equation (2.20) highlights an important difference with real liquids, where τ scales exponentially close to T_g . Finally, another remarkable prediction of MCT is that the exponents a and b are not independent, but are both determined by the exponent parameter λ

$$\frac{\Gamma^2(1-a)}{\Gamma(1-2a)} = \frac{\Gamma^2(1+b)}{\Gamma(1+2b)} \equiv \lambda, \quad (2.21)$$

with $1/2 \leq \lambda < 1$, from which $0 < a \leq 0.395\dots$ and $0 < b \leq 1$. The parameter λ is model dependent. For instance, in the case of the fully connected spherical p -spin, one obtains for any value of p that $\lambda = 1/2$, *i.e.* $a = 0.395\dots$ and $b = 1$. Remarkably, besides direct dynamic computation [29], also a static route is sometimes possible [20, 39].

2.1.2 Thermodynamics of the non-ideal glassy phase $T_k < T < T_d$

Below T_d , thermodynamics is dominated by an exponential subset of the exponentially many states, rather than by a simply ergodic paramagnetic (liquid) solution. Each of them possesses an internal free energy f^* bigger than the one obtained by prolonging the liquid solution. Nevertheless, the total free energy density Φ^* , which is essentially proportional to the logarithm of the thermodynamic partition function, takes into account all the dominating states simultaneously, so that

$$-\beta\Phi^* = -\beta f^* + \Sigma(f^*, T). \quad (2.22)$$

From this in particular follows that $\Phi^* < f^*$, since $\Sigma > 0$. Moreover, one can show that Φ^* is the analytic continuation of the paramagnetic (liquid) solution down to T_s , so that no thermodynamic indication of a transition comes with T_d .

To elucidate this point, we recall that the free energy Φ can be directly computed via the replica method [81] in the so called 1RSB scheme, which for RFOT models (as the p -spin $p \geq 3$) is found to be exact. In the case of the p -spin [22], a generalized free energy can be written as a function of the external parameters (such as the inverse temperature β , and the external magnetic field h which we neglect here) and three other parameters, $\Phi_{1\text{RSB}}(q_0, q_1, m; \beta)$. The physical values of these parameters as a function of temperature are extracted from $\Phi_{1\text{RSB}}$ by finding its stationary points,

$$\left. \frac{\partial \Phi_{1\text{RSB}}}{\partial q_0} \right|_{q_0^*, q_1^*, m^*} = \left. \frac{\partial \Phi_{1\text{RSB}}}{\partial q_1} \right|_{q_0^*, q_1^*, m^*} = \left. \frac{\partial \Phi_{1\text{RSB}}}{\partial m} \right|_{q_0^*, q_1^*, m^*} = 0, \quad (2.23)$$

so that $\Phi^* = \Phi_{1\text{RSB}}(q_0^*, q_1^*, m^*; \beta)$. Of course, these parameters have a very important physical interpretation. To this end, it is convenient to define the *overlap* as a measure of the similarity between two generic configurations $\{\sigma_i\}$ and $\{\tau_i\}$, for instance in the case of scalar spin variables

$$q_{\sigma\tau} = \sum_{i=1}^N \sigma_i \tau_i. \quad (2.24)$$

One can define also the overlap between two mean-field states α and β

$$q_{\alpha\beta} = \sum_{i=1}^N \langle \sigma_i \rangle_{\alpha} \langle \sigma_i \rangle_{\beta}, \quad (2.25)$$

where $\langle \cdot \rangle_{\alpha, \beta}$ denotes the average over the Boltzmann-Gibbs distributions restricted to state α or β respectively. The physical intuition underlying the replica formalism and the 1RSB scenario is the following. Let us consider two independent replicas of the same system. If only one state is present, as for $T > T_d$, both replicas will be in the same state, and we need only one parameter to describe the equilibrium overlap distribution, the self-overlap $q_{\alpha\alpha}$ of *that* state. If there are many states, the situation could in principle be very complicated. In the simplest case of 1RSB systems, the dominant states happen to be all equivalent, and one need only three parameters: the overlap between two different states, $q_{\alpha\beta} = q_0 \forall \alpha, \beta \neq \alpha$, the self-overlap of any of the states, $q_{\alpha\alpha} = q_1 \forall \alpha$, and the probability $1 - m$ ($0 \leq m \leq 1$) for two independent replicas to be in the same state (this in turn depends on the statistical weight of each individual state). The distribution of the overlap between two independent replicas, or, which is the same, of the overlap between states over the whole pairs of states, is then

$$P(q) = (1 - m)\delta(q - q_1) + m\delta(q - q_0). \quad (2.26)$$

In the high temperature paramagnetic phase, the system is ergodic and there is no need of replica symmetry breaking: one has simply $q_1^* = q_0^*$ and m is left undefined. In particular, all the dependence from q_1 and m (“in-states” quantities) disappears: $P(q) = \delta(q_0)$ and $\Phi_{1\text{RSB}}(q_1 = q_0) = \Phi_{\text{RS}}(q_0; \beta)$. In the whole range $T_s < T < T_d$, on the other hand, $m^* = 1$. This implies, again, both $P(q) = \delta(q_0)$ and $\Phi_{1\text{RSB}}(m = 1) = \Phi_{\text{RS}}(q_0; \beta)$: thermodynamics is completely blind to the presence of metastable states. The underlying physical reason is that, being exponentially many, the statistical weight of each single state is negligible. For $T < T_s$, finally, a non-trivial solution appears with $q_1 > q_0$ and $m < 1$. While the value q_1 appears discontinuously, since at T_s one has already $q_1 \neq q_0$, the behaviour of m and hence of $P(q)$ is continuous.

A thermodynamic description of the non-trivial phase for $T_s < T < T_d$ is still possible, but one has to be smarter. Following Monasson [84], we imagine a situation in which m real copies of the same system are coupled by a small attractive term of strength ϵ . By performing the limit $\epsilon \rightarrow 0$ after $N \rightarrow \infty$, we can select and study the original metastable states in a very elegant, clever way, by constraining all the copies to be in the same state, while keeping their mutual independence. The free energy for the system formed by the m copies then reads

$$e^{-\beta N \tilde{\Phi}(m, \beta)} = \sum_{\alpha} Z_{\alpha}^m = \sum_{\alpha} e^{-\beta N m f_{\alpha}} = \int_{f_{\min}}^{f_{\text{th}}} df e^{N[\Sigma(f, T) - \beta m f]} \sim e^{N[\Sigma(\tilde{f}, T) - \beta m \tilde{f}]}, \quad (2.27)$$

where the index α runs over all the possible states and $\tilde{f} = \tilde{f}(m, T)$ satisfies the saddle point equation

$$\left. \frac{\partial \Sigma(f, T)}{\partial f} \right|_{\tilde{f}(m, T)} = \frac{m}{T}. \quad (2.28)$$

The free energy $\tilde{\Phi}(m, \beta)$ can be computed via the replica method, thus obtaining $\tilde{\Phi}(m, \beta) =$

$m\Phi_{\text{1RSB}}(m, q_0^*(m, T), q_1^*(m, T); \beta)$. The crucial point is that now we can actually recognize the relevant physical content hidden in the m dependence of the replicated free energy Φ_{1RSB} . Promoting m to real values, equation (2.27) implies that $-\beta\tilde{\Phi}(m, \beta)$ is the Legendre transform of $\Sigma(f, T)$, which is concave. One can invert the relation to obtain Σ as a function of f . In practice, it is convenient to make an intermediate step, that is to compute $\Sigma(m, T)$ from the sole knowledge of $\tilde{\Phi}(m, \beta)$,

$$-\beta\tilde{f}(m, T) = \frac{\partial(-\beta\tilde{\Phi})}{\partial m}, \quad (2.29)$$

$$\Sigma(m, T) = \beta m\tilde{f}(m, T) - \beta\tilde{\Phi}(m, \beta) = -m\frac{\partial(-\beta\tilde{\Phi}(m, \beta))}{\partial m} - \beta\tilde{\Phi}(m, \beta). \quad (2.30)$$

At this point, a parametric plot of $\Sigma(m, T)$ versus $\tilde{f}(m, T)$ gives the complexity $\Sigma(f, T)$ as a function of the internal free energy f for each temperature. A positive complexity $\Sigma > 0$ is the thermodynamic signature of a non-trivial glassy phase, characterized by the emergence of an exponential in N number of well defined metastable states. States dominating thermodynamics for $T_s < T < T_d$ satisfy eq. (2.28) with $m = 1$, and are hence characterized by $f^* = \tilde{f}(m = 1, T)$.

Chapter 3

Finite connectivity mean-field: sparse graphs

This thesis is devoted to the study of a model defined on random sparse graphs. With the term *sparse*, one refers to the fact that the system is diluted, with respect to the fully connected limit, for what concerns the total number M of interactions in the system. We are interested in the limit $M = \alpha N$ for $N \rightarrow \infty$, with α kept constant. In general, interactions can involve an arbitrary number k of variables (for simplicity, we do not let here k fluctuate in the system). We consider the case in which the interactions are homogeneously distributed between variables, meaning that the degree of each variable is taken to follow the same distribution $P_{\text{degree}}(d)$, with average degree $c = k\alpha$. Notice that in general the local degree can fluctuate. However, as long as α is finite, then each variable has a finite number of first neighbours. This is a first crucial difference with respect to the more “naive” fully-connected mean-field.

Given a set V of N vertices (variable nodes) and a set E of M (hyper)-edges (interactions or function nodes), a sample is given by the (hyper)-graph¹ $G(V, E)$ and a particular realisation of additional quenched disorder, if present. Graphs are extracted according to a specific distribution. For instance, two popular choices are the Random Regular Graph (RRG) or the Erdős–Rényi ensemble (ER) [38]. In the first case, one has uniform probability over all the graphs having fixed degree c on each node, $P_{\text{degree}}(d) = \delta(d - c)$. This has the advantage of being more easily tractable from the analytical point of view, but forbid one to vary c (or α) continuously. To this end, one can introduce a family of graphs with heterogeneous local degree, the Erdős–Rényi ensemble, for which the degree follows a Poisson distribution of average $c = k\alpha$, $P_{\text{degree}}(d) = \frac{c^d}{d!} e^{-c}$, where d is integer but c is real. ER random graphs are usually defined by drawing for each interaction $a = 1, \dots, M$ a k -uplet of variables with uniform probability among the $\binom{N}{k}$ possibilities. Of course, these draws cannot be completely independent², if one wishes to avoid “repeated” interactions between the same k variables (this does not cure the eventuality of isolated variables, especially for low connectivities).

Sparse graphs share the crucial property of being *locally tree-like*, this meaning that there is absence of “small” loops in the network of interactions (as opposed to finite dimensional systems).

¹A graph structured in this way, separating variable from function nodes, is also called *bipartite* or *factor* graph. The use of factor graphs is very convenient when one wishes to study the situation $k > 2$. In that case, one usually speaks about hyper-graphs.

²What one has in mind with this definition is actually to extract with uniform probability one of the possible $\binom{N}{k}$ graphs obtained from selecting M edges among the $\binom{N}{k}$ possible ones.

On the other hand, for finite sizes $N < \infty$, long loops are always present. This can be understood in the following way [76]. The number of sites at a distance r from a reference variable scales³ in a tree as $\mathcal{N}(r) \sim [f(c)(k-1)]^r$, where r is the distance along the tree and $f(c) > 0$ is a generic function of the average degree $c = k\alpha$. This implies that the number of sites grows exponentially with the distance. At some point, the number of sites on the “boundary” (from site 0 perspective) becomes comparable with the total number of sites N (of course one has to stop before). If we imagine to construct a graph by iteratively adding edges at random between nodes, in this situation the probability that a new edge will involve (at least) two sites on the boundary becomes soon predominant: a loop is formed. One then argues that the typical size l of a loop scales with $N \rightarrow \infty$ as $\mathcal{N}(r=l) \sim N$, so that

$$l \sim \ln N. \quad (3.1)$$

The almost absence of small loops has very important technical consequences, the effect of long loops being often negligible since correlation typically decays on long paths. First of all, exact trees are factorizable, meaning that neighbouring variables become independent once the interaction between them is removed, the system itself resulting partitioned into independent sub-trees. This property is at the basis of very powerful, and substantially equivalent, iterative approaches that go under the names of *belief propagation* [76] and *cavity method* [78, 79]. In the next section the belief propagation formalism will be discussed, being a central tool in the development of this thesis.

The unavoidable presence in real random graphs of loops, even if very large, is however something we are worried about, since it may in principle spoil the effectiveness of these techniques. In particular, one has to be careful that correlations decay fast enough. The first obstacle one encounters is the onset T_d of a diverging point-to-set correlation length, which can be bypassed by solving the belief propagation or cavity equations at an appropriate replica symmetry breaking level (here we consider systems where the 1RSB picture is indeed correct). Notice that, however, loops have also an important physical implication, since they are both essential in order to introduce frustration into a system with open boundary conditions, and, even when the system is unfrustrated, they can be nonetheless a fundamental source of *self-induced* disorder⁴, independent from the eventual quenched one.

3.1 Belief Propagation

Belief propagation (BP) provides a smart way of computing (or approximating) marginals of the probability distribution of a statistical system, which in great generality can be written as

$$P(\underline{x}) = \frac{1}{Z} \prod_{i=1}^N \psi_i(x_i) \prod_{a=1}^M \psi_a(\underline{x}_{\partial a}), \quad (3.2)$$

³The reference variable at $r = 0$ is connected to an average number $c = k\alpha$ of function nodes, each of them being connected to variable 0 itself and $k-1$ other variables at distance $r = 1$. If the graph is RRG with fixed c , then each of the variables from any generation $r \geq 1$ is in turn connected to exactly $c-1$ other function nodes, each of them connecting other $k-1$ variable nodes, such that $f(c) \approx c-1$. In the case of ER graphs, one has simply $f(c) = c$, since the number d of “new generation” function nodes seen by any variable node i , which already participate in an edge, follows the distribution $P_{\text{edge}}(d)$ of average equal to c . For a definition of $P_{\text{edge}}(d)$ see eq. (3.16).

⁴Paradigmatic, in this regard, is the case of the Ising p -spin with ferromagnetic couplings (hence unfrustrated). Even on a RRG, which may appear not disordered on any non-diverging length scale, this model shows a glassy phase of the 1RSB kind [40].

where ∂a is the set of nodes $i \in V$ which participate in the a -th interaction, ψ_i represents the on-site coupling with an eventual external field, while ψ_a represents a k -body interaction term for the Hamiltonian, such that $|\partial a| = k$ is the function node degree (considered fixed for simplicity).

The most crude approximation one can resort to when dealing with systems with many interacting variables is to consider a factorized form for the probability distribution $P(\underline{x}) \approx \prod_{i=1}^N b_i(x_i)$, where the beliefs $b_i(x_i)$ are our proxies for the single-variable marginals of the system. Remarkably, this choice ceases to be an approximation if one considers the fully connected limit of the model, and for this reason it is often called (naive) mean-field approximation. A slight improvement would be to account at least for the most obvious correlations, *i.e.* the ones regarding variables directly entering some interaction. This is called the Bethe approximation:

$$P(\underline{x}) \approx P_{\text{Bethe}}(\underline{x}) = \prod_{i=1}^N b_i(x_i) \prod_{a=1}^M \frac{b_a(\underline{x}_{\partial a})}{\prod_{i \in \partial a} b_i(x_i)} = \prod_{a=1}^M b_a(\underline{x}_{\partial a}) \prod_{i=1}^N b_i(x_i)^{1-|\partial i|}, \quad (3.3)$$

where ∂i is the set of interactions $a \in E$ which include variable node i and $|\partial i|$ is their number, *i.e.* the degree of variable node i . The normalization is chosen in such a way that, if variables are uncorrelated, one recovers the factorization of the probability. Furthermore, it can be shown that P_{Bethe} is *exact* on trees, with b_i and b_a assuming the meaning of true marginals⁵. Unfortunately this property is not verified for arbitrary graphs; however, since the beliefs are intended to approximate the true marginals, we require them to satisfy normalization and local consistency equations

$$\sum_{x_i} b_i(x_i) = 1 \quad \forall i \in V, \quad (3.4)$$

$$\sum_{\underline{x}_{\partial a \setminus i}} b_a(\underline{x}_{\partial a}) = b_i(x_i) \quad \forall (a, i) : a \in E, i \in \partial a, \quad (3.5)$$

where we denote with $\partial a \setminus i$ the set of all the indices $j \in \partial a$ but i . These conditions are actually equivalent to three requirements: normalization of b_i 's, local consistency under marginalization over $|\partial a| - 1$ variables, and also normalization of b_a 's, which descends from the other two. A warning of caution is at this point required: even if the true marginals of any probability distribution satisfy these conditions, the converse it is not always verified. In particular, the fact of having found a set of locally consistence marginals does not imply that P_{Bethe} given in eq. 3.3 is actually a distribution possessing the b 's as marginals. This is very different from the simple case of the factorized probability, for which any normalized belief b_i represents also the ‘‘correct’’ marginal for the approximated distribution, that is always well normalized.

However, we will work with eq. 3.3, assuming the the underlying graph is actually a tree. We

⁵It can be proved by recursion [76]. We start from a leaf of the tree, that is from an interaction a that is connected to $|\partial a| - 1$ dangling variables of degree 1 (that we identify with the notation $\partial a \setminus i$), and only one variable i (since there are no loops in a tree) with degree strictly bigger than 1. Calling $P_T(\underline{x})$ the marginal probability distribution of all the variables in the tree *but* the dangling ones we have considered, we can write then $P(\underline{x}) = P_T(\underline{x}) b_a(\underline{x}_{\partial a \setminus i} | x_i) = P_T(\underline{x}) b_a(\underline{x}_{\partial a}) b_i(x_i)^{-1}$, where the first equality is a consequence of the fact that dangling variables interact with the rest of the tree only through variable i , while the second equality follows from Bayes' rule. At this point, we can adsorb the effect of the dangling ends as an auxiliary external field $\tilde{\psi}_i(x_i) = \sum_{\underline{x}_{\partial a \setminus i}} \psi_a(\underline{x}_{\partial a}) \prod_{k \in \partial a \setminus i} \psi_k(x_k)$ acting on variable i and repeat the procedure on $P_T(\underline{x})$, where the interaction a is removed together with the external fields (‘‘factor leaves’’) $\psi_k(x_k)$ for $k \in \partial a \setminus i$ and having redefined $\psi_i(x_i)$ as $\psi_i(x_i) \tilde{\psi}_i(x_i)$, until the tree is consumed. In the end, each interaction contributes with a factor b_a , while each variable node contributes with a factor b_i for every link it is connected to as long as it is non-dangling, hence the factor $b_i^{|\partial i| - 1}$.




Figure 3.1: Left: the lowest vertical arrow, from variable node j to function node a , is associated with the cavity marginal $\nu_{j \rightarrow a}(x_j)$, representing the marginal distribution of variable j in the absence of interaction a . BP local recursion eq. (3.7) allows one to express $\nu_{j \rightarrow a}(x_j)$ in terms of the product of the messages $\hat{\nu}_{b \rightarrow j}(x_j)$ incoming in j from all the function nodes $b \in \partial j \setminus a$ (upper arrows). Right: the top vertical arrow, from function node a to variable node j , is associated with the cavity marginal $\hat{\nu}_{a \rightarrow j}(x_j)$, representing the marginal distribution of variable j in the absence of all the interactions $b \in \partial j \setminus a$. According to BP eq. (3.8), the messages $\nu_{k \rightarrow a}(x_k)$ for $k \in \partial a \setminus j$ entering function node a (lower arrows) interact via $\psi_a(\underline{x}_{\partial a})$ to produce $\hat{\nu}_{a \rightarrow j}(x_j)$. Reprinted from [76].

introduce the free entropy functional F_{Bethe} (defined as $-\beta$ times the free energy, so that for $E = 0$ it gives back the entropy)

$$F_{\text{Bethe}}(\underline{b}) = -\beta E(\underline{b}) + S(\underline{b}) = \sum_a \sum_{\underline{x}_{\partial a}} b_a(\underline{x}_{\partial a}) \ln \psi_a(\underline{x}_{\partial a}) + \sum_i \sum_{x_i} b_i(x_i) \ln \psi_i(x_i) - \sum_a \sum_{\underline{x}_{\partial a}} b_a(\underline{x}_{\partial a}) \ln b_a(\underline{x}_{\partial a}) - \sum_i \sum_{x_i} (1 - |\partial i|) \ln b_i(x_i), \quad (3.6)$$

where $S(\underline{b}) = -\langle \ln P_{\text{Bethe}} \rangle_{P_{\text{Bethe}}}$. In analogy with naive mean field, one is tempted to associate to stationary points (in this case maxima, since F_{Bethe} is proportional to the free energy functional with the sign changed) of the Bethe free entropy, with some Lagrange multipliers imposing normalization and local consistency, the physical solutions for the b 's.

A more convenient way to carry out this task is by introducing the BP or cavity marginals $\nu_{j \rightarrow a}(x_j)$, $\hat{\nu}_{a \rightarrow j}(x_j)$ associated to each directed variable-to-function edge $j \rightarrow a$ or function-to-variable edge $a \rightarrow j$. These represent the marginal distributions of variable j respectively in the absence of interaction a (see left of fig. 3.1), or in the absence of all the interactions $b \in \partial j \setminus a$ (see fig. 3.1 on the right). They are also called BP or cavity *messages*, since represent the “message” which exits (arrives to) variable node j from an entire branch of the tree. In the absence of loops, different branches become independent once their connection to variable j is broken. This is used to build the following recursive scheme, also known as *Belief Propagation* equations

$$\nu_{j \rightarrow a}(x_j) = \frac{1}{Z_{j \rightarrow a}} \psi_j(x_j) \prod_{b \in \partial j \setminus a} \hat{\nu}_{b \rightarrow j}(x_j), \quad (3.7)$$

$$\hat{\nu}_{a \rightarrow j}(x_j) = \frac{1}{Z_{a \rightarrow j}} \sum_{\underline{x}_{\partial a \setminus j}} \psi_a(\underline{x}_{\partial a}) \prod_{k \in \partial a \setminus j} \nu_{k \rightarrow a}(x_k), \quad (3.8)$$

where $Z_{j \rightarrow a}$ and $Z_{a \rightarrow j}$ ensure the normalization of messages. They can be solved numerically on any graph by adding time indices and iterating the relations until a fixed point is (eventually) reached. On a real tree, one can start from the leaves and propagate the messages up to the root, hence reconstructing all the correct cavity marginals in a time that scales linearly with N . Of course, a simple relation connecting cavity messages to the true general marginals b_i, b_a can be obtained exactly on a tree

$$b_i(x_i) = \frac{1}{Z_i} \psi_i(x_i) \prod_{a \in \partial i} \hat{\nu}_{a \rightarrow i}(x_i), \quad (3.9)$$

$$b_a(\underline{x}_{\partial a}) = \frac{1}{Z_a} \psi_a(\underline{x}_{\partial a}) \prod_{i \in \partial a} \nu_{i \rightarrow a}(x_i). \quad (3.10)$$

What about if the underlying graph is not exactly a tree? In general, one can show that the stationary points of the Bethe free entropy (3.6) are in one-to-one correspondence with the fixed points of the belief propagation equations exactly through eqs. (3.9)-(3.10). Solving for a BP fixed point on an arbitrary graph is then equivalent to search for a complete set of Bethe beliefs. Of course, this procedure can be at the least rather dangerous, since we are not even assured that such beliefs correctly identify a proper probability distribution. We will anyway run the machinery, find a fixed point of BP and cross our fingers. The underlying reasonable idea is that, if correlations decay “fast enough”, one can hope BP or Bethe beliefs to be close to the correct marginals, when the underlying graph structure is locally tree-like.

3.1.1 Pairwise models

In the case of pairwise models the notation can be further simplified by noticing that function node a uniquely identifies a pair (ij) of site nodes. We can rewrite the BP equations in terms of only one species of outgoing messages $\nu_{j \rightarrow a}(x_j)$, which will be called hereafter as $\nu_{j \rightarrow i}(x_j)$

$$\nu_{j \rightarrow i}(x_j) = \frac{1}{Z_{j \rightarrow i}} \psi_j(x_j) \prod_{k \in \partial j \setminus i} \left(\sum_{x_k} \nu_{k \rightarrow j}(x_k) \psi_{kj}(x_k, x_j) \right), \quad (3.11)$$

having a simple diagrammatic interpretation if one imagines to combine the two subplots of figure 3.1 for $|\partial a| = 2$, or equivalently equations (3.7)-(3.8). The Bethe free entropy (3.6), rewritten in terms of the cavity messages, becomes in this case

$$F_{\text{Bethe}}(\underline{\nu}) = \sum_{i \in V} \ln Z_i - \sum_{(ij) \in E} \ln Z_{ij}, \quad (3.12)$$

$$Z_i = \sum_{x_i} \psi_i(x_i) \prod_{j \in \partial i} \left(\sum_{x_j} \psi_{ij}(x_i, x_j) \nu_{j \rightarrow i}(x_j) \right), \quad (3.13)$$

$$Z_{ij} = \sum_{x_i, x_j} \nu_{i \rightarrow j}(x_i) \psi_{ij}(x_i, x_j) \nu_{j \rightarrow i}(x_j). \quad (3.14)$$

In the rest of the thesis we will always work with pairwise notation.

3.1.2 Replica Symmetric Population Dynamics

The equations discussed so far only work for a given instance of the problem. In statistical physics, one is often interested in averages over the disorder, which in this case also comprehends the ensemble average over the graph distribution. Both the average on the graph ensemble and on any external randomness can be dealt with by promoting the BP equations to have a distributional meaning.

This essentially works as follows. In the Replica Symmetric (RS) approximation, one considers a randomly picked message to be a random quantity subjected to a well defined probability distribution $P(\nu)$, which is unique since just one state for the system is contemplated. Interpreting equation 3.11 as defining an update function $\Phi_{\text{update}}^J(\{\nu_k\}_d)$, taking as input d independent random messages and returning a new message, then $P(\nu)$ satisfies the self-consistency equation

$$P(\nu) = \mathbb{E}_J \mathbb{E}_d \int \prod_{k=1}^d d\nu_k P(\nu_k) \delta(\nu - \Phi_{\text{update}}^J(\{\nu_k\}_d)), \quad (3.15)$$

where \mathbb{E}_J represents the average over the quenched parameters of the Hamiltonian (possibly entering the definition of ψ_i and ψ_{ij} , and through them also Φ_{update}), while \mathbb{E}_d is the average over the edge-wise degree distribution, *i.e.* the probability (according to the generative graph ensemble) for a randomly extracted edge to insist on a site with exactly d other edges.

The probability $P(\nu)$ can be estimated numerically via the *population dynamics* algorithm. One takes a large but finite (say $U = 10^6$) collection of U messages $\{\nu\}_U$ to approximate $P(\nu)$. Starting from an initial condition $\{\nu\}_U^{t=0}$, one uses eq. (3.15) to obtain $\{\nu\}_U^{t+1/U}$ by replacing a random exemplar ν from the population $\{\nu\}_U^t$ with $\Phi_{\text{update}}^J(\{\nu_k\}_d)$, where the disorder J and the edge-wise degree d are extracted according to their independent probability distributions, and the random messages $\{\nu_k\}_d$ are chosen with uniform probability from $\{\nu\}_U^t$. The “fixed point” $\{\nu^*\}_U$ of this procedure, in a sense that we now specify, provides a numerical estimate of $P(\nu)$. The definition of a fixed point may appear indeed rather tricky. Of course, we are interested in comparing the evolution of the population between times t and $t + 1$, *i.e.* after $O(U)$ elementary steps. At this point, it is usually sufficient to study the behaviour with t of some physical observable, that is an average over the whole population $\{\nu\}_U^t$, and understand if it converges. This provides an operative definition for a fixed point of population dynamics.

It is finally convenient to make a last remark, which will be useful in the rest of the thesis. \mathbb{E}_d was defined as the average with respect of what we have called the edge-wise degree distribution, and that we will name $P_{\text{edge}}(d)$ from now on. The fact that $P_{\text{edge}}(d)$ is generally different from the degree distribution $P_{\text{deg}}(d)$ can be immediately recognized. For instance, for random regular graphs, where the degree is fixed to c , one trivially has $P_{\text{deg}}(d) = \delta(d - c)$, while $P_{\text{edge}}(d) = \delta(d - (c - 1))$. A very important exception is the case of Erdős–Rényi graphs, which will be studied throughout this thesis, for which one has, from the definition of edge-wise degree distribution

$$P_{\text{edge}}^{\text{ER}}(d) = \frac{(d+1)P_{\text{deg}}^{\text{ER}}(d+1)}{\sum_d (d+1)P_{\text{deg}}^{\text{ER}}(d+1)} = \frac{c^d e^{-c}}{d!} = P_{\text{deg}}^{\text{ER}}(d), \quad (3.16)$$

where we used the fact that $P_{\text{deg}}^{\text{ER}}(d) = \text{Poisson}(c)$. The extra $(d+1)$ factor in the definition of P_{edge}

(valid for any graph) follows from the fact that, even if a node with degree $(d+1)$ is present only with probability $P_{\text{deg}}(d+1)$, then it contributes $(d+1)/k$ times to the total number of interactions, $k = 2$ for pairwise models (an edge insisting on a high-degree site is in some sense over-represented, when picking edges at random). An important consequence is that single variable marginal probabilities b_i 's then become, in the population dynamics algorithm on ER graphs, random marginals with the same distribution of the cavity messages ν . This can be understood by noticing that the equations for the local marginals (3.9) are essentially identical to those for the messages ν eq. (3.7), apart from the fact that in the former equations *all* the edges insisting on each site are considered; the distribution of single site marginals is then formally equivalent to the right hand side of eq. (3.15), where ϕ_{update}^J has exactly the same definition in both cases, while the degree d is now extracted according to the site-perspective $P_{\text{deg}}(d)$; finally, since $P_{\text{deg}}^{\text{ER}}(d) = P_{\text{edge}}^{\text{ER}}(d)$, the distribution of single site marginals is found actually to satisfy the same equation (3.15). One can then use directly $\{\nu^*\}_U$ to compute averages of single site quantities.

3.1.3 Long range correlations: need for 1RSB

Below T_d , the emergence of properly defined global metastable states of infinite lifetime induces non-trivial correlations among the variables that may spoil the “naive” replica symmetric approach of equation (3.11). These correlations are of the *point-to-set* kind and involve a number of variables that diverges with the system size [88]. Notice that the physical intuition is in some respect very close to the one underlying the *Gedanken* experiment of [13], predicting the divergence of a static point-to-set correlation length at T_k in real liquids; in both cases, this divergence is associated to the very formation of global states of infinite lifetime, this happening at T_k (presumably) or at T_d , depending on the finite dimensional or mean-field nature of the model.

The decomposition into many states can be possibly dealt with by resorting to the 1RSB formalism [76]. In this approach, one identifies each different state α with a different fixed point $\underline{\nu}^\alpha$ of the belief propagation equations. Then, for each directed edge $i \rightarrow j$, the associated message $\nu_{i \rightarrow j}$ becomes a random variable obeying its own probability distribution $Q_{i \rightarrow j}(\nu)$ over the states. Since we want to weight states according to their internal free entropy, one has essentially that

$$Q_{j \rightarrow i}(\nu) = \sum_{\alpha} e^{m F_{\text{Bethe}}(\underline{\nu}^\alpha)} \delta(\nu - \nu_{j \rightarrow i}^\alpha). \quad (3.17)$$

A relation analogous to the BP equations can be derived for the Q 's, reading

$$Q_{j \rightarrow i}(\nu) \propto \int \prod_{k \in \partial j \setminus i} d\nu_{k \rightarrow j} Q_{k \rightarrow j}(\nu_{k \rightarrow j}) \delta\left(\nu - \Phi_{\text{update}}^{j \rightarrow i}(\{\nu_{k \rightarrow j}\}_{k \in \partial j \setminus i})\right) \left[Z_{j \rightarrow i}(\{\nu_{k \rightarrow j}\}_{k \in \partial j \setminus i})\right]^m, \quad (3.18)$$

where $\Phi_{\text{update}}^{j \rightarrow i}$ is given by eq. (3.11). Albeit this equations being formally very similar to the RS population dynamics one eq. (3.15), there are two important differences. The first point is that the distributions Q 's, apart from specific cases such as for random regular graphs (due to their homogeneity), are in general different for each edge; solving numerically the above equations would then involve a much bigger effort than the simple RS population dynamics⁶. Secondly, an explicit

⁶Equation (3.18) can be still averaged over the ensemble of random graphs and the external disorder. This leads to the definition of a probability distribution (over the random choice of an edge) of the probability distribution (over the states) of messages, $\mathcal{P}[Q(\nu)]$. Breaking one step of replica symmetry has added one distributional level to the

reweighting term $Z_{j \rightarrow i}^m$ has appeared in the right hand side of the equation.

The 1RSB equations can be drastically simplified if one limits him/herself to $m = 1$, as was firstly shown in [75] in the context of the tree reconstruction problem. For this reason, the 1RSB belief propagation equations at $m = 1$ that we are going to discuss are also often known as *reconstruction equations*. As a first step, let us consider the average messages

$$\nu_{i \rightarrow j}^{\text{av}}(x_i) \equiv \int d\nu Q_{i \rightarrow j}(\nu) \nu(x_i) = \sum_{\alpha} e^{F_{\text{Bethe}}(\nu^{\alpha})} \nu_{i \rightarrow j}^{\alpha}(x_i). \quad (3.19)$$

A remarkable result is that $\nu_{j \rightarrow i}^{\text{av}}$ are then found to satisfy the simple replica symmetric BP equations (3.11). This is actually not surprising, since the 1RSB formalism is based onto the equivalence among all the states, so that by averaging over all the solutions to the BP equations (or states) we still obtain a solution to the same BP equations.

A similar simplification is obtained at the level of the population dynamics equations implementing the average over the structure of the graph and the quenched disorder, if present. In this case, one introduces the probability distribution $\mathcal{P}[Q(\nu)]$, over the random choice of an edge, of $Q(\nu)$. It is also convenient to define the average distribution of messages $\bar{P}(\nu) \equiv \int dQ(\nu) \mathcal{P}[Q(\nu)] Q(\nu)$. It turns out that, under a ‘‘local uniformity condition’’ (essentially, that the RS solution is paramagnetic, $\nu_{i \rightarrow j}^{\text{av}} = 1/q$, where $q = |\mathcal{X}|$ is the size of the alphabet), the equations really simplify if one considers a new set of q distributions over the messages

$$R_x(\nu) = q \nu(x) \bar{P}(\nu), \quad \forall x \in \mathcal{X}. \quad (3.20)$$

$R_x(\nu)$ can be interpreted as the distribution of the messages $\nu(y)$ that are ‘‘biased’’ towards $y = x$. This choice in the end allows one to eliminate the reweighting term inside the population dynamics recursion for $R_x(\nu)$ induced by (3.18), obtaining (see Appendix A for details on the derivation)

$$R_x(\nu) = \mathbb{E}_J \mathbb{E}_d \int \prod_{k=1}^d \left[d\nu_k \sum_y \pi(x, y; J) R_y(\nu_k) \right] \delta(\nu - \Phi_{\text{update}}^J(\{\nu_k\}_d)), \quad (3.21)$$

where

$$\pi(x, y; J) = \frac{\psi_J(x, y)}{\sum_y \psi_J(x, y)}, \quad (3.22)$$

and $\psi_J(x, y)$ is the (pairwise) interaction term between neighbours on the graph. At this point, we have obtained two remarkable results: the density evolution BP equations do not involve distributions of distributions, and we have eliminated the reweighting term. A third crucial theorem can be proved [75], affirming that eq. (3.21) admits a non-trivial solution if and only if it converges to a non-trivial distribution when initialized according to

$$R_x^{t=0}(\nu) = \delta(\nu - \nu_x^{\text{plant}}), \quad \nu_x^{\text{plant}}(y) = \delta_{xy}. \quad (3.23)$$

Finally, if translational invariance holds, *e.g.* if $\psi_J(x, y) = \psi_J(|y - x|)$ and there is no external field, then $R_x(\nu) = R_0(T(-x)\nu)$, where $T(-x)$ is the operator that shifts the argument of the function $\nu(y)$ by $-x$, and it is sufficient to study the following recursion in terms of the sole $R \equiv R_0$

problem.

$$R(\nu) = \mathbb{E}_J \mathbb{E}_d \int \prod_{k=1}^d \left[d\nu_k \sum_y \pi(y; J) R(T(-y)\nu_k) \right] \delta(\nu - \Phi_{\text{update}}^J(\{\nu_k\}_d)), \quad (3.24)$$

where $\pi(y; J) \equiv \pi(0, y; J) \propto \psi_J(|y - 0|)$. In the end, for $m = 1$ we have brought back the technical complexity of 1RSB equations to the same level of RS ones. The typical internal free entropy of $m = 1$ states can be computed from the fixed point of (3.24) as

$$N^{-1} F_{\text{Bethe}} = \mathbb{E}_J \mathbb{E}_d \int \prod_{k=1}^d d\nu_k \sum_y \pi(y; J) R(T(-y)\nu_k) \ln Z(\{\nu_k\}_d; J) + \quad (3.25)$$

$$-\alpha \mathbb{E}_J \mathbb{E}_d \int d\nu_1 d\nu_2 R(\nu_1) \sum_y \pi(y; J) R(T(-y)\nu_2) \ln Z_2(\nu_1, \nu_2; J), \quad (3.26)$$

where Z and Z_2 are essentially given by equations (3.13)-(3.14),

$$Z(\{\nu_k\}_d; J) = \sum_x \bar{\psi}_J(x) \prod_{k=1}^d \sum_y \psi_J(|y - x|) \nu_k(y), \quad (3.27)$$

$$Z_2(\nu_1, \nu_2; J) = \sum_{x,y} \nu_1(x) \psi_J(|y - x|) \nu_2(y), \quad (3.28)$$

having defined as $\bar{\psi}_J(x)$ the equivalent of the single-site term $\psi_i(x_i)$, in order to differentiate the notation with respect to ψ_J , which here stands for the pairwise interaction also called ψ_a before. The reconstruction equations will be used to study the *typical* states of a system in its liquid and glassy phase above T_k , where the solution $m = 1$ is thermodynamically consistent.

3.2 Constraint satisfaction problems

Constraint satisfaction problems (CSP) can be defined as collections of M conditions or formulas (constraints) among N variables. Typical questions one may be interested in are for example: establishing whether a particular instance has a solution or not (decision problem), counting the number of its solutions (counting problem), minimizing the number of possibly violated constraints and eventually finding a configuration of the variables satisfying this condition (optimization). Computational complexity theory is interested in the worst-case scenario [41]. A problem is called *tractable* with respect to one of these questions if it is solvable by a deterministic⁷ algorithm in a time scaling polynomially in the system size for *all* the possible instances. *Intractable* problems are on the contrary only solvable in nonpolynomial running time by any known deterministic algorithm, at least for some choice of particularly “hard” instances (for a physicist’s introduction to computational complexity theory and some more refined classification see [74]). Worst-case analysis is not the most generally informative perspective one can adopt. In particular, nonpolynomial behaviour may be sometimes triggered by some pathological instances only, while a problem still maintains to be tractable under practical conditions.

For this kind of reasons, besides the worst-case scenario, the attention has recently moved also to

⁷A standard classical algorithm, as opposed to a nondeterministic algorithm (this is limited to a theoretical concept in the world of conventional, non-quantum computers), which via the branching of the computational flow can perform an exponential number of polynomial time computations in parallel.

the study of *typical* behaviour, as defined by considering the average properties of random instances drawn from a given ensemble. One is then interested to understand how the scaling of algorithmic running time depends on the choice of the parameters of the model, such as the constraints per variable ratio $\alpha = M/N$, and how it relates to the location of actual phase transitions in the standard statistical physics approach to these problems. Among the many CSPs that have been investigated in this direction, and particularly appreciated also from the physics community, we have the q -coloring (q -COL) [89, 99, 101], random k -satisfiability (k -SAT) [83, 85, 86], and random k -XORSAT [28, 65, 93]. A great deal of interest in these problems stems from the fact that they can be viewed as equivalent to the study of the groundstates of sparse glassy Hamiltonians. For instance, a straightforward mapping is possible between q -coloring and the sparse antiferromagnetic q -states Potts model, or between random k -XORSAT and the sparse Ising k -spin model. Notice that most of the CSPs usually studied in the literature deals with discrete variables.

This thesis is devoted to the study of an extension of the q -coloring to continuous variables from the statistical physics' typical-case perspective. For this reason, it seems appropriate to spend just a few words about this CSP, among all the others. The original discrete problem is readily defined: given a graph with pairwise interactions and a set of q colors, it consists in finding a proper coloring of the vertices such that no pair of adjacent variables has the same color. Although it may seem rather an extravagant question to dally with, it turns out that many real-world combinatorial optimization problems can be easily mapped onto the q -coloring. For instance, a practical application concerns the optimal scheduling of registers in the central processing unit of computers (see the introduction of [17] and references therein). The problem in this case can be stated as follows: one wants to physically assign variables, changing their values in a same time interval during the execution of a program, to different registers. This can be schematized by assigning each variable to a vertex and placing edges between variables that are updated during the same interval. Since we do not want variables connected by an edge to belong to the same register (*i.e.* to take the same color), we can recast the problem of scheduling the registers as that of finding a proper q -coloring to the generated graph, where q corresponds also to the number of employed registers (possibly minimum).

3.2.1 Phase transitions in random CSPs

In this context, physics methods and intuition can be very effective. The current picture emerged is that the slowing down of searching algorithms in certain regions of the phase diagram is deeply connected to the presence of actual phase transitions. The typical scenario is given in figure 3.2. Increasing the average degree c of variable nodes in the graph (or equivalently what we call the average connectivity $\alpha = M/N = c/k$, where $k = 2$ for pairwise interactions), the system undergoes several transitions, that can be usually interpreted through the 1RSB scheme. In the following, we will make an extensive use of the concept of *clusters* of solutions. Roughly speaking, clusters are akin to connected domains in the space of solutions, meaning that any two solutions in a cluster can be connected by a path of other solutions belonging to the same cluster, and differing only for a non-extensive number of variables. This provides a useful specialization to zero temperature of the concept of thermodynamic states.

At α_d , the *clustering* or *dynamic* transition, phase space fragments into an exponential in N number of clusters, each of them having a nonzero internal entropy (d 1RSB). At α_c , the *condensation* transition, equivalent to the Kauzmann transition (s 1RSB), the number of clusters dominating




Figure 3.2: Pictorial representation of the space of solutions in random constraint satisfaction problems, when varying the average vertex degree c . For $c_d < c < c_c$ solutions are split into an exponential in N number $e^{N\Sigma^*}$ of clusters of typical internal size e^{Ns^*} . Some clusters may appear also before α_d , but they only comprehend an exponentially small fraction of the solutions. For $c_c < c < c_s$ the measure condenses over the few largest clusters, until no solution exists beyond the SAT/UNSAT threshold c_s . At the rigidity point c_r , which might come before or after the condensation threshold c_c as well, dominating clusters start to contain an extensive number of frozen variables, *i.e.* variables taking the same value in all the solutions belonging to the given cluster (frozen clusters are colored in black). Another interesting threshold, not depicted in the figure, is the freezing point c_f , where all the clusters, not only the dominating ones, become frozen. Reprinted from [101].

the measure becomes sub-exponential in N (their complexity vanishes). Above the satisfiability threshold, generally known as SAT/UNSAT transition (also COL/UNCOL in the case of coloring as in figure 3.2), typical instances cannot be satisfied by any choice of the variables with probability going to one as $N \rightarrow \infty$. Finally, another important transition occurs at the *rigidity* threshold α_r , where typical clusters (*i.e.* the ones dominating the measure) develop a finite fraction of *frozen* variables. A variable is said to be frozen if it takes the same value on all the configurations belonging to a given cluster. One can also be interested in the *freezing* threshold $\alpha_f > \alpha_r$ for which *all* the clusters, even the atypical ones, are frozen (*i.e.* they contain an extensive number of frozen variables). If one “flips” a frozen variable, the system is forced to move to another cluster, thus implying that an extensive number of variables has to rearrange [98]. The presence of frozen variables may play a major role in determining the behaviour of some classes of local algorithms. Indeed, one expects to be impossible to construct in a polynomial time a solution belonging to a frozen cluster, for which an extensive number of highly correlated variables has to be set in a consistent way. For this reason, α_f (which obeys a rather difficult large deviation computation, contrarily to the typical-case α_r) is usually proposed as an upper bound to the algorithmic threshold for finding a solution [15, 59].

The statistical physics approach to random CSPs consists in studying the associated Hamiltonian model, where the energy is taken to be the number of violated constraints, in the limit $\beta \rightarrow \infty$. With this choice of Hamiltonian, solutions to the CSP have zero energy, and are the only relevant configurations for $\beta \rightarrow \infty$. Furthermore, all the solutions are weighted *uniformly*. When the space of solutions breaks, though, not all the clusters may contribute in the same way to the total partition function, and, as we said, we can distinguish between typical clusters (which dominate thermodynamics, since they contain all together a number of solutions that is exponentially in N bigger than the number of the remaining solutions), and atypical clusters. Since we are at zero temperature and zero energy, we can work with entropies instead of free energies. This is very handy: the statistical weight of clusters of solutions is simply given by the balance between their size (internal entropy) and their number (complexity).

This can be formalised via the *entropic zero temperature limit* approach [77]. In order to study




Figure 3.3: Complexity function for the 6-coloring of l -RRGs. Clusters are divided in classes and counted according to their internal entropy density s ; the number of clusters in each class, to the leading exponential order, is simply given by $e^{N\Sigma(s)}$. Empty circles indicate typical states (dominating as a whole the partition function) of entropy s_* . For $l = 18$ they satisfy $\Sigma'(s_*) = -m^* = -1$, while $\Sigma'(s_*) = -0.92$ and $\Sigma(s_*) = 0$ for $l = 19$. The dynamic transition is hence $l_d(6) = 18$, the condensation one is $l_c(6) = 19$ (the degree l cannot be varied continuously on RRGs). Reprinted from [61].

zero energy configurations, we fix $E = 0$, so that $-\beta f \equiv s$ (internal entropy density). By defining $-\beta\tilde{\Phi}(m; \beta)|_{\beta \rightarrow \infty} \equiv \Phi_s(m)$, equations (2.29)-(2.30) become

$$\Phi_s(m) = \Sigma(m) + ms, \quad (3.29)$$

$$s(m) = \frac{\partial \Phi_s(m)}{\partial m}, \quad (3.30)$$

$$\left. \frac{\partial \Sigma(s)}{\partial s} \right|_{s(m)} = -m, \quad (3.31)$$

where the complexity $\Sigma(s)$ represents, to the leading exponential order in N , the number $e^{N\Sigma(s)}$ of clusters of internal entropy s . The free entropy $\Phi_s(m)$ can be computed on sparse graphs by solving the 1RSB belief propagation or cavity equations for generic m . This provides a tool for studying the presence of both typical and atypical clusters, constrained to their internal size s , by computing the whole complexity curve $\Sigma(s)$ as the Legendre transform of $\Phi_s(m)$. An example is given in figure 3.3.

Part II

Continuous coloring problem

Preamble

In this second part we study⁸ in detail a constraint satisfaction problem with continuous variables defined on a sparse random graph, and that is called *continuous coloring*. The problem is essentially the following: given a sparse random graph $G = (V, E)$, which can be *e.g.* Erdős–Rényi or random regular, the aim is to assign a “color” $x_i \in [0, 2\pi)$ to each vertex in V , such that every pairs of connected vertices $(ij) \in E$ are different enough, that is they satisfy $\cos(x_i - x_j) \leq \cos \theta$ for some threshold value θ .

This model can be regarded as a continuous version of the well-known problem of discrete coloring of random graphs [101]. Moreover, it also corresponds to the 1-dimensional version of the Mari-Kurchan-Krzakala (MKK) model [70], a mean-field approximation to models of hard spheres showing jamming. In the MKK model, d -dimensional hard spheres living in a d -dimensional box with periodic boundary conditions interact with only a finite number of other particles, according to an underlying sparse graph network. From this point of view, continuous coloring can thus be interpreted also as an “angle-packing” problem, with the obvious identification of θ with the diameter of the particles [59].

The MKK model for $d = 2$ was numerically studied in [70], exhibiting the presence of a RFOT when increasing the diameter of the spheres (or equivalently the packing fraction), for sufficiently high connectivities. A generalized version of the MKK model, accounting also for p -body interactions, was studied in [80] for different values of p and d through the cavity (BP) formalism. The case $d = 1$ and $p = 2$, in which we are interested, is shown to undergo a RFOT for large connectivities. However, the dynamical or clustering threshold was not computed.

Let us start explaining why we looked for this model and why we think it is very useful, worth studying in detail. The continuous coloring possesses *all* the following features:

- it is a constraint satisfaction problem (CSP),
- having continuous variables,
- defined on a sparse random graph,
- showing a random first order transition (RFOT).

The most famous and well-studied model showing a RFOT is the spherical p -spin model: in this case one has N real (unbounded) variables x_i , subject only to a global constraint $\sum_i x_i^2 = N$. Unfortunately, this model is well defined only on very dense graphs, because as soon as one makes the interactions slightly sparser the model ground state condensates on a small subset of the variables, becoming meaningless from the physical point of view. In order to have a very sparse model,

⁸I acknowledge the valuable contribution of Thibault Lesieur to some aspects of this research.

one needs to avoid the above condensation phenomenon, and this can be achieved either adding a Lagrange multiplier to each real variable, or more simply by using variables defined on a bounded domain. Among the latter models we have all CSPs with discrete variables (*e.g.* q -col, k -SAT, k -XORSAT). Willing to use continuous variables with bounded domain, the simplest way is to choose models with vector spins, *e.g.* XY or Heisenberg spins. The continuous coloring problem we study can indeed be seen as an XY model where variables are unit norm vectors of 2 components, $\vec{s}_i = (\cos(x_i), \sin(x_i))$, the constraints that we enforce being written as $\vec{s}_i \cdot \vec{s}_j \leq \cos(\theta)$. A recent work addresses the glass transition of this kind of rotational degrees of freedom, but focusing on the opposite limit of a large number of vector components [100].

The results of [80] are particularly encouraging, since it is not obvious at all that taking a discrete CSP, like q -coloring, and transforming it to a version with continuous variables, the physics, in this case the RFOT, is preserved. The change in the symmetry of variables, from \mathbb{Z}_q to $O(2)$, is drastic and may change the nature of the phase transition. In the successful case we do find a model with all the above features (and this is what we are going to confirm in Chapter 4), we have in our hand a model which is very useful in many aspects. Let us list those aspects we find more interesting:

- being defined on a locally tree-like graph, the model can be solved analytically via the cavity method, and the location of phase transitions can be computed with high accuracy;
- having bounded variables interacting in a sparse way, the numerical simulation of the model can be made very efficiently, and thus a meaningful comparison between numerical results and analytical solution can be done;
- being a CSP with continuous variables, and where the density of constraints can be varied continuously, the model would show a jamming transition, for which it can be considered as a very simple sparse mean-field model for jamming (please notice that dense mean-field models for jamming, as the perceptron, have problems⁹ when generalized to the sparse case);
- having a RFOT, we expect the model to have a dynamical phase transition and, below the dynamic transition temperature, the energy relaxation to get stuck at some threshold energy connected in some way to the topological properties of the energy landscape (*e.g.* to the spectrum of the energy Hessian at the stationary points) that can be computed, for sufficiently smooth interaction potentials, thanks to the continuous nature of the variables;
- being defined on a sparse random graph, one can change the mean degree, thus testing very different regimes from the dense one to the very sparse one, thus better understanding how much of the classical (dense) mean-field physical behaviour is preserved in the sparse regime (this is particularly relevant to understand why it is so difficult to find good glass models in low dimensions where the number of nearest neighbours becomes small).

To the best of our knowledge, the above aspects have not been investigated in detail in a *single* model before.

⁹The delicate point is once more given by the fact that one generally considers continuous models with unbounded variables. In the case of the perceptron, if a variable is subjected to K constraints, the probability for all the random obstacles' components relative to that variable to have the same sign, *e.g.* to be all positive, is 2^{-K} . In this case that variable can satisfy all of its constraints by taking arbitrarily large values (irrespective of other variables) and the model is ill defined. While in the perceptron $K = O(N)$ and this event does never occur, in the diluted case K is finite and the model is ill defined with a finite probability.

Chapter 6 will focus on a very interesting aspect, namely the possibility of optimizing the dynamical threshold by reweighting properly the solutions of the problem. It is well known that in complex CSP, like the one we are studying, solutions may have very different features. For example, in discrete random CSP, solutions organize in clusters of very different sizes and of different nature (*e.g.* with or without frozen variables) [61, 86, 101]. Solutions that dominate the thermodynamics may be very different from the ones which are found via the best solving algorithms [59], and this makes difficult the connection between thermodynamic description and the behaviour of dynamical processes searching for solutions. Reweighting the solutions, that is giving to solutions a non-uniform statistical weight, is a simple way to count the atypical solutions that would not weigh enough in the uniform measure. We will perform an optimization of the reweighting potential, with the aim of postponing as more as possible the dynamical phase transition. This has been done recently in Refs. [18, 19] for the discrete CSP of bicoloring random hypergraphs, and was done in Refs. [69, 97] for models of hard spheres in infinite dimensions. The idea behind this optimization is that, working with the reweighted potential, the long range correlations leading to the ergodicity breaking at the dynamical phase transition will appear later, and algorithms should find solutions in an easier way in these “biased landscapes”. An innovative aspect of the present work is that the optimization of the potential is performed in a semi-automatic way by modifying the interaction potential which is a function in $[0, 2\pi)$, so formally with an infinite number of parameters (in practice, we discretize it with a very large number of points).

Thanks to the sparse nature of the model, we are going to support with extensive numerical simulations most of the analytical results. In particular, this will allow us to make in Chapter 7 the following interesting observation: a class of Monte Carlo based local search algorithms converges to solutions whose features are similar to the ones of typical solutions for the optimized potential. This observation leads us to conjecture that the ultimate algorithmic threshold for this class should be represented by the most optimized dynamical threshold.

Chapter 4

The model

4.1 Definitions

4.1.1 CSP definition

In the continuous coloring, N real angular variables $x_i \in [0, 2\pi)$, $i = 1, \dots, N$ interact via a sparse graph $G = (V, E)$, composed by a set V of variable nodes (vertices), with size $|V| = N$, and a set E of M pairwise interactions between variables, $|E| = M$. In this work, graphs are random instances drawn from the Erdos-Renyi ensemble, with average connectivity $\alpha = M/N$ (the average degree of variable nodes is hence equal to 2α) playing the role of an external tuning parameter. Neighbours on the graph have to satisfy the excluded-volume constraint $\cos(x_i - x_j) \leq \cos \theta$, for some excluded angle (or “diameter”) $\theta \in [0, \pi)$.

The statistical physicist’s favorite way of dealing with a constraint satisfaction problem, is usually by defining a Hamiltonian which exactly counts the number of violated constraints, that in the case of continuous coloring reads

$$\mathcal{H}_{\text{flat}}(\{x\}) = \sum_{(ij) \in E} \mathbb{I}(\cos(x_i - x_j - \omega_{ij}) > \cos \theta), \quad (4.1)$$

where $\mathbb{I}(\dots)$ is equal to 1 if the argument is satisfied, and 0 otherwise. The subscript “flat” refers to the fact that $\mathcal{H}_{\text{flat}} = 0$ on all the solutions to the problem. In the limit $\beta \rightarrow \infty$, the associated Boltzmann-Gibbs measure $Z^{-1}e^{-\beta\mathcal{H}_{\text{flat}}}$ is hence uniform over the space of solutions, while configurations not satisfying all the constraints have zero statistical weight.

The parameters $\omega_{ij} \in [0, 2\pi)$ are quenched random shifts living on each directed edge $i \rightarrow j$, with $\omega_{ji} = 2\pi - \omega_{ij}$. The original problem, which is recovered for $\omega_{ij} = 0$, is found to undergo a “crystallization” phenomenon when increasing connectivity (or packing fraction, by increasing θ), or equivalently when lowering temperature, as already noted in [80]. In the crystal phase, variables condense around a discrete set of values in the interval $[0, 2\pi)$, mimicking a q' -coloring packing, with $2\pi/q' > \theta$. On a real tree, with open boundary conditions, the homogeneous $\omega_{ij} = 0$ and non-homogeneous model are equivalent, since the latter can be mapped onto the former by the following transformation: $\forall (ij) \in E$, $\omega_{ij} \rightarrow 0$, $x_j \rightarrow x_j + \omega_{ij}$ and $x_k \rightarrow x_k + \omega_{ij}$ for all the $\{x_k\}$ belonging to the subtree starting from j (that is excluding the whole branch containing i). However, in the case of random graphs, where loops are always present (even if very large), or of a finite tree with fixed boundary conditions on the leaves, this construction is not guaranteed to hold for every




Figure 4.1: Discretization of precision $d = 3$ for the continuous coloring. The excluded angle (on each side of the particle) is equal to $\theta = 2\pi/q$ in the continuous limit, with $q = 8$. As a consequence of discretization, $2d - 1$ states are forbidden (red points), among the total $p = qd = 24$ possible ones. Big dots, together with the dashed lines, are meant only to identify the multiples of θ along the circumference. The increase of precision d with q fixed corresponds to the proliferation of small dots inside each sector.

choice of the ω_{ij} 's: loops, together with random shifts, induce frustration of the periodic order.

In order to suppress the crystal phase, [70] adopted a small degree of polydispersity. Here we will follow two strategies. In the analytical derivation of the 1RSB belief propagation equations for the homogeneous case $\omega_{ij} = 0$, we will impose translational (rotational) invariance: this allows to simplify the equations when the replica parameter $m = 1$, obtaining a RS-like scheme at the price of introducing *planted* random shifts in the messages (see the next section for a discussion about planting). Imposing translational invariance, one can disregard the ordered solution while maintaining the tree-like recursive structure of the BP equations. On the other hand, we will also support the BP predictions with direct numerical Monte Carlo simulations of the non-homogeneous model with uniform random shifts.

The parameter $\theta \in [0, \pi)$ represents half the excluded angle around each “particle”, playing the same role of the radius in hard sphere systems. In order to make direct contact with the very well known q -coloring, it is useful to consider “discretized” values of θ defined as $\theta = \frac{2\pi}{q}$, with $q \in \mathbb{N}_{>1}$. The meaning of this relation is straightforward: in the q -coloring, colors $u \in \{1, \dots, q\}$ can be associated to q -states Potts angular variables $x = \frac{2\pi}{q}(u - 1)$, naturally satisfying $\cos(x_i - x_j) \leq \cos(2\pi/q)$ if $x_i \neq x_j$. In the following, we will express the diameter θ in terms of integer q 's. This mapping allows for the following observation: given a solution to the discrete coloring with q colors, it also represents a solution to the continuous coloring with excluded angle $\theta = \frac{2\pi}{q}$ (actually, the solution is valid for any $\theta' \leq \frac{2\pi}{q}$). Conversely, the set of solutions to the continuous coloring for $\theta = \frac{2\pi}{q}$ contains all the solutions to the q' -coloring with $q' \leq q$.

Discretization

In order to numerically solve the BP equations, a discretization of the interval $[0, 2\pi)$ is needed. To this end, we introduce a clock approximation with p states by imposing

$$x_i = \frac{2\pi}{p}t_i, \quad t_i \in \{0, \dots, p-1\}, \quad \forall i \in V. \quad (4.2)$$

Calling d the number of clock states inside the excluded region $0 \leq x < \theta$, *i.e.* the discretization precision, then it follows $p = qd$ (as we said, we restrict to integer q for convenience). Each variable precludes $2d - 1$ values to its neighbours, as depicted in figure 4.1.

4.1.2 Planted model

In the planting procedure, one first draws a planted configuration $\{x_i^0\}$ according to some prior probability distribution, which is assumed to be factorized, $P_X(\{x_i\}) = \prod_i P_X(x_i)$. Then, a graph can be constructed in two ways, depending on whether random shifts are explicitly considered or not. In the first case, one can run through¹ all the pairs of vertices (ij) and add an edge with probability $2\alpha f(x_i^0, x_j^0; \beta)/(N - 1)$, where $f(x_i, x_j; \beta) = e^{-\beta\mathcal{H}(x_i, x_j)} / \int dx dy P_X(x) P_X(y) e^{-\beta\mathcal{H}(x, y)}$.

This ensures that the average degree c of variable nodes is equal to 2α

$$c = (N - 1) \int dx dy P_X(x) P_X(y) \frac{2\alpha f(x, y)}{(N - 1)} = 2\alpha. \quad (4.3)$$

However, we would also like different nodes to follow the same degree distribution. To this end, one can derive the average degree $c(x)$ conditional on the knowledge that $x_i^0 = x$,

$$c(x) = (N - 1) \int dy P_X(y) \frac{2\alpha f(x, y)}{(N - 1)} = 2\alpha \int dy P_X(y) f(x, y). \quad (4.4)$$

Then, we must also ask for $\int dy P_X(y) f(x, y) = 1$, so that $c(x) = 2\alpha$ independently from x .

In the presence of random shifts, we can plant the system in a slightly different way: given the configuration $\{x_i^0\}$ and a graph realisation extracted from the proper distribution, we associate a random variable $\omega_{ij} \in [0, 2\pi)$ to each directed edge $i \rightarrow j$, and $\omega_{ji} = 2\pi - \omega_{ij}$ on $j \rightarrow i$, with probability $p_{ij}(\omega_{ij}) = e^{-\beta\mathcal{H}(x_i - x_j - \omega_{ij})} / \int d\omega e^{-\beta\mathcal{H}(x_i - x_j - \omega)}$. In the following, we will use this second method in order to plant the system when running Monte Carlo numerical simulations (we indeed require random shifts to suppress crystallization). Notice that one can choose, without loss of generality, $\{x_i^0\} = 0$, this being equivalent to a set of local gauge transformations (valid for any i): $x_i \rightarrow x_i - x_i^0$, $\omega_{ij} \rightarrow \omega_{ij} - x_i^0 \quad \forall j \in \partial i$.

What is the point in this construction? In general, both the planted models are *different* from the original ones (with or without the random shifts), since planting a solution changes the graph ensemble, and, generally, its properties. This is particularly evident in the $\beta \rightarrow \infty$ limit: via planting, we can always construct for arbitrary α , even beyond the SAT/UNSAT threshold, a pair graph-configuration of exactly zero energy. However, before the condensation transition α_c , the planted models are equivalent² to the original ones: this is known as *quiet planting* [63]. We can understand this fact by considering that the superimposition of a single planted cluster in a region $\alpha_d < \alpha < \alpha_c$ which is already dominated by exponentially many other clusters, should be thermodynamically undetectable, see figure 4.2. The crucial point in the quiet planting technique is that the planted cluster exhibits the properties of typical $m = 1$ clusters from the random ensemble. This provides a very powerful tool to study the average properties of the whole clustered

¹For $f = 1$, this can be regarded as an alternative definition for the Erdős–Rényi ensemble, meaning that in the limit $N \rightarrow \infty$ most of the relevant properties do coincide (see [76] for details).

²Provided that the original random model displays a *uniform* BP fixed point: *e.g.* $\nu(x) = 1/q$, $x = 1, \dots, q$ for the q -coloring, or $\nu(x) = 1/(2\pi)$, $x \in [0, 2\pi)$ for the continuous coloring. We call it the paramagnetic (liquid) fixed point.

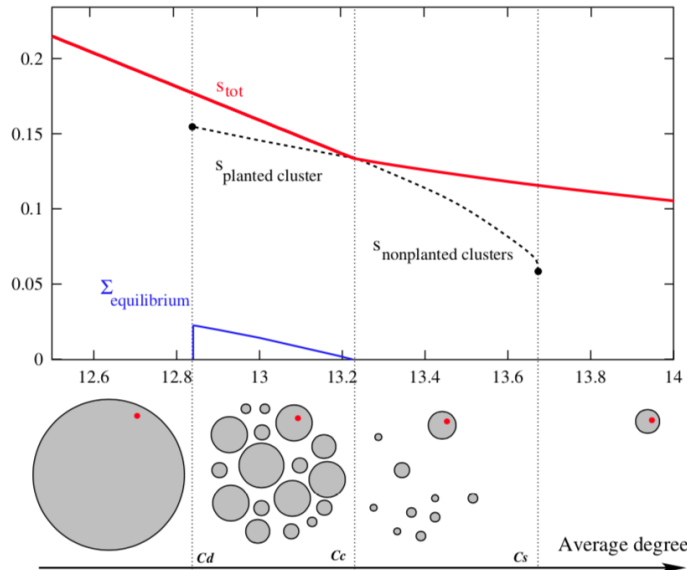


Figure 4.2: Phase diagram of the 5-coloring in the planted ensemble. Bottom: clustering phenomena. Red dots represent the planted configuration. Between the dynamic and condensation transition, the ergodic phase breaks into an exponential number of clusters which share the same properties of the $m = 1$ clusters from the random ensemble. The planted cluster is (just one) among them, hence it does not possess enough statistical weight to alter the partition function of the model, which is still globally RS or liquid-like. Above c_c , the planted cluster contains the majority of solutions and becomes dominant: the planted and random model become (in principle) distinguishable. Top: the total entropy of solutions s_{tot} (red line) is insensitive to the presence of the additional planted cluster up to c_c . The equilibrium complexity (blue line) is the same for the random and the planted ensemble, $\Sigma = s_{\text{tot}} - s^*$, where $s^* = s_{\text{planted cluster}}$ is the internal entropy of typical $m = 1$ clusters. Reprinted from [63]

phase $\alpha_d < \alpha < \alpha_c$ by focusing on a single, easy to build planted cluster.

4.1.3 Community detection definition

In the planting procedure (without random shifts), one adds an edge (ij) to the graph according to some probability³ $2\alpha f(x_i^0, x_j^0; \beta)/N$, which depends on the values taken by i and j in the planted configuration $\{x_i^0\}$. This generative model for random graphs is also known as the stochastic block model [37]. We may interpret the different values each variable can take as signaling the membership to a specific “community”, and the stochastic block model (planting) as a rule to establish connections between these communities. A natural question is to understand under what circumstances one is able to recover some knowledge on $\{x_i^0\}$ (original community structure) from the observation of the generated graph A . In the simplest (Bayes optimal) case, one also knows the parameters of the model, namely the prior P_X from which $\{x_i^0\}$ was extracted, and the function $f(x, y; \beta)$.

Following [37], this problem can be addressed through Bayesian inference. The conditional probability for the adjacency matrix of the graph ($A_{ij} = 1$ if an edge is drawn between nodes i and j , $A_{ij} = 0$ otherwise), reads

$$P(A|\{x_i^0\}) = \prod_{1 \leq i < j \leq N} (2\alpha f(x_i^0, x_j^0; \beta)/N)^{A_{ij}} (1 - 2\alpha f(x_i^0, x_j^0; \beta)/N)^{1-A_{ij}}. \quad (4.5)$$

The function $2\alpha f(x, y; \beta)$ is usually called the affinity function, and in our case is a symmetric

³In the following, we approximate $N - 1$ in the denominator with N , which is reasonable as $N \rightarrow \infty$.

function. The Bayesian belief for the original group assignment, conditional on the particular realisation of the graph, is thus

$$\begin{aligned}
P(\{x_i\}|A) &\propto P(A|\{x_i\})P_X(\{x_i\}) = \\
&= \prod_{i=1}^N P_X(x_i) \prod_{1 \leq i < j \leq N} \left(\frac{2\alpha f(x_i, x_j; \beta)}{N} \right)^{A_{ij}} \left(1 - \frac{2\alpha f(x_i, x_j; \beta)}{N} \right)^{1-A_{ij}}. \quad (4.6)
\end{aligned}$$

We can finally define an extensive Hamiltonian by considering

$$\begin{aligned}
\beta \hat{\mathcal{H}}(\{x_i\}|A) &= -\ln P(\{x_i\}|A) - M \ln N = \\
&= -\sum_{i=1}^N \ln P_X(x_i) - \sum_{1 \leq i < j \leq N} \left[A_{ij} \ln \frac{2\alpha f(x_i, x_j; \beta)}{N} + (1 - A_{ij}) \ln \left(1 - \frac{2\alpha f(x_i, x_j; \beta)}{N} \right) \right]. \quad (4.7)
\end{aligned}$$

We notice, in addition to the expected edge term proportional to A_{ij} , the presence of a weak $1/N$ interaction between non-edges making the graph fully connected, and an external field given by the prior P_X . However, as was shown in [37], the non-edge interaction can be treated in the large N limit in a mean-field way, thus recovering the sparsity of the graph. The BP equations associated to the planted Hamiltonian model (4.7), or community detection inference problem, are discussed in section 4.2. They are formally very close to the standard BP equations for the random model. From model (4.7), we can derive phenomenological bounds on inference based on thermodynamics thresholds, which we expect to be the same for the random and the planted model: for $\alpha < \alpha_d$, there is only one ergodic state which does not carry any information about the original assignment; for $\alpha_d < \alpha < \alpha_c$, the planted cluster is indistinguishable from all the others, and inference is impossible (*i.e.* there is equivalence between the planted and the random ensemble, at the leading order in N); for $\alpha > \alpha_c$, inference is in principle possible, since the planted cluster can be identified as the one with the largest internal entropy. However, as long as the BP equations admit a trivial liquid-like fixed point, to which a random initialization does converge, one should perform an exhaustive search of phase space in order to detect the planted cluster. Finally, inference becomes possible *and* easy, with BP achieving a non-trivial fixed point with positive overlap with the original assignment in linear time even when randomly initialized, only above the Kesten-Stigum threshold α_{KS} (also known as the de Almeida-Thouless local instability of the RS solution [34], or robust reconstruction threshold [49]), where the trivial paramagnetic (liquid) solutions becomes unstable.

4.1.4 Mixed model

Using the community detection formalism, we can recast both the discrete and continuous coloring as an inference problem, also defining a mixed model that allows us to interpolate between the two.

Discrete and continuous coloring as inference problems

Both the discrete and continuous coloring can be treated in a unified way by defining the following affinity function

$$f(x - y; \beta) = \frac{\exp(-\beta \mathbb{I}(\cos(x - y) > \cos \theta))}{\int dx dy P_X(x) P_X(y) \exp(-\beta \mathbb{I}(\cos(x - y) > \cos \theta))}, \quad (4.8)$$

the only difference being in the prior distributions: $P_X(x) = \frac{1}{2\pi}$ for the continuous coloring and $P_X(x) = \frac{1}{q} \sum_{k=0}^{q-1} \delta\left(x - \frac{2\pi k}{q}\right)$ for q -coloring, with $x_i \in [0; 2\pi)$. In the limit $\beta \rightarrow \infty$, this reads

$$f_{q\text{-col}}^{T=0}(x-y) = \begin{cases} 0, & \text{if } \cos(x-y) > \cos(2\pi/q) \\ \frac{q}{q-1}, & \text{otherwise,} \end{cases} \quad (4.9)$$

and

$$f_{\text{CCP}}^{T=0}(x-y) = \begin{cases} 0, & \text{if } \cos(x-y) > \cos(2\pi/q) \\ \frac{q}{q-2}, & \text{otherwise.} \end{cases} \quad (4.10)$$

Interpolating (mixed) model

We also consider a mixed model that interpolates between the discrete and the continuous coloring for each value of $\theta = \frac{2\pi}{q}$, $q \in \mathbb{N}_{>1}$. To this end, we modify the prior by introducing a continuous parameter $\rho \in [0, 1]$, so that for $\rho = 0$ one recovers the q -coloring, while for $\rho = 1$ one recovers the continuous version. We will be interested in the $\beta \rightarrow \infty$ limit of the interaction function, which reads

$$P_X(x) = \frac{\rho}{2\pi} + \frac{1-\rho}{q} \sum_{k=0}^{q-1} \delta\left(x - \frac{2\pi k}{q}\right), \quad \rho \in [0, 1] \quad (4.11)$$

$$f_{\text{Mixed}}^{T=0}(x, y) = \begin{cases} 0, & \text{if } \cos(x-y) > \cos(2\pi/q) \\ \frac{q}{q-1}, & \text{if } x, y \in \{2\pi k/q, k \in \mathbb{Z}\} \text{ and } x \neq y \\ \frac{q}{q-2}, & \text{otherwise.} \end{cases} \quad (4.12)$$

This peculiar shape for $f(x, y)$ is due to the fact that, in order to ensure the normalization condition $\int dx P_X(x) f(x, y) = 1 \forall y \in [0; 2\pi)$, a simpler translationally invariant form is not enough. This can be understood by considering that variables are effectively divided into two classes, depending on whether they can or cannot be written as $y = 2\pi k/q$, for $k \in \mathbb{Z}$. If variable y belongs to the former class, then it is compatible with $(q-1)$ of the q possible “discrete” values that neighbours assume with probability $(1-\rho)$. Conversely, if variable y belongs to the continuous background, by geometrical reasons it is then compatible with only $(q-2)$ of the q possible discrete angles. If the linking probability was a function of only the difference $x-y$, this would induce a difference in the average degree of variable nodes depending on the class, and hence the normalization condition would not hold. One can still work with a uniform average degree by “increasing” the linking probability of variables belonging to the continuous class with respect to the discrete one. Notice that when $\rho = 0$ or $\rho = 1$, there is just one kind of population (the probability for continuous variables with uniform distribution to take a discrete set of specific values being negligible), and the linking probability correctly matches the $\beta \rightarrow \infty$ limits $f_{q\text{-col}}^{T=0}(x-y)$ and $f_{\text{CCP}}^{T=0}(x-y)$ given above.

Discretized version of the inference problems

In the case of discretization, variables x 's can take only p values. We will denote the prior $P_X(x)$ as $\{n_a\} \in \mathbb{R}^p$, with $0 \leq n_a \leq 1$ and $\sum_{a=0}^{p-1} n_a = 1$. The function $f(x, y)$ becomes now a matrix

$\{f_{ab}\} \in \mathbb{R}^{p \times p}$, with $\sum_b n_b f_{ab} = 1 \quad \forall a \in \{0, \dots, p-1\}$.

- Discretized version of the continuous model

$$n_a = \frac{1}{p}, \quad (4.13)$$

$$f_{ab}^{\text{CCP}}(\beta) = \frac{\exp(-\beta \mathbb{I}(\cos(2\pi(a-b)/p) > \cos(2\pi/q)))}{[(2d-1)e^{-\beta} + p - (2d-1)]/p} \quad (4.14)$$

- Discretized q -coloring

$$n_a = \frac{1}{q} \sum_{k=0}^{q-1} \delta_a^{kd} \quad (4.15)$$

$$f_{ab}^{q\text{-col}}(\beta) = \frac{\exp(-\beta \mathbb{I}(\cos(2\pi(a-b)/p) = 1))}{[\exp(-\beta) + (q-1)]/q} \quad (4.16)$$

- Discretized mixed model ($\beta \rightarrow \infty$)

$$n_a = \frac{\rho}{p} + \frac{(1-\rho)}{q} \sum_{k=0}^{q-1} \delta_a^{kd}, \quad \rho \in [0, 1] \quad (4.17)$$

$$f_{ab}^{\text{Mixed}} = \begin{cases} 0, & \text{if } \cos(2\pi(a-b)/p) > \cos(2\pi/q) \\ A = \frac{q[(1-\rho)(q-2) + \frac{\rho}{d}(q-1)]}{(q-1)(q-2 + \frac{\rho}{d})(1-\rho + \frac{\rho}{d})}, & \text{if } \sin(\pi a/d) = \sin(\pi b/d) = 0 \text{ and } a \neq b \\ B = \frac{q}{q-2 + \frac{\rho}{d}}, & \text{otherwise.} \end{cases} \quad (4.18)$$

In the limit $d \rightarrow \infty$ the continuous formulation is correctly recovered, while for $\rho = 0$ and $\rho = 1$ we obtain the $\beta \rightarrow \infty$ limit of $f_{ab}^{q\text{-col}}(\beta)$ and $f_{ab}^{\text{CCP}}(\beta)$, respectively. Expression (4.18) is obtained in the same way as (4.12) by solving for A, B the system $\sum_a n_a f_{ab}^{\text{Mixed}}(A, B) = 1 \quad \forall b$.

4.2 Belief propagation for CCP

In this section, we derive the belief propagation equations for the continuous coloring in the Bayes optimal inference setting of section 4.1.3, following the discussion of [37]. We denote $\nu^{i \rightarrow j}$ the message sent from i to j and ν^i the marginal probability of variable i . Both messages and local marginals are in principle distributions over the continuous interval $[0, 2\pi)$, but thanks to discretization we treat them as p -components arrays $\nu_a^{i \rightarrow j}$ and ν_a^i , with $a \in \{0, \dots, p-1\}$. The belief propagation equations (3.11) for $\nu^{i \rightarrow j}$ associated to the planted Hamiltonian model (4.7) read then

$$\nu_a^{i \rightarrow j} = \frac{1}{Z_{i \rightarrow j}} n_a \prod_{k \neq j} \left[\sum_{b=1}^p (2\alpha f_{ab})^{A_{ik}} (1 - 2\alpha f_{ab}/N)^{1-A_{ik}} \nu_b^{k \rightarrow i} \right], \quad (4.19)$$

while local marginals are computed in terms of the messages as

$$\nu_a^i = \frac{1}{Z_i} n_a \prod_{k=1}^N \left[\sum_{b=1}^p (2\alpha f_{ab})^{A_{ik}} (1 - 2\alpha f_{ab}/N)^{1-A_{ik}} \nu_b^{k \rightarrow i} \right], \quad (4.20)$$

where $Z_{i \rightarrow j}$ and Z_i ensures that $\nu^{i \rightarrow j}$ and ν^i are normalized probability distributions. In the following, we choose to get rid of a term 2α multiplying f_{ab} in the edge-term since it factorizes and just shifts the free energy by a fixed amount.

Since the factor graph of this system is fully connected, one should in principle have to keep track of $N(N-1)$ messages. However, one can notice that the messages $\nu^{i \rightarrow j}$ depend weakly on the indices j as long as $A_{ij} = 0$,

$$\nu_a^{i \rightarrow j} - \nu_a^i = O(1/N) \quad \forall (ij) \notin E. \quad (4.21)$$

In the end [37], this allows to simplify the cavity equations to only keep track of the messages between connected nodes, by introducing an average external field h_a which accounts for the global message coming on each site from non-edges,

$$\nu_a^{i \rightarrow j(t+1)} = \frac{1}{Z_{i \rightarrow j}} n_a e^{-h_a^{(t)}} \prod_{k \in \partial i \setminus j} \left[\sum_{b=1}^p f_{ab} \nu_b^{k \rightarrow i(t)} \right] \quad (4.22)$$

$$\nu_a^{i(t+1)} = \frac{1}{Z_i} n_a e^{-h_a^{(t)}} \prod_{k \in \partial i} \left[\sum_{b=1}^p f_{ab} \nu_b^{k \rightarrow i(t+1)} \right] \quad (4.23)$$

$$h_a^{(t+1)} = \frac{2\alpha}{N} \sum_{k=1}^N \sum_{b=1}^p f_{ab} \nu_b^{k(t+1)}, \quad (4.24)$$

where we have added time indices in order to define update equations. For future convenience, we can also introduce an update function $\Phi^{\text{update}}(\{\nu_k\}_{1 \leq k \leq d}, f, h) \in \mathbb{R}^p$ and a normalization factor $Z(\{\nu_k\}_{1 \leq k \leq d}, f, h) \in \mathbb{R}$, that both take as inputs d generic *messages* and return

$$Z(\{\nu_k\}_{1 \leq k \leq d}, f, h) = \sum_{a=1}^p n_a e^{-h_a} \prod_{1 \leq k \leq d} \left[\sum_{b=1}^p f_{ab} \nu_{k,b} \right], \quad (4.25)$$

$$\Phi_a^{\text{update}}(\{\nu_k\}_{1 \leq k \leq d}, f, h) = \frac{n_a e^{-h_a}}{Z(\{\nu_k\}_{1 \leq k \leq d}, f, h)} \prod_{1 \leq k \leq d} \left[\sum_{b=1}^p f_{ab} \nu_{k,b} \right]. \quad (4.26)$$

Once a fixed point to the cavity equations is found, its corresponding Bethe free entropy (3.12) can be computed as

$$\frac{1}{N} F_{\text{Bethe}}(\{\nu^{i \rightarrow j}\}, \{\nu^i\}) = \frac{1}{N} \sum_{i=1}^N \log(Z_i) - \frac{1}{N} \sum_{1 \leq i < j \leq N, A_{ij}=1} \log(Z_{ij}) - \frac{1}{N} \sum_{1 \leq i < j \leq N, A_{ij}=0} \log(\tilde{Z}_{ij}), \quad (4.27)$$

where,

$$Z_i = \sum_{a=1}^p n_a e^{-h_a} \prod_{k \in \partial i} \left[\sum_{b=1}^p f_{ab} \nu_b^{k \rightarrow i} \right], \quad (4.28)$$

$$Z_{ij} = \sum_{a,b} \nu_a^{i \rightarrow j} f_{ab} \nu_b^{j \rightarrow i}, \quad (4.29)$$

$$\tilde{Z}_{ij} = \sum_{a,b} \nu_a^i (1 - 2\alpha f_{ab}/N) \nu_b^j = 1 - \frac{2\alpha}{N} \sum_{a,b} \nu_a^i f_{ab} \nu_b^j. \quad (4.30)$$

The extra term \tilde{Z}_{ij} comes from the contribution of non-edges, and has the same form of Z_{ij} (messages are replaced with local marginals thanks to relation (4.21)). The sum $\sum \log(\tilde{Z}_{ij})$ can be further simplified. By expanding the logarithm to order $1/N$ and extending the summation over all the $j \in V$, one gets in the limit $N \rightarrow \infty$

$$\frac{1}{N} \sum_{1 \leq i < j \leq N, A_{ij}=0} \log(\tilde{Z}_{ij}) = -\frac{2\alpha}{2} \sum_{a,b} \left[\frac{1}{N} \sum_{1 \leq i \leq N} \nu_a^i \right] f_{ab} \left[\frac{1}{N} \sum_{1 \leq i \leq N} \nu_b^i \right] + o(1) = -\alpha + o(1). \quad (4.31)$$

The last equality comes from the fact that $N^{-1} \sum_i \nu_a^i$ converges to the prior n_a (the numerical solutions to the BP equations from the planted initialization being either $\nu_a^i = n_a \forall i$, or a non-trivial one characterized by a large overlap with the planted configuration, $\nu_a^i \approx \delta_{a,x_i^0}$, where x_i^0 is extracted according to n_a) and from the normalization condition $\sum_a n_a f_{ab} = 1$. Moreover, from this also follows that the fixed-point auxiliary field h_a defined in (4.24) will always be, up to finite size correction, equal to a constant $h_a = 2\alpha, \forall a \in \{0, \dots, p-1\}$. The contribution to the Bethe free entropy from \tilde{Z}_{ij} and h_a (inside Z_i) is hence only an additive constant globally amounting to $-\alpha$, and we could disregard it, as already done before. We choose, however, to keep it in the definition of F_{Bethe} , therefore having

$$\frac{1}{N} F_{\text{Bethe}}(\{\nu^{i \rightarrow j}\}, \{\nu^i\}) = \frac{1}{N} \sum_{1 \leq i \leq N} \log(Z_i) - \frac{1}{N} \sum_{1 \leq i < j \leq N, A_{ij}=1} \log(Z_{ij}) - \alpha. \quad (4.32)$$

The presence of the site-independent field h_a can be useful to cure some instability of the BP equations on a given graph. Since h_a becomes equal to a constant in the thermodynamic limit, we will in the following just remove it from every equation and assume $h_a = 0$ everywhere.

4.2.1 RS population dynamics

In order to analyze the properties of a system in the thermodynamic limit, one is usually interested in computing the quenched average over an ensemble of random graphs. In the Replica Symmetric cavity equations (see also section 3.1.2) one assumes that messages incoming to a variable node are independent, and one can therefore describe the properties of a typical fixed point in the large N limit using a unique distribution of messages $\mathbb{P}(\nu)$. The update equation (4.22) can then be generalized to

$$\mathbb{P}^{(t+1)}(\nu) = \sum_d P_{\text{Edge}}(d) \int \prod_{k=1}^d d\nu_k \mathbb{P}^{(t)}(\nu_k) \delta \left[\nu - \Phi^{\text{update}}(\{\nu_k\}_d, f) \right], \quad (4.33)$$

where Φ^{update} is given by eq. (4.26), and for Erdős-Rényi random graphs one has

$$P_{\text{Node}}(d) = \exp(-2\alpha) \frac{(2\alpha)^d}{d!} \quad (4.34)$$

$$P_{\text{Edge}}(d) = \frac{(d+1)P_{\text{Node}}(d+1)}{\sum_{k=0}^{+\infty} k P_{\text{Node}}(k)} = P_{\text{Node}}(d). \quad (4.35)$$

From the equivalence of $P_{\text{Node}}(d)$ and $P_{\text{Edge}}(d)$, it also follows that for Erdős-Rényi random graphs messages $\nu^{i \rightarrow j}$ and variable marginals ν^i are subjected to the same statistics $\mathbb{P}(\nu)$ given by eq. (4.33).

In the following we will thus neglect the superscripts and work with a unique family $\{\nu\}$ of messages. The Bethe free energy (4.32) can be rewritten as

$$\begin{aligned} \frac{1}{N} F_{\text{RS}}(\mathbb{P}) &= \sum_d P_{\text{Node}}(d) \int \prod_{k=1}^d d\nu_k \mathbb{P}(\nu_k) \log [Z(\{\nu_k\}_d, f)] \\ &\quad - \alpha \int d\nu_1 d\nu_2 \mathbb{P}(\nu_1) \mathbb{P}(\nu_2) \log \left[\sum_{ab} \nu_{1,a} f_{ab} \nu_{2,b} \right] - \alpha, \end{aligned} \quad (4.36)$$

where Z is given by (4.25).

In practice, one represents $\mathbb{P}(\nu)$ by a finite population of U messages, with U large enough (typically, $U = O(10^6)$ messages), and estimates (4.33) by sampling messages from this set. The following algorithm (Algorithm 1) updates the starting population $\mathbb{P} = \{\nu_k\}_{1 \leq k \leq U}$ according to equation (4.33).

Algorithm 1: Replica symmetric population dynamics

```

 $\mathbb{P}^{\text{input}} = \{\nu_k\}_{1 \leq k \leq U};$ 
 $\mathbb{P}^{\text{output}} \leftarrow \{\};$ 
for  $i = 1 \cdots U$  do
    sample  $d$  according to  $P_{\text{Edge}}(d);$ 
     $S \leftarrow \{\};$ 
    for  $j = 1 \cdots d$  do
        sample  $\nu_j$  uniformly from  $\mathbb{P}^{\text{input}};$ 
         $S \leftarrow S \cup \{\nu_j\};$ 
    end
     $\nu \leftarrow \Phi^{\text{update}}(S, f);$ 
     $\mathbb{P}^{\text{output}} \leftarrow \mathbb{P}^{\text{output}} \cup \{\nu\};$ 
end
return  $\mathbb{P}^{\text{output}};$ 

```

4.2.2 Reconstruction equations for $m = 1$

The assumption that neighbouring variables on a cavity graph (when the corresponding function node is pruned) become independent can end up being false if the density of constraints α is too large. In a factor graph with locally tree-like structure, this is usually due to the fact that variables start to develop correlations on large distances, together with the presence of long loops. When this is the case, one has to resort to the 1RSB cavity equations in order to study the system (see also section 3.1.3). The 1RSB population dynamics equations generally describe the distribution of distributions of messages, and are therefore computationally much heavier to implement. Luckily, in the case of $m = 1$ (the m parameter is a Lagrange multiplier with which one can select fixed points of different free energy), and when the BP equations exhibit a para-magnetic (liquid) fixed point, one can perform a simplification of the 1RSB equations that gives them a RS structure (see also appendix A). Since the 1RSB computation at $m = 1$ is correct up to α_c , this allows us to derive the dynamical and Kauzmann transitions.

As the result of this simplification, rather than having only one $\mathbb{P}(\nu)$ distribution of messages,

one keeps track of p distributions $\mathbb{P}_a(\nu)$, $a \in \{0, \dots, p-1\}$, where p is the size of the (discrete) alphabet that variables can assume. The update equation (4.33) takes the following form

$$\mathbb{P}_a^{(t+1)}(\nu) = \sum_d P_{\text{Edge}}(d) \int \prod_{k=1}^d \left[d\nu_k \sum_b n_b f_{ab} \mathbb{P}_b^{(t)}(\nu_k) \right] \delta\left(\nu - \Phi^{\text{update}}(\{\nu_k\}_d, f)\right). \quad (4.37)$$

The Bethe free entropy becomes now

$$\begin{aligned} \frac{1}{N} F_{\text{Rec}}(\{\mathbb{P}_a\}) &= \sum_d P_{\text{Node}}(d) \int \prod_{k=1}^d \left[d\nu_k \sum_b n_b f_{ab} \mathbb{P}_b(\nu_k) \right] \log [Z(\{\nu_k\}_d, f)] \\ &\quad - \alpha \sum_{a,b} n_a n_b f_{ab} \int d\nu_1 d\nu_2 \mathbb{P}_a(\nu_1) \mathbb{P}_b(\nu_2) \log \left[\sum_{r,s} \nu_{1,r} f_{rs} \nu_{2,s} \right] - \alpha \end{aligned} \quad (4.38)$$

To iteratively find a non-trivial fixed point to (4.37), one starts with p populations $\mathbb{P}_a^{\text{Init}} = \{\nu_k\}_{1 \leq k \leq U}$ initialized in the planted solution

$$\nu_{k,b} = \delta_{a,b} \quad \forall \nu_k \in \mathbb{P}_a^{\text{Init}}, \quad a \in \{0, \dots, p-1\} \quad (4.39)$$

The algorithm implementing eq. (4.37) is given in Algorithm 2.

Algorithm 2: General reconstruction equation population dynamics

```

 $\mathbb{P}_a^{\text{input}} = \{\nu_k\}_{1 \leq k \leq U} \quad \forall a \in \{0, \dots, p-1\};$ 
 $\mathbb{P}_a^{\text{output}} \leftarrow \{\}$   $\forall a \in \{0, \dots, p-1\};$ 
for  $a = 0 \dots (p-1)$  do
    for  $i = 1 \dots U$  do
        sample  $d$  according to  $P_{\text{Edge}}(d)$ ;
         $S \leftarrow \{\}$ ;
        for  $j = 1 \dots d$  do
            sample  $b \in \{0, \dots, p-1\}$  with probability  $f_{ab} n_b$ ;
            sample  $\nu_j$  uniformly from  $\mathbb{P}_b^{\text{input}}$ ;
             $S \leftarrow S \cup \{\nu_j\}$ ;
        end
         $\nu \leftarrow \Phi^{\text{update}}(S, f)$ ;
         $\mathbb{P}_a^{\text{output}} \leftarrow \mathbb{P}_a^{\text{output}} \cup \{\nu\}$ ;
    end
end
return  $\mathbb{P}_a^{\text{output}}$ ;
    
```

As long as one works with a flat prior, as $P_X(x) = \frac{1}{2\pi}$ or $n_a = \frac{1}{p} \equiv n \quad \forall a \in \{0, \dots, p-1\}$, and the matrix f_{ab} depends only from the modulus of the distance between variables $f_{ab} = f_{0, (b-a+p) \bmod p}$, the system has a global rotational symmetry and one can avoid keeping track of p populations P_a , while can just focus on messages centered around zero and then shift them. The presence of a rotational symmetry implies

$$\mathbb{P}_b(\nu) = \mathbb{P}_a(T(a-b)\nu), \quad (4.40)$$

where $T(\Delta x)$ is the linear transformation that shifts a message $\nu(x)$ by an amount Δx . This allows us to simplify the equations, by keeping track of one distribution of messages \mathbb{P} only, as in the RS case. For convenience, we fix the index a in equation (4.37) to $a = 0$, corresponding to messages centered around $x = 0$. We also name $f_a \equiv f_{0a}$ and $f_{b-a} \equiv f_{0,(b-a+p) \bmod p}$ (the matrix f_{ab} being a single-variable function f_a for $a \in \{0, \dots, p-1\}$, with parity $f_a = f_{p-a}$). The update equation for the distribution of messages \mathbb{P} , along with the planted initial condition, then becomes

$$\mathbb{P}^{(t+1)}(\nu) = \sum_d P_{\text{Edge}}(d) \int \prod_{k=1}^d \left[d\nu_k \sum_b n f_b \mathbb{P}^{(t)}(T(-b)\nu_k) \right] \delta\left(\nu - \Phi^{\text{update}}(\{\nu_k\}_d, f)\right), \quad (4.41)$$

$$\nu_{k,b} = \delta_{0,b} \quad \forall \nu_k \in \mathbb{P}^{\text{Init}}. \quad (4.42)$$

The implementation of the precedent equation is displayed in Algorithm 3. The procedure formally consists in a RS population dynamics, where a planting is enforced through the introduction of random shifts for the messages extracted according to the interaction probability $n f_b$. Finally, the 1RSB $m = 1$ Bethe free entropy takes the form

$$\begin{aligned} \frac{1}{N} F_{\text{Rec II}}(\mathbb{P}) &= \sum_d P_{\text{Node}}(d) \int \prod_{k=1}^d \left[d\nu_k \sum_b n f_b \mathbb{P}^{(t)}(T(-b)\nu_k) \right] \log [Z(\{\nu_k\}, f)] \\ &\quad - \alpha \int d\nu_1 d\nu_2 \mathbb{P}^{(t)}(\nu_1) \sum_b n f_b \mathbb{P}^{(t)}(T(-b)\nu_2) \log \left[\sum_{r,s} \nu_{1,r} f_{s-r} \nu_{2,s} \right] - \alpha. \end{aligned} \quad (4.43)$$

Algorithm 3: Simplified reconstruction equation population dynamics (II)

$\mathbb{P}^{\text{input}} = \{\nu_k\}_{1 \leq k \leq U}$;

$\mathbb{P}^{\text{output}} \leftarrow \{\}$;

for $i = 1 \dots U$ **do**

sample d according to $P_{\text{Edge}}(d)$;

$S \leftarrow \{\}$;

for $j = 1 \dots d$ **do**

sample $b \in \{0, \dots, (p-1)\}$ with probability $n f_b$;

sample ν_j uniformly from $\mathbb{P}^{\text{input}}$;

$S \leftarrow S \cup \{T(-b)\nu_j\}$;

end

$\nu \leftarrow \Phi^{\text{update}}(S, f)$;

$\mathbb{P}^{\text{output}} \leftarrow \mathbb{P}^{\text{output}} \cup \{\nu\}$

end

return $\mathbb{P}^{\text{output}}$;

4.2.3 Stability analysis

The stability of the BP fixed points can be analyzed by using the population dynamics equations. One way to do this is by creating two copies of the same population and then by slightly perturbing one of the two, *i.e.* $P_1 = \{\nu_k\}_{1 \leq k \leq U}$ and $P_2 = \{\nu_k + \epsilon_k\}_{1 \leq k \leq U}$, where the ϵ_k 's are small perturbations to the fixed point messages. One then keep updating in parallel both of them, using the same

choice of messages, and tracks whether and how the difference between the two populations grows or decreases. A cleaner way to perform this task is by only keeping track of the difference between the original and perturbed messages to first order. This leads to some modified version of the reconstruction equation, where one simultaneously evolve both the messages ν_k and their linear perturbations ϵ_k ,

$$\begin{aligned} \mathbb{P}^{(t+1)}(\nu, \epsilon) &= \sum_d P_{\text{Edge}}(d) \int \prod_{k=1}^d \left[d\nu_k d\epsilon_k \sum_b n f_b \mathbb{P}^{(t)}(T(-b)\nu_k, T(-b)\epsilon_k) \right] \\ &\quad \delta\left(\nu - \Phi^{\text{update}}(\{\nu_k\}, f)\right) \delta\left(\epsilon - \sum_{j=1}^d \frac{\partial \Phi^{\text{update}}(\{\nu_k\}, f)}{\partial \nu_j} \epsilon_j\right). \end{aligned} \quad (4.44)$$

In the argument of the second δ function we have used a shorthand notation for the following matrix product

$$\begin{aligned} \left(\frac{\partial \Phi^{\text{update}}(\{\nu_k\}, f)}{\partial \nu_j} \epsilon_j \right)_a &= \sum_b \frac{\partial \Phi_a^{\text{update}}(\{\nu_k\}, f)}{\partial m_{j,b}} \epsilon_{j,b} = \\ &= \frac{n}{Z(\{\nu_k\}, f)} \left(\prod_{1 \leq k \leq d, k \neq j} \hat{\nu}_{k,a} \right) \hat{\epsilon}_{j,a} - \frac{\Phi_a^{\text{update}}(\{\nu_k\}, f)}{Z(\{\nu_k\}, f)} \sum_c n \left(\prod_{1 \leq k \leq d, k \neq j} \hat{\nu}_{k,c} \right) \hat{\epsilon}_{j,c}, \end{aligned} \quad (4.45)$$

where we have introduced for convenience the auxiliary (unnormalized) messages

$$\hat{\nu}_{k,a}(\nu_k, f) = \sum_b \nu_{k,b} f_{b-a}, \quad (4.46)$$

$$\hat{\epsilon}_{k,a}(\epsilon_k, f) = \sum_b \epsilon_{k,b} f_{b-a}. \quad (4.47)$$

Algorithm 4: Simplified reconstruction equation with tracking of first order difference

$\mathbb{P}^{\text{input}} = \{(\nu_k, \epsilon_k)\}_{1 \leq k \leq U}$;

$\mathbb{P}^{\text{output}} \leftarrow \{\}$;

for $i = 1 \cdots U$ **do**

sample d according to $P_{\text{Edge}}(d)$;

$S \leftarrow \{\}$;

for $j = 1 \cdots d$ **do**

sample $b \in \{0, \dots, p-1\}$ with probability $n f_b$;

sample (ν_j, ϵ_j) uniformly from $\mathbb{P}^{\text{input}}$;

$S \leftarrow S \cup \{T(-b)\nu_j, T(-b)\epsilon_j\}$;

end

$\nu \leftarrow \Phi^{\text{update}}(S, f)$;

$\epsilon \leftarrow \sum_{k=1}^d \frac{\partial \Phi^{\text{update}}(S, f)}{\partial \nu_k} \epsilon_k$;

$P^{\text{output}} \leftarrow P^{\text{output}} \cup \{\nu, \epsilon\}$;

end

return $\mathbb{P}^{\text{output}}$;

4.3 Continuous vs q -coloring

In this section we discuss the phase diagram of the continuous coloring, as obtained by numerically solving the $m = 1$ BP equations up to α_c . This means that we will focus on the detection of the dynamic and condensation thresholds $\alpha_d = \alpha_d(T = 0)$ and $\alpha_c = \alpha_c(T = 0)$, by studying the formation of *typical* clusters. The same analysis is performed in temperature, allowing the calculation of the transition lines $\alpha_d(T)$ and $\alpha_c(T)$, or equivalently $T_d(\alpha)$, $T_c(\alpha)$. All the transition lines for the continuous coloring here presented are computed using a discretization $d = 10$; we refer the reader to section 5.3.1 for an analysis on the corrections to the continuous limit, that are shown to scale as $1/d$.

We also perform an analytical computation of the Kesten-Stigum transition line $\alpha_{\text{KS}}(T)$, which signals the point where the BP para-magnetic fixed point becomes unstable (it is equivalent to de Almeida-Thouless local stability condition in the context of spin glasses [34]). Above α_{KS} inference is *easy*, because BP equations converge towards the planted fixed point even if initialized with a random condition. Following [37], we can relate the stability of the trivial solution with the top eigenvalue λ of the “transfer” matrix

$$T_{ab} = n_a(f_{ab} - 1), \quad (4.48)$$

which satisfies the stability criterion

$$2\alpha_{\text{KS}}\lambda^2 = 1. \quad (4.49)$$

For $2\alpha\lambda > 1$ a local perturbation to the fixed point is exponentially amplified along the graph. For $\rho = 0$ (q -coloring) and $\rho = 1$ (continuous coloring) the top eigenvalue λ can be derived explicitly, thus yielding (at $T = 0$)

$$\alpha_{\text{KS}}^{q\text{-col}} = (q - 1)^2/2, \quad (4.50)$$

$$\alpha_{\text{KS}}^{\text{CCP}} = \frac{1}{2} \left[\frac{\sin\left(\pi\left(1 - \frac{2}{q}\right)\right)}{\pi\left(1 - \frac{2}{q}\right)} \right]^{-2} \underset{q \rightarrow \infty}{=} O(q^2/8). \quad (4.51)$$

In the case of the continuous coloring, it is also interesting to consider the effects of discretization with precision d (recall that $p = qd$ is the number of clock states)

$$\alpha_{\text{KS}}^{d\text{-CCP}} = \frac{1}{2} \left[\frac{\sin\left(\left((q - 2)d + 1\right)\frac{\pi}{dq}\right)}{\sin\left(\frac{\pi}{dq}\right)\left((q - 2)d + 1\right)} \right]^{-2}. \quad (4.52)$$

By taking $d = 1$ or $d \rightarrow \infty$ one recovers $\alpha_{\text{KS}}^{q\text{-col}}$ or $\alpha_{\text{KS}}^{\text{CCP}}$ respectively. It is worth noticing the presence of a factor $1/4$ between the Kesten-Stigum bounds for the continuous and the discrete coloring, $\alpha_{\text{KS}}^{\text{CCP}} \approx \frac{1}{4}\alpha_{\text{KS}}^{q\text{-col}}$ for $q \gg 1$.

The three $T = 0$ transition lines α_d , α_c and α_{KS} are shown in figure 4.3 as a function of q for both the q -coloring on the left and the discretized version of the continuous coloring ($d = 10$) on the right. Looking at the values of α in the figure, it is evident how all the transitions appear in the continuous coloring before (*i.e.* at lower α) than in the discrete one. This is interesting and somehow surprising, as it results that despite *enlarging* the set of solutions for each value of α (all

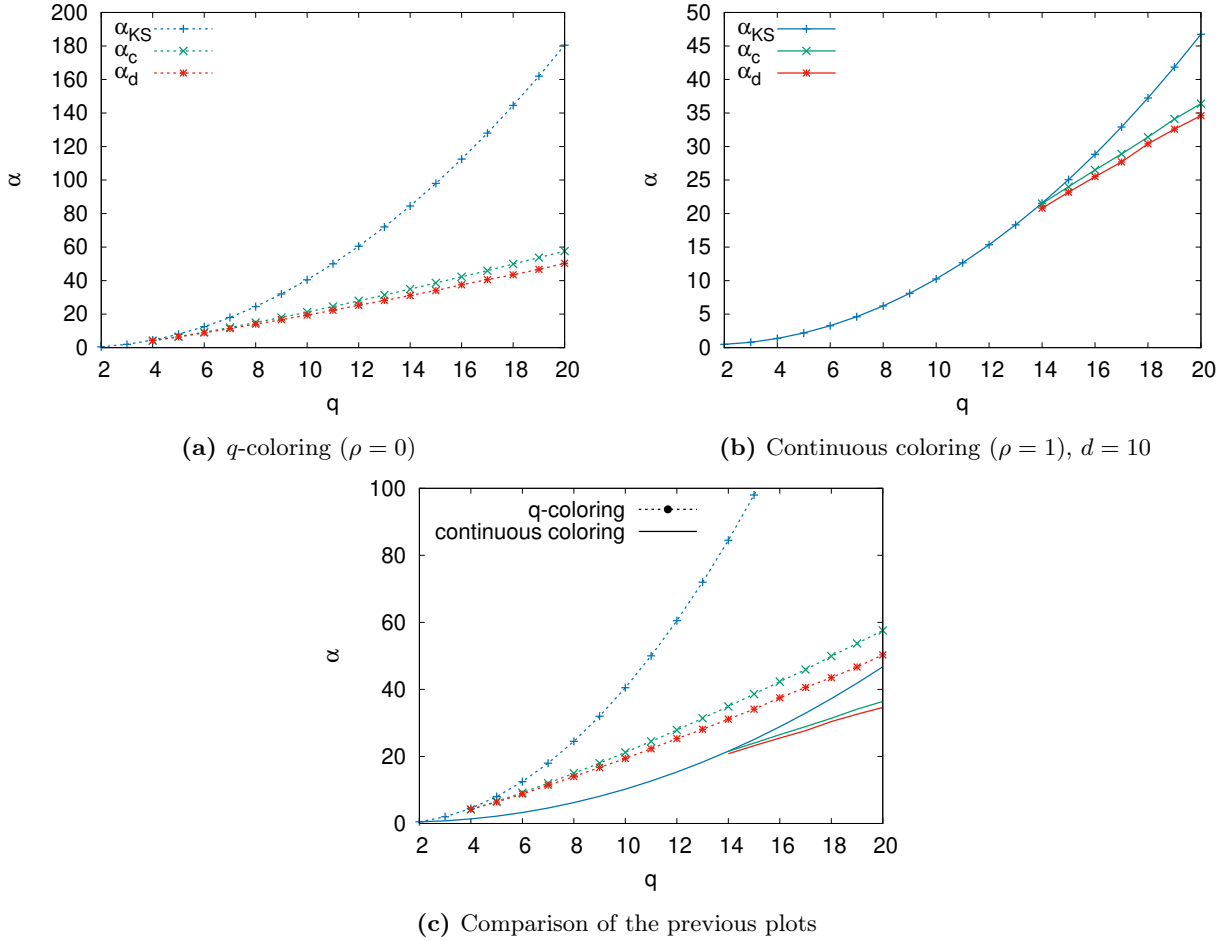


Figure 4.3: $T = 0$ phase diagram in the (q, α) plane for both the discrete (a) and continuous coloring (b) on Erdős-Rényi random graphs. Third figure (c) serves as a direct comparison between the two models. In the discrete coloring, q can only take integer values, the dotted transition lines in the figure being meant to guide the eye. The discrete model exhibits for $q \geq 4$ a random first order transition (the three transition lines merge at $\alpha = 2$ only for $q = 3$). In the continuous case, a random first order transition is found for $q \geq 14$. This is however just an upper bound to the actual point where the three transition lines are expected to merge. This point is likely to be located at a value of q that is not an integer. We notice that, for a given q , all the transitions in the continuous model exhibit a lower value of α than in the q -coloring, see also fig. 4.4a and the discussion in the text.

the solutions to the q -coloring also satisfy the continuous problem with $\theta = \frac{2\pi}{q}$ by definition), the clustering point is anticipated. As a practical consequence, since the ergodicity breaking point α_d is generally connected (even if it does not rigorously coincide) with the onset of hardness in the constraint satisfaction optimization problem, this would imply for $q > 14$ the existence of a whole α region where local algorithms searching for solutions to the continuous coloring get stuck and fail, while solutions to the discrete model are less and yet still easy to find.

We interpret this fact by arguing that in the q -coloring the discrete prior forces the solutions to satisfy all the constraints in a more tight, *i.e.* efficient, way. In this case, variables can only take q values, and the angular distance between two neighbours is either 0 (violated constraint) or a multiple of $2\pi/q$. The minimum allowed distance between variables is hence $2\pi/q$, which corresponds to a *contact* in the particle-system jargon. The fraction of contacts is in this case very relevant: the equilibrium distribution of (discrete) angular differences in the para-magnetic phase $\alpha < \alpha_c$ is by definition proportional to f_{ab} , which is 0 for $a = b$ and uniform otherwise, so that the

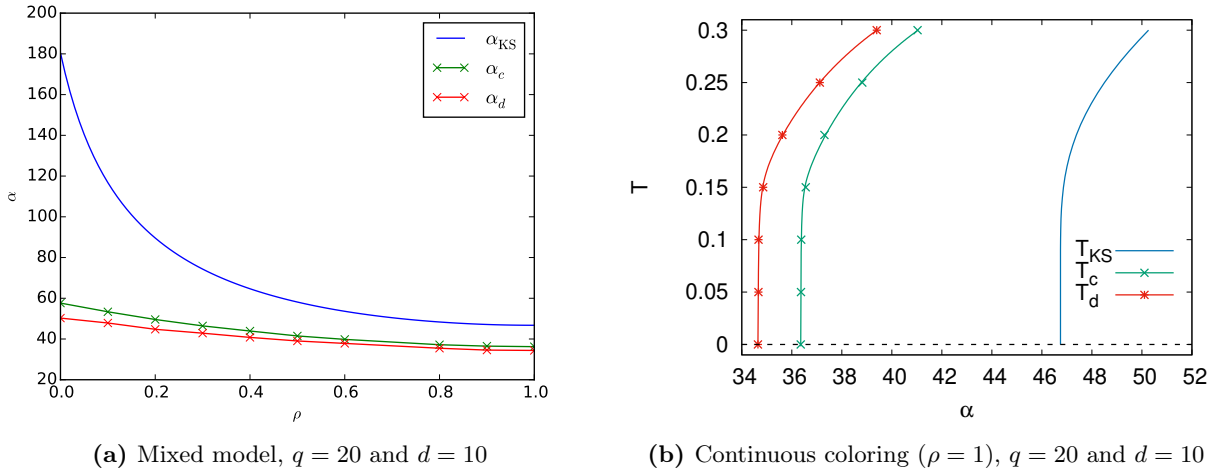


Figure 4.4: Left: $T = 0$ thresholds in the mixed model as a function of ρ , interpolating between the q -coloring ($\rho = 0$) and the continuous coloring ($\rho = 1$) with $\theta = 2\pi/q$. All the transitions appear earlier in the continuous model. Right: Phase diagram in the (α, T) plane for the continuous coloring with fixed $q = 20$. The dynamic and condensation thresholds are obtained by numerically solving the BP equations for different values of T (points).

probability of having a contact is simply equal to $2/(q-1)$, where the factor 2 comes from taking into account both the possibilities to form a contact in one dimension, $x \pm 2\pi/q$. This same analysis can be repeated in the case of continuous coloring. For clarity, let us consider the discretized version with $p = qd$ possible states. The fraction of exact contacts now becomes $2/(p-2d+1) \sim 2/[d(q-2)]$, which is heavily suppressed for $d \rightarrow \infty$. This is totally expected, since in this limit f_{ab} is promoted to a probability distribution on real values. In this situation, one should rather ask for an angular interval $\Delta \equiv n \frac{2\pi}{p}$ with associated probability $1/(q-1)$ on each side of the “particle” by integrating the previous probability. In our discretized notation we can write $\frac{n}{d(q-2)+1} = \frac{1}{q-1}$, from which it follows $\Delta \approx \frac{2\pi}{q} \frac{q-2}{q-1} \approx \frac{2\pi}{q}$. In the continuous coloring, typical solutions satisfy the constraints in a very loose way: exact contacts are “smeared” over an interval comparable to the diameter of the particles. Finally, notice that reversing this perspective the discrete model can be considered, for some qualitative aspect, akin to a continuous sphere system where particles are encouraged to stick together. It is indeed very well known in the literature that, by adding a short range attraction to a hard core repulsion, one can greatly extend the liquid phase of the system [95]. This phenomenon is at the heart of the biased thermodynamics approach discussed in chapter 6.

From figure 4.3 it is also evident that the region where the transition is continuous (small q region) results much more extended in the continuous coloring than in the discrete problem: the discrete q -coloring on Erdős-Rényi random graphs is known to display a random first order transition already for $q > 3$ [101], while in the continuous coloring the continuous transition ranges up to $q \approx 14$. The actual point where the transition changes its nature is in this case likely located at a value of q which is not an integer, close to but smaller than $q = 14$ (we only have points for integer q ’s). Below this point the three transition lines do coincide. In particular, since $\alpha_d = \alpha_c$, the phase dominated by an exponential number of clusters is missing.

In figure 4.4 we fix $q = 20$, for which the transition is appreciably first order, and study how the thresholds vary with ρ at $T = 0$ according to our mixed model (see section 4.1.4), or with temperature at fixed $\rho = 1$ (continuous coloring). In particular, figure 4.4a highlights the fact that

the continuous model has a much lower α_{KS} thresholds, making the associated inference problem hence easier. It is also interesting to observe that the variation of α_{KS} with ρ is much larger than for the other thresholds.

Chapter 5

Simulations

5.1 Numerical determination of α_d

In this section, the question of how to obtain the most accurate numerical estimation of α_d is addressed. The typical way of detecting α_d is by verifying the existence of a non-trivial solution to the planted population dynamics BP equations: usually, one starts from $\alpha > \alpha_d$ and decreases α until the non-trivial solution is lost. This method, however, is prone to inaccuracies. On the one hand, because of the finite number of BP iterations one can perform, one may simply not give the algorithm enough time to forget the planted solution and conclude that the non-trivial fixed point survives for values of α slightly smaller than α_d , thus leading to an underestimation of α_d . On the other hand, fluctuations due to the finite size of the population may become relevant close to α_d , where the non-trivial fixed point loses its stability, this leading to an early loss of the fixed point and to a consequent overestimation of α_d .

For these reasons, a more robust approach taking into account how things scale *around* α_d is very desirable. Following [18], we give independent, compatible estimations of α_d coming from both below and above the dynamic threshold, by studying respectively the (power-law) divergence of BP relaxation time and the birth of instability in the non-trivial fixed point.

The situation is more complicated when one considers real instances of the problem, where the finite size of the system determines the presence of loops in the graph and in general smooths the transitions, which are sharp only in the thermodynamic limit. For large enough sizes, however, we certainly expect the population dynamics BP predictions to be of some use. This can be directly tested through Monte Carlo simulations, where we observe a MCT-like behaviour of the dynamic correlation (the overlap) compatible to $\alpha_d \equiv \alpha_d^{\text{BP}} = 34.63(2)$, at least up to timescales where finite size effects are negligible. In this section we present the results of a detailed analysis at $T = 0$, the same procedure being straightforwardly applicable to any temperature.

5.1.1 Population dynamics

In figure 5.1 we show the data from the numerical solution of population dynamics equations for the Erdős–Rényi ensemble. The size of the population here considered is $U = 10^6$, each message being a probability array of $p = 200$ values ($q = 20$, $d = 10$). We chose the value $q = 20$ because, from figure 4.3, the transition is appreciably first order. On the left panel the overlap decay $q(t)$ is




Figure 5.1: Left: overlap as a function of the number t of BP iterations. As long as $\alpha < \alpha_d$, the overlap decays to zero, thus implying that any memory of the planted initialization is lost and one recovers the paramagnetic solution with uniform messages. Right: ratio $\lambda(t)$ between the magnitude of perturbations to the non-trivial solution at times $t + 1$ and t for $\alpha > \alpha_d$. After a transient, stationarity is reached, thus identifying the stability parameter $\lambda(\alpha) < 1$.

displayed for different values of $\alpha < \alpha_d$. Given a set of U marginals $\{\nu_i\}^t$, the overlap is defined as

$$\begin{aligned} q(t) &= \frac{1}{U} \sum_{i=1}^U \int dx_1 dx_2 \nu_i(x_1) \nu_i(x_2) \cos(x_1 - x_2) \\ &= \frac{1}{U} \sum_{i=1}^U [\langle \cos(x) \rangle_i^2 + \langle \sin(x) \rangle_i^2], \end{aligned} \quad (5.1)$$

where $\langle \cdot \rangle_i$ represents the average over the marginal distribution ν_i . We considered 10 different initializations of the random generator, which we refer to as samples. Each sample has a very smooth behaviour, almost identical to the others except for a slightly difference in the time it leaves the forming plateau. For this reason, the most natural way to present the data is by averaging over t at fixed q , once a spline interpolation of the single sample is performed.

On the right panel, we consider the case $\alpha > \alpha_d$, where the BP equations admit a non-trivial solution with $\lim_{t \rightarrow \infty} q(t) \neq 0$. Once convergence in the overlap is reached, we compute the ratio of the magnitude between two subsequent perturbations around the fixed point (by using the linearized version of the BP equations derived in section 4.2.3), as a function of time

$$\lambda(t) = \frac{\|\epsilon(t+1)\|}{\|\epsilon(t)\|}, \quad (5.2)$$

where we adopted the L_2 norm of the perturbation parameters averaged over the population, $\|\epsilon\| = \frac{1}{pU} \sum_{i=1}^U \sum_{a=0}^{p-1} \epsilon_{i,a}^2$. After a sufficient but finite number of steps, $\lambda(t)$ is actually selecting the contribution from the biggest in modulus eigenvalue of the linearized matrix which enforces the BP iterative procedure (see [18] for a detailed discussion), being the slowest mode to decay. We define the stability parameter $\lambda(\alpha)$ as essentially the modulus of this eigenvalue, with $\lambda(\alpha) < 1$ for the solution to be stable. In practice, we average over t the $\lambda(t)$ data of figure 5.1b once the plateau is reached. In this case we consider just one sample for each value of $\alpha > \alpha_d$.

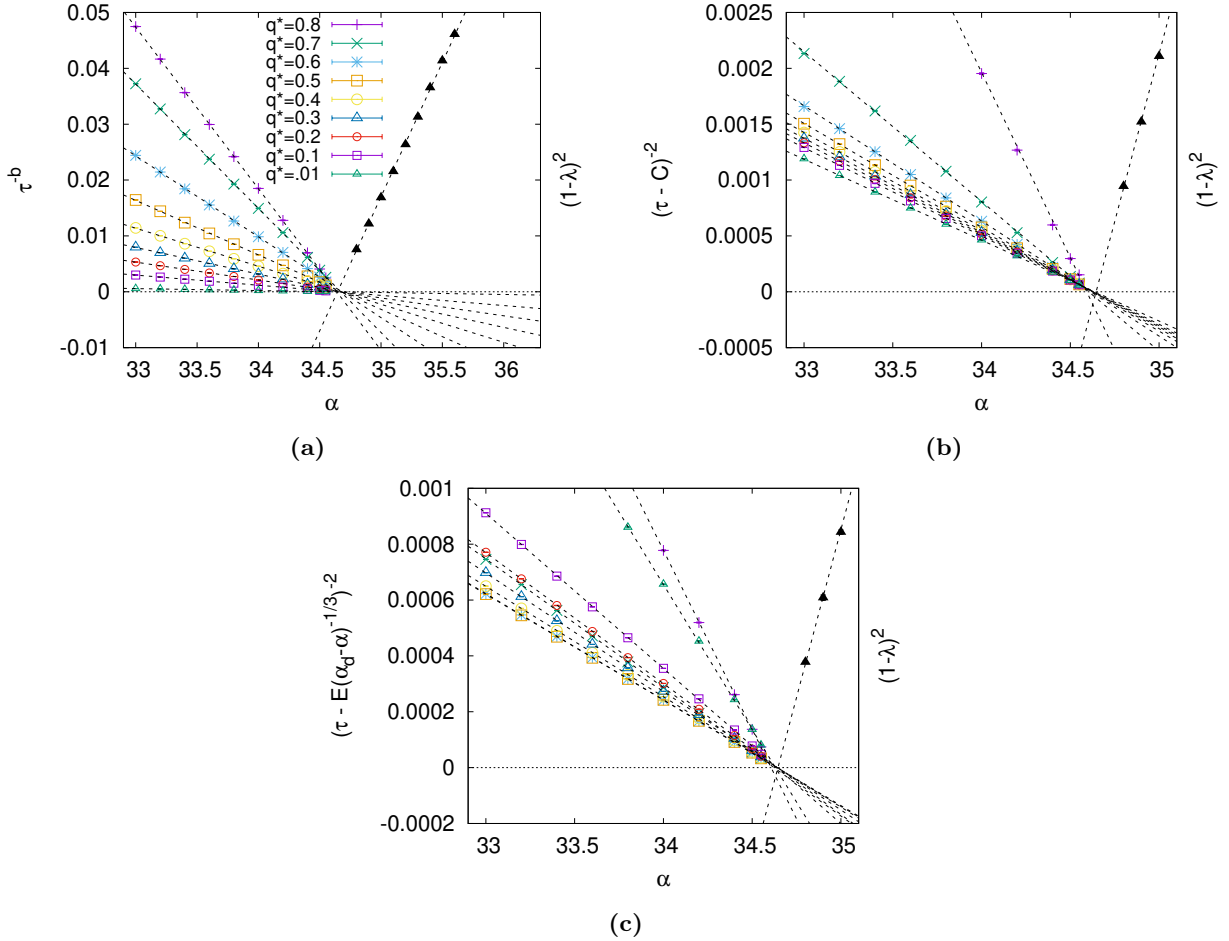


Figure 5.2: Colors: divergence of the overlap relaxation time $\tau_{q^*}(\alpha)$, defined as the time needed for $q(t)$ to decrease down to q^* . Figures (a), (b) and (c) show the data rescaled according to functions (5.3), (5.4) and (5.5), respectively. Black: behaviour of $(1 - \lambda(\alpha))^2$, where $\lambda(\alpha)$ is the stability parameter of the non-trivial solution for $\alpha > \alpha_d$. The correct scale for λ is given in the first figure (in both figures (b) and (c), black points follow a different y-scale than colored ones and are just meant to guide the eye towards the transition point). The dynamical transition α_d is located at the intersection of the color and black lines with the x-axis.

Let us first consider the region $\alpha < \alpha_d$. From figure 5.1a, it is clear that approaching α_d the overlap relaxation becomes increasingly slower. Following [18], this behaviour can be directly connected to the discontinuous bifurcation occurring to the fixed-point solution of the BP equation, which is of the form $\mathbb{P} = \mathcal{F}(\mathbb{P}, \alpha)$. To fix the ideas, it is convenient to consider an analogous relation in the scalar case, $q = f(q, \alpha)$, where q is real (the similarity with the overlap, which is indeed our scalar representative for \mathbb{P} , will be readily evident). Assuming that $q \geq 0$ and that $q = 0$ is always a solution, the dynamic threshold α_d is defined as the subitaneous birth of a second solution $q_d = f(q_d, \alpha_d)$, with $q_d > 0$. Our problem reduces to the study of the dynamic process $q^{(t+1)} = f(q^{(t)}, \alpha)$, with initial condition $q^{(t=0)} > q_d$ (in analogy with the planted initialization). In this context, it can be proved that the number of steps t around the plateau q_d diverges as $\tau \sim K(\alpha_d - \alpha)^{-\frac{1}{2}}$ to leading order for $\alpha \rightarrow \alpha_d^-$, where K is fixed by some derivative of $f(q, \alpha)$ at the special point (q_d, α_d) .

In order to test this prediction on our real data, we operatively define a family of relaxation times $\tau_{q^*}(\alpha)$ representing, for any given threshold $q^* < q_d = 0.838(3)$ (obtained in section 5.1.2), the

time needed for the overlap $q(t)$ to decrease down to q^* . By looking at the data in figure 5.1a, one notices that different sets of $\tau_{q^*}(\alpha)$ for different q^* should essentially differ by only a finite amount of steps, which is reasonably constant in α for any given q^* since the curves have approximately the same behaviour (same slope) on a linear x-scale. However, if we also consider the relatively small values of t needed to complete the decay and hence of $\tau_{q^*}(\alpha)$, it follows that relevant subleading corrections to the asymptotic scaling may be due just to the choice of q^* . This is evidenced by the behaviour of the first function that we use to fit the data, neglecting any scaling correction,

$$\tau^{(1)}(\alpha) = \frac{A}{(\alpha_d - \alpha)^{1/b}}. \quad (5.3)$$

This function depends on three free parameters: α_d , b and A . Even though at first glance it may appear to reasonably fit the data, see figure 5.2a (colors), all the values of the parameters, presented in table 5.1 (left), exhibit some dependence on q^* , which is very noticeable in the case of the exponent b (that was argued to be $b = 2$). To account for subleading corrections, we also consider the two following power laws, still depending on three parameters,

$$\tau^{(2)}(\alpha) = \frac{B}{(\alpha_d - \alpha)^{1/2}} + C, \quad (5.4)$$

$$\tau^{(3)}(\alpha) = \frac{D}{(\alpha_d - \alpha)^{1/2}} + \frac{E}{(\alpha_d - \alpha)^{1/3}}. \quad (5.5)$$

The rescaled data is presented in figures 5.2b and 5.2c. For what concerns the results of fit $\tau^{(2)}$, we obtain good χ^2 and a much less q^* -dependent slope (*i.e.* B) especially for intermediate values of $0.2 \leq q^* \leq 0.5$. We conclude that in this range the scalar bifurcation scaling form with exponent

q^*	α_d	b	A	χ^2/dof	α_d	B	C	χ^2/dof
0.8	34.639(2)	1.384(5)	12.94(2)	4.39	34.592(1)	17.4(1)	-4.77(5)	22.81
0.7	34.672(2)	1.280(4)	19.57(4)	10.83	34.604(1)	27.3(1)	-8.58(6)	30.19
0.6	34.680(2)	1.345(4)	23.31(4)	21.21	34.615(1)	31.1(1)	-8.76(7)	6.62
0.5	34.677(2)	1.427(4)	25.65(4)	23.81	34.621(1)	32.8(1)	-7.99(7)	1.46
0.4	34.671(2)	1.507(4)	27.43(4)	23.41	34.624(1)	33.7(1)	-7.05(7)	0.42
0.3	34.664(2)	1.588(5)	28.99(4)	20.64	34.627(1)	34.4(1)	-6.01(7)	0.80
0.2	34.657(2)	1.679(5)	30.63(4)	17.18	34.629(1)	35.0(1)	-4.75(7)	1.68
0.1	34.647(2)	1.807(5)	32.85(3)	12.21	34.632(1)	35.6(1)	-2.94(8)	3.42
0.01	34.627(1)	2.164(5)	39.48(3)	3.83	34.638(1)	37.2(1)	2.51(8)	9.80
q^*	α_d	D	E	χ^2/dof				
0.8	34.607(1)	27.9(2)	-15.2(2)	2.89				
0.7	34.623(1)	46.6(2)	-27.7(2)	3.20				
0.6	34.636(1)	51.3(3)	-28.6(2)	1.69				
0.5	34.641(1)	51.4(3)	-26.3(3)	4.37				
0.4	34.643(1)	50.3(3)	-23.3(3)	6.64				
0.3	34.643(1)	48.6(3)	-19.9(3)	7.79				
0.2	34.642(1)	46.2(3)	-15.8(3)	8.26				
0.1	34.639(1)	42.4(3)	-9.7(3)	7.99				
0.01	34.631(1)	31.0(3)	8.5(3)	5.34				

Table 5.1: Results from fit (5.3) (on the left), (5.4) (on the right) and (5.5) (bottom) for the divergence of the overlap relaxation time $\tau_{q^*}(\alpha)$ as obtained from population dynamics.

α_d	χ^2/dof
34.645(2)	1.39

Table 5.2: Results from the fit on the stability parameter $\lambda(\alpha)$.

$b = 2$ adequately describes our $\tau_{q^*}(\alpha)$ data. On the other hand, the function $\tau^{(2)}$ is shown to fail for values of q^* close to the plateau, $q^* = 0.8, 0.7$ (compare the values of χ^2 in table 5.1, right). In this case fit $\tau^{(3)}$ does a better job, thus suggesting that the scaling $b = 2$ is just apparently altered by some non-trivial subleading corrections. In general, we observe that function $\tau^{(3)}$, even though outperformed by $\tau^{(2)}$ when q^* takes the central values, always performs better than fit $\tau^{(1)}$, apart for the case $q^* = 0.01$ for which already $\tau^{(1)}$ produces a value of b close to $b = 2$. This corroborates the argument supporting $b = 2$; a cautious estimate for α_d from our data would be at this point $\alpha_d = 34.625(19)$.

A second independent estimation of α_d can be obtained by studying the loss of stability for the non-trivial solution when approaching α_d from above. A quantitative prediction from scalar bifurcation theory [18] states that $(1 - \lambda(\alpha))^2 \sim Y(\alpha - \alpha_d)$ vanishes linearly for $\alpha \rightarrow \alpha_d^+$. This is shown in figure 5.2a in black points (beware that points reproduced in figures 5.2b and 5.2c are out of scale), while the estimated α_d from the fit is given in table 5.2. Comparing it with the result from the analysis of the overlap decay for $\alpha < \alpha_d$, we conclude $\alpha_d = 34.63(2)$. The quoted uncertainty takes into account the systematic errors affecting the two independent estimations.

5.1.2 Zero Energy Monte Carlo

Monte Carlo (MC) simulations have the advantage of being much less eager of computer memory as compared to solving the BP equations, where one needs to work with entire distributions (as each of the messages) even in our RS-like case. On the other hand, relaxation times can be very long (we are approaching a MCT-like glass transition, indeed). This can be noticed by looking at figure 5.3a, where we plot the overlap decay for several orders of magnitude in the number of MC sweeps. Before analyzing the results, we give a precise definition of the overlap and discuss the details of the MC simulations at zero temperature.

The overlap between two configurations of angles \underline{x}_a and \underline{x}_b needs to be maximized over rotations (see also section 5.3.2) due to the global rotational symmetry of the model, which becomes continuous in the limit $d \rightarrow \infty$ of infinite discretization precision,

$$q(t) = \max_{A \in [0, 2\pi)} \langle \cos(x_a - x_b - A) \rangle = \sqrt{\langle \cos(x_a - x_b) \rangle^2 + \langle \sin(x_a - x_b) \rangle^2}, \quad (5.6)$$

where $\langle \cdot \rangle = \frac{1}{N} \sum_{i=1}^N$ stands for the average over the system. We consider both the overlap $q_{ab}(t)$ between two independent replicas of the same system, in which case $x_a = x_i^{(a)}(t)$ and $x_b = x_i^{(b)}(t)$, and the overlap $q(t)$ of each replica with the starting configuration, in which case $x_a = x_i(t)$ and $x_b = x_i(t = 0)$. The latter situation can be further simplified by initially planting the system in a configuration with all the variables identically equal to zero, which is possible thanks to the introduction of random shifts associated to each edge of the graph (that are required in the first place in order to suppress crystallization). Of course, we expect the equilibrium long time limit of both the overlaps $q(t)$ and $q_{ab}(t)$ to be identical. However, it is also clear that $q_{ab}(t)$ will in general

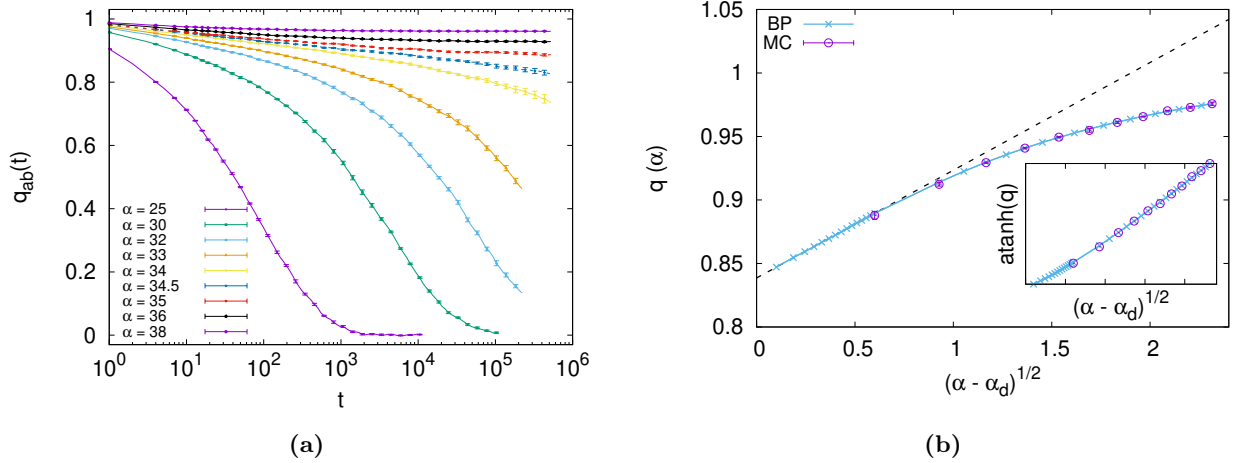


Figure 5.3: Left: overlap as a function of time between two independently evolving replicas starting from the same equilibrated (planted) initial condition. Different curves correspond to different values of α . Dashed lines for $\alpha = 34.5$ and $\alpha = 35$ indicate the region where the overlap supposedly starts to develop a plateau with $q = q_d$. Right: long time value of the overlap for $\alpha > \alpha_d$ from MC and BP simulations. The birth of the non-trivial solution follows a square root singularity, allowing us to extrapolate from the BP data $q_d = 0.838(3)$. The data is scaled according to $\alpha_d = 34.63$ from BP population dynamics. Inset: the trivial bending of data points for high α can be straightened by plotting $\tanh^{-1}(q(\alpha))$.

decorrelate faster.

For what concerns the algorithm, we adopt in this case a simple zero temperature heat bath rule: at each step, a randomly selected spin is updated with uniform probability among the Potts states which satisfy all the constraints with the neighbours. Since we start from a planted configurations of exactly zero energy, at least one such a state is guaranteed to exist for every spin, and the procedure is well posed. The advantage with respect to standard Metropolis is that moves are always accepted, this coming at the price of computing the local energy for each of the p states a variable can assume. Notice, however, that this can be performed in a smart way by running just one time over the neighbours of the variable to update and increasing by 1 the energy of the $2d - 1$ states inside the excluded region.

Obtaining a precise estimate of α_d from MC data is more difficult than from population dynamics. Figure 5.3a shows as a function of time (MC sweeps) the average of the overlap over 10 different systems of size $N = 10^4$. We checked finite size corrections to be smaller than the statistical errors by simulating a few samples of size $N = 10^5$. For $\alpha > \alpha_d$, planting allows us to initialize the system *inside* the planted cluster, which is by all means indistinguishable from a typical cluster for the random model, up to α_c . Unfortunately, the $\alpha > \alpha_d$ data is of little use in the estimation of α_d , since we do not have in this case an analogous of the BP stability parameter. A possible way round would be to consider the long time limit $q(\alpha)$ of the overlap at the plateau, which is expected [18] to approach a non zero value $q_d > 0$ with a square root singularity at α_d , as shown in figure 5.3b. Obtaining α_d from the $q(\alpha)$ data, on the other hand, would imply a fit where also the parameter q_d is unknown. For this reason we prefer first to estimate α_d from the overlap decay for $\alpha < \alpha_d$, and then use it to verify the square root singularity and eventually obtain q_d . Notice that the behaviour $(q(\alpha) - q_d) \sim (\alpha - \alpha_d)^{1/2}$ can be valid only in the vicinity of α_d , since we know by definition that $q(\alpha) < 1$. The strong deviation from a straight line behaviour in figure 5.3b is in this sense totally expected. What is inconvenient, though, is the fact that MC points are located




Figure 5.4: Left: power-law divergence of $\tau_{q^*}(\alpha)$, for $q^* = 0.8$. Both the overlap with the initial configuration and the one between two independent replicas are considered (the correlation times are named respectively τ_q and τ_{ab}). In this case we obtain $b \approx 0.26$, corresponding to a MCT-like exponent $\gamma_0 \approx 3.85$. Notice however that the value of b strongly depends on q^* , see table 5.3. Right: the value of α_d resulting from the fit depends both on q^* and on α_{max} , corresponding to the biggest value of α included in the fit. The red line indicates in both figures the estimate of α_d as obtained from BP in the previous section.

in the region where these corrections are already relevant. This is due to the fact that obtaining a reliable, size independent estimation of $q(\alpha)$ by MC simulations becomes very difficult close to α_d , where one has to increase both the time of the simulation and N . For this reason, we can only rely on the BP data, which fit reasonably well the square root singularity, to obtain $q_d = 0.838(3)$.

The last approach we can resort to is the study of the overlap correlation time for $\alpha < \alpha_d$. To this end, we once again define $\tau_{q^*}(\alpha)$ as the time needed for $q(t)$ and $q_{ab}(t)$ to decrease down to q^* , with $q^* < q_d \approx 0.83$. The $T = 0$ approach to α_d is essentially governed by MCT-like laws [87], predicting a power-law divergence $\tau(\alpha) \sim (\alpha_d - \alpha)^{-\gamma_0}$ for $\alpha \rightarrow \alpha_d^-$, where γ_0 is a specific $T = 0$ exponent. Here we prefer to use instead the parameter $b = 1/\gamma_0$, and fit the data according to eq. (5.3). The outcome is given in figure 5.4. In the left panel we show the rescaled data for $q^* = 0.8$, showing a value of α_d compatible with the BP prediction. On the other hand, from figure 5.3a it is already evident that in order to decrease q^* one may be forced to exclude some of the biggest values of α . For this reason, in 5.4b we study how the choice of the α_{max} included in the fit affects the

q^*	α_{max}	b (τ_q)	b (τ_{ab})
0.8	34	0.260(2)	0.263(2)
	33	0.262(3)	0.266(3)
	32	0.266(4)	0.274(5)
0.7	33	0.223(2)	0.225(3)
	32	0.221(3)	0.226(4)
0.6	33	0.206(2)	0.204(2)
	32	0.205(2)	0.200(3)
0.5	33	0.197(2)	0.197(2)
	32	0.197(2)	0.194(2)
0.4	32	0.188(2)	0.186(2)

Table 5.3: Value of the MCT-like exponent $b = 1/\gamma_0$ from Monte Carlo simulations.

estimation of α_d , for different values of q^* . From the picture, one can conclude that estimates on α_d are the lesser precise the lower α_{\max} ; in particular, the points dispersion over q^* at fixed α_{\max} and the size of error bars visibly decrease from $\alpha_{\max} = 32$ to $\alpha_{\max} = 33$ (for $\alpha_{\max} = 34$ we have only one point, so we cannot make any strong assertion). It is interesting, by the way, to highlight how the high dependence of α_d from q^* , for α_{\max} fixed, is accompanied by a relevant variation of also the exponent b with q^* , see table 5.3, similarly to what happened when analyzing BP data in the previous section. In this case, however, we lack at the present moment an analytical prediction for b , and an accurate treatment of subleading corrections to the simple power-law behaviour is beyond our reach. For this reason, we can only conclude from figure 5.4b that the behaviour of real MC data is reasonably well described by the BP prediction, considering the level of the systematic biases affecting our determination of α_d from MC simulations.

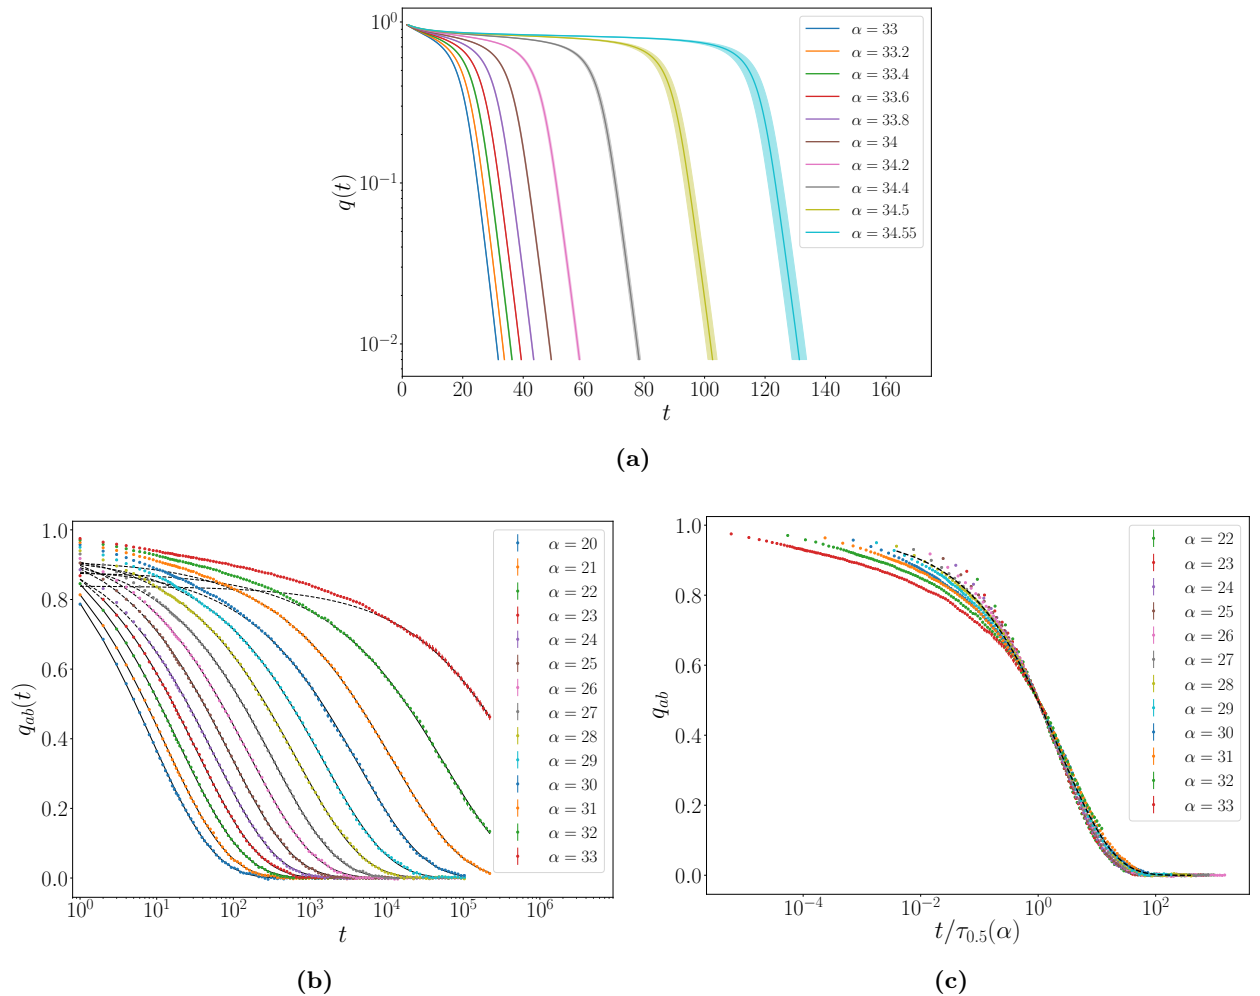


Figure 5.5: Top: overlap decay from population dynamics BP equations for $\alpha < \alpha_d$, using the same data of figure 5.1a. The log-scale on the y-axis highlights a purely exponential tail for the decay at small overlap. Bottom left: overlap decay from MC data. Black lines represent fits according to a stretched exponential of the form $Ae^{-(t/\hat{\tau})^\beta}$. The values of β vary between 0.48 and 0.53. Bottom right: the same data is plotted against $t/\tau_{0.5}(\alpha)$. Black dashed line is a stretched exponential of exponent $\beta = 0.48$.

Finally, it is interesting to consider in figure 5.5 the long time tail of the overlap decay. In the top panel, the BP data is shown to possess a simple exponential behaviour. The MC dynamics of real systems, on the contrary, exhibits a stretched exponential form $Ae^{-(t/\hat{\tau})^\beta}$, typical of MCT

systems close to the MCT transition. The value of the exponent β from the different fits is reasonably constant on the whole considered α -interval, varying between 0.48 and 0.53 without any recognizable trend. The parameter $\hat{\tau}(\alpha)$ from the fit is in principle proportional to the structural relaxation time of the system $\tau(\alpha)$. However, we believe the definition of $\tau(\alpha)$ in terms of $\tau_{q^*}(\alpha)$ to be somewhat more reliable, since it does not depend on the quality of a phenomenological fit and on the considered fitting interval. In the right panel we show the correlation for different values of α as a function of the rescaled time $t/\tau_{0.5}(\alpha)$, resulting in a decent collapse of all the curves. The collapse is not perfect as the tail becomes slightly more pronounced when increasing α .

5.2 Numerical determination of α_c

The condensation transition line $\alpha_c(T)$ is defined by the condition $\Sigma(\alpha, T) = 0$, where Σ is the complexity of the states dominating the partition function at temperature T . Using eq. (2.22), Σ can be written as the difference of Bethe free entropies

$$\Sigma = \frac{1}{N} [F_{\text{Rec II}}(\text{para}) - F_{\text{Rec II}}(\text{planted})], \quad (5.7)$$

where $F_{\text{Rec II}}$ is obtained for any BP fixed point from eq. (4.43), the temperature dependence being encoded into the affinity function f_a . In particular, $F_{\text{Rec II}}(\text{para}) = -\alpha$. A number of curves $\Sigma(\alpha, T)$ for different values of T are plotted in figure 5.6. The $T = 0$ limit is under control thanks to the




Figure 5.6: Complexity curves $\Sigma(\alpha, T)$. Dashed lines are a second order polynomial fit. The intersection of each curve with the x-axis returns the value of the condensation transition $\alpha_c(T)$.

fact that the curves for the two lowest values of T are practically indistinguishable. Notice that the step $\Delta\alpha = 0.2$ in the data is quite large, and the first point in each curve does not exactly correspond to $\alpha_d(T)$. The estimation of Σ is on the contrary much more precise and allows for a reliable determination of $\alpha_c(T)$. We also observe that the maximum of the complexity appears to decrease when increasing temperature. We cannot exclude the possibility that it goes to zero for some finite temperature, *i.e.* that the transition becomes continuous for high enough α .

5.3 Further details

5.3.1 Extrapolation of the continuous limit

The most prominent novelty of the coloring problem discussed in this thesis is, beyond any doubt, the continuous nature of its variables. Yet, as the reader may have already noticed, all the analysis performed so far was only limited to a discretized version of the model, with precision $d = 10$ to be more exact. This coincides, having chosen for definiteness to focus on the fixed $q = 20$ case, for which the transition is appreciably first order, with approximating the continuous interval $[0, 2\pi)$ with $p = qd = 200$ discrete clock-states. It is thus essential to assess how much of the quantitative predictions given so far, in particular for what concerns the computation of the transition lines, is still valid in the limit $p \rightarrow \infty$. The answer is provided in figure 5.7, where we plot the $T = 0$




Figure 5.7: Corrections to the $T = 0$ transition thresholds scale as $1/p$.

thresholds $\alpha_d(p)$ and $\alpha_c(p)$ for $p = 200, 240, 280, 320$, as obtained from the BP numerical procedure discussed in this chapter. The corrections, for the values of p considered, appear to be very relevant. Even worse, they show a very slow $1/p$ scaling. For these reasons, we believe a precise extrapolation of the continuous limit of the thresholds, starting only from the solution of the discretized BP equations, to be rather delicate. In this spirit we report the results from the fits of figure 5.7, to be taken with caution: $\alpha_d^\infty \approx 32.37$ and $\alpha_c^\infty \approx 33.86$. On a more bright side, however, some way to better control the extrapolation error could be hopefully provided by the study of the p dependence of α_{KS} , which is known analytically (4.52), and that also shows an approximate $1/p$ behaviour.

The $1/p$ scaling of discretization corrections for the model at hand was already reported in [80]. Unfortunately, the situation is very different from other planar spin-models on random graphs, exhibiting a very fast exponential convergence [67]. This is presumably due to the peculiarity of the excluded volume interaction, that sharply depends on the particle diameter θ . Among the effects of discretization, we have indeed that the forbidden states are contained by all means into an effective diameter which is smaller than θ by a quantity of order $1/p$ (see for example the red points in figure 4.1). When we increase discretization, we are hence actually increasing also the effective diameter (*i.e.* decreasing the effective q), which may play a direct role in the $1/p$ decrease of the transition thresholds, since the system is more constrained for bigger diameters, see for instance figure 4.3.

5.3.2 Global rotational symmetry

The most naive definition of the overlap for spin systems is given by the dot product between different configurations, which for XY models reads $\frac{1}{N} \sum_i \cos(x_i(t) - x_i(0))$. A clear insight on the reasons why this is not a good observable to describe the state of our system can be obtained by a direct inspection of the configurations visited during the MC evolution. To this end, we plot in figure 5.8 the behaviour of the naive overlap for a single sample along with some “screenshots” of the system at different marked times. The adopted convention is the following. Each pixel in the square boxes has two coordinates, the y -one corresponding to a label for each variable (in this case $N = 10^3$), while the x -axis represents the possible values that each variable can take (as usual, we considered in practice a discretization in $p = 200$ clock states). Due to the rescaling of the axes, in order to have a square ratio for the figures, pixels are actually horizontally stretched. The color legend associates a white pixel to the value taken by each variable at a given time. Then for each row (*i.e.* for each variable) we mark in black the x -values that are compatible with the constraints exerted by its neighbours, considered as frozen in the current configuration, and in grey the values that are not allowed.

Initially, the system is planted in the all zeros configuration. Our graphical expedient allows us to recognize at glance if and how the system loses memory of the initial condition. In particular, for the value of $\alpha = 36$ here considered, that is way beyond α_d , we observe how the system keeps an internal coherence reminiscent of the planted initialization, which is only apparently spoiled at large times when the zero mode associated to a global rotation kicks in. This explains why the top panel of figure 5.8 is so much different from the behaviour of the proper overlap we have already studied in this chapter, and fails at large times to follow the BP prediction (black dashed line).

The decorrelation contribution due to global rotations can be dismissed by working in the “center of mass” of the system. By placing variables on the unit circle, the center of mass is identified by the vector $(\langle \cos(x(t)) \rangle, \langle \sin(x(t)) \rangle)$, where $\langle \cdot \rangle$ is the average over the system. The modulus of this vector measures the remaining degree of polarization (it is equal to 1 when all the variables take the same value, as in the planted configuration where all $x = 0$, and to 0 when variables are uniformly spread over the circle) and coincides exactly with our rotationally invariant overlap definition (5.6). Equivalently, we can interpret the angle A^* maximizing the second term in (5.6) as the one identified by the vector of the center of mass. Notice that formula (5.6) is completely general, since it accounts for any initial condition $x_i(t = 0)$ and also for the overlap q_{ab} between different replicas (in that case, the underlying idea is again to globally shift back replicas on top of each other at every time in order to correctly evaluate the internal correlation of the state).

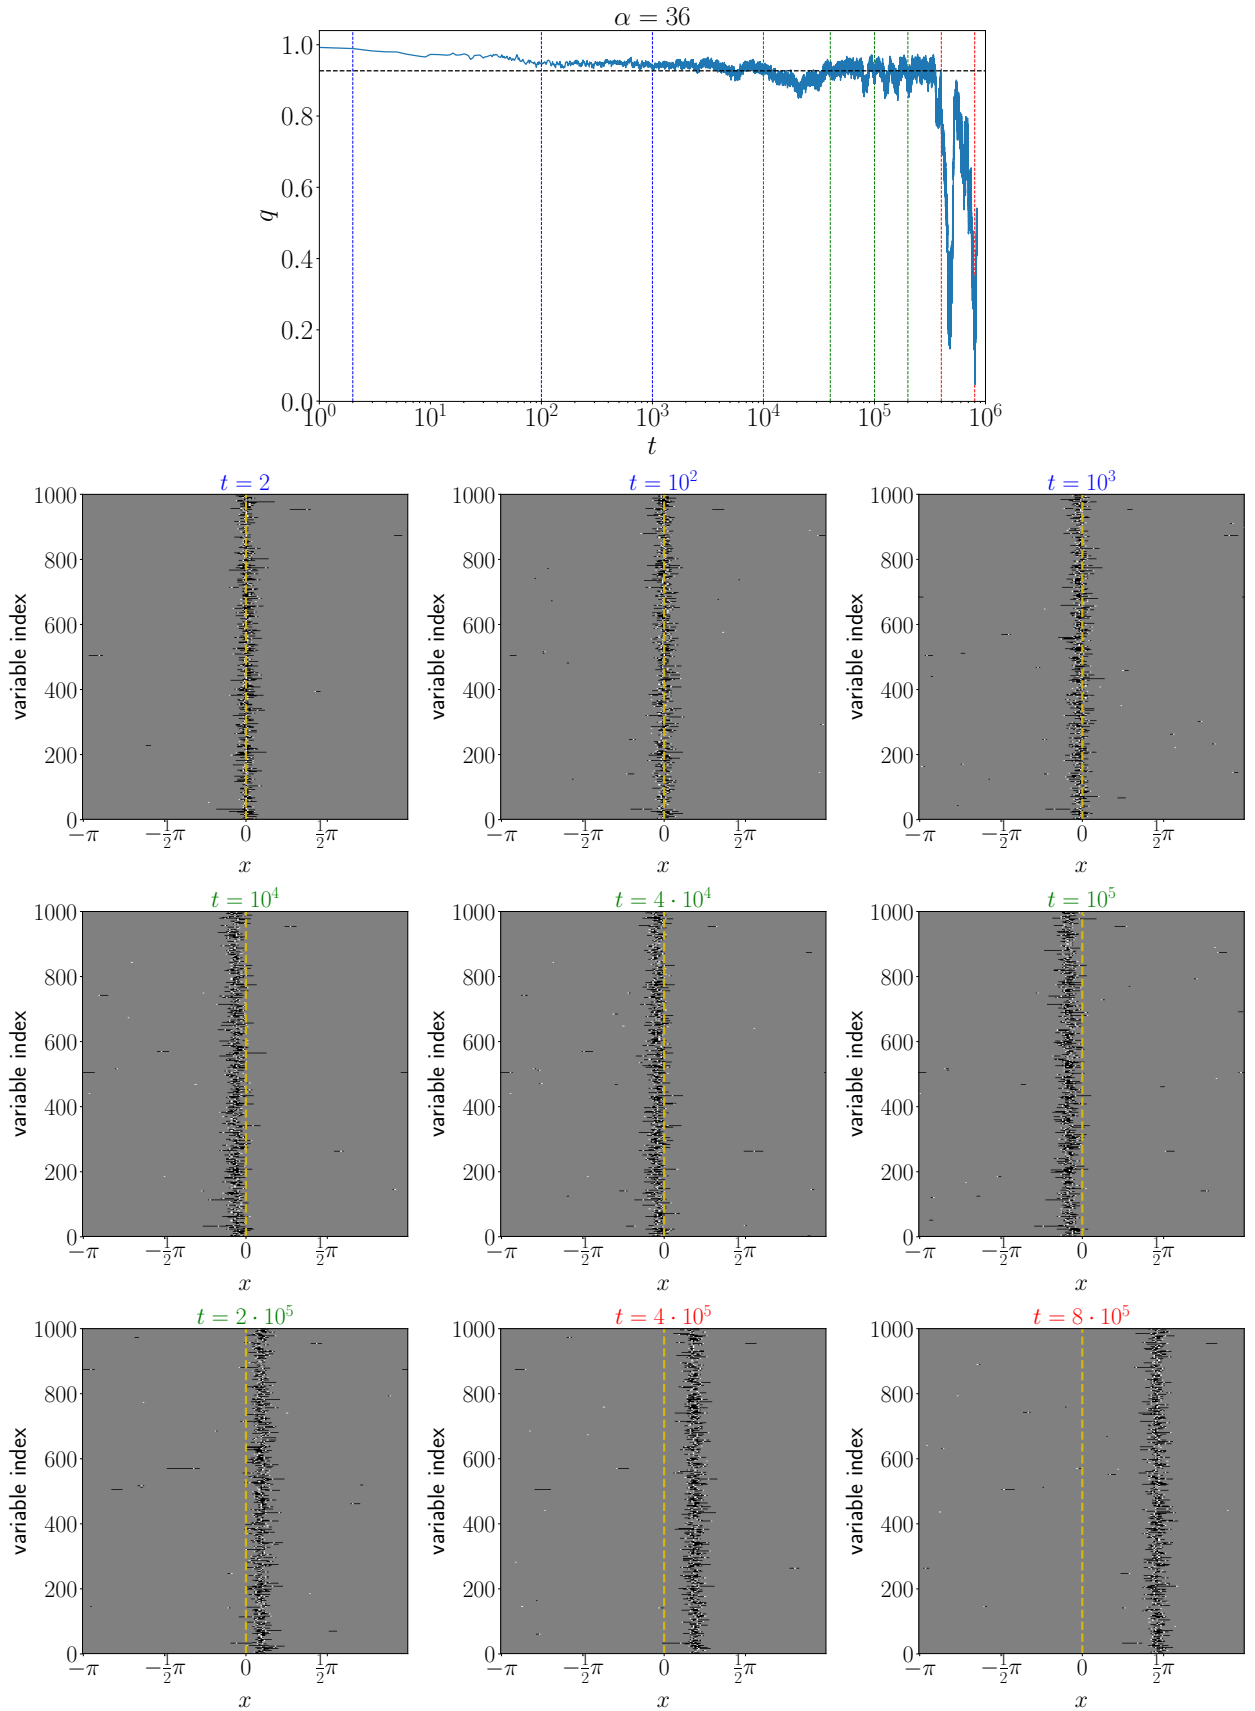


Figure 5.8: Top: $\alpha > \alpha_d$ decay of the “naive” overlap, defined as the plain dot product between different configurations, not accounting for the global rotational symmetry of the model. Bottom: a direct inspection on the state at time t of the system, initially planted in the configuration with all the variables equal to $x = 0$. (white: current configuration; black: values allowed by the neighbours; grey: non-allowed region)

Chapter 6

Postponing the dynamical transition

6.1 Biased potential

So far we have investigated our random CSP from the point of view of a purely hard-sphere interaction (in the limit $\beta \rightarrow \infty$): $\beta\mathcal{H}_{\text{flat}} \rightarrow \infty$ if some constraint is unsatisfied, $\beta\mathcal{H}_{\text{flat}} = 0$ otherwise. The associated Boltzmann-Gibbs distribution is hence *uniform* over the space of solutions to the CSP, and as a consequence the statistical weight associated to each cluster is simply proportional to their size. This situation can be generalized by introducing some bias in the probability measure over the solutions, namely by favouring some of them and disfavouring others, while leaving untouched the constraint definition. This can be viewed as a way to move some statistical weight to solutions belonging to clusters which are atypical, *i.e.* subdominant in the limit $N \rightarrow \infty$, with respect to the uniform measure. The characterization of the large deviation properties of the highly non-trivial phase space of random CSPs is indeed quite a fundamental question, since it is still not entirely clear to what extent equilibrium thresholds can be useful in order to understand and predict the behaviour of general algorithms looking for solutions, or even which are the correct threshold(s) to look at (to this subject is devoted the last part of this thesis).

This *biased thermodynamics* approach has been at the center of very recent investigations in the context of discrete random CSPs [6, 15, 18, 19]. Here we will follow the most simple strategy [18], which is to introduce a bias directly in the “soft part”, *i.e.* the one pertaining to satisfied constraints only, of the Hamiltonian $\mathcal{H}_{\text{flat}}$. The new equilibrium Boltzmann-Gibbs distribution can be written as

$$P(\{x_i\}) = \frac{1}{Z(\beta)} e^{-\beta[\mathcal{H}_{\text{flat}}(\{x_i\}) - \beta^{-1}H_1(\{x_i\})]}, \quad (6.1)$$

where the scaling β^{-1} on the soft part $H_1 \geq 0$, $H_1 = O(1)$ is introduced in order to have a finite limit for $\beta \rightarrow \infty$. Despite not being the most general definition, we will consider H_1 as independent of β : this is sufficient as long as one is mainly interested in the zero temperature limit.

The location of the SAT/UNSAT transition is by definition left untouched. However, the other transition thresholds are sensitive to the bias. This is well known in the case of hard sphere particle systems, where the hard core potential can be perturbed by adding short-range finite arbitrary potentials, such as square shoulders [32] or wells [33, 95], giving rise to very rich and unexpected phase diagrams. This is very relevant for soft matter colloidal systems, where such interaction




Figure 6.1: Dynamical transition line for a mean-field (infinite spatial dimensions) hard-sphere system with an additional square well short-range attraction of depth U_0 . On the x axis, $\hat{\varphi}$ is a rescaled packing fraction (playing a role similar to α in our CSP), while $\hat{T} = 1/U_0$. Notice that in this case the hard-core potential is chosen to be infinite even for finite temperatures, and the temperature can be adsorbed in the definition of the energy scale U_0 . The dotted vertical line represents the dynamical threshold in the absence of attraction. Reprinted from [69].

potentials can be experimentally engineered. We are particularly interested, for reasons that will be clarified in the following, to the case of a short-range attraction next to the hard core repulsion, which is known to usually determine a reentrancy phenomenon in the phase diagram, with a liquid cusp stretching into the glass phase as in figure 6.1.

The reentrancy of the liquid region can be further extended by considering more complicated biasing potentials. This was done in ref. [69] for infinite dimensional hard-spheres, where a generic function of the interparticle distance was approximated with n square wells (and shoulders) of vary-

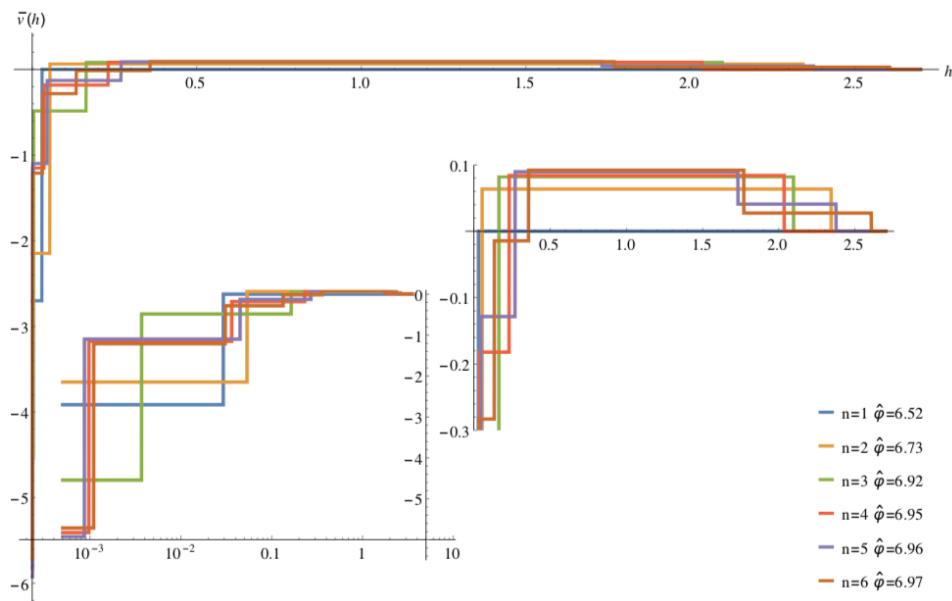


Figure 6.2: Optimized potentials, as a function of a rescaled interparticle gap h , maximising the dynamical packing fraction threshold $\hat{\varphi}_d$ for each value of the number of steps $n = 1, \dots, 6$ in the potential. This generalizes the single well bias of fig. 6.1. Reprinted from [69].

ing width and depth (height). By optimizing the dynamical transition (which is known analytically) via a gradient descent on the finite set of parameters, the generic potential of figure 6.2 was obtained, the main feature being a very close-range sharp attraction followed by a longer-ranged weak repulsion.

6.1.1 Optimized interaction for the Continuous Coloring

We perform a computation analogous in spirit to the one of ref. [69] in order to maximize the dynamical transition α_d . Unfortunately, we lack an analytical expression for α_d . A possible way round is to consider the gradient descent with respect to an observable which is hopefully related with the location of α_d : we consider the maximization of the complexity Σ , as explained in the next section. We will call the resulting optimized threshold α_d^{opt} , and the optimized interaction f^{opt} . We stress, however, that α_d^{opt} should be more safely interpreted as a lower bound to the actual




Figure 6.3: Soft part of the optimized interaction at $T = 0$ for the continuous coloring ($\theta = 2\pi/q$, with $q = 20$) as a function of the interparticle angular distance. The function is symmetric for $x \rightarrow 2\pi - x$. The difference $f^{\text{opt}} - f^{\text{flat}}$ does not go precisely to zero at large angles because we subtracted two normalized probabilities, both satisfying $\sum_a f_a = p$ in the discretized case. Inset: detail of the interaction.

optimal threshold, since there is still no rigorous proof linking the maximization of the dynamical transition with the one of Σ . On the other hand, the gradient of the complexity will be shown to be relatively simple to compute, and also accessible through Monte Carlo sampling. If the complexity-maximization approach actually proved to be exact, then we would have in our hands a very powerful and versatile tool to compute the optimized interaction also in more computer-memory consuming settings, such as in higher spatial dimensions, where BP really struggle. We leave a more detailed discussion on the conceptual and technical aspects of this procedure to section 6.2.

The optimized interaction f^{opt} is given in figure 6.3 for the usual choice of $q = 20$ and a discretization of precision $d = 10$ ($p = 200$ clock-states). The interaction f is in our setting proportional to the Boltzmann-Gibbs weight, rather than to an energy, for this reason a peak in the probability for $x = \theta$ is interpreted as a very short range attraction favouring contacts between angles/particles. With this choice for the potential, the zero temperature dynamic threshold




Figure 6.4: Transition lines in the continuous coloring with $q = 20$ and discretization $d = 10$, for both the choices of the interaction f^{flat} (solid) and f^{opt} (dashed). The dynamical transition is moved from $\alpha_d^{\text{flat}} = 34.63(2)$ to $\alpha_d^{\text{opt}} = 37.71(1)$.

estimated from BP is moved from $\alpha_d^{\text{flat}} = 34.63(2)$ to $\alpha_d^{\text{opt}} = 37.71(1)$.

The numerical optimization has been performed for $T = 0$, when the hard part of the interaction f (for $|x| < \theta$) is exactly zero. By taking the logarithm of the soft part of f (for $\theta \leq x \leq 2\pi - \theta$), we can obtain the pairwise contribution to the biasing Hamiltonian H_1 up to an additive constant, due to the arbitrariness in the normalization. When switching on temperature, the hard core excluded volume interaction is softened, the cost for violating a constraint being now proportional to $\exp(-\beta)$. The soft part of the interaction should instead be independent from β in our approach, apart for a global β -dependent normalization. In the following, we will always adopt implicitly the definition $f_\beta^{\text{opt}}(x) \propto Ae^{-\beta}\mathbb{I}[|x| < \theta] + f_{T=0}^{\text{opt}}\mathbb{I}[\theta \leq x \leq 2\pi - \theta]$. We fix the relative amplitude A of the hard part with respect to the soft part of the interaction f by requiring

$$\frac{\sum_{(-\theta; \theta)} f_a^{\text{opt}}}{\sum_{[\theta; 2\pi - \theta]} f_a^{\text{opt}}} = \frac{\sum_{(-\theta; \theta)} f_a^{\text{flat}}}{\sum_{[\theta; 2\pi - \theta]} f_a^{\text{flat}}} = \frac{(2d-1)e^{-\beta}}{p - (2d-1)}. \quad (6.2)$$

With this choice, the annealed energy $e_{\text{ann}}(\beta)$, counting the number of violated constraints per spin as a function of temperature in the paramagnetic phase, is equal both for f^{flat} and f^{opt} to

$$e_{\text{ann}}(\beta) \equiv \frac{M \sum_a f_a \mathbb{I}[\cos(\frac{2\pi a}{p}) > 0]}{N \sum_a f_a} = \frac{\alpha(2d-1)e^{-\beta}}{(2d-1)e^{-\beta} + p - (2d-1)}. \quad (6.3)$$

The phase diagram in the (α, T) plane is shown in figure 6.4, solid lines corresponding to f^{flat} , while dashed ones referring to f^{opt} . We can notice it is quite different from the one in figure 6.1. This is expected: in our model, temperature has a completely different meaning, since it tunes the strength of the constraints (we recover the infinite hard sphere potential only in the limit $\beta \rightarrow \infty$). If we were to interpret the diagram of 6.1 in our setting, we would recognize that the entire $(\hat{\varphi}, \hat{T})$ plane corresponds in fact to $T = 0$ (the hard constraint is never released), while the role of \hat{T} is

simply to parameterize different potentials (in the case of a single well of fixed width, we need only one parameter U_0 , which can adsorb also the physical temperature). In conclusion, we do not seek nor expect the complex reentrancy phenomenon typical of hard sphere like physical systems as in figure 6.1, but we rather place ourselves on the tip of the reentrancy and we explore another dimension, namely the one in which unsatisfied constraints are accepted with some small probability.

Finally, it is interesting to notice that also the Kesten-Stigum bound α_{KS} is sensitive to the modification of the interaction, and it actually increases the most with respect to the other thresholds, going from the $T = 0$ value $\alpha_{KS} = 46.743$ to $\alpha_{KS} = 77.972$, with an increase of more than 60%. The location of α_{KS} is particularly relevant when considering the associated Bayes optimal inference problem, since it signals the α -point below which inference is practically unfeasible. We conclude that, in order to improve the performance of optimization on the one hand or inference protocols on the other hand, one should optimize the interaction in opposite directions.

6.1.2 Discussion: the role of competing interactions

The two main features of the optimized potentials in figure 6.2 are, as we said, an extremely short-range and strong attraction followed by a much weaker, longer-ranged repulsion. How to disentangle the two contributions is not in general completely clear. The physical understanding is that the attraction is the key ingredient, while the repulsion, by effectively deepening the well, can contribute to further increase the liquid reentrance (the improvement over a purely short-range attraction for hard spheres in the limit $d \rightarrow \infty$ is about 3%). Very recently, the authors of [26] wondered to

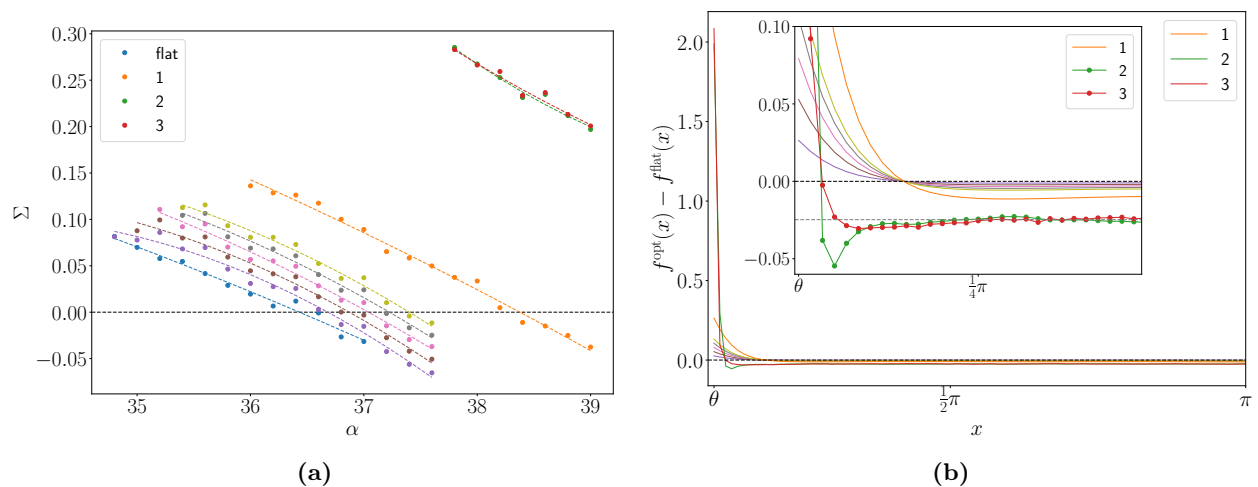


Figure 6.5: Left: complexity curves as a function of α for different interactions (depicted on the right panel) along the optimization procedure. Dotted lines are meant only to guide the eye. Curve number one (orange) is obtained after the very first step in the gradient descent starting from f^{flat} (blue). Curves in between are obtained by using a family of functions f' interpolating between f_{flat} and f_1 , *i.e.* $f' \propto f^{\text{flat}} + a(f_1 - f^{\text{flat}})$, for five values of $0 < a < 1$. Curves 2 and 3 refer respectively to an intermediate function along the procedure and to the result of our optimization f^{opt} .

what extent this result is still valid in finite dimensions. They directly verified that in a 3- d system the location of the (avoided) mode coupling transition can be moved by tuning the attraction, but it cannot be further improved by adding a suitable repulsion (at most, they identify a branch of parameters over which the optimum extends). A legitimate question would be to understand our model at hand which situation belongs or is closer to, hoping that this could also provide some more

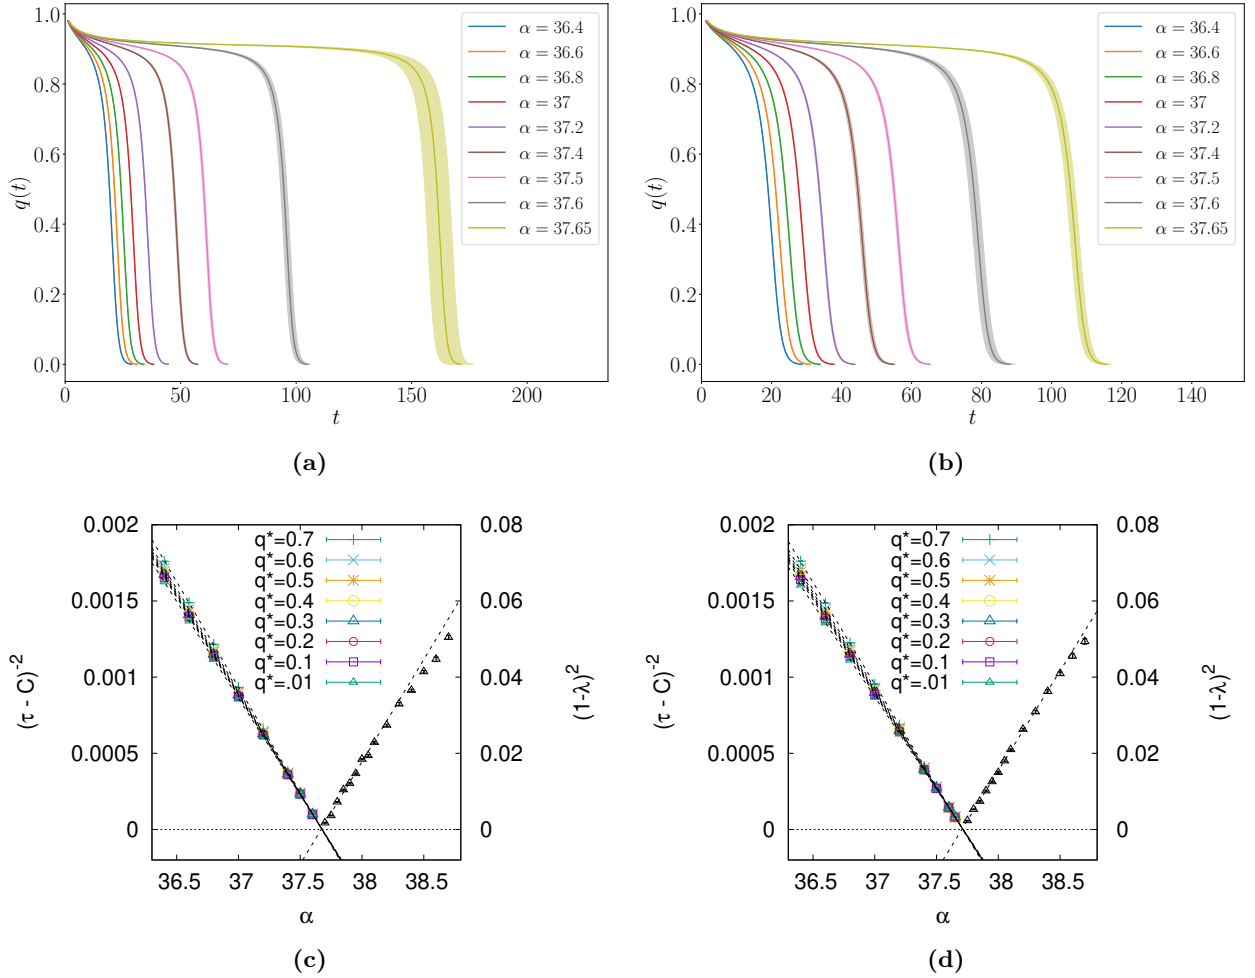


Figure 6.6: Estimation of α_d from BP population dynamics simulations using interactions 2 (left) and 3 (right) of figure 6.5, the latter corresponding to f^{opt} . Bottom figures (colors): divergence of the overlap relaxation time $\tau_{q^*}(\alpha)$, defined as the time needed for $q(t)$ to decrease down to q^* . The data is fitted according to eq. (5.4). (Black): behavior of $(1 - \lambda(\alpha))^2$, where $\lambda(\alpha)$ is the stability parameter of the non-trivial BP solution for $\alpha > \alpha_d$. The dynamical transition α_d is located at the intersection of the color and black lines with the x-axis.

general intuition about the role of competing interactions. Unfortunately, it would appear to be rather difficult to extract a definitive answer from our data, as we are going to discuss here below.

Coming back to the model at hand, by looking at figure 6.3, the short-range attraction can be immediately recognized in the peak of f , and yet also a hint of a very weak repulsion can be discriminated. The identification is somehow complicated not only from the weakness of the contribution, but also from its not so extended range (one should keep in mind, however, that the domain is in this case compact). In our numerical strategy, we pursued a discretization of the continuous, bounded domain into $p = 200$ states. As figure 6.3 makes pretty clear, though, the most relevant part of interaction, comprising both the attraction and the repulsion, is confined to a limited region of extension approximately $\pi/5 = 2\theta$ around $\theta = 2\pi/20$, *i.e.* containing about twice the number of discrete states in the excluded region $x < \theta$, which is anyway a small number (~ 20) compared to p . Most of the peak is even contained in the very first bin of extension $2\pi/p$, suggesting the possibility of a delta interaction in the continuous limit. This makes the determination of the details of the interaction around the peak a very delicate matter.

q^*	α_d	B	C	χ^2/dof	α_d	B	C	χ^2/dof
0.7	37.676(1)	26.9(1)	-6.85(7)	3.66	37.712(1)	27.3(1)	-7.06(8)	1.76
0.6	37.677(1)	27.3(1)	-5.90(7)	2.83	37.713(1)	27.7(1)	-6.11(9)	1.21
0.5	37.677(1)	27.5(1)	-5.09(7)	2.76	37.714(1)	27.9(1)	-5.32(9)	1.68
0.4	37.677(1)	27.6(1)	-4.36(7)	3.00	37.714(1)	28.1(1)	-4.65(9)	2.46
0.3	37.677(1)	27.7(1)	-3.67(7)	2.84	37.715(1)	28.2(1)	-3.96(9)	2.47
0.2	37.678(1)	27.7(1)	-2.88(7)	2.98	37.715(1)	28.2(1)	-3.13(9)	2.79
0.1	37.678(1)	27.8(1)	-1.79(7)	2.34	37.715(1)	28.4(1)	-2.07(9)	2.80
0.01	37.678(1)	28.1(1)	1.11(7)	4.44	37.716(1)	28.7(1)	0.82(8)	3.10

Table 6.1: Results from fit (5.4) for the divergence of the overlap relaxation time (color points in the bottom parts of fig. 6.6).

α_d	χ^2/dof	α_d	χ^2/dof
37.669(1)	15.52	37.705(1)	2.60

Table 6.2: Results from the fit on the stability parameter λ .

However, a curious phenomenon takes place during our optimization, see figure 6.5b, namely the fact that at intermediate steps in the procedure (green line in the figure) the interaction shows a repulsion term being more pronounced than in the final function which we call f^{opt} (red line). Notice that for the former (green) function the limits of the discretization precision are evident in the repulsive part of the interaction. As displayed in the left panel, an approximate computation of the complexity is not enough to distinguish the performance of the two interactions. For this reason, we repeat the same analysis of chapter 5 in order to obtain an accurate estimate of α_d for the two potentials, see figure 6.6 and the results in tables 6.1-6.2. From this we conclude that differences are present, although very little compared to the effects due to the leading attractive peak. As a result, any strong claim about the role of competing interactions in the present setting would risk to be utterly unsupported, since as we said a) the repulsive term of f^{opt} in figure 6.3 is very weak, b) the dynamic threshold α_d seems as well to weakly depend on the details of the repulsive part of the interaction and c) the very determination of these fine details is on the other hand limited by effect of discretization.

6.2 Complexity maximization

The “typical” complexity Σ (*i.e.* the one associated to dominant states) is an observable that can be estimated via BP as a difference of free entropies, see eq. (5.7). The behaviour of Σ as a function of α , as shown in figure 5.6, is the following: $\Sigma = 0$ below α_d , then Σ has a jump at α_d to a finite (maximum) value, then it is a monotonically decreasing function of α up to α_c , where $\Sigma = 0$. One may then wonder if there is a way to postpone the dynamical threshold α_d by directly acting on the shape of the complexity. A strategy that we found to be successful¹ is to maximize the value of the complexity close to its maximum at α_d : the idea is that by shifting the complexity curve to higher values of α , also the value of Σ at fixed $\alpha > \alpha_d$ will increase, see figure 6.5a. However, since the maximum value of the complexity can take any strictly positive value, there is no reason to exclude the existence of particular functions f for which the complexity curve is still postponed, but it is also lowered. In this sense, we lack a physical intuition (not to mention a rigorous proof)

¹Based on a joint work with Thibault Lesieur.

supporting the fact that the interaction f (not to be confused with a free energy) for which $\alpha_d(f)$ has its maximum should coincide with the function f for which also $\Sigma(\alpha_d(f); f)$ has its maximum.

Having said that, this method is the best chance to our knowledge to automatically optimize the potential, and it yields promising results. We compared it with other strategies, such as the maximization of the stability parameter λ of the BP non-trivial solution, which is $\lambda = 1$ at the dynamical transition and then decreases for $\alpha > \alpha_d$ (the solution being more stable). Since it takes a fixed value at α_d for any f , then the maximization procedure is in this case justified. Another strategy is the minimization of the so called Kullback-Leibler divergence, measuring the distance of the non-trivial distribution of messages $\mathbb{P}^{\text{eq}}(\nu)$ from the paramagnetic one, which in the continuous notation reads

$$\mathbb{E}_{\mathbb{P}^{\text{eq}}(\nu)} [D_{\text{KL}}(\nu^{\text{para}} || \nu)] = \mathbb{E}_{\mathbb{P}^{\text{eq}}(\nu)} \left[-\frac{1}{2\pi} \int_0^{2\pi} dx \log \left(\frac{\nu(x)}{(2\pi)^{-1}} \right) \right]. \quad (6.4)$$

For $\alpha < \alpha_d$ this quantity is zero, since the only solution is $\mathbb{P}^{\text{eq}} = \mathbb{P}^{\text{para}}$, it takes a non-trivial value at α_d and then increases for $\alpha > \alpha_d$. Following the same line of thought as in the case of the complexity, one should minimize it. Furthermore, the minimization of this quantity has a nice physical implication: since (6.4) increases the more the fixed point messages depart from the uniform solution, this implies that solutions containing frozen fields (which are highly irregular distributions) are penalized by the procedure. However, the numerical estimation of the gradient of λ and of the Kullback-Leibler divergence is much more involving than computing $\partial\Sigma/\partial f$, which reduces to computing the derivative of a free energy. In particular, the λ estimator becomes very noisy as the optimization proceeds, while the Kullback-Leibler divergence appears to decrease the condensation threshold α_c , making convergence difficult near the tri-critical point. We did not further investigate this phenomenon.

In the end, the best result is obtained through the complexity maximization. Even if its interpretation is still not completely clear to our eyes, it has some practical advantages. In particular, this method does not rely exclusively on BP. Being a derivative of a difference of free energies, the gradient $\partial\Sigma/\partial f$ can be computed also as some two-point correlation function, *e.g.* via Monte Carlo sampling, as outlined at the end of the next section. Moreover, the method is completely general, and it can in principle be applied also to the mean-field hard-spheres system studied in [69], where an exact method can be independently developed, hopefully contributing to shed some light on the physics behind the complexity maximization.

The maximization is implemented through a gradient descent: we move in the space of f (subjected to the normalization constraint $\sum_a f_a = p$) along the gradient $\partial\Sigma/\partial f$ at fixed $\alpha = \alpha_d + \delta\alpha$, compute the new α_d and repeat the procedure until stationarity is reached. Thanks to discretization, the gradient descent in the space of functions f can be numerically treated as an optimization in the finite set of parameters $\{f_a\}$. The procedure is the following

$$f_a^{(t+1)} = f_a^{(t)} + \mu \left(\frac{\partial\Sigma(\alpha_d(f^{(t)}) + \delta\alpha; f^{(t)})}{\partial f_a} + \Lambda \right) \mathbb{I}(\cos(2\pi a/p) \leq \cos(2\pi/q)), \quad (6.5)$$

where $\Lambda = -[p - (2d-1)]^{-1} \sum_{a=d}^{p-(d-1)} \partial\Sigma^{(t)}/\partial f_a$ is a Lagrange multiplier that ensures that $\{f_a^{(t+1)}\}$ is still normalized, and μ is the speed rate of the gradient descent. We found empirically that




Figure 6.7: Gradient of the complexity for different steps along the optimization process, computed through BP simulations for the model with $q = 20$ and discretization $d = 10$.

computing the gradient close to the dynamical transition results in less noise. In all the computations we have used $\delta\alpha = 0.5$. As the optimization runs (see fig. 6.7), the support of the gradient will concentrate on the forbidden values that encode the constraints of the problem, until the point, at the end of the optimization procedure, for which the only way to further increase the complexity will be to relax the constraints, that we forbid thanks to the factor $\mathbb{I}(\cos(2\pi(a-1)/p) \leq \cos(2\pi/q))$ in (6.5).

Finally, as also α_{KS} is postponed when α_d is maximized, one may wonder if by maximizing the stability of the paramagnetic solution at a given α one could end up also pushing α_d to the right. We cannot rule out this possibility, even though it does not appear to be a general strategy in our opinion. In particular, there are systems for which the paramagnetic solution is always stable (α_{KS} is virtually infinite), and as a consequence this method is ill defined. On a more general ground, it is true that the point-to-set correlations emerging at α_d can be realized in an exponential number of ways (states), and thus they can also be destroyed in many nonequivalent directions. It is essentially this freedom that one exploits in order to avoid “early” clusters and postpone as much as possible α_d . In correspondence of α_{KS} , on the contrary, symmetry is broken towards a finite number of states. If it is reasonable to expect that destroying any correlations close to α_d may as well result in a weakening of fluctuations out of the paramagnetic state at α_{KS} , it is not at all obvious to us how stabilizing the system towards these fluctuations could in general select the exact direction that destroys early clusters among the others, up to the point of really *maximizing* α_d .

6.2.1 Computation of the gradient of the complexity

The complexity we want to maximize is defined as

$$\Sigma = \frac{1}{N} [F_{\text{Rec II}}(\mathbb{P}^{\text{Para}}, f) - F_{\text{Rec II}}(\mathbb{P}^{\text{eq}}, f)], \quad (6.6)$$

where $\mathbb{P}^{\text{Para}}(\nu) = \prod_{a=1}^p \delta(\nu_a - n)$, $n = \frac{1}{p}$, is the paramagnetic fixed point and \mathbb{P}^{eq} is the non-trivial one arising for $\alpha > \alpha_d(f)$. In our discretized setting, $f = \{f_a\}$ is a symmetric function under

$a \rightarrow p - a$ and periodic of period p . The population dynamics definition of the free entropy is given by eq. (4.43), which we recall here for convenience as

$$\begin{aligned} \frac{1}{N} F_{\text{Rec II}}(\mathbb{P}, f^{\text{Fixed}}, f^{\text{State}}) &= \sum_d P_{\text{Node}}(d) \int \prod_{k=1}^d \left[d\nu_k \sum_b f_b^{\text{State}} n \mathbb{P}(T(-b)\nu_k) \right] \log \left[Z(\{\nu_k\}_d, f^{\text{Fixed}}) \right] \\ &\quad - \alpha \int d\nu_1 d\nu_2 \mathbb{P}(\nu_1) \sum_b f_b^{\text{State}} n \mathbb{P}(T(-b)\nu_2) \log \left[\sum_{r,s} \nu_{1,r} f_{s-r}^{\text{Fixed}} \nu_{2,s} \right]. \end{aligned} \quad (6.7)$$

The distinction between f^{State} and f^{Fixed} is here only introduced in order to discuss separately the different terms coming from the derivation with respect to f_a . Of course in the end one has to evaluate every derivative at $f^{\text{State}} = f^{\text{Fixed}} = f$, with $\sum_a n f_a = 1$, $n = \frac{1}{p}$. Finally, we also recall that

$$Z(\{\nu_k\}_d, f) = \sum_{a=0}^{p-1} n \prod_{k=1}^d \sum_{b=0}^{p-1} f_{b-a} \nu_{k,b}. \quad (6.8)$$

The derivative $\partial \Sigma / \partial f_a$ contains terms coming from the explicit dependence of Σ from both f^{State} and f^{Fixed} , but also from the change in the non-trivial fixed point distribution \mathbb{P}^{eq} . The paramagnetic fixed point distribution, on the contrary, is given by $\mathbb{P}^{\text{Para}} \equiv \prod_a \delta(\nu_a - n)$ for any choice of f satisfying $\sum_a f_a n = 1$, as can be immediately noticed by a direct inspection of the BP fixed point equation (4.41). We can thus write

$$\begin{aligned} \frac{N \partial \Sigma}{\partial f_a} &= \frac{\partial F_{\text{Rec II}}(\mathbb{P}^{\text{Para}}, f^{\text{Fixed}}, f^{\text{State}})}{\partial f_a^{\text{Fixed}}} - \frac{\partial F_{\text{Rec II}}(\mathbb{P}^{\text{eq}}, f^{\text{Fixed}}, f^{\text{State}})}{\partial f_a^{\text{Fixed}}} \\ &\quad + \frac{\partial F_{\text{Rec II}}(\mathbb{P}^{\text{Para}}, f^{\text{Fixed}}, f^{\text{State}})}{\partial f_a^{\text{State}}} - \frac{\partial F_{\text{Rec II}}(\mathbb{P}^{\text{eq}}, f^{\text{Fixed}}, f^{\text{State}})}{\partial f_a^{\text{State}}} \\ &\quad - \frac{\partial F_{\text{Rec II}}(\mathbb{P}^{\text{eq}}, f^{\text{Fixed}}, f^{\text{State}})}{\partial \mathbb{P}} \left[\frac{\partial \mathbb{P}^{\text{eq}}}{\partial f_a^{\text{Fixed}}} + \frac{\partial \mathbb{P}^{\text{eq}}}{\partial f_a^{\text{State}}} \right]. \end{aligned} \quad (6.9)$$

Most of the terms either simplify or are equal to zero. The first two terms cancel, as we are going to comment in a moment. The third term is trivially equal to zero, since both the logarithms in eq. (6.7) vanish when one considers $\nu_a = n$ and the fact that $n f_a^{\text{Fixed}}$ is correctly normalized. The fifth term is zero since the distribution \mathbb{P}^{eq} extremizes the Bethe free entropy. In the end, the only remaining term is

$$\frac{\partial \Sigma}{\partial f_a} = - \left. \frac{\partial N^{-1} F_{\text{Rec II}}(\mathbb{P}^{\text{eq}}, f^{\text{Fixed}}, f^{\text{State}})}{\partial f_a^{\text{State}}} \right|_{f^{\text{State}}=f^{\text{Fixed}}=f}. \quad (6.10)$$

Before proceeding with the computation, let us show that the first two terms indeed cancel. The derivative $\partial N^{-1} F_{\text{Rec II}} / \partial f_a^{\text{Fixed}}$, for generic \mathbb{P} , can be written after some manipulations² as

$$\left. \frac{\partial N^{-1} F_{\text{Rec II}}}{\partial f_a^{\text{Fixed}}} \right|_f = \alpha \int d\nu_1 d\nu_2 \mathbb{P}(\nu_1) \sum_b f_b^{\text{State}} n \mathbb{P}(T(-b)\nu_2) \frac{\sum_r \nu_{1,r} \nu_{2,r+a}}{\sum_{r,s} \nu_{1,r} f_{s-r} \nu_{2,s}}. \quad (6.11)$$

²One can use the decomposition (6.13) in order to recast the derivative of the first term in the rhs of (6.7) in the form reported in the text (the derivative of the second term is trivially proportional and they can be added).

If we multiply and divide by the same amount f_a , we can make a term $\nu_{1,r} f_a \nu_{2,r+a}$ appear, which is proportional by the definition of messages (see also eq. (3.10)) to the marginal probability $p(x_1 = r, x_2 = r + a)$, as a function of r , of two neighbouring variables at a fixed distance a . The whole integral thus represents the total equilibrium probability for two neighbours to be at distance a , which is essentially the pair correlation function of the system. In the entire range $\alpha < \alpha_c$, though, the pair correlation function of the paramagnetic and of the planted fixed point is the same and is simply equal (whenever short loops are absent, as in the related MK model [25]) to the normalized interaction $n f_a \propto e^{-\beta v(2\pi a/p)}$, where $v(x)$ is the pairwise interaction potential. Simplifying the extra $1/f_a$ factor, we obtain that the first two terms in eq. (6.9) cancel as they are both equal to a constant α/p . As a final remark, this derivation is clearly valid for any $T \neq 0$, when also $f_a \neq 0 \forall a$. The case $f_a = 0$ can nevertheless be recovered from the previous one by taking the limit $T \rightarrow 0$.

Coming back to our computation, plugging into eq. (6.10) the definition of F_{RecII} given by eq. (6.7), we obtain

$$\begin{aligned} \frac{\partial \Sigma}{\partial f_a} &= \alpha n \int d\nu_1 d\nu_2 \mathbb{P}^{\text{eq}}(\nu_1) \mathbb{P}^{\text{eq}}(T(-a)\nu_2) \log \left[\sum_{r,s} \nu_{1,r} f_{s-r} \nu_{2,s} \right] + \\ &- 2\alpha n \sum_d P_{\text{Node}}(d) \int d\nu_{d+1} \mathbb{P}^{\text{eq}}(T(-a)\nu_{d+1}) \prod_{k=1}^d \left[d\nu_k \sum_b f_b n \mathbb{P}^{\text{eq}}(T(-b)\nu_k) \right] \log [Z(\{\nu_k\}_d \cup \nu_{d+1}, f)], \end{aligned} \quad (6.12)$$

where we have used the fact that $d \cdot P_{\text{Node}}(d) = 2\alpha P_{\text{Node}}(d-1)$ and replaced a dummy index d in the second term with $d+1$. The quantity $\log(Z)$ appearing in the previous equation can be iteratively expressed in terms of the update function Φ^{update} given by eq. (4.26) as

$$\begin{aligned} \log [Z(\{\nu_k\}_d \cup \nu_{d+1}, f)] &= \log \left[\sum_a n \prod_{1 \leq k \leq d+1} \left(\sum_b f_{b-a} \nu_{k,b} \right) \right] = \\ &= \log \left[\sum_a \Phi_a^{\text{update}}(\{\nu_k\}_d, f) Z(\{\nu_k\}_d, f) \sum_b f_{b-a} \nu_{d+1,b} \right] = \\ &= \log [Z(\{\nu_k\}_d, f)] + \log \left[\sum_{ab} \Phi_a^{\text{update}}(\{\nu_k\}_d, f) f_{b-a} \nu_{d+1,b} \right]. \end{aligned} \quad (6.13)$$

Plugging (6.13) into (6.12), one recognizes³ that Φ^{update} can be interpreted as a random message ν distributed according to $\mathbb{P}^{\text{eq}}(\nu)$. The second term in (6.13) has then the same form of the first term in (6.12) and they can be directly added. Therefore one obtains

$$\begin{aligned} \frac{\partial \Sigma}{\partial f_a} &= -\alpha n \int d\nu_1 d\nu_2 \mathbb{P}^{\text{eq}}(\nu_1) \mathbb{P}^{\text{eq}}(T(-a)\nu_2) \log \left[\sum_{r,s} \nu_{1,r} f_{s-r} \nu_{2,s} \right] + \\ &- 2\alpha n \sum_d P_{\text{Node}}(d) \int \prod_{k=1}^d \left[d\nu_k \sum_b f_b n \mathbb{P}^{\text{eq}}(T(-b)\nu_k) \right] \log [Z(\{\nu_k\}_d, f)]. \end{aligned} \quad (6.14)$$

We notice that the second term does not depend on the index a . This term is trivially constant

³Formally, one can multiply (6.12) by a factor $1 = \int d\nu \delta(\nu - \Phi_a^{\text{update}}(\{\nu_k\}_d, f))$ in order to make appear the BP expression (4.41) for $\mathbb{P}^{\text{eq}}(\nu)$.

and can hence be disregarded, since at each step we already enforce the normalization condition $\sum_a f_a n = 1$. Finally one gets

$$\frac{\partial \Sigma}{\partial f_a} = -\frac{\alpha}{p} \int d\nu_1 d\nu_2 \mathbb{P}^{\text{eq}}(\nu_1) \mathbb{P}^{\text{eq}}(T(-a)\nu_2) \log \left[\sum_{rs} \nu_{1,r} f_{s-r} \nu_{2,s} \right] \quad (6.15)$$

Remarkably, this formula can also be estimated using Monte Carlo sampling. This is very convenient, since it might be the only numerically feasible approach in the case of higher spatial dimensions and/or high discretization. To this purpose, one needs to generate a planted graph A along with a planted configuration, noted as $\{x_i^0\}_{1 \leq i \leq N}$, and then run a Monte Carlo algorithm initialized in the planted configuration to create samples. Let us assume we have generated T samples $\{x_i^t\}_{1 \leq i \leq N, 1 \leq t \leq T} \in \{0, \dots, p-1\}^{N \times T}$ of the system. The two cavity messages ν_1 and ν_2 in eq. (6.15) are random variables drawn independently from $\mathbb{P}^{\text{eq}}(\nu)$. We can use at this point the fact that the distribution of cavity messages coincides, for the Erdős-Rényi ensemble, also with the distribution of single-variable local marginal probabilities for the complete graph. This allows one to recast an average over ν_i of the kind of $\sum_a \nu_{i,a} g(a)$ inside the integral in eq. (6.15), where g is a generic function, as the time average over the generated samples $\{x_i^t\}_{1 \leq t \leq T}$, for each fixed site i (eventually averaging over the choice of site i): $\frac{1}{T} \sum_{t=1}^T g(x_i^t)$.

Finally, notice that the presence of the shift $\mathbb{P}^{\text{eq}}(T(-a)\nu_2)$ translates into considering for ν_2 a new set of generated samples given by $\{x_i^t + a\}_{1 \leq i \leq N, 1 \leq t \leq T}$. The gradient of the complexity is then estimated using the following formula

$$\frac{\partial \Sigma}{\partial f_a} = -\frac{\alpha}{p} \frac{1}{U} \sum_{(i,j)} \log \left[\frac{1}{T^2} \sum_{1 \leq t_1, t_2 \leq T} f_{(x_j^{t_2} - x_j^0 + a) - (x_i^{t_1} - x_i^0)} \right], \quad (6.16)$$

where U is the number of pairs (i, j) in the sum. One can take $U = N$, by running one time over the index $i = 1, \dots, N$, and considering $j = i$ to belong to a second different system (since we want ν_1 and ν_2 to be extracted independently). Since the MC simulations are subject to coherent rotations of the spin variables due to the global rotational symmetry of the model, which becomes continuous in the $p \rightarrow \infty$ limit, one should also maximize at each step the overlap with the planted configuration over a global rotation of the system.

Chapter 7

Algorithmic behaviour

7.1 *Smart* algorithms beating the threshold

The phase diagram of random constraint satisfaction problems, as already briefly discussed in section 3.2.1 (see in particular fig. 3.2), is animated by the presence of multiple thresholds preceding the SAT/UNSAT transition: these are the clustering or dynamical threshold α_d , the condensation or Kauzmann threshold α_c , the rigidity threshold α_r and the freezing threshold α_f . Apart from particular cases in which some of them may coincide, as it happens when the transition is continuous and $\alpha_d = \alpha_c$ (for instance, in the 3-coloring of Erdős–Rényi random graphs [101], one has $\alpha_d = \alpha_c = 2$), or when $\alpha_c = \alpha_{\text{sat}}$ as in the XORSAT [93] (this comes from the fact that all the clusters of solutions have the same internal entropy due to the linear structure of the problem, so that they disappear altogether at α_c), these thresholds are generally distinct. Anyhow, the following relations hold by definition: $\alpha_d \leq \alpha_c \leq \alpha_{\text{sat}}$ and $\alpha_d \leq \alpha_r \leq \alpha_f \leq \alpha_{\text{sat}}$. The ordering of α_r and α_f with respect of α_c is on the other hand a priori undefined, depending on the specific parameters of the model (such as the number of colors q or the function node degree k in k -SAT).

Despite such a variety of well characterized thermodynamic thresholds, the comprehension of their consequences on the performance of algorithms searching for satisfied configurations is still poorly understood. A theoretical lower bound to the onset of computational hardness is given by the dynamic transition threshold α_d . Above α_d , simple local Monte Carlo (MC) procedures, such as simulated annealing (SA), are expected to get trapped at a threshold energy level $e_{\text{th}} > 0$, due to the proliferation of metastable states. One can alternatively devise annealing strategies working exactly at zero energy, as for instance using the MC update proposed in section 5.1.2, and where the external tuning parameter is the number of constraints α instead of temperature. Also in this case, the procedure is expected to fail approaching α_d due to the divergence of MC equilibration time, which prevents the system to actually find its way inside a cluster¹.

Why then is α_d considered only as a lower bound to the onset of hardness? The answer to this question lies in the following “experimental” observation: the dynamical threshold α_d appears to be irrelevant to the behaviour of most local algorithms, that are typically capable of finding

¹Generally, with the locution “equilibration time” one refers to the time needed to restore ergodicity, *i.e.* the time needed to jump from cluster to cluster in a uniform manner. However, it is also true that, if the time to find a way out from a cluster diverges, the same should happen to the time for entering a cluster from the outside. For this reason, the “planting trick” was needed in order to “equilibrate” (here in an ergodicity broken sense) the system inside a typical (but fictitious) cluster.




Figure 7.1: Uncontrollably *smart*: even naive Monte Carlo based algorithms can find solutions beyond α_d . We plot the number of MC sweeps $t_{\Delta\alpha}$ that is necessary in order to increase the number of links in the graph by $\Delta M = \Delta\alpha N$ without spoiling a solution, with $\Delta\alpha = 0.2$ a fixed smoothing parameter. See section 7.2.1 for a precise definition of the algorithm.

solutions with easiness even beyond α_d [4, 59, 96]. This circumstance is not surprising by itself, since many different algorithms and heuristic solvers do exist not satisfying detailed balance, and for which the location of α_d assumes in principle no special meaning. To this class belong the best-performing algorithms, such as survey propagation (SP) guided decimation [16, 17, 82], which is based on a message passing strategy. However, even in the case of simple MC procedures, where the algorithmic threshold is naively expected to coincide with α_d , this thermodynamic prediction is actually overridden, see figure 7.1. We refer to such simple algorithms succeeding beyond α_d , despite the energy landscape becoming very complex and rich of fatal traps due to the mean-field nature of the problems, as being *smart*.

A lot of work has been devoted to the characterization of this smartness in terms of the underlying visited energy landscape and of the structure of the space of solutions. The current understanding is that, even though the properties of typical solutions are correctly described by standard thermodynamics, the fact that arbitrary out-of-equilibrium procedures should always end up finding typical solutions is not guaranteed at all: this encourages the study of large deviation properties of the landscape as a tool to interpret algorithmic behaviour [15, 31, 102]. A prominent role in this program has been played by the notion of frozen variables [98, 101]. It has indeed been observed that all efficient algorithms return unfrozen solutions, *i.e.* solutions which do not contain an extensive number of frozen variables [59, 72], thus supporting the conjecture that only unfrozen solutions can be found in linear time [101, 103, 104]. The interpretation of this fact is straightforward: frozen clusters induce very strong correlations among the variables, an extensive number of which has to be set in a very coherent, precise way in order to avoid inconsistencies, and as a consequence local search algorithms “attacking” them would undergo an exponential slowing down due to the need for global rearrangements whenever an error in the search process is made.

The resulting picture would then be the following: efficient algorithms are capable of finding solutions in atypical, unfrozen clusters, which are guaranteed to exist up to the freezing threshold

α_f (the rigidity one α_r , on the contrary, accounts for the appearance of frozen variables in typical clusters only). The freezing transition $\alpha_f \geq \alpha_r$ is hence proposed as the ultimate algorithmic threshold [15]. This perspective is in agreement with numerical results from the literature, see for instance [31], where a message passing strategy dubbed belief propagation reinforcement was able to find solutions to the bicoloring problem for connectivities greater than α_r .

One may be worried, however, for the fact that the concept of “atypicality” as discussed so far would be still rather vague and of unspecified utility for what concerns the understanding of *individual* algorithmic behaviour: exponentially many atypical regions are indeed expected to exist in the space of solutions, each with its own features (such as a specific rigidity point α_r^*), but we are generally unable to predict to which class of them a certain algorithm will be attracted, and hence to compute the relevant α_r^* . Despite this limitations, the perspective here depicted induces us to hope that some algorithm actually saturating the bound α_f does indeed exist. It also suggests a natural ally in the design of new efficient algorithms: by appropriately biasing the measure over the set of solutions, one could hope to transform the longest lasting unfrozen atypical configurations into typical ones, and employ the resulting gain of knowledge in order to build the *ultimate* solver.

Along this line of work can be found the results of [5], observing that solutions obtained by efficient algorithms belong to regions of configuration space characterized by a *density* of solutions higher than in the typical case. This can be formalized by introducing the notion of a local entropy measuring the logarithm of the number of solutions as a function of distance in configuration space (*e.g.* the Hamming distance counting the number of different bits for models with binary variables) from a certain reference configuration. The reference configuration does not need actually to be a solution: the local entropy information at small distances was used in [6] to build the so called Entropy-driven Monte Carlo (EdMC) procedure, which starting from a random configuration is able to find the regions of the highest density. Remarkably, configurations maximizing the local entropy are likely found to actually be solutions.

In this last chapter, we will focus on the behaviour of MC based annealing procedures. Even though, as we said, this is not the most general class of algorithms, nor the best performing one, we believe understanding the very reasons underlying its *smartness* to be desirable and possibly useful. A fundamental tool to our program will be the notion of optimized interaction f^{opt} , defining a probability measure for which α_d is maximum. As figure 7.1 clearly shows, there is an evident similarity in the behaviour of our $T = 0$ algorithm when running with f^{flat} or f^{opt} . This and other similarities, such as the one regarding the observed pair correlation function (gap distribution) of the system along the procedure, point in the direction of attributing the observed *smartness* of f^{flat} to the fact that the algorithm is in this case attracted towards atypical regions of phase space which are “close” (in the sense that they share similar properties) to equilibrium configurations for f^{opt} , which have not yet clustered.

7.2 A novel insight into a CSP phase diagram

What is the origin and the meaning of the observed postponement of the dynamical transition when switching to f^{opt} ? In the most studied case of a uniform measure, it is generally accepted that clusters have an intrinsically geometric nature: they are, roughly speaking, “islands” of solutions separated in configuration space by extensive lengths (according to some proper definition of




Figure 7.2: Cartoon of the landscape of solutions around α_d , as seen from the uniform measure f^{flat} . Color legend: red regions represent solutions belonging to typical clusters, grey areas are atypical solutions instead. Clusters may be connected to the rest of solutions by atypical (entropically suppressed) paths. These paths can merge into an atypical connected component that traps the dynamics of local algorithms moving from solution to solution. Standard MC cannot equilibrate in this region. However, basing the acceptance probability over f^{opt} enables the algorithm to equilibrate and decorrelate up to the optimized clustering point α_d^{opt} , which is supposed to be the point where also this connected component finally breaks down.

distance) of “unsatisfied sea”. Let us suppose, in first approximation, that clusters are such purely geometric entities and that their definition as a consequence does not depend from the form of the soft potential. By changing the potential, then we should be able to focus our attention on atypical regions of the original landscape: the simplest, quite naive explanation for the increase of α_d is that these regions are clustering later, *i.e.* at bigger α .

This assumption results to be a too drastic simplification. In particular, a system that appears to be trapped inside the planted cluster according to f^{flat} , for $\alpha > \alpha_d$, is actually able to escape the cluster if we promote the dynamical MC rule to be based on the optimized interaction (this, and the other numerical results here discussed, are presented in detail in subsection 7.2.2). The reason is that, considering the very high dimensionality of phase space, clusters may as well be connected by some rare paths which are entropically suppressed, according to the uniform measure. The effect of the biased potential is to “widen” (give probability to) these channels. A pictorial real space representation can be given for our model in terms of caging: the switch-on of a short range attraction between particles effectively opens void channels that may favour the escaping of particles from an apparently locked situation. In this sense, a direct comparison or identification of clusters between the models obeying f^{flat} or f^{opt} is at the very least risky.

A possibly useful insight can come by interpreting clustering as the progressive, inhomogeneous rarefaction of the high dimensional space of solutions, which hence should appear as a sort of infinite dimensional porous material, or “sponge”. Forming clusters are connected to the rest of the solutions by increasingly rare paths, which nonetheless may still survive even beyond the clustering point α_d . Moreover, this picture also closely reminds us of the local entropy approach of [5], proposing exactly an inhomogeneous density profile in the space of solutions. In the present scenario, it is reasonable to expect local algorithms to be attracted towards atypical, highly connected (say denser) regions of solutions when approaching and crossing the dynamic transition threshold α_d . This means that different regions may become disconnected at different α ’s and that, in general, a connected component may still survive beyond α_d . The naive intuition behind α_d^{opt} is hence that it coincides with the point where even this atypical, connected component finally breaks down.

The nature of this connected component can be related once again to the real space interpretation of the optimized measure. It is indeed reasonable to assume that, starting from a “locked” configuration belonging to any cluster (as seen from the original problem), we can open void channels by switching on the optimized interaction, that allows particles to move freely and thermalize according to the biased measure. We should then be able to get arbitrarily close in the space of solutions to any other cluster, which is likely connected to the same component by analogous channels. Of course, we cannot use this “trick” to recover ergodicity in the original clustered measure by switching back to the f^{flat} interaction, since the way inside a cluster should coincide with its way out, which as we said is entropically suppressed. These ideas are pictorially sketched in figure 7.2.

The proposed picture is supported by the following numerical results concerning the out of equilibrium behaviour of a MC local search algorithm working at fixed zero energy (see section 7.2.1), that will be addressed in detail in the rest of this chapter. Its performance was already shown in figure 7.1: even working with the flat potential the algorithm is able to go beyond α_d . A closer look at the configurations found by the algorithm in this case allows to give a quantitative description of their atypicality: the pair correlation function between neighbours on the graph, which we call gap distribution, deviates from its equilibrium value², that is from proportionality to f^{flat} . When initialized in one of these solutions, the BP iterative procedure converges to the trivial fixed point. This was linked, see *e.g.* [102], to the fact that replica symmetry breaking at the one-step level is not sufficient to describe the totality of clusters. Likewise, if we perform a zero energy MC starting from these solutions, the behaviour is very different from that of a simple cluster: dynamics is able to find a “way out” allowing the overlap to decorrelate (this is why we prefer to speak of a connected component, rather than a cluster), but also dynamics slows down if we increase the waiting time. This behaviour resembles, at least at the qualitative level, the one of aging: the solution space visited by f^{flat} should be rather intricate and dominated by a decreasing number (with time) of rare directions or channels allowing the system to decorrelate.

Instead of addressing directly this complicated structure from the point of view of f^{flat} , we obtain some insight by studying a dynamics starting from these configurations but evolved according to f^{opt} . In this case the system is able to explore much more freely the space of solutions. We observe that after a transient the system reaches what appear to be typical solutions for f^{opt} . This has two important lines of interpretation. On the one hand, we observe that solutions found by the algorithm working with f^{flat} are “close” to typical solutions for f^{opt} , where by close we refer to two aspects: the fact that the non-trivial observed gap distribution shares some similarity with f^{opt} , and that the system can apparently still equilibrate using f^{opt} . The smartness of the searching algorithm, even when working with f^{flat} , hence corresponds to the fact that it is able to spontaneously optimize distances between variables, to some extent, in order to favour the dynamics of the system. The second point we would like to stress is that, conversely, typical configurations from the point of view of f^{opt} appear to be related to the anomalous connected component in the landscape of f^{flat} . We are thus led to hypothesize that α_d^{opt} is closely related to the maximum value of α for which such a connected component, *i.e.* with zero internal overlap, survives in the landscape of the unbiased potential (since otherwise a better potential could possibly be found further increasing α_d). The behaviour of the search algorithm when running according to f^{opt} , showing no anomaly

²Whenever the effect of short loops is negligible, the pair correlation functions obeys $g(r) \propto e^{-\beta v(r)}$, where $v(r)$ is the pair potential.




Figure 7.3: Constructing a pseudoenergy landscape from a constraint satisfaction problem. With this orientation of the y-axis, the pseudoenergy is defined as the smallest value of α , according to the ordered list of edges that was initially drawn, for which a given configuration (on the horizontal plane) ceases to be a solution. Reprinted from [59].

and exhibiting a power-law divergence of the timescale to increase α exactly in correspondence of α_d^{opt} , can be read in this direction.

7.2.1 $T = 0$ Monte Carlo algorithm

The algorithm, looking for satisfied configurations at increasing values of α , proceeds as follows. At first, an ordering of the possible $N(N - 1)/2$ edges (ij) is drawn in a list identifying an instance of the problem, and variables are initialised in a random configuration. Edges are iteratively proposed one at a time, starting from the situation $M = 0$. If variables (ij) associated to the $(M + 1)$ -th edge satisfy the constraint, that edge is added to the graph. Otherwise, a $T = 0$ heat bath Monte Carlo sweep as described³ in section 5.1.2 (an update of N variables among the solutions to the problem with only M interactions) is performed until the $(M + 1)$ -th edge is satisfied. Notice that since we add a link only when the relative clause is satisfied, and since $T = 0$ MC moves cannot increase the energy, during all the procedure the system only visits proper solutions of the CSP. This same procedure can be applied to different interactions by changing the heat bath weights. We do not worry about acceptance rates when comparing f^{flat} and f^{opt} , since heat bath comes with no rejections: the time here simply corresponds to the number of accepted moves (we allow though for the possibility that a variable chooses to “flip” to its current position).

A pictorial interpretation of this kind of algorithms can be obtained by following the ideas of [59]. First one introduces a pseudoenergy landscape as represented in figure 7.3. It can be defined as following: if we imagine to cut the pseudoenergy horizontally (at fixed α), each slice would then correspond to one of the bidimensional snapshots of figure 3.2, with solutions in this case living in the interior of the pseudoenergy surface. In a more precise way, we associate for each point in configuration space (the plane) a value of the pseudoenergy corresponding to the smallest value of α , according to the list we have fixed, for which that configuration ceases to be a solution. Our algorithm performs a stochastic descent in such a complex, glassy landscape. Notice, however,

³We use once again the clock version of the continuous coloring, since we want to compare the numerical data from optimization algorithms with the most accurate determination of the thermodynamic thresholds α_d and α_d^{opt} .

that the reason for algorithmic failure should not be attributed to the fact that the system gets trapped into a local minimum of the pseudoenergy lying above α_{sat} , since we believe the system never to enter a proper cluster. In analogy with the topological interpretation of standard glassy dynamics [64], the slowing down of the search procedure should be connected to the approaching to the (pseudo)energy value (that is the value of α) defining the dynamic transition, where many local minima form in first place and there is abundance of flat directions. This appears to be in agreement with the behaviour of the algorithm when running according to the optimized interaction f^{opt} , since in this case we observe a smooth power-law slowing down approaching α_d^{opt} . One may be tempted to associate α_d^{opt} with the location of the relevant topological saddles-to-minima transition in the portion of the pseudoenergy landscape *actually visited* by the algorithm (according to our definition, the pseudoenergy surface is independent from the choice of f^{flat} or f^{opt} , the only difference being encoded in the statistical weight we attribute to stationary points to determine what we call the typical properties of the landscape). We should warn the reader, however, that these ideas are at the moment just a qualitative suggestion. In particular, no rigorous analysis of the kind of [24] has been proposed in the literature for the topology of the pseudoenergy landscape here described, nor it is obvious whether the above construction for the pseudoenergy would result in a continuous, differentiable surface (consider for instance that α is truly continuous only in the thermodynamic limit).

7.2.2 Numerical results

The following points will be addressed in detail in this section:

1. Comparison of the algorithmic performance using f^{flat} or f^{opt} , pointing α_d^{opt} as the correct algorithmic threshold for f^{opt} , and as an upper bound to the algorithmic threshold for any other generic interaction.
2. Gap distributions along the search procedure. When running according to f^{opt} , the gap distribution stays close to its equilibrium expectation f^{opt} . Conversely, when running f^{flat} , the gap distribution becomes noticeably out-of-equilibrium even before α_d . The gap distribution develops in this case some qualitative feature similar to f^{opt} . This may be connected to the similarity in the performance between the two algorithms, and gives a practical explanation to the observed algorithmic *smartness*.
3. Further exploration around the solutions found by the algorithms. These solutions do not appear to belong to what we would properly define as a “cluster”. We study the behaviour of the overlap, which always decays independently from the interaction used even above α_d , and of the gap distribution as a function of time.
4. Finally, we show the evolution from the cluster planted according to f^{flat} . A zero energy MC dynamics exploiting f^{opt} is able to escape the cluster and to apparently converge to equilibrium in the biased measure (which has not yet clustered).

Comparison of algorithmic performance

We will refer to the data already presented in figure 7.1. We call $t_{\Delta\alpha}$ the number of MC sweeps that are needed in order to increase the number of edges in the graph from M to $M + \Delta\alpha N$. The




Figure 7.4: Algorithmic performance: the time needed to increase the connectivity α by $\Delta\alpha = 0.2$ diverges as a power law approaching $\alpha_d^{\text{opt}} > \alpha_d$ when f^{opt} is used. The extrapolated algorithmic threshold corresponds to $\alpha \approx 36.72$, to be compared to $\alpha_d^{\text{opt}} = 37.71(1)$ obtained from BP in section 6.1.2. Inset: detailed view of the long time region. Color dashed lines are meant to guide the eye.

time for a single link addition is indeed a very irregular quantity, discontinuously jumping from 0 to very large values (up to order $10^6 \sim 10^7$). The purpose of $t_{\Delta\alpha}$ is to quantify the amplitude of these fluctuations, which increases with α . We choose $\Delta\alpha = 0.2$ so that in each interval we likely find a “hard link”, requiring a time much greater than the average for that level of connectivity: with this choice, $t_{\Delta\alpha}$ is thus comparable to the time needed for adding such hard links. For smaller $\Delta\alpha$, on the contrary, one is still left with a consistent amount of noise due to the presence of many intervals without any relevant hard link inside; this situation can of course be dealt with by averaging over some samples, but at the price of defining a time scale $t_{\Delta\alpha}$ that is smaller than the actual difficulty of the problem the smaller is $\Delta\alpha$. Notice that we stop when the time to add an edge exceeds a threshold equal to 10^7 MC sweeps.

We averaged the logarithm of $t_{\Delta\alpha}$ over 55 (respectively 30) instances of the problem for $N = 10^4$ ($2 \cdot 10^4$). In order to avoid trivial biases, we only considered data up to the smallest value of α that was reached by all the samples in each case. For the f^{opt} interaction and size $N = 10^4$, this bound is equivalent to $\alpha \approx 36.4$, while solutions are found for some sample even up to $\alpha \approx 36.8$. The average over the samples of $t_{\Delta\alpha}$ at fixed α , however, is particularly sensitive to the presence of rare samples diverging earlier. A safer quantity to consider then would be the average of α at $t_{\Delta\alpha}$ fixed. We followed of course also this route, obtaining compatible results. This second approach, on the other hand, is subjected to a cutoff on $t_{\Delta\alpha}$ given by the smallest value of $t_{\Delta\alpha}$ that was reached. We observe that for our data this further limits the explored range of α , so we decided to present the data in the first way. Notice, however, that the very last points for $N = 10^4$ are presumably dominated by such a few most diverging samples (see figure 7.4, inset).

The graph instances used are of course the same for the two potentials. The f^{opt} data is found to be compatible with a power law diverging very close to α_d^{opt} , as shown in figure 7.4. The f^{flat} interaction seems to asymptotically perform worse. Nevertheless, the behaviour at high α is very

similar and there appears to be no sharp advantage in using f^{opt} at the time scale we are able to investigate. This is a first hint that the origin of the algorithmic smartness of f^{flat} should be due to some self-generated similarity with the f^{opt} case: as we are going to discuss in the next paragraphs, from our analysis it emerges that the algorithmic strategy based on f^{flat} explores more frequently than expected regions that would have been favored by f^{opt} , giving effectively higher “weight” to them. These findings support the idea that the very existence of α_d^{opt} could affect and possibly explain also the behaviour of a simple algorithm based on f^{flat} .

Gap distributions along the search procedure

A precise characterization of the *atypicality*, with respect to the uniform measure, of the solutions found above α_d can be given in terms of the distribution of angular distances between neighbouring variables. We plot in figure 7.5 the histogram of these distances as extracted from the configurations (one from each sample) obtained for $N = 10^4$ at different values of α along the procedure. The histogram is normalized to the value of $p = 200$, in order to follow the definition of f . A first comparison between the behaviour of the observed gap distribution for the two different potentials f^{flat} (on the left) and f^{opt} (on the right), shows that the former strongly depends on α even close but before α_d , while the latter remains stationary in α . This observation matches the fact that the behaviour for f^{opt} in figure 7.4 was extraordinarily consistent with the thermodynamic prediction α_d^{opt} : we are possibly exploring a quasi-equilibrium regime.

From the point of view of the algorithm running f^{flat} , figure 7.5 shows that the system ends up in configurations characterized by an excess of *contacts* between neighbouring particles. We already know that such a shape for the interaction is responsible for an increase of α_d at the thermodynamic




Figure 7.5: Left: atypical configurations found by the algorithm running with f^{flat} are characterized by a peculiar gap distribution displaying a peak at small gaps, already forming before $\alpha_d = 34.63(2)$. The equilibrium distribution in this case would have coincided with plain f^{flat} instead. Right: if we run the same algorithm but using f^{opt} , the gap distribution always stays close to its equilibrium shape. Since $\alpha_d^{\text{opt}} = 37.71(1)$, we are still well before the clustering transition for the biased measure.

level. A suggestive interpretation of this result is that our search algorithm is, to some extent, capable of self-optimizing distances between variables in order to surpass α_d . However, the MC rule is still built over f^{flat} at any moment: this suggests that the good performance of the algorithm should be strictly and mostly linked to the properties of the visited configurations. Finally, the non-trivial gap distribution observed in the left panel of figure 7.5, even if *smarter* than f^{flat} (meaning that can be associated to a measure reasonably clustering later than α_d), is still less optimized than f^{opt} . This hints α_d^{opt} as a theoretical (upper) bound to the asymptotic performance of our MC algorithm, irrespective of the interaction potential used. Remarkably, however, there appears to be no practical advantage in preferring f^{opt} over f^{flat} , at the timescales that have been numerically accessed in our simulations, see figure 7.4.

Further exploration around the found solutions

The atypicality of the solutions found by our algorithm can also be assessed by directly investigating the surroundings space of solutions. To this end, we perform a zero energy MC dynamics starting from the obtained solutions. We consider $\alpha = 35.5$, that is reasonably bigger than $\alpha_d = 34.63(2)$, and for which the MC evolution starting from the state planted according to f^{flat} is stuck inside the planted cluster, see figure 7.6. In all the cases we consider $N = 10^4$ and average over 10 samples.

We consider in different colors the four possible combinations between the interaction used in the search algorithm and the one used in the subsequent MC evolution. The first observation we can draw from figure 7.6 is that the $f^{\text{flat}} + f^{\text{flat}}$ data (in purple) is able to escape the plateau, even

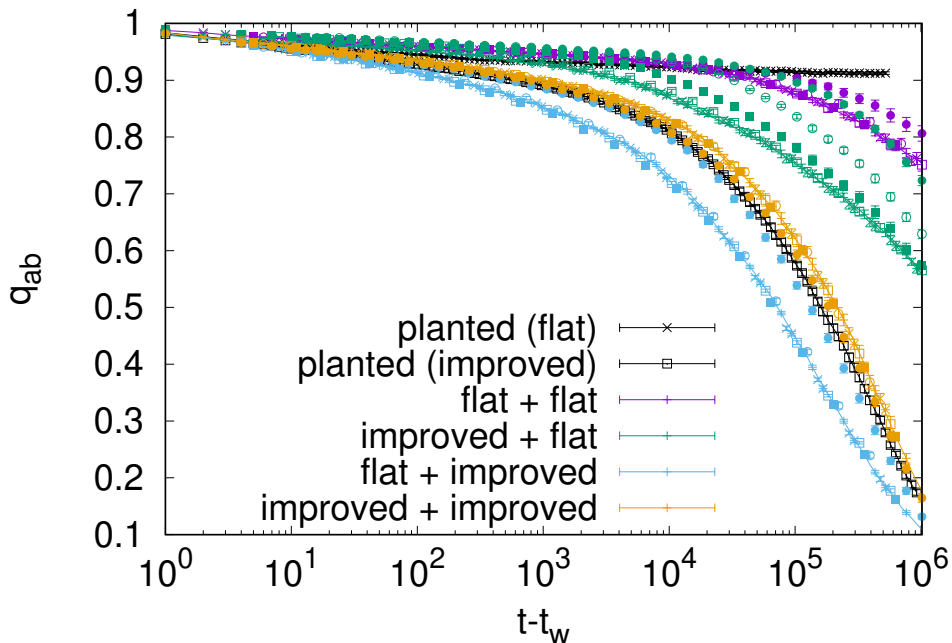


Figure 7.6: Zero energy MC evolution starting from solutions found by our search algorithm at a given value of $\alpha = 35.5 > \alpha_d$, for $N = 10^4$. Different colors are named “x+y”, where “x” is the interaction used in the search procedure, while “y” is the one used in the subsequent MC evolution starting from those solutions. For each type of curve, seven waiting times are displayed, $t_w = 0, 10, 10^2, 10^3, 10^4$ (full square), 10^5 (empty circle), 10^6 (full circle). Black: equilibrium behaviour of the overlap. For $\alpha = 35.5$, f^{flat} is stuck into the planted cluster, while the improved interaction f^{opt} is still able to decorrelate.

though typical configurations for f^{flat} are not. The overlap decay suggests us that the system is in an aging-like regime, with an effective age of order 10^6 MC sweeps. For smaller waiting times, we observe a practically stationary behaviour, as it is also confirmed by the gap distributions in the first plot of figure 7.7. Interestingly, it looks like the system has a hard time trying to get back to a gap distribution more similar to f^{flat} .

The behaviour of the $f^{\text{opt}} + f^{\text{opt}}$ data (in yellow) is diametrically different. In this case, solutions found by the search algorithm at the value of $\alpha = 35.5$ under consideration (which is by the way rather smaller than the improved clustering threshold $\alpha_d^{\text{opt}} = 37.71(1)$), are found to behave very closely to the expected equilibrium configurations. Once again, we can identify a time of order 10^6 MC sweeps, after which the overlap data seem, this time, to collapse onto the reference planted curve.

The two remaining combinations display specular behaviours. The idea, also supported by




Figure 7.7: Gap distributions at different t_w 's during the experiments of figure 7.6. Upper left corner: $f^{\text{flat}} + f^{\text{flat}}$. Upper right: $f^{\text{opt}} + f^{\text{opt}}$. Bottom left: $f^{\text{opt}} + f^{\text{flat}}$. (In dashed black, the gap distribution from the search procedure using f^{flat}). Bottom right: $f^{\text{flat}} + f^{\text{opt}}$.

the behaviour of the gap distribution (bottom line of figure 7.7) is that solutions found by the f^{flat} search algorithm can evolve into equilibrium configurations for f^{opt} , and viceversa equilibrium configurations for f^{opt} (or solutions found by the f^{opt} search algorithm, which is almost equivalent) can evolve into solutions very similar to the ones found by the f^{flat} search procedure. This supports our proposal of a connected component, disjoint from proper clusters (apart for rare channels essentially invisible to f^{flat}), that is accessed by local search procedures.

Evolution from the original clusters

We conclude this discussion by showing that clusters are an interaction dependent (*i.e.* not purely geometric) concept, and that well formed clusters for f^{flat} may be felt not the same way by f^{opt} , since the entropic barriers underlying the very physical definition of clusters as responsible for dynamical confinement are different for the two interactions. To this end, and similarly in spirit to an analogous experiment carried out in [60], we plant the system according to f^{flat} at $\alpha = 35.5 > \alpha_d$, and evolve it through the usual zero energy MC following f^{opt} . As shown in figure 7.8, the system is able to escape from the planted cluster and moves towards configurations similar to those that are typical according to f^{opt} .

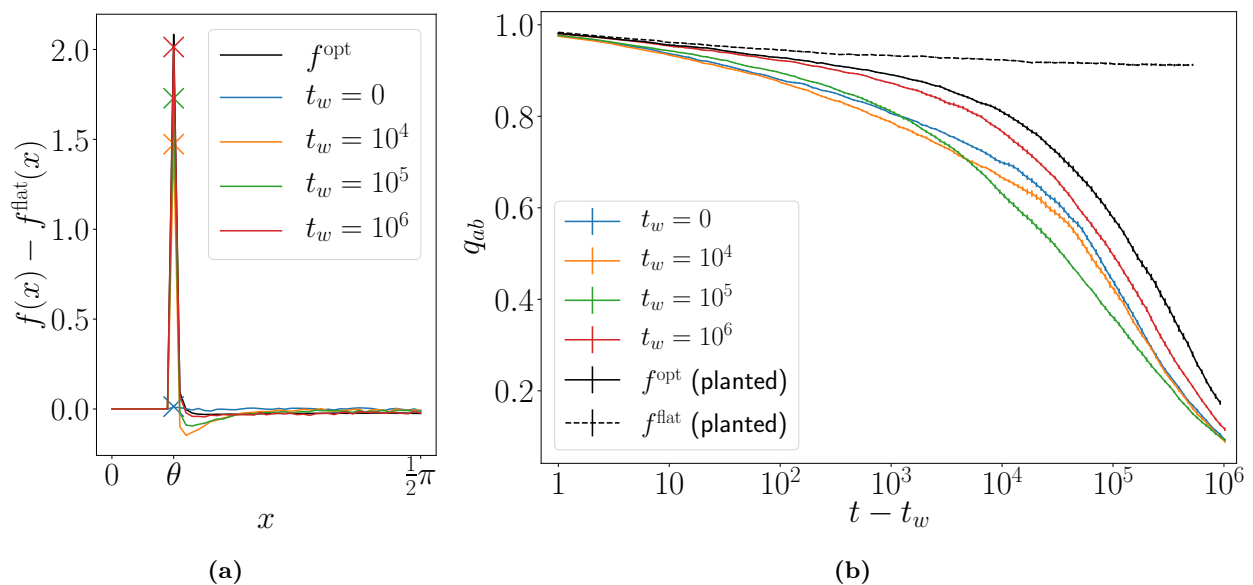


Figure 7.8: Right (color lines): basing the dynamical rule over f^{opt} , the system is not trapped inside a cluster, even though it was initialized by planting according to f^{flat} at $\alpha = 35.5 > \alpha_d$. For sufficiently large waiting times, the overlap curve gets close to the equilibrium behaviour for f^{opt} (black solid line) at the same value of α . Black dashed line represents the reference planted cluster according to f^{flat} . Left: also the gap distribution tends to the equilibrium behaviour according to f^{opt} .

7.3 Quench and simulated annealing

The phase diagram for the improved interaction shown in figure 6.4 takes also into account temperature, as the external parameter relaxing the hardness of the excluded volume constraint. A natural question is whether α_d^{opt} could be interpreted, in analogy with the previous $E=0$ situation, as the easy-hard transition for optimization procedures based on f^{opt} MC dynamics in temperature, and as an upper bound for those based on f^{flat} . Our data, unfortunately, point in a different direction.

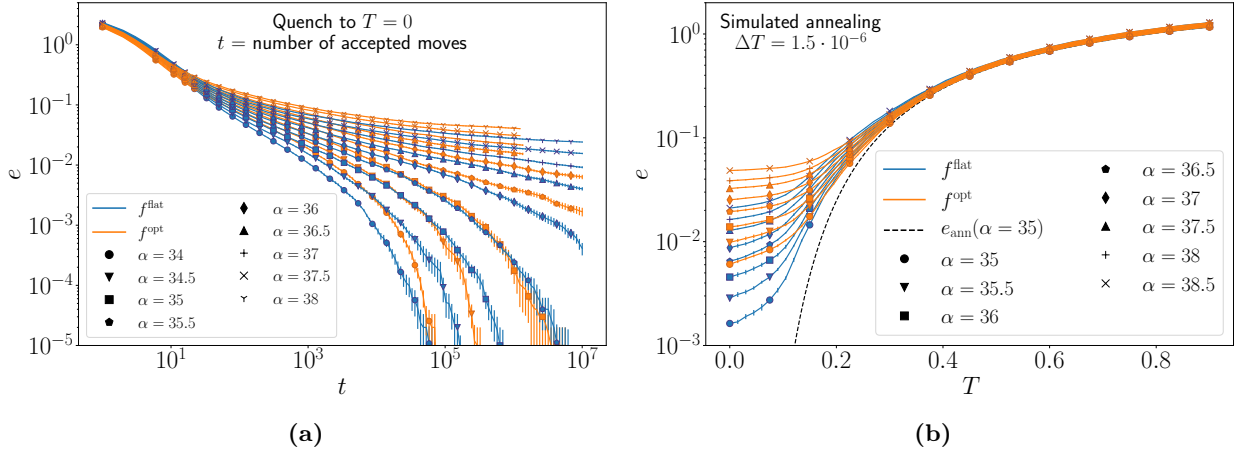


Figure 7.9: Optimization by decreasing temperature is easier using the interaction f^{flat} . Data is displayed for $N = 10^4$ and averaged over 10 samples (the same for both f^{flat} and f^{opt} , and for the two procedures). The energy e counts the number of violated constraints divided by N , as a function of time after the quench or of temperature during the SA. Left: for each value of α , the f^{flat} procedure is able to reach lower energies in any given time window, where t is defined as the number of accepted attempts to update the variables, divided by N . Asymptotically, f^{flat} appears to succeed for values of α at least as big as $\alpha = 36 > \alpha_d$ (5th blue curve from the bottom), while the determination of the behaviour of f^{opt} becomes more uncertain already for $\alpha = 35.5$ (4th orange curve from the bottom). Right: we show the behaviour of SA for a fixed value of the cooling rate, in order to display different values of α . The asymptotic limit $\Delta T \rightarrow 0$ will be studied in the following; however, f^{flat} always appears to perform better. Black dashed line represents the equilibrium annealed (paramagnetic) energy (6.3) for $\alpha = 35$.

The first intriguing observation is that, for both a quench to $T = 0$ and simulated annealing (SA), f^{flat} performs *better* than f^{opt} , see figure 7.9. Both procedures are initialized from random configurations. In the case of a quench to $T = 0$ we adopt a rejection free version of the Metropolis algorithm: first we compute the local energy contribution (number of violated edges) for each of the $p = 200$ states of variable i coming from the interaction with the neighbours (this can be done by cycling one time over the neighbours and increasing the local energy by 1 to all the states of variable i that are inside an interval of $\pm\theta$, random shifts included); then we propose a new state for variable i uniformly among those having a local energy equal or smaller than the current one.

In the f^{flat} case the proposal is automatically accepted, and one is left with exactly a rejection free version of the usual Metropolis algorithm at zero temperature. In the f^{opt} case, on the contrary, if the proposed local energy is equal to the current one, the move is accepted with the Metropolis rule applied to the soft part of the interaction, *i.e.* with probability $\min(1, w(x_{\text{new}})/w(x_{\text{old}}))$, where $w(x) = \prod_j f_{x-x_j-\omega_{ij}}^{\text{opt}}$ and the product is carried over the neighbours j for which the constraint is satisfied. It is evident that the f^{opt} procedure is identical to the previous one, with the exception that some of the moves that leave the energy constant are in this case rejected. Even accounting for this difference in the acceptance rates of the two procedures, by defining in the second case a MC sweep as the *successful* update of N variables, the energy decreases faster when using f^{flat} .

In the simulated annealing we decrease temperature from $T_{\text{max}} = 1.5$ to $T = 0$ with steps of $\Delta T = T_{\text{max}}/\tau$ for $\tau = 10^4, 10^5, 10^6, (10^7)$ and $N = (5 \cdot 10^3), 10^4$. A single MC sweep is performed at each temperature according to the heat bath rule (for the last step at $T = 0$, we force variables to select one of the states minimizing the local energy contribution, even if nonzero). The amplitude of the hard-core term proportional to $e^{-\beta}$ of f^{opt} is fixed by equation (6.2); the equilibrium annealed energy (6.3) for $\alpha = 35$ is shown in figure 7.9 as a black dashed line.

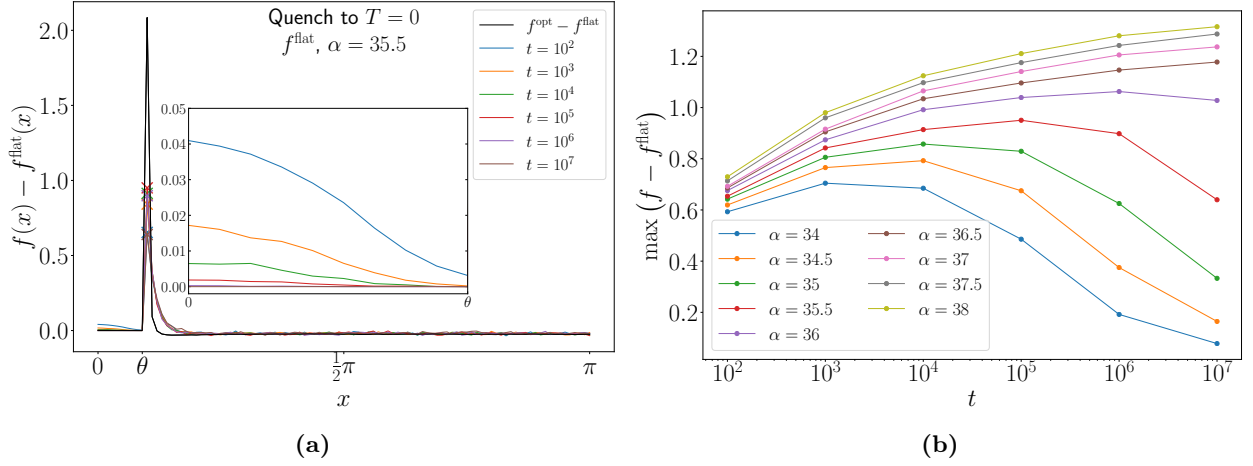


Figure 7.10: Behaviour of the gap distribution during a quench to $T = 0$ using the flat potential. Left: gap distribution for a fixed value of $\alpha = 35.5$. Between $t = 10^6$ and $t = 10^7$ all the samples have reached zero energy. The optimization process can be schematized in two regimes: in the initial time sector, when the number of violated constraint to be minimized is still consistent, the system favours exactly closed gaps in order to efficiently decrease the energy. This corresponds to an initial increase of the peak at $x = \theta$ with time. At longer times, when the energy is zero or it is close to be zero, the system is able to relax some of the closely satisfied constraints, and the peak decreases. Right: evolution of the peak with time for all the values of α . Curves for $\alpha \leq 35.5$ display a clear decrease in the peak for long times; for all of them the system has reached zero energy before $t = 10^7$. The curve $\alpha = 36$ is reasonably about to start decreasing. The corresponding energy, see figure 7.9, starts bending towards zero with a faster than a power law behaviour.

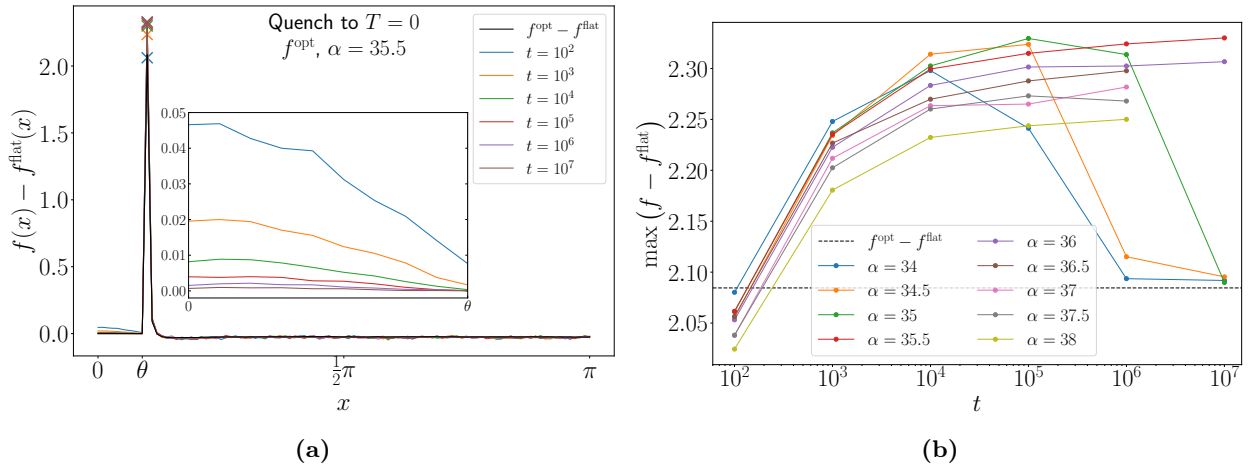


Figure 7.11: Same as figure 7.10, but using the potential f^{opt} . In this case the energy optimization performs worse, and we are not able to reach zero energy for $\alpha = 35.5$ in $t = 10^7$. The red curve on the right panel, corresponding to $\alpha = 35.5$, has not yet started decreasing. In the left panel, the gap distribution shows marked non-equilibrium features: the $x = \theta$ peak consistently exceeds f^{opt} , while, and above all, the distribution of violated gaps (inset) has a non constant shape.

An inspection of the gap distributions at different times along the quenching procedure is proposed in figures 7.10 and 7.11. In this case, a non-trivial distribution may arise in both the soft region $|x| \geq \theta$ and in the hard one $|x| < \theta$. For what concerns the soft region, we observe once more the spontaneous formation of a peak at $x = \theta$, also when running f^{flat} . Interestingly, the height of this peak seems to possess a dynamics on its own: it is increasing with time as long as the energy is consistently different from zero, then it starts decreasing once the energy has reached some threshold of order 10^{-3} and steadily approaches zero energy. A similar phenomenology is observed in the case

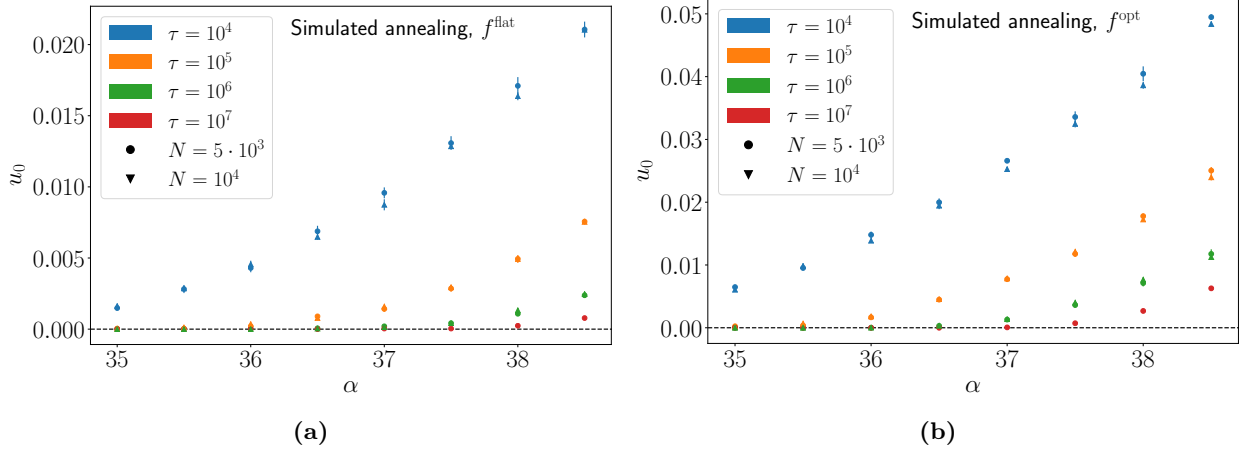


Figure 7.12: Energy reached at the end of simulated annealing, averaged over 10 different samples for each size. The SA cooling rate τ is defined by the relation $\Delta T = T_{\text{max}}/\tau$, where $T_{\text{max}} = 1.5$. The behaviour is qualitative similar for the two interactions, but the energy scale is bigger in the f^{opt} case. However, in both the cases, we are able to reach zero energy in all 10 samples up to $\alpha = 36$ for $N = 10^4$ and $\tau = 10^6$, and up to $\alpha = 36.5$ for $N = 5 \cdot 10^3$ and $\tau = 10^7$.

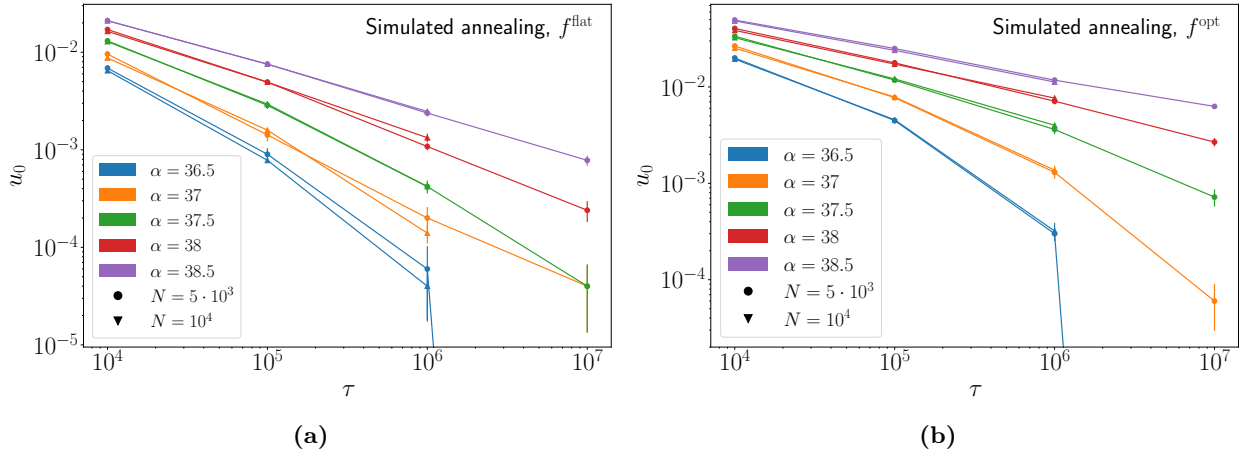


Figure 7.13: Energy reached at the end of simulated annealing. Figure 7.12 data is presented as a function of the cooling rate τ . For $\alpha < 36.5$ all the samples reach zero energy using a finite cooling rate (points are not displayed). For $\alpha = 36.5$ (blue curve), zero energy is reached with both f^{flat} and f^{opt} for $N = 5 \cdot 10^3$ and $\tau = 10^7$ (we plotted it as a very low value outside the frame due to the log scale on u_0).

of the improved interaction f^{opt} . This pronounced out-of-equilibrium character is testified also by the non-triviality of the hard part for $|x| < \theta$ (inset on the left panel). In particular, comparing the observed distributions of negative gaps for f^{flat} and f^{opt} (which we recall to have been optimized only for positive gaps), we notice a different slope close to $x = \theta$. We believe that the shape of the distribution of negative gaps may play a decisive role in understanding and predicting optimization performance.

The endpoint energy u_0 from simulated annealing is shown in figure 7.12 as a function of α for different sizes and cooling rates. A first remark is that the energy scale for f^{opt} is greater than that for f^{flat} for all the cooling rates we considered. The asymptotic behaviour for an infinitely slow cooling rate is however hard to extrapolate. Figure 7.13 somewhat suggests that u_0 should go to zero faster than a power law at least up to $\alpha = 37.5$, for both the interactions. The behaviour for bigger α 's is unfortunately less predictable, given the available points.

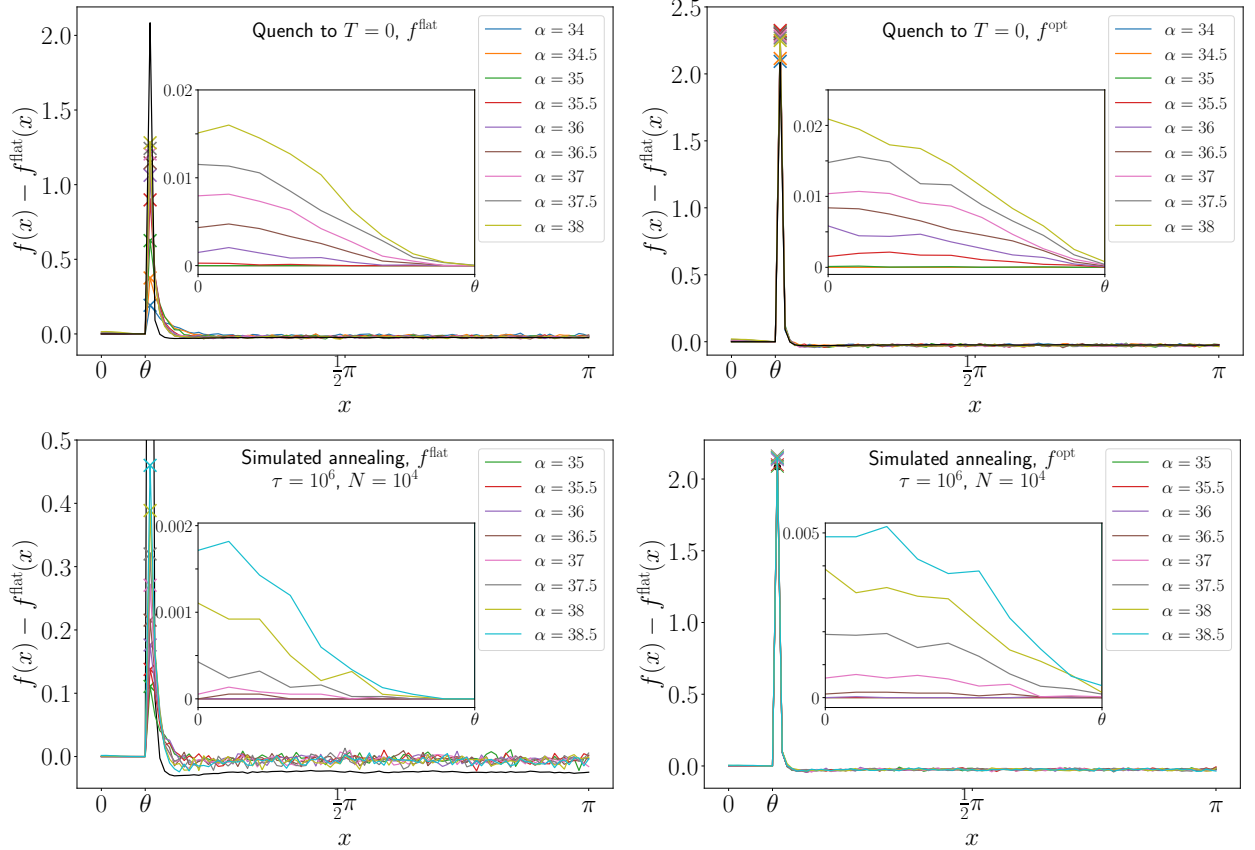


Figure 7.14: Comparison of the gap distributions at the end of the run after a quench to $T = 0$ (top) or simulated annealing (bottom). Even using the improved interaction f^{opt} , we observe the emergence of a non-trivial gap distribution in the range $|x| < \theta$ (inset), characterized by an abundance of slightly unsatisfied constraints $x \approx \theta$ with respect to the f^{flat} case.

The gap distribution at the end of simulated annealing for the slowest cooling rate available for $N = 10^4$ is shown in figure 7.14 (bottom line), and compared to the case of a quench to $T = 0$ (top line). As expected, the soft part of the distribution appears to be more “equilibrated” after SA than after the quench, meaning that it is more similar to f^{opt} when using the improved interaction and displays lower peaks when using f^{flat} . However, the hard part for $|x| < \theta$ shows a similar non-trivial, uncontrolled behaviour as in the quench. A qualitative difference between the results for f^{flat} and f^{opt} is that in the former violated gaps appear to vanish approaching $x = \theta$ with a higher power. This suggests that a higher concentration of particles at small positive gaps, as induced by f^{opt} , may also lead to an excess of small negative gaps, possibly disadvantaging optimization. The conclusion we can draw at the present stage is that, as we have already stressed, the lack of control on this part of the distribution, which is the most relevant to optimization, might be at the origin of the difference in the performance between the two interactions. In this perspective, it would be very interesting to perform again the optimization of the potential for $T > 0$, *i.e.* considering this time also the negative gaps region.

In figure 7.15 we adopt a different SA schedule, that is to fix a grid of temperatures ($\Delta T = 10^{-2}$) and to perform n MC sweeps at each temperature. The overlap between the first and last configurations at each temperature is plotted against the average energy (left panel). We observe that at any given energy level the improved interaction f^{opt} is able to decorrelate better. It however gets

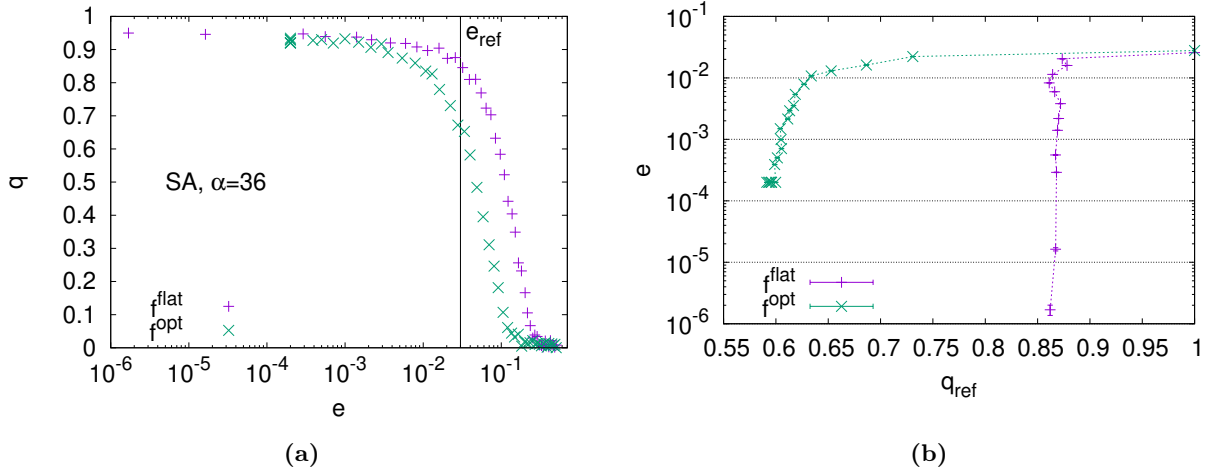
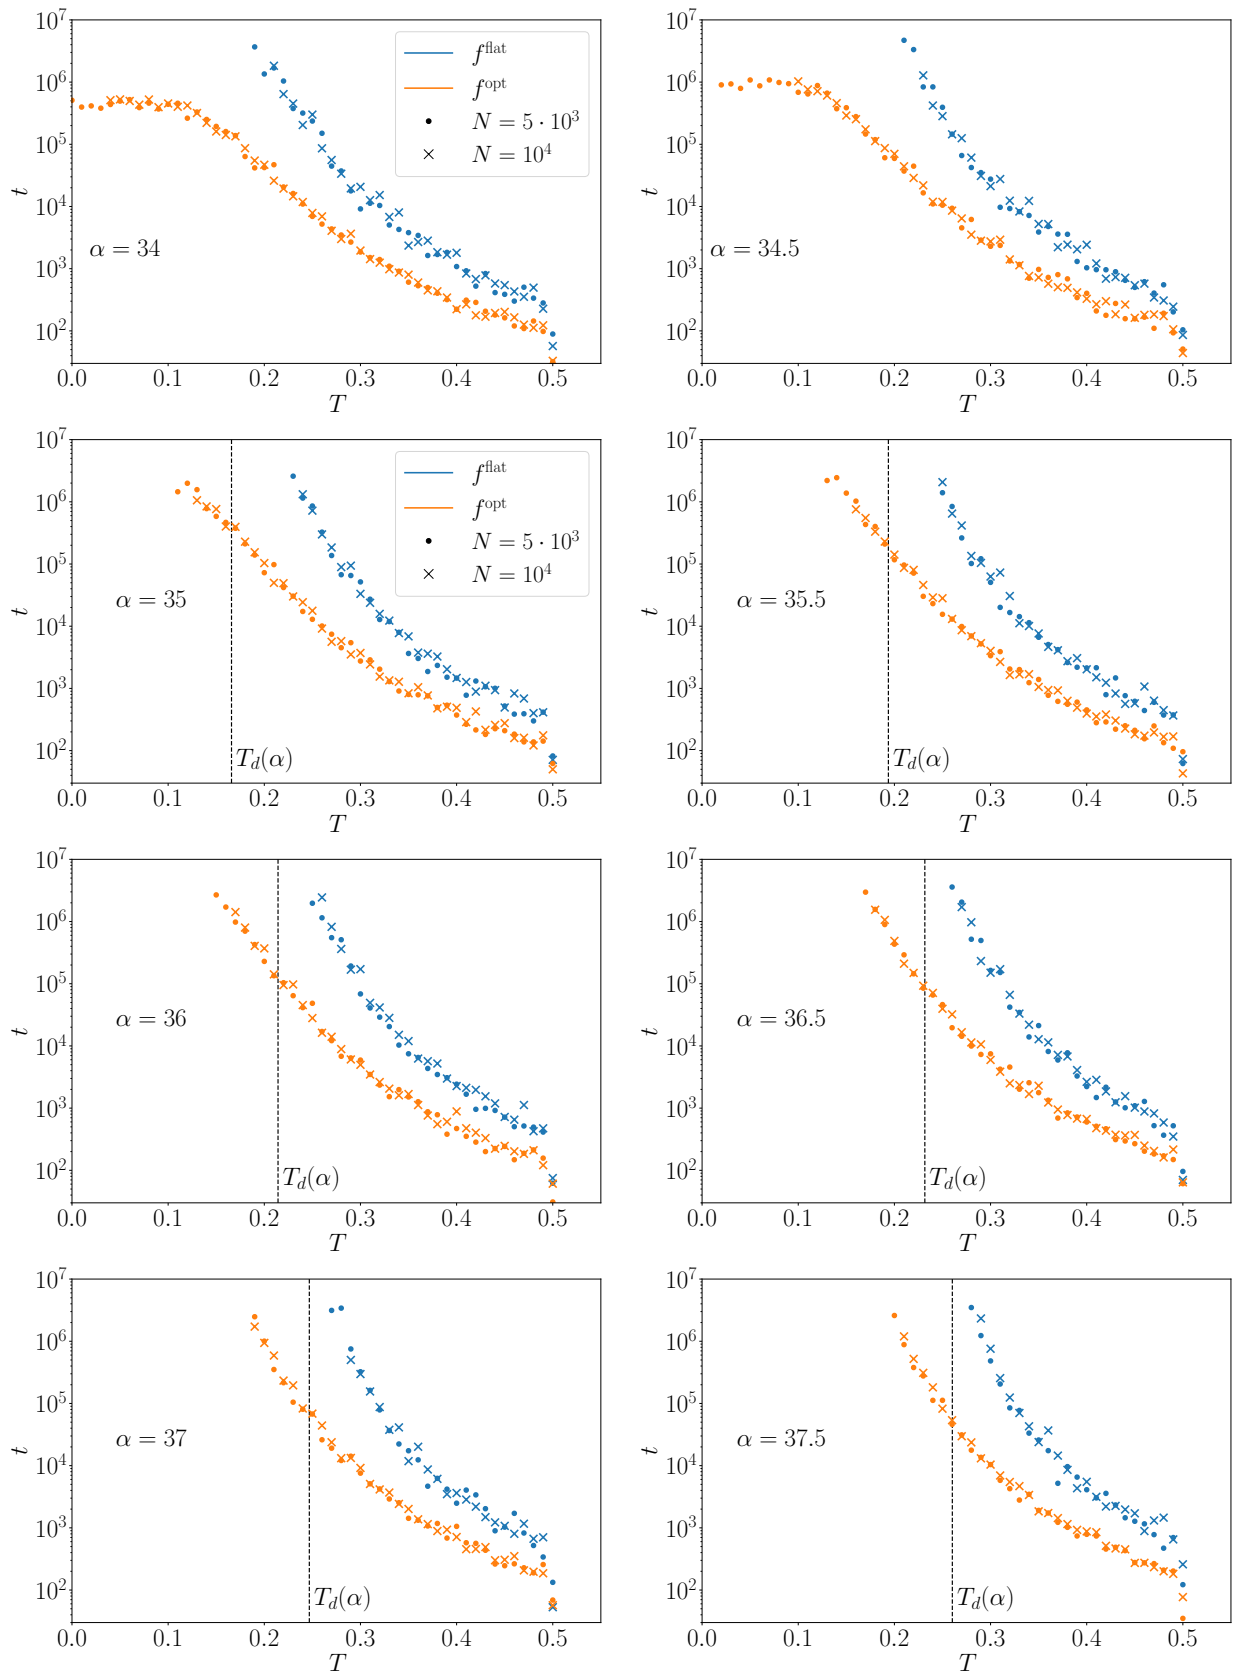


Figure 7.15: Simulated annealing running f^{flat} is able to reach zero energy without moving much in configuration space. On the contrary, using f^{opt} , the system can decorrelate more, but gets stuck at some threshold energy bigger than zero. In both panels we consider SA performed on a $N = 10^4$ system at $\alpha_d < \alpha = 36 < \alpha_d^{\text{opt}}$. The SA procedure is modified as follows: we set $\Delta T = 10^{-2}$ and perform n MC sweeps for each temperature, where here $n = 3.2 \cdot 10^3$. The energy is averaged over the last $n/2$ sweeps in each step. Left: overlap q between the first and last configurations in each SA step. Running f^{opt} , the system always decorrelate faster, before eventually getting stuck. Right: we compute q_{ref} as the overlap with a reference configuration chosen as the first one after that $e < e_{\text{ref}} = 3 \cdot 10^{-2}$. From that point on, the flat potential dives almost “vertically” in the energy landscape.

stuck at some local minimum of the energy which is different from zero. In the right panel we plot the overlap q_{ref} between the final configuration at each temperature and a reference configuration chosen as the first one after that $e < e_{\text{ref}} = 3 \cdot 10^{-2}$. We observe that the algorithm running f^{flat} is able to reach very low energy configurations without moving much in configuration space. These results are consistent with the *canyons versus valleys* picture proposed in [102], connecting the efficiency of SA to the fact that equilibrium configurations around T_d may belong to the basin of attraction of energy minima going down to zero energy (canyons). Our data suggests on one hand that it should be possible to tune the bias in the interaction in order to explore different canyons or valleys dominated scenarios, even though it is not obvious at all to us for the moment how to control this strategy. On the other hand, it also raises some questions about the true nature of the phase (in temperature) for $\alpha < \alpha_d^{\text{opt}}(T)$ as seen by the optimized interaction f^{opt} , which still possesses glassy features as soon as the equilibrium setting is abandoned.

7.3.1 Increasingly slow annealing

In order to better understand this point, we build a variant of simulated annealing in which the number of MC sweeps at each temperature is not a constant. Instead, we measure after each MC sweep the overlap with the last configuration at the previous temperature, calling t the number of sweeps needed in order for such an overlap to decrease below a very low threshold $q_{\text{max}} = 10^{-2}$. The system is then further evolved for a time t at the same temperature, computing the gap distribution and the average energy, before decreasing temperature following the same schedule as in the previous experiment of figure 7.15, $\Delta T = 10^{-2}$. This adaptive procedure allows us to maintain the system closer to equilibrium. This is testified from the observed gap distributions, which reproduce the interaction f^{flat} or f^{opt} used in the MC update in both the soft and hard region, and from the average






Figure 7.17: Increasingly slow annealing. We plot the number t of MC sweeps necessary at each temperature in order to decorrelate to an overlap smaller than $q_{\max} = 10^{-2}$ with respect to the last configuration at the previous temperature. Both the gap distribution and the energy follow at each step the equilibrium prediction, see *e.g.* figure 7.18. Vertical lines, whenever present, represent the thermodynamic estimate of the dynamic transition temperature $T_d(\alpha)$ for f^{flat} (above $\alpha_d = 34.63$) and $T_d^{\text{opt}}(\alpha)$ for f^{opt} (above $\alpha_d^{\text{opt}} = 37.71$).

energy, which exactly follows the equilibrium paramagnetic prediction (6.3) for each temperature that was reached. In particular, since the temperature dependence (6.3) is the same for the two interactions f^{opt} and f^{flat} , we can use interchangeably temperature or energy in order to compare the final points of the two algorithms, see figure 7.18.

This situation is summarized in figure 7.17 (extending on two pages). We have considered just two samples, one for each different size $N = 5 \cdot 10^3$ and $N = 10^4$. For all the values of α , the SA running f^{opt} is able to get to consistently lower temperatures, and hence lower energies, with respect to f^{flat} , before the time needed to switch to a lower temperature becomes numerically unfeasible. The first two panels show respectively $\alpha = 34$ and $\alpha = 34.5$, thus both below the clustering point α_d for f^{flat} . The f^{flat} algorithm is however unable in these cases to reach zero energy, and rather shows a divergence of times with an asymptotic energy/temperature likely to be strictly bigger than zero. Comparing with the f^{opt} data at the same connectivities, though, we observe a bending of the curve at low temperatures which is hardly predictable by the first part of the curve only. Some signs of a similar bending can be hinted by looking at the f^{opt} data for $\alpha = 35, 35.5$. However, the bending happens at bigger times t when approaching the relevant clustering threshold, so that at the timescales numerically achievable, and especially in the vicinity of α_d (for f^{flat}) and of α_d^{opt} (for f^{opt}), we are restricted to data in the first, non-predictive part of the curve, and hence we cannot distinguish any qualitative or quantitative change of behaviour when crossing the dynamical threshold. What we can say, though, is that above their clustering thresholds both algorithms behave in reasonable agreement with the thermodynamic expectation.

In conclusion, our data suggest that the region $\alpha < \alpha_d^{\text{opt}}$ should not be considered too “simple”, when temperature is introduced, even if working with the optimized f^{opt} . While reasonably reversible procedures, such as the one discussed in this last subsection, are still at least consistent with the thermodynamic expectations, more markedly out-of-equilibrium algorithms such as standard SA show a behaviour less easy to understand. In particular, even if the f^{flat} interaction leaves the reference equilibrium paramagnetic energy curve earlier in temperature, as also testified by the results in this last subsection, when implementing standard SA the flat potential appears to be able to outperform f^{opt} in the long run, see *e.g.* figure 7.9b. From our point of view, the fact that




Figure 7.18: Endpoint energies for the SA experiments of figure 7.17 and system size $N = 10^4$. Black dashed lines are the paramagnetic energy (6.3) for $\alpha = 34$ (lower curve) and $\alpha = 38.5$ (upper one).

standard SA using f^{opt} gets stuck in high-lying local energy minima where instead f^{flat} succeeds to reach zero energy came rather unexpected. All things considered, if on the one hand the present scheme is not sufficient in order to better control the behaviour of out-of-equilibrium procedures in the whole energy landscape when temperature is concerned, on the other hand we now realize there is no reason at all it should have to, the crucial point possibly being our lack of command onto the very distribution of violated gaps (the optimization of f^{opt} was indeed carried out only for $T = 0$). Understanding the peculiar atypical properties of glassy energy landscapes and the way they affect out-of-equilibrium behaviour still remains an intriguing, compelling open question.

Conclusions and Perspectives

In this thesis a relatively recent model [70, 80] has been explored from the perspective of random constraint satisfaction problems. Its connection with the well known q -coloring problem of random graphs has been at the center of our whole setting; in particular, we effectively showed that this model can be considered as a continuous variables extension of the former discrete random CSP. Similarly to the q -coloring, the continuous problem belongs to the RFOT class (random first order or discontinuous mean-field glass transition) for appropriate value of the parameters of the model. The relevant parameter is in this case the excluded angle (or particle diameter) θ , which is proportional to the inverse of the number of colors q . This implies that for small θ the transition becomes random first order. Transition lines can be precisely located through belief propagation numerical method, but are strongly dependent on the discretization precision d adopted. For $d = 10$ the model is shown to be RFO at least for values of θ smaller than $\theta = \frac{2\pi}{14}$.

What is the point in sponsoring a model as continuous, if one allows discretization to sneak in through the back door, the scrupulous reader might wonder. First, notice that discretization is only needed in order to numerically solve the BP equations and extract precise estimates of the thermodynamic quantities of interest. Even though the corrections decay fairly slowly, as $1/d$, it is reasonable to hope that some less precise but reliable extrapolation could still be obtained. However, the focus of this thesis is mostly about the physics around the dynamic or clustering transition. In this setting, there is no reason to believe that the effects induced by (a high) discretization should extend from the mere quantitative to also the qualitative level. For this reason, we preferred to support our numerical-analytical predictions from BP with direct numerical simulations of the discretized version of the continuous model. In doing so, we constantly kept as a privileged perspective the traditional field of discrete random CSPs, that is far from being settled. In particular, the aim to understand and possibly control the algorithmic easy-hard transition for optimization procedures cannot prescind from an accurate determination of the various thermodynamic thresholds, which may then prove to be relevant or not in order to explain algorithmic performance.

Our results allow to interpret and characterize the atypicality of solutions found by a simple zero energy MC procedure in terms of the observed distribution of angular distances between neighbours on the graph. This is particularly convenient, since the model is also equivalent to a 1-dimensional realization of hard-spheres, and the angular distance between variables directly enters the definition of the Hamiltonian. At any rate, the same *biased thermodynamics* approach followed in this thesis can be applied also to more traditional discrete random CSPs such as k -uniform hypergraph bicoloring [18]. In this case, the relevant quantity playing the role of the angular distance in our model, is the number of variables in each clause taking the same reference value, *e.g.* -1 . The common point between the two models, is that in both cases it is possible to quantify how much a given

set of variables participating in a same interaction is far from satisfying the constraint in a tight way, being this quantity a physical distance or not. We believe that where this is possible, then the following result from this thesis should apply: out-of-equilibrium MC procedures moving from solution to solution while increasing the variable node connectivity are naturally attracted towards configurations that satisfy the constraints in a tighter way, and for this reason they can succeed even beyond α_d .

This opens up to at least two complementary lines of development. On the one hand, there is the question of how our results compare with other, usually more efficient, optimization strategies, such as message passing (*e.g.* SP guided decimation [16]) or Entropy-driven MC [6]. Is the gap distribution still relevant in order to describe the solutions found by these algorithms? How does their performance relate with what we have called the optimized or optimum dynamic threshold α_d^{opt} ? This kind of questions should apply to both the continuous coloring and the hypergraph bicoloring. As a first step in this direction, it would be essential to compute the SAT/UNSAT threshold for the continuous coloring and compare it to α_d^{opt} . On the other hand, an important role in the comprehension of usual random CSPs is played by the notions of rigidity and freezing transitions. We somewhat avoided to analyze this aspect, that would require a non-immediate generalization of the concept of frozen variables to the continuous case. A perhaps better candidate to investigate the relation between rigidity, biased measure (according to the definition we have adopted) and algorithmic performance would be for a start a discrete model such as hypergraph bicoloring [15, 18].

Another interesting perspective regards the exploration of the energy landscape by the use of procedures that admit some fraction of violated constraints. In this respect, unfortunately, our f^{opt} interaction, that was optimized only at $T = 0$, proved not to be a resolutive choice. However, we can still learn something important from our analysis: out-of-equilibrium procedures, such as a MC quench to zero temperature or simulated annealing, appear to visit configurations whose atypicality we are able to qualify and quantify through the distribution of both positive and negative gaps between nearest neighbours on the graph. This suggests that by appropriately biasing the measure we might be able to study the (topological) properties of the landscape actually relevant to out-of-equilibrium behaviours, *e.g.* to mean-field aging. An interesting strategy in this respect would be to try to optimize the interaction during the simulation itself by measuring the gap distribution along the run. More generally, our analysis highlights the need of combining a soft bias, as the one adopted throughout this thesis, with also a modification of the hard-core energy term, which could in turn be optimized in a similar way to the one discussed in this thesis for the soft bias. Studying the typical properties in this modified setting might cast some light on the atypical properties of the original landscape, possibly improving algorithmic performance.

In Chapter 6 the maximization of the complexity has been proposed as a handy powerful method to automatize the optimization of the bias in order to postpone as much as possible the dynamic transition α_d . We have managed to approximate the soft bias with order 10^2 independent parameters by discretizing the interval $[0, 2\pi)$. Despite the very promising results, we lack a rigorous proof connecting the complexity maximization with that of α_d . Ideally, the infinite dimensional hard-spheres model, where an analytic expression for α_d is known and can be used to build a rigorous optimization procedure [69], could contribute to shed some light on the validity or the limitations of the proposed approach.

In sections 4.1.3 and 4.1.4 we have given the definition of a simple Bayes optimal inference problem derived from the continuous coloring Hamiltonian. This model could be very profitably studied from the point of view of inference. In particular, our analysis showed that the continuous coloring, having a smaller α_{KS} threshold with respect to its discrete counterpart, should represent an easier inference problem. On the other hand, our improved f^{opt} , which was obtained in order to postpone α_d and favour the optimization problem, is disadvantageous to inference, as it also increases α_{KS} . It would be hence interesting in this perspective to repeat the optimization of the interaction in the opposite direction with the aim of minimizing α_{KS} .

Finally, the continuous coloring, analyzed in this thesis around the dynamic or clustering point α_d , should as well present a jamming transition in correspondence of its zero temperature SAT/UNSAT threshold. The MKK model in $d = 2$ was precisely introduced in [70] as a sparse model for jamming. This led the authors of [80] to address this point in a similar model (they show results for the continuous coloring at connectivities where the transition is continuous, or still for the $d = 1$ case but with multi-body interactions) by the use of the BP methodology: on the one hand, they introduced a couple of approximate schemes in order to solve analytically the equations, but needing numerical support; on the other hand, a numerical solution of the BP equations suffered from the necessary presence of discretization, spoiling the continuous nature of the model. A progress in this direction could be likely made by a direct numerical simulation with continuous variables, which is particularly feasible in low dimensions thanks to the sparsity of the interactions.

Bibliography

- [1] G. Adam and J. H. Gibbs. On the Temperature Dependence of Cooperative Relaxation Properties in Glass-Forming Liquids. *The Journal of Chemical Physics*, 43(1):139–146, July 1965.
- [2] A. Altieri, M. C. Angelini, C. Lucibello, G. Parisi, F. Ricci-Tersenghi, and T. Rizzo. Loop expansion around the Bethe approximation through the M-layer construction. *Journal of Statistical Mechanics: Theory and Experiment*, 2017(11):113303, Nov. 2017.
- [3] F. Arceri, F. P. Landes, L. Berthier, and G. Biroli. Glasses and aging: A Statistical Mechanics Perspective. *arXiv:2006.09725 [cond-mat]*, Oct. 2020.
- [4] J. Ardelius and E. Aurell. Behavior of heuristics on large and hard satisfiability problems. *Physical Review E*, 74(3):037702, Sept. 2006.
- [5] C. Baldassi, A. Ingrosso, C. Lucibello, L. Saglietti, and R. Zecchina. Subdominant Dense Clusters Allow for Simple Learning and High Computational Performance in Neural Networks with Discrete Synapses. *Physical Review Letters*, 115(12):128101, Sept. 2015.
- [6] C. Baldassi, A. Ingrosso, C. Lucibello, L. Saglietti, and R. Zecchina. Local entropy as a measure for sampling solutions in Constraint Satisfaction Problems. *Journal of Statistical Mechanics: Theory and Experiment*, 2016(2):023301, Feb. 2016.
- [7] H. Bässler. Viscous flow in supercooled liquids analyzed in terms of transport theory for random media with energetic disorder. *Physical Review Letters*, 58(8):767–770, Feb. 1987.
- [8] U. Bengtzelius, W. Gotze, and A. Sjolander. Dynamics of supercooled liquids and the glass transition. 1984.
- [9] L. Berthier, P. Charbonneau, and S. Yaida. Efficient measurement of point-to-set correlations and overlap fluctuations in glass-forming liquids. *The Journal of Chemical Physics*, 144(2):024501, Jan. 2016.
- [10] L. Berthier and W. Kob. Static point-to-set correlations in glass-forming liquids. *Physical Review E*, 85(1):011102, Jan. 2012.
- [11] G. Biroli, J.-P. Bouchaud, A. Cavagna, T. S. Grigera, and P. Verrocchio. Thermodynamic signature of growing amorphous order in glass-forming liquids. *Nature Physics*, 4(10):771–775, Oct. 2008.

- [12] G. Biroli and M. Mézard. Lattice Glass Models. *Physical Review Letters*, 88(2):025501, Dec. 2001.
- [13] J.-P. Bouchaud and G. Biroli. On the Adam-Gibbs-Kirkpatrick-Thirumalai-Wolynes scenario for the viscosity increase in glasses. *The Journal of Chemical Physics*, 121(15):7347–7354, Oct. 2004.
- [14] J. P. Bouchaud and M. Mézard. Self induced quenched disorder: A model for the glass transition. *Journal de Physique I*, 4(8):1109–1114, Aug. 1994.
- [15] A. Braunstein, L. Dall’Asta, G. Semerjian, and L. Zdeborová. The large deviations of the whitening process in random constraint satisfaction problems. *Journal of Statistical Mechanics: Theory and Experiment*, 2016(5):053401, May 2016.
- [16] A. Braunstein, M. Mezard, and R. Zecchina. Survey propagation: An algorithm for satisfiability. *arXiv:cs/0212002*, Apr. 2006.
- [17] A. Braunstein, R. Mulet, A. Pagnani, M. Weigt, and R. Zecchina. Polynomial iterative algorithms for coloring and analyzing random graphs. *Physical Review E*, 68(3):036702, Sept. 2003.
- [18] L. Budzynski, F. Ricci-Tersenghi, and Guilhem Semerjian. Biased landscapes for random constraint satisfaction problems. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(2):023302, Feb. 2019.
- [19] L. Budzynski and G. Semerjian. Biased measures for random constraint satisfaction problems: Larger interaction range and asymptotic expansion. *Journal of Statistical Mechanics: Theory and Experiment*, 2020(10):103406, Oct. 2020.
- [20] F. Caltagirone, U. Ferrari, L. Leuzzi, G. Parisi, F. Ricci-Tersenghi, and T. Rizzo. Critical Slowing Down Exponents of Mode Coupling Theory. *Physical Review Letters*, 108(8):085702, Feb. 2012.
- [21] C. Cammarota and A. Cavagna. A novel method for evaluating the critical nucleus and the surface tension in systems with first order phase transition. *The Journal of Chemical Physics*, 127(21):214703, Dec. 2007.
- [22] T. Castellani and A. Cavagna. Spin-glass theory for pedestrians. *Journal of Statistical Mechanics: Theory and Experiment*, 2005(05):P05012, May 2005.
- [23] A. Cavagna. Supercooled Liquids for Pedestrians. *Physics Reports*, 476(4-6):51–124, June 2009.
- [24] A. Cavagna, I. Giardina, and G. Parisi. Stationary points of the Thouless-Anderson-Palmer free energy. *Physical Review B*, 57(18):11251–11257, May 1998.
- [25] P. Charbonneau, Y. Jin, G. Parisi, and F. Zamponi. Hopping and the Stokes–Einstein relation breakdown in simple glass formers. *Proceedings of the National Academy of Sciences*, 111(42):15025–15030, Oct. 2014.

- [26] P. Charbonneau and J. Kundu. Postponing the dynamical transition density using competing interactions. *Granular Matter*, 22(3):55, May 2020.
- [27] M. P. Ciamarra, M. Tarzia, A. de Candia, and A. Coniglio. Lattice glass model with no tendency to crystallize. *Physical Review E*, 67(5):057105, May 2003.
- [28] N. Creignou, H. Daudé, and O. Dubois. Approximating the Satisfiability Threshold for Random k-XOR-formulas. *Combinatorics, Probability and Computing*, 12(02), Mar. 2003.
- [29] A. Crisanti, H. Horner, and H. J. Sommers. The spherical-p-spin interaction spin-glass model. *Zeitschrift für Physik B Condensed Matter*, 92(2):257–271, June 1993.
- [30] A. Crisanti and H. J. Sommers. The spherical-p-spin interaction spin glass model: The statics. *Zeitschrift für Physik B Condensed Matter*, 87(3):341–354, Oct. 1992.
- [31] L. Dall’Asta, A. Ramezanpour, and R. Zecchina. Entropy landscape and non-Gibbs solutions in constraint satisfaction problems. *Physical Review E*, 77(3):031118, Mar. 2008.
- [32] G. Das, N. Gnan, F. Sciortino, and E. Zaccarelli. Unveiling the complex glassy dynamics of square shoulder systems: Simulations and theory. *The Journal of Chemical Physics*, 138(13):134501, Apr. 2013.
- [33] K. Dawson, G. Foffi, M. Fuchs, W. Götze, F. Sciortino, M. Sperl, P. Tartaglia, T. Voigtmann, and E. Zaccarelli. Higher-order glass-transition singularities in colloidal systems with attractive interactions. *Physical Review E*, 63(1):011401, Dec. 2000.
- [34] J. R. L. de Almeida and D. J. Thouless. Stability of the Sherrington-Kirkpatrick solution of a spin glass model. *Journal of Physics A: Mathematical and General*, 11(5):983–990, May 1978.
- [35] C. De Dominicis. Dynamics as a substitute for replicas in systems with quenched random impurities. *Physical Review B*, 18(9):4913–4919, Nov. 1978.
- [36] P. G. Debenedetti and F. H. Stillinger. Supercooled liquids and the glass transition. *Nature*, 410(6825):259–267, Mar. 2001.
- [37] A. Decelle, F. Krzakala, C. Moore, and L. Zdeborová. Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. *Physical Review E*, 84(6):066106, Dec. 2011.
- [38] Erdős, P. and Rényi, A. On Random Graphs I. *Publicationes Mathematicae Debrecen*, 6, 1959.
- [39] U. Ferrari, L. Leuzzi, G. Parisi, and T. Rizzo. Two-step relaxation next to dynamic arrest in mean-field glasses: Spherical and Ising p-spin model. *Physical Review B*, 86(1):014204, July 2012.
- [40] S. Franz, M. Mézard, F. Ricci-Tersenghi, M. Weigt, and R. Zecchina. A ferromagnet with a glass transition. *Europhysics Letters (EPL)*, 55(4):465–471, Aug. 2001.
- [41] M. R. Garey and D. S. Johnson. *Computers and Intractability; A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., USA, 1990.

- [42] J. H. Gibbs. Nature of the Glass Transition in Polymers. *The Journal of Chemical Physics*, 25(1):185–186, July 1956.
- [43] J. H. Gibbs and E. A. DiMarzio. Nature of the Glass Transition and the Glassy State. *The Journal of Chemical Physics*, 28(3):373–383, Mar. 1958.
- [44] M. Goldstein. Viscous Liquids and the Glass Transition: A Potential Energy Barrier Picture. *The Journal of Chemical Physics*, 51(9):3728–3739, Nov. 1969.
- [45] Gotze. The essentials of the mode-coupling theory for glassy dynamics. *Condensed Matter Physics*, 1(4):873, 1998.
- [46] W. Götze. *Complex Dynamics of Glass-Forming Liquids: A Mode-Coupling Theory*. Oxford University Press.
- [47] T. Hecksher, A. I. Nielsen, N. B. Olsen, and J. C. Dyre. Little evidence for dynamic divergences in ultraviscous molecular liquids. *Nature Physics*, 4(9):737–741, Sept. 2008.
- [48] A. Ikeda and K. Miyazaki. Mode-coupling theory as a mean-field description of the glass transition. *Physical Review Letters*, 104(25):255704, June 2010.
- [49] S. Janson and E. Mossel. Robust Reconstruction on Trees Is Determined by the Second Eigenvalue. *The Annals of Probability*, 32(3):2630–2649, 2004.
- [50] Jean-Pierre Hansen and Ian R. McDonald. *Theory of Simple Liquids*. Elsevier, 2013.
- [51] G. P. Johari. A resolution for the enigma of a liquid’s configurational entropy-molecular kinetics relation. *The Journal of Chemical Physics*, 112(20):8958–8969, May 2000.
- [52] G. P. Johari and M. Goldstein. Viscous Liquids and the Glass Transition. II. Secondary Relaxations in Glasses of Rigid Molecules. *The Journal of Chemical Physics*, 53(6):2372–2388, Sept. 1970.
- [53] W. Kauzmann. The Nature of the Glassy State and the Behavior of Liquids at Low Temperatures. *Chemical Reviews*, 43(2):219–256, Oct. 1948.
- [54] T. R. Kirkpatrick and D. Thirumalai. Dynamics of the Structural Glass Transition and the p -Spin—Interaction Spin-Glass Model. *Physical Review Letters*, 58(20):2091–2094, May 1987.
- [55] T. R. Kirkpatrick and D. Thirumalai. p -spin-interaction spin-glass models: Connections with the structural glass problem. *Physical Review B*, 36(10):5388–5397, Oct. 1987.
- [56] T. R. Kirkpatrick, D. Thirumalai, and P. G. Wolynes. Scaling concepts for the dynamics of viscous liquids near an ideal glassy state. *Physical Review A*, 40(2):1045–1054, July 1989.
- [57] T. R. Kirkpatrick and P. G. Wolynes. Connections between some kinetic and equilibrium theories of the glass transition. *Physical Review A*, 35(7):3072–3080, Apr. 1987.
- [58] T. R. Kirkpatrick and P. G. Wolynes. Stable and metastable states in mean-field Potts and structural glasses. *Physical Review B*, 36(16):8552–8564, Dec. 1987.

- [59] F. Krzakala and J. Kurchan. Landscape analysis of constraint satisfaction problems. *Physical Review E*, 76(2):021122, Aug. 2007.
- [60] F. Krzakala, M. Mézard, and L. Zdeborová. Reweighted belief propagation and quiet planting for random K-SAT. *arXiv:1203.5521 [cond-mat]*, Feb. 2014.
- [61] F. Krzakala, A. Montanari, F. Ricci-Tersenghi, G. Semerjian, and L. Zdeborova. Gibbs states and the set of solutions of random constraint satisfaction problems. *Proceedings of the National Academy of Sciences*, 104(25):10318–10323, June 2007.
- [62] F. Krzakala, M. Tarzia, and L. Zdeborová. Lattice Model for Colloidal Gels and Glasses. *Physical Review Letters*, 101(16):165702, Oct. 2008.
- [63] F. Krzakala and L. Zdeborová. Hiding Quiet Solutions in Random Constraint Satisfaction Problems. *Physical Review Letters*, 102(23):238701, June 2009.
- [64] J. Kurchan and L. Laloux. Phase space geometry and slow dynamics. *Journal of Physics A: Mathematical and General*, 29(9):1929–1948, May 1996.
- [65] M. Leone, F. Ricci-Tersenghi, and R. Zecchina. Exact Solutions for Diluted Spin Glasses and Optimization Problems. *Physical review letters*, 87:127209, Oct. 2001.
- [66] L. Leuzzi and T. M. Nieuwenhuizen. *Thermodynamics of the Glassy State*. CRC Press, Oct. 2007.
- [67] C. Lupo and F. Ricci-Tersenghi. Approximating the XY model on a random graph with a q-state clock model. *Physical Review B*, 95(5):054433, Feb. 2017.
- [68] T. Maimbourg, J. Kurchan, and F. Zamponi. Solution of the Dynamics of Liquids in the Large-Dimensional Limit. *Physical Review Letters*, 116(1):015902, Jan. 2016.
- [69] T. Maimbourg, M. Sellitto, G. Semerjian, and F. Zamponi. Generating dense packings of hard spheres by soft interaction design. *SciPost Physics*, 4(6):039, June 2018.
- [70] R. Mari, F. Krzakala, and J. Kurchan. Jamming versus Glass Transitions. *Physical Review Letters*, 103(2):025701, July 2009.
- [71] R. Mari and J. Kurchan. Dynamical transition of glasses: From exact to approximate. *The Journal of Chemical Physics*, 135(12):124504, Sept. 2011.
- [72] R. Marino, G. Parisi, and F. Ricci-Tersenghi. The backtracking survey propagation algorithm for solving random K-SAT problems. *Nature Communications*, 7(1):12996, Oct. 2016.
- [73] P. C. Martin, E. D. Siggia, and H. A. Rose. Statistical Dynamics of Classical Systems. *Physical Review A*, 8(1):423–437, July 1973.
- [74] S. Mertens. Computational complexity for physicists. *Computing in Science Engineering*, 4(3):31–47, May 2002.
- [75] M. Mezard and A. Montanari. Reconstruction on trees and spin glass transition. *Journal of Statistical Physics*, 124(6):1317–1350, Oct. 2006.

- [76] M. Mézard and A. Montanari. *Information, Physics, and Computation*. Oxford Graduate Texts. Oxford Univ. Press, Oxford, 2012.
- [77] M. Mezard, M. Palassini, and O. Rivoire. Landscape of Solutions in Constraint Satisfaction Problems. *Physical review letters*, 95:200202, Dec. 2005.
- [78] M. Mezard and G. Parisi. The Bethe lattice spin glass revisited. *The European Physical Journal B*, 20(2):217–233, Mar. 2001.
- [79] M. Mézard and G. Parisi. The Cavity Method at Zero Temperature. *Journal of Statistical Physics*, 111(1):1–34, Apr. 2003.
- [80] M. Mezard, G. Parisi, M. Tarzia, and F. Zamponi. On the solution of a ‘solvable’ model of an ideal glass of hard spheres displaying a jamming transition. *arXiv:1011.5080 [cond-mat]*, Feb. 2011.
- [81] M. Mezard, G. Parisi, and M. A. Virasoro. *Spin Glass Theory And Beyond: An Introduction To The Replica Method And Its Applications*. World Scientific Publishing Company, Nov. 1987.
- [82] M. Mézard, G. Parisi, and R. Zecchina. Analytic and algorithmic solution of random satisfiability problems. *Science (New York, N.Y.)*, 297(5582):812–815, Aug. 2002.
- [83] M. Mézard and R. Zecchina. Random K-satisfiability problem: From an analytic solution to an efficient algorithm. *Physical Review E*, 66(5):056126, Nov. 2002.
- [84] R. Monasson. Structural Glass Transition and the Entropy of the Metastable States. *Physical Review Letters*, 75(15):2847–2850, Oct. 1995.
- [85] R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, and L. Troyansky. Determining computational complexity from characteristic ‘phase transitions’. *Nature*, 400(6740):133–137, July 1999.
- [86] A. Montanari, F. Ricci-Tersenghi, and G. Semerjian. Clusters of solutions and replica symmetry breaking in random k -satisfiability. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(04):P04004, Apr. 2008.
- [87] A. Montanari and G. Semerjian. On the dynamics of the glass transition on Bethe lattices. *Journal of Statistical Physics*, 124(1):103–189, July 2006.
- [88] A. Montanari and G. Semerjian. Rigorous Inequalities between Length and Time Scales in Glassy Systems. *Journal of Statistical Physics*, 125(1):23–54, Nov. 2006.
- [89] R. Mulet, A. Pagnani, M. Weigt, and R. Zecchina. Coloring Random Graphs. *Physical Review Letters*, 89(26):268701, Dec. 2002.
- [90] G. Parisi, P. Urbani, and F. Zamponi. *Theory of Simple Glasses: Exact Solutions in Infinite Dimensions*. Cambridge University Press, Cambridge, 2020.
- [91] G. Parisi and F. Zamponi. On the High Density Behavior of Hamming Codes with Fixed Minimum Distance. *Journal of Statistical Physics*, 123(6):1145–1167, June 2006.

- [92] G. Parisi and F. Zamponi. Mean-field theory of hard sphere glasses and jamming. *Reviews of Modern Physics*, 82(1):789–845, Mar. 2010.
- [93] F. Ricci-Tersenghi, M. Weigt, and R. Zecchina. Simplest random K-satisfiability problem. *Physical Review E*, 63(2):026702, Jan. 2001.
- [94] O. Rivoire, G. Biroli, O. C. Martin, and M. Mézard. Glass models on Bethe lattices. *The European Physical Journal B*, 37(1):55–78, Jan. 2004.
- [95] F. Sciortino. One liquid, two glasses. *Nature Materials*, 1(3):145–146, Nov. 2002.
- [96] S. Seitz, M. Alava, and P. Orponen. Focused local search for random 3-satisfiability. *Journal of Statistical Mechanics: Theory and Experiment*, 2005(06):P06006, June 2005.
- [97] M. Sellitto and F. Zamponi. A thermodynamic description of colloidal glasses. *EPL (Europhysics Letters)*, 103(4):46005, Aug. 2013.
- [98] G. Semerjian. On the Freezing of Variables in Random Constraint Satisfaction Problems. *Journal of Statistical Physics*, 130(2):251–293, Jan. 2008.
- [99] J. van Mourik and D. Saad. Random graph coloring: Statistical physics approach. *Physical Review E*, 66(5):056120, Nov. 2002.
- [100] H. Yoshino. Disorder-free spin glass transitions and jamming in exactly solvable mean-field models. *SciPost Physics*, 4(6):040, June 2018.
- [101] L. Zdeborová and F. Krzakała. Phase transitions in the coloring of random graphs. *Physical Review E*, 76(3):031131, Sept. 2007.
- [102] L. Zdeborová and F. Krzakala. Generalization of the cavity method for adiabatic evolution of Gibbs states. *Physical Review B*, 81(22):224205, June 2010.
- [103] L. Zdeborova and M. Mezard. Constraint satisfaction problems with isolated solutions are hard. *Journal of Statistical Mechanics Theory and Experiment*, 2008, Nov. 2008.
- [104] L. Zdeborová and M. Mézard. Locked Constraint Satisfaction Problems. *Physical Review Letters*, 101(7):078702, Aug. 2008.

Appendix A

1RSB equations for $m = 1$

In this appendix we show how to obtain equation (3.21) when the replica parameter $m = 1$ for a general model with pairwise interactions, and under the additional assumption that its RS solution is paramagnetic, *i.e.* the RS fixed point messages are uniform over the alphabet. Firstly, let us recall some useful definitions. The BP iteration on a given graph eq. (3.11) can be framed into an update function

$$\Phi_{\text{update}}(x; \{\nu_k\}_d) = \frac{1}{Z_0(\{\nu_k\}_d)} \prod_{k=1}^d \sum_{x_k} \psi(x, x_k) \nu_k(x_k) \quad (\text{A.1})$$

which takes as input d arbitrary *messages* and returns a new message (it is hence a function of x , $x \in \mathcal{X}$ with $q = |\mathcal{X}|$ the size of the alphabet). The single-site term $\psi_j(x_j)$ in (3.11) has been disregarded since we do not want to break the symmetry of the paramagnetic state. Furthermore, in order to lighten the notation, we omit in the following the explicit dependence from the disorder (possibly entering through the pairwise term ψ_{ij} , see for instance eq. (3.15)). The normalization $Z_0(\{\nu_k\}_d)$ reads then

$$Z_0(\{\nu_k\}_d) = \sum_x \prod_{k=1}^d \sum_{x_k} \psi(x, x_k) \nu_k(x_k). \quad (\text{A.2})$$

Analogously, the 1RSB BP iteration eq. (3.18) defines a second update function for the Q 's

$$\mathcal{F}_2(\nu; \{Q_k\}_d) = \frac{1}{Z_1(\{Q_k\}_d)} \int \prod_{k=1}^d d\nu_k Q_k \delta(\nu - \Phi_{\text{update}}(\{\nu_k\}_d)) Z_0(\{\nu_k\}_d), \quad (\text{A.3})$$

which takes as input d arbitrary probability distributions $\{Q_k(\nu_k)\}_d$ and returns a new one. The normalization constant $Z_1(\{Q_k\}_d)$ can be computed as the integral over $d\nu$ of the numerator of expression (A.3) and simply reads

$$Z_1(\{Q_k\}_d) = \int \prod_{k=1}^d d\nu_k Q_k(\nu_k) Z_0(\{\nu_k\}_d). \quad (\text{A.4})$$

Equation (3.18) can still be averaged over the ensemble of random graphs (and the external disorder which we here omit). This leads to the definition of a probability distribution (over the random choice of an edge) of the probability distribution (over the states) of messages, $\mathcal{P}[Q(\nu)]$,

satisfying

$$\mathcal{P}[Q(\nu)] = \mathbb{E}_d \int \prod_{k=1}^d dQ_k(\nu_k) \mathcal{P}[Q_k(\nu_k)] \delta(Q(\nu) - \mathcal{F}_2(\nu; \{Q_k\}_d)), \quad (\text{A.5})$$

where \mathbb{E}_d stands for the average over $\mathbb{P}_{\text{edge}}(d)$ defined in (3.16).

It is useful at this point to define, similarly to eq. (3.19), the average message $\bar{\nu} \equiv \int d\nu Q(\nu)\nu$ which itself is a random (scalar) variable since $Q(\nu)$ is random. One can show that $\bar{\nu}$ is a solution of the simple RS population dynamics equation (3.15). To this end, let us compute the induced probability distribution $\mathbb{P}(\bar{\nu})$ for $\bar{\nu}$,

$$\begin{aligned} \mathbb{P}(\bar{\nu}) &= \int dQ(\nu) \mathcal{P}[Q(\nu)] \delta\left(\bar{\nu} - \int d\nu Q(\nu)\nu\right) = \\ &= \int dQ(\nu) \mathbb{E}_d \int \prod_{k=1}^d dQ_k(\nu_k) \mathcal{P}[Q_k(\nu_k)] \delta(Q(\nu) - \mathcal{F}_2(\nu; \{Q_k\}_d)) \delta\left(\bar{\nu} - \int d\nu Q(\nu)\nu\right) = \\ &= \mathbb{E}_d \int \prod_{k=1}^d dQ_k(\nu_k) \mathcal{P}[Q_k(\nu_k)] \delta\left(\bar{\nu} - \int d\nu \mathcal{F}_2(\nu; \{Q_k\}_d)\nu\right). \end{aligned} \quad (\text{A.6})$$

The integral inside the last delta function reads

$$\begin{aligned} \int d\nu \mathcal{F}_2(\nu; \{Q_k\}_d)\nu &= \frac{1}{Z_1(\{Q_k\}_d)} \int \prod_{k=1}^d d\nu_k Q_k(\nu_k) \Phi_{\text{update}}(\{\nu_k\}_d) Z_0(\{\nu_k\}_d) = \\ &= \frac{1}{Z_1(\{Q_k\}_d)} \Phi_{\text{update}}(\{\bar{\nu}_k\}_d) Z_0(\{\bar{\nu}_k\}_d) = \Phi_{\text{update}}(\{\bar{\nu}_k\}_d), \end{aligned} \quad (\text{A.7})$$

having used in the last equality the fact that

$$\begin{aligned} Z_0(\{\bar{\nu}_k\}_d) &= \sum_x \prod_{k=1}^d \sum_{x_k} \psi(x, x_k) \int d\nu_k Q_k(\nu_k) \nu_k(x_k) = \\ &= \int \prod_{k=1}^d d\nu_k Q_k(\nu_k) Z_0(\{\nu_k\}_d) = Z_1(\{Q_k\}_d). \end{aligned} \quad (\text{A.8})$$

One should also keep in mind that at this point $\Phi_{\text{update}}(\{\bar{\nu}_k\}_d) = \Phi_{\text{update}}(\{\bar{\nu}_k(Q_k)\}_d)$, with $\bar{\nu}_k(Q_k) \equiv \int d\nu_k Q_k(\nu_k)\nu_k$. By substituting (A.7) into (A.6) and inserting a factor $1 = \int \prod_k d\bar{\nu}_k \delta(\bar{\nu}_k - \int d\nu_k Q_k(\nu_k)\nu_k)$ inside the integral, we finally obtain

$$\begin{aligned} \mathbb{P}(\bar{\nu}) &= \mathbb{E}_d \int \prod_{k=1}^d \left[dQ_k(\nu_k) \mathcal{P}[Q_k(\nu_k)] d\bar{\nu}_k \delta\left(\bar{\nu}_k - \int d\nu_k Q_k(\nu_k)\nu_k\right) \right] \delta\left(\bar{\nu} - \Phi_{\text{update}}(\{\bar{\nu}_k(Q_k)\}_d)\right) = \\ &= \mathbb{E}_d \int \prod_{k=1}^d \left[d\bar{\nu}_k dQ_k(\nu_k) \mathcal{P}[Q_k(\nu_k)] \delta\left(\bar{\nu}_k - \int d\nu_k Q_k(\nu_k)\nu_k\right) \right] \delta\left(\bar{\nu} - \Phi_{\text{update}}(\{\bar{\nu}_k\}_d)\right) = \\ &= \mathbb{E}_d \int \prod_{k=1}^d d\bar{\nu}_k \mathbb{P}(\bar{\nu}_k) \delta\left(\bar{\nu} - \Phi_{\text{update}}(\{\bar{\nu}_k\}_d)\right). \end{aligned} \quad (\text{A.9})$$

Since the RS solution is paramagnetic by hypothesis, one has simply $\mathbb{P}(\bar{\nu}) = \delta(\bar{\nu} - \frac{1}{q})$, and hence $\bar{\nu}(x; Q) \equiv \bar{\nu}(x) = q^{-1} \quad \forall x$ independently on Q . This allows one to simply in this case the

expression (A.4) for $Z_1(\{Q_k\}_d)$, that ceases to depend on the d distributions $\{Q_k\}_d$

$$\begin{aligned} Z_1(\{Q_k\}_d) &= Z_0(\{\bar{\nu}_k\}_d) = \sum_x \prod_{k=1}^d \sum_{x_k} \psi(x, x_k) \frac{1}{q} = \\ &= q^{-d} \sum_x \prod_{k=1}^d \sum_{x_k} \psi(x, x_k) \equiv Z_1. \end{aligned} \quad (\text{A.10})$$

We are now ready to define the averaged probability distribution $\bar{P}(\nu)$ as the average of $Q(\nu)$ over $\mathcal{P}[Q(\nu)]$. Using expressions (A.5) and (A.3) one obtains

$$\begin{aligned} \bar{P}(\nu) &\equiv \int dQ(\nu) \mathcal{P}[Q(\nu)] Q(\nu) = \mathbb{E}_d \int \prod_{k=1}^d dQ_k(\nu_k) \mathcal{P}[Q_k(\nu_k)] \mathcal{F}_2(\nu; \{Q_k\}_d) = \\ &= \mathbb{E}_d \int \prod_{k=1}^d dQ_k(\nu_k) \mathcal{P}[Q_k(\nu_k)] \frac{1}{Z_1} \int \prod_{k=1}^d d\nu_k Q_k(\nu_k) \delta(\nu - \Phi_{\text{update}}(\{\nu_k\}_d)) Z_0(\{\nu_k\}_d) = \\ &= \mathbb{E}_d \frac{1}{Z_1} \int \prod_{k=1}^d d\nu_k \bar{P}(\nu_k) \delta(\nu - \Phi_{\text{update}}(\{\nu_k\}_d)) Z_0(\{\nu_k\}_d). \end{aligned} \quad (\text{A.11})$$

In order to get rid of the term Z_0 in (A.11), it is convenient to introduce q distribution functions

$$R_x(\nu) = q\nu(x)\bar{P}(\nu), \quad x \in \mathcal{X}. \quad (\text{A.12})$$

Notice that they are correctly normalized since $\int d\nu R_x(\nu) = \int d\nu dQ(\nu) \mathcal{P}[Q(\nu)] q\nu(x) Q(\nu) = q\bar{\nu}(x) = 1$. Substituting eq. (A.11) it follows

$$\begin{aligned} R_x(\nu) &= \mathbb{E}_d \frac{q}{Z_1} \int \prod_{k=1}^d d\nu_k \bar{P}(\nu_k) \delta(\nu - \Phi_{\text{update}}(\{\nu_k\}_d)) Z_0(\{\nu_k\}_d) \Phi_{\text{update}}(x; \{\nu_k\}_d) = \\ &= \mathbb{E}_d \frac{q}{Z_1} \int \prod_{k=1}^d [d\nu_k \bar{P}(\nu_k)] \delta(\nu - \Phi_{\text{update}}(\{\nu_k\}_d)) \prod_{k=1}^d \sum_{x_k} \psi(x, x_k) \nu_k(x_k) = \\ &= \mathbb{E}_d \frac{q^{1-d}}{Z_1} \int \prod_{k=1}^d [d\nu_k \sum_{x_k} \psi(x, x_k) R_{x_k}(\nu_k)] \delta(\nu - \Phi_{\text{update}}(\{\nu_k\}_d)). \end{aligned} \quad (\text{A.13})$$

By inspecting (A.10), one realizes that $\sum_{x_k} \psi(x, x_k)$ should not depend on x in order to preserve the symmetry of the paramagnetic state. We can thus rewrite (A.10) as $Z_1 = q^{1-d} (\sum_{x_k} \psi(x, x_k))^d$ for any arbitrary x . This allows to simplify eq. (A.13) into

$$R_x(\nu) = \mathbb{E}_d \int \prod_{k=1}^d \left[d\nu_k \sum_{x_k} \frac{\psi(x, x_k)}{\sum_{x'} \psi(x, x')} R_{x_k}(\nu_k) \right] \delta(\nu - \Phi_{\text{update}}(\{\nu_k\}_d)), \quad (\text{A.14})$$

which is the formula we wanted to prove.