



HAL
open science

A statistical physics approach to different problems in network theory

Alberto Guggiola

► **To cite this version:**

Alberto Guggiola. A statistical physics approach to different problems in network theory. Physics [physics]. Ecole normale supérieure - ENS PARIS, 2015. English. NNT: 2015ENSU0034. tel-01668590

HAL Id: tel-01668590

<https://theses.hal.science/tel-01668590>

Submitted on 20 Dec 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

THÈSE DE DOCTORAT

En vue de l'obtention du grade de

DOCTEUR DE L'ÉCOLE NORMALE SUPÉRIEURE

École Doctorale 564 : Physique en Île de France
Discipline : Physique



Présentée et soutenue le 26 Octobre 2015 par
ALBERTO GUGGIOLA

Une approche physique-statistique à différents problèmes dans la théorie des réseaux

Laboratoire de Physique Théorique
de l'École Normale Supérieure de Paris

Membres du jury :

Alain BARRAT	CPT Marseille	Rapporteur
Florent KRZAKALA	LPS-ENS Paris	
Matteo MARSILI	ICTP Trieste	
Rémi MONASSON	LPT-ENS Paris	Invité
David SAAD	Aston University	Rapporteur
Guilhem SEMERJIAN	LPT-ENS Paris	Directeur de Thèse

Numéro identifiant de la thèse : 79840

Acknowledgements

First of all, a particular thank goes to Guilhem Semerjian, for having accepted to supervise my doctoral studies, and for having helped me during these three tough but exciting years. The fact of literally finding his office door always open has been a wonderful incitement to fruitful discussion throughout my Ph.D., and his help has proven to be fundamental in many aspects. In particular, I would like to thank him for the very attentive reading of this manuscript (even in the middle of the summer!) and for the comments on it that enabled me to hopefully improve it a bit from its original version. From a more practical point of view, the help of Viviane Sebillé and Pascale Searle have been essential as regards the administrative tasks, not always straightforward.

Thanks to the Netadis network, I have had during these years the possibility of working on several projects, by collaborating with some of the best research groups at a European level. Even if this superposition has proven sometimes to be not so easy to manage, I have found it an extremely enriching experience both from a scientific and from a personal point of view. As regards this, I warmly thank the PIs with which I have had the opportunity to work with, namely Rémi Monasson (again here at the ENS Paris), Reimer Kühn (at the King's College in London), Matteo Marsili (at the ICTP in Trieste) and Riccardo Zecchina and his group at the Politecnico in Torino.

In the Netadis network, also the people with whom I did not have the chance of directly collaborating played a central role. From the informal discussion while having a beer, to the possibility of following one of their lectures during the wonderful schools and conferences we could take part into, the range of interesting interactions has proven to be quite variegated. With the other Netadis students, moreover, we little by little started to feel as schoolmates, and some good friendships were built; great thanks in particular to Caterina and Silvia, especially for the perfect organisation of the barbecues in Orsay. Of a similar importance for the survival during

these years have been the other students and researchers met in the ENS, with which I sincerely hope to be able and keep in contact in the future; in rigorous random order, I want to thank at least Charles, Alice, Alaa, Gaia, Corrado, Hugo and Dario.

I warmly thank Florent Krzakala and Matteo Marsili for having accepted to participate to the jury of this defence, as well as David Saad and Alain Barrat for carefully reading my manuscript.

Special thanks also go to the professors who were able to transmit their enthusiasm for physics and research activity to me; among them, Michele Caselle, Ciro Cattuto and Andrea Pagnani.

I also thank my family for their long-range support they offered me during these last three years, and for the much shorter-range support for the previous 24. The same goes for all my friends, magically gathering each time I manage to go back to Italy for a couple of days, and in some cases reaching me here in Paris.

Last but not least, thanks to Eleonora. One adventure is (almost) over; the next is ready to start.

A statistical physics approach to different problems in network theory

Statistical physics, originally developed to describe thermodynamic systems, has been playing for the last decades a central role in modelling an incredibly large and heterogeneous set of different phenomena taking for instance place on social, economical or biological systems. Such a vast field of possible applications has been found also for networks, as a huge variety of systems can be described in terms of interconnected elements. After an introductory part introducing these themes as well as the role of abstract modelling in science, in this dissertation it will be discussed how a statistical physics approach can lead to new insights as regards three problems of interest in network theory: how some quantity can be optimally spread on a graph, how to explore it and how to reconstruct it from partial information. Some final remarks on the importance such themes will likely preserve in the coming years conclude the work.

Keywords: Network theory, Statistical physics, Disordered systems, Inference, Spreading dynamics, Extreme events

Contents

1	Models and frameworks	5
1.1	The role of models	5
1.2	Forward problems, inverse problems	6
1.2.1	Same phenomena, different questions	6
1.2.2	A more precise definition	8
1.3	A physics modelling framework: the Ising model	10
1.3.1	Motivation	10
1.3.2	The Ising model	11
1.3.3	Universality of critical phenomena	14
1.3.4	The generalisations	15
1.3.5	Inverse problems in disordered systems	18
1.4	A computer science modelling framework: the constraint satisfaction problems	19
1.4.1	Definitions and examples	19
1.4.2	Infeasibility of the straightforward solutions	21
1.4.3	The worst case scenario	22
1.4.4	The typical case scenario	24
1.5	Spin glasses, rCSP and large deviations	27
1.5.1	Spin glasses and CSP: two languages, same problems	27
1.5.2	Looking for atypical events	28
1.5.3	The large deviations in a simple case	30
1.6	Statistical mechanics for non-physics problems	31
1.6.1	A common framework, many fields of interest	31
1.6.2	Social systems	32
1.6.3	Economics and finance	33
1.6.4	Neural networks	33
1.6.5	Other biological systems	34
1.6.6	Signal processing	35

1.6.7	Computer science	35
2	The network theory	37
2.1	Motivations and definitions	37
2.1.1	Ubiquity of networks	37
2.1.2	The birth of graph theory	38
2.1.3	The main features of a graph	39
2.1.4	Enriching the network metaphor	41
2.2	A path towards more realistic networks	42
2.2.1	The comparison between observed and reproduced features	42
2.2.2	Reproducing the observed average distance and clustering coefficient	44
2.2.3	Reproducing the observed degree distribution	46
2.3	Studying graphs as a branch of probability theory	48
2.3.1	A mathematical description	48
2.3.2	Combinatorial problems on graphs	49
2.4	Phase transitions, critical phenomena and statistical physics approach	51
2.4.1	Dynamics on networks, dynamics of networks	51
2.4.2	Statistical physics and network theory	52
2.4.3	Critical phenomena on networks	53
2.5	Examples of other problems	55
2.5.1	The detection of communities	55
2.5.2	Networked versions of other problems	57
3	Contagion dynamics on graphs	59
3.1	General framework	59
3.1.1	The problem and its range of applicability	59
3.1.2	Epidemic processes and compartmental models	61
3.1.3	The choice of the dynamics	61
3.1.4	The characteristics of the graph	63
3.1.5	One spreading phenomenon, many possible questions	64
3.1.6	Different approaches	66
3.2	Minimal contagious sets in random regular graphs	67
3.2.1	Definition of the problem	67
3.2.2	The energy function	69
3.2.3	Mapping to other standard problems in graph theory	69
3.3	The cavity method treatment	70

3.3.1	The replica symmetric RS formalism	70
3.3.2	The breaking down of the RS assumption	71
3.3.3	1RSB and energetic 1RSB formalism	73
3.3.4	The “energetic” 1RSB formalism	74
3.4	Main results	74
3.4.1	The solutions of the problem	74
3.4.2	Analytical results	75
3.4.3	Numerical results	76
3.5	Future perspectives	77
4	Exploring networks	79
4.1	Graph exploration and random walks	79
4.1.1	Exploring a graph, a very general problem	79
4.1.2	Interest and applications in different fields	79
4.1.3	Random walks on graphs	81
4.1.4	Joining the two: graph exploration through random walks	82
4.2	Rare event statistics on random walks on networks	84
4.2.1	The problem	84
4.2.2	The model	85
4.2.3	The large deviations	86
4.3	Main results	87
4.3.1	A degree-based approximation	87
4.3.2	Localisation transition	88
4.3.3	Mode-switching transition	90
4.4	Summary and future perspectives	94
4.4.1	Summary	94
4.4.2	Better understanding of the critical behaviour	94
4.4.3	Different functions	95
4.4.4	Functions taking value on the links	96
4.4.5	Different topologies	96
4.4.6	Improvement of clustering techniques	97
5	Inferring the graph	99
5.1	Motivation, applications and connection to other problems	99
5.1.1	Motivation	99
5.1.2	The matrix completion problem	100
5.1.3	The collaborative filtering	103

5.2	The inference of a network of interactions from a partial knowledge of the correlations	105
5.2.1	Definition of the problem	105
5.2.2	Motivation, applications and previous work	106
5.3	The model	107
5.3.1	Statistical context of the project	107
5.3.2	Prior, posterior and optimal estimate of the model	108
5.4	Choosing the couplings and evaluating the performances	110
5.4.1	Artificial models	110
5.4.2	The geometry of J	111
5.4.3	Generating synthetic J	113
5.4.4	The observables	114
5.5	Main analytical results	117
5.5.1	Reasons justifying the study of the Hessian	117
5.5.2	Derivation of the Hessian	118
5.5.3	Derivation of the Hessian in the eigenbasis of C	120
5.5.4	Perturbative approximation of the Hessian	121
5.6	The heuristics	122
5.6.1	Choices based on inference	122
5.6.2	Heuristics making use of the Hessian	123
5.6.3	Intractability of the complete Hessian	124
5.6.4	Diagonal elements of the Hessian	125
5.6.5	Validity of the perturbative approximation	126
5.7	Initial knowledge and effects of measurements	127
5.7.1	Choice of the initial condition	127
5.7.2	Update of C after a measurement	129
5.8	Results	129
5.8.1	Dependency on the initial condition	130
5.8.2	The exploration-exploitation trade-off	130
5.8.3	Different heuristics on a fixed geometry	132
5.8.4	Different graphs, different features	136
5.9	Future perspectives	138
5.9.1	Biological applications	138
5.9.2	Theoretical challenges	139
5.9.3	More complicated heuristics	140
5.10	Appendices	140

5.10.1	Caveat on the procedure for creating synthetic J	140
5.10.2	The value of the couplings	143
5.10.3	Regularisation and noise	144
5.10.4	Uniqueness of the minimum	145
5.10.5	Completion of the initial condition	147
6	Conclusions and perspectives	151
6.1	The network theory in the “big data” era	151
6.1.1	The end of theory?	151
6.1.2	An example of “rich data” analysis: Quora	152
6.1.3	New solutions to old questions	154
6.1.4	The future of network theory	155
6.2	More data, same problems	155
6.2.1	Optimisation of spreading on networks	155
6.2.2	Extreme events in network exploration	156
6.2.3	Inferring networks with partial information	157
	References	159
A	Minimal contagious sets in random regular graphs	173
B	Rare events statistics of random walks on networks: localization and other dynamical phase transitions	233

Introduction

This dissertation will discuss how several interesting and somehow classical problems in network theory can be addressed by using tools and techniques coming from the statistical physics. This choice is due to the fact that my doctoral program has been part of the European project NETADIS - Statistical Physics Approaches to Networks Across Disciplines, funded by the People Programme – Marie Curie Actions – of the European Union’s Seventh Framework Programme. This project has involved nine leading European institutions including among the others the King’s College in London, Sapienza University in Rome, the International Center for Theoretical Physics in Trieste and the Technische Universität in Berlin, and it has aimed at exploring different research themes such as the inference of networks from (potentially incomplete or noisy) data and the optimisation and control of processes taking place on a (possibly evolving during time) network.

Because of the ubiquity of systems that can be modelled as networks, the insights gained by these projects, among which the ones I have faced during my Ph.D. studies, have contributed to achieve a better understanding of phenomena taking place on a variety of fields such as system biology and neuroscience, socio-economical systems and finance, information technology. The statistical physics, whose main task has historically been the analysis of macroscopic system behaviours in terms of the properties and of the interactions of their microscopic components, has been the main conceptual framework unifying such projects. The same framework has been also used, as it has been commonly happening since the last few years, in order to face the so-called *inverse problems* that deal with the inference of the microscopic properties of a system having access to its macroscopic behaviour.

The plan of this work is the following. In chapter 1 the general ideas of the statistical physics of disordered systems are briefly reviewed. One of the main focuses is on the possibility of describing very different systems in a common way by using abstract models that catch only their most basic features, but which are able nevertheless to produce useful qualitative and quantitative prediction on them. In particular, it is

discussed how a precise microscopic description of a system is typically not needed in order to reproduce some apparently complicated processes occurring on it such as collective behaviours, avalanches, abrupt epidemic outbreaks; this point justifies the usage of very simplified models throughout the work. Two general frameworks having been used to set against a background a variety of different situations and phenomena are introduced. The first, originally coming from the computer science community, allows to model systems whose components have to satisfy a certain number of constraints; the second, coming from physics, enables to straightforwardly understand how a system whose elements interact only at a local level can show the emergence of collective, global phenomena. A short review of the fields in which the application of tools and concepts coming from the statistical physics proved to be fruitful is going to be proposed.

Chapter 2 is devoted to network theory and the terminology about it, extensively used in the following chapters, is therein introduced. A particular focus is set on the variety of systems which can be modelled as networks, and on the insights one can get by doing this. The need for defining more and more complex models in order to be able and reproduce more detailed features observed in real networks is discussed. Two different approaches to network theory are introduced, being respectively mostly connected to mathematics and probability theory on the one hand, and with statistical physics and critical phenomena on the other.

The three following chapters extensively describe the projects I have mainly worked on during my Ph.D. All of them refer to very legitimate questions one can ask about networks and processes taking place on them: respectively, what rules govern the spreading of something on a graph, how can we explore it and, finally, how can we reconstruct it by having access only to few and possibly noisy measurements. Chapter 3 refers to the paper [1], reprinted in appendix A, on contagion dynamics on networks, and in particular on how to find and statistically characterise minimal sets of initially active nodes able to infect all the network in a given time.

Chapter 4 discusses how a graph exploration can be performed by making use of a random walk on it. The project [2], reprinted in appendix B and developed in collaboration with professor Reimer Kühn at the King's College in London, aims in particular at statistically characterising the extreme, rare events that can take place while exploring different types of graphs in this way.

Chapter 5 will face the problem of reconstructing a network of interactions among a set of elements, having access just to a partial knowledge of how much correlated are the repeated measurements performed on them. After having defined the problem

in a “static” sense (how to optimally infer the network having been given a certain knowledge about it), the situation is seen under a more innovative and challenging perspective, by searching for some heuristics suggesting which element one should choose to measure in order to improve his knowledge on the system as fast as possible.

Finally, in chapter 6 some conclusions and perspectives on the work are discussed. In particular, I explain therein why I think that the problems faced in this thesis and more specifically a theoretical, abstract modelling of them, will still be topics of crucial interest in the coming years.

Chapter 1

Models and frameworks

1.1 The role of models

This thesis will be mainly devoted to a theoretical and numerical study of models. Because of their large range of validity, the insights obtained will be usable in many different contexts; on the other hand, specific details of real-world systems will in general be neglected. The aim of this kind of modelling is then rather the understanding of a common framework in which different phenomena can be put than the precise description of one of them.

According to a joke famous in the scientific community, the physicists try to solve complex problems by making oversimplifying assumptions (the punchline of one of the versions being “I’ve found a solution, but it works only for perfectly spherical horses in a vacuum”). As usual, a seed of truth is nested in this quip, as the models developed by physicists usually tend to focus on some aspects of the phenomenon under study, neglecting many others; these oversimplified models often lead nevertheless to precious insights about the world we live in, and enable people to make predictions (accurate in some cases, not so much in other) about it. For example, the idea of neglecting shape and size of an object (to be seen, hence, as a single point mass) in order to facilitate the theoretical study of its motion, could have seemed at first as detached from the real-world as the spherical horse of the joke; however, the Newtonian mechanics having been originated in this framework guaranteed a very precise quantitative explanation for a wide range of natural phenomena.

Coming to an example closer to the topic of this thesis, it will be discussed how a simple model of magnetic systems such as the Ising model can lead to non trivial and somehow unexpected behaviours, similar to the ones observed in many real-world situations. The physical modelling of a system can be seen as a search for its simplest possible description able to reproduce some of its interesting features. The modelling

through magnetic systems of social ones let the community understand that there is nothing special in obtaining a macroscopic ordering as an effect of local interactions, as under some conditions an order autonomously emerges without taking into account any specific details. The reasons why the collective behaviours occurring in very different systems can be modelled in sort of a unified framework will be briefly and non-rigorously discussed in the following section 1.3.3, dedicated to the universality of critical phenomena.

In this thesis, I will mainly deal with a sort of “extreme” modelling, similar to considering social systems as magnetic ones as discussed above; the centrality of modelling also for searching a more detailed description of a phenomenon should however not be underestimated. According to the classical paper by Arturo Rosenblueth and Norbert Wiener [3] “no substantial part of the universe is so simple that it can be grasped and controlled without abstraction.” Two different procedures can hence go under the name of modelling. The first one, topic of this dissertation, consists in finding different systems displaying similar qualitative features; the second, in looking for apt abstractions (i.e. systems with similar but simpler structures) in an attempt of precisely describing the problem under study.

As said by Jean-Philippe Bouchaud in [4], “to describe necessarily means to simplify, and even sometimes betray.” If this betrayal allows us to progressively approach the reality, through successive improved approximations, this can be considered a worthy compromise. At the end of the day, George Box’s consideration in [5,6] represents really the central point of this discussion: the right question about a model is not whether it is true (as, in some sense, the answer is always no), as one should rather ask whether it is useful.

1.2 Forward problems, inverse problems

1.2.1 Same phenomena, different questions

When describing a given system, several possible questions can be addressed. From a general point of view they can be split in two big groups. On the one hand, one may be interested in predicting the outcome of an experiment (or, more generally, in predicting the value of a certain observable) knowing the parameters of the model being observed: this corresponds to solve what is called a *direct* or *forward problem*. On the other hand, one can use as a starting point an observation or a measurement performed on the system to infer the parameters of the model having generated it

(*backward*, or *inverse problem*). This second type of question will be the one at the center of this dissertation.

Inverse problems are of great interest because of several different reasons. From an engineering perspective, they are useful for designing systems that will behave in a desired way; the framework under which these problems are addressed goes usually under the name of *control theory*. Secondly, one can think of several systems (for example, biological ones) in which the variables of interest are not easily measurable. In such cases, being able to infer them by using the information content of other, more convenient variables can be of the greatest importance, as it can enable the researchers to get the same insights with much less expensive or time-consuming experiments. This is for instance the case for the *protein folding problem*, as the biological functions of a protein are mainly due to its 3D structure, which is experimentally difficult to measure as it requires the use of techniques such as the x-ray diffraction. Due to the advancements in sequencing techniques, on the other hand, knowing its primary structure (i.e. the linear sequence of its amino acid structural units) has become since the last years a much easier task. A long-standing inverse problem on which statistical physics has given a fundamental contribution is how to infer 3D structure (interesting insight) knowing the linear sequence (easier measurement).

Another possible application is to reconstruction problem, where one has access to a low-dimensional representation of the data and wants to reconstruct the input having generated it. These problems are typically very well solved by our brain: apart from some extreme cases of optical illusion, for example, we are very good in inferring the 3D world around us by using the 2D visual information that our eyes are able to collect. A prototypical example of this kind of situations can be found in signal processing theory, and goes under the name of *source separation problem*. In this case, one has access to a signal which is the combination of several components and aim at correctly identifying these latter, assigning them to the different sources. Also this task is brilliantly performed by our brain, which is typically able to assign each components of the sounds that we hear to its source; in particular, a very clear example is our capability of separating the words pronounced by several people in a crowded room (these situations going usually under the name of *cocktail party problems*).

In the same framework, a topic of interest in medical research is the *discrete tomography*, whose aim is to reconstruct a multi-dimensional binary image from a certain number of lower dimensional projections. This number ought to be as small as possible, in order to let the examination be both as fast, and as non-invasive as possible.

Important advancements, both theoretical and practical, have been recently obtained in this field thanks to methods developed in the statistical physics community, going under the name of *compressed sensing* [7].

Apart from these applications, it should be stressed that inverse problems play a central role also in situations more strictly connected with the physics (and with the statistical physics in particular). For example, a prototypical model for magnetic systems is the Ising model that will be discussed in greater detail in the following of this chapter. In this case, the first problem (and yet already highly non-trivial) people tried to solve was the forward one, namely to determine correlations between spins having access to information such as the external fields acting on each spin, and their couplings; more recently, however, also the harder inverse version of the problem was addressed. Because of the large range of applicability of Ising models, the advances in its understanding originated very important progresses in many fields among which the ones described above.

1.2.2 A more precise definition

In a very abstract way, a physical system can be thought of as an operator G (the equations governing it) transforming a given state of a system (a model m defined by a set of parameters) into a set of observables (the data d). In the simplest case, such a relationship between parameters and data is linear, such that we can write $d = Gm$. The inverse problem is however far from trivial even in such a simple case, as one has typically to face *ill-posed* problems in which the number of equations is smaller than the number of unknowns; if this is the case, straightforwardly solving the inverse problem by calculating $m = G^{-1} d$ is not a viable strategy, as the matrix G is in general non invertible. Even more complicated are the non-linear inverse problems, in which the physical law G connecting the model m with the outcome d is a generic function, such that $d = G(m)$.

A general framework for addressing inverse problems makes use of the fact that for any given model m^* we can solve the forward problem, leading to a new set of data $d^* = G(m^*)$. By iteratively changing the model, we can hence try to reproduce better and better the original data d , aiming at the minimum of a functional measuring the distance between the observed and the recovered data (i.e. the error of the inferred model m^*); this latter goal can be addressed by performing a gradient descent on the parameters.

Another issue frequently encountered in inverse problems is the risk of building up a very complicated model perfectly reproducing the available data, but unable

to generalise to new ones. This is true in particular for the cases in which either we have uncertainties on the physical laws governing the process G or the data we have access to are noisy. This phenomenon goes under the name of *overfitting*, and is usually addressed by adding a penalty for the complexity of the model (i.e. by *regularising* it). In some sense, it can be thought of as an implementation of the Occam's razor principle, according to which we should prefer a more complex model only if we have strong enough arguments favouring it with respect to simpler ones. This general statement can be quantitatively implemented in a probabilistic perspective by imposing a prior distribution on the model parameters. In this way, the functional to be minimised becomes a linear combination of the error on the reconstructed data on the one hand, and of the complexity of the model (calculated for instance as the ℓ_2 norm over the model parameters) on the other:

$$\|d - Gm^*\|_2 + \alpha \|m^*\|_2$$

The value of α , in the former equation, governs the trade-off between the explanation and the generalisation capabilities of the model: $\alpha = 0$ corresponds to the absence of prior on the model, whereas larger and larger α lead to the determination of simpler and simpler models.

Another order of difficulties can be understood thinking again about the protein structure inference problem presented in the previous section. Even if a given set of couplings and fields produces under the model a unique set of site frequencies and 2-points correlations (well-posed forward problem), the converse does not hold, as these latter can be produced by many different couplings and fields (ill-posed inverse problem). This is due to the fact that we are for the time being disregarding the higher-order statistics (3-points correlations etc.) that would differentiate them. However, it is usual very risky to try and fit such information from the data, as the number of parameters increases very fast and the risk of overfitting becomes high. A typical way out of this situation is to choose, among all the models producing the same 1-point and 2-point statistics observed in the data, the most probable one. The methods following this idea are again, in some sense, a quantitative formulation of the Occam's razor principle, and they are usually referred to as *maximum likelihood* if they concern the choice of the parameters of a model and *maximum entropy* if they concern the choice of the model itself [8].

1.3 A physics modelling framework: the Ising model

1.3.1 Motivation

In the next sections, two frameworks playing a crucial role in providing simple enough models for describing a variety of different phenomena are going to be discussed. Both of them were originally designed for dealing with quite specific situations, but turned out to be of a much wider interest because of their generality. The very common characteristics such two frameworks are going to be able and appropriately describe are the possibility of stating a problem as a set of requests the variables of a system have to satisfy, and the interest of global, collective behaviours emerging under certain conditions in a system even if its elementary components interact only in a local, decentralised fashion. The frameworks of interest for dealing with systems showing such features are respectively the one of *constraint satisfaction problems*, originally coming from the computer science community and described in the following section 1.4, and the Ising model with its possible generalisations addressed in the rest of this one.

The Ising model, originally developed in the context of ferromagnetism, has experienced an impressive increase in its range of applicability in the last decades in reason of the ubiquity of real-world systems that, even if their components interact just on a local scale, exhibit collective behaviours. Another important characteristic of this model depends on the physics of the systems it had to describe in the first place: the variables σ_i , originally representing magnetic dipole moments of atomic spins, are binary (i.e. they can assume just two values, typically $\sigma_i \in \{-1, +1\}$ after a rescaling). The model is defined on a network, structure which is fully described by the list of elements or nodes composing it and by the set of connections existing among couples of them. Each node is allowed to interact with its neighbours; in particular, the Ising model on a two-dimensional lattice is one of the simplest models to show a phase transition (collective effect) originated by local, pairwise interactions. Because of this, it has become paradigmatic for understanding how a system can develop an ordered structure starting from a disordered one as soon as some parameter is conveniently tuned.

At first, the condition of having to deal with binary variables may seem very strict. After some thought, one sees that in a large number of systems the feature describing the state of their elements can be chosen out of just two possibilities. Even if a more detailed list will be postponed to the following section, two examples coming from completely different fields will clarify the point. In opinion dynamics, people may

be regarded as binary variables in a lot of situations in which an infinite range of possible feelings about a topic has to be translated into the choice of one out of two possibilities: the choice between Republicans and Democrats the US electors have to make is an example of this, any referendum (where the only degree of freedom is a yes/no answer) is another. Changing field, the Ising model has become a standard in neuroscience. On the one hand, the brain is one of the most complex systems known, on which very little is understood; on the other, at any time the state of the neuron can be characterised, at least at a very coarse level, according to whether it is *spiking* (i.e. emitting an electrical impulse) or not. Many other examples will be discussed in the following; as will be seen, moreover, slight variants of the Ising model will enable to study systems whose variables may assume more than two values.

1.3.2 The Ising model

The Ising model is formally defined on a set of N binary variables (i.e. *spins*) $\sigma_i \in \{+1, -1\}$ localised at the vertices of a d -dimensional regular lattice of total volume L^d . A configuration $\bar{\sigma}$ corresponds to the assignment of a value to each spin, and its energy is given by the following relation:

$$E(\bar{\sigma}) = - \sum_{\langle i,j \rangle} \sigma_i \sigma_j - h \sum_i \sigma_i \quad (1.1)$$

The first term on the right-hand side implements the ferromagnetic interaction, as any couple of neighbouring nodes i, j is energetically pushed towards being aligned (i.e. $\sigma_i = \sigma_j = 1$ or $\sigma_i = \sigma_j = -1$) over being unaligned (i.e. $\sigma_i = 1$ and $\sigma_j = -1$ or vice versa). The second term takes into account the possibility of a global field to act on the system, in such a way to energetically push all the spins to align with it. Both these aspects coming from the description of the magnetic systems are of interest also for the applications. In sociology, for instance, a qualitatively well-known social influence phenomenon called *homogenisation* or *social pressure* [9] describes the social incentives people usually feel in order to behave similarly to their acquaintances. In such a framework, the effect of the external field can be simply thought of as a homogenisation to some impulse coming from the outside and acting on everybody. Also cases in which this tendency is reversed (i.e. people preferring to behave differently from their friends) can be studied in the same framework, being the equivalent of antiferromagnetic systems in physics.

As usual in statistical mechanics models, the probability for the system to be in a given configuration depends on the energy of the latter and on the temperature T of the system:

$$P_T(\bar{\sigma}) \propto e^{-\frac{E(\bar{\sigma})}{T}} \quad (1.2)$$

The role played by the temperature becomes clearer by analysing two extreme cases. For $T \rightarrow \infty$, the energy of a configuration does not have any effect on its probability of being realised by the system: at high temperature, all states are equally probable and the spins are independent both among each other and on the external field. At very low temperature, instead, only the states with the lowest possible energy (i.e. the *ground states* of the system) have a probability larger than zero of being seen. If the external field is zero, two such states exist, consisting in having either all the spins up or all of them down. If $h \neq 0$, this symmetry is broken and just one ground state survives (all spins up if $h > 0$, all spins down otherwise). Because of its effect on the system, the temperature can be abstractly seen as a parameter tuning the tendency to the order or to the disorder: this can be understood by noticing that a change in temperature is completely equivalent to a rescaling of the couplings and of the external fields that determine the energy of the configurations. Such a parameter will be present also in the other applications of statistical physics, and for instance the reference to the “temperature” of a social system will have to be interpreted in this sense.

This model has become paradigmatic as, even if very simple, it exhibits an interplay between two contrasting contributions. On the one hand, the system is pushed towards one of the two ordered states (either all spins up or all spins down when $h = 0$) because of the ferromagnetic interactions; on the other, the thermal fluctuations originate an entropic contribution that favours the disordered, more numerous configurations. Depending on the value of T , their respective importance is tuned, and for instance as the temperature goes to zero the entropic contribution becomes negligible.

The previous difference between an ordered and a disordered phase can be stated in a more quantitative way by looking at the average value of the spins (i.e. at the *average magnetisation*), defined as

$$m = \frac{1}{N} \left\langle \sum_{i=1}^N \sigma_i \right\rangle_T \quad (1.3)$$

where the average is taken over all the configurations, each of them weighted with its probability of being found at the temperature T .

In one dimensional lattices, the model can be quite directly solved, and the average magnetisation turns out to be zero at all $T > 0$: the system is in its disordered

phase as soon as some thermal fluctuations are inserted. For larger dimensions, an explicit solution is much harder to obtain, but the existence of transitions between a low temperature, ordered phase showing a positive average magnetisation and a high temperature, disordered one in which $m = 0$ have been proved for $d \geq 2$. For $d = 2$, in particular, an exact solution has been found, whereas for larger dimensions the main features of the transition have been described by using the *renormalisation group* that will be briefly introduced in section 1.3.3. The transition between the two regimes is smooth for finite N , but gets steeper and steeper as N grows, so to generate a true phase transition in the infinite size limit; in this latter case a finite value of the temperature T_c exists such that the system is ordered for any $T < T_c$ and disordered for any $T > T_c$.

The model proposed has been up to now studied by looking at the probability of finding any given state when the system is at thermodynamic equilibrium at a given temperature T . A slightly different approach is to let the system evolve starting from an initially out of equilibrium state. By doing so, one can explicitly appreciate the way by which the order is established if the temperature at which the system evolves is under the critical one; in order to see this, a dynamics determining the evolution of such a magnetic system at any given temperature has to be defined. A very natural one is the *Metropolis algorithm* [10], according to which the probability of a spin to flip is proportional to $e^{-\frac{\Delta E}{T}}$, ΔE being the difference in energy that it would obtain by flipping. This approach takes into account the thermal fluctuations, as any energetically favourable flipping ($\Delta E < 0$) is accepted, but at $T > 0$ also the unfavourable ones are accepted with some probability: this latter is taken from the Boltzmann distribution at temperature T , explaining the similarity with the previous at-equilibrium formulation. By running this dynamics, some insights are reached on how an ordered configuration may eventually emerge from a disordered one below the critical temperature. Even when the average magnetisation (i.e. the global ordering) is still zero, a *coarsening* process can be seen, by which larger and larger ordered clusters of both signs start to appear and merge together. Finally, this merging process involves the system as a whole, and the globally ordered configuration having $m > 0$ is reached. A slightly more mathematically detailed description of this phenomenon will follow in section 1.3.3.

This case is really interesting to understand the importance of physical modelling. An open problem in social systems was to understand how interactions among acquaintances could eventually (even if not always) lead to the emergence of global trends. By the Ising model, it was shown that to explain that no specific hypothesis

on the system was requested: just the presence of local interactions may indeed be a sufficient condition to generate a long-range ordered configuration. Also the fact that such an ordered state is not always reached can be simply understood by imagining that certain systems are at a “temperature” higher than others (i.e. their fluctuations are more important).

1.3.3 Universality of critical phenomena

As discussed, the Ising model for ferromagnets has become a standard for modelling systems in which a transition from a disordered to an ordered state can be obtained by appropriately tuning its temperature (or some other quantities playing its role in different systems). The similarity existing among all these transitions is however not just qualitative, as it has been shown that such *critical phenomena* can be grouped into very few classes, the elements belonging to each of them being in some sense equivalent at the critical point. The reason why this is possible is that when getting nearer to the phase transition, involving by definition the system as a whole, larger scales start to come into play and it is no more requested to explicitly take into account the details of the different systems, since the microscopic details are sort of averaged out. The properties of such systems at the phase transition, therefore, becomes independent on their dynamical details, and can be expressed as a function of some very general features such as the dimension of the space in which they live and the symmetry broken by the phase transition [11].

In order to be more quantitative, let us consider a magnetic system showing the emergence of an ordered state from a disordered one; this phenomenon can be seen as a progressive enlargement of the typical length ξ over which its microscopic variables are correlated, i.e. the typical size of the *clusters* by recalling the terminology used in the previous section. In an infinite size system such a length has to diverge, by the definition itself of phase transition, as the temperature approaches its critical value T_c . In particular, one observes that for temperatures near enough to the phase transition, a power-law relation of the form $\xi \propto |T - T_c|^{-\nu}$ holds, where ν , called *critical exponent* of the transition, is the parameter governing the system behaviour in the proximity of the critical point.

The interest in such a quantity is that many phase transitions turn out to be described by the very same set of critical exponents; they are said to belong to the same *universality class* and they display identical scaling behaviour as they approach criticality. The fundamental contribution for the understanding of these common aspects among apparently very different systems was given by the beautiful unifying

framework of the *renormalisation group*, proposed by Kenneth G. Wilson [12]. I will not enter here into the details of what has constituted one of the greatest achievements in the modern statistical physics; one can just say that by using this method, very appropriate in order to study systems nearby their phase transitions, it becomes clear how different starting systems eventually fall into a few universality classes, and all the members of one of those show an identical critical behaviour as they share the fundamental characteristics of their dynamics.

This universality is really impressive thinking about the general interest that a small number of paradigmatic models may take on: the Ising model, for instance, turns out to belong to the same universality class as the liquid-gas phase transition, and a description of the critical features of the former is therefore perfectly appropriate also for the latter. Anticipating the main theme of this thesis, it can be already said that the Ising model will because of this be extremely useful for analysing in these terms a series of critical phenomena taking place on networks, i.e. on systems composed by a set of interconnected elements [13].

As a final remark of this section, let us stress that even if two systems belong to the same universality class, this should not at all be considered a proof of a general equivalence between them; what they share is indeed only their large-scale behaviours near enough to a phase transition, whereas their microscopic characteristics clearly have still to be described differently [14].

1.3.4 The generalisations

The prototypical Ising model has become during the years central in the evolution of the statistical mechanics. In reason of its somewhat restrictive hypothesis, many different generalisations have been proposed and studied in order to relax some of them. Such generalisations, being an interesting physical problem on their own, have also been studied in depth because of their larger range of applicability. Here some of them will be reported: the focus will however be mostly on their importance with respect to the possibility of modelling more and more general situations. The technical aspects connected to the search for solutions or even to a better phenomenological understanding would require a large review on their own and will be therefore neglected here.

A strict condition of the Ising model is that all the couplings have the same sign, the interactions between spins being either all ferromagnetic or all antiferromagnetic. One is on the other hand interested (especially for a series of applications where the interactions between the elementary components have to be considered as a degree

of freedom of the system) also in studying cases in which the interactions between particles have different signs, going generically under the name of *disordered systems*. Historically, the field of research of these so-called *spin glasses* emerged in the seventies [15] for understanding a class of magnetic alloys characterised by the coexistence of interactions of different signs. A widely accepted mathematical abstraction for describing such systems is the Edwards-Anderson model [16], which generalises the energy of a configuration in the Ising model as follows:

$$E(\bar{\sigma}) = - \sum_{\langle i,j \rangle} J_{ij} \sigma_i \sigma_j - B \sum_i \sigma_i \quad (1.4)$$

The sum is here again on the couples of neighbouring nodes, and the J are, for instance, extracted randomly from a given probability distribution. This generalisation led to an enormous number of works and this field is still very active today as many features have still to be fully understood. What is important from a modelling point of view, however, is that an extremely broad spectrum of phenomena can be studied in this framework, as a huge number of problems can be stated in terms of such a combination of different sign interactions, from biological systems (such as neural networks, or gene regulation networks) to random lasers. A particularly clear application is to constraint satisfaction problems. Even without entering into the details of these latter, that will be postponed to section 1.4, it can be already said that a constraint asking two variables to take the same value will be straightforwardly modelled with a ferromagnetic interaction $J_{ij} > 0$, whereas the request for two elements to take opposite value will be implemented by imposing an antiferromagnetic interaction between them.

The ground states, trivial to find in the Ising model, are on the contrary very difficult to determine in spin glasses; they will also have, in general, an energy larger than zero. This can be true even for small systems: if the interaction among three spins $\sigma_i, \sigma_j, \sigma_k$ is for example $J_{ij} > 0, J_{ik} > 0, J_{jk} < 0$, it is easy to see that the minimum energy configuration would have $\sigma_i = \sigma_j$ and $\sigma_i = \sigma_k$, but $\sigma_j \neq \sigma_k$ which is impossible. Such a situation is referred to as a *frustration* of the system and will also have a very straightforward analogous in the theory of the constraint satisfaction problems, corresponding to the case in which a set of constraints are in some sense contradictory, as no assignment of the variables is able to satisfy all of them.

The difficulty in finding ground states comes from the *complex energy landscapes* of these systems. A system which starts in an excited configuration and evolves according to the Metropolis algorithm will for instance strive to go to lower energy

configurations, as for reaching them it should change many spins states at the same time: this feature gives raise to the presence of *metastable states* in which the dynamics gets stuck for long times. The prototypical example of such systems are the glasses, whose properties have been massively studied for the last decades [17] but are still far from being comprehensively understood.

Another feature of the Ising model that one can be interested in relaxing is the fact that its variables are binary. For studying similar systems, but characterised by elements which can assume more than two states, the *Potts model* has been introduced. In this latter, each spin takes value in a finite set of q elements $\sigma_i \in \{1, \dots, q\}$ and, in the original version, the energy of a configuration generalises eq. 1.1 as follows:

$$E(\bar{\sigma}) = - \sum_{\langle i,j \rangle} J_{ij} \delta(\sigma_i, \sigma_j) - \sum_i h_{\sigma_i} \quad (1.5)$$

With positive couplings J , as in the Ising model, two spins have still lower energy if they are aligned. The difference is that now they can be unaligned in many different ways; in this case, the energy contribution depends just on the condition $\sigma_i \neq \sigma_j$ and not on the specific values taken by the two spins. Such an energy is directly connected to the so-called *q-colouring* problem: in this example of a CSP, we require any node in a graph to take a value that is different from the one taken by any of its neighbours; the interaction is hence antiferromagnetic as the neighbouring nodes are energetically pushed to take different values, but no “preference” is given on the colours by which this misalignment is obtained. This application explains also why the q values a spin may take are usually referred to also as “colours.”

A further generalisation of this case is the so-called *clock model*, in which the q values each spin may take are thought of as angles uniformly distributed along a circle, with possible values $\theta_n = \frac{2\pi n}{q}$, n being an integer between 1 and q . In this case, the spins may be represented as vectors $\boldsymbol{\sigma}_i \equiv (\cos \sigma_i, \sin \sigma_i)$. With a positive J , again, two spins are energetically favoured to be aligned; their energetic contribution will however more generally depend on the angle between them, and reporting for simplicity the equation in absence of external fields, one has:

$$E(\bar{\boldsymbol{\sigma}}) = - \sum_{\langle i,j \rangle} J_{ij} \boldsymbol{\sigma}_i \cdot \boldsymbol{\sigma}_j = - \sum_{\langle i,j \rangle} J_{ij} \cos(\theta_i - \theta_j) \quad (1.6)$$

If one is interested in modelling systems in which the energy explicitly depends on the two colours the neighbouring nodes take and not just on the conditions of being or not equal or on the “angular difference” between them, a *generalised Potts model*

may be defined. In this latter, all the energy levels can be independently defined and J_{ij} becomes a q by q matrix. This generalisation can be useful in many contexts, as for example in the biological applications that will be discussed in the following.

Also the discreteness of the variables may be relaxed. The most common way by which this can be achieved is by just sending the number of possible angles in a clock model $q \rightarrow \infty$ so that the spins may take continuous values; the *XY model* obtained in this way is formally described by exactly the same equation 1.6, the difference between the two being in the values each spin is allowed to take.

In many cases of interest, the interactions do not involve just two sites at a time. The so called *generalised Ising models*, defined to take into account this possibility, are described in the most general form by the following energy function:

$$E(\bar{\sigma}) = - \sum_{p=1}^{p_{max}} \sum_{i_1 < \dots < i_p} J_{i_1 \dots i_p} \sigma_{i_1} \dots \sigma_{i_p} \quad (1.7)$$

where the index p represents the order of the interaction. In all the cases discussed until now, the maximum order of the interaction had been implicitly fixed to $p_{max} = 2$, considering just one-body interactions (the external fields, so far called h_i) and the two-body interactions, so far called couplings J_{ij} . This formulation enables to study also higher order of interactions. This generalisation, being an interesting feature on its own, is also very useful from the applicative point of view: the constraints in many of the CSP that will be introduced in the following involve indeed three or more variables, and could not be modelled just by using two-body interactions.

1.3.5 Inverse problems in disordered systems

In all the models described before, a lot of efforts have been devoted to the understanding of the direct problems. In this context, they typically consist in calculating thermodynamic quantities such as magnetisations and correlation functions from the knowledge of the interactions acting among the spins (in the simplest cases, fields and couplings). This problem is highly non trivial on its own, as is shown by the fact that a solution for the $d = 3$ Ising model is still lacking. No simple relation can be generally established connecting fields and couplings with the thermodynamic quantities they induce, and an exact estimate is not computationally feasible as it would require a number of calculations exponential in the size of the system (because of the sum over all the possible configurations).

Because of the variety of systems that can be described as Ising models, in the last years much research has been devoted to the study of the far more complicated inverse

Ising problem [18]. In this latter, one has access to the thermodynamic quantities and wants to estimate the interactions having generated them. Formally, the problem can be relatively easily stated: knowing that each set of J and h will generate certain thermodynamic quantities, one has to find out which of them is more likely to have generated the ones we observe; an exact, exhaustive search for it would again require an exponential number of operations.

A brute force approach, giving typically good results but very slow in its convergence, is a *Monte Carlo simulation* that, starting from an initial arbitrary choice of fields and couplings, iteratively calculates the thermodynamic quantities they imply and compares them to the true ones. The parameters are then changed in order to diminish the discrepancy between the two, for instance by using a gradient descent on the parameters, and the procedure is repeated until a set of J and h producing thermodynamic quantities similar enough to the observed ones is produced. This approach is made more difficult by the fact that in this case even the solution of the direct problem, which has to be performed multiple times, is not trivial at all.

A comprehensive review of the methods developed in this field would go beyond the scope of this thesis, and the variety of methods proposed to address this central problem will therefore just be stressed: they go from different orders of mean-field approximations [19] to adaptive clustering approaches [20], from methods making use of the linear response theory [21] to expansions for small values of the correlations [22] to approximations of the probabilistic features of the system making use of the *pseudo-likelihood* [23].

1.4 A computer science modelling framework: the constraint satisfaction problems

1.4.1 Definitions and examples

A modelling framework that turns out to be of particular interest for the type of problems addressed in this thesis is the one of the *constraint satisfaction problems* [24]. Such problems involve a certain number of objects and a set of conditions or limitations they have to satisfy. Each of these conditions involves only few variables, in a number which remains in particular finite and typically small even in the infinite system size limit. One such problem is the already briefly mentioned q -colouring problem, in which each constraint involves exactly two variables (i.e. two neighbouring sites) and asks them not to take the same colour in the assignment proposed.

The most natural solution one can look for in these situations is an assignment of the variables such that all the constraints are satisfied; in the previous example, one would like to look for an assignment such that to none of the couples of neighbouring nodes is assigned the same colour.

For the sake of generality, a more general definition is usually given. The problem is assumed to be defined on a set of N variables $\sigma_1, \dots, \sigma_N$, each of them taking values in a finite set χ ; a global assignment will be indicated as $\bar{\sigma} = (\sigma_1, \dots, \sigma_N) \in \chi^N$. Because of the constraints, each assignment $\bar{\sigma}$ will be characterised by a cost, or an energy (by respectively using a computer science or a physics language) $E(\bar{\sigma})$ corresponding to the number of constraints it violates. The natural request expressed in the previous paragraph can hence be quantitatively stated as the search for a $\bar{\sigma}$ such that $E(\bar{\sigma}) = 0$; however, other legitimate questions can be also asked. A slight generalisation of the previously discussed request is the *optimisation problem*, the output of which is a configuration $\bar{\sigma}$ with minimum total cost. In the cases in which no assignment can satisfy all the constraints at the same time, this algorithm will return an *optimal* configuration $\bar{\sigma}^* = \arg \min_{\bar{\sigma}} [E(\bar{\sigma})]$; in a q-colouring problem, for instance, it will return an assignment of colours such that the number of neighbouring couples sharing the same colour (i.e. the number of *unsatisfied* constraints) is the smallest possible. In the *decision problems*, the output is a boolean variable telling whether an assignment such that $E(\bar{\sigma}) < E_0$ exists for the problem under analysis, without explicitly finding it. If we set $E_0 = 0$, these problems can be seen as a sanity check on the constraints involving the variables: indeed, depending on the output of this problem being positive or not, the CSP is respectively said to be *satisfiable* or *unsatisfiable*. An *evaluation problem*, finally, will determine the cost of an optimal solution; formally, the output of an algorithm implementing it will be $E^* = \min_{\bar{\sigma}} [E(\bar{\sigma})]$ and this result will tell us how many constraints the best possible assignment will be anyway forced to violate, without anyway explicitly proposing such an assignment.

Because of the generality of this framework, a lot of situations can be stated as constraint satisfaction problems. In more applicative fields, for instance, one is often interested in optimising something under some constraints. To make one out of many possible examples, in a *resource allocation problem* one has to decide how to distribute a fixed amount of resources in a system in order to optimise some outcome: in a public transportation system, this may correspond to the goal of maximising the number of people able to use the offered service and minimising their waiting time (i.e. optimising the service in some sense) by distributing in the best possible way

the overall availability of buses. As usual, we will mainly be concerned about more abstract versions of such situations, nevertheless possibly generating insights about such specific applications.

Many CSP are defined on networks and will hence be discussed in chapter 2 devoted to graph theory: the already discussed graph colouring is one such problem, and other examples will be proposed in the following. Another class of CSP having played a central role in the theory of computational complexity is the one of *SAT* problems. Here the variables x_i are boolean, meaning that their value can be either **TRUE** or **FALSE** or, for brevity, respectively 1 or 0; each variable can also be negated, the negation of $x_i = 1$ being $\bar{x}_i = 0$ and vice versa. The variables have to satisfy M constraints, taking the form of *clauses*: these latter are logical **OR** of a certain number of variables (which is in general different for any clause) or of their negations. Because of the definition of the logical **OR**, in order for a clause to be satisfied, at least one of the variables has to appear with its correct value. Let us consider for instance a clause involving three variables x_1, x_3, x_7 as follows:

$$x_1 \quad \mathbf{OR} \quad x_3 \quad \mathbf{OR} \quad \bar{x}_7$$

Such an assignment is satisfied by any choice of the three variables involved apart from $x_1 = 0, x_3 = 0, x_7 = 1$: with this assignment, the clause is said to be violated. As expected, for a fixed number of variables the problem gets harder and harder to satisfy as the number of clauses increases; for a large enough number of such clauses, no possible assignments will be able to satisfy all of them any more, and the problem will enter its *unsatisfiable* phase. Of particular interest are the *k-SAT* problems, where all the clauses the system have to satisfy involve the same number k of variables; in 3-SAT for instance all the constraints will be similar to the one discussed in the previous example.

1.4.2 Infeasibility of the straightforward solutions

A common feature of the CSP is that finding a solution in a non-efficient way is easy enough. If one thinks about the *k-SAT* problem, for instance, for any given assignment of the variables the number of clauses (if any) that it violates can be evaluated quite easily, i.e. in a time which grows just polynomially with the system size. If one systematically tried all the possible assignments, he would be able to answer all the possible questions on the CSP under exam in an exact way. For some of them, one could even stop in advance, for example as soon as an assignment violating no more than a given number of clauses would be found. Unfortunately, this turns

out to be practically infeasible, as the number of possible assignments is exponential in the size of the instance: for a k -SAT problem (and in general for any problem in which the variables may take one out of two values) defined on N variables, there are 2^N possible assignments. This exponential law is far from being a technical detail, as it practically makes this exhaustive approach completely infeasible even for quite small problems.

Smarter algorithms avoiding us to check all the possible solutions can be thought of; a class of such algorithms incrementally builds potential solutions as follows. At the beginning of the procedure all the variables are unassigned and at each step one of them is picked up at random, and it is temporarily assigned first to 0 and then to 1. If just one of the assignments does not lead to any clause violation, the variable is assigned to that value; if none of the assignments produces contradictions, the variable is assigned at random to 0 or 1. In some cases, however, both the assignments may lead to contradictions; without getting into the details, I will simply say that refined versions of this original algorithm have been designed so to face such situations. One of the main elements of these refinements is the *backtracking*, mechanism by which it is possible to get out of a contradiction as the one described above. In such a situation, the backtracking consists in de-assigning some of the variables previously fixed, and to restart. This iterative procedure is repeated until one is able to get to a complete assignment of the variables without incurring into any contradiction.

The procedures discussed can be refined over and over, so to accelerate the search for a solution, or to enlarge the probability of finding one. For reasons that will be clarified in the following section, however, all these refinements are in general not sufficient to be sure to find a solution of a problem big enough in a reasonable amount of time.

1.4.3 The worst case scenario

Given the immense number of CSP one can think of, a question of great importance is how to design a way of comparing them. More precisely, one would like to build a hierarchy of difficulties (i.e. of the number of operations requested) in problems, so to be able and assign each one of them to the appropriate level. Being this framework largely independent on the details, one could imagine to use as proxies of the “difficulty” of a problem its running time, or the number of operations requested for solving it. The hardness of a problem may however differ from an instance to another, and the notion of difficulty of a CSP problem (and not of one of its specific realisations) has hence still to be defined.

People in the theoretical computer science community, who addressed this issue first, were mostly concerned about being able to obtain performance guarantees for each problem. In order to have them, they adopted a *worst-case scenario*, according to which the difficulty of each problem was to be stated analysing its most difficult instance [25]. In order to let this scenario (known as *computational complexity theory*) be even more universal, they organised the hierarchy in quite huge classes, looking at the scaling law of the running time for increasing sizes N of the problem; also in this case, a precise definition of what is meant by “size” of the problem is not needed, as the following classification will work for any “reasonable”¹ choice.

The **P** class is the one of the easiest problems, that are solvable in polynomial time. More precisely, an algorithm exists such that the number of operations requested to solve a problem of size N is bounded by N^k for some k . For another big class of problems called **NP** the former statement cannot be proven, as all the known algorithms are not polynomially bounded (scaling typically as the exponential of N); but if a solution is given, its correctness can be verified in polynomial time. Demonstrating whether the two former classes are distinct is one of the major open problems in theoretical computer science; however, it is widely believed that this is the case (**P** ≠ **NP** conjecture) and that, according to this, being able to verify a solution in polynomial time should not be a guarantee of being able to design a polynomial algorithm for finding such a solution.

The **NP-complete** class is a subclass of **NP** whose members are characterised by the fact that any other problem in **NP** can be polynomially reduced to each of them. Without entering into details going beyond the scope of this thesis, I will just say that, because of this condition, if a polynomial time solution for one of the problems in **NP-complete** were found, all the problems in **NP** would be demonstrated to be solvable in polynomial time. The k -SAT problem above introduced is particularly interesting in this regards: its $k = 2$ version indeed is a simple **P** problem, whereas the $k \geq 3$ cases were the first to be rigorously assigned to the **NP-complete** class [26]. A huge number of other problems have joined them in the following years; among them, the *travelling salesman problem* dealing with the exploration of a network that will be extensively discussed in chapter 4.

The hierarchy of difficulties goes on with a class of problems which are at least as difficult as the ones in **NP**, but are not themselves in **NP** as a straightforward way of verifying a solution has not been found yet: this class goes under the name

1. By reasonable here it is meant that the size should not be defined in such a way to be a quantity exponential in the number of elements that contributes to finding a solution

of **NP-hard**. Of particular interest from a physicist point of view is the **#P** class. In all the previous cases, indeed, the focus was on a single solution, separating the problems according to the hardness of finding one, or of verifying the correctness of one. The **#P** class, on the other hand, typically gathers problems focused on the entire space of the solutions. Even in cases for which finding a solution can be an easy task (i.e. solvable in polynomial time), asking how many possible assignments satisfy a given property can be much harder, as the former is just a basic component of the procedure needed for solving the latter.

1.4.4 The typical case scenario

From a modelling perspective, the classification of CSP described in the previous section has two very strong positive points. Because of the worst case scenario adopted, the guarantees it ensures on a given problem are completely independent on the realisation: for a problem demonstrated to be in **P**, a polynomial time solution can always be found, no matter what. The second point is that, as the hierarchy is organised in large classes, this classification is widely independent of any specificity of the problem under exam (implementation, details of the CSP, precise question asked on it). A strict parallel among very different problems is established as a plus, and for example an (improbable) polynomial-time solution to be discovered for 3-SAT would immediately originate a polynomial-time solution for the graph colouring problem.

In the 80's, however, people started to question this classification. In fact, it turned out that many instances of hard problems (for instance, belonging to the **NP** or to the **NP-complete** classes) were straightforwardly solvable in polynomial time by using quite standard algorithms. Paradoxically, one can think of a CSP built in such a way to have just one specific instance only solvable in exponential time, whereas all the others, much easier, are addressable in polynomial time. Theoretical computer scientist would tell you that such a problem is in any case “hard”; but is it?

The former question is of course a bit provocative. The answer is clearly “yes”, as soon as we accept the notion of worst case scenario hardness. If we think that the answer should have been “no”, on the contrary, an alternative definition of hardness has to be thought of. This alternative should address a series of questions which the worst case scenario leaves (purposely) aside; which is the fraction of truly hard instances in a problem classified as “hard”, how difficult is a typical instance, and how the difficulty increases as we add more and more constraints are examples of such unanswered questions [25]. The approach used to address these points consists in generating a large number of instances so to look at their statistics (at their average

hardness, for instance); for the sake of generality, these instances have to be generated at random. In the case of k-SAT, for example, one fixes at first the number of variables and the number of clauses they have to satisfy; then, an explicit instance is obtained by selecting at random which variables have to appear in each clause, and their sign (i.e. whether or not they are negated in the clause). Because of the stochasticity of the generation procedure, these problems go under the name of *random constraint satisfaction problems* or *rCSP*.

This new situation under exam having been clarified, it is now possible to give a formal definition of the word “typical”, previously used in quite a generic way. In the context of rCSP, a given characteristic is said to be typical if the probability for it to occur goes to 1 as the system size grows ($N \rightarrow \infty$). Interestingly enough, these problems are not at all new in the scientific community. An entire branch of the physics, called statistical physics, has indeed been dealing with systems characterised by the average behaviour of their components since the end of the XIX century. The methods developed in two very far apart fields of the science, hence, turned out to be applicable to the same set of problems. Another even more striking similarity can be drawn among these problems and the physics of spin glasses introduced in section 1.3.4, and will be the main focus of the following chapter 1.5.

After this shift in the types of questions people were interested in, it became clear that a set of concepts and tools previously developed in the community of the statistical physics of disordered systems could have been applied to rCSP [27]. In order to clarify this statistical approach to rCSP, I will refer again to the previously discussed k-SAT example. By requiring a set of variables to satisfy more and more clauses the problem will clearly become harder and harder, as any clause added to the system at fixed N diminishes (or at most leaves unchanged) the number of satisfying assignments. If we define $\alpha = \frac{M}{N}$ as the ratio between the number of constraints and the number of variables, hence, we expect that the larger α is, the harder the problem gets; in particular, for α large enough the instances will often be unsatisfiable. In order to quantitatively characterise this common sense statement, a thermodynamic limit has to be performed so to be able and tell whether a rCSP with a given α is typically satisfiable or not: the correct thermodynamic limit will correspond in this case to sending M and N to ∞ while keeping their ratio $\alpha = \frac{M}{N}$ (i.e. the “difficulty” of the problem) fixed. Formally, calling $P(\alpha, N)$ the probability for a random k-SAT with N variables and clause density α to be satisfiable, one can study the behaviour of such a quantity as a function of α for several N (see figure 1.1).




Figure 1.1: Probability for a random 3-SAT to be satisfiable as a function of the clause density α for several system size n . As one can see, these curves are monotonically decreasing for any n and their steepness increases for systems getting larger and larger. In the infinite size limit one would expect a sharp transition between a “satisfiable with high probability” and an “unsatisfiable with high probability” phase in correspondence of a critical value α_s . The figure is taken from [25]

As expected, for very low (respectively high) α almost all the instances are satisfiable (respectively unsatisfiable) even for quite small system sizes. The most interesting part of the figure is however in between these two easily understandable regimes. Indeed, the part of the curve connecting the “solvable with high probability” with the “unsolvable with high probability” regions becomes steeper and steeper as N increases. For an infinite system size, it is expected that a sharp transition between a SAT and an UNSAT phase will emerge in correspondence of a threshold value: an α_s will hence exist such that:

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

By using the physics language, this is a phase transition analogous to the ones occurring in the magnetic systems that have been discussed in section 1.3. In particular, the satisfiability threshold α_s plays in this case the same role of the critical temperature T_c that was defined in such a way for an infinite size system to be ordered for any $T < T_c$ and disordered for any $T > T_c$.

The discussed SAT-UNSAT transition is the most natural phase transition one can think of for this kind of problems. The application of the statistical physics methods, however, let people discover other very interesting phase transitions inside the SAT phase, where solutions of the problem typically exist. These latter involve the structure of the space of the solutions, and their discover gave both a theoretical

and a practical contribution to the understanding of these problems (for example, for justifying the abrupt change in the algorithmic performances in correspondence of some values of $\alpha < \alpha_s$). The discussion of such transitions will be delayed to chapter 3, devoted to the problem of spreading dynamics optimisation on graphs, where explicit examples of such phase transitions will be shown.

1.5 Spin glasses, rCSP and large deviations

1.5.1 Spin glasses and CSP: two languages, same problems

In the previous sections I presented two frameworks apparently very far apart from each other. On the one hand, the constraint satisfaction problems, originally studied in the computer science community; on the other, the Ising model for magnetic systems and its generalisation to spin glasses and disordered systems. People realised at a certain point that the two were intrinsically similar, and that all the machinery developed in one community could have been fruitfully applied by the other (for instance, the study of the worst case scenario to spin glasses and the search for a ground state in CSP).

In some cases the connection between statistical mechanics and CSP emerges in a very clear way, as a mapping between problems is possible: the parallel between solving a colouring problem and finding the ground state of an antiferromagnetic Potts model has been for instance discussed in section 1.3.4. Even more generally, CSP can be described in terms of spin glasses: in both cases, indeed, a configuration of the binary variables is chosen depending on an assigned cost function. The interactions among the variables are in both cases randomly chosen and constitute a *quenched disorder* of the system; the procedure of averaging over the quenched disorder represented by the explicit realisation of couplings and fields in a spin glass model is analogous to performing ensemble averages over the instances while considering rCSP. As it has been discussed, the energy level of the ground state of a spin glass is not at all trivial to determine, and it is in general larger than zero even for small systems because of the frustration emerging among the couplings. From a CSP point of view, this situation is equivalent to affirm that the system is unsatisfiable as not all of its constraints can be matched at the same time; also in this case, one may incur in such a situation even in very small systems, such as a three variable a, b, c CSP in which we require $a = b$, $a = c$ and $b \neq c$. These three conditions clearly cannot be satisfied together and the system is said to be geometrically frustrated because of the presence of such contradictory conditions.

Having clarified the deep similarities between these two topics, a simple example will conclude this discussion by showing how a CSP problem can be precisely translated into the language of spin glasses, so to be able and explicitly tackle it with the methods developed in this latter field. The *maximum cut problem* is defined on a graph in which to each link is assigned a given weight. If the nodes of such a graph are partitioned in two groups V_+ and V_- , a certain number of links (called the *cut* of the graph) will connect a node in V_+ with one in V_- . The solution of this problem is a partition characterised by the sum of the weights of such links being the largest possible. Now let us imagine the weight of any link between a couple (i, j) of sites to be the opposite of a magnetic coupling J_{ij} , and to represent the fact of a node belonging to V_+ or V_- by respectively assigning it a variable $\sigma_i = +1$ or $\sigma_i = -1$. The previous problem turns hence out to be exactly the same as searching the ground state configuration $\bar{\sigma}$ minimising the spin glass energy function $E(\bar{\sigma}) = -\sum_{\langle i, j \rangle} J_{ij} \sigma_i \sigma_j$.

1.5.2 Looking for atypical events

In a sense, both of the approaches to CSP presented before are missing or neglecting a very important part of the problem: on the one hand the worst-case scenario does not tell us anything about what we could reasonably expect from a given, randomly chosen instance of the CSP under exam. On the other, focusing exclusively on typical, most probable cases flows away all information about the extreme situations one could encounter in such problems.

In particular, a more precise evaluation of how improbable are the non typical events is of the greatest interest in finance [28], where an extreme, unlikely event could be the one triggering a global crisis. The large deviation principles that will be presented in this section aim at exploring and characterising the fluctuations of a random system around its most probable state. This can be thought of as a generalisation of the typical-case scenario, as also the non-typical phenomena are therein addressed in a probabilistic sense.

This framework is based on the so-called *large deviation principle*. In an intuitive, non-rigorous form this principle says that the probability of an event P_n indexed by an integer n scales for large n as $P_n \approx e^{-nI}$ where I is some positive constant. More precisely, the probability of a random variable A_n to take value in a set B satisfies a large deviation principle if a *rate function* can be found such that

$$\lim_{n \rightarrow \infty} \left[-\frac{1}{n} \log P(A_n \in B) \right] = I_B \quad (1.8)$$

If this latter holds and $I_B \neq \{0, \infty\}$, $P(A_n \in B)$ has a leading exponential decaying behaviour in n (what was meant by the first, non rigorous form of the principle). If $I_B = 0$, the piece of information we are getting is that the set B considered contains the typical event; on the other hand, if $I_B = \infty$ the theorem still holds, but it only implies that the leading decaying behaviour is super-exponential in n and, in some sense, the large deviation *ansatz* turns out not to be fully justified in this case.

A convenient form of this principle is stated in terms of the probability density of the variable to assume a given value a as

$$P(A_n = a) \approx e^{-nI(a)} \quad (1.9)$$

A rigorous derivation of these statements from a mathematical point of view, which goes beyond the scope of this work, can be obtained in the context of the theory of probability [29].

A typical problem can in this context be split in two parts. At first, a variable of interest for which a large deviation principle holds has to be found; secondly, a precise form of the rate function quantitatively describing its exponentially depressed fluctuations around the typical behaviour has to be derived.

The most natural way to address this problem is to directly calculate the probability distribution of the random variable of interest, so to be able and explicitly study its asymptotic behaviour as $n \rightarrow \infty$. In many cases, this approach turns out to be very hard, or even impossible: this is the case for continuous random variables, or for variables which are not independently, identically distributed in which one cannot use the standard approximation techniques, such as the Stirling's one.

A different route can however be taken to derive the rate function in many cases of interest, the difference with the previously described method being the same differentiating the use of the *microcanonical* and of the *canonical ensembles* in statistical physics, i.e. by respectively fixing either its energy or its temperature.

Using the same notations as before, let us define the *scaled cumulant generating function* of A_n as

$$\lambda(k) = \lim_{n \rightarrow \infty} \left[\frac{1}{n} \log \langle e^{nkA_n} \rangle \right] \quad (1.10)$$

where k is a real number and the average on the right-hand side is defined as

$$\langle e^{nkA_n} \rangle = \int_{\mathbb{R}} e^{nka} P(A_n \in da)$$

Under the condition of $\lambda(k)$ existing and being differentiable for any value of k , the Gärtner-Ellis theorem ensures that a large deviation principle in the form 1.9

exists for the variable A_n ; this latter can be easily evaluated by what is called a *Legendre-Fenchel transform* as $I(a) = \sup_{k \in \mathbb{R}} \{ka - \lambda(k)\}$.

The former requests on $\lambda(k)$ are far away from trivial. In chapter 4 a case in which they do not hold will be discussed; the consequences of these hypotheses not being verified will be very interesting in themselves from a physical point of view, as they will shed light on how the system behaves under some specific conditions.

1.5.3 The large deviations in a simple case

The previously discussed ideas can be clarified by looking at a standard, very basic problem in probability theory. Let us consider a sum of n independent, identically distributed variables extracted from a Gaussian probability distribution with average μ and variance σ^2 ; their *sample average* is defined by:

$$S_n = \frac{1}{n} \sum_{i=1}^n X_i \quad (1.11)$$

The role played in the previous section by the system size is played in this case by the number of extracted variables, and quantities of particular interest are the asymptotic values holding in the $n \rightarrow \infty$ limit.

This problem can be looked at according to different perspectives. At a first level, one could be interested in something reminiscent of the typical case scenario above described, i.e. which value the quantity S_n is more likely to take as $n \rightarrow \infty$. The answer, even being in this case quite simple, is a result of major importance in probability theory and goes under the name of *law of large numbers*: it says that S_n converges *in probability* to its mean, such that formally:

$$\lim_{n \rightarrow \infty} P\left(S_n \in [\mu - \delta, \mu + \delta]\right) = 1 \quad (1.12)$$

If one is, on the other hand, interested in determining also the probability with which S_n will take an arbitrary value s , in general different from its mean, one can study its complete probability distribution. As expected, the average of Gaussian i.i.d. variables is still a Gaussian variable, and more precisely one obtains:

$$P(S_n = s) = \sqrt{\frac{n}{2\pi\sigma^2}} e^{-\frac{n(s-\mu)^2}{2\sigma^2}} \quad (1.13)$$

Having obtained this result, one can quite straightforwardly follow the route described in the previous section so to determine the large deviation function of S_n . By

neglecting the sub-dominant for $n \rightarrow \infty$ term of order \sqrt{n} , one gets a large deviation formula for the quantity of interest of the same form as eq. 1.9:

$$P(S_n = s) \approx e^{-nJ(s)} \quad \text{with } J(s) = \frac{(s - \mu)^2}{2\sigma^2} \quad (1.14)$$

where $J(s)$ is therefore the rate function. The possibility of deriving the rate function governing the exponential decay was in this case particularly simple, as it made use of the gaussianity of the variables X_i entering into play. However, a classical and crucial result of probability theory going under the name of *Cramér's theorem* states that the same is true for partial sums of i.i.d. variables extracted from some arbitrary probability distribution (which has nevertheless to satisfy some not-so-restrictive hypothesis) and that in all such cases the probabilities of realising “large deviations” away from the mean decay exponentially with a rate given by the so-called *Cramér function* [29].

1.6 Statistical mechanics for non-physics problems

1.6.1 A common framework, many fields of interest

Such a connection between different fields has during the years become more and more common. In particular, some of the most exotic phenomena in many applicative fields such as biology, finance, social systems may be quite easily understood in a physical framework, letting the statistical physics become a common language for very diverse problems. In the theory of complexity developed over the last forty years, indeed, physics has developed several models allowing one to understand how small perturbations can lead to large effects in systems in which the natural notion of linearity (small actions leads to small effects) does not hold. In some optimisation problems, for instance, the algorithmic performances were known to fall abruptly for some values of the parameters, and the experience of physicists in glassy systems for which the optimal solution may either not exist or not being reachable in a finite amount of time turned out to be useful. Many fields started to show a non trivial and almost unexpected connection with fundamental issues in statistical physics like the presence of phase transitions and the existence of glassy phases. The former, in particular, were good representatives of how a macroscopic, collective behaviour may emerge in a large system following simple microscopic laws [27]. These considerations becoming more and more evident led to the birth and to the rapid expansion of a new, interdisciplinary field of research at the crossroad between statistical physics,

information theory and combinatorial optimisation [18]. The most advanced tools and concepts coming from the statistical physics community, in particular, allowed to solve and to gain fruitful insights on very large constraint satisfaction problems like random satisfiability, colouring, or error correction codes. Among them, techniques such as message passing, survey propagation, cavity method will be presented in reason of their application to the specific case of the spreading dynamics in chapter 3.

1.6.2 Social systems

A field in which the usage of prototypical Ising-like models has been particularly successful is the one of the social sciences. A lot of hardly understood collective behaviours turned out to be quite simple phase transitions if thought in terms of magnetic systems. The emergence of global ordered states in the way crowds [30] or flocks of birds [31] move showed that the same principles hold both for humans and for animals, as soon as the number of elements is large enough to justify a statistical treatment. According to the Newton maxim reported in [28], indeed, “modelling the madness of people is more difficult than the motion of planets”, but as soon as large populations are involved statistical regularities begin to emerge. In other interesting studies, the understanding of vehicular and pedestrian traffic has been addressed [32] by making very few “realistic” assumptions on the characteristics of the basic elements.

As already partly discussed, many interesting results have been obtained by making use of statistical physics methods as regards the opinion dynamics [30]; in particular, how a global agreement on a certain topic may emerge even if every person has the possibility of discussing only with a very limited number of friends, or acquaintances. In this case, the finite number of states a component may assume appears a natural choice, and the model obtained in this way is maybe less abstract than in other cases.

Another situation in which the use of discrete values is very natural is in the models of segregation. By using variables taking two values as in the Ising model, Thomas Schelling showed in his foundational paper [33] how even a small preference of any person to have a minimal percentage of neighbours similar to him may result in a globally ordered structure (i.e. a complete segregation of the positive and negative “spins”). This model, at first carried out at a sociological level, received afterwards much more rigorous treatments from a statistical physics point of view [34]. The important insights obtained by these abstract models are confirmed by the Nobel prize in Economics attributed to T. Schelling for “having enhanced our understanding

of conflict and cooperation.” In particular, the absence of a simple correspondence between individual incentives and collective results was shown for the first time.

Finally, also the linguistics has been approached with these methods [30] in order to try and quantitatively understand how an agreement about the words to be used to describe a given concept can be eventually reached, as well as how languages evolve during time.

1.6.3 Economics and finance

The possibility of studying economical systems according to statistical physics should not be surprising, as they can be seen in a sense as an extension of the previously discussed social systems. Even if the consequences are typically much more severe, crisis are for instance in the simplest modelling framework not far away from the emergence of an agreement in an opinion dynamics context. The necessity to start using in economics studies physical models able to explain how wild consequences could be obtained by apparently innocuous perturbations on the system was claimed by J.P. Bouchaud in his provocative paper [28]. The modelling of heterogeneous agents interacting in a financial market in the framework of the statistical mechanics of disordered systems has been proposed and deeply studied [35].

Also the financial system has been addressed by the statistical physics community. The usefulness of a quantitative approach is in this case out of discussion as, according to the L. Bachelier maxim reported in [4], “the market, without knowing it, obeys to a law which overwhelms it: the law of probability.” The possibility of using statistical physics methods has been relatively straightforward in this case as the finance has been a data-rich system for many decades now. A quantitative approach to the financial instruments gave rise to the field known as *econophysics* [36], and a more precise study of the statistical properties of financial time series has been proposed for example in [28].

1.6.4 Neural networks

The biological application of statistical mechanics models are several and variegated. This depends probably on the fact that in many cases the precise functioning of the systems underlying biological processes is not well understood and simple models reproducing some features can be therefore of interest. As regards the Ising model, the most natural and already briefly discussed application is to the neural networks; a statistical physics approach has been used also in order to address other topics of

interest in biological systems, which are for simplicity discussed in the following section 1.6.5. A neuron, that at a given time is mainly characterised by the fact of either being firing or not, can be easily modelled as a an indicator variable (a spin) taking value $+1$ or 0 . In his classical paper [37] J. J. Hopfield tried to explain the ability of a large network of neurons to perform computational tasks just as a spontaneous, collective consequence of local, pairwise interactions; that is, by the same mechanisms that in magnetic systems are able to originate collective, global phenomena such as stable magnetic orientation as an outcome of local interactions among large number of elementary components. The study of the brain according to such models, that never stopped, forked into two main fields of research. On the one hand, people tried to understand better and better the functioning of the brain (*natural neural networks*), for instance by making use of data coming from experiments. On the other, the basic, abstract models of neural networks have become a standard in the machine learning community, for automatically performing tasks such as image recognition and natural language processing (i.e. *artificial neural networks*). Many results have been obtained in this field, and more and more complex frameworks have been analysed (the most recent and fashionable being probably the one of *deep learning* [38], on which relatively little is known for now).

1.6.5 Other biological systems

The amount of sequencing data available has been constantly and dramatically increasing over the last twenty years; in reason of that, the interest of being able to statistically modelling such data has more and more increased. In protein folding [39] there is a discrepancy between the data biologically and structurally important (the three dimensional folded structures) and the ones we can more easily have access to (the linear sequences of amino acids known as primary structures). A statistical connection between the two has been found, as by looking at the substitution patterns of the amino acids in all the members of a given protein family, one should in principle be able to infer the couples of sites most stable during the evolution, stability which is in turn a proxy for the structural importance of such couples (i.e. for finding the contact points in the folded structure). This problem is intrinsically inverse, as the quantities of interest are the couplings between sites, whereas the observable quantities are the frequency counts generated by them. For solving it, one would like to build up the most likely statistical model for the complete sequence able to reproduce the correct (measured) marginals.

A lot of study has also been devoted to the evolutionary dynamics from a statistical physics perspective [40]: among the others, topics such as the statistics of adaptive processes, the population genetics and the evolution of fitness landscapes have been touched. These fields have observed a real revolution during the last century, going from a qualitative to a quantitative approach. Even if already in 1924 J.B.S. Haldane began his classical paper [41] saying that “a satisfactory theory of natural selection must be quantitative”, this approach has become more and more fundamental with the avalanche of data that has started to be available for statistical analysis since the last couple of decades.

1.6.6 Signal processing

The problem of optimally reconstruct high dimensional signals using a limited number of measurements is of great interest in many different fields such as image processing, astronomy or biology. More generally, a central topic in signal processing theory is how to minimise the sampling of a *sparse* signal (i.e. with many components equal to zero) in such a way to be nevertheless able, later on, to reconstruct it either exactly or approximately. *Compressed sensing* triggered in this context a major evolution, as it enabled to record only the information needed for the reconstruction [42]. If with more traditional techniques one was forced to perform an acquisition not as sparse as the true signal, by using statistical physics methods on large systems it was possible to take a number of measurements approaching the theoretical limit [43]. Both the algorithms implemented for the reconstruction of the signal and for the design of the measurement matrix rely on statistical physics, the first being message-passing inspired, the second coming from the theory of crystal nucleation.

More generally, statistical physics methods have given a major contribution in analysing and understanding the performances of a class of algorithms used in signal processing [24] in order to decide how to perform the encoding, these latter methods going under the name of *error correcting codes* [44].

1.6.7 Computer science

The connection between statistical physics and problems of interest for the computer science community [24], and in particular to CSP [45], has already been discussed in some details: I will just remind that, in this context, statistical mechanics methods turned out to be of extreme utility in understanding the performance of well-known algorithms, in determining the theoretical thresholds for these problems

and in designing new procedures able to approach such limits. These discussion will be further detailed in chapter 3, where a statistical mechanics approach to the solution of the spreading dynamics on networks (which is generally speaking a constraint satisfaction problem) and the results obtained by it will be described.

Chapter 2

The network theory

2.1 Motivations and definitions

2.1.1 Ubiquity of networks

Networked systems consisting in a set of interconnected elements are really ubiquitous. They go from the intricate food webs of central importance in the functioning of ecosystems to the neurons which compose the brain, from the cell that performs its tasks thanks to complicated interactions among proteins, genes and other molecules to social system with individuals either collaborating or being in competition with the people they are in contact with. Depending on the context, the network can be embedded in the real space, with edges representing spatial proximity, or in a more generally defined space in which two nodes do not need to be really near to each other in order to be in contact (two far apart airports being in this sense in contact if at least one flight between them exists).

An ambitious task is to better understand the laws governing all these systems. In order to do this, one firstly has to imagine a way of consistently describing all of them as special cases of a more abstract model. This can be achieved by inserting into the model only the features common to all the networked systems and neglecting all the ones specific of some realisations. What one obtains by doing this is a *graph* or a *network* consisting in a set of elements called *nodes*. The possible interconnection between a couple of them is represented by a *link* or an *edge* joining them.

This very abstract modelling framework is a huge field on its own, and a comprehensive discussion of it is far beyond the scope of this thesis. However, very complete books and papers addressed the topic both from a more rigorous and mathematical perspective [46] and from a more applicative, physical one [14]. The aim of this chapter will be mainly to introduce some specific terminology that will be used in the




Figure 2.1: On the left hand side, a realistic representation of the position of the landmasses composing the town of Königsberg and of the bridges connecting them. On the right hand side: abstract graph representation of the situation under exam. Figure taken from [25].

following when discussing explicit cases of problems on networks and to give some intuition about the variety of interesting features one can look at when analysing a networked structure.

2.1.2 The birth of graph theory

The birth of the graph theory is very illuminating from the perspective of the importance of abstract models to improve our knowledge on apparently complicated systems.

In the XVIII century, the city of Königsberg was composed by four landmasses connected by seven bridges as can be seen in the left panel of figure 2.1. A curious question intriguing the inhabitants was whether it was possible to find a path through the city so to use each bridge exactly once. Leonard Euler understood that in order to solve this problem, the system was exactly describable just in terms of the landmasses and of the information telling which of them were connected by a bridge (see right panel of figure 2.1). All the other details such as the explicit position of the bridges, their length, the dimension and the location of the landmasses etcetera were completely irrelevant. The solution was indeed found in terms of such synthetic description, as Euler showed [47] how the path with the desired characteristics was impossible in Königsberg and, more generally, in any city in which an odd number of bridges converged on more than two nodes.

It is worthy to stress that in this case, and in general in this kind of combinatorial problems, the description obtained via the graph is exact. In other cases this is no

more possible, but the more or less refined approximations one obtains by abstractly studying the geometry of the problem are nevertheless useful for better understanding the situation under exam. In some situations (for instance in biological systems) the simple collection of nodes and links is a very rude description; even if in such cases it is typically hard to produce quantitatively accurate estimates of the outcomes of the system, the use of the *network metaphor* [48] allows one to have an idea about the kind of processes and phenomena that will possibly take place on it. More than the possibility of modelling a system through a network, however, one should ask himself whether this modelling really enriches our knowledge on the system and, in particular, which questions it will help us answering. Applying a highly simplifying network metaphor to a real system is something to do with caution and having clear in mind what features of this simplification one is going to look at.

2.1.3 The main features of a graph

An elegant way of abstractly defining networked structures is by using graph theory which specifies relationships among a collection of items. Communication networks (for instance set of routes, interconnections among computers or infrastructures for transporting energy), social networks, information networks (as web pages connected by hyperlinks) are all examples of systems describable in this way. One of the first characterisation of such systems one can think of is based on the *reachability* of its vertices, meaning by this the possibility of going from a node to another by travelling over the network links; if any node is reachable from any other, the graph is said to be *connected*. More quantitatively, one can look at the minimum number of steps needed to connect any couple of nodes on the graph. The mean of these values goes under the name of *average shortest path length*, and it is nothing but the *distance* that typically separates two vertices; the maximum among all the distances is instead usually referred to as the *diameter* of the network.

One of the most important local features of a graph is probably the *degree* k of each of its nodes, defined as the number of links converging on it. The average of all these values gives again a very useful global characterisation of the graph, as it is equivalent, after a rescaling by the system size N , to the density of connections among the network elements. Two extreme cases in this respect are of particular interest. Firstly, one can be interested in graphs featuring very few connections, and in particular with the lowest possible number of links. If one restricts to connected graphs, the smallest possible number of edges is $N - 1$ and corresponds to the graph being a *tree*. In this case, no multiple paths exist between any couple of nodes and

hence no closed loops can be drawn by following the edges. As will be discussed in the following, on trees the exact solution of many problems can be found by making use of iterative methods. Approximated versions of these methods can be nevertheless used also on graphs that, even without being strictly speaking trees, show very rare loops so that any node sees, around it, a structure that is very similar to a tree; they are called *locally tree-like* networks and constitute a class much broader than the one comprehending only trees. They can indeed be obtained by randomly choosing a small enough number of connections among the nodes, whereas the probability of getting a tree with such a procedure is negligible.

At the other extreme, one has the graphs where all the nodes are connected with all the others; they are said to be *complete* or *fully connected* or *cliques*, and for such a network of N nodes, the number of edges is $\binom{N}{2}$. Also this case is important from a modelling perspective. A well-known approximation of networked systems is indeed the so-called *mean field* where one makes the approximation of everybody being connected with everybody else [49]; in other words, the underlying graph is supposed to be a complete network. This approximation turns out to be extremely valuable in many situations; for instance, the *Curie-Weiss model* of ferromagnetism, relying on such assumption, enabled to exactly study the properties of the thermodynamic functions in the neighbourhood of the critical temperature. Many real-world networks are nevertheless very far away from this assumption; they are said to be *sparse* and only few of the possible links actually exist. Several quantitative definitions can be proposed for precisely saying what one means by “few”. However, since an N -nodes connected graph has always a number of edges between $N - 1$ (if it is a tree) and $\binom{N}{2}$ (if it is fully connected), every network having a number of edges E scaling as $E \sim N^\alpha$ with $\alpha < 2$ can be considered sparse; according to other works, instead, a graph is considered sparse if it has a number of edges linear in the system size, i.e. $E \sim N$.

The global description taking into account only the average degree of the nodes is clearly very partial. In particular, it is not able to distinguish two very different situations, in which the same mean connectivity is obtained in one case by nodes having all the same degree, and in the other by a network where highly and lowly connected nodes coexist. In order to see this difference, one has to look at more detailed quantities, and in particular at the complete degree distribution or at the least at higher moments concerning it.

2.1.4 Enriching the network metaphor

The description of a networked system as a set of interconnected elements has to be seen as a basic framework that can be enriched when dealing with systems with specific characteristics. Even without entering into the huge variety of such details that can be inserted in a network definition, a couple of features are worthy to be recalled. Depending on the interaction one wants to model being symmetrical or not, one can model the system via respectively an undirected or a directed graph. In this latter case, the relationship existing between a couple of nodes is specified by an arrow going from one of the two to the other; if the graph is undirected, instead, they are simply joined by an edge. Put another way, in an undirected graph the connections between nodes can be listed according to unordered pairs of vertices, whereas in directed graphs they are specified by ordered pairs such that (a, b) and (b, a) are not equivalent. The network of scientific collaborations is for instance symmetrical (i.e. undirected) whereas the links telling whom eats whom in a food web necessarily need to be considered as directed. In biological systems, the gene regulation networks are typically asymmetrical as the fact of gene a regulating gene b does not imply the other way round. When inferring the protein structure, on the other hand, the couplings are by definition symmetrical as the link is considered as a proxy of the spatial proximity of two sites, this latter being a symmetrical quantity. If the graph is directed, the metrics by which a network is characterised have to be appropriately modified, for instance by defining an *in-degree* and a *out-degree* for any node telling the number of links coming into and going out of a vertex, instead of a single quantity representing the number of neighbours it has.

Another hypothesis one can be interested in relaxing is the equivalence among the role played by any of the edges. Thinking about the brain, it has been experimentally shown how two neurons can be connected more or less strongly; the same holds for social systems, in which some of the people we are linked to are great friends of ours, whereas others are just acquaintances. This feature is modelled by assigning a weight to each of the links in the graph that can even be, in general, a negative quantity; this is for instance the case if one wants to take care of inhibitory relationships occurring in neural networks.

2.2 A path towards more realistic networks

2.2.1 The comparison between observed and reproduced features

From a modelling perspective, very interesting is the situation in which the real data are found to be dramatically different from the ones one obtains by generating a graph according to a simple model. For instance, in absence of strong a priori elements against this hypothesis, one can suppose that the degrees of two neighbouring nodes are not correlated. If this was the case, however, in a heterogeneous network the probability of a very highly connected node (also called a *hub* in the following) being connected to another one would be almost negligible; there are indeed largely more low-degree than high-degree nodes to which it could link. Looking for instance at the Internet graph, nevertheless, one discover what has been called a *rich-club phenomenon*, as the hubs are very well connected among each other [50]. This contradiction is enlightening as it shows that the former independence hypothesis has to be relaxed: in order for the description to be reliable, one has to allow a high-degree node to preferentially attach to another one of the same kind. If this is true, the network is said to be *assortative*. Also *disassortative* graphs where a high-degree node preferentially links to low-degree ones and vice versa are found in applications. This is the case for instance of the network composed by the personal computers and by the machines working as servers; the disassortativity can be seen as a *soft* version of a *bipartite graph* where the nodes are divided into two groups and all the links connect a node of one type to one of the other.

Another such case of discordance of a simple model with the real data appears when analysing triplets of nodes. If nodes A and B are both connected to a node C , a legitimate question to ask is the probability for A and B of being themselves linked to each other. In a random graph where all links exist with probability p , the answer is trivially p again. Even before proposing a rigorous description, one can see how this feature is unrealistic in many real cases; in social networks, for instance, two people being both friends to a third one are much more likely to be friend to each other than two randomly selected individuals.

A way for quantitatively analysing the structure of local neighbourhoods is by studying the *clustering coefficient* of a node i . Supposing it to have k_i neighbours, the maximum possible number of connections among them is $\frac{k_i(k_i-1)}{2}$; if this is the case, the subgraph constituted by i and its neighbourhood is a clique. The clustering coefficient is defined as the ratio between the number of actually existing links e_i

connecting one of the first neighbours of i to another and the maximum possible value such a number can take:

$$C(i) \equiv \frac{e_i}{\frac{k_i(k_i-1)}{2}}$$

As discussed, this quantity appears very naturally when analysing social networks and it was indeed firstly proposed in the sociology community under the name of *transitivity* [51]. The fact that the clustering coefficient of a real network is typically much higher than the one of a random graph in which all the links exist with the same probability is a sign of some non-trivial mechanism regulating the construction of the graph, and in particular of the impossibility of modelling the local neighbourhood of a node by a completely stochastic model. More complicated models able to reproduce such a feature will have to be thought of if one wants a reliable description of the system.

The principle discussed so far can be generalised. Any characteristic significantly deviating from the one obtained by a simplified model is a mark of some rule governing the system that has been neglected. The same remark exposed on the clustering coefficient could be applied to any geometric feature being over-represented (or even under-represented) with respect to what one would expect on random graphs. Such local structures are referred to as *network motifs* [52], and they have been particularly used to gain more insights on the mechanisms governing the formation of biological networks. When analysing the transcriptional regulating network of *E. Coli*, for instance, several patterns were found to be atypically represented in many distinct components of the network. Their study enabled to better understand which biological process each of them was associated to, and in reason of which it was preserved and selected during the evolutionary process [53].

In some sense, the iterative enrichment of the model needed to reproduce some observed characteristics has to be understood as a journey towards a better comprehension of the most important structural features in a network; at each step, one looks at the features its *null model* is not able to reproduce, and tries to modify it so to obtain them. The most important this deviation from randomness is, the most one should be convinced of the role played by it in the system. However, one has to be careful about the possible overfitting of the data, resulting in the determination of an extremely complex model able to perfectly reproduce the data already gathered, but not to generalise to yet unseen ones.

2.2.2 Reproducing the observed average distance and clustering coefficient

A comprehensive review of the networks topologies proposed during the last decades both as simple mathematical models and in order to reproduce different systems is far beyond the scope of this introductory chapter. Following the discussion of the previous section, it is interesting to see how more and more complicated graphs were progressively introduced so to be able and understand some features found in the real data that were impossible to reproduce with the simpler models. In relation to chapter 1, the question usually asked is what is the minimal set of conditions one has to impose on the model so to be able and reproduce some features observed in real systems.

When thinking about networks, the two most simple cases one can imagine are characterised by either a completely regular or a completely random structure; both these extremes have been extensively studied, but it turned out that none of them was able to reproduce at the same time the clustering coefficient and the average path length typical of real-world networks, and more complex networks had hence to be proposed.

A completely determined structure one can think of is a regular lattice. Given the loopiness typically characterising these graphs, their clustering coefficient is usually high. The mean length of the shortest path connecting two nodes is also large, as no long-range link exists and for connecting two far apart regions of the network many small steps are needed. The degree distribution is in this case a trivial δ function, as by definition all the nodes have the same number of neighbours.

The completely stochastic model, on the other hand, corresponds to a Erdős–Rényi graph in which any couple of nodes is connected with the same probability p [54]. This model has been widely used, and it is in some sense the most appropriate choice for modelling systems for which any knowledge of the principles guiding the creation of edges between the nodes is lacking. All the edges being independently drawn, the probability of a node to have degree k corresponds in this case to the probability of “extracting” k of the links (each of these events having probability p) and at the same time “not extracting” the other ones (events with probability $1 - p$). By considering all the possible permutations among nodes one easily sees the degree distribution for such a graph of size N to take the following form:

$$P(k) = \binom{N-1}{k} p^k (1-p)^{N-1-k} \quad (2.1)$$




Figure 2.2: In order to obtain “small-world” networks, one can start from a regular ring lattice and redirect the links with a certain probability. If the rewiring happens with probability 1, a Erdős–Rényi graph is obtained. The figure is taken from [14].

A case that is often considered is the large system size limit $N \rightarrow \infty$ where the average connectivity of the nodes is kept fixed to a given $\langle k \rangle$ by scaling the probability for a link to be drawn according to $p = \frac{\langle k \rangle}{N}$. In this case, the binomial degree distribution of eq. 2.1 can be approximated by a Poisson law as follows:

$$P(k) = e^{-\langle k \rangle} \frac{\langle k \rangle^k}{k!} \quad (2.2)$$

Any two couples of nodes are typically linked by quite a short path, of order $\log(N)$ for large systems. The clustering coefficient is in this case almost a meaningless quantity, as the probability for two nodes of being connected is always p independently on them being both linked to a third one or not.

Any graph for which the rule determining the presence of links deviates from the exposed very simple examples can be generically referred to as a *complex* network. Almost all the real-world networks fall into this class, as for instance they show the coexistence of nodes with very high and with very low degree that neither a lattice nor an Erdős–Rényi graph features. Studying the features of real networks, it appears that they show a mixture of the previously described characteristics. Networks of social interactions, for instance, were known to exhibit high clustering coefficient (because of the transitivity phenomenon [51]) and short paths connecting couples of nodes, as shown already in the original S. Milgram’s experiment [55].

A slightly more complex topology able to explain the discussed features was proposed in their seminal paper [56] by Duncan Watts and Steve Strogatz. They noticed that the existence of short paths seemed to depend on the presence of long-range links connecting far apart regions of the network, whereas the high clustering coefficient was originated by the locally highly ordered structure. The *small-world network* they




Figure 2.3: For an intermediate range of the rewiring probability the small-world networks feature the coexistence of two of the characteristics observed in real networks: the high clustering coefficient and the typically short path connecting any couple of nodes. The figure is taken from [14].

proposed aimed at joining these two features, by starting from a regular lattice and rewiring some of the links at random, so to create the looked-for long-range connections. As can easily be seen in figure 2.2, the two previously discussed ensembles are just special cases of this, as for a rewiring probability $p = 0$ one stays in the lattice case, whereas for $p = 1$ all the links are randomly reassigned and one obtains again a Erdős–Rényi graph.

In figure 2.3 the curves of the clustering coefficient and of the mean distance between nodes are reported as a function of the rewiring probability p . As one can see, the networks obtained when p is in an intermediate range have both the features one was looking for (short paths and high clustering). A minimal model able to reproduce these features observed in real networks was hence found.

2.2.3 Reproducing the observed degree distribution

As one can imagine, this is not the end of the story. Just two features were considered, and more and more complex models have to be considered as soon as one is interested in reproducing more detailed characteristics. At a certain point, a trade-off between the simplicity and the generality of the model on the one hand, and its ability to accurately simulating complex networks has to be evaluated according to the principles exposed in chapter 1. Another step in this progressive refinement is anyway worthy of being recalled.




Figure 2.4: A variety of real-world networks features a scale-free degree distribution. Starting from the top left corner and proceeding clockwise, the distribution refers to the worldwide airport network, to the actors' collaboration network, to a mapping of systems connected to the Internet and to data referring to the WWW structure. The figure is taken from [14].

The degree distribution of a small-world network cannot be written in a form as simple as the one of the Erdős–Rényi graph of eq. 2.1. Its shape is however quite similar to this latter as it features a pronounced peak in correspondence of the degree K of the original lattice and it exponentially decays for degrees different enough from this value, i.e. for large values of $|K - k|$ [57]. In real networks, instead, one usually observes the very different *power-law* degree distribution of the form:

$$P(k) \propto k^{-\nu} \quad (2.3)$$

The fact of such a distribution holding for a variety of different real system can be seen in figure 2.4 where very clear linear relations appear by plotting the distribution on a double logarithmic scale. This distribution is *heavy-tailed*, where this latter feature can be intuitively connected to the fact that it generates a coexistence of few very high degree nodes, called *hubs*, with a multitude of low degree ones. This heterogeneity leads to quite dramatic outcomes: the average degree is for instance typically no more a representative quantity of the system, as most identified power laws in nature have exponents $2 < \nu < 3$ such that the mean is well-defined but the variance diverges [58].

Again, a minimal model able to produce this feature was needed. In [59] A. L. Barabási and R. Albert showed how a more than reasonable mechanism governing the wiring of new nodes joining the network was able to generate graphs with the desired heavy-tailed degree distribution. This mechanism, called *preferential attachment*, can be intuitively stated by saying that each new-comer links to one of the already existing nodes with a probability that is proportional to the degree of this latter. This “rich-get-richer” dynamics, originally referred to also as *Matthew effect* or *accumulated advantage* [60] had been conjectured in sociology since many decades as a process governing the relative easiness by which already rich or famous people could further increase their richness or popularity. This very reasonable rule for studying a growing network was then quantitatively related to a specific feature one wanted the model to be able to reproduce, as it was shown that the degree distribution of networks growing according to such a rule spontaneously evolved towards a stationary power-law distribution of the form $P(k) \sim k^{-3}$, particular case of the more general equation 2.3.

2.3 Studying graphs as a branch of probability theory

2.3.1 A mathematical description

The possibility of directly looking at the graphical representation of a graph is useful as it enables to intuitively have an idea about several of its features; how dense are its connections, whether non-trivial structures can be found and whether the degree distribution is homogeneous or highly connected *hubs* and low-degree nodes coexist. For more quantitative analysis, however, it turns out that a synthetic description of nodes and links via a unique matrix A is fundamental. This latter is called *adjacency matrix* and it is in its simplest form boolean, its generic A_{ij} element being either 1 if a link between i and j exists or 0 otherwise. In particular, if the network one is considering is undirected A will be symmetrical as by definition $A_{ij} = A_{ji}$ for any couple (i, j) . The adjacency matrix representation can be naturally extended so to deal with weighted graphs: in this case, it is no more boolean as the value of the element A_{ij} represents the strength of the link connecting i and j . A non-existing link can be formally interpreted in this framework as a link with 0 weight. The usefulness of this mathematical representation will be discussed as regards two different classical problems.

The most basic way of exploring the structure of a network is by a *random walk* that starts from a node, follows one of its out-going links chosen at random and then iterates the procedure for a given number of steps. In this simplest case, a random step leaving from any given node will drive the walker to one of its first neighbours with equal probability and the sum of these latter probabilities has to be set equal to 1, as we impose the random walker to actually go somewhere at each time step. This process, that will be also at the core of the work discussed in chapter 4 and reprinted in appendix B, can be naturally described in terms of the adjacency matrix above introduced by defining a *transition matrix* $W_{ij} = A_{ij}/k_j$ whose W_{ij} element tells the probability of a transition from the node j to the node i .

Always in the context of the network exploration, one may be interested in knowing whether it is possible to reach a given node starting from another one in p steps. In some sense, this can be seen as a generalisation of the information contained in the adjacency matrix, as the links can be defined as paths of length 1 connecting couples of nodes. For $p > 1$ one has also to consider that if the graph is not a tree multiple paths can be found. Also this information can be obtained very easily in terms of the adjacency matrix: after some simple algebraic passages, indeed, one understands how the value of the $(i, j)^{th}$ element of the p^{th} power of A corresponds to the number of different paths of length p connecting i and j . In particular, if $(A^p)_{ij} = 0$ no such path exists.

2.3.2 Combinatorial problems on graphs

The first massive usage of graph theory was as a substrate for combinatorial problems naturally living on networks. These latter, very numerous, can be split into different classes; some examples will be given so to clarify their heterogeneity. A problem of interest in many real applications is to determine the maximum possible flow between a source and a destination that an infrastructure is able to sustain. This situation can be modelled in graph theory by choosing two nodes playing the role of the source and of the sink, and by associating each edge to a *capacity* which can be seen as the maximum amount of flow that can pass through it. The *maximal flow* problem aims at finding the maximum flow that can be distributed among allowed paths going from the source to the sink. Any acceptable solution to the problem has to satisfy two conditions. First of all, the capacity constraints have to be satisfied so that the flow passing through any of the links is not larger than its capacity. Secondly, the flows have to be conserved, meaning by this that the sum of the flows entering into any node must equal the sum of the flows exiting it (the sink and the source

nodes being clearly excepted from this constraint). A big class of problems deals with routing: among them, the determination of a Eulerian cycle as the one discussed in section 2.1.2 or the *travelling salesman problem*, where one looks for the shortest path connecting a set of vertices. A number of variations can be proposed around the problem of finding so-called *spanning trees* with given characteristics, the simplest one aiming at finding the minimal set of edges on a graph that touches all of its vertices. One can also look at subgraphs with given features, for instance by searching the *largest clique* in the graph. Other combinatorial problems aim at finding *matchings* on a graph, these latter being defined as set of non-adjacent edges (i.e. no two edges belonging to a matching share a common vertex). In particular, optimisation problems are typically interested in finding *maximum matchings* characterised by the fact that the addition of any single edge to it is going not to be a matching any more, as any possible extra edge is adjacent to some already chosen one.

One of the first questions about these problems was to understand their computational complexity. This analysis showed that a large number of them are difficult in the sense of section 1.4.3: in [61] the authors report a list of problems demonstrated to belong to the NP-complete class, and more than one hundred are listed in the “network theory” or “network design” categories. Because of the principle of problem reduction, the importance of graph theory in this context is even greater, as very different combinatorial problems can be shown to be equivalent to problems defined on networks. One of the simplest such examples is the mapping existing between the request of finding a perfect matching on a bipartite graph and the apparently very far away determination of the maximal flow on a network [25]. This latter is also equivalent to the problem of finding a *minimum cut*, corresponding to the smallest set of edges that one has to remove in order to prevent going from a source node to a target by moving on the links of the network.

Even if these problems are very abstract, they are strictly connected to real-world situations such as the optimal design of large scale IT systems, the minimisation of congestion in road traffic or the increase in efficiency of electrical grids [48]. Following the parallel introduced in chapter 1, a statistical physics approach aiming at understanding the features of the typical cases has become more and more studied in addition to the original worst-case perspective proposed in the theoretical computer science. Because of the practical interest of such problems, moreover, fast even if suboptimal solutions can be of interest in several contexts. Just to discuss one of the clearest examples, it is evident how matching problems are crucial for many applications [62], as they can be seen for instance as the search of the optimal allocation of

a set of available resources. This problem is worst-case intractable in its stochastic variant, in which the decisions have to be taken knowing just probabilistically the features of the graph one is working on. The statistical physics approach allowed to obtain numerically effective approximations, these latter being also generalisable to other problems of optimisation under uncertainty [63].

2.4 Phase transitions, critical phenomena and statistical physics approach

2.4.1 Dynamics on networks, dynamics of networks

When discussing the connection between dynamics and network theory, one has first of all to clarify that two different meanings can be given to the former in this context. A vast field of interest is in the dynamical processes occurring on networks. Of peculiar interest is the understanding of how the features of such processes change by changing the structure of the networks on which they take place. The results of quite classical dynamical processes models on *complex networks* have recently received a great deal of attention. For instance, one can remember the works on the *voter model* [64] and on the *agreement dynamics* [65] on small-world networks, or the epidemic spreading [66] and the random walks [67] on scale-free networks. The works reprinted in appendix A and B, that will be extensively discussed respectively in chapters 3 and 4, are also of this kind.

Another type of problems regards the dynamical evolution of the network itself [68]. Many different situations can be studied in this context. One can study a fixed number of nodes, and look at how characteristics such as connectivity, average path length, clustering coefficient and so on change as more and more links are added to the system. On the other hand, one may study the features of networks growing in size, where each new-comer links to some already existing node according to the preferential attachment rule discussed in 2.2.2 or following some other mechanism. Also a better understanding of whether a network can break down in disconnected components as a certain number of links (or of nodes) are removed from the system is a crucial point, especially because it can enable to design efficient strategies apt to protect critical infrastructures both from random failures and from targeted attacks. This topic was comprehensively tackled in [69], a comparative study carried out to verify the resilience of networks with different topologies to either random or selective

removal of nodes; the degree distribution being homogeneous or heterogeneous was found to be an element leading to very different results in this context.

These two aspects of the dynamics (on networks and of networks) should be considered as coexisting in real cases, as a dynamical process takes place on a network which typically, far from having been fixed once and forever, is itself changing over time. This latter evolution, however, takes usually place on longer time scales; the study of a process on a *quenched* network can hence be seen as aiming to understand the behaviour of the system on a shorter time scale on which the network structure can be thought of as fixed.

2.4.2 Statistical physics and network theory

As discussed, the graph theory was born as an abstract mathematical topic, mostly connected to combinatorial problems. However, it turned out more recently that principles and tools of the statistical physics could be useful in several aspects. The reason of this interchange can be understood by thinking about a Erdős–Rényi network where each link exists with probability p ; an instance of this type is just a randomly extracted element of the *ensemble* comprehending all such graphs. A parallel can be easily drawn with the *canonical ensemble* in thermodynamics representing the possible states of a system in thermal equilibrium with a heat bath at a fixed temperature; the role of the fluctuating energy is played in this case by the number of edges in the graph, which is not fixed. By slightly changing the rule according to which the random graph is generated, one is also able to obtain a network resembling the state of a mechanical system having an exactly specified total energy (i.e. extracted from the *microcanonical ensemble*): when dealing with a N -nodes network, it is indeed enough to randomly choose which couples of nodes to join by drawing a fixed number M of edges [70]. If one is interested in understanding the features of the networks with given characteristics rather than a specific realisation, the notion of statistical ensemble taken from the statistical physics is hence really powerful. Secondly, the ability discussed in chapter 1 of the statistical physics to focus on the characterisation of emergent phenomena in terms of the dynamical evolution of the basic components of the system is clearly of great interest in many of the applications naturally living on the networks that have been described so far. Thirdly, many real-world networked systems are very large, and the asymptotic properties that the statistical physics usually tries to understand (i.e. the ones holding in the *thermodynamic limit*) are usually the crucial ones.




Figure 2.5: Second order phase transition in a Erdős–Rényi graph signalling the birth of a giant component including a finite fraction of the nodes for a large enough connectivity.

Even more generally, networks are composed by a huge number of interacting microscopic units, and are hence clear examples of systems in which a statistical mechanics approach may lead, as discussed in chapter 1, to interesting insights about the global, collective behaviour emerging from such local interactions. As extensively discussed in [48], many problems on networks can be considered as disordered systems, as the interactions between variables (i.e. the links) are quenched and do not evolve in time, whereas the dynamics involves the system variables. Because of this parallel, all the considerations introduced in chapter 1 may be more or less directly translated so to be of interest also in the context of network theory.

2.4.3 Critical phenomena on networks

According to the discussion of the previous section 2.4.1 on the different meanings the word “dynamics” can take when talking about networks, two classes of critical phenomena may take place on graphs. On the one hand, one can find structural phase transitions in the network architecture when this latter is modified, for instance by adding or by removing nodes or edges; on the other, sharp transitions in cooperative models living on networks can be observed [13]. These two classes will be discussed separately.

One of the first structural changes on evolving networks to be described was found on Erdős–Rényi graphs when edges are added little by little while keeping fixed the number of nodes. The two limits of this process are trivial, as for very few links the graph is fragmented in many components, each of them connecting very few nodes; on the other hand, for a large enough number of edges the probability for the graph of being disconnected will become smaller and smaller. What is interesting,

however, is the transition between these two regimes. To see this, one has to define the probability for a node to belong to the largest component $\gamma_1(c)$, equivalent to the size of the network largest component divided by the system size N , as a function of the average connectivity of the graph c . In the large size limit, indeed, one finds a critical value c_p such that $\gamma_1(c^*)$ is equal to zero for any $c^* < c_p$ but becomes larger than zero for $c^* > c_p$ (see figure 2.5). This phenomenon is usually referred to as the *emergence of a giant cluster* in an Erdős–Rényi graph with connectivity density large enough. By using the language of thermodynamics, this is a *second order phase transition*, as the quantity under exam is continuous at the critical point but its derivative is not.

A similar problem is the *site percolation*, in which one activates a certain fraction p of the nodes in a regular lattice and wants to determine whether a path touching only activated nodes and connecting the top with the bottom of the lattice exists [71]. Also in this case, in the large system limit a critical behaviour is found in correspondence of a p_c such that the probability of such a path to exist is 0 for $p < p_c$ and 1 otherwise. It is worthy to stress how this problem is interesting also from a physical point of view, as it can be considered an abstract model for determining the probability for a liquid particle to entirely cross a porous material.

Another example of criticality emerging during the structural modification of a network is the emergence of scale-free architectures in graphs growing according to the preferential attachment principle. As discussed in 2.2.2, if this growing rule is applied the degree distribution converges to a power-law that can be written down as $f(x) = ax^{-k}$ with $k = 3$. The systems for which such a functional relation holds (for a generic value of the exponent k) are critical in the sense exposed in chapter 1, as by rescaling the variable one obtains the following:

$$f(x) = ax^{-k} \Rightarrow f(cx) = c^{-k}f(x) \Rightarrow f(cx) \propto f(x) \quad (2.4)$$

meaning that all power laws sharing a particular scaling exponent are equivalent up to a scaling factor. In terms of the critical phenomena language introduced in section 1.3.3, they are said to belong to the same universality class; this is also the reason for these distributions to be called “scale-free”.

The second class of critical behaviours is found while analysing processes taking place on a fixed network. The transition between different regimes in cooperative models such as the Schelling model for the segregation, introduced in the previous chapter 1, are examples of such phenomena. Other cases that that will be extensively described in the next chapter 3 are the phenomena occurring near epidemic thresholds in spreading dynamics.

In the study of abstract models of *avalanches*, critical behaviours were found in the distribution both of the size and of the duration of such events [72]. Also models enabling to understand how synchronisation phenomena can happen in large population of interacting elements have received a great deal of attention in the last decades and they have been investigated in the context of physical, biological, chemical, and social systems [73]. A successful approach to the problem of synchronisation consists in modelling each member of the population as a phase oscillator. By introducing strong enough coupling among these latter (for instance by using the Kuramoto model [74]) one is able to reproduce the looked-for phase transition leading to their synchronisation [75].

It is also worthy to stress that recent advancements in this field [13] showed how much critical phenomena taking place on complex networks [76] may differ from the ones, having been studied for many decades, happening on regular lattices [77] or on random graphs.

2.5 Examples of other problems

2.5.1 The detection of communities

Understanding whether a graph is composed by a combination of fairly independent compartments and being able to separate them is a task of primary relevance in network theory. Many real networks are indeed organised according to such a *community structure*, which when unveiled is likely to reveal some useful information about the internal structure of the elements, as vertices belonging to the same community probably share some common properties or play a similar role in the overall network.

From a more quantitative point of view, the community structure is defined as a natural partition of network nodes into subgroups such that the nodes inside a group are much more densely connected among them than they are with nodes belonging to other groups [78]. One of the difficult parts of this problem is that typically the number of communities (also known as *clusters* or *modules*) is not known in advance, and one has to choose it according to some rational which has a component of arbitrariness. Many community detection methods have been proposed [79] and estimating their accuracy is not a trivial task. Ideally, one would like to have access to some examples of networks with a well defined and known in advance community structure so to test whether the proposed algorithms are able to recover it. Since this is usually not possible, one is forced to use *generative models* defined in such a way to build modular networks [80]. One of the most popular generative model

is the *stochastic block model* [81], a generalisation of Erdős–Rényi graphs where all the links are randomly assigned, but the probability for two nodes to be connected is higher if they belong to the same community (or “block”) and lower otherwise. This model is of particular interest as critical phenomena were detected on it; as expected, indeed, if the probabilities of having in-group and out-group links become very similar the detection of the community structure becomes harder and harder. This common sense consideration has been recently shown to correspond to a phase transition between a region of the parameters in which one can hope to recover the modules, and another one in which any algorithm is doomed to fail in detecting the underlying block model [82].

The topicality of this field is proven by the amount of work still proposed on it; recent improvements [83] on well-established spectral methods were for instance able to recover clusters almost up to the theoretical detectability transition rigorously determined. The detection of communities is moreover of central importance in many different applications. In marketing, for instance, one is typically interested in grouping users with similar characteristics so to be able and propose them tailored special offers. In the so-called recommender systems [62], that will be extensively discussed in chapter 5, one has an extra degree of freedom; the graph underlying such systems is indeed bipartite, as one has a set of users on one side and a list of product on the other. Based on explicit or implicit ratings given by somebody, one could be interested both in extracting a set of users similar to each other, and a set of product similar to the ones a given user already has shown to appreciate.

The modular structure can even help one to understand how scientific communities evolve during in time. By looking at the communities in the network of the citations in scientific papers [84], one can for instance remark the emergence of interdisciplinary new fields: a higher and higher number of crossed references between the computer science and the biology community can be used as a threshold for signalling the birth of the new field known as computational biology. In a similar way, the most advanced research engines exploit the network topology of the web so to infer which pages are connected to a same topic [85].

Finally, one has to recall that in some cases the clusters correspond to a real structure underlying the system. In neuroscience, running algorithms related to the community detection on data coming from neural activity recordings can shed light on clusters of neurons spiking at the same time, this latter information being useful in order to identify the functional modules of the brain. More generally, because of the structural and functional systems having features of complex networks that

constitutes the brain, advancements in quantitative analysis of complex networks rapidly translate to studies of brain organisation [86].

2.5.2 Networked versions of other problems

Network theory can be also considered as a substrate for attacking problems of other branches of the science. A very interesting generalisation of *game theory*, for instance, is to a case in which each player, instead of being able to interact with anybody else, is allowed to play only with his neighbouring nodes on the network structure on which the problem is defined [87]. Possible applications of these problems are quite easy to find. Especially when thinking about real-world networks, indeed, it is clear how nowadays a comprehensive, centralised information is typically lacking and the unique possible control on the system has to emerge as a result of the agents acting independently and even, in most of the time, selfishly (i.e. having in mind their own objective rather than the well functioning of the system). The interaction of these multiple independent decision-makers on a networked structure defines a new interesting field at the boundary of game theory and graph theory, whose results may be of interest in a variety of different fields. Just to make a couple of examples, they could be used to optimally design social or economical incentives in a framework more realistic than the ones obtained by applying the standard economical theoretical principles, or to study the mechanisms regulating the resource allocation (for instance in communication networks) among selfish agents and how to avoid congestions in realistic traffic conditions.

The previous is just one out of many possible such examples. Apart from being considered a specific field per se, hence, network theory can be also seen as a way to quantitatively characterise the geometry of the systems on which other processes (also studied on their own) occur. Processes such as the spreading dynamics or the synchronisation among elements have indeed been studied on several types of graphs leading sometimes to very different results depending on the features of the underlying network. This generated a profitable interchange between the results obtained in network theory and the ones reached in the general study of the dynamical processes one is interested in.

Chapter 3

Contagion dynamics on graphs

3.1 General framework

3.1.1 The problem and its range of applicability

A point of great interest in the context of the network theory is to improve our knowledge of the mechanisms regulating the spreading dynamics on graphs. A better understanding of how some characteristics of the nodes can propagate through a network via its edges would be indeed applicable to several domains: because of the level of abstraction of such a problem, by appropriately defining the graph and the features spreading on it one could get fruitful insights in many different processes.

The first field of interest one can think of is the study of epidemic processes, where a disease spreads on a network whose links represent the spatial proximity of individuals. As this is in some sense really the most natural application, all these dynamical processes are sometimes referred to *epidemic processes*, even when they refer to very different systems.

A very rich variety of applications for such considerations can be found in the social sciences. In this case, again, the nodes represent the individuals. The links, on the other hand, may be more generally interpreted than in the previous case. Two people do not need to be physically near in order for a rumour to circulate: any connection by phone or mail is enough. It is also enough for them to be connected to the same “information hub” such as a newspaper, an Internet page, a television channel. The very same dynamics may concern how people start to use a new technology, or more generally a new product; because of the mechanism known as *social contagion*, indeed, people tend to conform to what their friends and acquaintances do [88]. In this sense, a better understanding of these processes would be of great interest for developing

better marketing strategies. More recently, also the problem of how emotions can propagate through the links of a social network has been addressed [89].

Several completely different fields can be touched just by changing the definition of the spreading quantity. As well as being in contact with a certain number of ill people let us become more likely to get ourselves the disease, an infrastructural element in touch with several damaged ones is likely not to work properly. These processes are connected to *avalanche* phenomena, as the failure of a very small number of elements may in the end lead to major, global system malfunctioning. What makes the most simple models unreliable in these cases is that they do not take into account the interdependence of the components. Even if the failure of many elements at the same time is extremely unlikely, indeed, this is no more necessarily true if one considers the non-linearity of the system (each failure making many elements much more fragile). In general, these effects are difficult to forecast; an accurate estimate of their potential impact is nevertheless fundamental as the two following examples will demonstrate.

The infrastructural networks, such as the transportation or the energetic ones, are typical cases in which such avalanche failures may occur; a good understanding of them is fundamental in order to know in advance what needs to be done to prevent large scale damages. The so-called North-East blackout of 2003, a cascade failure of the electric grid affecting more than 50 million people in the US and in Canada, is a spectacular example of such damages. In some sense, one has to forecast also second-order avalanches: the systemic breakdown of the electric network temporarily disrupted in that case also the communication network, blocked the industries, paralysed the transportations; in some regions even the restaurants were affected, because of possible contaminations of their water supply.

Another system in which these phenomena may (and do) occur is the economical one, in which they potentially lead to cascade defaults [90]. Because of the complex network of reciprocal exposures, the bankruptcy of any institution puts lots of others in a much riskier position. The analysis of such *systemic risk* is further complicated by the presence of highly non trivial connections among the elements. The *credit default swaps* are for instance, in their simplest form, financial agreements saying that the seller will compensate the buyer (usually the creditor of a loan) if the debtor having subscribed this latter defaults. Even without entering into technical details, it appears clear how the introduction of such three body interactions involving insurers, insurance buyers and reference entities leads to an extremely rich phenomenology; the events these interconnections may generate can be thought of as a factor of risk

themselves, as a precise forecast of what will happen under certain conditions becomes very difficult.

3.1.2 Epidemic processes and compartmental models

The focus of this chapter will be on epidemic processes, technically defined as dynamical evolutions of the states of the nodes of a graph [91]. The rules according to which a node changes its state may assume different forms, but they generally depend in some way on the state of its neighbours. As discussed in chapter 1, our aim is to study quite abstract models: the specific features of each node will be generally neglected, and all the information about an element will be condensed in a very schematic characterisation, according to which a node will for instance be referred to as *infected*, *recovered*, *vaccinated* or *susceptible*. In such *compartmental models*, each node can at each time take one in a finite (and usually small) set of possible values: using the formalism already introduced, $\sigma_i^t = \{1, \dots, q\}$. This definition is very near to the one of the *cellular automata* on grids in the computer science community, originally introduced as simple models to study self-organization, and in particular self-replicating systems [92].

In the simplest version, each element may take two values and its state evolves deterministically according to a definite rule involving its nearest neighbours; the rule for updating the state of cells, applied to all the network at the same time, is the same for each cell and does not change over time. Several generalisations have been proposed, for instance relaxing the binary states to finite sets, allowing a time dependency of the rules or even defining different rules for different nodes, studying asynchronous updates of the states or dynamics occurring on complex topologies [91], approaching more and more the compartmental models and the cellular automata studies. All these models have been faced during the years also from a statistical physics point of view [93].

3.1.3 The choice of the dynamics

In order to obtain quantitative results, one specific dynamics has to be defined. In a compartmental model framework, this is equivalent to specify the rules according to which a node changes its state [14]. The first big class of such rules is the one of the *irreversible dynamics*, so called as once a node goes from a state to another it will nevermore be able to return to the former. These dynamics are also called

monotonous, as the number of infected elements is forced to be non-decreasing over the time, or *unidirectional* since the state of a variable does not show any loop.

The simplest case is the susceptible-infected **SI model**, where an infected node will stay so for all the following part of the dynamics. This request on the dynamics may at the beginning seem very restrictive, but enables nevertheless to study very interesting problems. Thinking about the disease spreading, this condition is equivalent to study the evolution of the contagion on a time-scale shorter than the typical recovery time. For avalanche damages in infrastructural network, a logic hypothesis is that the cascade failures will occur in a very short time delay: the recovery of the system (for instance the arrival of people trying to fix the damages) will typically start when the avalanche has already reached all the available targets.

More general irreversible dynamics may be defined by considering variables taking more than two values. In the **SIR model**, a node can recover some time after having been infected; when this happens, it will no more be susceptible and it will never contract the disease any more. This aspect is of interest for studying diseases guaranteeing an immunity after the first time they are contracted. In the **SEIR model** it is introduced an *exposed* state one needs to be into in order to be infected. In the same context, one can also introduce vaccinated nodes, not participating to the dynamics, in order to see how their presence influences the spreading of the disease.

A different class of models considers the possibility for a node to go back to a state it has already left: they go under the name of *reversible processes*. The prototypical example is the **SIS model**, useful for modelling the long-term spreading of diseases not giving immunity; in such cases, after some time spent in the infected compartment an element becomes susceptible again.

Depending on the features of the system one wants to model, different combinations of the previous basic elements may be proposed. Just to let the idea be clearer, one can think a mixture of the SIS and the SIR model, in which the nodes finishing their infected period may or not acquire an immunity to a further contraction of the disease.

Another big difference in the evolution rules stands between *deterministic* and *stochastic* ones. In the former, each node in a given condition (defined by the state it and its neighbours are in) will behave in the same way; in the latter, the evolution from such a condition will be probabilistically defined, and may hence lead to different outcomes for undistinguishable nodes.

3.1.4 The characteristics of the graph

Having in mind the dynamics one wants to study, the features of the networks on which such a process takes place have to be defined. In particular, different choices for these latter will lead to different situations one may study. The graph being weighted, for instance, enables to describe the situation in which not all the links between nodes have the same importance; the same holds for the choice of a directed or a symmetric graph. If one aims at having reasonable numerical forecasts for systemic risks in the financial market, the graph has clearly to be supposed asymmetric as an institution being exposed to another one does not necessarily imply the other way round: even if this is the case, the two exposures should in general be considered different and quantitatively estimated separately.

The topology of the network underlying the dynamics has to be defined. When the objective is to obtain an accurate numerical description of the epidemic process, one is forced to explicitly run the dynamics either on the real network, or on a reconstruction as precise as possible if the former is not possible. In economics, the cascade failures of the institutions are simulated by using all the available information on who is exposed to whom, and how. During the last years, data are starting to be available also for many fields in which such a precise description has historically not been possible. This is the case of the social sciences, in which up to some years ago very controversial hypothesis on the interaction networks had to be made in absence of detailed information. The *SocioPatterns* project [94] offered reliable data on the structural characteristics of the face-to-face proximity network in real-world environments by using wireless wearable sensors. This project enabled to gain insights about how these networks differ by changing environments: among these, large medical conferences [95] and hospitals [96]. This collection enabled also to quantitatively compare the spreading dynamics in situations as different as a scientific conference and a museum exhibition [97]. Having access to these data allowed a better understanding of some phenomena. In [98] the authors noticed that in real-world networks the spreading is slower than one would expect in graphs in which couples of nodes are typically connected by short paths (small-world effect); by creating modified datasets without, for instance, clusters of nodes or temporal correlations among the links they were able to isolate the structural features the most important as causes of this slowing-down.

From a more abstract point of view, one can study spreading dynamics taking place on several types of topologies. A lot of work has been devoted to the study of spreading on examples of the so-called *complex networks* [99] and in particular on scale-free [100] and on small-world networks [101]. Also the dynamics occurring on

simpler topologies such as random regular graphs, lattices, or Erdős–Rényi graphs is far from trivial, and may lead to a rich phenomenology.

The focus on different topologies is useful to stress some qualitative differences in the spreading phenomena. The scale-free graphs are, for instance, very fragile if the most connected nodes are touched during a targeted attack, but very resistant to random attacks that will most likely damage peripheral nodes. In more homogeneous networks such as Erdős–Rényi graph, on the other hand, it is very difficult to plan a targeted attack as all the nodes have more or less the same structural importance in the overall network [14].

3.1.5 One spreading phenomenon, many possible questions

After having defined the dynamical rules and the topology, one still has to think about what to investigate. In such a general context, many different problems, both direct and inverse, may be of interest. One can for instance try to determine how a contagion will evolve from a given initial condition. This forward problem is central in public health studies: accurately forecasting the fraction of the population likely to be touched by the Ebola outbreak having access to data regarding the beginning of the contagion in western Africa has been recently a topic of the greatest importance. To the theoretical interest in understanding how an epidemics evolves one has to add a more practical one, as a quantitative estimate of the evolution of a spreading knowing a snapshot at a given time is essential in many cases.

The same holds also for the economical predictions. The stress tests performed by governmental bodies on some financial institutions may be seen as instances of the direct formulation of the spreading problem so to assess whether they would be able to cover losses induced by extreme, but still plausible, situations. These “what-if” simulations assign a certain initial condition on the system (for instance, the default of one country) and look at the possible outcomes of such a situation. If with some non-negligible probability one obtains an avalanche of credit defaults, for instance, a very dangerous element of systemic risk has been discovered, and some actions need to be taken in order to stabilise it.

A very wide spectrum of inverse problems can be addressed. A first typology makes use as before of the snapshot of the infection at a given time; instead of trying to determine its evolution in time, one searches in this case to reconstruct, for instance, the location and the time at which the spreading has originated. These *zero patient problems* has been recently addressed in many variants, for instance dealing

with the possibility of having access just to partial observation [102] or to noisy measurements [103].

Another group of interesting questions can be found in the general context of the optimisation problems. Since the aim of the dynamics introduced is to describe the spreading of a given characteristic on a network, the first big distinction is among the cases in which we want to hinder this propagation and the ones in which we want to facilitate it; the studies on marketing or on information dissemination belong to this latter group, the ones on the disease epidemics to the former. In order to understand the types of questions addressable it is useful to stress the two extreme and complementary forms an optimisation problem may take in both cases. It is worthy to stress that all these inverse problems are typically hard from a computational point of view, requiring a number of operations exponential in the system size to be exactly solved; ways of obtaining approximated solutions needs hence to be thought of.

In the first typology of problems one has some constraints on the initial state of the system, and wants to realise them in such a way to make the system be afterwards led by the dynamics to a desired and “optimal” in some sense final state. This is the case, for instance, when in a spreading disease process we have the possibility of vaccinating only a finite fraction of the population: the solution of this type of optimisation problems in this framework will tell us whether we can select the elements to vaccinate so to block a global outbreak, and how to choose them so to obtain such a result. The very same holds also for the situation in which we wish to reach the largest possible audience by addressing a message to a fixed and much smaller number of initial spreaders.

The second typology consists in the cases in which the desired final state is fixed. Thinking again about a public health problem, one may wish to prevent the outbreak of a disease; in this case, the central issue is to evaluate the minimal number of people to vaccinate (and how to choose them) in order to obtain such a result. The same in a viral marketing campaign, in which the interest is in the understanding of how many initial targets should be addressed so to propagate the message to the desired fraction of the population.

These problems should not be seen as juxtaposed. Because of the generality of the statistical physics approach, they can be addressed in very similar way. Even more interestingly, also intermediate cases can be studied, in which *penalties* may be assigned for instance both to the nodes chosen as initial spreaders (so to look for configurations with a low enough number of them) and to the nodes not touched by the spreading (so to find the initial targets of a marketing campaign able to efficiently

propagate the message). In some sense, the problem of interest is how to find the best possible trade-off between these two components; a very natural way of addressing this problem will be discussed in the following.

Another parameter of central importance (especially when thinking about applications) is the time requested by the spreading to occur. In marketing campaigns, for instance, one is interested in reaching the largest possible number of customer in a short enough time delay. For finding initial configurations satisfying such a principle, one may wish to add a constraint fixing the maximal time during which the spreading has to take place.

Using a constraint optimisation terminology, all the aspects described so far may be imposed either as *soft* or as *hard constraints*. In the former case, penalties may be imposed to the number of seeds or to the number of untouched nodes at the end, and also initial configurations leading to slow spreading may be penalised in favour of faster ones. In the latter, for instance, we restrict to the spreading dynamics reaching a given fraction of the nodes in no longer than a fixed time period T , or to the “percolating” configurations (i.e. the ones leading to a complete activation of the network).

3.1.6 Different approaches

As discussed, the spreading dynamics is of great interest in many context, and on which many different questions may be asked. Not surprisingly, it has been addressed in many different ways.

In [104] a large deviation approach has been proposed to quantitatively describe the credit contagion in economical system, and in particular to better assess the systemic risk at given conditions. Because of feedbacks, indeed, the interactions among firms may in situations of economic stress be much stronger than the typical ones.

For obtaining the most possible accurate estimate of the spreading of a given disease, one can run stochastic dynamics on real-world networks, taking into account both long-range (i.e. airline communications) and short-range (i.e. commuting) transportation systems, as well as precise demographic information on the different regions of interest [105]. These studies typically rely more on numerical simulations than on analytical results.

A somewhat completely opposite approach is to completely ignore the real networked structure of the systems, and to make the simplifying hypothesis of *fully mixed* populations. In this case, an infective individual is equally likely to spread the disease

to any other one. This strong (and questionable) assumption allows to write down systems of differential equations governing the dynamics of large systems; by solving these latter, one can monitor the number of infected as a function of time, search for the equilibria of the system at given conditions, or infer the typical fraction of the population touched by the infection [106].

In the spreading optimisation problem, one can try to characterise the nodes which are going to be the most effective initial spreaders. Intuitively, one may infer such a piece of information from a local topological analysis: a node with many neighbours should be a better spreader than a very isolated one. It turns out that such a local information may not be enough, and to precisely identify the best spreaders a global information as a whole is needed. The differentiation of a *core* and a *periphery* of the network proposed by the authors of [107] does not lead to the same results as a local, degree-based description. The global topological importance of each node can be assessed also in other ways, for instance by counting the number of shortest paths passing through it as done in [88] by defining the *betweenness centrality* of each element.

As will be briefly described in the following, finally, many specific instances of spreading phenomena can be mapped to problems of interest in probability theory. In such a context, mathematically rigorous results (for instance giving bounds on the minimal fraction of spreaders needed to obtain a complete contagion on a given network) can be looked for [108].

3.2 Minimal contagious sets in random regular graphs

In the following sections, the work presented in A will be discussed. Since all the technical details can be found therein, the aim of the rest of this chapter will mostly be to give some keys to interpretation, especially in connection with the general concepts discussed so far.

3.2.1 Definition of the problem

The work reprinted in A is focused on how to find minimal contagious sets on random graphs. The dynamical rule of the **SI** type, each node being either susceptible $\sigma_i = 0$ or infected $\sigma_i = 1$. The rule chosen to determine when the state of a node has to switch from susceptible to infected goes under the name of *bootstrap percolation* or *threshold model* [109]. According to it, each susceptible node i is infected as soon as a *sufficient* (i.e. larger than a fixed threshold l_i) number of its neighbours are infected. The evolution of the system under such a dynamical rule, occurring in discrete time,

is deterministically fixed for any assigned initial configuration of spreaders (also called *seeds* of the infection in the following), and the state of each node i at time t is a function of the state that i itself and its first neighbours were into at time $t - 1$:

$$\sigma_i^t = \begin{cases} 1 & \text{if } \sigma_i^{t-1} = 1 \\ 1 & \text{if } \sigma_i^{t-1} = 0 \text{ and } \sum_{j \in \partial i} \sigma_j^{t-1} \geq l_i \\ 0 & \text{otherwise} \end{cases} \quad (3.1)$$

In this case, a very useful simplification in how to describe the state of each node during the process can be introduced; instead of studying the complete trajectory formed by the σ_i^t for all t , one can equivalently describe each node by a single number, its activation time $t_i(\bar{\sigma}) = \min\{t : \sigma_i^t = 1\}$ at which it moves from the susceptible to the infected compartment. A second crucial simplification is that the interactions are in this formalism local, as the activation time of a node will depend just on the activation times of its neighbours (and, generally speaking, on the value of its threshold). The extension to more involved irreversible processes such as **SIR** models may be proposed [110]; on the other hand, the reversible dynamics such as **SIS** are much harder to deal with, as such simplifications cannot be used and one is forced to explicitly study the complicate relationships regulating the state of each node at each time, as a function of (in principle) all the other nodes.

These dynamical processes are very generally defined, and the features of such spreading phenomena may be studied on any type of networks. An interesting extension is to weighted graphs: in such cases each active node sends a given activation contribution to its neighbours. A susceptible node is hence activated if the sum of its incoming contribution exceeds its threshold. This model is of particular interest in neuroscience, the weights being representatives of the strengths of the links connecting couples of neurons in the *integrate-and-fire* model [37]. Another interesting possible extension is to directed graphs, as in many applications the networks are intrinsically asymmetric: the fact that a node i may contribute to the infection of a node j does not always imply the other way round.

In order to obtain analytical results, infections on random regular graphs where all the nodes have the same number of neighbours and the same activation threshold has been addressed in this work. This case is very simple, and yet it leads to very interesting phenomena; some possible ways to relax such conditions are nevertheless discussed as future work perspectives.

3.2.2 The energy function

Dealing with such a defined dynamics on a graph, one can very naturally introduce an optimisation problem where the variable to be optimised upon is the initial condition of the system and the *objective function*, minimised by the optimal choice of the initial configuration, that depends both on the initial and on the final state of the system. Both of them will have to be the most possible similar to a conditions defined as “optimal” for the problem, and the relative strength by which we require such similarities to hold will lead to slightly different formulations of the problem. In order for the equations to be more easily readable, one can define a shorthand for the initial condition of the dynamics by calling $(\sigma_1^{t=0}, \dots, \sigma_N^{t=0}) = \bar{\sigma}^{t=0} \equiv \bar{\sigma}$. Since the time evolution is deterministic, a cost can be assigned to any initial configuration $\bar{\sigma}$ as follows:

$$E(\bar{\sigma}) = \sum_i \left[\mu \sigma_i^0 - \varepsilon (1 - \sigma_i^T) \right] \quad (3.2)$$

where σ_i^t represents the state of the node i at the time t . Each node gives to the configuration cost a contribution μ if it is chosen as a seed, and ε if it is still inactive at time T , the maximum time allowed for the spreading to take place.

This formulation has various interesting limit cases. If $\varepsilon = 0$, no optimisation is performed on the seeds, which are independent identically distributed Bernoulli random variables parametrised by μ . If $\varepsilon = \infty$, on the other hand, the configurations in which not all the nodes get eventually activated have an infinite cost; this is equivalent to set the complete activation of the network as a hard constraint. The problem we mainly focused on in this paper consists in taking to take this $\varepsilon \rightarrow \infty$ limit, and then the $\mu \rightarrow -\infty$ one: in this way, just the minimal percolating configurations are considered as possible solutions. For finite values of ε and μ , on the other hand, the trade-off between number of initial spreaders and the fraction of the system activated may be analysed. Another limit of particular interest, even if not easily implementable in some cases, is $T \rightarrow \infty$: this corresponds indeed to finding the minimal configurations able to activate all the system without stopping the system evolution at some finite time T .

3.2.3 Mapping to other standard problems in graph theory

Many special cases of the spreading optimisation problem may be mapped to well-known problems in graph theory. This can be very useful as some rigorous results

obtained in the mathematical community may be easily translated in properties of the spreading phenomena, for instance in terms of minimal fraction of seeds needed for a complete activation of a graph under certain conditions. In some cases, moreover, one can confidently look for some interesting features of the spreading dynamics, by knowing that they hold in graph theoretical problems that are just special cases of the former. I will describe in this section some interesting examples, while referring to A for a complete list.

The first equivalence can be seen if, on an arbitrary graph, each node gets activated only if all of its neighbours are themselves active. In such a case, it is easy to see how, in order to reach a global contagion, no couple of neighbouring, inactive nodes has to be present in the initial configuration; if this is the case, moreover, the complete activation will take place in just one time step, regardless of the fixed temporal bound T . In the graph theory terminology, this condition is expressed by saying that the initially inactive nodes have to form an *independent set* [46]. The optimisation problem is hence equivalent to the problem of finding the largest independent set on the given graph.

In some sense opposite to this case is the one in which all the threshold are equal to 1 and we impose the spreading to take place in one time step. For a complete activation to be reached, all the nodes should either be seeds, or have at least one of their neighbours that is so. This condition is expressed in the graph theory by saying that the seeds have to form a *domination set* of the graph, and the spreading optimisation problem under these conditions is equivalent to the search for a minimal dominating set of a graph.

3.3 The cavity method treatment

3.3.1 The replica symmetric RS formalism

Without entering into technical details extensively discussed in A, I will intuitively explain how this optimisation problem may be addressed by using a cavity method approach [111]. In order to do that, one needs first of all to write down the situation under exam in a form more reminiscent of statistical mechanics problems. This can be achieved by introducing a probability distribution on the initial conditions $\bar{\sigma}$ with a weight given by the 3.2 such that

$$\eta(\bar{\sigma}) \propto e^{\sum_{i=1}^N [\mu\sigma_i - \varepsilon(1 - \sigma_i^T)]} \quad (3.3)$$

This expression is quite tough to deal with, as it contains the variables σ_i^T depending in a complicated way on all the variables σ_j at a distance from i smaller than T . By exploiting the fact that the dynamics is both deterministic and monotonous, the previous can however be rewritten in a simpler form in terms of the activation time t_i of each of the variables i . This having being done, one can introduce on each link of the graph a *message*, i.e. a probability distribution over the activation times of the two nodes involved $\eta_{i \rightarrow j}(t_{ij}, t_{ji})$. These quantities obey recursive relations, as any message is a function of the messages “incoming” from the other neighbouring nodes:

$$\eta_{i \rightarrow j} = \hat{g}(\{\eta_{k \rightarrow i}\}_{k \in \partial i \setminus j}) \quad (3.4)$$

If the iterative implementation of these conditions converges, from the self-consistent messages the quantities of interest can be calculated: among them, the probability for a node of being active at time $t = 0$ (i.e. its probability of being a seed in an optimal seed configuration) and the minimal density of seeds needed to activate all the network at time T . This approach, called **replica symmetric**, is based on the assumption of the $\{\eta_{k \rightarrow i}\}_{k \in \partial i \setminus j}$ being independent on j . This is true for all messages only if the underlying graph is a tree; however, if the graph is locally tree-like and the correlations between the messages decreases fast enough, this approximation is still good and the equations obtained this way go under the name of *loopy belief propagation*.

3.3.2 The breaking down of the RS assumption

The previously written equations in the previous sections are very general and could refer to any kind of graph, any node having an arbitrarily chosen degree and activation threshold. All the results reported in appendix A refer nevertheless to the specific case of a $k + 1$ random regular graph with homogeneous activation thresholds $l_i = l \forall i$.

The prediction of the density of seeds for different values of μ one obtains from the converged messages in the $k = l = 2, T = 3$ case are reported in the left panel of figure 3.1 and seems at first very reasonable. As μ becomes more and more negative, percolating initial configurations with fewer and fewer seeds are looked for, and the optimisation becomes stricter and stricter: the fraction of seeds θ is monotonically increasing with μ , but it seems possible to extrapolate a minimal θ for $\mu \rightarrow -\infty$. The situation is however quite different for other observables. For strong enough




Figure 3.1: The density of the seeds of the infection θ (in the left panel) and of the entropy s (in the right one) as a function of the parameter μ playing the role of a chemical potential. The plots refer to the results obtained by making use of the replica symmetric ansatz on a 3-random regular graph with homogeneous activation thresholds $l = 2$ and maximum time allowed for the dynamics $T = 3$.

optimisation, indeed, the entropy of the system becomes negative: this is a clear proof of the RS assumption not being satisfied any more for μ negative enough.

This feature, observed in several rCSP such as kSAT [112] and qCOL [113] for large enough values of k and q , is due to a phase transition leading to a drastic reorganisation in the structure of the space of the solutions in the satisfiable phase, in which solutions to the problem typically exist. Intuitively, one can think that as soon as the constraints get very hard to satisfy, long-range correlations among far-away variables start to appear because of the stronger and stronger frustration present in the system. In order to implement such constraints, hence, the knowledge of local information on the neighbourhood of a node are not enough any more, and a global knowledge of the system is requested.

More precisely, several phase transitions of this kind have been found in rCSP. They tend to show the same pattern (figure 3.2) regardless of the specific definition of the CSP under study. Using the terminology of the optimal spreading problem, if solutions with a large enough fraction of seeds θ are looked for, the problem is not very much frustrated, and the RS assumption of correlation decay holds. If this θ value is decreased, a first *dynamical* transition is found at θ_d . The space of the solution is split into clusters, but the thermodynamic predictions obtained by making use of the RS ansatz still hold. However, for even smaller values of θ (i.e. for even harder constraints) another threshold is encountered at θ_c : below this *condensation* transition, the RS prediction are no more true, and is in this phase that unphysical results such as the negative entropy described above may be found. In order to study this phase, hence, another approach able to take into account the long-range




Figure 3.2: Pictorial representation of the phase transitions pattern usually observed in random CSP, as they get harder to solve (in this case, for the initial density of seeds θ in a percolating configuration getting smaller). For $\theta > \theta_d$, the replica symmetric ansatz is correct and there is no non-trivial solution of the 1RSB equations. For $\theta_c < \theta < \theta_d$ the system enters a dynamics 1RSB phase characterised by an exponential number of clusters contributing to the Gibbs measure; the RS estimates are still correct in this parameter region. For an even smaller value of θ (i.e. even harder optimisation problem) the number of cluster contributing to the Gibbs measure becomes sub-exponential and the RS predictions of the thermodynamic quantities are no more correct. For $\theta < \theta_{min}$, finally, no solution can be found any more and the problem enter in its unsatisfiable phase.

correlations emerging in the system has to be used.

3.3.3 1RSB and energetic 1RSB formalism

The first extension of the RS assumptions able to take into account such correlations is the *one step symmetry breaking* formalism. In this latter, the configuration space is supposed to be broken into *clusters* such that the correlation decay assumption holds inside each of them separately. In this formalism, the iterative equations get harder compared to the RS ones, as each link is represented by a distribution of messages: it will assume different values according to the cluster it belongs to. As before, each of these probability distributions may be written as a function on the incoming ones, giving self-consistency equations of the form $\hat{P}_{i \rightarrow j} = \hat{G}(\{\hat{P}_{k \rightarrow i}\}_{k \in \partial i \setminus j})$.

The verification of the existence of a non-trivial 1RSB phase is one of the main contribution of the work in A with respect to previous works such as [114]. Such a result was however in some sense expected. As well as for the cases described in section 3.2.3, also for $T = 1$ the spreading problem can be mapped to another one. This latter, known as the Biroli-Mézard model [115], is used to calculate the closest possible packing of spheres on a lattice in presence of repulsive interactions among particles. Since this model was already demonstrated to show a 1RSB phase, such a

feature in the more general spreading dynamics was sure for the $T = 1$ case and very likely also for larger temporal boundaries.

3.3.4 The “energetic” 1RSB formalism

The resolution of the 1RSB equations is quite complicated, since as already discussed they refer to probability distributions. Practically, they can be addressed by using a *population dynamics* approach, corresponding to represent the probability distributions $P(h)$ as weighted samples of fields h_i [111]. It turns out that for a case of particular interest, a simplification going under the name of *energetic 1RSB formalism* is possible. The equations discussed above allow one to study a very general case of the spreading dynamics, i.e. to find the initial density of seeds θ as a function of arbitrarily chosen μ and ε . As discussed, however, the original optimisation problem consisted in the limit $\varepsilon \rightarrow +\infty$ and $\mu \rightarrow -\infty$, so to select just the minimal percolating configurations. After some technical passages described in A, the energetic formalism allows one to directly study such limits. In this way, it is no more possible to obtain for instance the complete curve of θ as a function of μ : the estimation at the 1RSB level of the most important quantity θ_{min} may on the other hand be derived in a much easier way.

3.4 Main results

3.4.1 The solutions of the problem

Different types of analysis can be performed on the presented problem. The first one goes under the name of *single link analysis* and makes use of the symmetry of the system (all the nodes of the graph having the same degree and the same threshold). Thinking about the RS case, we saw how each message can be written as a function of the incoming ones; this leads to recursive equations of the type 3.4, formally expressible as $\eta = f(\eta_1, \eta_2, \dots, \eta_k)$ in the case of a $k + 1$ regular graph. Because of the symmetry of the system, one can look for a single value of η satisfying a self-consistency equation of the type $\eta = f(\eta, \eta, \dots, \eta)$: if such a value can be found, the thermodynamic quantities of interest such as the fraction of seeds in minimal percolating configurations can be determined. The same approach can be used also within the 1RSB formalism, in which a unique probability distribution P able to satisfy an equation of the type $P = G(P, P, \dots, P)$ has to be looked for. In both cases, the final goal will be to obtain some analytical results in the large system size

limit. Some of the results obtainable in this way will be discussed in the following section 3.4.2; of particular interest is the determination of the minimal fraction of seeds needed to have a percolating configuration in the thermodynamic limit, and the comparison with rigorous results obtained following completely different paths in the mathematical community.

Another approach, discussed in section 3.4.3, is to algorithmically study finite size systems. In these cases, explicit configurations able to activate the whole network are eventually found, and different strategies for selecting the seeds to be added can be compared. Also in this case, both the RS and the 1RSB levels can be studied.

3.4.2 Analytical results

One of the major analytical achievements of this work has been the explicit solution in the $T \rightarrow \infty$ limit of the RS and of the 1RSB cases. This special case is of particular interest as it corresponds in some sense to the original maximisation problem, in which no constraint on the time taken to activate the whole graph is imposed. Considering $k+1$ random regular graphs in which every node has an activation threshold equal to l , it turns out that the $k = l$ and the $k > l$ cases are qualitatively different. Interestingly, this is the case also without optimising the choice, as discussed in A; also a much more diffuse discussion on the results can be found therein.

If $k = l$, every node needs all but one neighbours active in order to get activated. An already known [116] lower bound on the fraction of seeds $\theta_{min}(k)$ needed for a complete activation in this case is the following:

$$\theta_{min}(k) \geq \frac{k-1}{2k} \quad (3.5)$$

More recent works in the mathematical community [108] proved such a bound to hold tightly for $k = 2$, and the same property was conjectured also for $k = 3$. With our approach, we confirmed the conjecture for $k = l = 2$ and for $k = l = 3$, and we found 1RSB estimates of θ_{min} for larger values of k . Since RS and 1RSB results for θ_{min} are different in these latter cases, it is not clear yet whether these estimates are correct or further levels of replica symmetry breaking are needed.

Also for the $k > l$ case a lower bound for the initial seed density was known:

$$\theta_{min}(k, l) \geq \frac{2l - k - 1}{2l} \quad (3.6)$$

Also in this case, we managed to verify this bound to hold with an equal sign both for $k = 3, l = 2$ and for $k = 4, l = 3$, results that were not even conjectured before; we were also able to obtain 1RSB estimates for other values of k and l .

3.4.3 Numerical results

An alternative approach to determine the thermodynamic limit of quantities such as the minimal fraction of seeds in percolating configurations is the explicit construction of such configuration on finite size graphs. These are known as *single instance* results, as they do not refer to statistical ensembles of graphs, but to specific realisations.

The core of the procedure is as follows. Firstly an explicit instance of the graph of interest (for instance a $k+1$ random regular graph with 10000 nodes) is generated. Then the first seed is chosen according to some rule, and the direct dynamics is performed so to see which final state it leads to; for instance, whether or not it leads to the complete activation in a given number of time steps. If one is looking for minimal percolating configurations and the answer to the previous question is no, another seed is added according to the same rationale as before, and the procedure is repeated until an initial set of spreaders rich enough to percolate is found. In order to obtain more stable results, this single instance analysis can be repeated both on multiple realisations of graphs with fixed characteristics, and on several runs on the same instance. The interest of this framework is that it can be applied to different strategies, so to compare their results. In A two of them are taken into account, and the initial densities of seeds respectively obtained are compared.

The first one belongs to the class of the greedy algorithms, generally speaking used to solve problems step-by-step by iteratively doing what seems best in the short term [25]; more precisely, the rule according to which the nodes to be set as seeds are chosen is the following. Let us suppose that the current seed configuration is not able to activate all the network, and hence we have to add at least one extra seed. All the seeds configurations composed by all the previous seeds plus one are tried, and the number of nodes they are able to activate during the process is recorded. If one of the configurations built in this way percolates, the algorithm stops; otherwise, the node leading to the greater improvement in the final fraction of activated nodes is added to the seed configuration, and the procedure is iterated. An element of interest of this algorithm is its straightforward generalisation to $T \rightarrow \infty$ limit. In this case, indeed, it turns out that it is possible to avoid the study of the effect of the extra seed on the complete process, as it can be equivalently studied its effect on the configuration of the nodes activated at the end of the previous iteration. The densities found by this algorithm are not guaranteed to be optimal, and indeed in many cases the number of seeds needed to get a complete activation is quite larger than the one one would expect by looking at the bounds discussed above.

The other tested algorithm makes use of the statistical physics description of the problem. This is an extension of the *maxsum* message-passing procedure proposed in [114, 117] in which just the replica symmetric description was taken into account, obtained by considering also the replica symmetry breaking effects. Referring for the details to A, I will just briefly describe here the main idea. As described before, from the converged messages in the 1RSB formalism is possible to derive estimates of quantities such as the initial density of seeds. Calculating the contribution of each node to this latter, the tendency of every node to be a seed in a minimal percolating set can be evaluated. The decimation procedure [118] used consists hence in iteratively calculating this tendency for all the nodes not yet fixed to seeds, and eventually to set the one with the highest contribution as an initial spreader. This approach leads to very good results for small T , as the density of the percolating configurations built in such a way is very near to the theoretical bound; in particular such densities are smaller than the ones obtained by algorithms not taking into account replica symmetry breaking effects. For larger T some convergence issues start to appear, and expansions for directly studying the large T limit should be probably looked for in the future.

3.5 Future perspectives

The present study led to some interesting insights on the contagion dynamics on random graphs, and stressed some features that should be better studied and understood in the future.

Having shown that the assumption of replica symmetry does not hold under some conditions is not enough to affirm that the 1RSB approach is correct; indeed, one should verify the stability of such an ansatz with respect to further levels of replica symmetry breaking. This is true in particular for the cases in which the estimates of the minimal density of seeds at the RS and 1RSB are different, as the value obtained by the 1RSB description could again be just a lower bound on the true value.

Secondly, some asymptotic expansions should be tried, so to be able and analytically address some limits of interest. Among these latter, one can think about the large connectivity one $k \rightarrow \infty$, performed so far just in the special case $k = l$. Several non-trivial cases can be thought of in such a limit, for instance by keeping finite the threshold l , or by keeping fixed either the threshold $\frac{k}{l}$ or the difference $k - l$.

Even if we restricted our study to the random regular graphs, the single sample equations derived are very general and can be applied to any type of network with

arbitrarily fixed degree distribution and activation thresholds. In particular, one would like to assess the effect on the minimal density of seeds induced by fluctuations on these quantities. Completely different types of graphs, for instance showing scale-free degree distribution, could be studied with the same methods; in principle, one could also run the very same equations on real-world networks, for example to better understand the systemic risk in economic systems.

Finally, one can address different questions of interest in the same framework, as discussed in section 3.1.5. Many of these extensions can be very straightforwardly obtained from the results obtained in this work: one such example is how to optimally choose a fixed number of spreaders in order to maximize the fraction of nodes eventually getting activated. A better understanding of this problem could be for instance central for designing viral-marketing campaigns in which just a fixed budget can be spent for the kick-off of the message.

Chapter 4

Exploring networks

4.1 Graph exploration and random walks

4.1.1 Exploring a graph, a very general problem

When thinking about the variety of systems modelable as networks, a large set of possible questions about them could go under the name of the ways by which it is possible to explore them. How long does it take to reach a point from another one, how is probable to fall again on the starting point when randomly travelling over the links or what is the best possible route connecting a set of points are all specific instances of such questions.

Understanding the characteristics of paths on graphs is a very classical task. Both the forward and the inverse versions of such a problem are of great interest. Among the former we could insert the problem faced in appendix B, in which the statistics of rarely observed paths are studied. Among the latter, one can think of many important optimisation problems; the design of a transportation network able to minimise the average time needed to join two locations is just one out of many possible examples [48], some others will be discussed in the following sections.

The comprehension of how to navigate through a network is all but a precise task; indeed, it has been addressed in many different ways [91]. Some historical examples of these diverse approaches will be presented: the main goal of this chapter will nevertheless be to stress the richness of problems of interest in this context more than to give an exhaustive review of all such problems.

4.1.2 Interest and applications in different fields

The desire of understanding how the exploration of a network can be performed has been present since the very beginning of this field. The Königsberg bridges problem,

having been presented in chapter 2 as the one which gave birth to the network theory, is essentially one such problem since it asks whether it is possible to completely explore a given graph by following a route with some specific characteristics. The problem of searching for Hamiltonian cycles over a graph is similar but much harder to solve, being an NP-complete problem in the worst case scenario analysis discussed in section 1.4.3.

A problem whose applications to real-world situation are even simpler to imagine is the so-called *traveling salesman problem*, aiming at finding the shortest path connecting a series of nodes in a network. This problem has been deeply studied in theoretical computer science [119], and it has been shown to belong to the NP-complete class. Since the exact solutions of such problems is practically impossible even for quite small system sizes, several heuristics have been proposed to obtain at least approximate solutions [120]. As could be expected after the discussion on the connection between constraint satisfaction problems and statistical physics of chapter 1, this problem has been deeply analysed also under this perspective [121].

As already stated, one can try to reach a better understanding of how networks may be explored following different paths. One of the most interesting inverse problems in such a framework is how to build a system so to minimise the time spent for moving from one point to another: this is of central interest in the planning of infrastructures and of IT systems. A very rich phenomenology is known to emerge in these situations, including some counterintuitive effects. By analysing the problem as a game theoretical one, it has been for instance observed that adding extra routes to a system may lead to an increase in the average time needed for moving, this latter having supposed to be for each road a function of some fixed features (its length or its quality for instance) and of the number of cars travelling on it (in order to take into account possible congestion effects). In particular, if all the drivers are selfishly performing an optimal self-interested decision as regards which route to choose, the presence of a short cut could convince a very large part of them to use it, so to obtain as a final collective result the emergence of traffic jams along the short cut itself. The paradox is in the fact that by removing this latter and by therefore forcing the drivers to more homogeneously split among the slower routes, the average time requested for the travel can significantly decrease. This feature, going under the name of Braess's paradox [88], demonstrated that adding resources to a transportation network without performing an a priori forecast of their effects can in fact create unwanted incentives seriously undermining its efficiency. Such counterintuitive phenomena may happen also on different types of networks such as electrical power grids or biological systems

and show that, in principle, a malfunctioning network could be improved by removing certain (apparently fundamental) parts of it.

The last example of exploration in networks that I will propose is particularly clear when thinking about social systems. The mechanisms regulating the search, in this case, are the ones by which “people can explore chains of social contacts for information or referrals to others [88].” The first quantitative analysis of these aspects was probably tried by the American psychologist Stanley Milgram in the sixties [55]. In his famous experiment, 296 randomly chosen people (the *starters*) were given a letter and some details such as the name and the address of a unique target to reach. Everybody participating in the experiment had to forward the letter to somebody they knew on a first-name basis in order to try and reach the target in the smallest possible number of steps. Many of the letters never arrived to the final target; however, the median length of the 64 chains which closed successfully was six. The presence of short paths connecting any pair of nodes facilitating a very fast navigation of the network has been confirmed by more recent experiments of this kind; these latter have become much easier during the years, as for instance the list of personal contacts of each of us is nowadays explicitly recorded in many databases such as our friendships on the social networks.

4.1.3 Random walks on graphs

A well studied statistical problem is to understand the features of a *random walk*. This is defined as the path covered by a particle hopping on a lattice, randomly choosing at each time step which of its neighbours to jump onto. The easiest case of such a path consisting of random steps takes place on a single axis, on which the random walker decides at each time step whether to move rightwards or leftwards by one length unit. This problem, far from being trivial, has continuously received a great deal of attention since its first appearance [122]. Formally, indeed, it is equivalent to the stochastic process built by successively adding independent, identically distributed random variables [123]. The random walk was hence just another way of looking at one of the most basic and well-studied topics in probability theory. If in the previously defined process one assigns a variable $+1$ to each step on the right and -1 to each step on the left, it is easy to see how the final position of the random walker on the axis will correspond to the sum of such randomly extracted variables.

Even this very simple model can be applied to situations of interest. The successive bets of a gambler having originally a certain amount of money, indeed, can be seen as events leading to the winning or to the losing of one unit of its total wealth. A

well-posed question is how long will the gambler survive, i.e. how many bets will it take on average to exhaust his money reservoir. This parallel is not as naive as it could seem: random walks are indeed one of the basic elements for modelling the price of a fluctuating stock in the financial market [124].

The applications of random walks are really ubiquitous. This is probably due to the fact that by them it is possible to model a dynamics occurring on a given graph without needing any precise assumptions on the reasons because of which such a dynamics takes place: at a first level of abstraction, hence, many dynamics can be modelled as random walks, even if it is very reasonable to imagine that they are not purely governed by randomness. I will briefly discuss two other examples, very different from the financial one already mentioned. Firstly, random walks are largely used for modelling biological systems [125]. In particular, they are used to model how bacteria can move to regions where the density of nutrients is higher, process known as *chemotaxis* [126]. Secondly, they may be used to understand how gas particles move. The path traced by a travelling molecule can indeed be thought of as a random walk as, in this case, the interactions with all the other particles are extremely complicated. A useful though simplifying approach to this situation is hence to model such interactions as random events changing the direction of the molecule under observation.

Especially this second application is important in order to understand the relation of random walks to diffusion models known as *Wiener processes*: in the limit of infinitely small (and infinitely numerous) steps, the random walk is indeed equivalent to a Brownian motion [127] occurring in continuous time. Many other insightful extensions have been proposed and studied. Steps of variable size, trends or bias pushing the system in one direction preferentially, quantum implementations of the random walks algorithms [128] are all deeply studied generalisations.

4.1.4 Joining the two: graph exploration through random walks

In many applications, the two ideas presented in the previous sections are strictly related, as random walks have become a privileged way of exploring networks. In many cases, indeed, one does not have access to global information on the system under exploration, and has hence to rely only on local characteristics. In many contexts, even having such a global knowledge would not be enough to explore the graphs according to system-wide heuristics, unmanageable from a computational point of view for large enough system sizes.

As usual in this context, also as regards search inefficient solutions may be implemented quite easily. For instance, the *breadth-first search* [129] algorithms start from a randomly selected node, and examine all of its neighbours. All this information is stored, and the neighbourhood of the previously reached nodes is explored (getting an exact knowledge of the portion of the graph at distance smaller or equal than two from the starting node). This solution is potentially very slow and very demanding in terms of the quantities that one needs to keep into memory, especially for the gigantic graphs (such as the Web) one would like to analyse by such algorithms. The problems get even worse if the graph under exam is changing over time: in this case, because of the long time requested by the procedure, this exhaustive search is not even guaranteed any more to produce the correct result. These strategies can be refined if one has access to slightly less local information. For instance, if one knows the degree of all the neighbours of the node examined at a given time, a so-called *maximum-degree-strategy* [130] may be implemented, by exploring only the part of the graph we can reach by going to the neighbour with the largest degree. This heuristics is particularly efficient on scale-free graph, as few steps are typically needed for reaching a hub, connected with almost anybody else in the network.

Searching with very local information is hence useful in two different types of situations. In the first, including for instance the optimal routing in IT infrastructure, it can be shown to be sufficiently efficient [48]. In the second, for instance on the Internet, it is just impossible to do anything else as nobody knows the entire network structure, and even if somebody knew it, it would not be manageable to use all the information.

The PageRank algorithm [85], introduced in the Google research engine in order to measure the importance of the web pages, is basically an application of random walks to the Internet graph. An intuitive explanation of this algorithm can be given in terms of *random surfers*, which move through the web pages by following random links on each page they fall onto. If, as it appears reasonable, we accept the hypothesis that the most reliable pages are the one with more incoming links from other ones, then the time spent on each page by the random surfer is a proxy of the importance of the page. This very basic principle has been refined over and over again, for example by implementing an iterative procedure following the idea that a page is important if it has many incoming links from other important pages, and so on. Other improvements were for instance the *teleportation* of the random surfer with some probability to any other web page, so to obtain significant results also in case of web pages with no out-link. Such an application enabled to drastically improve the performances of the

engines, at the same time significantly lowering the noise introduced by spam pages or links [62].

The principle of random surfers is used also for performing other functions. Among these, an iterative update of the knowledge on the structure of the Internet. In this case, since both the web pages and the links connecting them are appearing and disappearing continuously, one would in principle like to update our prior knowledge almost in real time, so to always have the most accurate possible vision of the network one has to deal with. In this sense, the random walk ideas are really used in order to explore the graph under exam. The same basic mechanisms are also useful for estimating the traffic on a given page, or on a given portion of the Web [85].

As usual for problems in graph theory, also the exploration of a network by random movements has been addressed in a more abstract way, trying to quantitatively determine some features of these phenomena. For instance, the return time after which a random walk returns on the node from which it originally left, the hitting time after which it is able to reach an arbitrarily selected node, the cover time after which it has been able to explore the entire network or the effort needed to reach some regions difficult to see with just local information are all problems having been addressed in a theoretical sense. The questions addressed in appendix B are also of this kind, not dealing directly to the solution of a practical problem such as the removal from the search results of spam pages. In this context is however particularly easy to see how an improvement in the understanding of these phenomena could immediately lead to the design of better algorithms or procedures in the many fields in which the random walks are used as basic elements.

4.2 Rare event statistics on random walks on networks

4.2.1 The problem

In the rest of this chapter the paper [2] reprinted in appendix B will be discussed. This work is somehow the first to investigate the large deviation phenomena occurring while networks are explored through random walks. As discussed, random walks are widely used for analysing, organising or performing tasks on networks; since rare events may lead to severe consequences in such situations, an accurate and quantitative forecast of the possible large deviations is of particular relevance for a comprehensive assessment of the performance of optimisation and search algorithms.

The understanding of the rare events is of interest in all the fields already mentioned in which the exploration of a graph via a random walk is used, from marketing to data transport. A case of peculiar interest is the cybersecurity, as worms and viruses typically perform random walks on networks. In this situation, a rare event could consist in a path hitting an atypically large number of sensible nodes, whose damage can block the functioning of the entire network. In this work, a statistical characterisation of such events has been addressed for the first time; even if much work has still to be done as will be discussed in the future perspectives section, the first results obtained are already very promising.

More specifically, we study systems in which a variable is associated to each node. These variables are in principle completely arbitrary: in particular, they can be uncorrelated or correlated with the degree of the node, or deterministically determined by this value. The quantity characterising each path will be the sum of all the variables encountered by the random walker during the path, and averages over all the possible paths will be performed. The rare paths can be more easily studied as the ones which are typical if the probabilities of hopping from one site to another are opportunely deformed. As will be seen, paths exploring preferentially highly connected or sparsely connected regions can be generated by making use of such deformations.

By tuning the level of deformation imposed to the paths, two different types of phase transitions have been observed. The first, somehow expected, is a *localisation transition* telling us that the very extreme fluctuations are events taking place on a small subset of the nodes in the systems (for instance, paths just exploring the core of the network, its most highly connected part). The second one, of even greater interest, is a *mode-switching transition*, telling that there can be abrupt changes in the way rare events are realised by changing the parameter regulating the deformations of the paths.

4.2.2 The model

As discussed in chapter 2, a generic graph can be represented in a compact form via its adjacency matrix A , whose (i, j) element is 1 if a link between i and j exists and 0 otherwise. This representation is particularly useful in the context of random walks. The transition matrix W , whose element W_{ij} is the probability of a transition from j to i during a random walk, is indeed simply obtained by dividing the adjacency matrix by the degree, such that $W_{ij} = \frac{A_{ij}}{k_j}$.

A path of length ℓ can be written down as the sequence of the nodes on which it passes, $\mathbf{i}_\ell = (i_0, i_1, \dots, i_\ell)$. By calling ξ_i the quenched random variable associated

with the i^{th} node of the graph, the variables of interest are empirical path-averages of the form $\hat{\phi}_\ell = \frac{1}{\ell} \sum_{t=1}^{\ell} \xi_{i_t}$.

In the paper in appendix B some simple cases are considered since the features one would like to look at already appear in these situations. The networks used are extracted from the Erdős–Rényi ensemble with fixed mean connectivity and the variable sitting on each node are just the degree of the node itself normalised by the average degree, such that

$$\xi_i = f(k_i) = \frac{k_i}{\langle k \rangle} \quad (4.1)$$

The formalism used is however very general, and possible extensions to other interesting cases, both as regards the value assigned to each variable and the topology of the graph, will be discussed. In what follows, the long path limit $\ell \rightarrow \infty$ for which the large deviation theory holds has to be understood whenever the value of ℓ is not specified.

4.2.3 The large deviations

Following the large deviation framework introduced in chapter 1, for studying in a convenient way the rare events is useful to define the cumulant generating function. With the previous definitions, and by calling $P(\mathbf{i}_\ell)$ the probability of observing a given path \mathbf{i}_ℓ , this latter is defined in this case as:

$$\psi(s) = \lim_{\ell \rightarrow \infty} \frac{1}{\ell} \log \sum_{\mathbf{i}_\ell} \left[P(\mathbf{i}_\ell) e^{s \sum_{t=1}^{\ell} \xi_{i_t}} \right] \quad (4.2)$$

The parameter s , playing the role of an inverse temperature, is very important. If $s = 0$ the probability distribution is undeformed: by evaluating the derivative of the cumulant in this point, the typical behaviour is obtained as

$$\bar{\phi}_\ell \equiv \left\langle \frac{1}{\ell} \sum_{t=1}^{\ell} \xi_{i_t} \right\rangle = \psi'_\ell(s)|_{s=0}$$

On the other hand, by using values of s larger than 0, the random walks are biased towards the paths characterised by a large value of ϕ . If the variable on each node is proportional to its degree, this allows us to study the walks touching nodes with degrees atypically large. The larger the value of s , the largest weight is given to the path with largest ϕ , up to eventually assign a non-zero weight only to the path with maximal ϕ for $s \rightarrow \infty$. The same holds for negative values of s , for which a larger weight is assigned to the paths with a smaller value of ϕ .

By tuning the value of s one can get the *rate function* containing the information about the fluctuations around the typical behaviour, as by formally writing $P(\phi) \sim e^{-\ell I(\phi)}$ one has $I(\phi) = \sup_s \{s\phi - \psi(s)\}$ [29].

Just to recap, it turns hence out that a simpler alternative to studying rare events according to the transition probability W is to study the typical events in the deformed ensemble $W_{ij}(s) = e^{s\xi_i} W_{ij}$. By tuning s , different levels of extreme events may be stressed. This deformed matrix enables also to a spectral-based approximation of the cumulant generating function, as by calling $\lambda_1(s)$ the leading eigenvalue of $W(s)$, one has $\psi(s) \sim \log \lambda_1(s)$.

4.3 Main results

4.3.1 A degree-based approximation

As discussed, in order to get the leading behaviour of ψ it is enough to calculate the largest eigenvalue of the deformed transition probability matrix $W(s)$. This is computationally important for large system sizes N , as one does not need the complete diagonalisation of the matrix: some procedures enable indeed to efficiently evaluate just the largest (or a certain number of the largest) eigenvalues [131]. For large enough N , nevertheless, even these optimised techniques are not enough to obtain the result in a reasonable amount of time. If the average degree of the graph is sufficiently large and the degree distribution is not too heterogeneous, however, one can use a simplifying degree-based approximation. This latter consists in assuming the i^{th} component of the leading eigenvector of W to depend just on the degree of the node. In this way, a sum over all the nodes is translated into a sum over all the possible values of the degree, computationally much easier if the set of such possible values is small enough; this is the reason behind the request of having a not too heterogeneous degree distribution. On the other hand, the average degree has to be large enough for this approximation, that makes use of the law of large numbers, to give sensible results.

The results of this approximation have been verified on an Erdős–Rényi graph with $N = 1000$ and $\langle k \rangle = 30$, where both the previous requests are known to hold: as can be seen in figure 4.1, the exact and the approximated curves are in very good accordance averaging over one thousand instances both for $\Psi(s)$ and for $I(\phi)$. Another type of approximation possibly allowing to study large systems in which the degree-based assumptions are not allowed will be discussed in the final section of this chapter.




Figure 4.1: Cumulant generating function $\Psi(s)$ for a $N = 1000$ Erdős–Rényi networks with $\langle k \rangle = 30$ and $f(k_i) = \frac{k_i}{\langle k \rangle}$, comparing the average results obtained by using the large-degree approximation and by running explicit numerical simulations on 1000 samples. In the inset, the same comparison is reported as regards the rate functions.

4.3.2 Localisation transition

Before discussing the localisation transitions found for this problem, one needs to precisely define what is meant by localisation. If one thinks of an N -component vector, a perfectly localised vector will have non-zero component on just one position (and this component will be equal to 1 after a normalisation). At the other extreme, a perfectly delocalised vector will be homogeneously spread among all its components, each of them taking value $\frac{1}{N}$. This idea is condensed in the *inverse participation ratio* of the vector \mathbf{v} , defined as

$$\text{IPR}[\mathbf{v}] = \frac{\sum_i v_i^4}{[\sum_i v_i^2]^2} \quad (4.3)$$

This ratio is by definition between 0 and 1, and the highest it is the most localised the vector is. In particular, for increasing system size, this quantity will scale as $\frac{1}{N}$ if \mathbf{v} is delocalised and as $\mathcal{O}(1)$ if it is not.

It is known [62] that the eigenvector associated to the largest eigenvalue (also called leading or principal eigenvector in what follows for brevity) of the transition matrix represents the equilibrium distribution of the unbiased random walk. This latter is generally delocalised, as it is free to explore the entire graph, and for $s = 0$ the component of the leading eigenvector associated to a node is just proportional to its degree. As the transition matrix is deformed by setting s to non-zero values, a larger and larger weight is assigned to specific paths, and for $|s|$ large enough the leading




Figure 4.2: Visualisation of the localisation transition occurring for large and negative values of s on a $\langle k \rangle = 3$ Erdős–Rényi graph. The nodes whose component in the leading eigenvector is at least 0.1 times the largest one are highlighted. The three figures refer, going from the left to the right, to $s = -0.42$, $s = -0.5$ and $s = -50$, values in correspondence of which one can respectively see a delocalisation of the leading eigenvector across the network, a localisation on a sparsely connected part of it and, finally, a localisation on just a single leaf node (corresponding to the values of s for which the IPR goes to 1).

eigenvector *may* localise on a portion of the graph. By considering the function defined above assigning to each node a value proportional to its degree, for instance, it may happen that for large enough s only the paths touching nodes with very high degrees are given a non-negligible weight; as a consequence, the most connected nodes will have a component much larger than the other ones in the equilibrium distribution represented by the component of the principal eigenvector.

This intuition is supported by the figure 4.2, in which the nodes whose component of the leading eigenvector is at least 0.1 times the largest one have been coloured in red. The figures refer to negative s , for which one expects the localisation to occur on nodes with low degree. Indeed, for s small enough in absolute value the leading eigenvector is still quite delocalised, with many components comparable to the largest one. As s becomes large and negative, however, the random walks localise on a subgraph characterised by atypically small degree, and for a further decreasing of s one component is dramatically larger than all the other ones, with a localisation becoming sharper and sharper. Analogous results can be obtained for positive s ; in this case, for stronger and stronger deformation of the transition matrix W a stronger and stronger localisation on the *core* (i.e. on the most connected part) of the graph and, eventually, on the node with highest degree is observed.

In order to have a more quantitative description, the average IPR over a large number of graph realisations is plotted as a function of the deformation parameter s . Such curves are reported for several system sizes N so to be able and extrapolate the asymptotic behaviour in the large system size limit $N \rightarrow \infty$.

As can be seen in the left panel of figure 4.3, in the $\xi_i \propto k_i$ case a localisation transition occurs both for positive and for negative values of s . For negative s , in correspondence of $s \sim -4$ the IPR value is around 0.5 for every system size: for such a deformation, the paths having non-negligible weight are the ones exploring chains of nodes, in which the random walker can land in one out of two nodes from any position. For even stronger deformations a further increasing IPR is observed. In the inset, the qualitatively different way by which the IPR scales for increasing N in the localised and in the delocalised regime is reported.

This double localisation is not completely trivial, as by using different quenched variables on the nodes different behaviours are observed. A case potentially very interesting also for applications is the one in which each node i take a value 1 if its degree is larger than the average degree of the graph, and 0 otherwise:

$$\xi_i = f(k_i) = \Theta(k_i - \langle k \rangle) \quad (4.4)$$

The interest of this (and similar) definitions is that the path averages $\bar{\phi}_\ell$ one obtains count the fraction of touched nodes with a degree atypically large. Of course also the symmetric version can be used, where paths on lower-than-average nodes assume a larger statistical weight. As can be seen in the right panel of figure 4.3, the localisation transition takes place just for negative values of s and the potential localisation on the core of the graph for positive s is not observed; in the inset, it is shown as before the region in which the localisation occurs, the average IPR switching from a $\frac{1}{N}$ to a $\mathcal{O}(1)$ scaling for increasing system sizes.

4.3.3 Mode-switching transition

The Gärtner-Ellis theorem discussed in chapter 1 can be used to obtain the real rate function only if this latter is convex, as the Legendre-Fenchel transforms used therein yields functions that are necessarily convex. The rate functions obtained by using this theorem when the differentiability hypothesis holds are necessarily strictly convex (i.e. convex and with no linear parts).

If this is not the case, and the cumulant generating function shows one or more non-differentiable points, the rate function one gets by applying the inverse transform will contain a linear region: this latter will not be interpretable as a part of the rate




Figure 4.3: Average IPR curves on 1000 sample of $N = 10^3$ Erdős–Rényi graphs with $\langle k \rangle = 6$. On the left, the quenched random variable associated to each node is proportional to its degree, $f(k_i) = \frac{k_i}{\langle k \rangle}$; a localisation transition is spotted for large enough both negative and positive values of the deformation parameter s . On the right panel the same is reported in the case where the random variable is $f(k_i) = \Theta(k_i - \langle k \rangle)$, for which a single localisation transition for negative values of s is observed. In the insets the curves of the IPR as a function of $1/N$ in correspondence of different s are reported in double logarithmic scales, so to let the difference between the two regimes be clearer.

function. As can be seen in figure 4.4, indeed, there is not any more a univocal relation between the rate function and the cumulant generating function, the cumulant generating function $\lambda(k)$ being the same for all the non-convex rate functions $I(\phi)$ sharing the same complex hull.

As can be seen in the left panel of figure 4.5, the emergence of a linear region of the rate function, becoming more and more pronounced for increasing N , has been spotted: as already said, the linear part is just the convex envelop of the actual rate function, that cannot be explicitly determined under these conditions. As expected, this linear region corresponds to a non-differentiable point of the cumulant generating function that can be seen in the inset of the same figure.

This phenomenon depends on a non-analyticity in the largest eigenvalue λ_1 , and can be better understood by looking also at the second one λ_2 . In particular, one can define a *correlation length* as follows:

$$\tau(s) \equiv \frac{1}{\log \frac{\lambda_1(s)}{\lambda_2(s)}} \quad (4.5)$$

These curves as a function of the deformation parameter are reported in the right panel of figure 4.5 for several system sizes, and it is clear how two points exist where this correlation length asymptotically diverges as $N \rightarrow \infty$. When looking




Figure 4.4: When a Legendre-Fenchel transform (represented by an arrow in the figure) is applied on a non-convex rate function $I(s)$ as the one on the right, a non-differentiable cumulant generating function $\lambda(k)$ is obtained (see central panel). This latter, however, is the same for all non-convex rate functions having the same convex hull as $I(s)$. The only information one can get from a $\lambda(k)$ such as the one reported in the central panel is hence the shape of the convex hull of the rate function, but not its precise form (see right panel). The figure is taken from [29].

at Erdős–Rényi graphs with the same average connectivity, the discontinuity in the derivative of the cumulant has been verified to correspond to the negative value of s for which $\tau(s)$ diverges. The presence of another, positive *critical* value of s is a clue that probably another linear regime in the rate function has to be looked for, even if it is not as clearly recognisable as the other one. Terms such as “correlation length” and “critical” value of s have not been used by chance: indeed, it has been seen that $\tau(s)|_{s^*}$ scales as $\log(N)$, this latter being the characteristic length of a N -nodes graph. A more precise characterisation of this critical behaviour has yet to be performed (for instance, by determining the critical exponents).

By looking at the definition of $\tau(s)$, one easily sees that its divergence in correspondence of s^* implies a degeneracy (at least) in the two largest eigenvalues $\lambda_1(s^*) = \lambda_2(s^*)$ in the infinite system size limit. The situation induced by this degeneracy is a well-studied situation in quantum mechanics called *level crossing*. The modes corresponding to the first and to the second eigenvalue indeed crosses, and the leading mode (the one corresponding to the largest eigenvalue) is a different one before and after the critical point. This is seen also in very simple two-level quantum systems [132], in which depending on the sign of $E_1 - E_2$ the ground state energy changes, and in particular for $E_1 = E_2$ the two eigenvalues are degenerate.

In finite systems, however, this crossing is impossible for the so-called *level repulsion*. In an intuitive way, a region around the critical point exists where the modes corresponding to λ_1 and to λ_2 are hybridised, so that one observes a so-called *avoided level crossing*. Nevertheless, before and after this mixed zone the leading modes still correspond to different eigenvalues; such a feature can be obtained in quantum system




Figure 4.5: On the left panel: rate function $I(\phi)$ for Erdős–Rényi graphs with $\langle k \rangle = 3$ and $f(k_i) = \frac{k_i}{c}$ for system sizes ranging from $N = 100$ to $N = 6400$. In the inset, the behaviour of $\Psi(s)$ in the vicinity of the non-differentiable point is shown. For the largest system size are plotted as well a linear fit of the convex envelope of the left branch of $I(\phi)$ and a quadratic fit of its the right branch. On the right panel the divergence in correspondence of two s values of the correlation length τ defined in eq. 4.5 is shown for $\langle k \rangle = 6$ Erdős–Rényi graphs of sizes ranging from $N = 200$ to $N = 3200$.

by inserting an off-diagonal perturbation, avoiding a degeneracy in correspondence of $E_1 = E_2$ [133].

We have been able to start and confirm this *mode-switching* transition in s^* as follows. The IPR of the eigenvectors v_1 and v_2 , respectively associated to λ_1 and to λ_2 , can be separately studied. If one looks at them before and after the hybridised zone, one can see that the IPR of v_1 before the transition is more similar to the IPR of v_2 after it, and vice versa. By looking at these kind of observables, hence, one can see that there has really been a switching between the leading mode and the second one in correspondence of s^* , and in a non-rigorous way one could say that the second eigenvector after s^* in some sense “is” the proper continuation of the first before that point.

The discussed level crossing implies that the mode which realises the typical rare event changes at the critical s^* to another mode, and the original one is no longer representative of the dominant rare events. Because of this, one has good reasons to call this a *dynamical phase transition* in the rare event ensemble, this being probably the most interesting result obtained so far as regards this project.

4.4 Summary and future perspectives

4.4.1 Summary

In this project several promising results have been obtained. In particular, the problem of understanding the statistics of extreme events in random walks on graphs was addressed by looking at the spectral properties of a transition matrix appropriately deformed. By doing that and studying the scaling with N of the leading eigenvector IPR, it has been found that localisation transitions may occur, so that the large deviations are realised by modes of the biased transition matrix living on small subsets of the nodes. Finally, a mode-switching transition has been detected, implying that by changing the large deviation scale (i.e. by tuning the deformation parameter s) the rare events may be realised in terms of different eigenvectors. In the rest of this chapters some possible extensions of this work (some of which already started, or quite straightforwardly implementable) are going to be presented.

4.4.2 Better understanding of the critical behaviour

As discussed, the mode-switching transition could be thought of as an actual critical behaviour. In order to do that, however, a more precise scaling with the system size has to be performed. In particular, one would like to reach larger system sizes so to have a more reliable $N \rightarrow \infty$ extrapolation. Doing that is computationally difficult even by studying just the leading eigenvector (i.e. by calculating only the largest eigenvalue instead of diagonalising the complete matrix). One could however try some approximation schemas such as the cavity-based method proposed in [134]. Having such a precise extrapolation for larger sizes, one could try and estimate quantities of interest such as the critical exponents of the phase transitions. Moreover, if one looks at the right panel of figure 4.5, two questions have still to be answered. Firstly, for increasing N the correlation length τ diverges, but also the position of the peak slightly moves towards smaller and smaller values of s in modulus: also this position should be extrapolated in the infinite size limit, so to obtain the value of s^* where the real phase transition occurs when $N \rightarrow \infty$. Secondly, another critical points is observed for positive s , and possibly another mode-switching is occurring: the linear regime corresponding to this second transition has not been spotted yet in the rate function. Finally, one would like to associate these mode-switching transitions to some observables associated to the modes. By doing that, one would have a clearer representation of how the rare events are realised in different ways before and after

these transitions; one possible idea is to use as such a variable the IPR, but even better observables to look at could be probably thought of.

4.4.3 Different functions

As discussed at the beginning of the chapter, the analysis proposed can be straightforwardly applied to any function assigning a value to each node of the graph. Both this chapter and the paper reprinted in B focused almost exclusively on such variable being for any node proportional to its degree.

Even in the context of variables deterministically assigned depending on the nodes of the network, other interesting cases have already been addressed, even if a comprehensive set of results has still to be produced. In particular, the binary function 4.4 taking value 1 on the nodes with degree larger than the average is of clear interest for biasing just towards paths connecting nodes with “atypically” large or small degrees, but without caring too much about the precise value of the degree itself. When studying different topologies, one could substitute the average with the median, so to obtain significant results also on graphs with very heterogeneous degree distributions such as the scale-free ones.

On the other extreme case, some results have been obtained by assigning the quenched variables at random, independently on the degree of the node they refer to. If the variable assigned in this way is binary, this situation is of interest in the cybersecurity context. One can imagine the nodes taking value 1 as *random nuggets* hidden on the graph, and one is trying to understand the statistics of rare events able to reach a very large number of such nuggets. In this case, the $\Psi'(s)$ appears to be in extremely good accordance with a straight line, and the parametric reconstruction of the rate function is hence very difficult to obtain.

In order to get a better understanding of this extreme case, one can firstly examine an intermediate situation where the variable assigned to a node is correlated to its degree. One could in particular assign the seeds by using a sigmoid function, smooth version of the step function 4.4 that is of great interest for biological applications [52].

$$\xi_i(k_i) = \begin{cases} 1 & \text{with probability } \frac{1}{2} \left(1 + \tanh(\beta(k_i - c)) \right) \\ 0 & \text{with probability } \frac{1}{2} \left(1 - \tanh(\beta(k_i - c)) \right) \end{cases} \quad (4.6)$$

β is in this case a control parameter playing the role of an inverse temperature in thermodynamic systems; its function is to allow the tuning of the level of correlation between the degrees and the values the function takes on the sites. In particular, one

can see that the two cases studied so far are just special cases of this formulation. In the “high temperature regime”, the variable is not correlated with the degree any more and one recovers the random nuggets case:

$$\beta \rightarrow 0 \Rightarrow \text{P}(\xi_i(k_i) = 1) = \text{P}(\xi_i(k_i) = 0) = \frac{1}{2}, \quad \forall k_i \quad (4.7)$$

On the other hand, the sigmoid function becomes steeper and steeper as β increases. Eventually, as $\beta \rightarrow \infty$ the step function deterministically assigning a value 1 to the nodes with degree larger than the average and 0 to the others is recovered:

$$\beta \rightarrow \infty \Rightarrow \xi_i(k_i) = \begin{cases} 1 & \text{if } k_i > \langle k \rangle \\ 0 & \text{if } k_i < \langle k \rangle \\ \frac{1}{2} & \text{if } k_i = \langle k \rangle \end{cases} \quad (4.8)$$

4.4.4 Functions taking value on the links

In all the cases discussed quenched value, fixed once for all before starting to explore the network, was assigned to each node. A very interesting extension, reachable by just some minor technical changes with our approach, would be on the other hand to study functions taking value on the links of the graph. As before, the observables of interest would be path-averages dependent on the links covered during the random walks.

The possibility of simply enough writing down this extension (and *a fortiori* the results it will lead to) has not to be taken for granted in general. The formulation and the solution of even very similar problems can be indeed very different depending on whether they concern the nodes or the edges of a network. In the context of graph exploration, for instance, determining whether a Eulerian cycle can be spotted on a network is a relatively easy task, whereas determining if a Hamiltonian cycle is possible is computationally hard.

Similarity and differences between the two cases would be interesting to look at on their own; moreover, by looking at observables taking values on the edges between nodes one can approach potential applications such as a better comprehension of the extreme event statistics in information or traffic flows.

4.4.5 Different topologies

All the results reported refers to Erdős–Rényi random graphs. However, the approach proposed can be directly applied to any kind of networks, and some preliminary results have already been obtained for instance on scale-free graphs, or on random

regular graphs. By looking at different topologies, one could try and imagine a more complete theoretical framework able to explain the critical behaviours observed, the characteristics of such behaviours that are observed regardless of the type of graph the random walk is performed onto and the ones that depends on some specific features, etcetera. The scale-free graphs are expected to produce particularly interesting results, as they show a very broad degree distribution; in some sense it is therefore reasonable to forecast that the rare events on such graphs will be qualitatively very different from the ones occurring on Erdős–Rényi networks. Some preliminary results confirming this idea have already been obtained.

4.4.6 Improvement of clustering techniques

As discussed in chapter 2, one of the problems mostly studied in network theory is the clustering of nodes according to some characteristics. In this work, by simply looking at the spectral properties of biased transition matrices we have been able to somewhat perform this task. As discussed in the localisation transition, indeed, the leading eigenvector lives under some conditions on a subset of the nodes (for instance, on the ones with large degree). Also thinking about the richness of both well-established [79] and more advanced [83] clustering techniques making use of spectral properties of the network, one could suspect that some improvements to these methods could be obtained by therein inserting in some way the insights obtained by our large deviation approach. This is however a very rough idea, and a much deeper understanding will be needed in order to see whether such approach could actually be doable and fruitful.

Chapter 5

Inferring the graph

5.1 Motivation, applications and connection to other problems

5.1.1 Motivation

In both the chapters 3 and 4, respectively devoted to the spreading dynamics and to rare events in network exploration, a complete knowledge of the graph on which the phenomenon under study is occurring has been taken for granted. This assumption is not at all always guaranteed, and in many interesting cases one wants to estimate some quantities of a process even without perfectly knowing its structure. A possible approach to this problem is to try and infer the graph from the partial knowledge one has at his disposal, and then to suppose the process to take place on this estimated network. The estimation may in general be refined over time for two different orders of reasons. Firstly, one can simply have access to more and more information on the network structure as more measurements are performed on it. It is also possible, secondly, that by simulating some process taking place on the real, unknown graph on the reconstructed network one obtains results disagreeing with the ones observed on the former; such a situation can be considered as a hint showing us that the previous inference was not accurate enough for obtaining reliable estimates of the process under exam. If this is true, an iterative improvement of the model so to approach more and more the features observed in the real system can be tried.

The object of this chapter is to discuss the task of inferring a graph from just a partial knowledge about it. It is expected that, in general, the more information one has, the better the reconstruction will be. A part of the work will consist in evaluating how the quality of the reconstruction depends on the amount of available information under different conditions. The most interesting part of the work, however, will make

use of the hypothesis that one is able to choose how to enrich the set of information currently available (i.e. what to measure on the system): in this case, one would like to define a strategy for choosing the information to add in an optimal way. For instance, one may look at the heuristics enabling to understand the graph structure with the smallest possible number of measurements.

The problem addressed in this project is of very general interest. In all the situations that can be modelled as networks, the hypothesis of perfectly knowing the graph on which some dynamics takes place is very strong. Developing methodologies for efficiently and reliably producing estimates of how elements are linked among each other is a high-level problem, and any advancement in this direction may be exploited in a variety of fields (as the ones discussed in chapter 2) and for addressing a variety of problems (the ones discussed in chapter 3 and 4 being just two examples among many possible others).

5.1.2 The matrix completion problem

As discussed in section , a graph can be comprehensively represented in a condensed form by making use of an adjacency matrix; all the information contained in figure 5.1, for instance, can be resumed in the following adjacency matrix A^{true} :

$$A^{true} = \begin{matrix} & \begin{matrix} A & B & C & D & E \end{matrix} \\ \begin{matrix} A \\ B \\ C \\ D \\ E \end{matrix} & \begin{pmatrix} 0 & 5 & 0 & 0 & 3 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 4 & 0 & 1 & 3 & 0 \\ 0 & 3 & 0 & 0 & 0 \end{pmatrix} \end{matrix} \quad (5.1)$$

The graph is completely described by the value of its $N^2 = 25$ adjacency matrix elements, some of which being equal to zero: because of this mapping, in the following of this section the terminology referring to matrices will be preferred to the one referring to graphs. The *matrix completion problem* consists in trying to optimally estimate the missing entries of a given matrix having just some partial information [135]. For instance, if one has already measured the value of 10 elements of A^{true} , the current knowledge he has about the system is something as the following A^{meas} :




Figure 5.1: Graphical representation of the directed graph whose adjacency matrix is the one of equation 5.1. The width of each link represents its weight.

$$A^{meas} = \begin{matrix} & A & B & C & D & E \\ \begin{matrix} A \\ B \\ C \\ D \\ E \end{matrix} & \begin{pmatrix} ? & ? & 0 & ? & 3 \\ ? & 0 & ? & 0 & ? \\ ? & ? & ? & 0 & 0 \\ 4 & ? & ? & ? & ? \\ ? & 3 & 0 & ? & 0 \end{pmatrix} \end{matrix}$$

In order to try a reconstruction of the missing elements, some hypothesis about the system have clearly to be formulated. In the most general situation, indeed, the value of each matrix element is completely independent from all the others, and nothing can be said about it before actually performing a measurement. [136]

A mathematical ansatz typically used is to assume the matrix that has to be reconstructed to be low-rank; from a general perspective, this idea can be understood in terms of the discussion of chapter 1 about minimal models. A natural question when dealing with high-dimensional data characterised by a large number of attributes, is whether they could be generated by a much simpler model (i.e. by a model in which the number of free parameters is much smaller than the total number of features of the system) [137]. The reasonableness of this hypothesis will be further discussed in the following section when introducing the collaborative filtering, specific example in which the matrix completion problem is of interest.

Even having to deal with a matrix for which the low-rank approximation is not justified, the discussed ideas are still of interest. How to find an as accurate as possible low-rank approximation for a generic matrix to obtain a more synthetic description of the system is a very classical problem [138]. The methods for tackling it, among which one can at least remember the *principal components analysis*, are especially needed for data-rich fields, such as for instance the genomics and, more generally, the bioinformatics [139]. In these cases, indeed, dealing with the true data matrices is computationally unaffordable. Even finding such approximations in an efficient way has become a topic of interest per se [137].

The matrix completion problem has straightforward applications in information theory, and in particular in signal processing. In image reconstructions, one can imagine situations where many of the pixels composing an image are corrupted, and they have to be reconstructed using the information one has. In this sense, the advancements in matrix reconstruction problems can be connected to the compressed sensing discussed in section 1.6.6. A prior knowledge of the characteristics of signals enabling to perfectly reconstruct it having just a partial set of measurements can directly contribute also in designing optimal procedures on how to take such measurements.

The main result in this framework is the determination of thresholds telling how many measurements are needed for the exact reconstruction of a matrix with a given rank [140]. Also of interest is the design of algorithms remaining tractable for large system sizes and able to reconstruct the matrices with a number of measurements approaching the information theoretical bound analytically determined.

Another aspect deeply studied in this context is the robustness of the reconstructed matrix in presence of noisy entries [141, 142]. Especially for real-world applications, indeed, any measurement taken on a system will unavoidably include some fluctuations. A procedure able to produce reasonably accurate estimates of the elements even in the presence of noise is hence preferable to another whose results strongly depend on the precise value of the measured elements. Spectral methods able to reliably estimate the rank of the matrix from few available entries have also been proposed [143].

5.1.3 The collaborative filtering

A task strictly connected to the matrix completion problem is the *collaborative filtering*. This idea was originally introduced as the key component of an experimental mail system able to select the most interesting messages for a given user in a stream of incoming electronic documents [144]. It assumed afterwards a much more general meaning, and any system aiming at automatically making predictions about the interests of a user by making use of similar information regarding other users is nowadays referred to as a collaborative filtering, even if no explicit “collaboration” among people comes into play.

The originally proposed active action of suggesting items to other users has progressively been substituted by an assumption that can be expressed in its simplest form as follows: two users having behaved (for instance, having bought, watched or listened) in a similar way in the past, or having rated similarly a set of items, are likely to agree also on the actions they will take or on the ratings they will give in the future [140]. One of the most celebrated examples of these ideas is the algorithm implemented by NetFlix in order to predict users’ film preferences just by using the few ratings each user typically gives [136]. The accuracy of such predictions is of the greatest importance for this and similar companies for augmenting their sales, increasing their incomings from advertisements and allowing the design of targeted offers. As a proof of this, it is enough to recall the so-called *NetFlix challenge* during which a one million dollars reward was assigned to the first contestant able to

	Shrek	Snow White	Spider-man	Super-man
Alice	Like	Like		Dislike
Bob		Like	Dislike	Like
Chris		Dislike	Like	
Tony	Like		Dislike	?

Figure 5.2: Example taken from [145] of a rating matrix whose reconstruction is of interest in the context of the recommendation systems. As one can see, in this case binary ratings have been given to some of the films by any user. In general, they can also be integer, for instance in a range between one and five “stars” or even *implicit* and based, for instance, on the pages previously visited and on the items already bought.

produce an algorithm improving by 10 % the NetFlix’s own system to predict film ratings given by users [62].

Looking at figure 5.2, it is clear how strictly this application is related to the matrix completion problem, regarding in its most general form the recovery of a data matrix from incomplete or corrupted information. In this context, moreover, the mathematical assumption of dealing with low-rank matrices assumes a clear meaning: it corresponds to the reasonable hypothesis of the taste of a user being describable at least approximately by a small enough number of factors. If this was not the case, the whole task described in this section would be doomed to failure, as the ratings will be more or less independent the ones on the others. The precise number of factors needed for describing the system has to be chosen so to build the simplest model able to catch the complexity of the data (i.e. it has to avoid both overfitting and underfitting).

When dealing with recommendations, it turns out that a central problem is how to choose between two different types of suggestions. On the one hand, one may suggest an item very close to the preferences already expressed by the user. Alternatively, one could make a riskier recommendation, choosing something further away from the previously chosen items; this way is potentially more profitable as the user could discover a new set of items he is interested into and start purchasing in this direction. This point is referred to as *exploration-exploitation trade-off* [146]; apart from this interesting application to recommendation systems, it is a long standing problem in the field of control in uncertain environments, where one has to choose between doing exploratory actions improving the knowledge on the system, and exploiting the information already collected in an optimal way [147].

5.2 The inference of a network of interactions from a partial knowledge of the correlations

5.2.1 Definition of the problem

The rest of this chapter will be devoted to the discussion of a problem on which a paper is currently under redaction; for this reason, the discussion will be a bit longer and technically developed compared to the ones given in the previous chapters. In order to better follow the main ideas of the project, in particular, the discussion on some further details will be postponed to the appendices 5.10. Leaving aside for the time being the technical details, one is interested in studying a set of variables that are coupled and whose values are measured multiple times. Because of the interactions acting in the system, the result of such measurements will not be independent, and correlations will start to appear among the values one gets during the different observations performed on the system. The aim of this project supervised by professor Rémi Monasson is to try and optimally infer the interactions acting on the system by knowing just a subset of the correlations between variables during the observations.

Two different orders of problems can be thought of in this framework. On the one hand, one is interested in the “static” problem: knowing a fixed subset of correlations among the values taken by the variables in the different measurements, one would like to obtain the optimal estimate of the unknown interactions. Even if this problem is connected to the matrix completion one, as the available information is fixed once and for all, many important differences between the two are anyway present: the main is that in this case the matrix one wants to infer and the one on which one has some partial information live in two separate spaces, and no direct measurement can be made on the interactions one wants to recover. A very interesting point is to analyse how the quality of the inference improves as more and more correlations among variables are known, as this could give some useful hints on the optimal number of correlations to measure before trying the inference of the interactions. One could for instance imagine a situation of “retarded learning” where up to a certain number of measurements the inference is really bad, and only after that transient the couplings can be reasonably predicted; on the other hand, one can find out that a number of measures are enough to obtain good predictions, and making any following measurement is not so useful any more as the quality of the prediction somehow saturates.

The most innovative part of this work is however connected to a “dynamic” version of the problem, where one has the possibility of choosing the couples of elements

about which one wants to know the correlation during different measures on the system. This is particularly important on sparse networks where the number of significant couplings between variables is small; in this case, by randomly choosing which measure to perform one is almost sure to select a couple of sites not linked, and hence not to improve that much the quality of the inference. We will discuss some possible heuristics aiming at augmenting the probability of performing “good” measurements (i.e. improving the most the quality of the inference) for any level of knowledge about the correlations.

5.2.2 Motivation, applications and previous work

The problem introduced is first of all of interest from a theoretical point of view, as it is at the boundary of several very well studied topics. Among them the graph theory, as the set of couplings constitutes a network of interactions, the non-trivial relationship between correlations and direct couplings, the optimal reconstruction of missing data. Nevertheless, it seems like this problem has not been studied in detail so far; the work the most similar to the one we discuss is probably [148] where the authors face the problem of completing a correlation matrix. The methodologies they propose, though, are quite near to the ones of the matrix completion works above described: the matrix to complete is supposed low-rank, and thresholds on the number of requested measures to reach an exact reconstruction are derived. There are at least two major differences with this project. Firstly, the matrix whose elements are measured is in their case the same one wants to complete; secondly, the dynamic part of the problem aiming at optimally choosing the measures to take is completely missing.

The problem as it has been defined is not just a mathematical exercise: in many situations, indeed, one has direct access only to correlations between variables whereas the quantities of interest are the direct couplings among them. One of the most natural examples of such systems is the finance, where one can see the correlations between the value of shares during time and is interested in understanding relationships among companies and sectors. Others can be found in biology: for instance, one would like to reliably infer the gene regulatory network controlling the gene expression levels, but what can be typically measured is the correlation between these latter. In this sense, this project is again strictly connected also to the discussion of the previous chapters on inverse problems.

5.3 The model

5.3.1 Statistical context of the project

Even without entering into details about the different possible approaches to statistics, a short reminder of the main features of Bayesian methods is worthy in order to better understand the following discussion. Such methods are particularly useful in the context of the inverse problems generically discussed in section 1.2, as they allow to express the *posterior probability* that a certain model θ is the one that describes a system having generated the data y (inverse problem) as a function of the *sampling distribution* telling the probability of such model to generate the data (direct problem) and on a *prior probability* we assign to the model itself, independently from the observations [149]. Such a probabilistic treatment of the problem allows to deal both with its possible underdetermination (if the number of equations is smaller than the number of unknowns) and with the presence of noise letting the relationship between the model θ and the data y writable in the form $y = G(\theta) + \eta$. θ, y and η are treated as random variables, and one considers as a “solution” of the inverse problem the determination of the probability of having θ given y .

The arbitrariness of the prior has been widely debated since the proposition of this approach. Even without entering into this discussion, one can just remark that, in a sense, all statistical methods are subjective as they rely on mathematical idealisation of the system examined [149]. Intuitively, this prior can be thought to be a quantity allowing one to specify in advance which kind of solutions one believes to be more likely, so to assign different weights to the multiple solutions that can be found, all of them able to explain the observed data [150].

The data y being considered as fixed, the three introduced quantities are related (apart from a normalisation) through the *Bayes' rule* as follows:

$$P(\theta|y) \propto P(\theta)P(y|\theta) \quad (5.2)$$

the $|$ symbol indicating as usual the probability of an event conditioned on another one. The data y affect the posterior only through $P(y|\theta)$; this latter, for fixed y , can be seen as a function of θ and it gives the *likelihood function* of the models.

Let us now make use of this framework to more precisely state the problem faced in this work. Dealing with a N -spin system, the role of the model θ is played by the N by N matrix J whose $(i, j)^{th}$ element specifies the coupling between sites i and j . A single measurement x on the system will consist in a vector of N elements (also called *spins* in the following for brevity) corresponding to the values taken by the spins in

presence of the interaction matrix J . The connection between the two is supposed to follow a Gaussian law, such that the probability of obtaining a measurement array $x = (x_1, \dots, x_N)^T$ can be written down as follows:

$$P(x|J) \propto e^{-\frac{1}{2}x^T J x} \quad (5.3)$$

A single measurement on such a system is expected not to be sufficiently informative in order to determine the couplings acting among the variables; one can nevertheless perform multiple measurements on it. The results obtained in this way can be gathered in a M by N matrix X , where each of the M rows is a different measurement on the system, i.e. an independently drawn N -dimensional array extracted from the distribution 5.3. As the interaction matrix J is the same for all measurements, the rows of X will be correlated. One can hence define an N by N empirical correlation matrix C in which each position (i, j) corresponds to the value of the empirical average among all the rows of X of the product of the i^{th} and of the j^{th} measurements on the system; by indexing with m the rows of X , one has $C_{ij} = \langle x_i^m x_j^m \rangle_m$. This matrix will play the role of the data θ in the previous general formulation 5.2 and will follow in this Gaussian framework a *Wishart distribution* [151], which is a generalisation of the χ_k^2 distribution describing the sum of the squares of k independently drawn standard normal variables. The Wishart distribution arises very naturally in the context of multidimensional Bayesian analysis [152], as it is the one followed by the *scatter matrix* $X^T X$, this latter being the maximum likelihood estimate for the covariance matrix of the model. As anticipated, by setting the dimension of the system p equal to 1 one recover the well-known χ^2 distribution.

In what follows, not all the elements of the correlation matrix will be known. In order to let this point be clear, the correlation matrix will be referred to as C_{true} , the set of its elements already having been measured will be called \tilde{C} . The ones not yet measured will be called on the other hand C_{\perp} , and by C we will refer to a matrix where the elements belonging to \tilde{C} correspond to their true value, and the others, still unknown, have been estimated in some way.

5.3.2 Prior, posterior and optimal estimate of the model

By explicitly writing down the equations regulating the direct problem in the Gaussian model before discussed, the sample distribution of the correlations C_{true} for

a given interaction matrix J turns out to be the following:

$$P(C_{true}|J) \propto e^{-\frac{M}{2} \text{Tr} [(-J)C_{true}] + \frac{M-N-1}{2} \log \det C_{true} + \frac{M}{2} \log \det(-J)} \equiv e^{-\frac{M}{2} F_{direct}(C_{true}, J)} \quad (5.4)$$

having called M the number of performed measurements in the N -spin system with couplings J examined. For brevity, a ratio between these parameters telling how well sampled the system is is defined as $\alpha \equiv \frac{N+1}{M}$.

The observed data is not in this case the matrix C_{true} , but just some elements of it. By calling \tilde{C} the subset of measured correlation, the sample distribution of interest is hence $P(\tilde{C}|J)$ which can be obtained by integrating eq. 5.4 over all the unknowns components C_{\perp} . This integration is highly non-trivial; however, under the reasonable hypothesis of a large enough number of measurements it can be approximated by using the *saddle-point method* so to keep only the dominant term:

$$P(\tilde{C}|J) = \int dC_{\perp} P[C = (\tilde{C}, C_{\perp})|J] \propto e^{-\frac{M}{2} \min_{C_{\perp}} F_{direct}[C=(\tilde{C}, C_{\perp}), J]} \quad (5.5)$$

In order to address the inverse problem, one has to reverse this relation; to do this, according to the Bayesian theory above introduced, a prior probability distribution over the model J has to be defined. This choice is quite arbitrary; there is however a natural option typically used in absence of strong elements favouring different ones. In the Gaussian world in which this problem lives, it is not surprising to imagine also the parameters to follow a Gaussian distribution around zero. This choice is implemented by using a so-called *weight decay regulariser* [153] penalising the parameters according to their ℓ_2 norm. This can be seen also as a *parameter shrinkage method* as by using this prior the parameters are shrunk towards zero; the risk of overfitting by selecting very large couplings is hence also avoided as the elements of J are forced not to be too large.

If this prior is used, the application of the Bayes' rule leads to the following posterior distribution:

$$P(J|\tilde{C}) \propto e^{-\frac{M}{2} \{ \min_{C_{\perp}} (F_{direct}) + \frac{\gamma}{2} \text{Tr}(J^2) \}} \quad (5.6)$$

γ being the parameter controlling the strength of the imposed regularisation. The model J having generated the observed \tilde{C} with the highest probability can hence be found by minimising with respect to C_{\perp} and to J the quantity:

$$F[C = (\tilde{C}, C_{\perp}), J] = \text{Tr}(JC) - (1 - \alpha) \log \det C - \log \det J + \frac{\gamma}{2} \text{Tr}(J^2) \quad (5.7)$$

The minimisation with respect to J enables to write down the couplings J as a function of the correlations C , as

$$\begin{aligned} \frac{\partial F}{\partial J} = 0 &\Rightarrow C - J^{-1} + \gamma J = 0 \Rightarrow \\ &\Rightarrow J(C) = \frac{1}{2\gamma}(C - \sqrt{C^2 + 4\gamma\mathbb{I}}) \end{aligned} \quad (5.8)$$

Because of this, the F can be more conveniently expressed as a function of C only, so to let the minimisation with respect to C_{\perp} being of easier solution; indeed, the result can be formally written down as:

$$\frac{\partial F}{\partial C_{\perp}} = 0 \Rightarrow J_{\perp} = -(1 - \alpha)(C^{-1})_{\perp} \quad (5.9)$$

What this relation means is that in order to obtain the couplings J_{\perp} associated to the correlations not measured yet, one has at first to complete somehow the matrix C . This completion can be performed in the simplest case as a gradient descent on the unknown elements of C , aiming at obtaining a minimal discrepancy between the two sides of eq. 5.9 J_{\perp} and $-(1 - \alpha)(C^{-1})_{\perp}$. The completed C has hence to be inverted and the corresponding J elements can be obtained after the rescaling given by the constant term $-(1 - \alpha)$.

A particularly interesting case of these general formulas is the limit $\alpha \rightarrow 0$ corresponding to an ideal situation in which an infinite number of measurements has been performed on a finite system: this assumption is equivalent to imagine that the measurements are completely noise-free. The following of this work will mainly deal with this limit, that even if quite unrealistic in real applications enables to obtain useful insights on the problem.

5.4 Choosing the couplings and evaluating the performances

5.4.1 Artificial models

Before starting to explain the details of the cases studied, a methodological remark is needed. As already said, the aim of this project is to infer an unknown set of couplings from correlation measurements. As usual in these cases, however, the heuristics and the algorithms proposed have first of all to be tested on artificial models whose characteristics are known in advance, so to verify whether they actually give

results which are better compared to the ones obtained by already existing methods. The situation has already been discussed in section 2.5.1 as regards the problem of community detection.

In all the cases that will be referred to in the rest of the chapter, hence, the inference procedure will be run knowing already the true answer from the beginning, as this latter has been itself generated by us. In particular, some of the proposed heuristics will be evidently self-referential and not applicable to real situations, as they will make use of the real value of the very same quantities one wants to estimate. These algorithms will be used as *null models* with which the more realistic heuristics are compared so to better understand their behaviour. Once the methods will be proven to reliably work at least under certain conditions on these artificial models, one will be ready to use them in more realistic scenarios where one really wants to improve in the most efficient way the knowledge on instances that are not known a priori.

5.4.2 The geometry of J

It is somehow expected that the results of the procedure will depend on the geometry of the interaction matrix. In order to quantitatively evaluate these differences, the proposed heuristics will be used on several artificial models; each of them will be characterised by a given structure of the *contact map*, meaning by this term the set of couples connected by a coupling different from zero. In other words, the interaction networks will be extracted from different statistical ensembles. The geometries will be chosen so to take into account the features most likely to affect the result of the inference procedure: cases with different degree distributions, with links being chosen at random or according to some overall structure, and with diverse loopiness (i.e. number of paths connecting any couple of nodes) will be considered.

In all the cases, we impose the network induced by the couplings that are different from zero to be connected on the set of nodes of the system. If this is not the case and for instance the J separately spans two components of the graph, then the problem can be without loss of generality split into two smaller ones concerning them separately. The only thing still to be verified as final check will be whether in this case the procedures proposed are able to fast enough see the presence of two disconnected subgraphs.

In order to comprehensively understand the behaviour of the heuristics, quite a wide range of topologies has been studied. At one extreme, one has a completely random *Erdős-Rényi graph*. On the other, one can study completely ordered *square*




Figure 5.3: 3-level hierarchical geometries with $N = 16$ on the left panel and $N = 33$ in the right one. Such contact topologies are composed by a highly connected hub, a set of leaves with degree 1 and an intermediate level of nodes connecting the other two. The nodes have been coloured so to let the three hierarchical levels be more easily visible.

lattices; for these latter, a feature expected to influence the results is their strong “loopiness”.

A different and potentially interesting case is the one in which the degree distribution is homogeneous, but a global structure is lacking; this can be modelled via a *random regular graph*, where each node has a fixed number of neighbours which are randomly chosen among all the other vertices. Of particular interest is the special case where the number of neighbours is fixed to two, as it corresponds up to a permutation of the vertices indexes to a *ring*.

All the geometries above introduced are characterised by a degree distribution that is either a δ function in correspondence of the average value, or is peaked around it. Since this more or less strict homogeneity in the connectivity is expected to be a crucial factor in determining the results, one wants to consider also graphs where nodes with very large and with very low degree coexist (i.e. network with a heavy-tailed degree distribution). In principle, and especially thinking about possible applications, one would like to study scale-free graphs. The heuristics used in this project are however for the time being quite computationally expensive, and one is forced to study small systems: generating them by a preferential attachment procedure originates a geometry not that far away from a *star*, with a central node attached to all the others and some very sparse connections among the remaining vertices. In order to

understand the effect of a more complex level of organisation, it has been chosen to run the algorithms on *hierarchical structures* as the one in figure 5.3. These latter can give precious insights on the way the heuristics work, for instance by telling us whether they are able to first infer a small part of the structure and then to enlarge such local knowledge or, vice versa, they approximatively catch the global structure and then refine the estimates on the actual value of the connections.

5.4.3 Generating synthetic J

From now on, whenever no misunderstanding is possible the non-zero elements of J will be simply referred to as to the couplings of the system. As discussed in 5.4.1 one is interested in testing the performances of the heuristics over already known J with different geometries so to check their effectiveness. In order to do that, one has to first of all find a J with a fixed geometry for which the C_{true} maximising $P(C_{true}|J)$ is a well defined correlation matrix. Both the two matrices entering into this relation have to satisfy some properties. First of all, the J has to be positive defined in order for the Gaussian law expressed in equation 5.3 to be well defined. The correlations between variables need moreover to be in the range between -1 and 1 and with diagonal elements C_{ii} (i.e. the correlation of a spin with itself) always equal to 1 .

As shown above, such correlation matrix has to satisfy the relation $C_{true} = -(1 - \alpha)J^{-1}$. In particular, one is interested in finding the C_{true} that would maximise the previous probability in the case of an infinite precision of the measurements; this is the $\alpha = 0$ case, in which the particularly simple relation $C_{true} = -J^{-1}$ is found.

A procedure to generate the looked for J is the following. First of all, one has to define an interaction network having an arbitrary geometry (i.e. with arbitrarily chosen non-zero elements) and to assign some values to the couplings of the couples belonging to this network. Then, one verifies whether such a matrix is positive defined; if not, one calculates the most negative eigenvalue $-\mu$ of J and modifies J according to $J \leftarrow J + (\mu + \epsilon) \mathbb{I}$, so to obtain a new matrix with minimum eigenvalue ϵ . The matrix obtained in this way is invertible, and one can hence calculate $C = -J^{-1}$; in general, this latter will not have all ones on the diagonals. In order to impose also this constraint, one can rescale the elements of J according to the following rule: $J_{ij} \leftarrow J_{ij} \sqrt{C_{ii} \cdot C_{jj}}$. One is at this point able to define the synthetic data to work on as $C_{true} = -J^{-1}$.

The described procedure for obtaining a synthetic J shows some drawbacks, in particular as regards the lack of control one has on the actual values taken by the element of the coupling matrix itself. Such drawbacks, potentially creating inference

problems whose difficulty is difficult to predict, and even generating convergence issues in cases where some of the interactions get very large, are presented in some details in section 5.10.1, to be considered as a final appendix to this discussion. In 5.10.2, some further discussions on the possible choices regarding the explicit values initially assigned to the couplings is finally given.

5.4.4 The observables

All the artificial models used in this project are designed in such a way to be able and assess the algorithmic performances on already completely known instances. One still needs to define however how to quantitatively perform this assessment; in some sense, this is an arbitrary choice, as depending on the features we are most interested in the same results could be considered good or not. The first possible observable to look at directly comes from the Bayesian approach discussed in section 5.3.2, as one considers the value of F of the most probable configuration, this latter having been obtained as a result of the gradient descent procedure. The highest this value is, the most probable is that the inferentially found coupling configuration corresponds to the true one.

A sanity check is done in order to verify this observable to lead to reasonable results. Three heuristics are defined for this. The first, null one consists in simply randomly choosing the next correlation element to measure; in some sense, this will be for all the project a natural baseline to which all the smarter heuristics will have to be compared. Other two strategies, on the other hand, can be used as a comparison only on artificial models, as they make use of the same elements one is aiming to retrieve. The order by which the measurements are taken corresponds in these cases to either a descending or an ascending order of the absolute value of the couplings (which will be unknown in a real situation). These strategies will be referred to respectively as *high couplings first* and *low couplings first* and, for brevity, often abbreviated in **HCF** and **LCF**. In order to check the reasonableness of the proposed observable, one has to verify these “fake” strategies to lead to very good (respectively, very bad) results.

At a first sight, it looks like the low couplings first and high couplings first strategies are respectively unbeatably good and bad. In some of the results reported in the following this apparently straightforward *extremality* feature will not be satisfied. After some thought, one sees that this is not an inconsistency. Because of the sparseness of the J , indeed, many couplings are degenerate and exactly equal to zero: under these conditions, the low coupling first and the high coupling first heuristics choose the next element to measure at random, and the possibility of another more




Figure 5.4: In order to evaluate the reasonableness of the chosen observable, the value of the minimum of F are reported for the two fake strategies LCF and HCF and for a random choice of the element to measure on a single instance as a function of the number of measurements already taken. In this latter case, the results obtained by averaging over 15 runs are reported. In all the cases, γ is fixed to 0.2 and no prior knowledge on the system is assumed. The results are very homogeneous: starting from the top left panel and proceeding clockwise, the figures refer to a $N = 15$ ring, to a $N = 15$ random regular graph with $k = 3$, to a $N = 15$ Erdős–Rényi network with average degree $\langle k \rangle = 2.2$, and to the smallest of the hierarchical structures proposed in figure 5.3.

complicated strategy to perform better than this random choice has not to be a priori excluded.

The results of these three strategies are found to be almost equivalent in all the analysed geometries. As can be seen in figure 5.4, when averaging over multiple runs the random choice corresponds to the straight line connecting the initial and the final point of F as a function of the number of measurements. These latter are by definition equal for all the strategies, as they corresponds to the inference respectively made on a completely unknown or on a completely known matrix, and so there is no difference among the different heuristics. The HCF leads instead to a very fast increase in F at the beginning, when the correlations corresponding to real couplings are measured, and virtually no further improvements when one starts to measure the positions of C_{true} associated to zero couplings. A specular situation holds for LCF, telling us that the chosen observable is, at least at this level of check, reasonable.

Other two observables are considered; they are a bit further away from an abstract, Bayesian definition of the problem but can be more easily related to the quantities one would be interested in in a real case. They can be considered as ways of better understanding the features of the artificial models, as they again are calculated by using information on the systems that one would not know in a real situation. As the final objective is to reconstruct the interaction matrix J , a very legitimate quantity to look at for assessing the performance of the different heuristics is the mean squared error on the elements of the J_{inf} which is the estimate of the coupling matrix at any level of knowledge of the correlations.

In some situations, however, one may be not so interested in recovering the actual values of the elements of J , as determining the positions having a different from zero coupling (i.e. the “geometry” of the interaction) can already be a sufficient insight on the system. In order to take into account that the mistakes made in the inference of the values of J are not so serious as soon as one has understood where the contacts are, one can look at the *true positive rate* of the inferred coupling matrix. For doing this on a network where the number of true contacts is M , at each step of the procedure one looks at the fraction of the largest M elements of J_{inf} corresponding to J elements different from zero. The value one obtains in this way is in a range between zero and one; when it is one, in particular, all the M true couplings are among the M largest values of J_{inf} , the order among them being however irrelevant.

The three previously defined observables have been seen to produce very consistent results in all the cases having been looked at. As can be seen in figure 5.5, in particular, when choosing at random the elements to measure it appears clear that, modulo




Figure 5.5: Linear relationships holding among the observables proposed to evaluate the performances of the different heuristics. All the results refer to a random choice of the elements to measure averaged over 50 runs on a $N = 30$ single instance of a ring and of a Erdős–Rényi graph. No initial knowledge on the system is assigned. On the left panel, the relationship between true positive rate and value of F_{min} is reported, whereas on the right panel the same is shown for the mean squared error on the reconstructed J_{inf} and, again, the value of F_{min} .

some minor fluctuations, simple linear relationships connect the values of the different observables. I will typically refer in the following to the true positive rate as it is what most directly connects to the physical problem one is interested in solving.

5.5 Main analytical results

5.5.1 Reasons justifying the study of the Hessian

The most innovative strategies that will be proposed will make use of what could be called a *geometrical characterisation* of the probability space under exam. Anticipating a discussion that will be more detailed in section 5.8.2, one can say that when partial information on a system have been used to infer some of its features, two ways of improving our knowledge about it can be thought of. Firstly, one can go to the less known region (i.e. by using a Bayesian language the one in which our prior is weaker) and perform some measurements therein so to be able to let the inferred quantities be more accurate. Another possible strategy is, on the other hand, to verify the correctness of the variables on which our prior is stronger, so to improve a lot our knowledge on the system if we find out that our previous estimate about them did not correspond to their actual value.

For implementing both of these strategies in our problem, it is needed to evaluate the N^2 by N^2 Hessian matrix, whose position $((i, j), (k, l))$ corresponds to the value

of $\frac{\partial^2 F}{\partial C_{ij} \partial C_{kl}}$. Such a matrix contains indeed the information about how the F one is trying to optimise so to obtain the most possible reliable estimate of J depends on the values of the elements of C , and more in particular it describes its local curvature as a function of these latter. Connecting this matrix to the intuitive reasoning of the previous paragraph, hence, if one chooses to measure the couple (i, j) minimising the Hessian, he is looking at the variable around which F is the “flattest”; in some sense, one is choosing to measure a correlation characterised by our lack of information or intuition on how F depends on it. On the other hand, measuring the element corresponding to the maximum of the Hessian can be thought of as going and checking the actual value of a correlation on which we have the strongest prior; the reasoning behind this choice is that if it turns out that our estimate was incorrect, our knowledge on the system suddenly improves.

At the beginning of the project, we suspected that one out of the two previous reasoning was correct whereas the other was not, and hence that one had better always choose to measure respectively the unknown element of C in correspondence of which the distribution of F was either the flattest or the most peaked. This ansatz has however been proven to be false, and which one out of the two ideas was the one leading to the best quality of the inference turned out to depend on specific features of the system under exam, and in particular on the geometry of its interactions.

5.5.2 Derivation of the Hessian

In what follows, several heuristics using some features of the geometry of $F(C)$ in order to decide which element to measure next will be defined and their results will be compared to the ones obtainable by using more naive methods. In order to implement these ideas we have first of all to derive the expression of the Hessian of F with respect to the elements of the correlation matrix C . The main elements of this calculations will be reported here, even if a precise description of all the passages will be neglected in favour of a more detailed discussion on the results.

In order to obtain the Hessian, one has first of all to write down the first derivative of F . The generic form of the equation already discussed in eq. 5.9 is the following:

$$\frac{\partial F}{\partial C_{ij}} = -J_{ij} - (1 - \alpha)(C^{-1})_{ij} \quad (5.10)$$

The second derivative with respect to a (in general) different element of C will then read:

$$\frac{\partial^2 F}{\partial C_{ij} \partial C_{kl}} = -\frac{\partial J_{ij}}{\partial C_{kl}} - (1 - \alpha) \frac{\partial (C^{-1})_{ij}}{\partial C_{kl}} \quad (5.11)$$

Both the terms on the right-hand side can be written in a more explicit way. The main ingredients of this somehow involved derivation are briefly summarised. First of all, one has to remember that the correlation matrix is symmetric; because of this, the *self-derivative relations* entering into the calculations assume the following non-trivial form:

$$\frac{\partial C_{bc}}{\partial C_{kl}} = \delta_{bk}\delta_{cl} + \delta_{bl}\delta_{ck} - \delta_{bc}\delta_{bk}\delta_{cl} \quad (5.12)$$

As before, one has to make use of the relation 5.8 connecting J and C so to eliminate the explicit dependency of the formulas on the correlation matrix.

The presence of a square root in 5.8 makes the matrix differentiation quite hard. This problem is bypassed by making use of the power expansion of such square roots so to be able to differentiate separately each term in the infinite sum. At the end of the algebraic passages, it turns out that these terms can be summed up again, and the final results will be written down in a much more convenient (also thinking about algorithmic implementations) way.

In order to more easily deal with the power expansion of the terms involving C , a eigendecomposition of this latter is intensively used. This consists in writing down the symmetric matrix C in terms of a diagonal matrix Λ whose entries are the eigenvalues of C and an orthogonal matrix Q whose m^{th} column v_m corresponds to the m^{th} eigenvector of C . The relation connecting these three quantities is $Q = V\Lambda V^T$ and in particular the (i, j) element of the correlation matrix is expressed as:

$$C_{ij} = \sum_{m=1}^N v_m^i \lambda_m v_m^j \quad (5.13)$$

The final formulation of the Hessian can be written in a synthetic way by introducing an auxiliary matrix I : the $I_{mm'}$ elements of this latter take different values depending on the corresponding eigenvalues to be equal to each other or not, and in particular:

$$I_{mm'} = \begin{cases} \frac{\lambda_m}{\sqrt{4\gamma + \lambda_m^2}}, & \text{if } \lambda_m = \lambda_{m'} \\ \frac{\sqrt{4\gamma + \lambda_{m'}^2} - \sqrt{4\gamma + \lambda_m^2}}{\lambda_{m'} - \lambda_m} & \text{if } \lambda_m \neq \lambda_{m'} \end{cases} \quad (5.14)$$

Resumming all the pieces together, one obtains the following formula for a generic element of the Hessian, characterised by four indices (i, j, k, l) :

$$\begin{aligned} \frac{\partial^2 F}{\partial C_{ij} \partial C_{kl}} = & -\frac{1}{2\gamma} \left\{ \left[(1 - \delta_{ij}) \delta_{ik} \delta_{jl} + \delta_{il} \delta_{jk} \right] + \right. \\ & - \sum_{m=1}^N \sum_{m'=1}^N \left\{ v_m^i v_{m'}^j \left[(1 - \delta_{kl}) v_m^k v_{m'}^l + v_m^l v_{m'}^k \right] I(m, m') \right\} \left. \right\} + \\ & + (1 - \alpha) \left[(1 - \delta_{kl}) (C^{-1})_{ik} (C^{-1})_{lj} + (C^{-1})_{il} C_{kj}^{-1} \right] \quad (5.15) \end{aligned}$$

Of particular interest for what concerns the heuristics that will be used in the following are the diagonal elements of the Hessian:

$$\begin{aligned} \frac{\partial^2 F}{\partial C_{ij}^2} = & -\frac{1}{2\gamma} \left(1 - \sum_{m=1}^N \sum_{m'=1}^N \left\{ v_m^i v_{m'}^j \left[(1 - \delta_{ij}) v_m^i v_{m'}^j + v_m^j v_{m'}^i \right] I(m, m') \right\} \right) + \\ & + (1 - \alpha) \left[(1 - \delta_{ij}) (C^{-1})_{ii} (C^{-1})_{jj} + (C^{-1})_{ij} (C^{-1})_{ij} \right] \quad (5.16) \end{aligned}$$

In order to obtain more readable equations, a change of basis turns out to be useful. Since the decomposition on the basis of the eigenvectors of C has been deeply used in the derivation, it is not unreasonable to imagine that also the final result will be easier to read in the same basis. Indeed, by moving to it with the following change of coordinates:

$$M_{m_1, m_2; m'_1, m'_2} \equiv \sum_{i, j, k, l} v_{m_1}^i v_{m_2}^j \frac{\partial^2 F}{\partial C_{ij} \partial C_{kl}} v_{m'_1}^k v_{m'_2}^l \quad (5.17)$$

and working again separately on the terms of eq. 5.15, one is able to reach the following very compact relation:

$$M_{m_1, m_2; m'_1, m'_2} \equiv \left(\delta_{m_1, m'_1} \delta_{m_2, m'_2} + \delta_{m_1, m'_2} \delta_{m_2, m'_1} \right) \left[-\frac{1}{2\gamma} + \frac{I(m_1, m_2)}{2\gamma} + \frac{1 - \alpha}{\lambda_{m_1} \lambda_{m_2}} \right] \quad (5.18)$$

5.5.3 Derivation of the Hessian in the eigenbasis of C

In the eigenbasis of C , the Hessian can be written in a block-diagonal form after a permutation of rows and columns. In particular, it is composed by some diagonal elements corresponding to the (m_1, m_1) elements, and some 2 by 2 blocks associated with a couple (m_1, m_2) ; these four elements are moreover equal to each other because of the symmetry properties of the Hessian. The eigenvalues of such a matrix are its diagonal elements, called for brevity f_{m_1, m_1} , and for any block one of the block

elements f_{m_1, m_2} and a zero. These values can hence be organised in a $N \times N$ matrix where the dimensions correspond respectively to m_1 and to m_2 .

One is able to recover the elements in the original basis by doing the following:

$$\begin{aligned} H_{ij,kl} &= \sum_{m_1, m_2} f_{m_1, m_2} \langle ij | m_1 m_2 \rangle \langle m_1 m_2 | kl \rangle = \\ &= \sum_{m_1, m_2} f_{m_1, m_2} (v_{m_1}^i v_{m_2}^j + v_{m_1}^j v_{m_2}^i) (v_{m_1}^k v_{m_2}^l + v_{m_1}^l v_{m_2}^k) \end{aligned} \quad (5.19)$$

A computationally convenient way of performing this calculation has been found in terms of *Kronecker products* [154]. This latter, not reported here for the sake of brevity, has been verified to produce the very same results of a direct, and much more inefficient, calculation of the Hessian in the original (i, j) basis.

5.5.4 Perturbative approximation of the Hessian

As will be discussed in the following section 5.6, the explicit use of the complete Hessian to select the next correlation to measure is extremely heavy from a computational point of view. Another strategy therein defined will need as a key component the evaluation of how the eigenvalues of the Hessian change as the line and the column associated to a given couple of sites (i, j) are removed. The perturbative calculation discussed in this section allows to efficiently get approximate estimates of such quantities.

Let us start by defining the modified Hessian $\hat{H}^{(i,j)}$ as the one in which all the elements in the row and in the column corresponding to (i, j) are set equal to zero, except the diagonal one which is set to one. The eigenvalues of $\hat{H}^{(i,j)}$ are the same as the ones of the restricted Hessian one gets by removing the (i, j) row and column, plus a 1. This calculation will enable the design of a much more efficient procedure that will substitute the explicit diagonalisation of all the N^2 reduced matrices with the calculation of the perturbed eigenvalues of the complete Hessian in absence of a row and of a column.

The previously introduced modification of the Hessian can be written down as follows:

$$\hat{H}_{ab,cd}^{ij} = H_{ab,cd} - \Delta_{ab,cd}^{ij} \quad (5.20)$$

where Δ is element-wise defined as follows:

$$\Delta_{ab,cd}^{(i,j)} = (\delta_{ab,ij} + \delta_{cd,ij})H_{ab,cd} - \delta_{ab,cd}\delta_{ab,ij}(H_{ab,cd} + 1) \quad (5.21)$$

The first order perturbative approximation of the eigenvalues $\hat{\lambda}^{(i,j)}$ of the modified Hessian, whose validity has still to be evaluated, can be derived as:

$$\hat{\lambda}_{m_1,m_2}^{(i,j)} \simeq \lambda_{m_1,m_2} - \langle v_{m_1,m_2} | \Delta_{ab,cd}^{(ij)} | v_{m_1,m_2} \rangle \quad (5.22)$$

One of the convenient features of working with the eigenvalues and not with the complete matrices is that the former are the same in the original basis (a, b, c, d) and in the eigenbasis (m_1, m_2, m'_1, m'_2) . One is hence not forced to explicitly move between the two.

The different contributions of Δ to 5.22 can be separately evaluated, leading to the following final result:

$$\hat{\lambda}_{m_1,m_2}^{(i,j)} \simeq \lambda_{m_1,m_2} + \frac{1}{2} \left(v_{m_1}^i v_{m_2}^j + v_{m_1}^j v_{m_2}^i \right) \left[H_{ij,ij} + 1 - 2\lambda_{m_1,m_2} \right] \quad (5.23)$$

This expression is quite convenient as the only information needed are the values of the original eigenvalues and the diagonal elements of the Hessian. Moreover, in the eigenvector basis the eigenvalues can be directly written, without having to explicitly write down and diagonalise the complete Hessian matrix.

5.6 The heuristics

5.6.1 Choices based on inference

First of all, two strategies for choosing which correlation to measure based in a very direct way on the inferential procedure described in section 5.3 are studied. They have been defined in order to be able and compare the more involved heuristics that will be discussed in the following sections to *null strategies* a bit smarter than a simple random selection of the elements to measure.

These heuristics make use of the results of the gradient descent performed on the correlation elements still unknown after each measurement on C_{true} . As a result of this part of the procedure, one obtains a C_{inf} and a J_{inf} corresponding to the minimum of F . The simplest thing one can do as a following step is hence to measure the maximum element of C_{inf} or of J_{inf} among the yet unmeasured correlation elements. In some sense, one relies on the inference procedure itself and chooses to measure the elements which are going to be the most informative according to it.

For brevity, these procedures will be referred to as *high inferred C* or *high inferred J* in the following, and sometimes abbreviated as **HIC** or **HIJ** respectively. As will be discussed in the following, both these strategies reach results typically slightly better than a random choice, but especially HIJ gives on some geometries (in particular on hierarchical structures) really impressive results; the reasons explaining this feature have still to be fully understood, but clearly they depend on the characteristics of the topology of J .

5.6.2 Heuristics making use of the Hessian

As discussed in section 5.5.1, a potentially interesting way of rapidly improving our knowledge on the system under exam is by exploiting the information contained in the Hessian matrix of F with respect to the correlation matrix C . Two slightly different heuristics can be proposed, based on the results of section 5.5.2 thanks to which one can easily enough have access to the Hessian values. The most intuitive is to write down the elements of the Hessian in the basis of the eigenvectors of C , in which they are particularly simple. By considering the leading eigenvector of this matrix, one should be able to identify the most important element. However, by doing this a couple (m_1, m_2) of eigencomponents of C is found, rather than a couple of sites (i, j) ready to give a hint on what correlation to measure next. One can nevertheless write again this vector in the original basis, and decide to measure the couple (i, j) contributing the most to the principal eigenvector. This strategy has a set of interesting possible extensions; k measurements at the same time could for instance be straightforwardly implemented by selecting the k largest components of this vector. This idea turned out, however, to be quite difficult to apply. When translating the leading eigenvector back into the original basis, its largest components typically correspond to already measured couples; as on the other hand one is forced to measure some still unknown correlations, in the end the couple which the heuristics suggests us to measure is one not contributing so much to the leading eigenvector of the Hessian. This strategy, hence, does not work as nicely as one would have naively expected.

An alternative route is proposed by making use of the results reported in section 5.5.3 on how to conveniently write down the elements of the Hessian matrix in the eigenbasis of the correlation matrix C . By exploiting that result, indeed, one can define a greedy strategy as follows. For all the couples (i, j) not having been measured yet, one removes the corresponding row and column from the Hessian and calculates

the determinant of the reduced matrix. A possible way for selecting the next measurement is hence to choose the couple which, when removed, leaves the determinant of the Hessian the largest or the smallest possible. They will be respectively referred to as *high Hessian first* and *low Hessian first*, and often abbreviated as **HHF** and **LHF**. For reasons still to be completely understood, it appears that in most cases the latter choice leads to better results; a brief, intuitive discussion on this point will be presented in section 5.8.2.

5.6.3 Intractability of the complete Hessian

Even after having introduced the simplification connected to studying the Hessian in the much simpler eigenbasis of the correlation matrix, the previously discussed strategy shows still two major drawbacks. The first one is that one needs to explicitly write down the complete $N^2 \times N^2$ Hessian, and hence runs into memory problems even for quite small values of N . The second one is that all the Hessians reduced by removing a row and a column have to be explicitly diagonalised, even if one reasonably expects the corresponding eigenvalues to be quite similar among each other. This task is terribly heavy from a computational point of view. At any time one needs to choose which correlation to measure next, indeed, one has to calculate the complete Hessian and then to diagonalise an $O(N^2)$ number of reduced Hessians of size $(N^2 - 1)$ by $(N^2 - 1)$ each. Recalling that the diagonalisation of matrix of size p is a task of computational complexity of order $O(p^3)$, each measurement would require approximatively the intractable number $O(N^2 \cdot (N^2)^3) \sim O(N^8)$ operations.

These issues are addressed by proposing in the following two different approaches. In the first, only the diagonal elements of the Hessian will be used. This approximation whose validity is not at all guaranteed has two important positive features. On the one hand, it reduces the size of the matrices one has to deal with, and on the other it shows a very straightforward mapping between the elements of such matrices and the couples (i, j) of sites; because of this, deciding which correlation to measure next will be particularly easy as no further operation will be needed to go back into that basis.

Secondly, the perturbative calculation discussed in section 5.5.4 will be exploited. Having at our disposal the perturbative expansion of the eigenvalues of the Hessian in absence of one row and column, one can again sum these values so to obtain an approximation of the determinant for each of the reduced Hessians, and then choose the next element to measure by following the same principle adopted when dealing with the explicit calculation of all the reduced Hessians.

5.6.4 Diagonal elements of the Hessian

Using the calculations of section 5.5.2, two more strategies to choose the measurements that could be implemented also during an inference procedure on real, unknown data and not only on artificial ones will be defined according to the diagonal values of the Hessian matrix. In section 5.5.1 the reason why the use of the Hessian can be of interest have been discussed. Two elements makes difficult an explicit implementation of such principles. On the one hand, the intractability of the complete Hessian from a practical point of view even for quite small systems has been discussed in section 5.6.3. Secondly, and at an even more fundamental level, the elements of the Hessian depend on four indices, whereas one is interested in determining a couple of sites (i.e. two indices) whose correlation is not known yet and which is more likely to improve the most our knowledge of the interactions among the sites.

A possible way to address both these problems, whose efficacy has to be tested, is to suppose that the best part of the information of the Hessian one wants to exploit is contained in its diagonal part, whose elements can be determined according to equation 5.16. The validity of this ansatz will depend on the relative importance of the out-of-diagonal terms, this latter being associated with the curvature of the Hessian as a function of two different elements of the correlation matrix C . We suspect this reasoning to be strongly dependent on the geometry of the system under exam; in some cases, indeed, it has been seen that analogous results can be obtained by using the information contained in all the Hessian or just the ones included in its diagonal part. As a future perspective, one would like to get a more quantitative characterisation of the networks for which the hypothesis holds, probably in terms of spectral properties of the graph defined by the couplings different from zero.

It is however easy to see why the validity of this ansatz would be a good news for approaching more realistic ways of solving the problem. This is first of all true from a computational point of view, as the diagonal of the Hessian is a vector of N^2 elements of the type $\frac{\partial^2 F}{\partial C_{ij}^2}$ that can even be reorganised in a very intuitive form in a N by N matrix where the two indices refer to the two sites i and j . The difference of this latter with the complete N^2 by N^2 Hessian matrix is really significant in terms of the size of the system one can hope to deal with. Secondly, each element of this N by N matrix is univocally associated to a couple of sites, and the elements maximising or minimising it can be very naturally and directly chosen as the ones to measure in the following iteration of the inference procedure.

The two strategies consist in recalculating after each measure all the diagonal elements of the Hessian matrix corresponding to couples (i, j) not yet measured, and




Figure 5.6: Validity of the approximation of the Hessian for rings of different sizes. Even if on the smallest $N = 20$ ring the strategy using the perturbed Hessian seems not to work as well as the one using the complete Hessian, by analysing a $N = 40$ ring (on which a direct comparison is not possible as it is too large for explicitly running the LHF heuristics) one sees that, after having rescaled the curves, the results obtained with the perturbed Hessian are almost equivalent to the ones obtained with the LHF strategy on the smaller ring. The true positive rates obtained by selecting at random the elements to measure in the two cases is reported for comparison.

then to select the couple for which $\frac{\partial^2 F}{\partial C_{i,j}^2}$ is respectively the largest or the smallest. These two strategies will be referred to in the following as *high Hessian diagonal* and *low Hessian diagonal*, and respectively abbreviated as **HHD** and **LHD**.

5.6.5 Validity of the perturbative approximation

As discussed in section 5.6.3, strategies making use of a perturbative expansion of the Hessian can be really convenient. In particular, after having calculated the approximated determinant of all the reduced Hessians obtained by removing one row and one column, two heuristics are defined choosing the following measurements in exactly the same way as the previously discussed HHD and LHD. They will be referred to as *low perturbed Hessian* and *high perturbed Hessian* and abbreviated as **LPH** and **HPH** respectively.

The perturbative approximation of the Hessian of F is expected to be more and more accurate as the system size increases. The truth of this statement is difficult to evaluate in this framework, as the calculation of the complete Hessian becomes computationally intractable for small enough values of N . An indirect proof has

nevertheless been obtained. When the inference is performed on couplings having a ring geometry, for small N it is possible to obtain very good results by using the strategies exploiting the complete Hessian, whereas by using the perturbative expansion the results are no more so impressive. For increasing N , on the other hand, it is not possible to calculate the complete Hessian any more, but as can be seen in figure 5.6 the approximated one leads to good results, probably because of its greater consistency with the analytical result for large enough system sizes.

5.7 Initial knowledge and effects of measurements

5.7.1 Choice of the initial condition

A last point needs to be defined before explicitly running the simulations, that is the knowledge on the system one already has before starting the inference procedure. The simplest case is to imagine that none such knowledge is a priori available. In this case, however, one has to deal with an initial transient in which the inference will concern only the components on which some information is available. After the very first correlation measurement involving sites i and j is performed, also the C_{inf} one obtains at the end of the gradient descent will be equal to zero everywhere apart from the position C_{ij} , because of the regularisation term. This is very reasonable, as we have in some sense no clue whatsoever letting us imagine that other parts of that matrix are different from zero.

The previous feature can be extended as follows. Let us suppose that the m taken measurements can be partitioned in p groups m_1, \dots, m_p such that each node is present in at most one of them. If this condition holds and the group m_1 includes at least one correlation involving node 1, for instance, in groups m_2, \dots, m_p there will be no correlations including it. Let us call I_k the set of nodes spanned by the correlations in m_k . Considering a $N = 15$ graph, one such situation can be found if the first 7 measurements can be partitioned as follows:

$$\begin{cases} m_1 = \{(C_{true})_{1,2}, (C_{true})_{1,4}, (C_{true})_{2,5}\} \\ m_2 = \{(C_{true})_{10,11}, (C_{true})_{11,15}, (C_{true})_{10,13}, (C_{true})_{13,15}\} \end{cases}$$

where $I_1 = \{1, 2, 4, 5\}$, $I_2 = \{10, 11, 13, 15\}$ and all the other nodes have not been touched by any correlation measurement yet.

Again in reason of the regularisation imposed on the system, the C found as a result of the gradient descent is found to show the following block structure:

$$C_{ij} \quad \begin{cases} \neq 0 & \text{if } i, j \in I_k \text{ for some } k \in \{1, p\} \\ = 0 & \text{otherwise} \end{cases}$$

Also in this case, thus, the results of the inference do not concern the entire system, but rather some of its components separately.

A possible solution in order to bypass this initial transient is to start from a \tilde{C} already having information on the system as a whole. For example one could include in the initial \tilde{C} a set of elements $(C_{true})_{i_p, j_p}$ such that the ensemble of links $\{(i_p, j_p)\}$ forms a spanning tree on the graph. Another choice that will be used in what follows is to begin the procedure from a *star-like* knowledge of the system with respect to a given node i , that is $\tilde{C} \ni (C_{true})_{ij} \quad \forall j$. In the geometries heterogeneous in the degree of the nodes as the Erdős–Rényi graphs, the dependence of the results on the choice of the node i selected for the initial knowledge is going to be discussed; it is possible, in fact, that a star-like knowledge around a node with a higher degree will be more informative, and will therefore induce a better inference of the coupling matrix J .

Also the tree-like initial knowledge has, nevertheless, some drawbacks that need to be carefully considered. In particular, if the tree over the nodes of the system is selected at random, it could consist in a set of couples of nodes none of which being linked by a coupling. In this case, the first measurements will probably modify very abruptly the estimated J and C , possibly creating some convergence issues in the gradient descent. Two other initial conditions have been studied to address this issue. In the first one, the initial knowledge on the correlation consists in a fixed percentage of the couples actually interacting through J . Also in this case, however, the initial estimates will be very different from the true matrices, as it is legitimate to expect a very dense interaction matrix if all the measured correlations are strong. Another initial knowledge potentially solving this problem is what will be referred to as a *rich* initial condition. In this latter, we imagine to have a good knowledge of the links in a given region of the system. In the network of the interaction, one selects one of the nodes at random and measures the correlations of this nodes with its neighbours and, if these latter are too few, iterates by also measuring some of the correlations they have with their own neighbours (i.e. arriving at distance 2 from the initially selected node).

In conclusion, no final answer has been reached on which initial condition is the most appropriate to assign. This study has to be considered as an attempt to understand the effects of different initial knowledge on the algorithmic performances of the

proposed heuristics on each geometry. Thinking about applications, moreover, one has to consider that in many cases one has to work with initial conditions of a given type, without having the possibility of arbitrarily choosing them.

5.7.2 Update of C after a measurement

A discussion on the issues emerging when trying to complete different types of initial conditions is delayed to appendix 5.10.5. Another moment in which a somehow arbitrary choice has to be made is after any given measurement. After that the gradient descent has led us to the previous estimate of C , one chooses which of the still unknown elements to measure next, by enriching of one element the current \tilde{C} . Before iterating the procedure by running again the gradient descent, however, the correlation matrix has again to be completed by assigning some value to the still unknown elements, i.e. to the positions still belonging to C_{\perp} . In principle, one could do this by treating at every iteration the \tilde{C} as the initial condition and by using the completion procedure described in 5.10.5. After some thoughts, however, it turns out that this is not probably the most efficient choice as one is wasting all the information obtained as a result of the previous gradient descent, even if the matrix has typically not changed very much and we hence expect the result of this new minimisation to be quite similar to the former. In order to make use of this intuition, one can use as a starting point of the new gradient descent on the still unknown elements the result of the previous one after having fixed the measured correlation to its true value. The matrix so obtained is strictly speaking no more guaranteed to be positive definite, and if it is not the gradient descent will not even start. This procedure has been however verified to be much more efficient than the other and to typically work. For the sake of generality, as a future perspective one would like to insert a switch so to use the least efficient procedure if the smartest has convergence issues.

5.8 Results

The problem introduced is far from being completely understood. In the following, the main results obtained are reported. Even if some of them can be given a non-rigorous justification or can intuitively be related the ones to the others, it is clear how a unified framework producing these results as particular cases is still missing.




Figure 5.7: True positive rate averaged over multiple instances and runs on $N = 33$ hierarchical networks for several strategies. The left panel refers to an initial complete ignorance about the system, whereas in the right panel a tree-like initial knowledge is assigned. For reasons still to be fully understood, some of the strategies appears to be much stabler than others with respect to the initial knowledge of the system.

5.8.1 Dependency on the initial condition

It turns out that the effectiveness of the different heuristics depends in a non-trivial way on the initial conditions.

In figure 5.7 the averaged curves of the true positive rates both for none initial knowledge and for a tree-like initial knowledge on hierarchical networks with $N = 33$ nodes are shown. As one can see, some strategies appears to be stabler as they almost do not change by incrementing the initial knowledge on the system; this is true for a random choice of the elements to measure, and for the heuristics previously called HHD and HPH. Others such as HIJ, instead, dramatically improve by enriching the initial knowledge. It is worthy to say that the fact of a heuristics being stable with respect to the initial condition appears to depend on the geometry on which the couplings are defined.

5.8.2 The exploration-exploitation trade-off

The way of calculating the complete Hessian, its diagonal part and a perturbative approximation of it have been discussed. One has still to decide how to practically make use of the information contained in such matrices. After having calculated the diagonal element of the Hessian, for instance, it appears reasonable to measure the couple that either corresponds to the maximum or to the minimum of this matrix. A rigorous understanding of this point is still missing, but the reasons why both these idea could be a reasonable choice have been intuitively presented in section 5.5.1.




Figure 5.8: Analysis of the opportuneness of measuring the couple corresponding to the largest or the smallest diagonal element of the Hessian respectively on a ring (left panel) or on a hierarchical network (right panel). The same is studied also for the strategies making use of the approximated version of the complete Hessian.

After some thought, one sees how the opportuneness of following the former or the latter principle depends on how often the inferences made on the system are wrong. If this happens often, indeed, verifying our estimates is going to improve the results more than given them for granted and going on to measure other parts of the system. From the simulations it turns out that for J with different geometries the situation may be very different, and in some cases almost specular. In figure 5.8, one can see that on the ring (left panel) the strategies measuring the elements corresponding to the smallest element of the Hessian lead to very good results, whereas if one measures the couples with largest Hessian the inference of the underlying interaction network is slower than by taking random measurements. On the right panel the results on a hierarchical structure are reported, and the effectiveness of the different strategies is reversed.

In some sense, this feature can be connected to the well-known problem in machine learning going under the name of *exploration-exploitation trade-off* that has already been briefly addressed in the previous section 5.1.3. Having already collected some information on a system, it is not trivial at all to decide whether the best option is to explore what one does not know yet about it, or rather to exploit what one already knows. A mixed strategy alternating the two approaches has been imagined and some preliminary results about it have been obtained; however, more work is needed to better understand the theoretical reason at the basis of this discrepancy that one gets by changing the topology of the couplings.




Figure 5.9: Average true positive rate curves for different heuristics on $N = 20$ rings, without any initial knowledge on the system (left panel) or with a tree-like initial knowledge (right panel).

5.8.3 Different heuristics on a fixed geometry

By looking at the performances of the different proposed heuristics on networks on which the topology of the couplings is fixed, we have not been able yet to draw a comprehensive picture. Some common features of the results have nevertheless been found and will be exposed in this section.

The ring and the random regular graphs seem to offer the most easily interpretable framework. In both the cases, indeed, the strategies preferring to measure the still unknown parts of the graph (for instance LHD) systematically lead to an inference that is faster and more reliable than the ones that go and verify the potential errors in what has already been estimated with a certain degree of confidence during the procedure (for instance HHD).

As can be seen in figure 5.9, already on a small enough graph this trend can be appreciated. By assigning a richer, tree-like initial knowledge on the system, moreover, the strategies can be ranked quite easily according to their performances already from the very first measures, whereas starting from no knowledge at all there is a more difficult to evaluate initial transient in which the performances are all similar. In the former case, it is worthy to remark how the results obtained via the two strategies more directly based on the previous inference results, and especially HIJ, improve a lot; nevertheless, they do not get as good as the heuristics based on the Hessian.

What has been said is still true for larger systems, even if in this case it is no more possible to run the strategies using the true Hessian because of computational intractability; for brevity, in this case only the results obtained by imposing a tree-like




Figure 5.10: Average true positive rate curves for different heuristics with a tree-like initial knowledge on a $N = 40$ ring (left panel) and on a $N = 40$ random regular graph with $k = 3$ (right panel).




Figure 5.11: Average true positive rate curves for different heuristics with (right panel) and without (left panel) a tree-like initial knowledge on a $N = 36$ square lattice.

initial knowledge are reported. In the left panel of figure 5.10 one can see how the two strategies based on the Hessian LHD and LPH are still the best. For the random regular graph the situation appears to be very similar; in the case reported in the right panel of figure 5.10, however, for a high enough number of measurement the inferential strategy HIJ gets almost as good as the one using the perturbed Hessian, whereas the LHD heuristics is still the most effective.

Also as regards square lattices, two different sizes have been studied in order to be able and evaluate the effect of using the complete Hessian on the smallest one; the $N = 16$ networks proposed for this reason where the couplings form a 4 by 4 square lattice led however to very noisy results. Leaving as a future perspective the comprehension of why for such small networks the results are so worse compared for instance to rings or random regular graphs, the results on 6 by 6 lattices are reported




Figure 5.12: Average true positive rate curves for different heuristics with (right panel) and without (left panel) a tree-like initial knowledge on a $N = 33$ hierarchical network of interactions.

in figure 5.11. The right panel, where an initial tree-like knowledge is assigned before starting to use the heuristics, seems indeed to be very similar to the equivalent figure on random regular graphs. What is interesting and still to be understood is that in absence of an initial knowledge, and differently from what happened on rings and random regular graphs, the perturbed Hessian strategy leads to results that are just equivalent to the ones obtained by measuring couples at random. Having a richer initial knowledge of the system seems thus to be particularly crucial for this topology.

As mentioned in section 5.8.2, when considering hierarchical structures the situation slightly changes. In particular, in this case the optimal trade-off between exploration and exploitation could be different from before; because of the structure having a highly connected node, a certain number of leaves and a second layers of elements connecting the two, as soon as one has some information about the system it is possibly more profitable to continue and sample the same region than to look for information elsewhere. By looking at the left panel of figure 5.12 that refers to the inference in absence of a prior knowledge on a $N = 33$ hierarchical network, one easily sees how measuring the couple with the highest diagonal element of the Hessian (HHD heuristics) is the unique strategy significantly improving the results obtainable by a trivial random choice of the measures; interestingly, the same strategy led on all the other topologies to results even worse than the ones got by such a trivial selection. Another feature of interest is the effect of adding an initial knowledge on the system (right panel of figure 5.12). By doing this, indeed, the performance of HHD seems to stay almost unchanged, whereas the simpler inferential heuristics HIJ improves dramatically and becomes the most efficient. Also the heuristics making use




Figure 5.13: Average true positive rate curves for different heuristics with (right panel) and without (left panel) a tree-like initial knowledge on a $N = 20$ Erdős–Rényi network of interactions.

of the highest element of the perturbed Hessian, whose results are compatible with a random choice if no initial knowledge on the system is given, remarkably improves with a tree-like initial condition.

The last case to analyse is the one of interaction networks defining Erdős–Rényi graphs on the nodes of the system. For reasons that will be discussed in appendix 5.10.1, one has some right to believe that the inference will be more difficult than the one on the other geometries. In particular, we found several convergence issues when looking at larger system size; for these latter, smarter procedures for defining the J_{true} to be recovered have probably to be thought of in such a way to obtain artificial models one can afterwards deal with from a computational point of view. For smaller networks, such problems were found to be not as severe and the results can be seen in figure 5.13. As in the other cases, in absence of a prior knowledge it is not easy to determine which strategy leads to the best results. For a large enough number of measurements the LHF and the LPH heuristics seem to be the most effective; this is in accordance, for instance, with the results found for random regular graphs and rings, even if in this case the improvement with respect to a random choice is much reduced. The first part of the curves, however, associated to the situation in which few correlation measures have been taken, reveals an apparently contradictory behaviour, as the HHD and HPH heuristics seem to be the most effective. This feature could depend in a non-trivial way on the geometry of the couplings, as it is not eliminated by using a tree-like initial knowledge on the system (see the right panel of figure 5.13) which gives as usual, much less noisy and much easier to read results. By looking at these latter, the good performances of the LHF and LPH strategies in a wide

range of number of measurements is confirmed; as before, if the initial knowledge of the system is richer also the simpler “inferential” strategies HIJ and HIC become significantly more effective than a random choice of which element to measure next.

5.8.4 Different graphs, different features

Even if one can somehow accept the fact that the proposed algorithms work differently depending on the geometry of the interaction, what is still missing from our framework is a more quantitative explanation of the features of the graphs that are the most significant in order to determine such differences. Various attempts have been made by trying to connect the efficacy of the heuristics to the sparseness of the network, to its loopiness or to its degree heterogeneity, but no satisfactory answer has been found yet in this regards.

What follows should be seen as the beginning of a work aiming at understanding this point. As has been seen, on some geometries the heuristics consisting in measuring the correlation connected to the smallest diagonal element of the Hessian appears to work very well. This implies that one could expect the couples for which this quantity is smaller to be more often associated to true couplings than the other ones. For letting this intuitive reasoning be more quantitative, one can study several instances supposing that a certain number of measurements has already been taken. For all of them, the elements of J are ordered according to the value of the diagonal part of the Hessian; all these permuted lists of couplings are finally averaged.

If the previous assumption is correct, one would expect the mean array of the interactions to be increasing; this would mean that, for instance, the couple whose diagonal element is the 10th smallest in magnitude is more likely to be a real coupling than the one associated to the 20th, and less likely to be so than the 3rd one. Unfortunately, by looking at the results it turns out that the situation is more involved than that.

In figure 5.14 it can be seen a situation in which the heuristics has been verified not to work so much better than a random choice. As expected, there is no monotonicity whatsoever; it almost looks like the elements that most frequently correspond to real couplings are the intermediate ones.

Even in cases where the heuristics works, however, it appears not to be a very clear global tendency (see figure 5.15 for the results on rings and on random regular graphs); still, the couple with the smallest diagonal element of the Hessian has typically a coupling larger than the others. Such a quantity clearly contains some information about the system and on the ring, in particular, it is associated to a value of average J




Figure 5.14: An analysis of the average relationship between value of the diagonal part of the Hessian and value of the coupling is proposed by looking at the average value of J of the couple having the x^{th} smallest diagonal element of the Hessian. The figure refers to 100 $N = 25$ Erdős–Rényi instances; on each, 50000 different initial conditions of the type called “rich” in section 5.7.1 have been studied, with the 20 percent of the correlations corresponding to true couplings known from the beginning. The two quantities plotted appear not at all to be monotonically related.

that is almost the double of the second largest one; since in this case all the couplings are quite homogeneous among each other, this feature can be translated in a more intuitive form by saying that the highest ranked couple corresponds to a real coupling twice more often than the second highest does.

The global trend one was hoping to see in this kind of graph is missing. This could potentially be a problem when one is interested in taking multiple measurements at the same time. Since the second ranked element on the couples does not seem to show qualitatively different features, one can suspect that for instance measuring the two couples associated to the two smallest diagonal elements of the Hessian would produce results similar to the ones obtained by measuring the first one plus another one randomly chosen among all the others.

By better looking at the results, it looks like the situation is not as hopeless as that. Even if a clear global trend is lacking, one finds that, on average, it is not only the highest ranked element to contain some information about the system. Out of 317 elements, on random regular graphs the four most highly ranked couples are all among the 19 having the largest average value of J . The same holds also for the ring, where out of 295 elements the four highest ranked are among the 17 with largest average coupling.

What has been discussed in this section has to be seen as a preliminary work in order to understand the influence of the topological characteristics of the couplings




Figure 5.15: The same as in figure 5.14, but related to random regular graphs (left panel) and rings (right panel), again of size $N = 25$. The information contained in the highest ranked couple is clear in both cases.

on the results of the different heuristics. As a perspective for future work, one would like to determine at least leading order analytical estimates of the Hessian (or of its diagonal) for several geometries; this would be of interest also as a way of getting some clues about which heuristics could work the best on the different topologies considered.

5.9 Future perspectives

5.9.1 Biological applications

As discussed in the general introduction of this chapter, the reconstruction of a coupling matrix from a partial knowledge of the correlations is a theme of great interest in several fields. Before imagining a real application, however, one would need a very abrupt improvement of the performances of the proposed heuristics, as we have been able so far to study systems composed by some tens of nodes, whereas a real network could comprehend thousands (for instance for a gene regulatory network) if not millions (for a recommendation system) of nodes. Some of the strategies, as the ones using the full Hessian, scale very badly with an increasing system sizes both in the memory needed and in the computational time requested for the algorithm to converge. For very large systems, one would probably need even more approximated heuristics based on the fundamental ideas discussed in this work. Just to make an example, one could exploit the fact that the global inference of a 10^6 by 10^6 coupling matrix will remain virtually unchanged by any single extra measurement: in such conditions, one could for instance try and locally update only the part of the

inference that has been touched by the last measurement taken, without running a global gradient descent on all the elements not measured yet.

For larger systems that are of interest for real-world applications, also the point of allowing multiple measurements at the same time becomes crucial. As briefly discussed, the methods described in this chapter are quite straightforwardly generalisable to such a situation; the goodness of the results on different geometries and by using different heuristics has however still to be carefully evaluated.

5.9.2 Theoretical challenges

The core of the project on which the heuristics proposed have been based is the analytical analysis of the problem described in the first sections of this chapter. Many theoretical questions have still to be answered. Firstly and expectedly, it has been verified that the coupling topology is crucial in determining the performances of the different strategies, even if some common features have been found and discussed. In order to better understand this point, a spectral analysis of different geometries could be tried so to find out more precisely which are the elements the most important in differentiating the results. Any improvement in this theoretical understanding of the problem could give very useful hints about how to define efficient heuristics on a given network being characterised by a set of features. Among the latter, the ones suspected to be the more influential in determining the effectiveness of the different procedures are the homogeneity or the non homogeneity in the degree distribution, the loopiness, the sparseness, the typical distance between any couple of nodes. All of these could also interact in a non trivial way, and determining a priori how a strategy will work on a given system is not at all expected to be an easy task.

Another quite separate part of the project that is still missing is to start and take into account the possibility of having noisy measurements, with $\alpha \neq 0$. All the analysis so far performed will need to be repeated for different levels of noise, and some of the difficulties that will arise can be already forecast. As shown in the appendix 5.10.4, for instance, for a high enough level of noise the F will no more have a unique minimum, and the result of the gradient descent will typically depend on its starting point. Even knowing since now these complications, it is clear how crucial is the generalisation to $\alpha \neq 0$, as the hypothesis of having at our disposal perfectly accurate measurements is clearly an absurd in many real-world contexts such as the biological one.

5.9.3 More complicated heuristics

It has been seen how a richer initial condition typically leads to much better results. This insight can be turned the other way round by saying that if the initial knowledge on the system is not rich enough, the best one can do is to perform some measurements designed in such a way to obtain a first understanding of the structure of the system and just in a second moment starting to use the real strategies discussed. Depending on the features of the system under study, an optimal subdivision between the measurements to be taken during this initial transient and the ones following the real heuristics could be determined.

Another aspect to go into in more depth is the effect of the regularisation parameter γ ; on this aspect some preliminary work has been done and the results referring to it are reported in appendix 5.10.3. In addition to those, it is reasonable to imagine that the more information on the system we have, the smaller a regularisation is needed for the algorithms to converge. A way of practically implementing such a principle would be to perform an *annealing* on γ , that is decreasing its value as more and more measurements are taken on the system. Some promising results have been obtained in this direction, but a comprehensive understanding of the effects of such an annealing on the inference quality of the different heuristics is still missing.

5.10 Appendices

5.10.1 Caveat on the procedure for creating synthetic J

The point one has to consider when using the strategy discussed in section 5.4.3 is that a rescaling of the values of J is performed, and this latter is not guaranteed to be uniform along the elements of the matrix. The procedure needed to obtain a $C = J^{-1}$ with the characteristics specified above will produce a final J_{true} with elements potentially very different from each other even if the starting J had all the elements equal. This is the case, for instance, on a severely inhomogeneous network as a Erdős–Rényi graph where a very strong variability among the couplings is introduced, whereas on a network with a higher symmetry such as a ring this effect is not observed.

This makes difficult the comparison among different geometries, as in some sense one loses the control on the actual value of the couplings. Even starting from analogous conditions, it has been verified that on some geometries one ends up with couplings much larger than the ones observed on others, and hence with an inference




Figure 5.16: Value of the average of the couplings (left panel) and of the difference between the largest and the smallest element of J (right panel) after the rescaling procedure. The reported results refer to single instance of the different topologies of different sizes N .

task that has to be expected easier. As a future perspective, one would like to test the possible use of an inverse problem also at this stage, for instance by appropriately assign the value of J from which the procedure described has to start in order to obtain similar J_{true} for different geometries.

In order to quantitatively estimate the impact of this issue, some analyses on single instances have been performed; in all the cases that will be discussed, the starting point has been set as a network with all the couplings equal to 0.7 in modulus, with the sign randomly assigned with probability 0.5. All this study has to be considered preliminary as more stable results will be obtained by considering ensemble averages for the different types of geometry. The more natural observable one can look at is the average value of J after the rescaling and at the range between the minimum and the maximum value. The results as a function of the system size N can be seen respectively in the left and in the right panel of figure 5.16.

As regards the ring, for any size the elements of J are very homogeneous and very similar to the originally assigned value of 0.7. For random regular graphs and two dimensional lattices, the situation appears to be quite similar. In both cases, the average value of the couplings is smaller than the input one, possibly affecting the performances of the algorithms as, all the other factors being constant, the inference is going to be more difficult if the difference between contacts and non-contacts is smaller. This worsening should be true in particular for the lattice, as the obtained average coupling value is systematically smaller if compared to the other geometries. The rescaling procedure, however, does not introduce a strong inhomogeneity for these geometry, probably because of their intrinsic symmetry; as can be seen in the




Figure 5.17: Value of the smallest (left panel) and of the largest (right panel) coupling after the rescaling procedure. For both, the absolute values are considered. The reported results refer to single instance of the different topologies of different sizes N .

right panel, the deviation between the largest and the smallest element is still quite small.

As expected, the Erdős–Rényi case shows some qualitatively different features. Looking at the average value, it looks like the couplings are even bigger than the ones of the already discussed geometries. However, one should not be too confident in the algorithm performances, as the couplings are for this topology much more diverse as can be seen by looking at the difference between the smallest and the largest element of J , bigger than the one found for all the other geometries by almost one order of magnitude.

The hierarchical structures, defined only for $N = 16$ and $N = 33$, seem to produce results intermediate between the two extremes so far discussed as regards the difference between minimum and maximum coupling. Especially for the larger system the average of J is much smaller than the original 0.7 value and one should not expect a very easy inference task.

Other interesting insights can be obtained by looking at the maximum and minimum absolute values of the couplings after the rescaling (see figure 5.17). As regards the maximum value, one can see that all the geometries apart from the Erdős–Rényi graphs behave similarly and have a maximum J not too far away from the input value of 0.7. The Erdős–Rényi case is again quite different, as its largest coupling is found to be around 3. This feature turns out to be quite central in analysing some convergence issues that emerged in the simulations. It has been seen, indeed, that a trade-off between the easiness of the problem from a statistical physics and from a numerical point of view has to be looked for. For chosen couplings that are too

small, the difference between *couplings* and *non-couplings* becomes tinier and tinier, and the former are detectable with more and more difficulty. If too large couplings are chosen, on the other hand, the gradient descent procedure is very unstable. In particular, by measuring an extra element there is a non-negligible probability that the technique that will be described in section 5.7.2 for updating the estimate of C will produce a matrix not invertible any more. In order to avoid convergence issues in such a case, one should set very small learning rates, leading in turn to an extremely slow convergence even for very small networks. The fact that in Erdős–Rényi graphs one indirectly obtains such extremely large couplings via the rescaling procedure can be regarded as one of the reasons of the observed convergence lack.

Also the minimum absolute value of J (see left panel in figure 5.17) gives us some useful insights as it can be connected to the detectability of the weakest contacts. The ring appears to be the easiest case also in this respect, all its couplings being practically identical with each other. Going towards more difficult cases, one found respectively random regular graphs, lattices and hierarchical structures. The most difficult geometry, also in this case, is the Erdős–Rényi, its smallest J being between 0.2 and 0.3. For a $N = 49$ network, for instance, this means that we should be able to detect couplings that are one third of the corresponding ones on the ring: it is not unexpected, then, that the performances on random graphs will be worse, all the other parameter being equal.

These differences can lead to very different results, especially as regards the ability of identifying where the true contacts are. In this case, indeed, all of them are considered in a sense on the same level, independently on the magnitude of the coupling. If some of them, nevertheless, are much smaller than the others, they will be much more difficult to identify; in an extreme case, a “true” contact whose coupling is exactly 0 is by definition impossible to detect. Such paradoxical effects should not be regarded as signs of poor performances of the algorithms.

5.10.2 The value of the couplings

Being interested in recovering a set of couplings, the first point to be fixed is whether these latter should be equal to each other. Since this possible symmetry is anyway broken by the procedure described in 5.4.3, the couplings will for the sake of generality not be assumed to be equal. Another interesting point is how to decide the sign of the couplings. In the first part of the project, the simplest case where all the couplings are positive has been considered. Especially when thinking about possible applications, however, it is important to be able and deal with situations in

which positive and negative couplings coexist so that some couples are energetically favoured to be aligned, and some are not. This is far from being a technicality, as this choice let the problem enter in the field of the disordered systems discussed in chapter 1. The reason why this case is expected to be harder than the first one is that the balancing effects between positive and negative couplings will lead to smaller effective correlations between sites, diminishing the information content of the taken measurements on C .

Since for the time being we are not yet thinking about a specific application, the couplings are extracted from a normal distribution with a given mean and standard deviation (small enough so to avoid too spread values), and afterwards a randomly chosen half of them is assigned a negative sign. Another possibility having been examined is to have bigger couplings on the selected geometry and smaller, but different from zero, couplings elsewhere. As one is typically more interested in coupling matrices that are really sparse, however, this possible generalisation has been considered not to be really enriching and has been neglected in the simulations.

5.10.3 Regularisation and noise

Considering the Bayesian approach to the problem proposed in section 5.3, before actually starting to run simulations and look at the results one needs to define the value of the regularisation parameter and the noise level of the system by assigning a value to the parameters previously called γ and α . As regards this latter, given that this work is quite preliminary in its approach and no established results were found in the literature, we restricted to the noiseless case where $\alpha = 0$; after having better understood this simpler case dealing with perfectly accurate measurements on the system, anyway, one would like to extend the analysis performed also to noisy case. This has to be considered as a future extension of this project.

The possible values of γ have been, on the other hand, more carefully studied. It has been found out that this regularisation on the model has to be neither too small nor too big. In the former case, the strategies start to show convergence issues, as a large enough regularisation on the matrices is needed in order to be able and invert them. If γ is too large, on the other hand, one is throwing away some information about the system, as the flat prior on the system given by the regularisation starts to take more and more importance in the inference procedure compared to the results of the measurements actually performed on the system. As can be seen in figure 5.18, however, it looks like a large enough range of possible γ values exists such that both these problems are avoided. In particular, the curves of the error on J_{inf}




Figure 5.18: For decreasing values of the regularisation γ , one is able and reconstruct more and more reliably the actual values of the J matrix. Even for a large enough value of γ , however, one is able to correctly determine the couples connected by a coupling different from zero, as can be seen by looking at the true positive rates on the right panel, very stable with respect to different choice of the regularisation parameter. Both the figures refer to a combination of geometry and heuristics producing good inference results, so to be able and more clearly pinpoint the feature under discussion. The geometry chosen is a $N = 25$ ring, and the elements to measure are chosen according to their value on the diagonal part of the Hessian matrix: this strategy will be introduced in section 5.6 and will be referred to as LHD, standing for *low Hessian diagonal*. The curves are averages over 5 instances of the network, on each having in turns performed 5 runs of the procedure.

for γ in this range appear to be simply rescaled; errors on the reconstruction of J smaller than the ones reported in the figure get however more and more difficult to obtain as one would need smaller and smaller γ , incurring in the convergence issues previously introduced. The true positive rates, on the other hand, seem to be practically independent on the choice of a regularisation in this range. This needed trade-off has anyway to be considered as a possible explanations for the bad results obtained under some conditions: especially on Erdős–Rényi graphs, where some of the couplings get very large because of the mechanism discussed in section 5.10.1, from the numerical studies it looks like the minimal γ for which the procedures converge is still too large to obtain reliable inferences.

5.10.4 Uniqueness of the minimum

As discussed in eqs. 5.8 and 5.9, the solution of the inference problem is found by looking at the model minimising the value of F . In order for this idea to be more rigorously defined and independent on the initial knowledge of the system, one would like to check under which conditions such a minimum is uniquely defined.

This task is achieved by looking at the quadratic form of the Hessian in the compact form found in eq. 5.18 defined as follows:

$$Q(\{X_{m_1, m_2}\}) \equiv \sum_{m_1, m_2; m'_1, m'_2} X_{m_1, m_2} M_{m_1, m_2; m'_1, m'_2} X_{m'_1, m'_2} \quad (5.24)$$

The uniqueness of the minimum will hence be verified for the parameters for which such a quadratic form is positive defined. After some calculations involving the quadratic form of the different components of M as they are defined in 5.18, one gets the following relation:

$$Q(\{X_{m_1, m_2}\}) = \sum_{m_1, m_2} (X_{m_1, m_2})^2 \left[\frac{1}{\gamma} (I_{m_1, m_2} - 1) + \frac{2(1 - \alpha)}{\lambda_{m_1} \lambda_{m_2}} \right] \quad (5.25)$$

First of all, one can study the small γ limit. The result one obtains in this case is the following:

$$\lim_{\gamma \rightarrow 0} Q(\{m_1, m_2\}) = \sum_{m_1, m_2} X_{m_1, m_2}^2 \left(-\frac{2\alpha}{\lambda_{m_1} \lambda_{m_2}} \right) \quad (5.26)$$

This limit is hence unfortunately quite uninformative. For $\gamma \rightarrow 0$, the Hessian is found to be negative definite for any value of $\alpha > 0$ and completely flat for $\alpha = 0$. This is consistent with the fact that, in order to run the gradient descent discussed above, one has to put a regularization parameter γ strictly larger than 0.

A direct study of the formulas for generic γ is hence needed. In order to obtain an even more compact form of the Hessian, one can call G the term multiplying X_{m_1, m_2} in eq. 5.25. By doing that, one obtains the following formally very simple relation for the quadratic form Q :

$$Q(m_1, m_2) \equiv \sum_{m_1, m_2} X_{m_1, m_2}^2 G(\lambda_{m_1}, \lambda_{m_2}, \gamma, \alpha) \quad (5.27)$$

The uniqueness of the minimum is therefore directly related to the positivity of the function G . In order to study how the positivity is connected to the value of the parameters, G is separately looked at as a function of γ for several values of α and vice versa. The values of the two eigenvalues λ_{m_1} and λ_{m_2} are in this case arbitrarily fixed, as the goal is to understand the general features of the curves. However, the true values of the eigenvalues could be inserted in the relations so to obtain the correct positivity thresholds in some specific cases of interest.

As one can see in the left panel of figure 5.19, the behaviour of G as a function of α for increasing γ appears to be very reasonable. In the limit of perfect sampling $\alpha = 0$,




Figure 5.19: On the left panel: curves of G as a function of γ for several values of the noise level α . On the right panel: G as a function of α for several regularisation parameters γ .

indeed, G is always positive for any $\gamma \geq 0$ (i.e. even in absence of regularisation). For larger α , on the other hand, it can be found a value of $\gamma_c(\alpha)$ such that G is positive for $\gamma > \gamma_c$ and negative otherwise. As expected, $\gamma_c(\alpha)$ is an increasing function of α ; the noisier the measurements, the stronger we need a regularisation to obtain a positive definite Hessian and to have a well-defined minimisation problem as regards the original function F .

By examining the dependency of G on the noise α , it turns out that this latter is linear for any value of γ . Indeed, by defining $A \equiv -\frac{2}{\lambda_{m_1} \lambda_{m_2}}$ and $B \equiv \frac{1}{\gamma} (I(\lambda_{m_1}, \lambda_{m_2}, \gamma) - 1) + \frac{2}{\lambda_{m_1} \lambda_{m_2}}$ one can write $G(\alpha) = A\alpha + B$.

Also in this case, a critical α_c in correspondence of which the quadratic form is equal to zero can be found for any choice of the regularisation γ . As one can see from the right panel of figure 5.19, such a value is increasing with γ . This feature can be intuitively understood, as by choosing larger and larger regularisations, one is able to study noisier and noisier systems.

5.10.5 Completion of the initial condition

In order to run the gradient descent, the initial knowledge \tilde{C} one has on the system has to be completed in some way, by assigning values also to the elements still belonging to C_{\perp} . The most simple idea we checked was to assign all the missing elements to zero, and to recover a more appropriate estimate of them as a result of the optimisation procedure. After some thoughts, however, one sees that this is actually a very strong assumption, creating some major convergence issues. For better understanding the point, let us imagine that two correlation measurements C_{ij} and C_{jl} have been taken, and that both of them were found to be strong. Even if C_{il} has not been measured yet, setting it equal to 0 is clearly straining, as the most

natural hypothesis one can make in this case is that also this element will be quite strong. From a more technical point of view, it turns out that the matrix obtained by setting equal to zero the unknown elements is not at all guaranteed to be positive defined, and this prevents the gradient descent to correctly run as the matrix cannot be inverted. Even in the parameter regions where the minimum of F is unique and the outcome of the gradient descent does not depend on the initial condition, this latter has nevertheless to be at least invertible for the procedure to work.

Following the idea discussed above, another way of completing \tilde{C} so to obtain an invertible C is introduced as follows. Let us focus first of all on the case where the initial knowledge on the correlation is star-like around the i^{th} node, in which the proposed procedure can be shown to give an invertible matrix. Without loss of generality, let us fix $i = 1$, imagining that the first row and column of C have been measured. In order to overcome the previously discussed difficulties, a reasonable way of guessing one of the missing correlations $C_{i,j}$ with $i, j \geq 2, i \neq j$ is to set $C_{i,j} = C_{1,i} C_{1,j}$. If the node 1 is strongly correlated both with i and with j , for instance, setting $C_{ij} = 0$ requests clearly a stronger hypothesis than setting it to a larger value. If on the other hand at least one of the two correlations $C_{1,i}$ or $C_{1,j}$ is weak, the prior guess about $C_{i,j}$ will be a weak value as well.

The elements of C are thus defined as follows:

$$\begin{cases} C_{i,i} = 1 & \forall i \in [1, N] \\ C_{i,1} = C_{1,i} \equiv C_i & \forall i \in [2, N] \\ C_{i,j} = C_{j,i} = C_i C_j & \forall i, j \in [2, N], i \neq j \end{cases} \quad \begin{array}{l} \text{: measured correlations} \\ \text{: unknown correlations} \end{array} \quad (5.28)$$

In order to show the matrix built in such a way to be invertible, one can demonstrate the quadratic form associated to it to be always positive. After some calculations and substituting the matrix elements one is able to obtain the following:

$$\begin{aligned} \sum_{j,k=1}^N x_j C_{j,k} x_k &= \\ &= \sum_{j=1}^N x_j C_{j,j} x_j + \sum_{j=2}^N x_j C_{j,1} x_1 + \sum_{k=2}^N x_1 C_{1,k} x_k + \sum_{\substack{j,k \geq 2 \\ j \neq k}} x_j C_{j,k} x_k = \dots = \\ &= \sum_{j=2}^N x_j^2 (1 - C_j^2) + \left(x_1 + \sum_{j=2}^N C_j x_j \right)^2 \geq 0 \quad (5.29) \end{aligned}$$

which is indeed positive for an arbitrary choice of x as $|C_j| \leq 1$. A possible generalisation of the previously described procedure of completion of the initial C is

proposed as follows, even if in the most general case this is no more guaranteed to generate a positive definite matrix.

The generalisation goes as follows: when the initial knowledge on C is a star around a node i , we discussed in the previous section that a sensible choice for setting a first estimate of the element C_{kl} , $k, l \neq i$ is $C_{kl} = C_{ki} \cdot C_{il}$. A possible interpretation of this choice makes use of the fact that, considering the graph induced by the initial knowledge on the N nodes graph, all the couples of nodes k, l with $k, l \neq i$ are connected by a unique path, that is the one that passes through i . The shortest (and unique) path from k to l is then formed by the two links (k, i) and (i, l) , and hence the correlation between k and l can be evaluated using this path according to eq. 5.28.

This reasoning can immediately be extended to a tree-like initial knowledge on C : also in this case, indeed, there is a unique path connecting any pair of nodes, even if it can be longer than in the previous case. However, if the path connecting nodes k and l passes through the nodes $k_1, k_2, \dots, k_{p-1}, k_p$ the estimate of $C_{k,l}$ will be defined as $C_{k,l} = C_{k,k_1} \cdot C_{k_1,k_2} \cdot \dots \cdot C_{k_{p-1},k_p} \cdot C_{k_p,l}$.

Even in a more general case, in which the initially known elements of C do not form a tree over the original graph, this method can be applied. In this case, however, there could be multiple paths connecting a pair of nodes: the missing C can be hence estimated by applying the above described procedure on the shortest one. In order to take into account also the case of an initial knowledge whose elements do not form a connected graph on the original system, the estimate for a $C_{j,l}$ with j and l belonging to disconnected components of the graphs is set equal to zero. This is also consistent with the idea above discussed of setting all the initial estimates equal to zero in absence of a prior knowledge of the system.

Chapter 6

Conclusions and perspectives

6.1 The network theory in the “big data” era

6.1.1 The end of theory?

In a provocative essay [155] published in 2008 Chris Anderson claims that, because of the deluge of data at our disposal (and of the increasing rate at which we are able to collect and analyse them), our need for models will irreversibly shrink during the next few years. When the accessible sets of data were small, a model for generalising them was fundamental, and the disentanglement between correlations and causations was an aspect of particular importance. This could be less evident when, from a practical point of view, *all* the data are accessible, as in this case he claims that knowing the correlations may be enough. He reports Google’s research director Peter Norvig paraphrasing the famous George Box maxim, saying that “All models are wrong, and increasingly you can succeed without them.” The example of what Google did in the online advertising market is enlightening: without pretending to know its mechanisms, conventions and culture, Google’s researchers and engineers were nevertheless able to conquer it by using the best analytical tools on the best and richest available data. No specific model was involved, no causal effect was claimed to have been rigorously determined. According to Anderson no “science”, in a traditional sense, was performed.

I believe this attitude to be way too extreme for several reasons. First of all, even if it is true that our ability to analyse data has exponentially (literally, according to Moore’s law) increased over the last decades, the same holds also for the quantity of collected data. In a sense, it is to be expected that also in the next years we will be forced to analyse just a small fraction of the data we will have, and that modelling will therefore keep its central place in science. Looking at one of the first truly *big data*

fields of research, the astronomy, one can see that the Large Synoptic Survey Telescope (LSST), planned to enter operation in 2022, will gather 30 TB a night [156, 157]. The same order of magnitude of data will be collected by the largest particle physics experiments, such as the ones running at the LHC [158]. For the time being, just the data that seem at a first sight to be the most informative are actually analysed, all the rest being either stored (waiting for increased analytical abilities) or directly discarded.

Apart from this technical inability, that as far as we know could even be overcome in the future, there is probably also a somewhat more fundamental aspect at play. I am indeed convinced that we will keep on feeling the necessity of having more simple and “understandable” models. Trying a risky metaphor, I would say that the modelling capabilities in the *big data* era will be as important as the Newtonian laws of motions after the discovery of the quantum mechanics and of the relativity theory. The frameworks discussed by C. Anderson, hence, will probably survive as two complementary approaches, more than being in a real juxtaposition.

Most of the problems described in this thesis will be central also in the growing field of the *data science*, even if they will probably be seen under a slightly different perspective. A first difference with what has been discussed will be quantitative: the solutions needed to deal with *big data* are already, and will be for the next years, of the greatest interest for the computer science and for the engineering communities [159]. More strictly connected with the modelling approach kept in this work will however be their qualitative features (that are sometimes referred to by talking about *rich data* [160]), that will request a lot of theoretical work to try and better understand the underlying mechanisms; the following example will clarify the point.

6.1.2 An example of “rich data” analysis: Quora

Let us consider the social network Quora [161], which main idea is to let the users ask more or less any kind of question, and to let other ones (experts or not in the field) propose their answers. Put this way, the situation could seem similar to the matching problem discussed in chapter 2, the objective being to build a mapping between the questions on the one hand, and the people able to answer them on the other. The first issue comes when one quantitatively analyses the situation, as there are several millions of users (and, therefore, of questions and answers proposed) and the efficiency of the proposed algorithm is therefore really crucial.

A fundamental part of the data scientists’ job at this company is to maintain high the signal over noise ratio, in order to keep the users interested; to do so, the “quality”

of each user both in asking and in answering questions has to be assessed. A system of up and down votes makes this possible by exploiting an iterative procedure similar to the PageRank used by the Google research engine; the quality of a user has to be understood as a dynamical quantity, as it is obviously not fixed once and for all for a user.

As regards the people aiming to answer questions, the system should propose good ones to them: question they can and they are willing to answer. Assigning a set of “interesting” questions to a user is a huge problem on its own: it depends on constraints both global (the topics considered interesting by the whole community in a given day) and local (what are my expertises?). Even grouping questions posed in a free form is a big problem, as it requires cutting-edge *natural language processing* capabilities applied to sort of a community detection problem as the ones discussed in chapter 2.

As usual in social networks, also in this case there is the possibility to *follow* other users, so to preferentially see the questions they propose and the answers they give; one can also *endorse* another user by confirming his competences in a field. The well known exploration-exploitation phenomenon discussed in chapter 5 is also a crucial point: if I say to be fond of music, for instance, the system should propose me several contents related to it in order to *exploit* the piece of information I have given to the system; nevertheless, to avoid the boredom that naturally would arise if the questions proposed to me are too homogeneous, it should also be able to change the topic now and then, guessing another field which I am likely to appreciate so to *explore* my personal tastes.

What at a first sight could have seemed just a standard, computer science textbook matching problem, turned thus out to be both quantitatively and qualitatively extremely involved. In the following section, it will be shown how also the problems tackled in chapters 3, 4 and 5 will still be of central importance in the next years (and probably even decades). In conclusion, it seems to me that we are still quite far away from “the end of theory”; probably, rather, a big issue will be how to try and unify the insights one can have by for instance making use of a machine learning approach on the one hand, and of a more theoretical, modelling framework on the other. This difficulty, however, has already been faced in other fields such as the biology, leading to the coexistence of several levels of abstraction; the same natural processes can indeed be described both by very abstract, mathematical models and by more precise, quantitatively accurate, data-oriented approaches.

6.1.3 New solutions to old questions

The introduced “data deluge” is going to let the traditional ways of solving a large number of problems rapidly become obsolete. Let us think about the clustering problem introduced in chapters 2 and more extensively discussed as regards the recommendation systems in chapter 5, having a clear application in the search for an accurate *market segmentation*. This latter is a long-standing goal, as it is needed in order to better understand the common features of the groups of people taking advantage of a given service.

The standard strategies adopted in marketing to address this task up to some years ago (and still now in some cases) made use of some coarse graining of the customers according to very basic and structured features such as their age, sex, level of instruction or city of residence. In a “data-rich” era, this kind of classification will appear extremely rough, as for instance any web service (from the ones for reserving hotel rooms to the ones for buying books or clothes) is already making personalised suggestions based on our previous buying history or on our activity on the social networks. Even the preferences therein expressed by some friends of ours can be profitably taken into account, as anybody is more likely to buy a product that his acquaintances have already shown to like, as already known since many years in the sociology community [51].

This switching from a regime in which one has to do the best he can with the few data available to another in which, on the other hand, the main question is which part out of the immense quantity of data continuously collected one is going to exploit, and according to which principles, has already been observed in some fields and is to be expected in many others. The heterogeneity of data will also be a very interesting aspect coming into play. Even in a historically data-rich field such as finance, the data to analyse have always been quite homogeneous and easy to assemble; this is not going to be true any more when for instance messages written on Twitter in a natural-language, unstructured form and regarding a company are going to be used as one of the factors determining its stock market value.

In other sectors, such as the sociology, the data were up to some years ago very hard and expensive to collect, as this operations required for instance surveys among the population. The dramatic fall in the cost and time needed for obtaining such information has already affected these fields, quite abruptly shifting their focus from a mainly qualitative to a much more quantitative perspective.

6.1.4 The future of network theory

When trying to forecast the evolution of network theory during the next years, one has to take into account at least three practical differences that will drastically modify the way it was approached in the past [76]. Firstly, the data acquisition is going to become more and more automatised, and very large databases regarding the topologies of several real-world networks will become almost effortless available. This dynamics has been already at least partially observed in economics and in biology, but many other fields are going to join this ever-increasing “datafication” [162]. Secondly, the computing power one will have access to will continue to increase; even if the “intractable” combinatorial problems discussed in chapter 2 will remain so up to a major revolution (the most probable seeming to be at the moment the reliable availability of quantum information technologies), larger and larger systems will be analysable, and better and better approximations will be obtained in a reasonable amount of time. Last but not least, the already started breakdown of boundaries among disciplines will very likely go on; the possibility of looking at features shared by very different phenomena and at other ones seeming to be, on the contrary, specific of a given field or process will surely shed light on the basic mechanisms regulating these systems.

6.2 More data, same problems

6.2.1 Optimisation of spreading on networks

Especially in the introductions of the previous chapters, the general interest of the problems faced during my Ph.D. has been discussed, even if in a way far from having been exhaustive. As a final conclusion of the work, I would like to briefly describe the importance such themes will still have in a world that will be change as forecast in the previous section 6.1.

In chapter 3, the problem of how to induce and control optimal spreading on networks has been extensively discussed. This optimality is going to be (and, partly, it already is) a very crucial task to realise in the practical implementation of the so-called *smart cities*. In order to minimise wastes, diminish the costs for the population, avoid cascade failures and not overcharge the city infrastructures, the flows on the energy or on the water supply grids will be automatically modified using data coming from smart meters that will continuously transmit data to be analysed and exploited. The very same will hold for the traffic regulation in which, recalling the discussion of

chapter 1, both the inverse and the direct problems will be of interest: designing new routes so to minimize the risk of traffic jam and handling extreme events in real time will be examples of the two typologies respectively.

The problem of optimal spreading information in a complex network will very likely remain of central interest also in the IT area, especially by considering the larger and larger importance taken by the study and the application of distributed computing systems [163]. These latter, composed by several autonomous computational entities each of which having its own local memory, have to communicate among each other and to coordinate their actions in order to achieve a common goal. It is therefore clear how the design of efficient methods to optimally transfer across the entities pieces of information needed for accomplishing such a coordination is going to be crucial for these systems to improve their performances. When dealing with very large networks of computing units, moreover, one will also need to take care of the possible failure of some of the nodes; again, the discussion of chapter 3 on cascade failures and on how to possible predict and avoid them is going to be of straightforward application to this field.

6.2.2 Extreme events in network exploration

In chapter 4 the exploration of graphs, and in particular the quantitative determination of how improbable extreme events are in these processes, has been discussed. The applications therein described will continue to be of interest, and a number of new ones will likely emerge as a result of the “datafication” process taking place in many fields. As soon as new or larger datasets concerning networks become available, the question on how it is possible to *navigate* through them by following their links very naturally arises.

A topic of particular relevance will emerge from a mixture of the problems having been faced, in this work, in chapters 4 and 5. Many systems one will want to explore, indeed, will be very rapidly changing over time, and therefore a complete knowledge of them at any given moment will be quite an unrealistic hypothesis. Google research engine, for instance, will need to explore a network whose size is becoming more and more impressive as time passes by. According to the data referring to 2013 reported in [164], in particular, every minute more than five hundred websites are created and therefore join the previously existing nodes. In this framework, one will need to determine algorithms and procedures able to optimally explore incredibly large networks which are just partially known, and for which the information one has at

his disposal are not completely reliable and are dramatically changing over quite small time scales.

6.2.3 Inferring networks with partial information

Also as regards the inference of a network of interactions from a sparse or unreliable knowledge about the correlations among its components, some of the reasons of interest in applicative fields have been described in chapter 5, for instance the importance this problem has in order to optimally suggest items to a user based on his tastes by making use of the recommendation systems. These problems were however thought to take place on a given dataset with a certain number of either noisy or even completely missing entries. The challenge one will have to face in the future is that one will need in some sense to play on different fields at the same time. This aspect is again connected to the increased richness of the data one will have to work with in the next years.

Referring again to the data collected by Quora, one sees how they refer to several distinct areas such as the topics of interests, the ratings given to other users and received by them or the number of days or hours elapsed on average between two connections to the website. All of these datasets are typically only sparsely known, but in principle one would like to exploit the information contained in any of them in order to improve the knowledge we have about any other; in some sense, one would therefore like to perform crossed inferences referring to different datasets. These latter are moreover highly heterogeneous, as some of them are structured and some are not, and the variables they store can be real, categorical or boolean. Achieving a theoretically rigorous comprehension on how to optimally perform this kind of crossed inferences is going to be probably out of reach for many years; one could expect however that the insights having been obtained in simplified frameworks may help in designing better heuristics so to improve the results obtainable in real-world, extremely involved cases.

The reason why these studies will likely remain of interest is that such systems are going to become more and more ubiquitous during the next few years. NetFlix was probably one of the first companies in which a more scientific approach to this problem was tried; nowadays, however, recommendation systems are implemented in various form by many different companies. Amazon, since many years suggesting our next book to buy depending for instance on which books we have already looked at or bought and on the ratings we have given to other items; companies offering musical streaming services such as Deezer or Spotify, interested in proposing us a

set of songs that we perhaps do not know yet, but that we are likely to appreciate; websites selling clothing such as Zalando, able to propose us a pair of shoes matching with the trousers we just bought; news sites regulating the feeds of news appearing when we open their homepages according to our interests and on the trending topics of the moment. All these companies have nowadays quite large “data science” groups and are developing and implementing cutting-edge machine learning techniques to improve their results in this crucial area.

These systems will very likely become also useful for improving current web search engines, that are going to get more and more personalised, by starting to increasingly take into account our previous search history and our similarity with other users. The pages that will be proposed to us will be in this case no more be the ones fitting the best, in an abstract way, to our research; rather, such systems will be able to consider also the typology of the pages we (and users similar to us) have most often opened in the past, and the results we are most likely going to appreciate are going to be the ones preferentially suggested to us.

References

- [1] Alberto Guggiola and Guilhem Semerjian. Minimal contagious sets in random regular graphs. *Journal of Statistical Physics*, 158(2):300–358, 2015.
- [2] Caterina De Bacco, Alberto Guggiola, Reimer Kühn, and Pierre Paga. Rare events statistics of random walks on networks: localization and other dynamical phase transitions. *arXiv preprint arXiv:1506.08436*, 2015.
- [3] Arturo Rosenblueth and Norbert Wiener. The role of models in science. *Philosophy of Science*, 12(4):316–321, 1945.
- [4] Jean-Philippe Bouchaud and Marc Potters. *Theory of financial risk and derivative pricing: from statistical physics to risk management*. Cambridge University Press, 2003.
- [5] George EP Box. Science and statistics. *Journal of the American Statistical Association*, 71(356):791–799, 1976.
- [6] George EP Box. Robustness in the strategy of scientific model building. *Robustness in statistics*, 1:201–236, 1979.
- [7] Emmanuelle Gouillart, Florent Krzakala, Marc Mezard, and Lenka Zdeborová. Belief-propagation reconstruction for discrete tomography. *Inverse Problems*, 29(3):035003, 2013.
- [8] Edwin T Jaynes. Information theory and statistical mechanics. *Physical review*, 106(4):620, 1957.
- [9] Leon Festinger, Kurt W Back, and Stanley Schachter. *Social pressures in informal groups: A study of human factors in housing*. Number 3. Stanford University Press, 1950.
- [10] Nicholas Metropolis, Arianna W Rosenbluth, Marshall N Rosenbluth, Augusta H Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953.
- [11] Shang-Keng Ma. Modern theory of critical phenomena. *Frontiers in Physics*.

- [12] Kenneth G Wilson. The renormalization group and critical phenomena. *Reviews of Modern Physics*, 55(3):583, 1983.
- [13] Sergey N Dorogovtsev, Alexander V Goltsev, and José FF Mendes. Critical phenomena in complex networks. *Reviews of Modern Physics*, 80(4):1275, 2008.
- [14] Alain Barrat, Marc Barthelemy, and Alessandro Vespignani. *Dynamical processes on complex networks*. Cambridge University Press, 2008.
- [15] David Sherrington and Scott Kirkpatrick. Solvable model of a spin-glass. *Physical Review Letters*, 35(26):1792, 1975.
- [16] Samuel Frederick Edwards and Phil W Anderson. Theory of spin glasses. *Journal of Physics F: Metal Physics*, 5(5):965, 1975.
- [17] Marc Mezard, Giorgio Parisi, and Miguel Angel Virasoro. Spin glass theory and beyond. 1987.
- [18] Marc Mezard and Thierry Mora. Constraint satisfaction problems and neural networks: A statistical physics perspective. *Journal of Physiology-Paris*, 103(1):107–113, 2009.
- [19] Federico Ricci-Tersenghi. The Bethe approximation for solving the inverse Ising problem: a comparison with other inference methods. *Journal of Statistical Mechanics: Theory and Experiment*, 2012(08):P08015, 2012.
- [20] Simona Cocco and Rémi Monasson. Adaptive cluster expansion for inferring Boltzmann machines with noisy data. *Physical review letters*, 106(9):090601, 2011.
- [21] Hilbert J. Kappen and FB Rodriguez. Efficient learning in Boltzmann machines using linear response theory. *Neural Computation*, 10(5):1137–1156, 1998.
- [22] Vitor Sessak and Rémi Monasson. Small-correlation expansions for the inverse Ising problem. *Journal of Physics A: Mathematical and Theoretical*, 42(5):055001, 2009.
- [23] Erik Aurell and Magnus Ekeberg. Inverse Ising inference using all the data. *Physical Review Letters*, 108(9):090201, 2012.
- [24] Marc Mezard and Andrea Montanari. *Information, physics, and computation*. Oxford University Press, 2009.
- [25] Cristopher Moore and Stephan Mertens. *The nature of computation*. Oxford University Press, 2011.

- [26] Stephen A Cook and David G Mitchell. Finding hard instances of the satisfiability problem. In *Satisfiability Problem: Theory and Applications: DIMACS Workshop, March 11-13, 1996*, volume 35, page 1. American Mathematical Soc., 1997.
- [27] Olivier C Martin, Rémi Monasson, and Riccardo Zecchina. Statistical mechanics methods and phase transitions in optimization problems. *Theoretical Computer Science*, 265(1):3–67, 2001.
- [28] Jean-Philippe Bouchaud. Economics needs a scientific revolution. *Nature*, 455(7217):1181–1181, 2008.
- [29] Hugo Touchette. The large deviation approach to statistical mechanics. *Physics Reports*, 478(1):1–69, 2009.
- [30] Claudio Castellano, Santo Fortunato, and Vittorio Loreto. Statistical physics of social dynamics. *Reviews of Modern Physics*, 81(2):591, 2009.
- [31] William Bialek, Andrea Cavagna, Irene Giardina, Thierry Mora, Edmondo Silvestri, Massimiliano Viale, and Aleksandra M Walczak. Statistical mechanics for natural flocks of birds. *Proceedings of the National Academy of Sciences*, 109(13):4786–4791, 2012.
- [32] Dirk Helbing. Traffic and related self-driven many-particle systems. *Reviews of Modern Physics*, 73(4):1067, 2001.
- [33] Thomas C Schelling. Dynamic models of segregation†. *Journal of Mathematical Sociology*, 1(2):143–186, 1971.
- [34] Luca Dall’Asta, Claudio Castellano, and Matteo Marsili. Statistical physics of the Schelling model of segregation. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(07):L07002, 2008.
- [35] Damien Challet, Matteo Marsili, and Riccardo Zecchina. Statistical mechanics of systems with heterogeneous agents: Minority games. *Physical Review Letters*, 84(8):1824, 2000.
- [36] John C Hull. *Options, futures, and other derivatives*. Pearson Education India, 2006.
- [37] John J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8):2554–2558, 1982.
- [38] Li Deng and Dong Yu. Deep learning: methods and applications. *Foundations and Trends in Signal Processing*, 7(3–4):197–387, 2014.

- [39] Martin Weigt, Robert A White, Hendrik Szurmant, James A Hoch, and Terence Hwa. Identification of direct residue contacts in protein–protein interaction by message passing. *Proceedings of the National Academy of Sciences*, 106(1):67–72, 2009.
- [40] Daniel Fisher, Michael Lässig, and Boris Shraiman. Evolutionary dynamics and statistical physics. *Journal of Statistical Mechanics: Theory and Experiment*, 2013(01):N01001, 2013.
- [41] John Burdon Sanderson Haldane. A mathematical theory of natural and artificial selection—i. *Bulletin of mathematical biology*, 52(1-2):209–240, 1990.
- [42] David L Donoho. Compressed sensing. *Information Theory, IEEE Transactions on*, 52(4):1289–1306, 2006.
- [43] F. Krzakala, M. Mézard, F. Sausset, Y. F. Sun, and L. Zdeborová. Statistical-physics-based reconstruction in compressed sensing. *Phys. Rev. X*, 2:021005, May 2012.
- [44] Tom Richardson and Ruediger Urbanke. *Modern coding theory*. Cambridge University Press, 2008.
- [45] Marc Mézard, Giorgio Parisi, and Riccardo Zecchina. Analytic and algorithmic solution of random satisfiability problems. *Science*, 297(5582):812–815, 2002.
- [46] Béla Bollobás. *Modern graph theory*, volume 184. Springer Science & Business Media, 1998.
- [47] Leonhard Euler. Solutio problematis ad geometriam situs pertinentis. *Commentarii Academiae Scientiarum Petropolitanae*, 8:128–140, 1741.
- [48] Chi Ho Yeung and David Saad. Networking—a statistical physics perspective. *Journal of Physics A: Mathematical and Theoretical*, 46(10):103001, 2013.
- [49] Mark Kac. Mathematical mechanisms of phase transitions. Technical report, Rockefeller Univ., New York, 1969.
- [50] Romualdo Pastor-Satorras and Alessandro Vespignani. *Evolution and structure of the Internet: A statistical physics approach*. Cambridge University Press, 2007.
- [51] Stanley Wasserman and Katherine Faust. *Social network analysis: Methods and applications*, volume 8. Cambridge University Press, 1994.
- [52] Uri Alon. *An introduction to systems biology: design principles of biological circuits*. CRC press, 2006.

- [53] Shai S Shen-Orr, Ron Milo, Shmoolik Mangan, and Uri Alon. Network motifs in the transcriptional regulation network of Escherichia Coli. *Nature Genetics*, 31(1):64–68, 2002.
- [54] P Erdős and A Rényi. On random graphs. *Publicationes Mathematicae Debrecen*, 6:290–297, 1959.
- [55] Jeffrey Travers and Stanley Milgram. An experimental study of the small world problem. *Sociometry*, pages 425–443, 1969.
- [56] Duncan J Watts and Steven H Strogatz. Collective dynamics of ‘small-world’ networks. *Nature*, 393(6684):440–442, 1998.
- [57] Alain Barrat and Martin Weigt. On the properties of small-world network models. *The European Physical Journal B-Condensed Matter and Complex Systems*, 13(3):547–560, 2000.
- [58] Mark EJ Newman. Power laws, Pareto distributions and Zipf’s law. *Contemporary Physics*, 46(5):323–351, 2005.
- [59] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.
- [60] Robert K Merton et al. The Matthew effect in science. *Science*, 159(3810):56–63, 1968.
- [61] Michael R Garey and David S Johnson. Computers and intractability: a guide to NP-completeness, 1979.
- [62] Anand Rajaraman and Jeffrey David Ullman. *Mining of massive datasets*. Cambridge University Press, 2011.
- [63] F. Altarelli, A. Braunstein, A. Ramezanzpour, and R. Zecchina. Stochastic Matching Problem. *Phys. Rev. Lett.*, 106:190601, May 2011.
- [64] Claudio Castellano, Daniele Vilone, and Alessandro Vespignani. Incomplete ordering of the voter model on small-world networks. *EPL (Europhysics Letters)*, 63(1):153, 2003.
- [65] Luca Dall’Asta, Andrea Baronchelli, Alain Barrat, and Vittorio Loreto. Agreement dynamics on small-world networks. *EPL (Europhysics Letters)*, 73(6):969, 2006.
- [66] Marc Barthélemy, Alain Barrat, Romualdo Pastor-Satorras, and Alessandro Vespignani. Velocity and hierarchical spread of epidemic outbreaks in scale-free networks. *Physical Review Letters*, 92(17):178701, 2004.

- [67] Lazaros K Gallos. Random walk and trapping processes on scale-free networks. *Physical Review E*, 70(4):046116, 2004.
- [68] Sergey N Dorogovtsev and Jose FF Mendes. Evolution of networks. *Advances in physics*, 51(4):1079–1187, 2002.
- [69] Réka Albert, Hawoong Jeong, and Albert-László Barabási. Error and attack tolerance of complex networks. *Nature*, 406(6794):378–382, 2000.
- [70] J Willard Gibbs. *Elementary principles in statistical mechanics*. Courier Corporation, 2014.
- [71] Simon R Broadbent and John M Hammersley. Percolation processes. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 53, pages 629–641. Cambridge Univ Press, 1957.
- [72] Eric Bonabeau. Sandpile dynamics on random graphs. *Journal of the Physical Society of Japan*, 64:327, 1995.
- [73] Juan A. Acebrón, L. L. Bonilla, Conrad J. Pérez Vicente, Félix Ritort, and Renato Spigler. The Kuramoto model: A simple paradigm for synchronization phenomena. *Rev. Mod. Phys.*, 77:137–185, Apr 2005.
- [74] Yoshiki Kuramoto. *Chemical oscillations, waves, and turbulence*, volume 19. Springer Science & Business Media, 2012.
- [75] Steven H Strogatz. From Kuramoto to Crawford: exploring the onset of synchronization in populations of coupled oscillators. *Physica D: Nonlinear Phenomena*, 143(1):1–20, 2000.
- [76] Réka Albert and Albert-László Barabási. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74(1):47, 2002.
- [77] Andrea Pelissetto and Ettore Vicari. Critical phenomena and renormalization-group theory. *Physics Reports*, 368(6):549–727, 2002.
- [78] M. E. J. Newman and M. Girvan. Finding and evaluating community structure in networks. *Phys. Rev. E*, 69:026113, Feb 2004.
- [79] Santo Fortunato. Community detection in graphs. *Physics Reports*, 486(3):75–174, 2010.
- [80] Andrea Lancichinetti, Santo Fortunato, and Filippo Radicchi. Benchmark graphs for testing community detection algorithms. *Physical review E*, 78(4):046110, 2008.

- [81] Aurelien Decelle, Florent Krzakala, Cristopher Moore, and Lenka Zdeborová. Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. *Physical Review E*, 84(6):066106, 2011.
- [82] Aurelien Decelle, Florent Krzakala, Cristopher Moore, and Lenka Zdeborová. Inference and phase transitions in the detection of modules in sparse networks. *Physical Review Letters*, 107(6):065701, 2011.
- [83] Florent Krzakala, Cristopher Moore, Elchanan Mossel, Joe Neeman, Allan Sly, Lenka Zdeborová, and Pan Zhang. Spectral redemption in clustering sparse networks. *Proceedings of the National Academy of Sciences*, 110(52):20935–20940, 2013.
- [84] Mark EJ Newman. The structure of scientific collaboration networks. *Proceedings of the National Academy of Sciences*, 98(2):404–409, 2001.
- [85] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The PageRank citation ranking: Bringing order to the web. 1999.
- [86] Ed Bullmore and Olaf Sporns. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*, 10(3):186–198, 2009.
- [87] Asu Ozdaglar and Ishai Menache. *Network Games: Theory, Models, and Dynamics*. Morgan & Claypool Publishers, 2011.
- [88] David Easley and Jon Kleinberg. *Networks, crowds, and markets: Reasoning about a highly connected world*. Cambridge University Press, 2010.
- [89] Adam DI Kramer, Jamie E Guillory, and Jeffrey T Hancock. Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, 111(24):8788–8790, 2014.
- [90] Pierre Paga and Reimer Kühn. Contagion in an interacting economy. *Journal of Statistical Mechanics: Theory and Experiment*, 2015(3):P03008, 2015.
- [91] Stefano Boccaletti, Vito Latora, Yamir Moreno, Martin Chavez, and D-U Hwang. Complex networks: Structure and dynamics. *Physics Reports*, 424(4):175–308, 2006.
- [92] John Von Neumann, Arthur W Burks, et al. Theory of self-reproducing automata. *IEEE Transactions on Neural Networks*, 5(1):3–14, 1966.
- [93] Stephen Wolfram. Statistical mechanics of cellular automata. *Reviews of Modern Physics*, 55(3):601, 1983.

- [94] Ciro Cattuto, Wouter Van den Broeck, Alain Barrat, Vittoria Colizza, Jean-François Pinton, and Alessandro Vespignani. Dynamics of person-to-person interactions from distributed rfid sensor networks. *PLOS ONE*, 5(7):e11596, 07 2010.
- [95] Martin Szomszor, Patty Kostkova, Ciro Cattuto, Wouter Van den Broeck, Alain Barrat, and Harith Alani. Providing enhanced social interaction services for industry exhibitors at large medical conferences. In *Proceedings of the 3rd International Conference on Developments in eSystems Engineering (DeSE 2010)*, 2010.
- [96] Alain Barrat, Ciro Cattuto, Vittoria Colizza, Lorenzo Isella, Caterina Rizzo, Alberto E. Tozzi, and Wouter Van den Broeck. Wearable sensor networks for measuring face-to-face contact patterns in healthcare settings. In Martin Szomszor and Patty Kostkova, editors, *Proceedings of the 3rd International ICST Conference on Electronic Healthcare for the 21st century (eHealth 2010)*, 2011.
- [97] Lorenzo Isella, Juliette Stehlé, Alain Barrat, Ciro Cattuto, Jean-François Pinton, and Wouter Van den Broeck. What’s in a crowd? Analysis of face-to-face behavioral networks. *Journal of Theoretical Biology*, 271(1):166–180, 2011.
- [98] Márton Karsai, Mikko Kivelä, Raj Kumar Pan, Kimmo Kaski, János Kertész, A-L Barabási, and Jari Saramäki. Small but slow world: How network topology and burstiness slow down spreading. *Physical Review E*, 83(2):025102, 2011.
- [99] Romualdo Pastor-Satorras and Alessandro Vespignani. Epidemic dynamics and endemic states in complex networks. *Physical Review E*, 63(6):066117, 2001.
- [100] Romualdo Pastor-Satorras and Alessandro Vespignani. Epidemic spreading in scale-free networks. *Phys. Rev. Lett.*, 86:3200–3203, Apr 2001.
- [101] Christopher Moore and Mark EJ Newman. Epidemics and percolation in small-world networks. *Physical Review E*, 61(5):5678, 2000.
- [102] Andrey Y. Lokhov, Marc Mézard, Hiroki Ohta, and Lenka Zdeborová. Inferring the origin of an epidemic with a dynamic message-passing algorithm. *Phys. Rev. E*, 90:012801, Jul 2014.
- [103] Fabrizio Altarelli, Alfredo Braunstein, Luca Dall’Asta, Alessandro Ingrosso, and Riccardo Zecchina. The patient-zero problem with noisy observations. *Journal of Statistical Mechanics: Theory and Experiment*, 2014(10):P10016, 2014.

- [104] JPL Hatchett and Reimer Kuehn. Credit contagion and credit risk. *Quantitative Finance*, 9(4):373–382, 2009.
- [105] Vittoria Colizza, Alain Barrat, Marc Barthélemy, and Alessandro Vespignani. The role of the airline transportation network in the prediction and predictability of global epidemics. *Proceedings of the National Academy of Sciences of the United States of America*, 103(7):2015–2020, 2006.
- [106] Mark EJ Newman. Spread of epidemic disease on networks. *Physical review E*, 66(1):016128, 2002.
- [107] Maksim Kitsak, Lazaros K Gallos, Shlomo Havlin, Fredrik Liljeros, Lev Muchnik, H Eugene Stanley, and Hernán A Makse. Identification of influential spreaders in complex networks. *Nature Physics*, 6(11):888–893, 2010.
- [108] Sheng Bau, Nicholas C Wormald, and Sanming Zhou. Decycling numbers of random regular graphs. *Random Structures & Algorithms*, 21(3-4):397–413, 2002.
- [109] John Chalupa, Paul L Leath, and Gary R Reich. Bootstrap percolation on a Bethe lattice. *Journal of Physics C: Solid State Physics*, 12(1):L31, 1979.
- [110] Andrey Y Lokhov, Marc Mézard, and Lenka Zdeborová. Dynamic message-passing equations for models with unidirectional dynamics. *Physical Review E*, 91(1):012811, 2015.
- [111] Marc Mézard and Giorgio Parisi. The Bethe lattice spin glass revisited. *The European Physical Journal B-Condensed Matter and Complex Systems*, 20(2):217–233, 2001.
- [112] Andrea Montanari, Federico Ricci-Tersenghi, and Guilhem Semerjian. Clusters of solutions and replica symmetry breaking in random k-satisfiability. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(04):P04004, 2008.
- [113] Lenka Zdeborova and Florent Krzakala. Phase transitions in the coloring of random graphs. *Physical Review E*, 76(3):031131, 2007.
- [114] Fabrizio Altarelli, Alfredo Braunstein, Luca Dall’Asta, and Riccardo Zecchina. Optimizing spread dynamics on graphs by message passing. *Journal of Statistical Mechanics: Theory and Experiment*, 2013(09):P09011, 2013.
- [115] Olivier Rivoire, Giulio Biroli, Olivier C Martin, and Marc Mézard. Glass models on Bethe lattices. *The European Physical Journal B-Condensed Matter and Complex Systems*, 37(1):55–78, 2004.

- [116] Lowell W Beineke and Robert C Vandell. Decycling graphs. *Journal of Graph Theory*, 25(1):59–77, 1997.
- [117] Fabrizio Altarelli, Alfredo Braunstein, Luca Dall’Asta, and Riccardo Zecchina. Large deviations of cascade processes on graphs. *Physical Review E*, 87(6):062115, 2013.
- [118] Marc Mézard and Riccardo Zecchina. Random k-satisfiability problem: From an analytic solution to an efficient algorithm. *Physical Review E*, 66(5):056126, 2002.
- [119] David L Applegate, Robert E Bixby, Vasek Chvatal, and William J Cook. *The Traveling Salesman Problem: A Computational Study*. Princeton university press, 2011.
- [120] Daniel J Rosenkrantz, Richard E Stearns, and Philip M Lewis, II. An analysis of several heuristics for the traveling salesman problem. *SIAM journal on computing*, 6(3):563–581, 1977.
- [121] Anirban Chakraborti and Bikas K Chakrabarti. Statistical physics of the Travelling Salesman Problem. *arXiv preprint cond-mat/0001069*, 2000.
- [122] Karl Pearson. The problem of the random walk. *Nature*, 72(1865):294, 1905.
- [123] Gregory F Lawler and Vlada Limic. *Random walk: a modern introduction*, volume 123. Cambridge University Press, 2010.
- [124] Eugene F Fama. Random walks in stock market prices. *Financial analysts journal*, 51(1):75–80, 1995.
- [125] Howard C Berg. *Random walks in biology*. Princeton University Press, 1993.
- [126] Wolfgang Alt. Biased random walk models for chemotaxis and related diffusion approximations. *Journal of Mathematical Biology*, 9(2):147–177, 1980.
- [127] Leo P Kadanoff. *Statics, Dynamics and Renormalization*. World Scientific, 2000.
- [128] Ben C Travaglione and Gerald J Milburn. Implementing the quantum random walk. *Physical Review A*, 65(3):032310, 2002.
- [129] Mark EJ Newman. The structure and function of complex networks. *SIAM review*, 45(2):167–256, 2003.
- [130] Lada A Adamic, Rajan M Lukose, Amit R Puniyani, and Bernardo A Huberman. Search in power-law networks. *Physical Review E*, 64(4):046135, 2001.

- [131] Cornelius Lanczos. *An iteration method for the solution of the eigenvalue problem of linear differential and integral operators*. United States Governm. Press Office, 1950.
- [132] Lev Davidovich Landau and Evgenii Mikhailovich Lifshitz. *Quantum mechanics: non-relativistic theory*, volume 3. Elsevier, 2013.
- [133] Claude Cohen-Tannoudji, Bernard Diu, and Franck Laloë. *Mécanique quantique*. 1986.
- [134] Yoshiyuki Kabashima, Hisanao Takahashi, and Osamu Watanabe. Cavity approach to the first eigenvalue problem in a family of symmetric random sparse matrices. In *Journal of Physics: Conference Series*, volume 233, page 012001. IOP Publishing, 2010.
- [135] Raghunandan H Keshavan, Andrea Montanari, and Sewoong Oh. Matrix completion from a few entries. *Information Theory, IEEE Transactions on*, 56(6):2980–2998, 2010.
- [136] Emmanuel J Candès and Terence Tao. The power of convex relaxation: Near-optimal matrix completion. *Information Theory, IEEE Transactions on*, 56(5):2053–2080, 2010.
- [137] Alan Frieze, Ravi Kannan, and Santosh Vempala. Fast Monte-Carlo algorithms for finding low-rank approximations. *Journal of the ACM (JACM)*, 51(6):1025–1041, 2004.
- [138] Carl Eckart and Gale Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936.
- [139] Michael E Wall, Andreas Rechtsteiner, and Luis M Rocha. Singular value decomposition and principal component analysis. In *A practical approach to microarray data analysis*, pages 91–109. Springer, 2003.
- [140] Emmanuel J Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717–772, 2009.
- [141] Raghunandan Keshavan, Andrea Montanari, and Sewoong Oh. Matrix completion from noisy entries. In *Advances in Neural Information Processing Systems*, pages 952–960, 2009.
- [142] Emmanuel J Candès and Yaniv Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925–936, 2010.

- [143] Alaa Saade, Florent Krzakala, and Lenka Zdeborová. Matrix completion from fewer entries: Spectral detectability and rank estimation. *arXiv preprint arXiv:1506.03498*, 2015.
- [144] David Goldberg, David Nichols, Brian M Oki, and Douglas Terry. Using collaborative filtering to weave an information tapestry. *Communications of the ACM*, 35(12):61–70, 1992.
- [145] Xiaoyuan Su and Taghi M Khoshgoftaar. A survey of collaborative filtering techniques. *Advances in Artificial Intelligence*, 2009:4, 2009.
- [146] Yateendra Deshpande and Alessandro Montanari. Linear bandits in high dimension and recommendation systems. In *Communication, Control, and Computing (Allerton), 2012 50th Annual Allerton Conference on*, pages 1750–1754. IEEE, 2012.
- [147] Alex Simpkins, Raymond De Callafon, and Emanuel Todorov. Optimal trade-off between exploration and exploitation. In *American Control Conference, 2008*, pages 33–38. IEEE, 2008.
- [148] Christian Kahl and Michael Günther. Complete the correlation matrix. In *From Nano to Space*, pages 229–244. Springer, 2008.
- [149] Andrew Gelman, John B Carlin, Hal S Stern, and Donald B Rubin. *Bayesian data analysis*, volume 2. Taylor & Francis, 2014.
- [150] Jérôme Idier. *Bayesian approach to inverse problems*. John Wiley & Sons, 2013.
- [151] John Wishart. The generalised product moment distribution in samples from a normal multivariate population. *Biometrika*, pages 32–52, 1928.
- [152] Ronald A Fisher. The statistical utilization of multiple measurements. *Annals of eugenics*, 8(4):376–386, 1938.
- [153] Christopher M Bishop. *Pattern recognition and machine learning*. Springer, 2006.
- [154] Roger A Horn and Charles R Johnson. *Matrix analysis*. Cambridge University Press, 2012.
- [155] Chris Anderson. The end of theory: The data deluge makes the scientific method obsolete, 2008.
- [156] Kirk Borne. The new lsst informatics and statistical sciences collaboration team.
- [157] http://www.theregister.co.uk/Print/2008/10/03/lsst_jeff_kantor/.
- [158] <http://home.web.cern.ch/about/computing>.

- [159] Viktor Mayer-Schönberger and Kenneth Cukier. *Big data: A revolution that will transform how we live, work, and think*. Houghton Mifflin Harcourt, 2013.
- [160] Nate Silver. *The signal and the noise: Why so many predictions fail-but some don't*. Penguin, 2012.
- [161] <https://www.quora.com/>.
- [162] Mark Lycett. 'Datafication': Making sense of (big) data in a complex world. 2013.
- [163] George Coulouris, Jean Dollimore, Tim Kindberg, and Gordon Blair. *Distributed systems*. Pearson Education, 1989.
- [164] <http://www.dailymail.co.uk/sciencetech/article-2381188/Revealed-happens-just-ONE-minute-internet-216-000-photos-posted-278-000-Tweets-1-8m-Facebook-likes.html>.

Appendix A

Minimal contagious sets in random regular graphs

AG, Guilhem Semerjian

Minimal Contagious Sets in Random Regular Graphs

Alberto Guggiola · Guilhem Semerjian

Received: 28 July 2014 / Accepted: 7 October 2014 / Published online: 21 October 2014
© Springer Science+Business Media New York 2014

Abstract The bootstrap percolation (or threshold model) is a dynamic process modelling the propagation of an epidemic on a graph, where inactive vertices become active if their number of active neighbours reach some threshold. We study an optimization problem related to it, namely the determination of the minimal number of active sites in an initial configuration that leads to the activation of the whole graph under this dynamics, with and without a constraint on the time needed for the complete activation. This problem encompasses in special cases many extremal characteristics of graphs like their independence, decycling or domination number, and can also be seen as a packing problem of repulsive particles. We use the cavity method (including the effects of replica symmetry breaking), an heuristic technique of statistical mechanics many predictions of which have been confirmed rigorously in the recent years. We have obtained in this way several quantitative conjectures on the size of minimal contagious sets in large random regular graphs, the most striking being that 5-regular random graph with a threshold of activation of 3 (resp. 6-regular with threshold 4) have contagious sets containing a fraction $1/6$ (resp. $1/4$) of the total number of vertices. Equivalently these numbers are the minimal fraction of vertices that have to be removed from a 5-regular (resp. 6-regular) random graph to destroy its 3-core. We also investigated Survey Propagation like algorithmic procedures for solving this optimization problem on single instances of random regular graphs.

Keywords Bootstrap percolation · Optimization problems · Cavity method · Random graphs

1 Introduction

Models of epidemic spreadings as dynamical processes occurring on a graph appear in various contexts besides epidemiology [15, 23, 34, 42, 63]; for instance social sciences study viral

A. Guggiola · G. Semerjian (✉)
LPTENS, Unité Mixte de Recherche (UMR 8549) du CNRS et de l'ENS, associée à l'UPMC
Univ Paris 06, 24 Rue Lhomond, 75231 Paris Cedex 05, France
e-mail: guilhem@lpt.ens.fr

marketing campaigns aimed at propagating new social trends, and in economy it is crucial to understand cascading effects potentially leading to the bankrupt of financial institutions. In these models individual agents are located on the vertices of a graph, and their state (healthy or contaminated for instance) evolve in time according to the state of their neighbours, the edges of the graph representing the contacts between agents that can possibly transmit the illness from one contaminated agent to an healthy one.

There is a great diversity in the details of these models: the dynamics can occur in continuous (asynchronous) or discrete time, according to deterministic or random rules, the state of an agent can be boolean (healthy or contaminated) or describe several levels of contamination, and finally the dynamics can be monotonous or not. To precise this last point, a dynamics is said monotonous if the states of an agent always occur in the same order in time, for instance in the Susceptible-Infected-Recovered (SIR) model the only allowed transitions are $S \rightarrow I$ and $I \rightarrow R$, a Recovered individual being immune forever, whereas in the SIS model an agent can become infected several times in a row. In this paper we will concentrate on a simple monotonous dynamics, that evolve deterministically in discrete time, with inactive (Susceptible) variables becoming active (Infected) when their number of active neighbours reach some threshold, and then remain active for ever. For this reason it is called the threshold model, see [39] for a version introduced in sociology with an underlying complete graph, and [27] for its first appearance in physics under the name of bootstrap percolation (on random regular graphs).

Given one specific dynamical model there are many different questions that can be asked. The first, a priori simplest, issue concerns the time evolution of the system from a random initial condition, taking the initial state of each agent as an independent random variable. For monotonous dynamics a stationary state is reached after some time, and one can wonder whether the epidemic has invaded the whole graph (in other words whether it percolates) in this final state. The probability of this event obviously depends on the fraction of infected vertices in the initial condition, and this may lead to phase transitions for certain class of graphs; see [5, 11, 43] for such a study of the bootstrap percolation on finite-dimensional lattices, and [12, 25, 27, 44–46, 50, 73] for various type of dynamics on random graphs. In particular one finds for the bootstrap percolation on random regular graphs a phase transition at some initial critical density θ_r (dependent on the degree of the graph and the threshold of activation): with high probability initial conditions with a fraction θ of active vertices (without correlations between the sites) are percolating if and only if $\theta > \theta_r$.

Besides these studies of the “forward” (or “direct”) time evolution, which are somehow simplified by the independence assumption for the initial state variables, one can also formulate more difficult inference and optimization questions. An example of the former type is to infer some information on the initial state given a snapshot of the epidemic after some time evolution [6, 51, 66, 72]; this “inverse problem” is particularly relevant in epidemiology in the search of the “zero patient” who triggered the spreading of an illness. For what regards the latter type of questions, the design of an efficient vaccination campaign can indeed be seen as an optimization problem: find the smallest set of nodes (to minimize the economical and social cost) whose vaccination will prevent the epidemic to reach a given fraction of the population [7]. We shall actually consider in this paper the somehow reverse optimization problem, namely targeting a small set of initially active sites that lead to the largest possible propagation of the contagion. This obviously makes more sense in the perspective of viral marketing, in which it was first considered [47] than in the epidemiological one; the initial adopters of a new product, that can be financially incited to do so, are expected to convince most of their acquaintances and progressively the largest possible part of the population.

From this point of view the additional constraint that the propagation should be as fast as possible is also a relevant one.

More precisely, one can define two versions of this optimization problem: (i) given a fixed number of initially active agents, choose them in order to maximize the number of active agents at some fixed later time, or in the final state of the propagation; (ii) find the minimal number of initially active agents such that all the agents are active, again after some time or in the final state. We will concentrate on the latter version of the problem but part of our analysis applies to both. These optimization problems are known to be hard from a (worst-case) computational complexity point of view [28,35,47], even to approximate. Exhibiting minimal percolating sets for bootstrap percolation on finite dimensional lattices is relatively easy thanks to their regular structures, but more refined extremal problems are also relevant in this case, see for instance [20,62]. The understanding of these optimization problems seems less advanced in the case of sparse random graphs. There exist upper and lower bounds on the size of minimal contagious sets [4,35,67], some based in particular on the expansion properties of such graphs [31]. One particular case of the optimization problem (when the threshold of activation is equal to the degree of the vertex minus one) is actually equivalent to the decycling number problem of graph theory [19] (also known as minimal Feedback Vertex Set), which was settled rigorously for 3-regular random graphs in [17] (this paper also contains bounds for higher degrees). As this last point unveils the notion of minimal contagious sets is connected in some special cases to many other problems in graph theory; one way to see this connection is to picture the inactive sites of the initial condition as particles to be put on the graph. One wants to pack as many as possible of them (to obtain a contagious set of minimal size), yet they do have some kind of repulsive interactions because of the constraint of complete percolation at a later time. This is particularly clear when the threshold of activation is equal to the degree for all vertices: the problem is then exactly equivalent to the hard-core particle model, also known as independent set or vertex cover.

The strategy we shall follow to determine the minimal size of contagious sets of sparse random graphs will be the same as in [8,9], namely a reformulation under the form of a statistical mechanics model which can be treated with the so-called cavity method [53–56]. This (heuristic) method yields predictions for any interacting model defined on a sparse random graph; its use in the context of random constraint satisfaction problems led to the discovery of a very rich phenomenology of phase transitions [48,56], with many of these predictions later confirmed rigorously [1,3,13,30,32,57]. Let us emphasise in particular the determination of the maximal size of independent sets of random regular graphs (which as we saw is a problem related to the present one), for which the predictions of the cavity method (see [14] and references therein) have been recently rigorously confirmed (for graphs of large enough but finite degree) in [33]. Another example in the context of graph theory is the study of matchings in random graphs, where the cavity method [75] has also been proved to be correct [26]. The main originality of our contribution with respect to [8,9] is the use of a more refined version of the cavity method (i.e. incorporating the effects of replica symmetry breaking), and an analytical study of the limit where the time at which the complete activation is required is sent to infinity.

The rest of the article is organized as follows. In Sect. 2 we define precisely the dynamics under study, recall briefly some known results for random initial conditions, formulate the optimization problem and propose various interpretations of it, and for the convenience of the reader we summarize the main results to be obtained in the following. In Sect. 3 we derive the cavity method equations, both at the replica symmetric and one step of replica symmetry breaking level. The solution of these equations for random regular graphs is presented in Sect. 4, which contains the main analytical results of this work. Section 5 is devoted to

single sample numerical experiments, where we confront the analytical predictions with the optimized initial configurations obtained with two kind of algorithms (a simple greedy one and a more involved procedure based on message passing). We finally draw our conclusions and present perspectives for future work in Sect. 6. The most technical parts of the computations are deferred to two Appendices.

2 Definitions and Main Results

2.1 Definition of the Dynamics

Let us consider a graph on N vertices (or sites), $G = (V, E)$, with the vertices labelled as $V = \{1, \dots, N\}$, and the number of edges denoted $|E| = M$. The dynamical process under study concerns the evolution of variables σ_i^t on the vertices, $\sigma_i^t = 0$ (resp. 1) if the vertex i is inactive (resp. active) at time t . We shall denote $\underline{\sigma}^t = (\sigma_1^t, \dots, \sigma_N^t)$ the global configuration at time t . The latter is determined by the initial condition $\underline{\sigma}$ at the initial time, $\underline{\sigma}^0 = \underline{\sigma}$, and then evolves subsequently in a deterministic and parallel way, in discrete time, according to the rules:

$$\sigma_i^t = \begin{cases} 1 & \text{if } \sigma_i^{t-1} = 1 \\ 1 & \text{if } \sigma_i^{t-1} = 0 \text{ and } \sum_{j \in \partial i} \sigma_j^{t-1} \geq l_i, \\ 0 & \text{otherwise} \end{cases}, \quad (1)$$

where ∂i is the set of neighbours of i on the graph, and l_i is a fixed threshold for each vertex; we will also use $d_i = |\partial i|$ to denote the degree of vertex i . The dynamics is monotonous (irreversible), an active site remaining active at all later times, an inactive site i becoming active if its number of active neighbours at the previous time crosses the threshold l_i . Note that the configuration $\underline{\sigma}^t$ at time t is a deterministic function of the initial condition $\underline{\sigma} = \underline{\sigma}^0$, and that by monotonicity one can define the final configuration $\underline{\sigma}^f = \lim_{t \rightarrow \infty} \underline{\sigma}^t$, this stationary configuration being reached in a finite number of steps for all finite graphs.

It turns out that the final configuration $\underline{\sigma}^f$ is also the one reached by a sequential dynamics in which at each time step only one site i with at least l_i active neighbours is activated; a moment of thought reveals the independence of the final configuration with respect to the order of the updates. $\underline{\sigma}^f$ is indeed the smallest configuration (considering the partial order $\underline{\sigma} \leq \underline{\sigma}'$ if and only if $\sigma_i \leq \sigma'_i$ for all vertices) larger than the initial condition $\underline{\sigma}$, such that no further site can be activated. It will sometimes be useful in the following to think of this process in a dual way, corresponding to the original presentation of bootstrap percolation in [27], namely to consider that inactive sites are sequentially removed if they have less than a certain number of inactive neighbours. An equivalent definition of $\underline{\sigma}^f$ is thus given by the inactive sites it contains, that form the largest set (with respect to the inclusion partial order) contained in the set of inactive sites of $\underline{\sigma}$, and such that in their induced graph the degree of site i is larger or equal than $d_i - l_i + 1$; they form thus a (generalized inhomogeneous version of the) core of the initially inactive sites.

2.2 Reminder of the Behaviour for Random Initial Conditions on Random Regular Graphs

To put in perspective the optimization problem to be studied in this paper it is instructive to first recall briefly some well-known results for the evolution from a random initial configuration [12, 27]. For the sake of simplicity let us consider G to be a $k + 1$ -random regular

graph (i.e. a graph drawn uniformly at random among all graphs in which every vertex has degree $k + 1$), with a uniform threshold for activation set to $l_i = l$ for all vertices. Suppose that the states of the vertices in the initial condition are chosen randomly, independently and identically for each vertex, with a probability θ (resp. $1 - \theta$) for a vertex to be active (resp. inactive). The probability for one vertex i_0 to be active at some time $t + 1$, denoted x_{t+1} , can be computed from the following equation:

$$x_{t+1} = \theta + (1 - \theta) \sum_{p=l}^{k+1} \binom{k+1}{p} \tilde{x}_t^p (1 - \tilde{x}_t)^{k+1-p} . \tag{2}$$

Indeed such a vertex was either active in the initial condition, or has seen at least l of its neighbours activate themselves before time t , and without the participation of i_0 . The probability \tilde{x}_t of this last event obeys the recursive equation

$$\tilde{x}_{t+1} = \theta + (1 - \theta) \sum_{p=l}^k \binom{k}{p} \tilde{x}_t^p (1 - \tilde{x}_t)^{k-p} , \tag{3}$$

with a number of participating neighbours reduced from $k + 1$ to k as i_0 has to be supposed inactive here. The initial condition for these equations is $x_0 = \tilde{x}_0 = \theta$. In the limit $t \rightarrow \infty$ of large times $\tilde{x}_t \rightarrow \tilde{x}_\infty(\theta)$, the smallest fixed-point in $[0, 1]$ of the recursion (3). For each $k \geq 2$ and l with $2 \leq l \leq k$ there exists a threshold $\theta_r(k, l)$ such that $\tilde{x}_\infty(\theta)$ is equal to 1 for $\theta > \theta_r$, strictly smaller than 1 for $\theta < \theta_r$. From Eq. (2) one realizes that the same statement applies to $x_\infty(\theta)$, hence θ_r is the threshold for complete activation (percolation) from a Bernoulli random initial condition with probability θ for each active site. Studying more precisely Eq. (3) one realizes that for $l = k$ the transition is continuous ($x_\infty(\theta_r^-) = 1$), with an explicit expression for the threshold, $\theta_r(k, k) = \frac{k-1}{k}$. For $2 \leq l \leq k - 1$ the transition is discontinuous ($x_\infty(\theta_r^-) < 1$), the threshold θ_r is obtained as the solution of the equations:

$$\begin{cases} \tilde{x}_r = \theta_r + (1 - \theta_r) \sum_{p=l}^k \binom{k}{p} \tilde{x}_r^p (1 - \tilde{x}_r)^{k-p} \\ 1 = (1 - \theta_r) l \binom{k}{l} \tilde{x}_r^{l-1} (1 - \tilde{x}_r)^{k-l} \end{cases} , \tag{4}$$

where $\tilde{x}_r = \tilde{x}_\infty(\theta_r^-)$ is the value of the fixed-point of (3) at the bifurcation where it disappears discontinuously. For $l = 2$ these equations can be solved explicitly and yield

$$\theta_r(k, l = 2) = 1 - \frac{(k - 1)^{2k-3}}{k^{k-1}(k - 2)^{k-2}} . \tag{5}$$

For generic values of the parameters k, l there is no explicit expression of θ_r , as (4) are algebraic equations of arbitrary degree; some numerical values of θ_r will be given in Table 4. For a given value of k the threshold $\theta_r(k, l)$ is growing with l : if an initial condition leads to complete activation for some parameter l it will also be activating under the less constrained dynamics with $l' < l$.

The relevant range for the threshold parameter l in this study of random initial conditions is $2 \leq l \leq k$. Indeed for $l = 0$ after one step the configuration is completely active regardless of σ^0 , for $l = 1$ a single active site (per connected component) in the initial configuration is enough to activate the whole graph, hence in these two cases $\theta_r = 0$. On the other hand if $l = k + 1$ one has $\theta_r = 1$: any pair of adjacent inactive sites in the initial condition will remain inactive for ever, and the number of such pairs is linear in N as soon as $\theta < 1$.

Note that the recursion equations (2, 3) are exact if the neighbourhood up to distance t of the vertex i_0 is a regular tree of degree $k + 1$. The limit $t \rightarrow \infty$ can be taken in this way

only if the graph considered is an infinite regular tree. A rigorous proof that this reasoning is in fact correct also for the large size limit of random regular graphs (that converge locally to regular trees) can be found in [12].

2.3 Definition of the Optimization Problem Over Initial Conditions

Let us now come back to a general graph G with some thresholds l_i for vertex activation, and consider the minimal fraction of active vertices in an initial configuration that activates the whole graph, i.e.

$$\theta_{\min}(G, \{l_i\}) = \frac{1}{N} \min_{\underline{\sigma}} \left\{ \sum_{i=1}^N \sigma_i \mid \sigma_i^f = 1 \ \forall i \right\}. \quad (6)$$

This corresponds to the minimal size of a contagious (or percolating) set, divided by the total number of vertices. Following [8,9] it will turn out useful to introduce another parameter T (a positive integer) in this optimization problem, and impose now that the fully active configuration is reached within this time horizon T :

$$\theta_{\min}(G, \{l_i\}, T) = \frac{1}{N} \min_{\underline{\sigma}} \left\{ \sum_{i=1}^N \sigma_i \mid \sigma_i^T = 1 \ \forall i \right\}. \quad (7)$$

Obviously for any finite graph $\theta_{\min}(G, \{l_i\}, T)$ decreases when T increases and has $\theta_{\min}(G, \{l_i\})$ as its limit for $T \rightarrow \infty$. To turn the computation of θ_{\min} into a form more reminiscent of statistical mechanics problems we shall introduce a probability measure over initial configurations:

$$\eta(\underline{\sigma}) = \frac{1}{Z(G, \{l_i\}, T, \mu, \epsilon)} e^{\sum_{i=1}^N [\mu \sigma_i - \epsilon(1 - \sigma_i^T)]}, \quad (8)$$

where $\underline{\sigma}^T$ is as above the configuration obtained after T steps of the dynamics starting from the configuration $\underline{\sigma} = \underline{\sigma}^0$, the μ and ϵ are for the time being arbitrary parameters, and the partition function Z ensures the normalization of this law. The parameter μ is a ‘‘chemical potential’’ that controls the fraction of initially active vertices (if $\epsilon = 0$ the measure η reduces to the Bernouilli measure), while ϵ is the cost to be paid for each site i inactive at the final time T . In particular if $\epsilon = +\infty$ one has

$$\eta(\underline{\sigma}) = \frac{1}{Z(G, \{l_i\}, T, \mu, \epsilon = +\infty)} e^{\mu \sum_{i=1}^N \sigma_i} \prod_{i=1}^N \mathbb{I}(\sigma_i^T = 1), \quad (9)$$

with $\mathbb{I}(A)$ is the indicator function of the event A , the measure is thus supported by activating initial configurations (within the time horizon T). It is then obvious that the knowledge of Z allows to deduce the sought-for minimal density θ_{\min} , as

$$\theta_{\min}(G, \{l_i\}, T) = \lim_{\mu \rightarrow -\infty} \frac{1}{\mu} \frac{1}{N} \ln Z(G, \{l_i\}, T, \mu, \epsilon = +\infty). \quad (10)$$

Actually one can gain more information from the whole dependency of the partition function on μ . Suppose indeed that the number of initial configurations with a fraction θ of active vertices that activate the whole graph in T steps is, at the leading exponential order, $e^{Ns(\theta)}$, with an entropy density $s(\theta)$ of order one with respect to N . Then this entropy density can

be computed, in the large N limit, as a Legendre transform of the logarithm of the partition function. More precisely, defining the free-entropy density ϕ as

$$\phi(G, \{l_i\}, T, \mu, \epsilon = +\infty) = \frac{1}{N} \ln Z(G, \{l_i\}, T, \mu, \epsilon = +\infty), \quad (11)$$

the evaluation of the sum over configurations in the definition of Z via the Laplace method yields in the large N limit:

$$\phi(G, \{l_i\}, T, \mu, \epsilon = +\infty) = \sup_{\theta \in [\theta_{\min}, 1]} [\mu \theta + s(\theta)], \quad (12)$$

hence $s(\theta)$ can be obtained by an inverse Legendre transform of $\phi(\mu)$, with $s(\theta) = \phi(\mu) - \mu \theta$ and $\theta = \phi'(\mu)$.

For completeness let us also make a similar statement when ϵ is finite, i.e. when one does not impose strictly the constraint of complete activation at time T . Denoting $s(\theta, \theta')$ the entropy density of initial configurations that have a fraction θ of initially active vertices and that lead after T steps of evolution to a configuration with a fraction θ' of active sites, one has

$$\phi(G, \{l_i\}, T, \mu, \epsilon) = \frac{1}{N} \ln Z(G, \{l_i\}, T, \mu, \epsilon) = \sup_{\theta, \theta'} [\mu \theta - \epsilon (1 - \theta') + s(\theta, \theta')]. \quad (13)$$

Varying the parameters μ and ϵ thus allows to reconstruct the function $s(\theta, \theta')$, and hence to solve the optimization problem denoted (i) in the introduction, namely for a fixed value of θ find the maximal reachable θ' . We will mainly concentrate in the following of the paper on the optimization problem denoted (ii) in the introduction, that is imposing the full activation of the graph at time T ($\theta' = 1$), which as explained above can be studied via the computation of $s(\theta) = s(\theta, \theta' = 1)$ from the inverse Legendre transform of the free-entropy with $\epsilon = +\infty$.

The definitions above were valid for any graph and any choice of the activation thresholds; we shall however be particularly interested in the case of large random regular graphs with uniform thresholds, we thus define

$$\theta_{\min}(k, l) = \lim_{N \rightarrow \infty} \mathbb{E}[\theta_{\min}(G, \{l_i = l\})], \quad \theta_{\min}(k, l, T) = \lim_{N \rightarrow \infty} \mathbb{E}[\theta_{\min}(G, \{l_i = l\}, T)], \quad (14)$$

where the average is over uniformly chosen regular graphs of degree $k+1$ on N vertices, with the same threshold for activation l on every vertex. The fact that the limit in the definition of $\theta_{\min}(k, l, T)$ exists could actually be proven rigorously using the method developed in [18], and it is expected that $\theta_{\min}(G, \{l_i = l\}, T)$ is self-averaging (i.e. concentrates around its average in the large N limit). The existence of $\theta_{\min}(k, l)$ might be a more difficult mathematical problem that we shall not discuss further; it is a reasonable conjecture that it coincides with the limit of $\theta_{\min}(k, l, T)$ when $T \rightarrow \infty$, i.e. that the large size and large time limits commute. We will see in Sect. 4.2.1 one argument in favour of this conjecture. Let us emphasize that $\theta_{\min}(k, l) < \theta_r(k, l)$, with a strict inequality. This is indeed a large-deviation phenomenon: even if most initial configurations with density smaller than θ_r do not activate the whole graph some very rare ones (with a probability exponentially small in N in the Bernoulli measure of parameter $\theta < \theta_r$) are able to do so. Note also that $\theta_{\min}(k, l)$ is growing with l at fixed k , for the same reasons as explained above in the discussion of θ_r . The computations of θ_{\min} we shall present will follow the strategy explained above on an arbitrary graph, namely the computation of a free-entropy density, that we define in the case of random regular graphs as the quenched average over the graph ensemble,

$$\phi(k, l, T, \mu, \epsilon) = \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}[\ln Z(G, \{l_i = l\}, T, \mu, \epsilon)]. \quad (15)$$

2.4 Equivalence with Other Problems and Bounds

As mentioned in the introduction the problem of minimal contagious sets can be related, for appropriate choices of the threshold parameters l_i , to other standard problems in graph theory.

Consider first the case of an arbitrary graph where the thresholds l_i are equal to the degrees d_i for all vertices. An inactive site in the initial configuration will be activated only if it is surrounded by active vertices, and it will do so in a single step. In other words in any percolating initial condition, whatever the time horizon T , the inactive vertices must form an independent set (no two inactive vertices are allowed to be neighbours). For regular random graphs one has thus $\theta_{\min}(k, k+1, T) = \theta_{\min}(k, k+1)$ for all T , and this quantity is equal to 1 minus the density of the largest independent sets of a $k+1$ -regular random graph.

Another correspondance with previously studied models arises when $T = 1$, for any choice of the thresholds l_i . Indeed in this case the vertex i can be inactive in a percolating initial configuration only if its number of inactive neighbours is smaller than some value (namely, $\leq d_i - l_i$). These generalized hard-core constraints (repulsion between inactive vertices) correspond exactly to the so-called Biroli–Mézard (BM) model [21, 70] (with the correspondance inactive vertex \leftrightarrow vertex occupied by a particle in the BM model, and $d_i - l_i \leftrightarrow \ell_i$ of the BM model). Hence for $T = 1$ the minimal density θ_{\min} is 1 minus the density of a close packing of the corresponding BM model. Further specializing this $T = 1$ case by setting $l_i = 1$ on each vertex leads to the constraint that every inactive site in a percolating initial configuration has to be adjacent with at least one active site, in other words that the active sites form a dominating set of G . The minimal density θ_{\min} is thus the domination number (divided by N) of G .

Consider now the thresholds of activation to be 1 less than the degrees, i.e. $l_i = d_i - 1$ on all vertices, with no constraint on the time of activation ($T = \infty$). As explained at the end of Sect. 2.1, the inactive vertices in the final configuration form the 2-core of the inactive ones in the initial configuration. A percolating initial configuration must be such that this 2-core is empty, in other words the subgraph induced by the inactive sites of the initial configuration must be acyclic (a tree or a forest), i.e. the active sites have to form a decycling set [19] (also known as a Feedback Vertex Set), and $N\theta_{\min}$ is the decycling number of G . This characterization leads to the following bound for every $k+1$ -regular graph with thresholds k of activation on every site,

$$\theta_{\min}(k, k) \geq \frac{k-1}{2k}. \quad (16)$$

Indeed if A denotes the number of active vertices in a percolating initial configuration, the $N - A$ other vertices induces a forest, the number of edges between inactive vertices is thus at most $N - A - 1$. On the other hand this number is at least $\frac{k+1}{2}N - (k+1)A$ (the first term being the total number of edges, and the number of edges incident to at least one active site being at most $(k+1)A$). The decycling number of random regular graphs was studied in [17], proving in particular that the bound (16) is actually tight for 3-regular large random graphs, i.e. $\theta_{\min}(2, 2) = 1/4$, and it was conjectured to be also the case for 4-regular ones (i.e. $\theta_{\min}(3, 3) = 1/3$). An asymptotic lowerbound on $\theta_{\min}(k, k)$ for large values of k was worked out in [41], we will come back on this result in Sect. 4.2.1. Note also that the decycling number of arbitrary sparse random graphs was studied with physics methods in [71, 76].

For general thresholds smaller than the degrees minus one the active sites of a percolating initial configuration must form a “de-coring” set instead of a “de-cycling” set (i.e. their removal has to provoke the disappearance of a q -core with $q > 2$). A generalization of the lower bound (16) to any $k + 1$ -regular graph with uniform threshold l was given in [35], and reads

$$\theta_{\min}(k, l) \geq \frac{2l - k - 1}{2l}. \quad (17)$$

Its proof goes as follows. Consider the sequential process explained at the end of Sect. 2.1 in which at each time t a single vertex gets activated, and denote $E(t)$ the number of edges between active and inactive vertices after t steps of this process. By definition of the activation rule $E(t + 1) - E(t) \leq k + 1 - 2l$. If as above A denotes the number of active sites in a percolating initial configuration, by definition $E(N - A) = 0$, hence $E(0) \geq (N - A)(2l - k - 1)$. On the other hand $E(0) \leq (k + 1)A$, which gives the lower bound (17) on the possible values of A .

We should also mention an upper bound on the minimal sizes of contagious sets derived in [4, 67] for graphs of arbitrary degree distributions, which yields in the case of $k + 1$ -regular graphs:

$$\theta_{\min}(k, l) \leq \frac{l}{k + 2}. \quad (18)$$

To conclude this discussion let us mention that the “de-coring” perspective on the minimal contagious set problem is somehow reminiscent (even if not directly equivalent), to the Achlioptas processes [2, 69] (more precisely of their offline version [24]) where one looks for an extremal event avoiding the appearance of an otherwise typical structure (a giant component in the Achlioptas processes, a core in the minimal contagious set case).

2.5 Main Analytical Results

Let us draw here a more detailed plan of the rest of the paper to make its reading easier and faster for someone not interested in the technical details of the statistical mechanics method (who can browse quickly over the next section and jump to the results announced in Sect. 4). In order to compute the minimal density θ_{\min} of contagious sets we shall rephrase this problem as a statistical mechanics model and apply to it the cavity method. The latter is based on self-consistent assumptions of various degrees of sophistication, parametrized by the so-called level of replica symmetry breaking. We will study the first two levels of this hierarchy, named replica symmetric (RS) and one step of replica symmetry breaking (1RSB). These two approaches will lead to two predictions for θ_{\min} , to be denoted respectively $\theta_{\min,0}(k, l, T)$ and $\theta_{\min,1}(k, l, T)$. From general bounds established in the context of disordered statistical mechanics models (first for the Sherrington-Kirkpatrick model [40, 64, 74] and later for some models on sparse random graphs [36, 37, 65]) it is expected that the different levels of the cavity method provide improving lower bounds on θ_{\min} , namely

$$\theta_{\min,0}(k, l, T) \leq \theta_{\min,1}(k, l, T) \leq \theta_{\min}(k, l, T). \quad (19)$$

Our computation of $\theta_{\min,0}(k, l, T)$ and $\theta_{\min,1}(k, l, T)$ relies on the resolution of a set of roughly $2T$ algebraic equations on $2T$ unknowns, explicit numbers will be given in Sect. 4. We managed to perform analytically the $T \rightarrow \infty$ limit and reduce the determination of $\theta_{\min,0}(k, l)$ and $\theta_{\min,1}(k, l)$ (their limit when $T \rightarrow \infty$) to a finite number of equations, that will also be presented along with numerical results in Sect. 4. We found four particular cases

in which the predictions of the first two levels of replica symmetry breaking coincide when $T \rightarrow \infty$, which led us to conjecture that they are the exact ones, namely:

$$\theta_{\min}(2, 2) = \frac{1}{4}, \quad \theta_{\min}(3, 3) = \frac{1}{3}, \quad \theta_{\min}(4, 3) = \frac{1}{6}, \quad \theta_{\min}(5, 4) = \frac{1}{4}, \quad (20)$$

all these cases saturating the lower bounds of (16, 17). The first (resp. second) equality was actually proven (resp. conjectured) in [17]. We have also performed a large degree expansion of the decycling number of random regular graphs, yielding the conjecture

$$\theta_{\min}(k, k) = 1 - \frac{2 \ln k}{k} - \frac{2}{k} + O\left(\frac{1}{k \ln k}\right). \quad (21)$$

3 Cavity Method Treatment of the Problem

3.1 Factor Graph Representation

As explained in Sect. 2.3 the central quantity to compute is the free-entropy density defined from the partition function normalizing the probability law (8), that for completeness we shall generalize to possibly site dependent chemical potentials μ_i and costs for non-activation ϵ_i :

$$\eta(\underline{\sigma}) = \frac{1}{Z(G, \{l_i\}, T, \{\mu_i, \epsilon_i\})} e^{\sum_{i=1}^N [\mu_i \sigma_i - \epsilon_i (1 - \sigma_i^T)]}. \quad (22)$$

This expression is not very convenient to handle directly because the variables σ_i have complicated interactions under this law: σ_i^T is indeed a function of all variables σ_j on the vertices j at distance smaller than T from i . A way to circumvent this difficulty and to turn the interactions of the model into local ones has been proposed in [8,9], and we shall follow the same approach here.

Let us first define $t_i(\underline{\sigma})$ as the time of activation of site i in the dynamical process generated by the initial configuration $\underline{\sigma}$, i.e. $t_i(\underline{\sigma}) = \min\{t : \sigma_i^t = 1\}$, with conventionally $t_i(\underline{\sigma}) = \infty$ if this time is strictly greater than the time horizon T . These variables obey the following equations:


$$t_i(\underline{\sigma}) = f(\sigma_i, \{t_j(\underline{\sigma})\}_{j \in \partial i}; l_i) \quad \forall i \in V, \quad (23)$$

where the function f is defined as

$$f(\sigma, t_1, \dots, t_n; l) = \begin{cases} 0 & \text{if } \sigma = 1 \\ 1 + \min_l(t_1, \dots, t_n) & \text{if } \sigma = 0 \text{ and } 1 + \min_l(t_1, \dots, t_n) \leq T. \\ \infty & \text{otherwise} \end{cases} \quad (24)$$

Here $\min_l(t_1, \dots, t_n)$ is the l -th smallest t_i , i.e ordering the arguments as $t_1 \leq t_2 \leq \dots \leq t_n$ one has $\min_l(t_1, \dots, t_n) = t_l$. This translates the dynamic rules (1) in terms of the activation times, a site i activating at the time following the first time where at least l_i of its neighbours are active. In the following $f(0, t_1, \dots, t_n; l)$ will be abbreviated in $f(t_1, \dots, t_n; l)$. Reciprocally one can show that if a set of $\{t_i\}_{i \in V}$ verifies the condition that for all i either $t_i = 0$ or $t_i = f(\{t_j\}_{j \in \partial i}; l_i)$, then they correspond to the activation times for the dynamics started from the initial condition $\underline{\sigma}$ such that $\sigma_i = 1$ if and only if $t_i = 0$. These two descriptions in terms of $(\sigma_1, \dots, \sigma_N)$ and (t_1, \dots, t_N) are thus equivalent, yet the great advantage of the

Fig. 1 A portion of the factor graph corresponding to the measure (25)



latter is that the conditions to enforce among the t_i are local along the graph, and that they contain in an obvious way the information on σ_i^T that was lacking to deal with (22).

Finally a last twist on Eq. (22) will be to “duplicate” the activation time t_i on all edges connecting i to one of its neighbour j , introducing redundant variables t_{ij} to be finally constrained to be all equal to t_i . Let us denote \underline{t} the collective configurations of all these $2M$ variables t_{ij}, t_{ji} on each edge $\langle i, j \rangle$ of the graph, that take values in $\{0, 1, \dots, T, \infty\}$, and consider the following probability measure on $(\underline{\sigma}, \underline{t})$:

$$\eta(\underline{\sigma}, \underline{t}) = \frac{1}{Z(G, \{l_i\}, T, \{\mu_i, \epsilon_i\})} \prod_{i=1}^N w_i(\sigma_i, \{t_{ij}, t_{ji}\}_{j \in \partial i}), \tag{25}$$

where the functions w_i are defined by

$$w_i(\sigma_i, \{t_{ij}, t_{ji}\}_{j \in \partial i}) = e^{\mu_i \sigma_i} e^{-\epsilon_i \mathbb{I}(f(\sigma_i, \{t_{ki}\}_{k \in \partial i}; l_i) = \infty)} \prod_{j \in \partial i} \mathbb{I}(t_{ij} = f(\sigma_i, \{t_{ki}\}_{k \in \partial i}; l_i)). \tag{26}$$

The above observations imply that the marginal of $\underline{\sigma}$ under $\eta(\underline{\sigma}, \underline{t})$ is precisely the desired one from Eq. (22), and that in the support of the law the \underline{t} are strictly constrained to be the activation times for the dynamics starting from $\underline{\sigma}$. This correspondance being one-to-one the partition function is the same in (22) and (25). A portion of the factor graph [49] associated to the probability law (25) is sketched in Fig. 1, with black squares representing the function nodes (interactions) w_i , black circles the variables σ_i , and white circles the variables t_{ij}, t_{ji} . One notes that if the original graph G is a tree (resp. is locally a tree) then the corresponding factor graph is a tree (resp. is locally a tree). This fact was the motivation for the “duplication” of the t_i variables on the surrounding edge, without it short loops of interactions would still be present in the factor graph.

3.2 Replica Symmetric (RS) Formalism

Let us now explain how the probability law (25) and its associated normalization Z can be handled within the cavity formalism, first at the simplest, so called Replica Symmetric (RS), level.

3.2.1 Single Sample Equations

If the graph G were a finite tree the factor graph associated to (25) would be a tree, hence Z and the marginals of η could be computed exactly via the recursive equations that we are about to write down. If the graph is only locally tree-like these equations are only approximate, they correspond to the (loopy) Belief Propagation equations, valid under some assumptions of long-range correlation decay under the measure η . This recursive computation of Z amounts to introduce on each directed edge $i \rightarrow j$ of the graph a “message” $\eta_{i \rightarrow j}(t_{ij}, t_{ji})$, which is a normalized probability distribution over a pair of activation times. These messages obey recursion relations of the form $\eta_{i \rightarrow j} = \widehat{g}(\{\eta_{k \rightarrow i}\}_{k \in \partial i \setminus j}; l_i, \epsilon_i, \mu_i)$, where the mapping $\eta = \widehat{g}(\eta_1, \dots, \eta_k; l, \epsilon, \mu)$ is given by

$$\eta(t, t') = \frac{1}{\widehat{z}_{\text{iter}}(\eta_1, \dots, \eta_k; l, \epsilon, \mu)} \left[\delta_{t,0} e^{\mu} \prod_{i=1}^k \left(\sum_{t''} \eta_i(t'', 0) \right) + e^{-\epsilon \delta_{t,\infty}} \sum_{t_1, \dots, t_k} \eta_1(t_1, t) \dots \eta_k(t_k, t) \mathbb{I}(t = f(t_1, \dots, t_k, t'; l)) \right]. \tag{27}$$

Here and in the following unprecised summations over a time index go along $\{0, \dots, T, \infty\}$. The function $\widehat{z}_{\text{iter}}(\eta_1, \dots, \eta_k; l, \epsilon, \mu)$ is defined by normalization, in such a way that $\sum_{t, t'} \eta(t, t') = 1$.

The knowledge of the messages $\eta_{i \rightarrow j}$ on all edges of the graph allows to compute the free-entropy density, according to the Bethe formula:

$$\phi = \frac{1}{N} \ln Z = \frac{1}{N} \sum_{i=1}^N \ln \widehat{z}_{\text{site}}(\{\eta_{j \rightarrow i}\}_{j \in \partial i}; l_i, \epsilon_i, \mu_i) - \frac{1}{N} \sum_{\langle i, j \rangle \in E} \ln \widehat{z}_{\text{edge}}(\eta_{i \rightarrow j}, \eta_{j \rightarrow i}), \tag{28}$$

where the second sum runs over the (undirected) edges of the graph, and the local partition functions are

$$\begin{aligned} \widehat{z}_{\text{site}}(\eta_1, \dots, \eta_{k+1}; l, \epsilon, \mu) &= e^{\mu} \prod_{i=1}^{k+1} \left(\sum_{t'} \eta_i(t', 0) \right) \\ &+ \sum_{t=1}^T \sum_{t_1, \dots, t_{k+1}} \eta_1(t_1, t) \dots \eta_{k+1}(t_{k+1}, t) \mathbb{I}(t = 1 + \min_l(t_1, \dots, t_{k+1})) \\ &+ e^{-\epsilon} \sum_{t_1, \dots, t_{k+1}} \eta_1(t_1, \infty) \dots \eta_{k+1}(t_{k+1}, \infty) \mathbb{I}(\min_l(t_1, \dots, t_{k+1}) \geq T) \end{aligned} \tag{29}$$

$$\widehat{z}_{\text{edge}}(\eta_1, \eta_2) = \sum_{t, t'} \eta_1(t, t') \eta_2(t', t). \tag{30}$$

The marginals of the law (25) can also be deduced from the messages, for instance the probability distribution of the activation time t_i for the vertex i reads $\eta(t_i) = \widehat{\eta}_{\text{site}}(\{\eta_{j \rightarrow i}\}_{j \in \partial i}; l_i, \epsilon_i, \mu_i)(t_i)$, where

$$\widehat{\eta}_{\text{site}}(\eta_1, \dots, \eta_{k+1}; l, \epsilon, \mu)(t) = \frac{1}{\widehat{z}_{\text{site}}(\eta_1, \dots, \eta_{k+1}; l, \epsilon, \mu)} \left\{ \delta_{t,0} e^{\mu} \prod_{i=1}^{k+1} \left(\sum_{t'} \eta_i(t', 0) \right) \right. \\ \left. + (1 - \delta_{t,0} - \delta_{t,\infty}) \sum_{t_1, \dots, t_{k+1}} \eta_1(t_1, t) \dots \eta_{k+1}(t_{k+1}, t) \mathbb{I}(t = 1 + \min_l(t_1, \dots, t_{k+1})) \right. \\ \left. + \delta_{t,\infty} e^{-\epsilon} \sum_{t_1, \dots, t_{k+1}} \eta_1(t_1, \infty) \dots \eta_{k+1}(t_{k+1}, \infty) \mathbb{I}(\min_l(t_1, \dots, t_{k+1}) \geq T) \right\}. \tag{31}$$

The probability that the vertex i is active in the initial condition is then deduced as $\eta(\sigma_i = 1) = \eta(t_i = 0)$. As explained above in Eq. (13), one can deduce from the above results the entropy density $s(\theta, \theta')$ for initial configurations with a fraction θ of active sites leading to a fraction θ' of active sites after T steps, taking $\mu_i = \mu$ and $\epsilon_i = \epsilon$ for all sites, with

$$s(\theta, \theta') = \phi(\mu, \epsilon) - \mu\theta + \epsilon(1 - \theta'), \quad \theta = \frac{1}{N} \sum_{i=1}^N \eta(t_i = 0), \quad \theta' = \frac{1}{N} \sum_{i=1}^N \eta(t_i \leq T). \tag{32}$$

Note that the derivatives of ϕ with respect to μ and ϵ can be taken only on the explicit dependence in (28), the recursion equations on the messages $\eta_{i \rightarrow j}$ being precisely the stationarity condition of ϕ with respect to the η 's.

3.2.2 A More Compact Parametrization of the Messages

Each probability distribution $\eta(t, t')$ is a priori described by $(T + 2)^2 - 1$ independent real numbers (the times run over $T + 2$ values, including ∞ , and there is a global normalization constraint). We shall see however that a much more compact parametrization is possible, which will be very useful for the further analytical treatment of the model. From now on we shall assume that $\mu_i = \mu$ and $\epsilon_i = \epsilon$ for all vertices. To unveil these simplifications let us first rewrite Eq. (27) more explicitly:

$$\eta(0, t') = \frac{1}{\widehat{z}_{\text{iter}}} e^{\mu} \prod_{i=1}^k (\eta_i(0, 0) + \eta_i(1, 0) + \dots + \eta_i(T, 0) + \eta_i(\infty, 0)) \tag{33}$$

$$\eta(t, t') = \frac{1}{\widehat{z}_{\text{iter}}} \sum_{t_1, \dots, t_k} \eta_1(t_1, t) \dots \eta_k(t_k, t) \mathbb{I}(t = 1 + \min_l(t_1, \dots, t_k, t')) \quad \text{for } t \in \{1, \dots, T\} \tag{34}$$

$$\eta(\infty, t') = \frac{1}{\widehat{z}_{\text{iter}}} e^{-\epsilon} \sum_{t_1, \dots, t_k} \eta_1(t_1, \infty) \dots \eta_k(t_k, \infty) \mathbb{I}(\min_l(t_1, \dots, t_k, t') \geq T) \tag{35}$$

where in all the three cases t' can take any value in $\{0, 1, \dots, T, \infty\}$. Now the condition “ $\min_l(t_1, \dots, t_k, t') = t - 1$ ” is easily seen to be equivalent to “at least l of the time arguments are $\leq t - 1$ and at most $l - 1$ of them are $\leq t - 2$ ”. Similarly the condition “ $\min_l(t_1, \dots, t_k, t') \geq T$ ” is equivalent to “at most $l - 1$ times are $\leq T - 1$ ”. This observation allows to rewrite the above equations under the following form:

$$\eta(0, t') = \frac{1}{\widehat{z}_{\text{iter}}} e^{\mu} \prod_{i=1}^k (\eta_i(0, 0) + \eta_i(1, 0) + \dots + \eta_i(T, 0) + \eta_i(\infty, 0)) \tag{36}$$

$$\eta(t, t') = \frac{1}{\widehat{z}_{\text{iter}}} \sum_{\substack{I, J, K \\ |I| + \mathbb{I}(t' \leq t-2) \leq l-1 \\ |I| + |J| + \mathbb{I}(t' \leq t-1) \geq l}} \prod_{i \in I} \left(\sum_{t''=0}^{t-2} \eta_i(t'', t) \right) \prod_{i \in J} \eta_i(t-1, t) \prod_{i \in K} \left(\sum_{t'' \geq t} \eta_i(t'', t) \right) \tag{37}$$

$$\eta(\infty, t') = \frac{1}{\widehat{z}_{\text{iter}}} e^{-\epsilon} \sum_{I, J} \prod_{i \in I} \left(\sum_{t''=0}^{T-1} \eta_i(t'', \infty) \right) \prod_{i \in J} \left(\sum_{t'' \geq T} \eta_i(t'', \infty) \right) \tag{38}$$

where the summation in the second (resp. third) line is over the partitions I, J, K (resp. I, J) of $\{1, \dots, k\}$. These expressions reveal a first simplification, as already noticed in [8, 9]: among the $(T + 2)^2$ elements of $\eta(t, t')$ only $3T + 2$ are distinct. Indeed $\eta(0, t')$ is independent of t' , for a given value of $t \in \{1, \dots, T\}$ $\eta(t, t')$ takes at most three distinct values, whether $t' \geq t, t' = t - 1$, or $t' \leq t - 2$ and finally $\eta(\infty, t')$ takes two values whether $t' \leq T - 1$ or $t' \geq T$. There is however a further simplification to perform: in the right hand sides of the above equations the η_i 's always appear under the form of particular linear combinations. In particular the elements under the diagonal of the matrices η_i , i.e. $\eta_i(t, t')$ with $t \geq t'$, always intervene under the form $\sum_{t \geq t'} \eta(t, t')$. This allows to reduce further the number of relevant linear combinations of elements of the η 's. A convenient parametrization of the messages η is thus provided by the numbers a_t for $t \in \{0, 1, \dots, T\}$ and b_t for $t \in \{1, \dots, T\}$, defined by:

$$e^{\mu a_t} = \frac{\eta(0, 0)}{\sum_{t'} \eta(t', t)}, \quad e^{\mu b_t} = \frac{\eta(0, 0)}{\sum_{t'=0}^t \eta(t', t)} = \frac{\eta(0, 0)}{\sum_{t'=0}^t \eta(t', t'')} \quad \forall t'' \geq t. \tag{39}$$

One can consistently extend these definitions with $b_0 = 0$, and it will be useful to adopt the convention $e^{-\mu b_{-1}} = 0$ in order to simplify some expressions. Let us denote $h = (a_0, a_1, \dots, a_T, b_{T-1}, \dots, b_1)$ the vector of $2T$ reals encoding in this way a matrix η ; h will be called a cavity field in the following (note that we excluded b_T which disappears from the final expressions). The recursion relations (36–38) should now be transformed into a relation between cavity fields, i.e. $h = g(h_1, \dots, h_k)$, with $h_i = (a_0^{(i)}, a_1^{(i)}, \dots, a_T^{(i)}, b_{T-1}^{(i)}, \dots, b_1^{(i)})$. Inserting the definitions (39) into the Eqs. (36–38) leads to the explicit form for g ,

$$\begin{aligned} e^{-\mu a_t} &= 1 + e^{-\mu} \sum_{t'=1}^T \sum_{\substack{I, J, K \\ |I| + \mathbb{I}(t' \geq t+2) \leq l-1 \\ |I| + |J| + \mathbb{I}(t' \geq t+1) \geq l}} \mathcal{P}_{t'}(h_1, \dots, h_k; I, J, K) + e^{-\mu-\epsilon} \\ &\quad \times \sum_{\substack{I, J, K \\ |I| + |J| + \mathbb{I}(t \leq T-1) \leq l-1}} \mathcal{P}_T(h_1, \dots, h_k; I, J, K) \\ e^{-\mu b_t} &= 1 + e^{-\mu} \sum_{t'=1}^t \sum_{\substack{I, J, K \\ |I| \leq l-1 \\ |I| + |J| \geq l}} \mathcal{P}_{t'}(h_1, \dots, h_k; I, J, K) \end{aligned} \tag{40}$$

where we defined

$$\begin{aligned} \mathcal{P}_t(h_1, \dots, h_k; I, J, K) &= e^{\mu \sum_{i=1}^k a_0^{(i)}} \prod_{i \in I} e^{-\mu b_{t-2}^{(i)}} \prod_{i \in J} (e^{-\mu b_{t-1}^{(i)}} - e^{-\mu b_{t-2}^{(i)}}) \\ &\quad \times \prod_{i \in K} (e^{-\mu a_t^{(i)}} - e^{-\mu b_{t-1}^{(i)}}). \end{aligned} \tag{41}$$

It can be checked that for $T = 1$ and $\epsilon = +\infty$ these equations correspond, as they should, to the one of the Biroli–Mézarad model (see Eqs. (108, 109) of [70]). One can also express the partial partition functions $\widehat{z}_{\text{site}}$ and $\widehat{z}_{\text{edge}}$ in terms of these fields. It will be more convenient to factor out a common part in the site and edge contributions to the free-entropy. Denoting $r(\eta) = \sum_t \eta(t, 0)$, we define z_{edge} as:

$$\begin{aligned} z_{\text{edge}}(h_1, h_2) &= \frac{\widehat{z}_{\text{edge}}(\eta_1, \eta_2)}{r(\eta_1)r(\eta_2)} \\ &= e^{\mu(a_0^{(1)}+a_0^{(2)})} \left\{ e^{-\mu(a_T^{(1)}+a_T^{(2)})} + \sum_{t=0}^{T-1} \left[(e^{-\mu a_t^{(1)}} - e^{-\mu a_{t+1}^{(1)}}) e^{-\mu b_t^{(2)}} \right. \right. \\ &\quad \left. \left. + e^{-\mu b_t^{(1)}} (e^{-\mu a_t^{(2)}} - e^{-\mu a_{t+1}^{(2)}}) \right] \right\}, \end{aligned} \tag{42}$$

where the explicit expression is obtained from Eq. (30). Similarly, exploiting Eq. (29), we get for the site term (factoring also a contribution from the chemical potential):

$$\begin{aligned} z_{\text{site}}(h_1, \dots, h_{k+1}; l, \epsilon; \mu) &= \frac{e^{-\mu \widehat{z}_{\text{site}}(\eta_1, \dots, \eta_{k+1}; l, \epsilon; \mu)}}{r(\eta_1) \dots r(\eta_{k+1})} \\ &= 1 + e^{-\mu} \sum_{t=1}^T \sum_{\substack{I, J, K \\ |I| \leq l-1 \\ |I|+|J| \geq l}} \mathcal{P}_t(h_1, \dots, h_{k+1}; I, J, K) + e^{-\mu-\epsilon} \\ &\quad \times \sum_{\substack{I, J, K \\ |I|+|J| \leq l-1}} \mathcal{P}_T(h_1, \dots, h_{k+1}; I, J, K) \end{aligned} \tag{43}$$

where as above in the summations I, J, K denotes a partition of $\{1, \dots, k + 1\}$.

To summarize the results of this reparametrization, on a given graph one has cavity fields $h_{i \rightarrow j}$ on each directed edge, obeying the Belief Propagation equations $h_{i \rightarrow j} = g(\{h_{k \rightarrow i}\}_{k \in \partial i \setminus j})$, with the g defined in Eq. (40), and the Bethe free-entropy density is computed from these cavity fields according to

$$\phi = \mu + \frac{1}{N} \sum_i \ln z_{\text{site}}(\{h_{j \rightarrow i}\}_{j \in \partial i}; l_i, \epsilon, \mu) - \frac{1}{N} \sum_{\langle i, j \rangle \in E} \ln z_{\text{edge}}(h_{i \rightarrow j}, h_{j \rightarrow i}), \tag{44}$$

with z_{site} and z_{edge} defined in Eqs. (43) and (42) respectively. Note that the factors r introduced in the definitions of z_{site} and z_{edge} compensate because in the expression of the Bethe free-energy of Eq. (28) the messages on each directed edge appear exactly once in the site term and once in the edge term. The marginals of the law $\eta(\underline{\sigma}, \underline{t})$ can also be computed from the cavity fields h , in particular from the expression (31) one obtains the marginal of one activation time from the incident cavity fields as

$$\eta_{\text{site}}(h_1, \dots, h_{k+1}; l, \epsilon; \mu)(t) = \frac{1}{z_{\text{site}}(h_1, \dots, h_{k+1}; l, \epsilon; \mu)} \times \left\{ \begin{aligned} &\delta_{t,0} + (1 - \delta_{t,0} - \delta_{t,\infty})e^{-\mu} \sum_{\substack{I, J, K \\ |I| \leq l-1 \\ |I|+|J| \geq l}} \mathcal{P}_t(h_1, \dots, h_{k+1}; I, J, K) \\ &+ \delta_{t,\infty} e^{-\mu-\epsilon} \sum_{\substack{I, J, K \\ |I|+|J| \leq l-1}} \mathcal{P}_T(h_1, \dots, h_{k+1}; I, J, K) \end{aligned} \right\}. \tag{45}$$

3.2.3 Random (Regular) Graphs

The replica symmetric cavity method, for generic models defined on sparse random graphs, postulates the asymptotic validity of the above computations, exact on finite trees, thanks to the local convergence of random graphs to trees and an assumption of correlation decay at large distance. The order parameter is then a probability distribution over cavity fields, the randomness arising from the fluctuations of the degrees of the vertices in the graph and/or the randomness in the local interactions.

In the case of random regular graphs with no disorder in the coupling the situation is even simpler, as one can look for a “factorized” solution with all cavity fields equal. In particular for the model at hand on a $k + 1$ random regular graph, with the same threshold of activation l for all vertices, the RS prediction for the typical free-entropy density in the thermodynamic limit defined in Eq. (15) reads

$$\phi(k, l, T, \mu, \epsilon) = \mu + \ln(z_{\text{site}}(h, \dots, h)) - \frac{k + 1}{2} \ln(z_{\text{edge}}(h, h)), \tag{46}$$

which is easily obtained from (44) noting that $2M = (k + 1)N$ in a regular graph. The field h is the fixed-point solution of the cavity recursion (40),

$$h = g(h, \dots, h). \tag{47}$$

The marginal law for the activation time is obtained from (45) by setting all the fields to h , which allows finally to compute the entropy density from the Legendre inverse transform discussed in (13).

We shall discuss the results obtained from this RS prediction in the next Section, more explicit formulas for the RS equation in this case, along with some technical details on their resolution being displayed in the Appendix 1. One can however anticipate that in some regime of parameters the RS hypothesis will be violated. This is for instance known for $T = 1, \epsilon = +\infty$, which corresponds to the Biroli–Mézarard model; it was indeed shown in [70] that for large negative values of μ the predictions of the RS ansatz are unphysical, the frustration arising from the contradictory constraints of putting as few active vertices in the initial condition as possible while imposing that all vertices become active at a latter time induces long-range correlations between variables that are incompatible with the RS ansatz. This limit $\mu \rightarrow -\infty$ being the interesting case for the computations of the minimal density of contagious sets, we shall now see how to include the effects of replica symmetry breaking in this model.

3.3 One Step of Replica Symmetry Breaking (1RSB) Formalism

The long-range correlation decay assumption underlying the RS cavity method breaks down for models with too much frustration. In this case one has to picture the configuration space as fractured into pure states, or clusters, that we shall index here by γ , such that the correlation decay assumption only holds for the Gibbs–Boltzmann probability law restricted to one pure-state. The partition function restricted to a given pure-state is denoted Z_γ , in such a way that $Z = \sum_\gamma Z_\gamma$. The replica symmetry breaking version of the cavity method then postulates some properties of this decomposition into pure states, which are compatible with the local convergence of the graph under study to a tree. In the first non-trivial version of the RSB formalism, so called one-step RSB (1RSB), one assumes the existence of a complexity function, also called configurational entropy in the context of glasses, $\Sigma(\phi)$, such that the number of pure states with an internal free-entropy density $\phi_\gamma = \frac{1}{N} \ln Z_\gamma$ close to some value ϕ is, at the leading exponential order, $e^{N\Sigma(\phi)}$. The computation of $\Sigma(\phi)$ is performed via the 1RSB potential with a parameter m (known as the Parisi breaking parameter), related to Σ through a Legendre transform structure [58]:

$$\Phi(m) = \frac{1}{N} \ln \sum_\gamma Z_\gamma^m = \sup_\phi [\Sigma(\phi) + m\phi]. \quad (48)$$

The function $\Sigma(\phi)$ can be reconstructed in a parametric way varying m , with

$$\Sigma(\phi_{\text{int}}(m)) = \Phi(m) - m\phi_{\text{int}}(m), \quad \phi_{\text{int}}(m) = \Phi'(m), \quad (49)$$

$\phi_{\text{int}}(m)$ denoting the internal free-entropy density of the clusters selected by the corresponding value of m . The value $m = 1$ plays a special role in this approach, as it corresponds a priori to the original computation of the free-entropy density of the model. However a so-called condensation (or Kauzmann) transition can occur, signaled by the vanishing of the complexity Σ associated to $m = 1$. In this case the Gibbs–Boltzmann measure is dominated by a sub-exponential number of pure-states, corresponding to a parameter $m_s < 1$ with $\Sigma(m_s) = 0$. In the following paragraphs we shall derive the 1RSB equations and the expression of the 1RSB potential for the model under study, before discussing the concrete results for random regular graphs in the next Section.

3.3.1 Single Sample Equations

Let us first discuss the 1RSB formalism with the basic messages represented in terms of the matrices $\eta(t, t')$. In the RS description one had a message $\eta_{i \rightarrow j}$ on each directed edge of the graph, solution of the recurrence equations $\eta_{i \rightarrow j} = \widehat{g}(\{\eta_{k \rightarrow i}\}_{k \in \partial i \setminus j}; l_i, \epsilon, \mu)$, see Eq. (27). At the 1RSB level one introduces instead a distribution $\widehat{P}_{i \rightarrow j}(\eta)$ on each directed edge, the randomness being over the choice of the pure-state γ with a weight proportional to Z_γ^m . These distributions are thus found to obey the recurrence equations $\widehat{P}_{i \rightarrow j} = \widehat{G}[\{\widehat{P}_{k \rightarrow i}\}_{k \in \partial i \setminus j}]$, where $\widehat{P} = \widehat{G}(\widehat{P}_1, \dots, \widehat{P}_k)$ means

$$\widehat{P}(\eta) = \frac{1}{\widehat{Z}_{\text{iter}}(\widehat{P}_1, \dots, \widehat{P}_k)} \int d\widehat{P}_1(\eta_1) \dots d\widehat{P}_k(\eta_k) \delta(\eta - \widehat{g}(\eta_1, \dots, \eta_k)) \widehat{z}_{\text{iter}}(\eta_1, \dots, \eta_k)^m, \quad (50)$$

with \widehat{g} and $\widehat{z}_{\text{iter}}$ defined in Eq. (27), and $\widehat{z}_{\text{iter}}$ normalizes the distribution \widehat{P} . The 1RSB potential $\Phi(m)$ defined above is then computed from the solution of these equations, according to

$$\Phi(m) = \frac{1}{N} \sum_{i=1}^N \ln \widehat{z}_{\text{site}}(\{\widehat{P}_{j \rightarrow i}\}_{j \in \partial i}; l_i, \epsilon_i, \mu_i) - \frac{1}{N} \sum_{(i,j) \in E} \ln \widehat{z}_{\text{edge}}(\widehat{P}_{i \rightarrow j}, \widehat{P}_{j \rightarrow i}), \quad (51)$$

where

$$\widehat{z}_{\text{site}}(\widehat{P}_1, \dots, \widehat{P}_{k+1}) = \int d\widehat{P}_1(\eta_1) \dots \widehat{P}_{k+1}(\eta_{k+1}) \widehat{z}_{\text{site}}(\eta_1, \dots, \eta_{k+1})^m, \quad (52)$$

$$\widehat{z}_{\text{edge}}(\widehat{P}_1, \widehat{P}_2) = \int d\widehat{P}_1(\eta_1) \widehat{P}_2(\eta_2) \widehat{z}_{\text{edge}}(\eta_1, \eta_2)^m \quad (53)$$

are weighted averages, over the pure-states distribution, of the site and edge contributions to the free-entropy defined in (29, 30). Similarly the marginal distribution of an activation time can be computed as a weighted average of the RS expression in the various pure-states, i.e.

$$\eta(t) = \frac{1}{\widehat{z}_{\text{site}}(\widehat{P}_1, \dots, \widehat{P}_{k+1})} \int d\widehat{P}_1(\eta_1) \dots \widehat{P}_{k+1}(\eta_{k+1}) \widehat{\eta}_{\text{site}}(\eta_1, \dots, \eta_{k+1})(t) \widehat{z}_{\text{site}}(\eta_1, \dots, \eta_{k+1})^m. \quad (54)$$

Note that the derivative $\Phi'(m)$, which plays an important role to compute the complexity from Eq. (49), can be taken in (51) on the explicit dependence on m only, the recursion relations on the $\widehat{P}_{i \rightarrow j}$ being the stationarity conditions of (51) with respect to the \widehat{P} 's.

As we have seen in the discussion of the RS cavity method the matrices η can be parametrized in a more economic way by the fields h (vectors of $2T$ real numbers). The expressions of the 1RSB quantities can also be rewritten using this parametrization. After a few lines of algebra one finds that the potential $\Phi(m)$ reads

$$\Phi(m) = \mu m + \sum_{i=1}^N \ln \mathcal{Z}_{\text{site}}(\{P_{j \rightarrow i}\}_{j \in \partial i}; l_i, \epsilon, \mu) - \sum_{(i,j) \in E} \ln \mathcal{Z}_{\text{edge}}(P_{i \rightarrow j}, P_{j \rightarrow i}), \quad (55)$$

with

$$\mathcal{Z}_{\text{site}}(P_1, \dots, P_{k+1}) = \int dP_1(h_1) \dots P_{k+1}(h_{k+1}) z_{\text{site}}(h_1, \dots, h_{k+1})^m, \quad (56)$$

$$\mathcal{Z}_{\text{edge}}(P_1, P_2) = \int dP_1(h_1) P_2(h_2) z_{\text{edge}}(h_1, h_2)^m, \quad (57)$$

the weighted averages of the quantities defined in (42, 43). The field distributions $P_{i \rightarrow j}(h)$ are solutions of the recurrence equations $P_{i \rightarrow j} = G(\{P_{k \rightarrow i}\}_{k \in \partial i \setminus j})$, where the mapping $P = G(P_1, \dots, P_k)$ is given explicitly by

$$P(h) = \frac{1}{\mathcal{Z}_{\text{iter}}(P_1, \dots, P_k)} \int dP_1(h_1) \dots dP_k(h_k) \delta(h - g(h_1, \dots, h_k)) z_{\text{iter}}(h_1, \dots, h_k)^m. \quad (58)$$

$\mathcal{Z}_{\text{iter}}$ is a normalizing factor ensuring that the left hand side is a probability distribution, g is the function defined in Eq. (40), and the reweighting factor reads

$$z_{\text{iter}}(h_1, \dots, h_k) = \frac{e^{-\mu} \widehat{z}_{\text{iter}}(\eta_1, \dots, \eta_k) r(\widehat{g}(\eta_1, \dots, \eta_k))}{r(\eta_1) \dots r(\eta_k)} = e^{-\mu a_0(h_1, \dots, h_k)}, \quad (59)$$

the last equality following from Eqs. (36, 39).

3.3.2 Random Regular Graphs

For the reasons explained in the context of the RS ansatz in Sect. 3.2.3 one can look for a simple factorized solution of the 1RSB equations in the case of a $k + 1$ regular random graph with all thresholds of activation equal to l . In this case one has to find a distribution $P(h)$ solution of

$$P(h) = \frac{1}{Z_{\text{iter}}} \int dP(h_1) \dots dP(h_k) \delta(h - g(h_1, \dots, h_k)) z_{\text{iter}}(h_1, \dots, h_k)^m, \quad (60)$$

where $m \in [0, 1]$ is the Parisi breaking parameter and the functions g and z_{iter} are the ones defined in Eqs. (40, 59). The 1RSB potential is then computed as

$$\begin{aligned} \Phi(m) = \mu m + \ln & \left(\int dP(h_1) \dots dP(h_{k+1}) z_{\text{site}}(h_1, \dots, h_{k+1})^m \right) \\ & - \frac{k + 1}{2} \ln \left(\int dP(h_1) dP(h_2) z_{\text{edge}}(h_1, h_2)^m \right), \end{aligned} \quad (61)$$

with the functions $z_{\text{site}}, z_{\text{edge}}$ defined in Eqs. (42, 43). As already mentioned above $\Phi'(m)$ can be computed by taking into account only the explicit dependence on m of (61).

3.4 “Energetic” 1RSB Formalism

Even within the simplified case of the factorized ansatz for regular graphs the 1RSB equations are relatively complicated, as they involve the resolution of a distributional equation on $P(h)$. However we are ultimately interested in a particular limit for the computation of the minimal density of contagious sets, namely the case where $\epsilon = +\infty$ (to take into account only the fully activating configurations), and in the limit $\mu \rightarrow -\infty$ (to select the initial configurations with the minimal number of active sites). It turns out that a simplified version of the 1RSB formalism can be devised in this case, corresponding to the “energetic” version of the 1RSB cavity method, first developed in [55, 56], see in particular Sect. 5 of [70] for such a treatment of the related Biroli–Mézard model. This simplified treatment amounts to take simultaneously the limit $m \rightarrow 0$ and $\mu \rightarrow -\infty$, with a fixed finite value of a new parameter $y = -\mu m$. To explain the meaning of this limit let us rewrite more explicitly the expression of the 1RSB potential of Eq. (48) in the case $\epsilon = +\infty$, introducing the complexity $\Sigma(s, \theta)$ counting the (exponential) number of clusters containing a number of order e^{Ns} of activating initial configurations with a fraction θ of active sites, hence with a free-entropy density $\phi = \mu\theta + s$:

$$\Phi(m) = \sup_{\theta, s} [\Sigma(s, \theta) + m(\mu\theta + s)]. \quad (62)$$

In the limit $m \rightarrow 0, \mu \rightarrow -\infty$ with $y = -\mu m$ this function becomes

$$\Phi_e(y) = \sup_{\theta} [\Sigma_e(\theta) - y\theta], \quad \Sigma_e(\theta) = \sup_s \Sigma(s, \theta). \quad (63)$$

The “energetic” complexity $\Sigma_e(\theta)$ can thus be computed via an inverse Legendre transform of the potential $\Phi_e(y)$,

$$\Sigma_e(\theta(y)) = \Phi_e(y) + y\theta(y), \quad \theta(y) = -\Phi'_e(y). \quad (64)$$

As we shall see the “energetic” 1RSB cavity equations leading to the computations of $\Phi_e(y)$ are much simpler than the initial 1RSB ones at finite values of μ and m . The price to pay for this simplification is the loss of information on the entropy of the clusters when going from

$\Sigma(s, \theta)$ to $\Sigma_e(\theta)$. However this is not a problem for the determination of θ_{\min} : its estimate at the 1RSB level, to be denoted $\theta_{\min,1}$, is the smallest value of θ with $\Sigma_e(\theta) \geq 0$. Indeed the least dense activating configurations have to be in some pure states, whatever their entropy.

3.4.1 Simplification of the Cavity Field Recursion (Warning Propagation Equations)

We want to simplify the Eq. (40) giving $h = g(h_1, \dots, h_k)$ with $\epsilon = +\infty$ and in the limit $\mu \rightarrow -\infty$. First let us make some remarks, valid when $\epsilon = +\infty$ for any value of μ . From the definition (39) of the fields b_t , or from their expressions in (40), it is obvious that

$$e^{-\mu b_T} \geq e^{-\mu b_{T-1}} \geq \dots \geq e^{-\mu b_1} \geq e^{-\mu b_0} = 1. \tag{65}$$

One can also notice that for $\epsilon = +\infty$ one has, for any μ , the equality $a_T = b_T$: this appears both from the definition (39) of the fields, as $\eta(\infty, t) = 0$ when $\epsilon = +\infty$, and from the recursion relations (40), the last term in the first line of (40) disappearing when $\epsilon = +\infty$. To continue the above chain of inequalities let us first compute from (40)

$$e^{-\mu a_{T-1}} - e^{-\mu a_T} = e^{-\mu + \mu \sum_{i=1}^k a_0^{(i)}} \sum_{\substack{I, J \\ |I|=l-1}} \prod_{i \in I} e^{-\mu b_{T-1}^{(i)}} \prod_{i \in J} (e^{-\mu b_T^{(i)}} - e^{-\mu b_{T-1}^{(i)}}), \tag{66}$$

where I, J forms a partition of $\{1, \dots, k\}$. This shows that $e^{-\mu a_{T-1}} \geq e^{-\mu a_T} = e^{-\mu b_T}$, because in the right-hand side $e^{-\mu b_T^{(i)}} \geq e^{-\mu b_{T-1}^{(i)}}$. These inequalities can then be continued by recurrence, as for $t \in \{0, \dots, T - 2\}$ one obtains from (40)

$$e^{-\mu a_t} - e^{-\mu a_{t+1}} = e^{-\mu + \mu \sum_{i=1}^k a_0^{(i)}} \sum_{\substack{I, J \\ |I|=l-1}} \prod_{i \in I} e^{-\mu b_t^{(i)}} \left(\prod_{i \in J} (e^{-\mu a_{t+1}^{(i)}} - e^{-\mu b_t^{(i)}}) - \prod_{i \in J} (e^{-\mu a_{t+2}^{(i)}} - e^{-\mu b_t^{(i)}}) \right), \tag{67}$$

hence

$$e^{-\mu a_0} \geq a^{-\mu a_1} \geq \dots \geq e^{-\mu a_{T-1}} \geq e^{-\mu a_T} = e^{-\mu b_T} \geq e^{-\mu b_{T-1}} \geq \dots \geq e^{-\mu b_1} \geq e^{-\mu b_0} = 1, \tag{68}$$

and for any $\mu \leq 0$:

$$a_0 \geq a_1 \geq \dots \geq a_{T-1} \geq a_T = b_T \geq b_{T-1} \geq \dots \geq b_1 \geq b_0 = 0. \tag{69}$$

Let us now take the limit $\mu \rightarrow -\infty$ in the Eq. (40), assuming that a_t and b_t have finite limits. Treating (40) at the leading exponential order one obtains

$$a_t = \max \left(0, \max_{t' \in [1, T]} \max_{\substack{I, J, K \\ |I| + \mathbb{I}(t' \geq t+2) \leq l-1 \\ |I| + |J| + \mathbb{I}(t' \geq t+1) \geq l}} S_{t'}(h_1, \dots, h_k; I, J, K) \right), \tag{70}$$

$$b_t = \max \left(0, \max_{t' \in [1, t]} \max_{\substack{I, J, K \\ |I| \leq l-1 \\ |I| + |J| \geq l}} S_{t'}(h_1, \dots, h_k; I, J, K) \right), \tag{71}$$

where

$$S_t(h_1, \dots, h_k; I, J, K) = 1 - \sum_{i \in I} (a_0^{(i)} - b_{t-2}^{(i)}) - \sum_{i \in J} (a_0^{(i)} - b_{t-1}^{(i)}) - \sum_{i \in K} (a_0^{(i)} - a_t^{(i)}). \tag{72}$$

Now from the inequalities (69) it appears that $S_t \leq 1$, hence that the a 's and b 's belong to the interval $[0, 1]$. It is however natural to assume that they are integers, as in the limit $\mu \rightarrow -\infty$ they can be interpreted as differences between number of particles in constrained groundstate configurations (see [55,70] for more details). Within this ansatz the a 's and b 's can only be equal to 0 or 1; using in addition the inequalities (69) one realizes that the fields h can only take $2T + 1$ possible values, that we shall call A_t for $t \in \{0, 1, \dots, T - 1\}$ and B_t for $t \in \{0, 1, \dots, T\}$. These are defined as follows; A_t denotes the case where $a_0 = \dots = a_t = 1$, all the other a 's and b 's vanishing. For $t \in \{2, \dots, T\}$, B_t means that $b_1 = \dots = b_{t-1} = 0$, all the other a 's and b 's being equal to 1. Finally B_1 corresponds to the case where all a 's and b 's are equal to 1, and B_0 to the case where they all vanish. Note that one can consistently extend these definitions to $A_T = B_T$, as by definition $a_T = b_T$.

It remains to determine the value of $h = g(h_1, \dots, h_k)$ in this $\mu \rightarrow -\infty$ limit, when all the fields h_1, \dots, h_k belong to the set $\{A_0, A_1, \dots, A_{T-1}, A_T = B_T, B_{T-1}, \dots, B_1, B_0\}$ of ‘‘hard fields’’, or Warning Propagation messages. Some algebra, sketched in Appendix 1, leads to:

$$g(B_{t_1}, \dots, B_{t_n}, A_{t_{n+1}}, \dots, A_{t_k}) = \begin{cases} B_{1+\min_l(t_1, \dots, t_n)} & \text{if } n \geq l \text{ and } \min(t_{n+1}, \dots, t_k) \geq 1 + \min_l(t_1, \dots, t_n) \\ A_{\min(t_{n+1}, \dots, t_k)-1} & \text{if } n \geq l - 1 \text{ and} \\ & 1 + \min_{l-1}(t_1, \dots, t_n) \leq \min(t_{n+1}, \dots, t_k) \leq \min_l(t_1, \dots, t_n) \\ B_0 & \text{otherwise} \end{cases} \tag{73}$$

where $t_1, \dots, t_n \in \{0, \dots, T - 1\}$ and $t_{n+1}, \dots, t_k \in \{0, \dots, T\}$. We assumed conventionally that $\min_l(t_1, \dots, t_{l-1}) = \infty$.

The Eq. (73) can be given a very intuitive interpretation. The messages $h \in \{A_0, \dots, A_{T-1}, B_0, \dots, B_T\}$ can be interpreted as ‘‘warnings’’ sent from one vertex of the graph to one of its neighbours, with the following meanings. A vertex i sends a message $h_{i \rightarrow j} = B_t$ to one of its neighbour j to say: ‘‘if j is kept inactive at all times the configuration of i and of its sub-tree (the one rooted at i and excluding j) leads to complete activation of the sub-tree within the time horizon T , and i activates itself at time t ’’. In particular $h_{i \rightarrow j} = B_0$ means that i is activated in the initial configuration. On the contrary i sends the message $h_{i \rightarrow j} = A_t$ to j to express: ‘‘the complete activation of i and its sub-tree requires that j becomes activated at time t ’’. The rules of Eq. (73) for the combination of these messages are then obtained by finding the configuration compatible with them, containing the minimal number of active sites in the initial configuration (because of the $\mu \rightarrow -\infty$ limit):

- if strictly less than $l - 1$ incoming messages are of the type B_{t_i} , with $t_i \in \{0, \dots, T - 1\}$, the central site i will never have more than l active neighbours (even with the participation of the receiving site j) if it is initially inactive, hence the only way for i to be active at time T is to be active in the initial configuration, which implies $h_{i \rightarrow j} = B_0$.
- if at least l of the incoming messages are of the type B_{t_i} , with $t_i \in \{0, \dots, T - 1\}$, say $(B_{t_1}, \dots, B_{t_n})$, the central site i will become active at time $t = 1 + \min_l(t_1, \dots, t_n)$, without the ‘‘help’’ of the activation of the site j receiver of the message. This situation thus leads to a message of type B_t , at the condition that all other incoming messages

of type $\{A_0, \dots, A_T\}$ do not require the activation of the central site i at a time strictly earlier than $t = 1 + \min_l(t_1, \dots, t_n)$.

- the participation of the activation of the receiving site j is required at some time t when the above condition is not fulfilled, i.e. when the incoming messages $(A_{t_{n+1}}, \dots, A_{t_k})$ require the activation of the central site at some time $t_{act} = \min(t_{n+1}, \dots, t_k) < 1 + \min_l(t_1, \dots, t_n)$. This mechanism is possible if at time $t_{act} - 1$ already $l - 1$ of the neighbours sending messages of type B are active, i.e. it requires $\min_{l-1}(t_1, \dots, t_n) \leq t_{act} - 1$. The “help” needed from the receiving site is that it is active at some time before $t_{act} - 1$; in the limit $\mu \rightarrow -\infty$ the least dense configurations, and thus the least stringent constraint on the time of activation is privileged, hence the message sent in this case is $h_{i \rightarrow j} = A_{t_{act}-1}$.
- all cases not fulfilling one of the conditions above require that i is active in the initial configuration to be active at time T , hence the message sent is $h_{i \rightarrow j} = B_0$.

3.4.2 Energetic 1RSB Single Sample Equations

Within this ansatz the 1RSB distributions $P(h)$ greatly simplify, as they are supported on the discrete set $h \in \{A_0, A_1, \dots, A_{T-1}, A_T = B_T, B_{T-1}, \dots, B_1, B_0\}$. We shall denote p_t the weight in $P(h)$ of the event $h = A_t$, and similarly q_t for $h = B_t$ (with again the convention $p_T = q_T$ to simplify notations), i.e.

$$P(h) = \sum_{t=0}^{T-1} p_t \delta(h - A_t) + \sum_{t=0}^T q_t \delta(h - B_t). \tag{74}$$

The 1RSB recursion relation (58) now reduces to a recursion between these finite-dimensional vectors of probabilities; inserting the definition (74) in the right hand side of (58) and exploiting the combination rule (73) between hard fields, one obtains the following limit for the recursion relation $P = G[P_1, \dots, P_k]$:

$$\begin{aligned} p_t &= \frac{1}{Z[P_1, \dots, P_k]} e^y \tilde{p}_t, \quad \tilde{p}_t = \sum_{\substack{I, J, K \\ |I|=l-1 \\ |J| \geq 1}} \prod_{i \in I} \left(\sum_{t'=0}^t q_{t'}^{(i)} \right) \prod_{i \in J} p_{t+1}^{(i)} \prod_{i \in K} \\ &\quad \times \left(\sum_{t'=t+1}^T q_{t'}^{(i)} + \sum_{t'=t+2}^{T-1} p_{t'}^{(i)} \right) \text{ for } t \in \{0, \dots, T-1\} \\ q_t &= \frac{1}{Z[P_1, \dots, P_k]} e^y \tilde{q}_t, \quad \tilde{q}_t = \sum_{\substack{I, J, K \\ |I| \leq l-1 \\ |I|+|J| \geq l}} \prod_{i \in I} \\ &\quad \times \left(\sum_{t'=0}^{t-2} q_{t'}^{(i)} \right) \prod_{i \in J} q_{t-1}^{(i)} \prod_{i \in K} \left(\sum_{t'=t}^T q_{t'}^{(i)} + \sum_{t'=t}^{T-1} p_{t'}^{(i)} \right) \text{ for } t \in \{1, \dots, T\} \\ q_0 &= \frac{1}{Z[P_1, \dots, P_k]} \left[1 - \sum_{t=0}^{T-1} \tilde{p}_t - \sum_{t=1}^T \tilde{q}_t \right], \\ Z[P_1, \dots, P_k] &= 1 + (e^y - 1) \left[\sum_{t=0}^{T-1} \tilde{p}_t + \sum_{t=1}^T \tilde{q}_t \right] \end{aligned} \tag{75}$$

the reweighting term of Eq. (59) becoming indeed $z_{\text{iter}}(h_1, \dots, h_k)^m = e^{ya_0(h_1, \dots, h_k)}$, hence the factor e^y multiplying the probabilities of all warnings except B_0 ; this is indeed the only case where an active site has to be inserted in the initial configuration.

To compute the 1RSB potential we have to study the limit of the contribution of site and edge terms in the limit $\mu \rightarrow -\infty, m \rightarrow 0$. We have from Eq. (43)

$$z_{\text{site}}(h_1, \dots, h_{k+1})^m \rightarrow \exp \left[y \max \left(0, \max_{t \in [1, T]} \max_{\substack{I, J, K \\ |I| \leq t-1 \\ |I|+|J| \geq t}} \mathcal{S}_t(h_1, \dots, h_{k+1}; I, J, K) \right) \right], \tag{76}$$

which can be simplified following the same reasoning than the one which led to (73). This yields

$$Z_{\text{site}}(P_1, \dots, P_{k+1}) \rightarrow 1 + (e^y - 1) \sum_{t=1}^T \sum_{\substack{I, J, K \\ |I| \leq t-1 \\ |I|+|J| \geq t}} \prod_{i \in I} \left(\sum_{t'=0}^{t-2} q_{t'}^{(i)} \right) \prod_{i \in J} q_{t-1}^{(i)} \prod_{i \in K} \left(\sum_{t'=t}^T q_{t'}^{(i)} + \sum_{t'=t}^{T-1} p_{t'}^{(i)} \right), \tag{77}$$

where I, J, K is a partition of $\{1, \dots, k+1\}$. This expression can be interpreted intuitively in terms of the warnings defined above; the factor multiplying $(e^y - 1)$ is indeed the probability of complete activation, at time $t \in \{1, \dots, T\}$, for an initially empty site receiving messages (h_1, \dots, h_{k+1}) from its neighbours, with their respective distributions P_1, \dots, P_{k+1} . As a matter of fact, for its activation to occur at time t at least l neighbours must have activated without any help from the central site at time $t - 1$, no more than $l - 1$ must be active at time $t - 2$ (otherwise the activation time would be strictly less than t), and the neighbours sending messages of type $A_{t'}$ should not require activation at a time $t' < t$.

For the edge term we obtain from Eq. (42)

$$z_{\text{edge}}(h_1, h_2)^m \rightarrow \exp \left[-y \min_{t \in [0, T]} \min((a_0^{(1)} - b_t^{(1)}) + (a_0^{(2)} - a_t^{(2)}), (a_0^{(1)} - a_t^{(1)}) + (a_0^{(2)} - b_t^{(2)})) \right], \tag{78}$$

hence

$$Z_{\text{edge}}(P_1, P_2) \rightarrow e^{-y} + (1 - e^{-y}) \left[\left(\sum_{t=0}^T q_t^{(1)} \right) \left(\sum_{t=0}^T q_t^{(2)} \right) + \sum_{t=0}^{T-1} p_t^{(1)} \sum_{t'=0}^t q_{t'}^{(2)} + \sum_{t=0}^{T-1} p_t^{(2)} \sum_{t'=0}^t q_{t'}^{(1)} \right]. \tag{79}$$

One can interpret the factor multiplying $(1 - e^{-y})$ as the probability of complete activation when two messages (h_1, h_2) drawn with the probabilities P_1, P_2 are sent in the two opposite directions of an edge.

Let us summarize the main findings of this subsection. In the limit $\mu \rightarrow -\infty, m \rightarrow 0$ with $y = -\mu m$ the 1RSB formalism simplifies in the following way. The cavity field distributions $P_{i \rightarrow j}(h)$ have now a discrete support with $2T$ possible values, each of them is thus described by a (normalized) vector of $2T$ probabilities denoted $\{p_t, q_t\}$. These vectors are solutions of

recurrence equations of the form $P_{i \rightarrow j} = G(\{P_{k \rightarrow i}\}_{k \in \partial i \setminus j})$, the mapping G being defined in Eq. (75). The energetic limit of the 1RSB potential is then computed as

$$\Phi_e(y) = -y + \frac{1}{N} \sum_{i=1}^N \ln(\mathcal{Z}_{\text{site}}(\{P_{j \rightarrow i}\}_{j \in \partial i})) - \frac{1}{N} \sum_{\langle i, j \rangle} \ln(\mathcal{Z}_{\text{edge}}(P_{i \rightarrow j}, P_{j \rightarrow i})), \quad (80)$$

with the expression of $\mathcal{Z}_{\text{site}}$ and $\mathcal{Z}_{\text{edge}}$ given in Eqs. (77, 79). This expression of Φ_e is variational, its derivative with respect to y (which is needed in the computation of the inverse Legendre transform in (64)) can be taken on the explicit dependence only.

3.4.3 Random Regular Graphs

For the reasons already exposed in the context of the RS and of the full 1RSB cavity formalism a factorized solution of the energetic 1RSB equations can be searched for when dealing with random $k + 1$ regular graphs with a constant threshold of activation l . One has thus a single vector of probabilities $P = (\{p_t, q_t\})$, fixed-point solution of Eq. (75), from which the energetic 1RSB potential is obtained as

$$\Phi_e(y) = -y + \ln(\mathcal{Z}_{\text{site}}(P, \dots, P)) - \frac{k+1}{2} \ln(\mathcal{Z}_{\text{edge}}(P, P)), \quad (81)$$

with $\mathcal{Z}_{\text{site}}$ and $\mathcal{Z}_{\text{edge}}$ defined in Eqs. (77, 79).

4 Results of the Cavity Method for Random Regular Graphs

We shall present now the results of the resolution of the cavity equations for random regular graphs of degree $k + 1$, with an activation threshold equal to l for all vertices. In all this discussion it will be understood that $\epsilon = +\infty$, i.e. we only consider initial configurations that activate the whole graph in T steps. We will first present in Sect. 4.1 the results for finite values of T , which are qualitatively the same for all values of k, l and T ; the behaviour of the replica symmetric cavity method are first displayed, then we turn to the effects of replica symmetry breaking, in particular in the “energetic” limit to compute the minimal density of initially active sites in activating configurations. In a second part (Sect. 4.2) we shall discuss the limit $T \rightarrow \infty$, in which some further analytical computations can be performed. In this case several qualitatively distinct phenomena emerge, depending on the values of k and l .

4.1 Finite T Results

4.1.1 Replica Symmetric Formalism

The technical details of the resolution of the RS equation $h = g(h, \dots, h)$, where g is given in Eq. (40), and of the computation of the free-entropy density, are deferred to the Appendix 1. From a numerical point of view it is an easy task, as it corresponds essentially to the resolution of a set of $2T$ equations on $2T$ unknowns. Let us discuss the numerical results obtained in this way. On the left panel of Fig. 2 we display the curve $\theta(\mu)$ of the average fraction of initially active sites as a function of the chemical potential μ ; the curve is for $k = l = 2$ and $T = 3$, the qualitative features are independent of these precise values. This function is increasing as it should, and reaches a finite limit when $\mu \rightarrow -\infty$, that would be the candidate value for θ_{\min} if the RS computation was correct in this limit. This however cannot be true,




Fig. 2 The density of initially active sites θ (left panel) and the entropy s (right panel) as a function of the chemical potential μ , computed from the replica symmetric cavity equations, for $k = l = 2$ and $T = 3$




Fig. 3 The RS entropy $s(\theta)$ of configurations with a fraction θ of initially active sites able to activate completely the graph within time T , for $k = l = 2$ (left panel) and $k = 3, l = 2$ (right panel). The curve labelled “random” is the binary entropy function $-\theta \ln \theta - (1 - \theta) \ln(1 - \theta)$ that counts all configurations with such an initial density. The curves in the limit $T \rightarrow \infty$ are computed analytically, from Eq. (82) for the left panel and (101) for the right panel, see Sect. 4.2 for a further discussion of this limit

as revealed from the computation of the entropy, displayed in the right panel of Fig. 2: for $\mu < \mu_{s=0}$ the RS entropy becomes negative, which is a certain indication of the inadequacy of the RS theory in this regime. In Fig. 3 we display the results for the entropy $s(\theta)$ of the number of configurations with a fraction θ of initially active sites, for the regime of positive entropies where the RS prediction cannot be ruled out at once (for the cases $k = l = 2$ and $k = 3, l = 2$). For increasing values of T these curves converge to a limit, this will be further discussed in Sect. 4.2.1. The numerical values of the chemical potential and of the fraction of active sites at the point of entropy cancellation, which would be the best guess from the RS computation of the value of θ_{\min} , denoted respectively $\mu_{s=0}$ and $\theta_{\min,0}$, can be found for various values of T in the Tables 1, 2, and 3 for the cases $k = l = 2, k = l = 3$ and $k = 3, l = 2$ respectively. For $T = 1$ they reproduce, as they should, the results of the Biroli–Mézarid model given in [70].

4.1.2 1RSB Results

As we have seen above the hypothesis underlying the RS computation must go wrong when μ is decreased towards $-\infty$, as the entropy computed within the RS scheme becomes negative for $\mu < \mu_{s=0}$; a 1RSB computation is thus required to investigate the limit $\mu \rightarrow -\infty$ and

Table 1 Numerical results from the cavity computations at finite T for $k = l = 2$; the results in the limit $T \rightarrow \infty$ are explained in Sect. 4.2

T	RS		1RSB				Energetic 1RSB	
	$\mu_{s=0}$	$\theta_{\min,0}$	μ_d	θ_d	μ_c	θ_c	y_s	$\theta_{\min,1}$
1	-7.403996	0.422251	-6.49	0.4292	-6.69	0.4275	5.563433	0.424257
2	-11.374979	0.325742	-9.89	0.3291	-11.23	0.3260	10.826348	0.325882
3	-17.292682	0.289093	-13.7	0.2922	-17.28	0.2890	17.232166	0.289097
4	-24.936318	0.271564	-20.9	0.2731	-24.93	0.2715	24.933659	0.271564
5	-34.966263	0.262167	-31.3	0.2628	-34.63	0.2622	34.966225	0.262167
6	-49.901175	0.256844					49.901175	0.256844
7	-74.984724	0.253779					74.984724	0.253779
8	-120.79085	0.252036					120.79085	0.252036
10	-378.44778	0.250553					378.44778	0.250553
15	-1.069×10^4	0.250018					1.069×10^4	0.250018
20	-3.4×10^5	0.250000					3.4×10^5	0.250000
∞	$-\infty$	$\frac{1}{4}$					$+\infty$	$\frac{1}{4}$

Table 2 Numerical results from the cavity computations at finite T for $k = l = 3$; the results in the limit $T \rightarrow \infty$ are explained in Sect. 4.2

T	RS		1RSB				Energetic 1RSB	
	$\mu_{s=0}$	$\theta_{\min,0}$	μ_d	θ_d	μ_c	θ_c	y_s	$\theta_{\min,1}$
1	-6.113951	0.479455	-5.35	0.4906	-5.39	0.4900	4.644980	0.482712
2	-8.175902	0.397326	-7.38	0.4027	-7.95	0.3988	7.485437	0.397922
3	-10.381917	0.366187	-8.63	0.3725	-10.33	0.3663	10.077681	0.366291
4	-13.140888	0.351221	-9.59	0.3583	-13.11	0.3513	13.037666	0.351234
5	-17.249334	0.343205	-10.3	0.3507	-17.36	0.3432	17.232334	0.343206
6	-24.322138	0.338721					24.321721	0.338721
7	-35.739653	0.336191					35.739653	0.336191
8	-54.198587	0.334760					54.198587	0.334760
∞	$-\infty$	$\frac{1}{3}$					$+\infty$	$\frac{1}{3}$

hence the properties of the least dense activating initial conditions, in particular their density θ_{\min} .


We have thus solved numerically the 1RSB equations (60) using population dynamics methods [54], i.e. representing $P(h)$ as a weighted sample of fields h_i . This method has become fairly standard and we shall not give more details on the procedure, see for instance [53, 54] for detailed presentations. In the particularly important $m = 1$ case we used a version of this procedure, inspired by the tree reconstruction problem, that allows to get rid of the reweighting terms in (60) and is thus much more precise and efficient numerically, see [52, 61] for more technical details.

The results of these investigations follow the usual pattern encountered in constraint satisfaction problems [48]: for large enough values of μ (i.e. for dense enough initial config-

Table 3 Numerical results from the cavity computations at finite T for $k = 3, l = 2$; the results in the limit $T \rightarrow \infty$ are explained in Sect. 4.2

T	RS		1RSB				Energetic 1RSB	
	$\mu_{s=0}$	$\theta_{\min,0}$	μ_d	θ_d	μ_c	θ_c	y_s	$\theta_{\min,1}$
1	-7.730059	0.362794	-7.06	0.3681	-7.38	0.3654	6.778540	0.363813
2	-10.21534	0.236821	-9.16	0.2416	-10.12	0.2372	9.873120	0.237009
3	-11.90150	0.182272	-10.38	0.1875	-11.85	0.1824	11.72892	0.182338
4	-13.03158	0.151659	-11.45	0.1563	-13.00	0.1517	12.92114	0.151693
5	-13.80059	0.132014	-12.47	0.1354	-13.78	0.1321	13.71834	0.132036
6	-14.33193	0.118324					14.26439	0.118341
7	-14.70251	0.108237					14.64332	0.108251
8	-14.96150	0.100498					14.90729	0.100510
10	-15.26375	0.089415					15.21429	0.089425
15	-15.42086	0.074242					15.37163	0.074251
20	-15.27922	0.066569					15.22489	0.066579
30	-14.85174	0.058995					14.78367	0.059008
∞	-12.72072	0.046283					12.54796	0.046328

Fig. 4 The complexity at $m = 1$ as a function of the chemical potential μ , for $k = l = 2$ and $T = 1$. The function is defined for $\mu < \mu_d \approx -6.49$, the complexity being positive for $\mu > \mu_c \approx -6.69$



urations) there is no non-trivial solution of the 1RSB equation with $m = 1$; decreasing μ a non-trivial solution appears discontinuously at a threshold μ_d (the “dynamic” transition). Its complexity (or configurational entropy) Σ is positive in an interval $\mu \in [\mu_c, \mu_d]$, which thus corresponds to a “dynamic 1RSB phase” with an exponential number of clusters contributing to the Gibbs measure, see Fig. 4 for an illustration in the case $T = 1$. The numerical values of μ_d and μ_c (as well as the associated densities of initially active sites θ_d and θ_c), can be found for several values of T in the Tables 1, 2, and 3. For the values of μ in the interval $[\mu_c, \mu_d]$ the thermodynamic predictions of the RS computations are correct. Note that in all the cases we investigated ($k = 2, 3, 2 \leq l \leq k$ and $T \leq 5$) we always found a discontinuous transition with $\mu_c < \mu_d$; we cannot rule out the possibility that for other values of the parameters the replica symmetry breaking transition is continuous with $\mu_c = \mu_d$ (as happens for instance in the independent set problem at low degrees [14]).




Fig. 5 Study of the condensed phase for $k = l = 2$ and $T = 1$. *Left panel* complexity as a function of m for $\mu = -7.5 < \mu_c$, the complexity vanishes for $m_s \approx 0.84$. *Right panel* Parisi parameter m_s as a function of $-1/\mu$, departing from 1 for $\mu < \mu_c$; the *dashed line* corresponds to the linear behaviour $-\mu m_s = 5.56$ that fits the $\mu \rightarrow -\infty$ limit

Lowering further the chemical potential, i.e. in the regime $\mu < \mu_c$, the complexity at $m = 1$ becomes negative. This is thus a true replica symmetry breaking phase with only a sub-exponential number of clusters contributing to the Gibbs measure; μ_c corresponds to the “condensation” transition. In this phase the thermodynamic properties of the model differ from the RS prediction and are given by the properties of the clusters selected by the static value of the Parisi parameter, $m_s(\mu)$, for which the complexity vanishes. This value can be determined by computing the complexity as a function of m , for a fixed value of μ , see left panel of Fig. 5 for an example.

To compute the minimal density $\theta_{\min}(T)$ one has to take the limit $\mu \rightarrow -\infty$; we have introduced above in Sect. 3.4 a simplifying ansatz in this limit, assuming in particular a finite value of $-\mu m$. To check the consistency of this ansatz we solved the complete 1RSB equations for $T = 1$ and several values of μ large and negative. The Parisi parameter m_s is plotted as a function of $-1/\mu$ in the right panel of Fig. 5; in the limit $\mu \rightarrow -\infty$ one obtains indeed a linear behaviour, corresponding to a finite limit of $-\mu m_s$.

4.1.3 Energetic 1RSB Results

We turn now to the results obtained via the energetic 1RSB cavity method, i.e. taking simultaneously the limits $\mu \rightarrow -\infty$ and $m \rightarrow 0$ with a finite value for $y = -\mu m$. The equations to solve in this case amounts to find the fixed point of Eq. (75), from which one obtains the 1RSB potential (81) and the energetic complexity $\Sigma_e(\theta)$ from the Legendre transform structure explained in (64), as a parametric plot varying the parameter y . The computational complexity of this problem is drastically reduced compared to the complete 1RSB equations: as in the RS case one has a set of (roughly) $2T$ equations on $2T$ real unknowns, instead of an equation on a probability distribution of fields. More technical details on the procedure to solve these equations can be found in Appendix 1.

Figure 6 displays the energetic complexity $\Sigma_e(\theta)$ for a few values of T , in the cases $k = l = 2$ and $k = 3, l = 2$. The expert reader will notice that we restricted the range of y used in this plot to the so-called physical branch, in such a way that Σ_e is a concave function of θ . The most important characteristics of these curves are the values of $\theta_{\min,1}$ where the complexity vanishes, and the corresponding values y_s of the parameter y ; these are reported for several values of T in the last columns of the Tables 1, 2, and 3. Indeed $\theta_{\min,1}$ is the 1RSB prediction for θ_{\min} , as it corresponds to the smallest density of active sites in initial




Fig. 6 The complexity $\Sigma_e(\theta)$ obtained from the energetic 1RSB cavity formalism, for $k = l = 2$ (left panel) and $k = 3, l = 2$ (right panel); see Sect. 4.2 for explanations on the $T \rightarrow \infty$ result

configurations belonging to clusters with a non-negative complexity. For $T = 1$ these values can be successfully cross-checked with the results of the Biroli–Mézard model [70], and the parameter y_s agrees with the fit of $-\mu m_s(\mu)$ in the limit $\mu \rightarrow -\infty$ obtained from the full 1RSB equations (cf. right panel of Fig. 5).

4.2 The Large T Limit

The limit case $T \rightarrow \infty$ is particularly interesting as it corresponds to the original influence maximization problem with no constraint on the time taken to activate the whole graph. This limit can be performed analytically for the RS and energetic 1RSB formalism; the technical details of these computations can be found in Appendix section “The Large T Limit”, we present here the results of these analytical simplifications. It turns out that the case $k = l$ is qualitatively different from the case $k > l$, we shall thus divide this section according to this distinction.

4.2.1 The Case $k = l$

Let us first recall that when $k = l$ the dynamics from a random initial configuration of density θ has a continuous transition at $\theta_r(k, k) = \frac{k-1}{k}$ (see Sect. 2.2); we also saw in Sect. 2.4 that minimal contagious sets (with no constraint on the activation time) correspond to minimal decycling sets, which led to the bound $\theta_{\min}(k, k) \geq \frac{k-1}{2k} = \frac{\theta_r(k, k)}{2}$. In the rest of this subsection we shall for simplicity abbreviate $\theta_r(k, k)$ by θ_r .

As suggested by the left panel of Fig. 3 in the case $k = l = 2$, the RS entropy $s(\theta)$ converges to a limit curve when $T \rightarrow \infty$. This limit curve can actually be computed analytically for all k ; we defer the details of the computation to Appendix section “Asymptotics for $l = k$ ” and only state here the properties of this limit curve. For $\theta \geq \theta_r$ it coincides with the binary entropy function $-\theta \ln \theta - (1 - \theta) \ln(1 - \theta)$; this is a posteriori obvious. Indeed by definition of θ_r typical configurations in this density range do activate the whole graph, hence the number of activating initial configurations coincide (at the leading exponential order) with the total number of configurations of this density. A non-trivial portion of the limit curve arises in the density range $[\theta_r/2, \theta_r]$, where it is given by

$$s(\theta) = -\frac{k}{2} (2\theta - \theta_r) \ln(2\theta - \theta_r) + k\theta \ln \theta + (1 - \theta) \ln(k - 1) - \frac{k + 1}{2} \ln\left(\frac{k - 1}{k}\right). \quad (82)$$

This function has the same value and the same first derivative than the binary entropy function in θ_r , while at the lower limit $\theta_r/2$ of its range of definition it has an infinite derivative with a finite value

$$s(\theta_r/2) = \ln k - \frac{k-1}{2k} \ln(k-1) - \frac{k-1}{2} \ln 2. \tag{83}$$

The parametric plot of $s(\theta)$ also contains a vertical segment for $\theta = \theta_r/2$, from $-\infty$ to the value given in (83).

The complexity $\Sigma_e(\theta)$ of the energetic 1RSB formalism also converges to a limit curve when $T \rightarrow \infty$, as shown in Fig. 6 and obtained analytically in Appendix section ‘‘Asymptotics for $l = k$ ’’. This limit curve has the same vertical segment in $\theta_r/2$ from $-\infty$ to the value (83); the non-trivial part of the curve is given in a parametrized form as follows:

$$\Sigma_e(\tilde{\lambda}) = \ln \mathcal{Z}_{\text{site}}(\tilde{\lambda}) - \frac{k+1}{2} \ln \mathcal{Z}_{\text{edge}}(\tilde{\lambda}) - y(\tilde{\lambda})(1 - \theta(\tilde{\lambda})), \tag{84}$$

$$\theta(\tilde{\lambda}) = 1 - \frac{e^{y(\tilde{\lambda})}}{e^{y(\tilde{\lambda})} - 1} \frac{\mathcal{Z}_{\text{site}}(\tilde{\lambda}) - 1}{\mathcal{Z}_{\text{site}}(\tilde{\lambda})} - \frac{k+1}{2} \frac{1}{e^{y(\tilde{\lambda})} - 1} \frac{1 - \mathcal{Z}_{\text{edge}}(\tilde{\lambda})}{\mathcal{Z}_{\text{edge}}(\tilde{\lambda})}, \tag{85}$$

where $\tilde{\lambda}$ is the positive parameter along the curve, the Parisi parameter

$$y(\tilde{\lambda}) = \ln \left(\frac{(1 + \tilde{\lambda})^k - k \tilde{\lambda}^{k-1} - \tilde{\lambda}^k}{(k-1) \tilde{\lambda}^k} \right), \tag{86}$$

is the slope of the tangent to the curve $\Sigma_e(\theta)$, and

$$\mathcal{Z}_{\text{site}}(\tilde{\lambda}) = 1 + \frac{(k+1+\tilde{\lambda})((1+\tilde{\lambda})^{k-1} - k \tilde{\lambda}^{k-1})}{(k-1)(1+\tilde{\lambda})^k}, \tag{87}$$

$$\mathcal{Z}_{\text{edge}}(\tilde{\lambda}) = \frac{\tilde{\lambda}}{1+\tilde{\lambda}} \left(1 + \frac{(1+\tilde{\lambda})^{k-1} - \tilde{\lambda}^{k-1}}{(1+\tilde{\lambda})^k - k \tilde{\lambda}^{k-1} - \tilde{\lambda}^k} \right). \tag{88}$$

When $\tilde{\lambda} \rightarrow 0^+$ this part of the curve connects with the vertical segment in $\theta_r/2$. The large values of $\tilde{\lambda}$ yield a non-concave branch of Σ_e that has to be discarded.

Depending on the value of k qualitatively different behaviours emerge from the analysis of the RS entropy and 1RSB energetic complexity:

- For $k = l = 2$ the entropy of the endpoint in $\theta_r/2$ given in (83) is strictly positive (it is equal to $(\ln 2)/2$); moreover the energetic complexity curve converges, in the $T \rightarrow \infty$ limit, to a vertical segment (the non-trivial part parametrized by $\tilde{\lambda}$ is convex and has thus to be discarded). This leads to the conclusion that $\theta_{\min} = \theta_r/2 = 1/4$ in this case, saturating the lowerbound of (16), and recovering the rigorous result of [17] on the decycling number of 3-regular graphs. This is a reassuring evidence in favour of the validity of the approach, in particular on the interversion of the $T \rightarrow \infty$ and $N \rightarrow \infty$ limit. It would be an even more challenging computation to determine the limit of θ_d and θ_c as T diverges; we are however tempted to conjecture that they both go to $1/4$ and that the effects of replica symmetry breaking are irrelevant in this limit. A numerical argument in favour of this conjecture will be presented in Sect. 5, where it is shown that a simple greedy algorithm is able to find contagious sets of these densities. Assuming this is true, the expression (82) would give for $k = 2$ the typical (quenched) entropy of the decycling sets of 3-regular random graphs in their non-trivial regime of densities $[1/4, 1/2]$. Note that the coincidence of the RS entropy and 1RSB energetic complexity at θ_{\min} is reminiscent of the phenomenology discussed for the matching problem in [75],

which might suggest that the minimal density activating configurations are at a large Hamming distance in configuration space one from the other.

- For $k = l = 3$ the expression (83) of the entropy in $\theta_r/2$ is still positive (equal to $\ln 3 - (4/3) \ln 2$), hence the endpoint of the non-trivial part of both the RS entropy and the 1RSB complexity curves occurs in $\theta_{\min,0} = \theta_{\min,1} = \theta_r/2 = 1/3$, saturating again the bound (16). This leads to the conclusion that $\theta_{\min} = 1/3$ in this case, as was also conjectured in [17]. However, at variance with the previous case, the energetic complexity curve has a non-trivial part for $\theta > \theta_{\min}$, as shown in the left panel of Fig. 7. We thus expect that the limits of θ_d and θ_c when $T \rightarrow \infty$ are strictly greater than $1/3$, hence that simple algorithms would have difficulties to find the minimal contagious sets (see Sect. 5 for a numerical check of this statement), and that the RS entropy (82) is incorrect for some regime of densities close to $1/3$.
- Finally when $k = l \geq 4$ the entropy in (83) is negative, the cancellation of s occurs at a value $\theta_{\min,0}$ strictly between $\theta_r/2$ and θ_r , see the right panel of Fig. 7. The energetic complexity vanishes on its non-trivial part parametrized by $\tilde{\lambda}$, at a value $\theta_{\min,1}$ slightly larger than $\theta_{\min,0}$, see Table 4 for some numerical values. Whether $\theta_{\min,1}$ should be taken as a conjectured exact value for θ_{\min} or simply as a lowerbound is dubious and might depend on the value of k . Indeed one should test the stability of the 1RSB ansatz against further levels of replica symmetry breaking. This computation is in principle doable along the lines of [59,60,70], but has not been performed yet. It is however relatively easy to set up an asymptotic expansion at large k of the thresholds $\theta_{\min,0}$ and $\theta_{\min,1}$ from the expressions (82, 84). One finds that the first terms of the expansion are equal at the RS and 1RSB level, it is thus natural to conjecture that they are indeed the correct expansion of θ_{\min} , namely

$$\theta_{\min}(k, k) = 1 - \frac{2 \ln k}{k} - \frac{2}{k} + O\left(\frac{1}{k \ln k}\right). \quad (89)$$

This conjecture is in agreement with the rigorous lowerbound proven in [41],

$$\theta_{\min}(k, k) \geq 1 - \frac{2 \ln k}{k} - \frac{4 - 2 \ln 2}{k} + o\left(\frac{1}{k}\right). \quad (90)$$

It can also be compared with the asymptotic expansion in the case $l = k + 1$ [38] where the inactive vertices have to form an independent set of the graph:

$$\theta_{\min}(k, k + 1) = 1 - \frac{2 \ln k}{k} + \frac{2 \ln \ln k}{k} + \frac{2 \ln 2 - 2}{k} + o\left(\frac{1}{k}\right). \quad (91)$$

The third term of this expansion is of a larger order; indeed the condition imposed on the graph induced by the inactive vertices is much more stringent when $l = k + 1$ (it has to be made of isolated vertices) with respect to the case $l = k$ (it only has to be acyclic).

Let us mention at this point that $\theta_{\min}(T)$, the minimal density of initial configuration percolating within T steps of the dynamics, reaches its asymptotic value θ_{\min} as $T \rightarrow \infty$ with different finite T corrections in the various cases listed above. The analysis of Appendix section “Asymptotics for $l = k$ ” shows that for $k = l = 2$ (resp. $k = l = 3$) these corrections are of order 2^{-T} (resp. 3^{-T}), which is in agreement with a numerical fit of the data in Table 1 (resp. Table 2). On the contrary for $k = l \geq 4$ these corrections are only polynomially small in T .

Finally, we could also compute analytically the distribution of activation times, within the RS formalism, for the initial configurations with a non-trivial density θ of active vertices in




Fig. 7 The RS entropy $s(\theta)$ and energetic 1RSB complexity $\Sigma_e(\theta)$ in the $T \rightarrow \infty$ limit, for $k = l = 3$ (left panel) and $k = l = 4$ (right panel). The binary entropy function is also plotted for comparison (the RS entropy coincides with it for $\theta \geq \theta_r$). The physical part of the complexity extends on a small range of θ , on which it is only slightly smaller than the RS entropy, the inset allows to see this small difference at the end of the domain of definition of Σ_e

Table 4 The predictions of the RS and energetic 1RSB cavity method in the $T \rightarrow \infty$ limit

k	l	θ_r	$\mu_{s=0}$	$\theta_{\min,0}$	y_s	$\theta_{\min,1}$
2	2	$\frac{1}{2}$	$-\infty$	$\frac{1}{4}$	∞	$\frac{1}{4}$
3	2	0.111111	-12.720727	0.046283	12.547960	0.046328
3	3	$\frac{2}{3}$	$-\infty$	$\frac{1}{3}$	∞	$\frac{1}{3}$
4	2	0.050781	-9.633812	0.013108	9.125975	0.013258
4	3	0.275158	$-\infty$	$\frac{1}{6}$	∞	$\frac{1}{6}$
4	4	$\frac{3}{4}$	-14.904539	0.378463	14.883293	0.378465
5	2	0.029096	-9.499859	0.005715	8.891066	0.005820
5	3	0.165116	-12.395257	0.076228	12.333754	0.076247
5	4	0.397212	$-\infty$	$\frac{1}{4}$	∞	$\frac{1}{4}$
5	5	$\frac{4}{5}$	-9.786306	0.422619	9.647302	0.422695
6	2	0.018854	-9.675930	0.003098	9.026488	0.003166
6	3	0.112870	-10.396651	0.042825	10.234248	0.042894
6	4	0.269022	-16.484079	0.150054	16.480311	0.150055
6	5	0.486312	-40.532392	0.300090	40.532392	0.300090
6	6	$\frac{5}{6}$	-8.403727	0.460014	8.191036	0.460228

the interval $[\theta_r/2, \theta_r]$. Their cumulative distribution function $P_t = \eta(t_i \leq t)$ obtained from the marginals of the law (25) reads in the $T \rightarrow \infty$ limit with t kept fixed:

$$P_{t+1} = \theta + \frac{(2\theta - \theta_r)(1 - \theta_r)}{\theta_r} w_t^{k+1} + (1 - \theta_r)(k + 1)w_t^k \left(\frac{\theta}{\theta_r} - \frac{2\theta - \theta_r}{\theta_r} w_t \right), \quad (92)$$

where w_t is a series defined recursively by

$$w_0 = \theta_r, \quad w_{t+1} = \theta_r + (1 - \theta_r)w_t^k. \quad (93)$$

Examples of this cumulative distribution are displayed in Fig. 8. As explained above the predictions of the RS cavity method are not expected to be correct for $\theta < \theta_c$; in the particular case $k = l = 2$ we however expect this result to be true down to $\theta = \theta_{\min} = 1/4$. Note that P_t




Fig. 8 The integrated distribution of activation times (92) for percolating initial conditions of density $\theta \in [\theta_r/2, \theta_r]$. The curves are presented in the case $k = l = 2$

goes to 1 when $t \rightarrow \infty$, in other words in the limit $T \rightarrow \infty$ the support of the distribution of activation times does not scale with T and remains of order 1. One can also check that when $\theta = \theta_r$, the prediction P_t of (92) coincides, as it should, with the distribution of activation times for random initial conditions of density θ_r given in Eq. (2); to see this one can notice that w_t is equal to the series \tilde{x}_t defined in Eq. (3) for the study of random initial conditions, when $k = l$ and $\theta = \theta_r$. At the lower limit of the interval of density, $\theta = \theta_r/2$, one obtains instead a simple expression,

$$P_{t+1} = \frac{\theta_r}{2} + (k + 1) \frac{1 - \theta_r}{2} w_t^k. \tag{94}$$

A straightforward analysis of (92, 93) reveals that for all $\theta < \theta_r$ the cumulative distribution P_t reaches 1 with corrections of order $1/t$, in other words the probability $P_t - P_{t-1}$ that a vertex activates precisely at time t has a power-law tail with exponent -2 . On the contrary the random initial conditions of density θ_r have $1 - P_t$ of order $1/t^2$, hence the exponent of the tail is -3 ; random initial conditions with $\theta > \theta_r$ have instead an exponentially decaying tail for their distribution of activation times.

4.2.2 The Case $k > l$

We shall now turn to a description of the limit as $T \rightarrow \infty$ of the RS and energetic 1RSB results when $k > l$, with again the technical details relegated in the Appendix section “Asymptotics for $l < k$ ”. The RS entropy $s(\theta)$ coincides with the binary entropy function for $\theta \geq \theta_r$, for exactly the same reasons as explained above in the case $k = l$ (here and in the rest of this subsection we denote θ_r the threshold $\theta_r(k, l)$). The non-trivial part of $s(\theta)$ and $\Sigma_e(\theta)$ are obtained in a parametric way, with unfortunately rather long expressions that we shall now progressively describe. We keep implicit below the dependency of all quantities on k and l when there is no risk of confusion.

This parametrization is given in terms of a real λ in the range $]0, \lambda_r]$, where this upper limit is expressed in terms of the threshold θ_r for activation from a random initial condition as $\lambda_r = (1 - \theta_r)\theta_r^{k-1}$. We need first to introduce some auxiliary functions $\widehat{u}(\lambda), \widehat{v}(\lambda), u_*(\lambda)$

and $v_*(\lambda)$. The first two are given explicitly as

$$\widehat{u}(\lambda) = \left(\frac{1 - \theta_r}{\lambda}\right)^{\frac{1}{k-1}}, \quad \widehat{v}(\lambda) = \widetilde{x}_r \left(\frac{1 - \theta_r}{\lambda}\right)^{\frac{1}{k-1}}, \tag{95}$$

where we recall that \widetilde{x}_r is the fixed-point of Eq. (3) at the bifurcation θ_r , see also (4). The last one, $v_*(\lambda)$, is defined as the smallest positive solution of

$$v = 1 + \lambda \sum_{p=l}^k \binom{k}{p} \left(\lambda l \binom{k}{l}\right)^{-\frac{k-p}{k-l}} v^{\frac{p(k-1)-k(l-1)}{k-l}}, \tag{96}$$

then $u_*(\lambda)$ can be deduced as the solution of

$$1 = \lambda l \binom{k}{l} v_*(\lambda)^{l-1} (u_*(\lambda) - v_*(\lambda))^{k-l} \quad \text{with } u_*(\lambda) \geq v_*(\lambda). \tag{97}$$

One can check that $u_*(\lambda) \geq \widehat{u}(\lambda) \geq \widehat{v}(\lambda) \geq v_*(\lambda)$ on the interval $\lambda \in]0, \lambda_r]$, and that $u_* = \widehat{u} = 1/\theta_r$ and $v_* = \widehat{v} = \widetilde{x}_r/\theta_r$ in $\lambda = \lambda_r$. We then define two functions $F_{\text{site}}(\lambda)$ and $F_{\text{edge}}(\lambda)$ through

$$F_{\text{site}}(\lambda) = \frac{\lambda}{u_*} \left[\widehat{u}^{k+1} + (k+1) \sum_{p=l}^k \binom{k}{p} \left[\frac{l-1}{k-l} I_{p-1} - I_p \right] \right] \tag{98}$$

$$F_{\text{edge}}(\lambda) = \frac{1}{u_*} \left[(\widehat{u} - \widehat{v})^2 + 2u_*v_* - v_*^2 + 2\lambda l \binom{k}{l} I_{l-1} \right] \tag{99}$$

where for clarity we kept implicit the λ dependency of \widehat{u} , \widehat{v} , u_* and v_* , and we introduced

$$I_p = \left(\lambda l \binom{k}{l}\right)^{-\frac{k-p}{k-l}} \int_{v_*}^{\widehat{v}} dv v^{\frac{p(k-1)-k(l-1)}{k-l}}, = \left(\lambda l \binom{k}{l}\right)^{-\frac{k-p}{k-l}} \tag{100}$$

$$\times \begin{cases} \ln\left(\frac{\widehat{v}}{v_*}\right) & \text{if } p=l-1 \text{ and } k=2l-1, \\ \frac{k-l}{(p+1)(k-1)-(k+1)(l-1)} \left(\widehat{v}^{\frac{(p+1)(k-1)-(k+1)(l-1)}{k-l}} - v_*^{\frac{(p+1)(k-1)-(k+1)(l-1)}{k-l}}\right) & \text{otherwise.} \end{cases}$$

We can finally give the parametric form of the RS entropy $s(\theta)$:

$$s(\lambda) = \ln(1 + F_{\text{site}}(\lambda)) - \frac{k+1}{2} \ln\left(\frac{F_{\text{edge}}(\lambda)}{u_*(\lambda)}\right) + \mu(\lambda)(1 - \theta(\lambda)),$$

$$\theta(\lambda) = \frac{1}{1 + F_{\text{site}}(\lambda)},$$

$$\mu(\lambda) = -\ln(\lambda u_*(\lambda)^k), \tag{101}$$

where $\mu(\lambda)$ is the opposite of the derivative of $s(\theta)$ in the point $\theta(\lambda)$. Thanks to the values \widehat{u} , \widehat{v} , u_* and v_* assume in λ_r this curve joins the binary entropy function in θ_r with a continuous slope.




Fig. 9 The RS entropy $s(\theta)$ and energetic 1RSB complexity $\Sigma_e(\theta)$ in the $T \rightarrow \infty$ limit, for $k = 3, l = 2$, from the analytical formulas given in (101, 102)

Similarly the 1RSB entropic complexity $\Sigma_e(\theta)$ is obtained parametrically as

$$\begin{aligned} \Sigma_e(\lambda) &= \ln \left(1 + \left(1 - \frac{1}{\lambda u_*(\lambda)^{k-1}} \right) F_{\text{site}}(\lambda) \right) - \frac{k+1}{2} \ln \left(\frac{1 + (\lambda u_*(\lambda)^{k-1} - 1) F_{\text{edge}}(\lambda)}{\lambda u_*(\lambda)^k - u_*(\lambda) + 1} \right) \\ &\quad - y(\lambda)(1 - \theta(\lambda)) , \\ \theta(\lambda) &= \frac{1 - \frac{1}{\lambda u_*(\lambda)^k} F_{\text{site}}(\lambda)}{1 + \left(1 - \frac{1}{\lambda u_*(\lambda)^{k-1}} \right) F_{\text{site}}(\lambda)} - \frac{k+1}{2} \frac{1 - \frac{1}{u_*(\lambda)} F_{\text{edge}}(\lambda)}{1 + (\lambda u_*(\lambda)^{k-1} - 1) F_{\text{edge}}(\lambda)} , \\ y(\lambda) &= \ln(\lambda u_*(\lambda)^k - u_*(\lambda) + 1) , \end{aligned} \tag{102}$$

with $y(\lambda)$ giving the slope of the tangent of $\Sigma_e(\theta)$ in the point $\theta(\lambda)$.

An example of the limit for the RS entropy can be found in the right panel of Fig. 3 for $k = 3, l = 2$, along with some finite T curves, and a similar plot for the energetic complexity is displayed in the right panel of Fig. 6. The entropy and energetic complexity for this case in the limit are compared in Fig. 9. The values $\theta_{\min,0}$ and $\theta_{\min,1}$ where $s(\theta)$ and $\Sigma_e(\theta)$ vanish are easily determined numerically from the above representation, and are collected in Table 4 for various values of k and l . For most of the cases one finds $\theta_{\min,1}$ to be slightly larger than $\theta_{\min,0}$; as explained above the exactness of this 1RSB prediction has still to be assessed from a computation of the stability with respect to further replica symmetry breaking.

There are however two special cases which stand on a different footing, namely $(k, l) = (4, 3)$ and $(k, l) = (5, 4)$. Indeed in these two cases one has the same phenomenology than for $k = l = 3$, namely a coincidence of $\theta_{\min,0}$ and $\theta_{\min,1}$ due to a vertical segment in the curves $s(\theta)$ and $\Sigma_e(\theta)$ extending to positive values. This phenomenon can be understood by studying the limit $\lambda \rightarrow 0$ of the above representation of these curves. After some algebra one finds indeed that for $k < 2l - 1$,

$$\begin{aligned} \lim_{\lambda \rightarrow 0} \theta(\lambda) &= \frac{2l - k - 1}{2l} , & \lim_{\lambda \rightarrow 0} s(\lambda) &= \lim_{\lambda \rightarrow 0} \Sigma_e(\lambda) = \frac{k+1}{2l} \ln \left(\frac{l^l}{(l-1)^{l-1}} \binom{k}{l} \right) \\ & & & - \frac{k-1}{2} \ln \left(\frac{2l}{2l - k - 1} \right) , \end{aligned} \tag{103}$$




Fig. 10 An example of the cumulative distribution of activation times for $k = 3, l = 2$, obtained with the numerical resolution of the RS equations for a large but finite value of $T = 400$, with a parameter $\lambda = 0.005$, corresponding to an initial density of active sites of 0.089. The *two horizontal lines* corresponds to $P(s = 0^+)$ and $P(s = 1^-)$ from Eq. (104), delimiting the fraction of vertices that activate within a finite time after the beginning of the process (resp. before its end)

the limiting value for θ being valid both for the RS (101) and 1RSB (102) expressions. It turns out that for $k = 4, l = 3$ and $k = 5, l = 4$, the latter expression for the entropy s and complexity Σ_e is strictly positive, hence the simple predictions $1/6$ and $1/4$ for θ_{\min} in these two cases respectively, that saturate the lowerbound of (17). We did not find any other values of k, l that produce the same phenomenon.

Finally the distribution of activation times in the RS formalism exhibits a very different pattern with respect to the case $k = l$ (see Fig. 10 for an illustration). As a matter of fact, in the limit $T \rightarrow \infty$ the activation times t of the vertices have to be divided in three categories, each of them comprising a finite fraction of the N vertices: (i) $t = O(1)$ (ii) $t = O(T)$ (iii) $t = T - O(1)$. The category (ii) of vertices can be described by a scaling function for the cumulative distribution, $P(s) = P_{t=sT}$, with $s \in]0, 1[$ a reduced time. One has $P(s = 0^+) > 0$ and $1 - P(s = 1^-) > 0$, these two numbers representing the fractions of vertices of type (i) and (iii) respectively. They can be computed following the techniques of the Appendix section “Asymptotics for $l < k$ ”, yielding for initial configurations with a fraction $\theta(\lambda) < \theta_r$ of active vertices:

$$\begin{aligned}
 P(s = 0^+) &= \theta + \theta \frac{\lambda}{u_*} \sum_{p=l}^{k+1} \binom{k+1}{p} v_*^p (u_* - v_*)^{k+1-p} , \\
 1 - P(s = 1^-) &= \theta \frac{\lambda}{u_*} \left(\frac{1 - \theta_r}{\lambda} \right)^{\frac{k+1}{k-1}} \left[1 - \sum_{p=l}^{k+1} \binom{k+1}{p} \tilde{x}_r^p (1 - \tilde{x}_r)^{k+1-p} \right] . \quad (104)
 \end{aligned}$$

5 Algorithmic Results

We shall present in this Section the results of numerical experiments performed on finite size random regular graphs, for which we have constructed explicitly some activating initial con-

figurations. We have used two strategies to do so, one based on a simple greedy heuristic, the other inspired by the results of the cavity method. Both of them build iteratively a percolating initial configuration, starting from the configuration with all vertices inactive, and adding one active vertex at a time (another route would be to start from the all active configuration and sequentially reduce the number of active vertices, but we did not investigate this alternative strategy). We shall denote τ the number of addition steps performed by the algorithm, and $\underline{\sigma}(\tau)$ the initial configuration considered at this point (that contains by definition τ active vertices). The configuration denoted $\underline{\sigma}^T(\tau)$ (resp. $\underline{\sigma}^f(\tau)$) is thus the configuration obtained after T (resp. an infinite) number of steps of the dynamics defined in (1) from the initial configuration $\underline{\sigma}(\tau)$; we will denote $|\underline{\sigma}^T(\tau)|$ the number of active vertices in this configuration. The algorithm stops when this number reaches N , as $\underline{\sigma}(\tau)$ is then the first percolating initial configurations encountered. The difference in the two algorithms to be presented below lies in the rule used to choose which additional active vertex to add in the initial configuration in a step $\tau \rightarrow \tau + 1$.

5.1 A Greedy Algorithm

Let us first consider the case of a finite time horizon T , i.e. the problem of finding an initial configuration $\underline{\sigma}$ with $\underline{\sigma}^T$ the fully active configuration and $\underline{\sigma}$ containing the smallest possible number of active vertices. The simplest strategy is to choose at each time step $\tau \rightarrow \tau + 1$ the inactive vertex of $\underline{\sigma}(\tau)$ whose activation leads to the largest possible value of $|\underline{\sigma}^T(\tau + 1)|$, and stop at the first time τ such that $\underline{\sigma}^T(\tau)$ is the fully active configuration. This can be immediately generalized to the case $T = \infty$ by including at each time step the vertex whose activation increases most $|\underline{\sigma}^f(\tau + 1)|$; this version of the greedy procedure was actually a tool in the rigorous bounds on θ_{\min} for graphs with good expansion properties of [31]. If several vertices lead to the same increase the ties can be broken arbitrarily. The time complexity of the greedy algorithm is a priori cubic in the number N of vertices: a linear number of steps $\tau \rightarrow \tau + 1$ have to be performed before finding a percolating initial configuration. For each of these steps a number of order N of candidate new configurations $\underline{\sigma}(\tau + 1)$ have to be considered, the computation of $\underline{\sigma}^T(\tau + 1)$ requiring itself a linear number of operations for each configuration. It is however easy to reduce significantly this complexity when $T = \infty$. As explained at the end of Sect. 2.1, in this case the final configuration of the dynamical process can be obtained sequentially, regardless of the order of the activations. By monotonicity the configuration $\underline{\sigma}^f(\tau + 1)$ can be computed by adding one active vertex to $\underline{\sigma}^f(\tau)$ (instead of $\underline{\sigma}(\tau)$) and determining the number (of order 1) of additional activations that can be triggered by this addition. This reduces the total complexity to a quadratic scaling with N .

In Fig. 11 we plot the fraction of active vertices in the configuration $\underline{\sigma}^T(\tau)$ as a function of the density τ/N of the active vertices in the initial configuration obtained after τ steps of this greedy procedure; when the curve reaches 1 we have thus obtained an initial configuration that percolates within T steps (note that the part of the curve for smaller τ corresponds to the alternative optimization problem labelled (i) in the introduction). The density of the contagious sets reached in this way are summarized in Table 5; as expected these densities are strictly greater than the prediction $\theta_{\min,1}$ of the 1RSB cavity method, and also than the ones reached by more involved message-passing algorithms (see the discussion in next subsection).

One can clearly see a qualitative difference between the cases $k = l$ and $k > l$ in the two panels of Fig. 11: in the latter case as T gets larger the last active vertices added in the initial configuration before finding a percolating one provoke a very steep increase in the final size




Fig. 11 The density of active vertices in the configuration $\underline{\sigma}^T(\tau)$ after τ steps of the greedy algorithm, for $k = l = 2$ (left panel) and $k = 3, l = 2$ (right panel). Each curve corresponds to a single run of the algorithm on a graph of $N = 10^4$ vertices

Table 5 The density of (finite time) contagious sets reached by the greedy and message-passing algorithms, compared to the predictions of the cavity method for their minimal size

T	$k = l = 2$				$k = 3, l = 2$			
	$\theta_{\min,1}$	θ_{sp}	θ_{maxsum} [10]	θ_{greedy}	$\theta_{\min,1}$	θ_{sp}	θ_{maxsum} [10]	θ_{greedy}
1	0.424257	0.426	0.427	0.482	0.363813	0.366	0.370	0.426
2	0.325882	0.328	0.330	0.376	0.237009	0.240	0.243	0.291
3	0.289097	0.291	0.293	0.335	0.182338	0.185	0.190	0.233
4	0.271564	0.273	0.275	0.311	0.151693	0.156	0.164	0.197
5	0.262167	0.263	0.266	0.296	0.132036	0.142	0.146	0.174
7	0.253779		0.257	0.278	0.108251	0.127	0.125	0.144
10	0.250553		0.251	0.265	0.089425		0.108	0.119

The data for the algorithmic results correspond to averages over ten graphs of size $N = 10^4$

of the activated set. As said above the greedy procedure can easily be generalized to $T = \infty$; the density of the smallest contagious sets constructed in this way are presented in Table 6 for various values of k and l . As these results demonstrate the greedy algorithm is able, in all cases we investigated, to find contagious sets with a density strictly smaller than θ_r , the density above which typical uncorrelated configurations are percolating. However in general the density reached by this simple procedure is strictly greater than the prediction $\theta_{\min,1}$ of the cavity method for their minimal size; this is in agreement with the interpretation of the replica symmetry breaking creating metastable states that trap simple local search procedures and prevent them from reaching global optima of the cost function landscape in which the search moves. The only exception is the case $k = l = 2$, for which the minimal density $1/4$ (corresponding to the decycling number of 3-regular random graphs [17]) is actually reached by the greedy procedure; this result is in line with the analysis of Sect. 4.2.1, which revealed a disappearance of the RSB phase in the large T limit for this peculiar case.

Further information on the minimal contagious sets produced by the greedy algorithm with $T = \infty$ can be obtained from the distribution of the activation times of the vertices they induce, which are plotted in Fig. 12. Of course as the graphs under study are finite the support of these distributions is bounded; in all cases we investigated we found that the time to reach total activation from these initial configurations scales logarithmically with the number of vertices of the graph (see also Fig. 13 for a comparison between two different

Table 6 The density of (infinite time) contagious sets reached by the greedy algorithm, compared to the predictions of the cavity method

k	l	θ_r	$\theta_{\min, l}$	θ_{greedy}
2	2	$\frac{1}{2}$	$\frac{1}{4}$	0.250
3	2	0.111111	0.046328	0.070
3	3	$\frac{2}{3}$	$\frac{1}{3}$	0.387
4	4	$\frac{3}{4}$	0.378465	0.482
5	5	$\frac{4}{5}$	0.422695	0.551

The algorithm was run on ten graphs of size $N = 10^4$, the last column is the average over these repetitions. Experiments with graphs of different sizes revealed a very clear $1/N$ dependency of the finite-size corrections of θ_{greedy} in the cases with $k = l$. We could not get such a clear dependency when $k > l$, slower finite-size corrections might be at play in these cases




Fig. 12 The “numerical” curves represent the distribution of activation times for the least dense activating initial configurations found by the greedy algorithm for $T = \infty$, for $k = l = 2$ (left panel) and $k = 3$, $l = 2$ (right panel). In both cases the graph studied contained $N = 8 \times 10^4$ vertices, in the left panel the complete activation is reached in 93 steps, in the right one it takes 367 steps. For comparison in the left panel the analytical prediction is plotted both for $T = \infty$ (see Eq. (94)) and for $T = 93$, in the right panel the analytical curve corresponds to $T = 367$

sizes of the graph). The qualitative difference between the cases $k = l$ and $k > l$ expected from the discussion of the $T \rightarrow \infty$ limit of Sect. 4.2 is indeed apparent on these curves; in the latter case a finite fraction of the vertices are activated at the very end of the dynamical process. However the activation time distributions induced by the configurations produced by the greedy algorithm are not in quantitative agreement with the RS analytical predictions (with a value of T and θ chosen to fit the numerical ones). A possible explanation for this discrepancy is that the greedy algorithm is a very “out-of-equilibrium” algorithm, hence the configurations it reaches are not the typical ones of the “equilibrium” measure (8).

5.2 Survey Propagation

The second algorithmic procedure we investigated is based on the insight provided by the statistical mechanics analysis on the structure of the configuration space of the problem; it corresponds indeed to the Survey Propagation algorithm introduced in [56] for the analysis of random satisfiability problem (and more precisely to its variant introduced in [16] for the energy minimization in the unsatisfiable phase of such problems). An idealized thought




Fig. 13 The distribution of activation times for the least dense activating initial configurations found by the greedy algorithm for $T = \infty$, for $k = 3, l = 2$, and two different sizes N of the graph. For $N = 4 \times 10^4$ the complete activation took place after $T = 248$ steps, while for $N = 8 \times 10^4$ it occurred at $T = 367$

experiment for the construction of minimal contagious sets would be to sequentially assign the values of the σ_i according to their marginal probabilities in the law (8), with $\epsilon = +\infty$ and $\mu = -\infty$; the exact determination of such marginals is in general a very hard computational tasks, and in practice one has to content oneself with approximations provided for instance by message passing procedures. This is the road we have followed here, by implementing the single-sample energetic 1RSB equations (75), i.e. assigning to each directed edge $i \rightarrow j$ of the graph under study a vector $P_{i \rightarrow j}$ of $2T$ probabilities. At each step τ of the algorithm the Eq. (75) are iterated several times to look for a global solution of these equations; the presence of τ active (decimated) vertices in the current configuration $\underline{\sigma}(\tau)$ is implemented as a boundary condition in these equations, easily seen to be $P_{i \rightarrow j}(h) = \delta(h - B_0)$ for the outgoing messages from an activated vertex i . The information contained in such a solution of the 1RSB equations can be a priori exploited in several ways; we chose to compute, for each vertex i not yet activated, the quantity

$$W_i = 1 - \frac{\partial}{\partial y} \ln \mathcal{Z}_{\text{site}}(\{P_{j \rightarrow i}\}_{j \in \partial i}) + \frac{1}{2} \sum_{j \in \partial i} \frac{\partial}{\partial y} \ln \mathcal{Z}_{\text{edge}}(P_{i \rightarrow j}, P_{j \rightarrow i}), \quad (105)$$

i.e. the contribution of the site i to the derivative of the potential Φ_e given in Eq. (80). This number measures indeed the tendency of i to be active in all configurations belonging to the clusters considered in the energetic 1RSB formalism. Accordingly we choose the vertex i with the largest value of W_i to be the new active vertex to be added to $\underline{\sigma}(\tau)$ in order to form $\underline{\sigma}(\tau + 1)$. For simplicity we fixed the value of y in the whole procedure to the value y_s determined analytically, that leads to a vanishing complexity before the decimation; we also tried to recompute this value of y during the course of the decimation but did not obtain significant improvement of the performances in the cases considered.

The values of the density of the percolating initial configurations we managed to construct in this way are presented in Table 5 for the two cases $k = l = 2$ and $k = 3, l = 2$, for several (relatively small) values of T . The results are better than the simple greedy algorithm, and in most of the cases also than the maxsum replica-symmetric algorithm [8–10], but in some cases deviate significantly from the prediction $\theta_{\min,1}$ for the density of minimal contagious sets.

An analytical understanding of the performances of such decimation procedures is actually a challenging open problem (see [29, 68] for partial results in the simpler case of the Belief-Propagation guided decimation). We did not study much larger values of T because we faced in this case convergence issues for the iterations of the Eqs. (75), that a simple damping did not seem to alleviate efficiently. A pragmatic, even if not completely satisfactory, position we adopted for the results at $T \geq 4$ for the case $k = 3, l = 2$, was to ignore somehow the convergence problems, stopping the iterations of (75) after a time fixed beforehand, and computing the value of W_i from these unconverged messages. As Table 5 demonstrates this attitude is not unreasonable as the densities reached are still better than the one of the greedy algorithm (yet can get worse than the maxsum procedure [8–10]).

6 Conclusions and Perspectives

In this paper we have continued the study initiated in [8, 9] of the minimal contagious sets for the bootstrap percolation (or threshold model) dynamics on random graphs. We have shown the importance of taking into account the phenomenon of replica symmetry breaking in the determination of the minimal density θ_{\min} of active vertices in percolating initial conditions, and could simplify analytically the equations determining θ_{\min} in the limit $T \rightarrow \infty$ where the constraint on the time to reach a complete activation of the graph disappears. Reformulating the problem as the minimal number of vertices to be removed in a graph in order to destroy some specific subgraphs (its cycles or more generically its q -core) we recovered a previously known result for the decycling number of 3-regular random graphs [17] as well as a conjecture for 4-regular ones [17], and proposed new quantitative conjectures for the sizes of the minimal “de-coring” sets for all pairs of degree of the graph and minimal degree of the targeted core. These take a particularly simple rational form for the removal of the 3-core in 5- and 6-regular random graphs.

Let us sketch now some possible directions for future study. A first project would be to test the stability of the 1RSB ansatz we used to compute $\theta_{\min,1}$, to assess for which values of (k, l) this number should be expected to be the exact value θ_{\min} and not only a lowerbound. This computation should be doable following the techniques of [59, 60, 70] for all finite T , and might even be simplified in the large T limit. By analogy with the independent set problem which is a marginal case of the problem investigated here one could surmise to find that the 1RSB ansatz is stable for large enough values of the degree k (and maybe also of the threshold l). This is also the regime where one can hope to see a mathematically rigorous proof of these predictions, as recently obtained for the independent sets in [33]. Asymptotic expansions of $\theta_{\min,0}(k, l)$ and $\theta_{\min,1}(k, l)$ in the large k limit for $k > l$ should also be performed, considering either l fixed in this limit, l proportional to k , or $k - l$ fixed.

For the sake of concreteness and simplicity we presented explicit results only for regular random graphs, however we gave the intermediate equations of the RS and 1RSB cavity method under a form that can be directly applied to any sparse random graph ensembles with arbitrary prescribed degree distribution, and possibly fluctuating thresholds for activation. The latter could naturally be correlated with the degree of the vertices, triggering for instance the activation if the fraction of active neighbours reaches some fixed proportion (instead of a fixed number). It would be interesting to see how the results presented here are qualitatively modified by the local fluctuations in the graph structure, which would be particularly severe in the case of power-law tails in the degree distribution.

We also concentrated exclusively in this paper on the problem of optimizing the number of initially active vertices, imposing that all vertices are active at a later time. The variant of this

problem where one puts a constraint on the maximal number of active vertices allowed in the initial configuration and try to maximize the level of activation at a later time is also relevant, in particular for applications to real-world situations. At the RS level we have sketched how to do this by controlling the parameter ϵ (the cost to be paid for finally inactive vertices) that we kept arbitrary in the first steps of the computations, a systematic study and the inclusion of the effects of replica symmetry breaking remains to be done.

Finally we believe that the message passing procedure inspired by the energetic 1RSB equations presented in Sect. 5.2 would be worth investigated further. One should try to study (and cure) the convergence issues that arise for larger values of T , maybe changing the way the information provided by the messages is used. One could in particular exploit them in a softer way by implementing a reinforcement technique [8,9] instead of a direct decimation. A more extensive comparison with the maxsum message passing procedure studied in [8,9] could also be interesting.

Acknowledgments We warmly thank Fabrizio Altarelli, Victor Bapst, Alfredo Braunstein, Amin Coja-Oghlan, Luca Dall’Asta, Svante Janson, Marc Lelarge and Riccardo Zecchina for useful discussions, and in particular FA, AB, LDA and RZ for sharing with us the unpublished numerical results [10] on their maxsum algorithm, and SJ for a useful correspondence and for pointing out the reference [17]. The authors acknowledge the support of the French Agence Nationale de la Recherche (ANR) under reference ANR-11-JS02-005-01 (GAP project) and of the People Programme (Marie Curie Actions) of the European Union’s Seventh Framework Programme FP7/2007-2013/ under REA Grant Agreement No 290038.

Appendix 1: The Limit $\mu \rightarrow -\infty$ of the Fields Recursion

We justify here the Eq. (73) for the recursion $h = g(h_1, \dots, h_k)$ between “hard fields” $h_i \in \{A_0, A_1, \dots, A_{T-1}, A_T = B_T, B_{T-1}, \dots, B_1, B_0\}$. We can first notice that in Eqs. (70, 71) the (constrained) maximum over the partitions I, J, K of S_t is always reached for $|I|+|J|$ and $|I|$ as small as possible (because $a_t^{(i)} \geq b_{t-1}^{(i)} \geq b_{t-2}^{(i)}$), which allows to rewrite

$$a_t = \max \left(0, \max_{t' \in [1, T]} \max_{\substack{J, K \\ |J|=l - \mathbb{I}(t' \geq t+1)}} S_{t'}(h_1, \dots, h_k; \emptyset, J, K) \right), \tag{106}$$

$$b_t = \max \left(0, \max_{t' \in [1, t]} \max_{\substack{J, K \\ |J|=l}} S_{t'}(h_1, \dots, h_k; \emptyset, J, K) \right), \tag{107}$$

where J, K forms a partition of $\{1, \dots, k\}$. In addition one realizes that

$$\begin{aligned} & \max_{\substack{J, K \\ |J|=l}} S_t(h_1, \dots, h_k; \emptyset, J, K) = 1 \\ & \Leftrightarrow \left(\sum_{i=1}^k \mathbb{I}(h_i \in \{A_0, \dots, A_{t-1}\}) = 0 \text{ and } \sum_{i=1}^k \mathbb{I}(h_i \in \{B_0, \dots, B_{t-1}\}) \geq l \right), \tag{108} \end{aligned}$$

which by logical negation leads to

$$\begin{aligned} & \max_{\substack{J, K \\ |J|=l}} \mathcal{S}_t(h_1, \dots, h_k; \emptyset, J, K) \leq 0 \\ \Leftrightarrow & \left(\sum_{i=1}^k \mathbb{I}(h_i \in \{A_0, \dots, A_{t-1}\}) \geq 1 \text{ or } \sum_{i=1}^k \mathbb{I}(h_i \in \{B_0, \dots, B_{t-1}\}) \leq l - 1 \right). \end{aligned} \tag{109}$$

Combining these logical rules leads after a short reasoning to

$$g(h_1, \dots, h_k) = A_t \Leftrightarrow (a_t = 1 \text{ and } a_{t+1} = 0) \tag{110}$$

$$\Leftrightarrow \begin{cases} \sum_{i=1}^k \mathbb{I}(h_i \in \{B_0, \dots, B_t\}) = l - 1 \\ \text{and } \sum_{i=1}^k \mathbb{I}(h_i \in \{A_0, \dots, A_t\}) = 0 \\ \text{and } \sum_{i=1}^k \mathbb{I}(h_i = A_{t+1}) \geq 1 \end{cases}, \tag{111}$$

and

$$g(h_1, \dots, h_k) = B_t \Leftrightarrow (b_t = 1 \text{ and } b_{t-1} = 0) \tag{112}$$

$$\Leftrightarrow \begin{cases} \sum_{i=1}^k \mathbb{I}(h_i \in \{B_0, \dots, B_{t-1}\}) \geq l \\ \text{and } \sum_{i=1}^k \mathbb{I}(h_i \in \{B_0, \dots, B_{t-2}\}) \leq l - 1. \\ \text{and } \sum_{i=1}^k \mathbb{I}(h_i \in \{A_0, \dots, A_{t-1}\}) = 0 \end{cases} \tag{113}$$

Considering the various possible cases leading to a field of type A_t or B_t yields finally (73).

Appendix 2: Technical Details on the Resolution of the Factorized RS and Energetic 1RSB Equations

We shall present in this Appendix the details of the RS and energetic 1RSB cavity equations in the particular case of random $k + 1$ regular graphs with an uniform threshold l of activations. It turns out that despite their different interpretations these two version of the cavity method can be treated in an unified way. We thus begin by introducing this common formulation, then we unveil the simplifications that arise in the case $l = k$, before finally discussing the limit $T \rightarrow \infty$, both in the case $l = k$ and $l < k$.

Common Formulation

RS Cavity Method

Consider the fixed-point RS equation $h = g(h, \dots, h)$, with g defined in Eq. (40); alternatively we saw in Eqs. (66, 67) an expression for the differences $e^{-\mu a_t} - e^{-\mu a_{t+1}}$. Setting $h_i = h$ in the right-hand sides of these equations, and using the identity

$$\sum_{\substack{I, J, K \\ |I| \leq l-1 \\ |I|+|J| \geq l}} f(I, J, K) = \sum_{\substack{I, J, K \\ |I|+|J| \geq l}} f(I, J, K) - \sum_{\substack{I, J, K \\ |I| \geq l}} f(I, J, K), \tag{114}$$

for any function f of a partition I, J, K , allows to show the equivalence of the fixed-point equation on $h = (a_0, \dots, a_T, b_{T-1}, \dots, b_1)$ with:

$$e^{-\mu a_t} - e^{-\mu a_{t+1}} = e^{-\mu + \mu k a_0} \binom{k}{l-1} e^{-\mu(l-1)b_t} \times \left[\left(e^{-\mu a_{t+1}} - e^{-\mu b_t} \right)^{k-l+1} - \left(e^{-\mu a_{t+2}} - e^{-\mu b_t} \right)^{k-l+1} \right], \tag{115}$$

$$e^{-\mu b_{t+1}} - e^{-\mu b_t} = e^{-\mu + \mu k a_0} \sum_{p=l}^k \binom{k}{p} \times \left[e^{-\mu p b_t} \left(e^{-\mu a_{t+1}} - e^{-\mu b_t} \right)^{k-p} - e^{-\mu p b_{t-1}} \left(e^{-\mu a_{t+1}} - e^{-\mu b_{t-1}} \right)^{k-p} \right]. \tag{116}$$

These equations are valid for $t \in \{0, \dots, T-1\}$, with the boundary conditions $e^{-\mu b_{-1}} = 0, b_0 = 1, a_T = b_T, a_{T+1} = b_{T-1}$. The thermodynamic quantities can also be simplified in this factorized case, the site contribution to the RS free-entropy reading from Eq. (43):

$$z_{\text{site}} = 1 + e^{-\mu + \mu(k+1)a_0} \sum_{t=1}^T \sum_{p=l}^{k+1} \binom{k+1}{p} \left[e^{-\mu p b_{t-1}} \left(e^{-\mu a_t} - e^{-\mu b_{t-1}} \right)^{k+1-p} - e^{-\mu p b_{t-2}} \left(e^{-\mu a_t} - e^{-\mu b_{t-2}} \right)^{k+1-p} \right], \tag{117}$$

while the edge contribution of Eq. (42) becomes

$$z_{\text{edge}} = e^{2\mu a_0} \left[e^{-2\mu a_T} + 2 \sum_{t=0}^{T-1} \left(e^{-\mu a_t} - e^{-\mu a_{t+1}} \right) e^{-\mu b_t} \right]. \tag{118}$$

Let us introduce some new notations and define a change of parameters on the unknowns a_t, b_t , as $u_t = e^{-\mu a_t}, v_t = e^{-\mu b_t}$. We also define a new parameter λ , with $\lambda = e^{-\mu + \mu k a_0}$. In terms of these new quantities the above set of equations becomes

$$u_t - u_{t+1} = D(u_{t+1}, v_t) - D(u_{t+2}, v_t), \tag{119}$$

$$v_{t+1} - v_t = S(u_{t+1}, v_t) - S(u_{t+1}, v_{t-1}), \tag{120}$$

with $v_{-1} = 0, v_0 = 1, u_T = v_T, u_{T+1} = v_{T-1}$, and

$$D(u, v) = \lambda \binom{k}{l-1} v^{l-1} (u - v)^{k-l+1}, \quad S(u, v) = \lambda \sum_{p=l}^k \binom{k}{p} v^p (u - v)^{k-p}. \tag{121}$$

In other words the u 's and v 's are solutions of a set of polynomial equations, and as such should be viewed as a function of λ and T (and of course of k and l). They also obey, on top of the boundary conditions, the inequalities $u_0 \geq u_1 \geq \dots \geq u_T = v_T \geq v_{T-1} \geq \dots \geq v_1 \geq v_0 = 1$. The chemical potential μ has disappeared from this set of equations, but actually it is now implicitly a function of λ and T , as from the definition of λ one recovers μ with $\mu = -\ln(\lambda u_0^k)$.

For future use we emphasize here an identity between the derivatives of D and S and introduce a new function $C(u, v)$:

$$C(u, v) = \frac{\partial D}{\partial u} = \frac{\partial S}{\partial v} = \lambda l \binom{k}{l} v^{l-1} (u - v)^{k-l}. \tag{122}$$

Let us also rewrite the thermodynamic quantities in terms of these new variables. The expressions (117) and (118) become

$$z_{\text{site}} = 1 + F_{\text{site}} , \quad z_{\text{edge}} = \frac{1}{u_0} F_{\text{edge}} , \tag{123}$$

where we introduced the two functions

$$F_{\text{site}}(\lambda, T) = \frac{\lambda}{u_0} \sum_{t=1}^T \sum_{p=l}^{k+1} \binom{k+1}{p} \left[v_{t-1}^p (u_t - v_{t-1})^{k+1-p} - v_{t-2}^p (u_t - v_{t-2})^{k+1-p} \right], \tag{124}$$

$$F_{\text{edge}}(\lambda, T) = \frac{1}{u_0} \left[v_T^2 + 2 \sum_{t=0}^{T-1} (u_t - u_{t+1}) v_t \right]. \tag{125}$$

We emphasize here the dependency on λ and T , which was kept implicit in the u_t and v_t 's. One has then the final expressions of all RS thermodynamic quantities as:

$$\phi = \mu + \ln(z_{\text{site}}) - \frac{k+1}{2} \ln(z_{\text{edge}}), \quad \mu = -\ln(\lambda u_0^k), \quad s = \phi - \mu\theta, \quad \theta = \frac{1}{z_{\text{site}}}. \tag{126}$$

One can also express the probability distribution of the activation times in terms of these new variables. Denoting P_t the cumulative distribution, i.e. the probability that the activation time of one vertex is smaller or equal than t , one has from Eq. (45):

$$P_t = \frac{1}{z_{\text{site}}} [1 + F_{\text{site}}(\lambda, T, t)] , \tag{127}$$

where we defined

$$F_{\text{site}}(\lambda, T, t) = \frac{\lambda}{u_0} \sum_{t'=1}^t \sum_{p=l}^{k+1} \binom{k+1}{p} \left[v_{t'-1}^p (u_{t'} - v_{t'-1})^{k+1-p} - v_{t'-2}^p (u_{t'} - v_{t'-2})^{k+1-p} \right]. \tag{128}$$

One can check that, as it should, $P_0 = \theta$ the fraction of initially active sites (summations over empty sets being equal to zero by convention), and $P_T = 1$ (as $\epsilon = +\infty$ all vertices are active at the final time).

Energetic 1RSB Cavity Method

We now turn to a similar study of the energetic 1RSB equations in the factorized case, namely the determination of the normalized vector of probabilities $P = (p_0, \dots, p_{T-1}, q_T, \dots, q_0)$, solution of the fixed-point equation $P = G(P, \dots, P)$, with the mapping G defined in Eq. (75).

Let us first note that in general the normalization $Z[P_1, \dots, P_k]$ of (75) can be expressed in terms of q_0 ,

$$Z = 1 + (e^y - 1)(1 - Zq_0) \Rightarrow \frac{e^y}{Z} = 1 + q_0(e^y - 1) . \tag{129}$$

This remark allows to rewrite the fixed-point equation $P = G(P, \dots, P)$ as

$$\begin{aligned}
 p_t &= (1 + q_0(e^y - 1)) \binom{k}{l-1} \left(\sum_{t'=0}^t q_{t'} \right)^{l-1} \\
 &\times \left[\left(\sum_{t'=t+1}^T q_{t'} + \sum_{t'=t+1}^{T-1} p_{t'} \right)^{k-l+1} - \left(\sum_{t'=t+1}^T q_{t'} + \sum_{t'=t+2}^{T-1} p_{t'} \right)^{k-l+1} \right], \\
 q_t &= (1 + q_0(e^y - 1)) \sum_{p=l}^k \binom{k}{p} \\
 &\times \left[\left(\sum_{t'=0}^{t-1} q_{t'} \right)^p \left(\sum_{t'=t}^T q_{t'} + \sum_{t'=t}^{T-1} p_{t'} \right)^{k-p} - \left(\sum_{t'=0}^{t-2} q_{t'} \right)^p \left(\sum_{t'=t-1}^T q_{t'} + \sum_{t'=t}^{T-1} p_{t'} \right)^{k-p} \right],
 \end{aligned}$$

where in the first line $t \in \{0, \dots, T - 1\}$ and in the second $t \in \{1, \dots, T\}$. These two sets of equations are supplemented by the normalization condition $q_0 + \dots + q_T + p_{T-1} + \dots + p_0 = 1$.

The site and edge contributions of the energetic 1RSB potential, defined in (77, 79), become in the factorized case:

$$\begin{aligned}
 \mathcal{Z}_{\text{site}} &= 1 + (e^y - 1) \sum_{t=1}^T \sum_{p=l}^{k+1} \binom{k+1}{p} \\
 &\times \left[\left(\sum_{t'=0}^{t-1} q_{t'} \right)^p \left(\sum_{t'=t}^T q_{t'} + \sum_{t'=t}^{T-1} p_{t'} \right)^{k+1-p} - \left(\sum_{t'=0}^{t-2} q_{t'} \right)^p \left(\sum_{t'=t-1}^T q_{t'} + \sum_{t'=t}^{T-1} p_{t'} \right)^{k+1-p} \right], \\
 \mathcal{Z}_{\text{edge}} &= e^{-y} + (1 - e^{-y}) \left[\left(\sum_{t=0}^T q_t \right)^2 + 2 \sum_{t=0}^{T-1} p_t \sum_{t'=0}^t q_{t'} \right].
 \end{aligned}$$

Now let us change variables and trade the unknowns p_t, q_t for some variables u_t, v_t , and the parameter y for some parameter λ , according to

$$u_t = \frac{1}{q_0} \left(\sum_{t'=0}^T q_{t'} + \sum_{t'=t}^{T-1} p_{t'} \right), \quad v_t = \frac{1}{q_0} \sum_{t'=0}^t q_{t'}, \quad \lambda = (1 + q_0(e^y - 1))q_0^{k-1}. \quad (130)$$

Inserting these definitions in the above equations one realizes that the quantities u_t and v_t are solutions of exactly the same set of Eqs. (119, 120) defined in the RS case, and obey the same boundary conditions and inequalities. From the solution of these equations, for a given value of the parameter λ , one recovers the parameter y noting that by the normalization condition one has $u_0 = 1/q_0$, hence $y = \ln(\lambda u_0^k - u_0 + 1)$. The expressions of $\mathcal{Z}_{\text{site}}$ and $\mathcal{Z}_{\text{edge}}$ within this parametrization are easily obtained from the above equations and read:

$$\mathcal{Z}_{\text{site}} = 1 + \left(1 - \frac{1}{\lambda u_0^{k-1}} \right) F_{\text{site}}, \quad \mathcal{Z}_{\text{edge}} = \frac{1 + (\lambda u_0^{k-1} - 1)F_{\text{edge}}}{\lambda u_0^k - u_0 + 1}, \quad (131)$$

with the same functions F_{site} and F_{edge} defined in Eqs. (124, 125) for the RS case. One has finally an expression for the thermodynamic quantities of the energetic 1RSB formalism as

$$\Phi_e = -y + \ln Z_{\text{site}} - \frac{k+1}{2} \ln Z_{\text{edge}}, \quad y = \ln(\lambda u_0^k - u_0 + 1), \quad \Sigma_e = \Phi_e + y\theta, \quad (132)$$

where θ is here the opposite of the derivative of Φ_e with respect to y , which after a short computation reads

$$\begin{aligned} \theta &= 1 - \frac{e^y}{e^y - 1} \frac{Z_{\text{site}} - 1}{Z_{\text{site}}} - \frac{k+1}{2} \frac{1}{e^y - 1} \frac{1 - Z_{\text{edge}}}{Z_{\text{edge}}} \\ &= \frac{1 - \frac{1}{\lambda u_0^k} F_{\text{site}}}{1 + \left(1 - \frac{1}{\lambda u_0^{k-1}}\right) F_{\text{site}}} - \frac{k+1}{2} \frac{1 - \frac{1}{u_0} F_{\text{edge}}}{1 + (\lambda u_0^{k-1} - 1) F_{\text{edge}}}. \end{aligned} \quad (133)$$

Simplifications for $l = k$

In the case $l = k$ further simplifications arise. Indeed the function $S(u, v)$ defined in (121) is in this case independent of u , and the Eqs. (119, 120) can be rewritten as:

$$v_0 = 1, \quad (134)$$

$$v_t = 1 + \lambda v_{t-1}^k \quad \text{for } t \in \{1, \dots, T\}, \quad (135)$$

$$u_{T-1} = v_T + \lambda k v_{T-1}^{k-1} (v_T - v_{T-1}), \quad (136)$$

$$u_t = u_{t+1} + \lambda k v_t^{k-1} (u_{t+1} - u_{t+2}) \quad \text{for } t \in \{0, \dots, T-2\}. \quad (137)$$

This set of equations is particularly simple to solve, and admits a single solution for each value of λ . One can indeed compute by recurrence the value of the v_t for increasing values of t from 0 to T , then deduce the value of u_{T-1} , and finally by a downward recurrence the values of u_t for t from $T-2$ to 0. The thermodynamic observables are then deduced from (126) in the RS case or (132) in the energetic 1RSB case, where the site contributions can be simplified from (124), yielding

$$F_{\text{site}}(\lambda, T) = \frac{\lambda}{u_0} \left[v_{T-1}^{k+1} + (k+1) \sum_{t=1}^T (u_t - u_{t+1}) v_{t-1}^k \right]. \quad (138)$$

These simplifications can also be performed for the function (128) giving the distribution of activation times, which reads in the case $k = l$:

$$F_{\text{site}}(\lambda, T, t) = \frac{\lambda}{u_0} \left[v_{t-1}^{k+1} + (k+1) v_{t-1}^k (u_{t+1} - v_{t-1}) + (k+1) \sum_{t'=1}^t (u_{t'} - u_{t'+1}) v_{t'-1}^k \right]. \quad (139)$$

Numerical Resolution for $l < k$

In the case $l < k$ we did not find a simple change of variables on the unknowns u_t, v_t that would put the system of Eqs. (119, 120) in the triangular form that appeared naturally when $k = l$ and led to a direct resolution by successive substitutions. We therefore resorted to the Newton-Raphson iterative method for solving (119, 120), taking care of choosing a good initial condition for the iterations to be convergent. This guess on the solution was provided

by analytical asymptotic expansions, either in the limit $\lambda \rightarrow 0$ or with $T \rightarrow \infty$ (see next paragraph). Depending on the values of λ and T we found either 0, 1 or 2 relevant solutions of (119, 120), but this multi valuedness has no physical meaning and comes only from the arbitrary choice of the parametrization in terms of λ . Indeed there is a single solution for each value of the chemical potential μ (or y in the energetic 1RSB formalism).

The Large T Limit

In the rest of this Appendix we shall justify analytically the claims made in Sects. 4.2.1 and 4.2.2 on the behaviour of the RS and energetic 1RSB solutions as T goes to infinity.

The Trivial Solution

As anticipated in Sect. 4, in the large T limit the portion of the curve $s(\theta)$ corresponding to $\theta > \theta_r$ should coincide with the entropy $-\theta \ln \theta - (1-\theta) \ln(1-\theta)$ counting all configurations with a fraction θ of initially active sites, as such configurations are typically activating (see the reminder on random initial configurations of Sect. 2.2). Let us see how to prove this statement. A moment of thought, considering for instance the form of the RS equations at $\epsilon = 0$, reveals that this situation should correspond to a solution of (119, 120) with $u_t = \tilde{u}$, independently of t . This ansatz is indeed consistent with Eq. (119), and with this substitution Eq. (120) becomes

$$v_{t+1} = 1 + S(\tilde{u}, v_t). \quad (140)$$

This last equation is a simple recursion on the v 's, with the initial value $v_0 = 1$. For the boundary condition $u_T = v_T$, $u_{T+1} = v_{T-1}$ to be asymptotically (when $T \rightarrow \infty$) verified one has to impose the values of \tilde{u} and λ such that the v_t solution of (140) converge to \tilde{u} when $t \rightarrow \infty$, in other words that the smallest fixed point solution $v \geq 1$ of $v = 1 + S(\tilde{u}, v)$ is precisely equal to \tilde{u} . The condition $\tilde{u} = 1 + S(\tilde{u}, \tilde{u})$ imposes the following relationship between \tilde{u} and λ , $\tilde{u} = 1 + \lambda \tilde{u}^k$. Using this condition one can then rewrite (140) as

$$\frac{v_{t+1}}{\tilde{u}} = \frac{1}{\tilde{u}} + \left(1 - \frac{1}{\tilde{u}}\right) \sum_{p=l}^k \binom{k}{p} \left(\frac{v_t}{\tilde{u}}\right)^p \left(1 - \frac{v_t}{\tilde{u}}\right)^{k-p}. \quad (141)$$

Comparing this equation with (3) one realizes that by definition of θ_r , all the values of \tilde{u} in the interval $[1, 1/\theta_r[$ are such that the condition $v_t \rightarrow \tilde{u}$ is fulfilled (with the value of λ fixed by $\tilde{u} = 1 + \lambda \tilde{u}^k$). Let us now compute the RS thermodynamic quantities associated with this solution. As the u_t are independent of t the summation in Eq. (124) can be performed with a telescopic identity, and yields after a short computation $F_{\text{site}} = \tilde{u} - 1$. Similarly one sees easily from (125) that $F_{\text{edge}} = \tilde{u}$ for this solution. This gives indeed the function $s(\theta) = -\theta \ln \theta - (1-\theta) \ln(1-\theta)$ for $\theta > \theta_r$ upon replacing in the expression of the RS thermodynamic potential (cf. Eq. (126)). In addition the cumulative distribution P_t of activation times defined in Eq. (127) coincides on this solution with the series x_t of Eq. (2) obtained as the activation time cumulative distribution of a random initial condition.

In the following we shall describe the non-trivial part of the resolution of the RS and energetic 1RSB equations in the large T limit, i.e. in the RS case the part of the curve $s(\theta)$ for $\theta < \theta_r$. The cases $l = k$ and $l < k$ are technically rather different, we shall thus divide the discussion according to this distinction.




Fig. 14 A graphical representation of the recursion $v_{t+1} = 1 + \lambda v_t^k$ (here for $k = 2$). The *dashed straight line* corresponds to $v_{t+1} = v_t$, the *three solid curves* are, from *bottom to top*, for $\lambda < \lambda_c$, $\lambda = \lambda_c$ and $\lambda > \lambda_c$

Asymptotics for $l = k$

As explained in Sect. 1 in the case $l = k$ the equations on v_t decouple, these quantities become independent of T and are solutions of the recurrence $v_{t+1} = 1 + \lambda v_t^k$. A straightforward study of this equation (see Fig. 14 for an illustration) reveals the existence of a critical value λ_c such that v_t converges to a finite value when $t \rightarrow \infty$ if $\lambda \leq \lambda_c$, while it diverges when $\lambda > \lambda_c$. This critical parameter and the associated fixed-point v_c of the recurrence are solution of the equations:

$$v_c = 1 + \lambda_c v_c^k, \quad 1 = \lambda_c k v_c^{k-1}, \tag{142}$$

which are easily solved and yield $\lambda_c = \frac{(k-1)^{k-1}}{k^k}$, $v_c = \frac{k}{k-1}$.

The case $\lambda < \lambda_c$ corresponds actually to the trivial solution already discussed above, let us thus consider the alternative situation, $\lambda > \lambda_c$. The divergence of v_t is then actually very steep, with a double exponential form. Indeed when $v_t \gg 1$ the recurrence becomes approximately $v_{t+1} \approx \lambda v_t^k$, which reveals that $(\ln \ln v_t)/t$ converges to $\ln k$. As $u_0 \geq v_T$ one also has a divergence of u_0 with T in this regime; from (126) (resp. (132)) this implies that the chemical potential μ of the RS formalism (resp. the parameter y of the energetic 1RSB one) go to $-\infty$ (resp. $+\infty$), i.e. that the parametric curve $s(\theta)$ (resp. $\Sigma_e(\theta)$) has a vertical tangent in this regime. Furthermore we shall prove now that the corresponding density θ of initially active sites converges to $(k - 1)/(2k)$ (both in the RS and energetic 1RSB cases), hence this branch corresponds to a vertical segment. This is actually a consequence of the following statement on the behaviour of the functions F_{site} and F_{edge} of Eqs. (138, 125):

$$\forall \lambda > \lambda_c, \quad \lim_{T \rightarrow \infty} F_{\text{site}}(\lambda, T) = \frac{k + 1}{k - 1}, \quad \lim_{T \rightarrow \infty} F_{\text{edge}}(\lambda, T) = \frac{2k}{k - 1}, \tag{143}$$

as can be easily deduced from the expressions of θ given in (123, 126) and (133), along with the divergence of u_0 in the latter case. To prove the claim of Eq. (143), let us first note that, iterating (137), one obtains

$$u_t - u_{t+1} = (u_0 - u_1) \frac{1}{k^t} \frac{1}{(\lambda v_0^{k-1})(\lambda v_1^{k-1}) \dots (\lambda v_{t-1}^{k-1})} \tag{144}$$

$$= (u_0 - u_1) \frac{1}{k^t} \frac{v_1 v_2 \dots v_{t-1}}{(v_1 - 1)(v_2 - 1) \dots (v_{t-1} - 1)}, \tag{145}$$

where we used (135) to go from the first to the second line. We can thus write

$$u_t - u_{t+1} = (u_0 - u_1) \frac{1}{k^t} \alpha_t \frac{1}{v_t}, \tag{146}$$

where we introduced the sequence α_t (note its independence on T) as

$$\alpha_t = \prod_{t'=1}^t \frac{v_{t'}}{v_{t'} - 1}, \quad \alpha_0 = 1. \tag{147}$$

We also have, in terms of this series,

$$u_0 - u_1 = k^T \frac{1}{\alpha_T} v_T (v_T - v_{T-1}). \tag{148}$$

Using these relations, along with the representation $u_0 = v_T + \sum_{t=0}^{T-1} (u_t - u_{t+1})$, allows to rewrite the definition of (125) as:

$$F_{\text{edge}} = \frac{\alpha_T \frac{1}{k^T} \frac{v_T}{v_T - v_{T-1}} + 2 \sum_{t=0}^{T-1} \alpha_t \frac{1}{k^t}}{\alpha_T \frac{1}{k^T} \frac{1}{v_T - v_{T-1}} + \sum_{t=0}^{T-1} \frac{\alpha_t}{v_t} \frac{1}{k^t}}. \tag{149}$$

The sum in the denominator can be transformed by noting that, from the definition of α_t , $\alpha_t/v_t = \alpha_t - \alpha_{t-1}$. This yields

$$F_{\text{edge}} = \frac{\alpha_T \frac{1}{k^T} \frac{v_T}{v_T - v_{T-1}} + 2 \sum_{t=0}^{T-1} \alpha_t \frac{1}{k^t}}{\alpha_T \frac{1}{k^T} \frac{1}{v_T - v_{T-1}} + \frac{1}{k^T} \alpha_{T-1} + \frac{k-1}{k} \sum_{t=0}^{T-1} \alpha_t \frac{1}{k^t}}. \tag{150}$$

Notice now that α_t has a finite limit when $t \rightarrow \infty$, thanks to the divergence of v_t (for the limit of α_t to exist it is actually enough that $v_t \gg t$). Hence the summations in the above equation converge when $T \rightarrow \infty$ thanks to the exponentially decaying factor $1/k^t$, and all other terms in the numerator and denominator are neglectible in this limit. This proves the limit $2k/(k - 1)$ for F_{edge} (one could also compute the main correction, of order k^{-T} , from this expression). The statement on F_{site} is proved with similar manipulations, that brings from (138) to the expression (exact for all T),

$$F_{\text{site}} = \frac{\alpha_T \frac{1}{k^T} \frac{v_{T-1}(v_T - 1)}{v_T(v_T - v_{T-1})} + \frac{k+1}{k} \sum_{t=0}^{T-1} \alpha_t \frac{1}{k^t}}{\alpha_T \frac{1}{k^T} \frac{1}{v_T - v_{T-1}} + \frac{1}{k^T} \alpha_{T-1} + \frac{k-1}{k} \sum_{t=0}^{T-1} \alpha_t \frac{1}{k^t}}. \tag{151}$$

As above the limit $T \rightarrow \infty$ can now be taken safely, the converging summations being the only non-vanishing terms of the numerator and denominator, hence the convergence of F_{site} to $(k + 1)/(k - 1)$, with corrections of order k^{-T} . These corrections actually contribute to the non-trivial dependence on λ of s and Σ_e (which are both finite) in this regime; we did not

push their determination further, and merely observe here that their order k^{-T} explains the statement on the finite T corrections to θ_{\min} for $k = l = 2$ and $k = l = 3$ made in Sect. 4.2.1.

We have just seen that in the $T \rightarrow \infty$ limit the cases $\lambda < \lambda_c$ and $\lambda > \lambda_c$ describe, respectively, the trivial branch $\theta > \theta_r$ of the RS entropy and its vertical segment at $\theta_r/2$. To describe the range $[\theta_r/2, \theta_r]$ of non-trivial densities of initially active sites one has thus to investigate a regime where λ is in a T -dependent scaling window around λ_c .

Let us denote \tilde{v}_t the solution of the recursion right at the critical point, i.e. $\tilde{v}_{t+1} = 1 + \lambda_c \tilde{v}_t^k$, with $\tilde{v}_0 = 1$. This series converges to v_c , with an asymptotic behaviour which is easily found to be

$$\tilde{v}_t = v_c - \frac{2k}{(k-1)^2} \frac{1}{t} + O\left(\frac{1}{t^2}\right). \tag{152}$$

Now if $\lambda = \lambda_c + \delta$, with an infinitesimal positive value of δ , the solution v_t of the recursion $v_{t+1} = 1 + \lambda v_t^k$ spends a time of order $\delta^{-1/2}$ around the avoided fixed-point v_c before crossing over to the doubly exponentially growing regime investigated above (this is a general feature of such recursive equations in the neighbourhood of a bifurcation, see for instance [22]). It is thus natural to investigate the scaling window parametrized by $\hat{\lambda}$ as

$$\lambda = \lambda_c + 2\pi^2 \frac{(k-1)^{k-2} \hat{\lambda}^2}{k^{k-1} T^2}, \tag{153}$$

the numerical prefactor and the square on $\hat{\lambda}$ being chosen to simplify the following expressions. One can then look for a solution of the recurrence equation under the form $v_t = v_c + \frac{1}{T} V(t/T)$, with $V(s)$ a scaling function. Expanding at the leading order in T one obtains a differential equation on V ,

$$V'(s) = \frac{2\pi^2 k \hat{\lambda}^2}{(k-1)^2} + \frac{(k-1)^2}{2k} V(s)^2. \tag{154}$$

The latter can be integrated into

$$V(s) = -\frac{2k}{(k-1)^2} \frac{\pi \hat{\lambda}}{\tan(\pi \hat{\lambda} s)}, \tag{155}$$

the constant in the solution of the differential equation being obtained by a matching argument between the regime $s \rightarrow 0$ and the large t asymptotics of the critical series \tilde{v}_t given in (152). Note that this form is only valid for $\hat{\lambda} < 1$, otherwise one enters the regime where v_T diverges with T . One can furthermore assume a similar scaling ansatz for the u_t , introducing a scaling function $U(s)$ under the form $u_t = v_c + U(t/T)$. Inserting these forms in Eq. (137) yields a differential equation on U ,

$$\frac{U''(s)}{U'(s)} = -\frac{(k-1)^2}{k} V(s), \tag{156}$$

which is integrated in

$$U'(s) = B \sin^2(\pi \hat{\lambda} s), \quad U(s) = A + \frac{B}{2} \left(s - \frac{\sin(2\pi \hat{\lambda} s)}{2\pi \hat{\lambda}} \right), \tag{157}$$

with A and B two constants of integration. These can be fixed by imposing the boundary conditions $u_T = v_T$ and $u_{T+1} = v_{T-1}$, which translates here in $U(1) = V(1)/T$ and $U'(1) = -V'(1)/T$. Solving these equations yield A and B ; considering in particular $u_0 = v_c + U(0)$ one obtains, at the leading order in a large T expansion,

$$u_0 = v_c + \frac{1}{T} \frac{\widehat{\lambda}^2}{\sin^4(\pi\widehat{\lambda})} \left(1 - \frac{\sin(2\pi\widehat{\lambda})}{2\pi\widehat{\lambda}} \right) \frac{\pi^2 k}{(k-1)^2} - \frac{1}{T} \frac{\widehat{\lambda}}{\tan(\pi\widehat{\lambda})} \frac{2\pi k}{(k-1)^2}. \tag{158}$$

One realizes at this point that for any fixed $\widehat{\lambda} < 1$, the limit of u_0 coincides with v_c , in other words we are describing in this regime the end of the trivial branch, with $\theta \approx \theta_r$. To describe the non-trivial regime of densities $[\theta_r/2, \theta_r]$ one has thus to further refine the scaling window, taking now $\widehat{\lambda}$ approaching 1 in a T -dependent way. The inspection of (158) reveals that the correct scaling that allows to obtain a non-trivial limit of u_0 corresponds to $\widehat{\lambda} = 1 - O(T^{-1/4})$. We shall thus set

$$\widehat{\lambda} = 1 - \frac{1}{\sqrt{\pi}} \left(\frac{\widetilde{\lambda}}{(k-1)T} \right)^{\frac{1}{4}}, \tag{159}$$

with $\widetilde{\lambda} > 0$ the new parameter describing this scale, the numerical prefactor being chosen for convenience. After a short computation one obtains the limit as $T \rightarrow \infty$ of the thermodynamic quantities in this scaling regime of λ as

$$u_0(\widetilde{\lambda}) = \frac{k}{k-1} \frac{1+\widetilde{\lambda}}{\widetilde{\lambda}}, \quad F_{\text{site}}(\widetilde{\lambda}) = \frac{1}{k-1} \frac{k+1+\widetilde{\lambda}}{1+\widetilde{\lambda}}, \quad F_{\text{edge}}(\widetilde{\lambda}) = \frac{k}{k-1} \frac{2+\widetilde{\lambda}}{1+\widetilde{\lambda}}, \tag{160}$$

the last two expressions being obtained by inserting the scaling ansatz on u_t and v_t in the definitions (125, 138); at the lowest order one can actually replace the v_t 's by v_c there. This yields a parametric representation of the thermodynamic quantities of the RS (resp. energetic 1RSB) formalism in terms of $\widetilde{\lambda}$, by inserting these last results in Eq. (126) (resp. (132, 133)). In the RS case one can check that $\widetilde{\lambda} \rightarrow 0$ corresponds to $\theta \rightarrow \theta_r/2$, while $\widetilde{\lambda} \rightarrow \infty$ yields $\theta \rightarrow \theta_r$, hence this scaling regime allows to cover the desired range $[\theta_r/2, \theta_r]$ for the densities of initially active sites. It is furthermore possible to invert the relation $\theta(\widetilde{\lambda})$, which yields finally the formula (82) announced in the main text for the entropy of activating initial configurations of density in the non-trivial interval $[\theta_r/2, \theta_r]$. In the energetic 1RSB case this last step does not seem possible and the final result (84) is presented in a form parametrized by $\widetilde{\lambda}$. We did not embark in a systematic study of the finite T corrections in this regime, it is however clear that they are polynomially small in T , which justifies the statement made in Sect. 4.2.1 on the corrections to $\theta_{\min}(T)$ for $k = l \geq 4$.

Let us finally justify the results presented at the end of Sect. 4.2.1 on the distribution of activation times. Assuming a finite value of t , the expression of (139) becomes in the regime parametrized by $\widetilde{\lambda}$:

$$F_{\text{site}}(\widetilde{\lambda}, t) = \frac{\lambda_c}{u_0(\widetilde{\lambda})} \left[\widetilde{v}_{t-1}^{k+1} + (k+1)\widetilde{v}_{t-1}^k (u_0(\widetilde{\lambda}) - \widetilde{v}_{t-1}) \right], \tag{161}$$

the last summation in (139) yielding a subdominant correction of order $1/T$. Note that $F_{\text{site}}(\widetilde{\lambda}, t)$ tends to $F_{\text{site}}(\widetilde{\lambda})$ as $t \rightarrow \infty$, which means that the support of the distribution of the activation times does not scale with T in this regime. The expression (92) for the cumulative distribution of activation times follows then easily from its generic definition given in Eq. (127), upon expressing all the quantities depending on $\widetilde{\lambda}$ as a function of the corresponding θ . In the main text we introduced for clarity the series $w_t = \theta_r \widetilde{v}_t$, to allow for an easier comparison with the distribution of activation times from a random initial condition.

Asymptotics for $l < k$

Let us now discuss the solution of the set of Eqs. (119, 120) in the limit $T \rightarrow \infty$, in the case $l < k$, and justify the statements made in Sect. 4.2.2; as we shall see their behaviour and the method of study is qualitatively different compared to the case $l = k$.

We shall first rephrase Eqs. (119, 120) as a single recursive equation, by introducing a four-dimensional vector w_t defined by

$$w_t = \begin{pmatrix} u_t \\ u_{t+1} \\ v_t \\ v_{t-1} \end{pmatrix}. \tag{162}$$

The recursive equations (119, 120) on the u_t 's and v_t 's become a single recursion on w_t , of the form $w_{t+1} = R(w_t)$ where the function R is given by

$$R \begin{pmatrix} u \\ u_+ \\ v \\ v_- \end{pmatrix} = \begin{pmatrix} u_+ \\ E(u, u_+, v) \\ v + S(u_+, v) - S(u_+, v_-) \\ v \end{pmatrix}. \tag{163}$$

The function S was defined in (121), while $E(u, u_+, v)$ is given implicitly as $D(E(u, u_+, v), v) = D(u_+, v) + u_+ - u$, with the function D of (121). Inverting this relation one obtains an explicit expression of E :

$$E(u, u_+, v) = v + \left((u_+ - v)^{k-l+1} + \frac{1}{\lambda \binom{k}{l-1}} \frac{u_+ - u}{v^{l-1}} \right)^{\frac{1}{k-l+1}}. \tag{164}$$

We have thus a representation of the time evolution of w as the flow of a discrete dynamical system in a four-dimensional space. The boundary conditions on the u_t 's and v_t 's translate into conditions on the allowed values of w_0 and w_T . The former must indeed lie in the two-dimensional manifold with $v = 1$ and $v_- = 0$, while the latter is restricted to the two-dimensional manifold defined by $u = v$ and $u_+ = v_-$. When $T \rightarrow \infty$, for a fixed value of λ , the solution w_t of the recursion $w_{t+1} = R(w_t)$ must find a way to go infinitely slowly from the first manifold at $t = 0$ to the second one at $t = T \rightarrow \infty$. It must in consequence remains as close as possible to the fixed points of the evolution map R .

The study of the equation $w = R(w)$ is very simple and shows that these fixed points span the two-dimensional subspace with $u = u_+$, $v = v_-$. One can then compute the Jacobian matrix of R on such a fixed-point, and realizes that this matrix has two eigenvalues equal to 1 (corresponding to the invariance of the fixed-point subspace under $u \rightarrow u + \delta u$ and $v \rightarrow v + \delta v$), and two eigenvalues $C(u, v)$ and $1/C(u, v)$, where C is the function defined in (122). All the fixed points have thus an unstable direction, except the one-dimensional set of fixed points obeying the further condition $C(u, v) = 1$, which constitutes a line of marginal fixed points. In the $T \rightarrow \infty$ limit the solution w_t is thus expected to remain close to this line, otherwise the flow along the unstable directions forbid to go from one boundary manifold at $t = 0$ to the other one at $t = T \gg 1$. This analysis is corroborated by the numerical results presented in Fig. 15, where we show the solution u_t, v_t determined numerically for some large but finite value of T . In particular the right panel demonstrate that for most values of t (i.e. excluding both t finite and $T - t$ finite in the large T limit), the couple (u_t, v_t) falls on the marginal fixed-point line $C(u, v) = 1$.




Fig. 15 The solution of the Eqs. (119, 120) for $k = 3, l = 2$, with $\lambda = 0.005$ and $T = 400$. *Left panel* the solid curves are u_t (top) and v_t (bottom) as functions of t ; the dashed horizontal lines correspond, from top to bottom, to u_*, \hat{u}, \hat{v} and v_* , solutions of (166, 167). *Right panel* parametric plot of the same data, with symbols instead of lines to appreciate the discreteness in t . Dashed line is the solution of the equation $C(u, v) = 1$, almost superimposed with most of the points (v_t, u_t) . The arrows point to the beginning (v_*, u_*) and end (\hat{v}, \hat{u}) of the scaling regime along the curve $C(u, v) = 1$

More precisely, the solution u_t, v_t can be described in the large T limit by two scaling functions $U(s)$ and $V(s)$, function of a rescaled time $s = t/T \in]0, 1[$, such that at the leading order,

$$u_t = U\left(\frac{t}{T}\right), \quad v_t = V\left(\frac{t}{T}\right). \tag{165}$$

Inserting this ansatz in the Eqs. (119, 120), one realizes that the condition $C(U(s), V(s)) = 1$, that we obtained intuitively above, is indeed precisely what is needed to enforce (119, 120) at the leading order in the large T limit. Note that the explicit dependency of U and V on s can be determined from the sub-dominant corrections in this limit; however we shall not need it in what follows. It will indeed be enough to compute the value of U and V for t small and t close to T , i.e. for s around 0 and 1. As revealed by the numerical data presented in Fig. 15, the matching between the scaling regime described by the functions U, V (i.e. for s strictly between 0 and 1) and the boundary conditions at $t = 0$ and $t = T$ affects the series v_t but not u_t . In other words, for t finite while $T \rightarrow \infty$ one has $u_t \rightarrow u_* = U(0)$ independently of t , where u_* is some (λ dependent) constant still to be determined, while v_t converges to the solution of the recursion $v_{t+1} = v_t + S(u_*, v_t) - S(u_*, v_{t-1})$ obtained from (120) by replacing u_t by its limit u_* . Equivalently one has in this regime $v_{t+1} = 1 + S(u_*, v_t)$. When $t \rightarrow \infty$ (after the large T limit) this series v_t converges to $v_* = V(0)$, the smallest fixed-point solution of this recursion on v ; for this behaviour to match the beginning of the scaling regime (i.e. $s \rightarrow 0$) one must impose simultaneously

$$C(u_*, v_*) = 1, \quad \text{and} \quad v_* = 1 + S(u_*, v_*). \tag{166}$$

The first equation allows to express u_* as a function of v_* ; replacing in the second one leads to the single equation on v_* given in Eq. (96), while (97) is nothing but an explicit version of the condition $C(u_*, v_*) = 1$. A similar reasoning in the regime $T - t$ finite reveals that $U(1) = \hat{u}$ and $V(1) = \hat{v}$ have to obey

$$C(\hat{u}, \hat{v}) = 1, \quad \text{and} \quad \hat{v} = S(\hat{u}, \hat{v}) + \hat{u} - S(\hat{u}, \hat{u}). \tag{167}$$

It is easy to check that the expressions of \hat{u} and \hat{v} given in (95) are indeed solutions of these two equations, using the equations on θ_r and \tilde{x}_r of Eq. (4). By definition for $\lambda \in]0, \lambda_r]$ one




Fig. 16 The functions u_* , \widehat{u} , \widehat{v} and v_* (from top to bottom) solutions of Eqs. (166, 167) as a function of λ for $k = 3, l = 2$. The upper two and lower two curves meet in $\lambda = \lambda_r$. When $\lambda \rightarrow 0$ the upper three curves diverge, while v_* converges to $l/(l - 1)$

has $u_* \geq \widehat{u} \geq \widehat{v} \geq v_*$, see Fig. 16 for a representation of the solution of the Eqs. (166, 167) as a function of λ . In λ_r , where one recovers the trivial solution studied in Appendix section “The Trivial Solution”, one has $u_* = \widehat{u} = 1/\theta_r$ and $v_* = \widehat{v} = \widetilde{x}_r/\theta_r$.

Let us now deduce the value of F_{site} and F_{edge} in the large T limit from the above characterization of the behaviour of the u_t ’s and v_t ’s. From Eq. (125) one has in this limit

$$\lim_{T \rightarrow \infty} F_{\text{edge}}(\lambda, T) = \frac{1}{u_*} \left[\widehat{u}^2 - 2 \int_0^1 ds U'(s)V(s) \right], \tag{168}$$

the matching regimes of t finite and $T - t$ finite having neglectible contributions to the summation. The integral above can be computed even if we have not determined the time-dependency of the scaling functions $U(s)$ and $V(s)$: using $ds U'(s) = du$ and the condition $C(U(s), V(s)) = 1$, one has

$$- \int_0^1 ds U'(s)V(s) = \int_{\widehat{u}}^{u_*} du v(u) = u_* v_* - \widehat{u} \widehat{v} + \int_{v_*}^{\widehat{v}} dv u(v), \tag{169}$$

where $u(v)$ (resp. $v(u)$) is the solution of $C(u(v), v) = 1$ (resp. $C(u, v(u)) = 1$). The equation $C(u(v), v) = 1$ can be explicitly solved into

$$u(v) = v + \left(\lambda l \binom{k}{l} \right)^{-\frac{1}{k-l}} v^{-\frac{l-1}{k-l}}. \tag{170}$$

This allows to compute the integral in (169) and to obtain (99).

We shall now compute similarly the limit of F_{site} that was defined in Eq. (124). In that equation we shall exploit the fact that $u_t - u_{t+1}$ is of order $1/T$ to perform the approximation

$$\begin{aligned} (u_t - v_{t-2})^{k+1-p} &= (u_{t-1} - v_{t-2})^{k+1-p} + (k + 1 - p) \\ &\quad (u_t - u_{t-1})(u_{t-1} - v_{t-2})^{k-p} + O\left(\frac{1}{T^2}\right). \end{aligned} \tag{171}$$

Within this approximation the first term leads to a telescopic summation, we then get

$$F_{\text{site}} \sim \frac{\lambda}{u_0} \sum_{p=l}^{k+1} \binom{k+1}{p} \times \left[v_{T-1}^p (u_T - v_{T-1})^{k+1-p} - (k+1-p) \sum_{t=1}^T v_{t-2}^p (u_t - u_{t-1})(u_{t-1} - v_{t-2})^{k-p} \right] \tag{172}$$

As $u_T = v_{T-1} + O(1/T)$ in the first summation only the term $p = k + 1$ survives; the second term can be rearranged as above in terms of integrals of the scaling functions, namely

$$\lim_{T \rightarrow \infty} F_{\text{site}}(\lambda, T) = \frac{\lambda}{u_*} \left[\widehat{u}^{k+1} - (k+1) \sum_{p=l}^k \binom{k}{p} \int_0^1 ds U'(s) V(s)^p (U(s) - V(s))^{k-p} \right] \tag{173}$$

$$= \frac{\lambda}{u_*} \left[\widehat{u}^{k+1} + (k+1) \sum_{p=l}^k \binom{k}{p} \int_{\widehat{u}}^{u_*} du v(u)^p (u - v(u))^{k-p} \right] \tag{174}$$

$$= \frac{\lambda}{u_*} \left[\widehat{u}^{k+1} + (k+1) \sum_{p=l}^k \binom{k}{p} \int_{v_*}^{\widehat{v}} dv (-u'(v)) v^p (u(v) - v)^{k-p} \right] \tag{175}$$

Inserting the expression of $u(v)$ given in Eq. (170) yields easily to the value of F_{site} written in (98). The parametric representations of $s(\theta)$ and $\Sigma_e(\theta)$ given in Sect. 4.2.2 are then direct consequences of Eqs. (126, 132, 133).

For what concerns the distribution of activation times, one has in the regime $t = sT$ with $s \in]0, 1[$ the following limit for the function F_{site} defined in (128):

$$\lim_{T \rightarrow \infty} F_{\text{site}}(\lambda, T, t = sT) = \frac{\lambda}{u_*} \left[\sum_{p=l}^{k+1} \binom{k+1}{p} V(s)^p (U(s) - V(s))^{k+1-p} - (k+1) \sum_{p=l}^k \binom{k}{p} \int_0^s ds' U'(s') V(s')^p (U(s') - V(s'))^{k-p} \right]. \tag{176}$$

Studying the limit $s \rightarrow 0+$ and $s \rightarrow 1-$ of this expression leads to the expressions (104) for the fraction of vertices which activate at the very beginning and at the very end of the process.

References

1. Achlioptas, D., Coja-Oghlan, A.: Algorithmic barriers from phase transitions. In: IEEE 49th Annual IEEE Symposium on Foundations of Computer Science, 2008. FOCS'08, pp. 793–802. IEEE (2008)
2. Achlioptas, D., D’Souza, R.M., Spencer, J.: Explosive percolation in random networks. *Science* **323**(5920), 1453–1455 (2009)
3. Achlioptas, D., Ricci-Tersenghi, F.: On the solution-space geometry of random constraint satisfaction problems. In: Proceedings of the 38th Annual ACM Symposium on Theory of Computing (2006)
4. Ackerman, E., Ben-Zwi, O., Wolfowitz, G.: Combinatorial model and bounds for target set selection. *Theor. Comput. Sci.* **411**(44–46), 4017–4022 (2010)
5. Aizenman, M., Lebowitz, J.L.: Metastability effects in bootstrap percolation. *J. Phys. A* **21**(19), 3801 (1988)

6. Altarelli, F., Braunstein, A., Dall'Asta, L., Lage-Castellanos, A., Zecchina, R.: Bayesian inference of epidemics on networks via belief propagation. *Phys. Rev. Lett.* **112**, 118701 (2014)
7. Altarelli, F., Braunstein, A., Dall'Asta, L., Wakeling, R.: Containing epidemic outbreaks by message-passing techniques. *Phys. Rev. X* **4**, 021024 (2014)
8. Altarelli, F., Braunstein, A., Dall'Asta, L., Zecchina, R.: Large deviations of cascade processes on graphs. *Phys. Rev. E* **87**, 062115 (2013)
9. Altarelli, F., Braunstein, A., Dall'Asta, L., Zecchina, R.: Optimizing spread dynamics on graphs by message passing. *J. Stat. Mech.* **2013**(09), P09011 (2013)
10. Altarelli, F., Braunstein, A., Dall'Asta, L., Zecchina, R.: Private communication (2014)
11. Balogh, J., Bollobás, B., Duminil-Copin, H., Morris, R.: The sharp threshold for bootstrap percolation in all dimensions. *Trans. Am. Math. Soc.* **364**(5), 2667–2701 (2012)
12. Balogh, J., Pittel, B.G.: Bootstrap percolation on the random regular graph. *Random Struct. Algorithms* **30**(1–2), 257–286 (2007)
13. Bapst, V., Coja-Oghlan, A., Hetterich, S., Rassmann, F., Vilenchik, D.: The Condensation Phase Transition in Random Graph Coloring. [arXiv:1404.5513](https://arxiv.org/abs/1404.5513) (2014)
14. Barbier, J., Krzakala, F., Zdeborova, L., Zhang, P.: The hard-core model on random graphs revisited. *J. Phys. Conf. Ser.* **473**(1), 012021 (2013)
15. Barrat, A., Barthelemy, M., Vespignani, A.: *Dynamical Processes on Complex Networks*. Cambridge University Press, Cambridge (2008)
16. Battaglia, D., Kolář, M., Zecchina, R.: Minimizing energy below the glass thresholds. *Phys. Rev. E* **70**, 036107 (2004)
17. Bau, S., Wormald, N.C., Zhou, S.: Decycling numbers of random regular graphs. *Random Struct. Algorithms* **21**(3–4), 397–413 (2002)
18. Bayati, M., Gamarnik, D., Tetali, P.: Combinatorial approach to the interpolation method and scaling limits in sparse random graphs. *Ann. Probab.* **41**(6), 4080–4115 (2013)
19. Beineke, L.W., Vandell, R.C.: Decycling graphs. *J. Graph Theory* **25**(1), 59–77 (1997)
20. Benevides, F., Przykucki, M.: Maximum Percolation Time in Two-Dimensional Bootstrap Percolation. [arXiv](https://arxiv.org/abs/1306.2465) (2013)
21. Biroli, G., Mézard, M.: Lattice glass models. *Phys. Rev. Lett.* **88**, 025501 (2001)
22. Biskup, M., Schonmann, R.: Metastable behavior for bootstrap percolation on regular trees. *J. Stat. Phys.* **136**(4), 667–676 (2009)
23. Boccaletti, S., Latora, V., Moreno, Y., Chavez, M., Hwang, D.-U.: Complex networks: structure and dynamics. *Phys. Rep.* **424**(4), 175–308 (2006)
24. Bohman, T., Frieze, A., Wormald, N.C.: Avoidance of a giant component in half the edge set of a random graph. *Random Struct. Algorithms* **25**(4), 432–449 (2004)
25. Bohman, T., Picollelli, M.: Sir epidemics on random graphs with a fixed degree sequence. *Random Struct. Algorithms* **41**(2), 179–214 (2012)
26. Bordenave, C., Lelarge, M., Salez, J.: Matchings on infinite graphs. *Probab. Theory Relat. Fields* **157**(1–2), 183–208 (2013)
27. Chalupa, J., Leath, P.L., Reich, G.R.: Bootstrap percolation on a bethe lattice. *J. Phys. C Solid State Phys.* **12**(1), L31 (1979)
28. Chen, N.: On the approximability of influence in social networks. In: *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '08*, pp. 1029–1037 (2008)
29. Coja-Oghlan, A.: On belief propagation guided decimation for random k -sat. In: *Proceedings of 22nd SODA*, p. 957 (2011)
30. Coja-Oghlan, A.: The Asymptotic k -sat Threshold. [arXiv:1310.2728](https://arxiv.org/abs/1310.2728) (2013)
31. Coja-Oghlan, A., Feige, U., Krivelevich, M., Reichman, D.: Contagious Sets in Expanders. [arXiv:1306.2465](https://arxiv.org/abs/1306.2465) (2013)
32. Daudé, H., Mora, T., Mézard, M., Zecchina, R.: Pairs of sat assignments and clustering in random boolean formulae. *Theor. Comput. Sci.* **393**, 260–279 (2008)
33. Ding, J., Sly, A., Sun, N.: Maximum independent sets on random regular graphs. [arXiv:1310.4787](https://arxiv.org/abs/1310.4787) (2013)
34. Dorogovtsev, S.N., Mendes, J.F.F.: Evolution of networks. *Adv. Phys.* **51**(4), 1079–1187 (2002)
35. Dreyer Jr, P.A., Roberts, F.S.: Irreversible k -threshold processes: graph-theoretical threshold models of the spread of disease and of opinion. *Discrete Appl. Math.* **157**(7), 1615–1627 (2009)
36. Franz, S., Leone, M.: Replica bounds for optimization problems and diluted spin systems. *J. Stat. Phys.* **111**(3–4), 535–564 (2003)
37. Franz, S., Leone, M., Toninelli, F.L.: Replica bounds for diluted non-poissonian spin systems. *J. Phys. A Math. Gen.* **36**, 10967–10985 (2003)
38. Frieze, A., Luczak, T.: On the independence and chromatic numbers of random regular graphs. *J. Comb. Theory Ser B* **54**(1), 123–132 (1992)

39. Granovetter, M.: Threshold models of collective behavior. *Am. J. Sociol.* **83**, 1420–1443 (1978)
40. Guerra, F.: Broken replica symmetry bounds in the mean field spin glass model. *Commun. Math. Phys.* **233**(1), 1–12 (2003)
41. Haxell, P., Pikhurko, O., Thomason, A.: Maximum acyclic and fragmented sets in regular graphs. *J. Graph Theory* **57**(2), 149–156 (2008)
42. Hethcote, H.W.: The mathematics of infectious diseases. *SIAM Rev.* **42**, 599 (2000)
43. Holroyd, A.E.: Sharp metastability threshold for two-dimensional bootstrap percolation. *Probab. Theory Relat. Fields* **125**(2), 195–224 (2003)
44. Janson, S., Luczak, M., Windridge, P.: Law of large numbers for the sir epidemic on a random graph with given degrees. [arXiv:1308.5493](https://arxiv.org/abs/1308.5493) (2013)
45. Janson, S., Luczak, T., Turova, T., Vallier, T.: Bootstrap percolation on the random graph $g_{n,p}$. *Ann. Appl. Probab.* **22**(5), 1989–2047 (2012)
46. Karrer, B., Newman, M.E.J.: Message passing approach for general epidemic models. *Phys. Rev. E* **82**, 016101 (2010)
47. Kempe, D., Kleinberg, J., Tardos, E.: Maximizing the spread of influence through a social network. In: *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '03*, pp. 137–146 (2003)
48. Krzakala, F., Montanari, A., Ricci-Tersenghi, F., Semerjian, G., Zdeborová, L.: Gibbs states and the set of solutions of random constraint satisfaction problems. *Proc. Natl. Acad. Sci.* **104**(25), 10318–10323 (2007)
49. Kschischang, F., Frey, B., Loeliger, H.: Factor graphs and the sum-product algorithm. *IEEE Trans. Inf. Theory* **47**(2), 498 (2001)
50. Lelarge, M.: Diffusion and cascading behavior in random networks. *Games Econ. Behav.* **75**(2), 752–775 (2012)
51. Lokhov, A.Y., Mézard, M., Ohta, H., Zdeborova, L.: Inferring the Origin of an Epidemic with Dynamic Message-Passing Algorithm. [arXiv:1303.5315](https://arxiv.org/abs/1303.5315) (2013)
52. Mézard, M., Montanari, A.: Reconstruction on trees and spin glass transition. *J. Stat. Phys.* **124**(6), 1317–1350 (2006)
53. Mézard, M., Montanari, A.: *Information, Physics and Computation*. Oxford University Press, Oxford (2009)
54. Mézard, M., Parisi, G.: The Bethe lattice spin glass revisited. *Eur. Phys. J. B.* **20**, 217 (2001)
55. Mézard, M., Parisi, G.: The cavity method at zero temperature. *J. Stat. Phys.* **111**(1–2), 1 (2003)
56. Mézard, M., Zecchina, R.: Random k -satisfiability problem: from an analytic solution to an efficient algorithm. *Phys. Rev. E* **66**(5), 056126 (2002)
57. Molloy, M.: The freezing threshold for k -colourings of a random graph. In: *Proceedings of the 44th Symposium on Theory of Computing*, p. 921. ACM (2012)
58. Monasson, R.: Structural glass transition and the entropy of the metastable states. *Phys. Rev. Lett.* **75**, 2847–2850 (1995)
59. Montanari, A., Parisi, G., Ricci-Tersenghi, F.: Instability of one-step replica-symmetry-broken phase in satisfiability problems. *J. Phys. A* **37**(6), 2073 (2004)
60. Montanari, A., Ricci-Tersenghi, F.: On the nature of the low-temperature phase in discontinuous mean-field spin glasses. *Eur. Phys. J. B.* **33**(3), 339–346 (2003)
61. Montanari, A., Ricci-Tersenghi, F., Semerjian, G.: Clusters of solutions and replica symmetry breaking in random k -satisfiability. *J. Stat. Mech.* **2008**(04), P04004 (2008)
62. Morris, R.: Minimal percolating sets in bootstrap percolation. *Electron. J. Comb.* **16**(1), R2 (2009)
63. Newman, M.: *The structure and function of complex networks*. SIAM Rev. **45**(2), 167–256 (2003)
64. Panchenko, D.: *The Sherrington–Kirkpatrick Model*. Springer, New York (2013)
65. Panchenko, D., Talagrand, M.: Bounds for diluted mean-fields spin glass models. *Probab. Theory Relat. Fields* **130**(3), 319–336 (2004)
66. Pinto, P.C., Thiran, P., Vetterli, M.: Locating the source of diffusion in large-scale networks. *Phys. Rev. Lett.* **109**, 068702 (2012)
67. Reichman, D.: New bounds for contagious sets. *Discrete Math.* **312**(10), 1812–1814 (2012)
68. Ricci-Tersenghi, F., Semerjian, G.: On the cavity method for decimated random constraint satisfaction problems and the analysis of belief propagation guided decimation algorithms. *J. Stat. Mech.* **2009**(09), P09001 (2009)
69. Riordan, O., Warnke, L.: Achlioptas process phase transitions are continuous. *Ann. Appl. Probab.* **22**(4), 1450–1464 (2012)
70. Rivoire, O., Biroli, G., Martin, O., Mézard, M.: Glass models on Bethe lattices. *Eur. Phys. J. B.* **37**, 55 (2004)

71. Qin, S.-M., Zhou, H.-J.: Solving the Undirected Feedback Vertex Set Problem by Local Search. [arXiv:1405.0446](https://arxiv.org/abs/1405.0446) (2014)
72. Shah, D., Zaman, T.: Rumors in a network: who's the culprit? *IEEE Trans. Inf. Theory* **57**(8), 5163–5181 (2011)
73. Shrestha, M., Moore, C.: Message-passing approach for threshold models of behavior in networks. *Phys. Rev. E* **89**, 022805 (2014)
74. Talagrand, M.: The parisi formula. *Ann. Math.* **163**, 221 (2006)
75. Zdeborová, L., Mézard, M.: The number of matchings in random graphs. *J. Stat. Mech.* **2006**(05), P05003 (2006)
76. Zhou, H.-J.: Spin glass approach to the feedback vertex set problem. *Eur. Phys. J. B.* **86**(11), 1–9 (2013)

Appendix B

Rare events statistics of random walks on networks: localization and other dynamical phase transitions

Caterina De Bacco, AG, Reimer Kühn, Pierre Paga

Rare events statistics of random walks on networks: localization and other dynamical phase transitions

Caterina De Bacco,¹ Alberto Guggiola,² Reimer Kühn,³ and Pierre Paga³

¹*LPTMS, Centre National de la Recherche Scientifique et Université Paris-Sud 11, 91405 Orsay Cedex, France.*

²*LPTENS, Unité Mixte de Recherche (UMR 8549) du CNRS et de l'ENS, associée à l'UPMC Université Paris 06, 24 Rue Lhomond, 75231 Paris Cedex 05, France.*

³*Department of Mathematics, King's College London*

Rare event statistics for random walks on complex networks are investigated using the large deviations formalism. Within this formalism, rare events are realized as typical events in a suitably deformed path-ensemble, and their statistics can be studied in terms of spectral properties of a deformed Markov transition matrix. We observe two different types of phase transition in such systems: (i) rare events which are singled out for sufficiently large values of the deformation parameter may correspond to *localized* modes of the deformed transition matrix; (ii) “mode-switching transitions” may occur as the deformation parameter is varied. Details depend on the nature of the observable for which the rare event statistics is studied, as well as on the underlying graph ensemble. In the present letter we report on the statistics of the average degree of the nodes visited along a random walk trajectory in Erdős-Rényi networks. Large deviations rate functions and localization properties are studied numerically. For observables of the type considered here, we also derive an analytical approximation for the Legendre transform of the large-deviations rate function, which is valid in the large connectivity limit. It is found to agree well with simulations.

Random walks are dynamical processes widely used to analyze, organize or perform important tasks on networks such as searches [1, 2], routing or data transport [3–5]. Their popularity is due to their cheap implementation, as they rely only on local information, such as the state of the neighborhood of a given node of the network. This ensures network scalability and allows fast data transmission without the need for large storage facilities at nodes, such as big routing tables in communication networks. These features make random walks an efficient tool to explore networks characterized by a high cost of information. Examples are sensor networks [6] where many signaling packets are needed to acquire wider networks status information. In peer-to-peer networks the absence of a central server storing file locations requires users to perform repeated local searches in order to find a file to download, and various random walk strategies have been proposed as a scalable method [7–9] in this context. Less attention has been paid to characterize rare events associated with random walks on networks. Yet the occurrence of a rare event can have severe consequences. In hide-and-seek games for instance [10], rare events represent situations where the seeker finds either most (or unusually many) of the hidden targets, or conversely none (or unusually few). In the context of cyber-security, where one is concerned with worms and viruses performing random walks through a network, a rare event would correspond to a situation where unusually many sensible nodes are successfully attacked and infected, which may have catastrophic consequences for the integrity of an entire IT infrastructure. Characterizing the statistics of rare events for random walks in complex networks and its dependence on network topology is thus a problem of considerable technological importance. A variant of this problem was recently analyzed for biased random walks in complex networks [11]. That paper addressed rare fluctuations in single node occupancy for an ensemble of independent (biased) walkers in the stationary state of the

system. By contrast, our interest here is in rare event statistics of *path averages*, or equivalently of time integrated variables. Rare event statistics of this type has been looked at for instance in the context of kinetically constrained models of glassy relaxation [12]; relations to constrained ensembles of trajectories were explored in [13] for Glauber dynamics in the 1d Ising chain. While these studies were primarily concerned with the use of large deviations theory as a tool to explore dynamical phase transitions in homogeneous systems, our focus here is on the interplay between rare event statistics and the heterogeneity of the underlying system.

In the present Letter we use large deviations theory to study rare events statistics for path averages of observables associated with sites visited along trajectories of random walks. Within this formalism, rare events are realized as typical events in a suitably deformed path-ensemble [12, 14]. Their statistics can be studied in terms of spectral properties of a deformed version of the Markov transition matrix for the original random walk model, the relevant information being extracted from the algebraically largest eigenvalue of the deformed transition matrix. Such deformation may direct random walks to subsets of a network with vertices of either atypically high or atypically low coordination. It also amplifies the heterogeneity of transition matrix elements for large values of the deformation parameter and we observe that, as a consequence, the eigenvector corresponding to the largest eigenvalue of the deformed transition matrix may exhibit a *localization transition*, indicating that rare large fluctuations of path averages are typically realized by trajectories that remain localized on small subsets of the network. Within localized phases, we also encounter a second type of dynamical phase transition related to *switching between modes* as the deformation parameter used to select rare events is varied. Our methods allow us to study the role that network topology and heterogeneity play in selecting these special paths, as well as

to infer properties of paths actually selected to realize extreme events.

The model. We consider a complex network with adjacency matrix A , with entries $a_{ij} = 1$ if the edge (ij) exists, $a_{ij} = 0$ otherwise. The transition matrix W of an unbiased random walk has entries $W_{ij} = a_{ij}/k_j$ where k_j is the degree of node j and W_{ij} is the probability of a transition from j to i .

Writing $\mathbf{i}_\ell = (i_0, i_1, \dots, i_\ell)$ a path of length ℓ , quantities of interest are empirical path-averages of the form

$$\hat{\phi}_\ell = \frac{1}{\ell} \sum_{i=1}^{\ell} \xi_i, \quad (1)$$

where the ξ_i are quenched random variables associated with the vertices $i = 1, \dots, N$ of the graph, which could be independent of, be correlated with, or be deterministic functions of the degrees k_i of the vertices. It is expected that the $\hat{\phi}_\ell$ are for large ℓ sharply peaked about their mean

$$\bar{\phi}_\ell = \frac{1}{\ell} \sum_{\mathbf{i}_\ell} P(\mathbf{i}_\ell) \sum_{i=1}^{\ell} \xi_i = \left\langle \frac{1}{\ell} \sum_{i=1}^{\ell} \xi_i \right\rangle \quad (2)$$

where $P(\mathbf{i}_\ell)$ denotes the probability of the path \mathbf{i}_ℓ .

The average (2) can be obtained from the *cumulant generating function* $\psi_\ell(s) = \ell^{-1} \ln \sum_{\mathbf{i}_\ell} P(\mathbf{i}_\ell) e^{s \sum_{i=1}^{\ell} \xi_i}$ as $\bar{\phi}_\ell = \psi'_\ell(s)|_{s=0}$. Here, we are interested in rare events, for which the empirical averages $\hat{\phi}_\ell$ take values ϕ which differ significantly from their mean $\bar{\phi}_\ell$. Large deviations theory predicts that for $\ell \gg 1$ the probability density $P(\phi)$ for such an event scales exponentially with path-length ℓ , $P(\phi) \sim e^{-\ell I(\phi)}$, with a *rate function* $I(\phi)$ which, according to the Gärtner-Ellis theorem [14] is obtained as a Legendre transform $I(\phi) = \sup_s \{s\phi - \psi(s)\}$ of the limiting cumulant generating function $\psi(s) = \lim_{\ell \rightarrow \infty} \psi_\ell(s)$, provided that this limit exists and that it is differentiable. We shall see that the second condition may be violated, and that the derivative $\psi'(s)$ may develop discontinuities at certain s -values, entailing that we observe regions where $I(\phi)$ is strictly linear and only represents the convex hull of the true rate function [14].

In order to evaluate $\psi_\ell(s)$, we express path probabilities using the Markov transition matrix W and a distribution $\mathbf{p}_0 = (p_0(i_0))$ of initial conditions as $P(\mathbf{i}_\ell) = [\prod_{i=1}^{\ell} W_{i_i, i_{i-1}}] p(i_0)$, entailing that $\psi_\ell(s)$ can be evaluated in terms of a deformed transition matrix $W(s) = (e^{s\xi_i} W_{ij})$ as $\psi_\ell(s) = \ell^{-1} \ln \sum_{i_\ell, i_0} [W^\ell(s)]_{i_\ell, i_0} p(i_0)$. Using a spectral decomposition of the deformed transition matrix one can write this as

$$\psi_\ell(s) = \ln \lambda_1 + \frac{1}{\ell} \ln \left[(\mathbf{1}, \mathbf{v}_1)(\mathbf{w}_1, \mathbf{p}_0) + \sum_{\alpha(\neq 1)} \left(\frac{\lambda_\alpha}{\lambda_1} \right)^\ell (\mathbf{1}, \mathbf{v}_\alpha)(\mathbf{w}_\alpha, \mathbf{p}_0) \right]. \quad (3)$$

Here the λ_α are eigenvalues of $W(s)$, the \mathbf{v}_α and \mathbf{w}_α are the corresponding right and left eigenvectors, $\mathbf{1} = (1, \dots, 1)$, and the bracket notation (\cdot, \cdot) is used to denote an inner product. Eigenvalues are taken to be sorted in decreasing order $\lambda_1 \geq |\lambda_2| \geq |\lambda_3| \dots \geq \lambda_N$, with the first inequality being a

consequence of the Perron-Frobenius theorem [15]. This concludes the general framework. For the remainder of this Letter, we will restrict our attention to the case where $\xi_i = f(k_i)$.

For long paths, the value of the cumulant generating function is dominated by the leading eigenvalue $\lambda_1 = \lambda_1(s)$ of the transition matrix $W(s)$, so $\psi(s) = \log \lambda_1(s)$. In the $s = 0$ case, the eigenvalue problem is trivial, as the column-stochasticity of the transition matrix yields a left eigenvector $w_i \equiv 1$ corresponding to the maximal eigenvalue $\lambda_1 = 1$. The associated right eigenvector is $v_i \propto k_i$. For nonzero s , such closed form expressions are in general not known. Performing a direct matrix diagonalization is quite daunting for large system sizes N , even if one exploits methods that calculates only the first eigenvalue [16]. Hence we are interested in fast viable approximations. Here we describe one such approximation expected to be valid for networks in which vertex degrees are typically large.

Degree-based approximation. We start by considering the left eigenvectors \mathbf{w} instead of the right eigenvectors, for which the eigenvalue equation can be written as

$$\lambda w_j = \frac{1}{k_j} \sum_{i \in \partial j} w_i e^{sf(k_i)}. \quad (4)$$

This system of equations can be simplified by considering a degree-based approximation for the first eigenvector, where one assumes that the values of w_i only depend on the degree of the node i : $w_i = w(k_i)$. If the average degree is large enough and the degree distribution is not too heterogeneous, we can write the eigenvalue equation (4) by appeal to the law of large numbers as

$$\lambda_1(s) w(k) = \sum_{k'} P(k'|k) w(k') e^{sf(k')} \quad (5)$$

where $P(k'|k)$ is the probability for the neighbor of a node of degree k to have degree k' .

In an Erdős-Rényi (ER) ensemble [17], and more generally in any configuration model ensemble, we have $P(k'|k) = P(k') \frac{k'}{\langle k \rangle}$. In this case the right-hand side of (5) does not depend on k and the $w(k)$ are in fact k -independent. The eigenvalue equation then simplifies to

$$\lambda_1(s) = \left\langle \frac{k}{\langle k \rangle} e^{sf(k)} \right\rangle, \quad (6)$$

where the average is over the degree distribution $P(k)$. This approximation yields excellent results for large mean connectivities $c = \langle k \rangle$ on ER graphs, and more generally for configuration models without low degree nodes. This is illustrated in figure 1, where we plot a comparison with numerical simulations for ER graphs with $c = 30$. In figure 1 and throughout the remainder of the paper simulation results are obtained as averages over 1000 samples.

Eigenvector localization. Because of the heterogeneity of the underlying system, one finds the random walk transition matrix to exhibit localized states, both for fast and slow relaxation modes [18], even in the undeformed system, although




FIG. 1. (Colour online) Cumulant generating function $\psi(s)$ for ER networks with $c = 30$ and $f(k_i) = k_i/c$, comparing the large-degree approximation (6) (blue line) with results of a numerical simulation (green line). The inset shows the corresponding rate functions.

the eigenvector corresponding to the largest eigenvalue (the equilibrium distribution) is typically delocalized. However, given the nature of the deformed transition matrix, one expects the deformed random walk for large $|s|$ to be localized around vertices where $sf(k_i)$ is very large; hence we anticipate that in the deformed system, even the eigenvector corresponding to the largest eigenvalue *may* become localized for sufficiently large $|s|$. In order to investigate this effect quantitatively we look at the inverse participation ratio of the eigenvector corresponding to the largest eigenvalue λ_1 of $W(s)$. Denoting by v_i its i -th component, we have

$$\text{IPR}[\mathbf{v}] = \frac{\sum_i v_i^4}{\left[\sum_i v_i^2\right]^2} \quad (7)$$

One expects $\text{IPR}[\mathbf{v}] \sim N^{-1}$ for a delocalized vector, whereas $\text{IPR}[\mathbf{v}] = O(1)$ if \mathbf{v} is localized.

Results on random graphs. We performed numerical simulations to evaluate $\lambda_1(s)$ and the $\text{IPR}[\mathbf{v}_1(s)]$ for several types of network, defined by their random graph topology. In the present letter we restrict ourselves to discussing results for ER networks. We found that other network ensembles such as scale-free random graphs give qualitatively similar results; we will report on these in an extended version of this letter.

We looked at various examples for the function $f(k_i)$ but in the present letter we only report results for the normalized degree $f(k_i) = k_i/c$; other deterministic types of degree-dependent functions exhibit similar behavior, thus focusing on the normalized degree is sufficient to capture the important aspects of this problem. We restrict our simulations to the largest (giant) component of the graphs, in order to prevent spurious effects of isolated nodes or small disconnected clusters (e.g. dimers) dominating $\lambda_1(s)$ and the IPR for negative s , as these would represent trivial instances of rare events, where a walker starts, and is thus stuck on a small disconnected component of the graph. From here on, the network size given

must be understood as the size of the networks from which the giant component is extracted.

Fig. 2 shows the existence of two localized regimes for sufficiently large values of $|s|$, with IPRs on the localized side of both transitions increasing with system size. Results can be understood, as for large $|s|$ the deformed random walk is naturally attracted to the nodes with the largest (resp. smallest) degrees for positive (resp. negative) s . Thus for large negative s the deformed walk tends to be concentrated at the end of the longest dangling chain, whereas for large positive s it will be concentrated at the site with the largest available coordination. On an ER network where the large-degree tail of the degree distribution decays very fast, such a high degree vertex is likely to be connected to vertices whose degrees are lower, even significantly lower, than that of the highest degree vertex in the network, which leads to IPRs approaching 1 in the large N limit. Conversely, for negative s , the deformed random walk will be attracted to the ends of dangling chains in the network, with the probability of escape from a chain decreasing with its length (with the length of the longest dangling chain increasing with system size). This can explain that IPRs initially saturate at $1/2$ for large systems. Only upon further decreasing s to more negative values will the asymmetry of the deformed transition matrices, to and away from the end of a dangling chain, induce that further weight of the dominant eigenvector to become concentrated on the end-site, leading to a further increase of the IPR.




FIG. 2. $\text{IPR}[\mathbf{v}]$ as a function of the deformation parameter s for ER graphs with $c = 6$, and $f(k_i) = k_i/c$. The inset exhibits the N^{-1} -scaling of IPRs for 4 different values of the deformation parameter s , chosen in pairs on either side of *two* localization transitions, one at negative, and one at positive s .

From the values of $\lambda_1(s)$ we also derived the large deviation rate functions for path averages of the normalized degree $f(k_i) = k_i/c$, for various systems sizes and average connectivities. In fig. 3 we report $I(\phi)$ for an ER network at a low connectivity of $c = 3$. While the right branch of $I(\phi)$ is for large N well approximated by a parabola, our results show the emergence of a linear region on the left branch, which

becomes more pronounced as the system size is increased. This is a signature of a non-differentiable point of $\psi(s)$ at a point s^* estimated to be at $s^* = -0.060 \pm 0.002$: at this point the Gärtner-Ellis theorem cannot be used to evaluate the rate function, and the linear branch only represents the convex envelope of the true $I(\phi)$ [14]. The latter can either coincide with its convex envelope, or it can indeed be non-convex. However this information cannot be accessed by the theorem. The emergence of a jump-discontinuity in $\psi'(s)$ is due to a level crossing of the two largest eigenvalues, where the system switches between two modes that correspond to the largest eigenvalue on either side of s^* . In finite systems the crossing is an ‘avoided crossing’ due to level repulsion, but the two largest eigenvalues become asymptotically degenerate at s^* in the $N \rightarrow \infty$ limit, leading to a divergence of the correlation length $\xi(s) = [\ln(\lambda_1(s)/\lambda_2(s))]^{-1}$ at s^* , in close analogy with phenomenology of second order phase transitions, the divergence being logarithmic in N in the present case.




FIG. 3. Rate function $I(\phi)$ for ER graphs with $c = 3$, and $f(k_i) = k_i/c$ for system sizes ranging from $N = 100$ to $N = 6400$. In the inset, we show $\psi(s)$ in the vicinity of the non-differentiable point. For the largest system size, a linear fit of the convex envelope of the left branch and a quadratic fit of the right branch of $I(\phi)$ are shown as well.

Conclusions and future perspectives. In this Letter we have analyzed rare events statistics for path averages of observables associated with sites visited along random walk trajectories on complex networks. Results are obtained by looking at spectral properties of suitably deformed transition matrices. The main outcome of our analysis is the possible emergence of two types of dynamical phase transitions in low mean degree systems: localization transitions which entail that large deviations from typical values of path averages may be realized by localized modes of a deformed transition matrix, and *mode-switching transitions* signifying that the modes (eigenvectors) in terms of which large deviations are typically realized may switch as the deformation parameter s and thus the actual scale of large deviations are varied. Results of numerical simulations consistently support these claims. We also developed an analytical approximation valid for networks in

which degrees are typically large.

Our work opens up the perspective to study a broad range of further interesting problems. On a technical level, one would want to implement more powerful techniques, such as derived in [19], to obtain the largest eigenvalue in the present problem class for larger system sizes. Then there is clearly the need to systematically study the dependence of the phenomena reported here on the degree statistics, and on the nature of the observables for which path averages are looked at. We have gone some way in this direction and will report results in an extended version of the present paper. In particular one might wish to look at observables which, rather than being deterministic functions of the degree, are only statistically correlated with the degree, or at observables taking values on *edges between* nodes [13, 14]. This could be of interest in applications such as traffic or information flows on networks subject to capacity constraints on edges. Moreover, given the nature of the mode-switching transition observed in the present letter, it is clearly conceivable that *several such transitions* could be observed in a single system, depending of course on the nature of the observables studied and on the topological properties of the underlying networks. Finally, critical phenomena associated with the localization transition and with mode-switching transitions also deserve further study. We believe that this list could go on.

This work was supported by the Marie Curie Training Network NETADIS (FP7, grant 290038).

- [1] L. A. Adamic, R. M. Lukose, A. R. Puniyani, and B. A. Huberman. Search in power-law networks. *Phys. Rev. E*, 64:046135, 2001.
- [2] R. Guimerà, A. Díaz-Guilera, F. Vega-Redondo, A. Cabrales, and A. Arenas. Optimal network topologies for local search with congestion. *Phys. Rev. Lett.*, 89(24):248701, 2002.
- [3] S. D. Servetto and G. Barrenechea. Constrained random walks on random graphs: routing algorithms for large scale wireless sensor networks. In *Proceedings of the 1st ACM international workshop on Wireless sensor networks and applications*, pages 12–21. ACM, 2002.
- [4] B. Tadić, S. Thurner, and G. J. Rodgers. Traffic on complex networks: Towards understanding global statistical properties from microscopic density fluctuations. *Phys. Rev. E*, 69(3):036102, 2004.
- [5] B. Tadić and S. Thurner. Information super-diffusion on structured networks. *Physica A: Stat. Mech. and its App.*, 332:566–584, 2004.
- [6] H. Tian, H. Shen, and T. Matsuzawa. Randomwalk routing for wireless sensor networks. In *Parallel and Distributed Computing, Applications and Technologies, 2005. PDCAT 2005. Sixth International Conference on*, pages 196–200. IEEE, 2005.
- [7] Y. Chawathe, S. Ratnasamy, L. Breslau, N. Lanham, and S. Shenker. Making gnutella-like p2p systems scalable. In *Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 407–418. ACM, 2003.
- [8] Q. Lv, P. Cao, E. Cohen, K. Li, and S. Shenker. Search and replication in unstructured peer-to-peer networks. In *Proceed-*

- ings of the 16th international conference on Supercomputing, pages 84–95. ACM, 2002.
- [9] N. Bisnik and A. Abouzeid. Modeling and analysis of random walk search algorithms in p2p networks. In *Hot topics in peer-to-peer systems, 2005. HOT-P2P 2005. Second International Workshop on*, pages 95–103. IEEE, 2005.
- [10] K. Sneppen, A. Trusina, and M. Rosvall. Hide-and-seek on complex networks. *Europhys. Lett.*, 69(5):853, 2005.
- [11] V. Kishore, M. S. Santhanam, and R. E. Amritkar. Extreme events and event size fluctuations in biased random walks on networks. *Phys. Rev. E*, 85:056120, 2012.
- [12] J. P. Garrahan, R. L. Jack, V. Lecomte, E. Pitard, K. van Duivendijk, and F. van Wijland. First-order dynamical phase transition in models of glasses: an approach based on ensembles of histories. *J. Phys. A*, 42:075007, 2009.
- [13] R. L. Jack and P. Sollich. Large deviations and ensembles of trajectories in stochastic models. *Progr. of Theor. Phys. Supplement*, 184:304–317, 2010.
- [14] H. Touchette. The large deviation approach to statistical mechanics. *Phys. Rep.*, 478(1):1–69, 2009.
- [15] F. R. Gantmacher. *Applications of the Theory of Matrices*. Interscience, New York, 1959.
- [16] Cornelius C. Lanczos. *An iteration method for the solution of the eigenvalue problem of linear differential and integral operators*. United States Governm. Press Office, 1950.
- [17] P. Erdős and A. Rényi. On the evolution of random graphs. *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, 5:17–ñ61, 1960.
- [18] R. Kühn. Spectra of random stochastic matrices and relaxation in complex systems. *Europhys. Lett.*, 109:60003, 2015.
- [19] Y. Kabashima, H. Takahashi, and O. Watanabe. Cavity approach to the first eigenvalue problem in a family of symmetric random sparse matrices. *J. Phys. Conf. Ser.*, 233:012001, 2010.

Abstract

Statistical physics, originally developed to describe thermodynamic systems, has been playing for the last decades a central role in modelling an incredibly large and heterogeneous set of different phenomena taking for instance place on social, economical or biological systems. Such a vast field of possible applications has been found also for networks, as a huge variety of systems can be described in terms of interconnected elements. After an introductory part introducing these themes as well as the role of abstract modelling in science, in this dissertation it will be discussed how a statistical physics approach can lead to new insights as regards three problems of interest in network theory: how some quantity can be optimally spread on a graph, how to explore it and how to reconstruct it from partial information. Some final remarks on the importance such themes will likely preserve in the coming years conclude the work.

Keywords: Network theory, Statistical physics, Disordered systems, Inference, Spreading dynamics, Extreme events

Résumé

La physique statistique, développée à l'origine pour décrire les systèmes thermodynamiques, a joué pendant les dernières décennies un rôle central dans la modélisation d'un ensemble incroyablement vaste et hétérogène de différents phénomènes qui ont lieu par exemple dans des systèmes sociaux, économiques ou biologiques. Un champ d'applications possibles aussi vaste a été trouvé aussi pour les réseaux, comme une grande variété de systèmes peut être décrite en termes d'éléments interconnectés. Après une partie introductive sur les thèmes abordés ainsi que sur le rôle de la modélisation abstraite dans la science, dans ce manuscrit seront décrites les nouvelles perspectives auxquelles on peut arriver en approchant d'une façon physico-statistique trois problèmes d'intérêt dans la théorie des réseaux: comment une certaine quantité peut se répandre de façon optimale sur un graphique, comment explorer un réseau et comment le reconstruire à partir d'un jeu d'informations partielles. Quelques remarques finales sur l'importance que ces thèmes préserveront dans les années à venir conclut le travail.

Mots clés: Théorie des réseaux, Physique statistique, Systèmes désordonnés, Inference, Dynamique de propagation, Événements extrêmes